

Spring 2015

Mitigation of magnetic interference and compensation of bias drift in inertial sensors

Eric Christopher Frick
University of Iowa

Copyright © 2015 Eric Christopher Frick

This thesis is available at Iowa Research Online: <https://ir.uiowa.edu/etd/5472>

Recommended Citation

Frick, Eric Christopher. "Mitigation of magnetic interference and compensation of bias drift in inertial sensors." MS (Master of Science) thesis, University of Iowa, 2015.
<https://doi.org/10.17077/etd.cetctkj>

Follow this and additional works at: <https://ir.uiowa.edu/etd>

Part of the [Biomedical Engineering and Bioengineering Commons](#)

MITIGATION OF MAGNETIC INTERFERENCE AND COMPENSATION OF BIAS DRIFT IN INERTIAL SENSORS

by

Eric Christopher Frick

A thesis submitted in partial fulfillment
of the requirements for the Master of
Science degree in Biomedical
Engineering in the Graduate College of
The University of Iowa

May 2015

Thesis Supervisor: Associate Professor Salam Rahmatalla

Copyright by
ERIC CHRISTOPHER FRICK
2015
All Rights Reserved

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

MASTER'S THESIS

This is to certify that the Master's thesis of

Eric Christopher Frick

has been approved by the Examining Committee for the
thesis requirement for the Master of Science degree in
Biomedical Engineering at the May 2015 graduation.

Thesis Committee:

Salam Rahmatalla, Thesis Supervisor

David Wilder

Tim Marler

Nicole Grosland

Edwin Dove

ACKNOWLEDGMENTS

I would like to acknowledge the people in the 3D Bio-Motion Research Lab (3DBMRL) at the Center for Computer Aided Design (CCAD) at the University of Iowa for their help and support with this thesis. I would like to especially thank Jonathan DeShaw for his assistance in the experimental design and process, data processing, and acting as a sounding board in our many discussions pertaining to this topic. I would also like to thank Steven Ewart for his help in the experimental process of motion capture and application of magnetic interference.

I would like to acknowledge the Office of Naval Research, for supplying the funds necessary to conduct the research and experimentation found in this document.

I would also like to acknowledge and thank the members of my thesis committee, Professors David Wilder, Tim Marler, Nicole Grosland, and Edwin Dove, for their advice and guidance in this research.

I would like to especially thank my best friend Lauren Imhoff, for her continued and unwavering support in my endeavors, both academic and not.

Furthermore, I would like to extend a special thank you to my advisor and mentor Salam Rahmatalla, for his direction and expertise, and insights in the area of motion capture and the research process.

Finally, I would like to thank my parents Robert Frick and Sheryl Klink for their enduring patience and support during the hectic times encountered in my short tenure at graduate school, always there to offer me a hand to steady me when I faltered.

ABSTRACT

Magnetic interference in the motion capture environment is caused primarily by ferromagnetic objects and current-carrying devices disturbing the ambient, geomagnetic field. Inertial sensors gather magnetic data to determine and stabilize their global heading estimates, and such magnetic field disturbances alter heading estimates. This decreases orientation accuracy and therefore decreases motion capture accuracy. The often used Kalman Filter approach deals with magnetic interference by ignoring the magnetic data during periods interference is encountered, but this method is only effective when the disturbances are ephemeral, and cannot retroactively repair data from disturbed time periods.

The objective of this research is to develop a method of magnetic interference mitigation for environments where magnetic interference is the norm rather than the exception. To the knowledge of this author, the ability to use inertial and magnetic sensors to capture accurate, global, and drift-free orientation data in magnetically disturbed areas has yet to be developed. Furthermore there are no methods known to this author that are able to use data from undisturbed time periods to retroactively repair data from disturbed time periods. The investigation begins by exploring the use of magnetic shielding, with the reasoning that application of shielding so as to impede disturbed fields from affecting the inertial sensors would increase orientation accuracy. It was concluded that while shielding can mitigate the effect of magnetic interference, its application requires a tedious trial and error testing that was not guaranteed to improve results. Furthermore, shielding works by redirecting magnetic field lines, increasing field complexity, and thus has a high potential to exacerbate magnetic interference.

Shielding was determined to be an impractical approach, and development of a magnetic inference mitigation algorithm began. The algorithm was constructed such that magnetic data would be filtered before inclusion in the orientation estimate, with the result that exposure in an undisturbed environment would improve estimation, but exposure to a disturbed environment would have no effect. The algorithm was designed for post-processing, rather than real-time use as Kalman Filters are, which enabled magnetic data gathered before and after a time point could affect estimation.

The algorithm was evaluated by comparing it with the Kalman Filter approach of the company XSENS, using the gold standard of optical motion capture as the reference point. Under the tested conditions of stationary periods and smooth planar motion, the developed algorithm was resistant to magnetic interference for the duration of testing, while the Kalman Filter began to degrade after approximately 15 seconds. In a 190 second test, of which 180 were spent in a disturbed environment, the developed algorithm resulted in 0.4 degrees of absolute error, compared to the of the Kalman Filter's 78.8 degrees.

The developed algorithm shows the potential for inertial systems to be used effectively in situations of consistent magnetic interference. As the benefits of inertial motion capture make it a more attractive option than optical motion capture, immunity to magnetic interference significantly expands the usable range of motion capture environments. Such expansion would be beneficial for motion capture studies as a whole, allowing for the cheaper, more practical inertial approach to motion capture to supplant the more expensive and time consuming optimal option.

PUBLIC ABSTRACT

Motion capture is the process of recording motion such that it can be later recreated and analyzed. The benefits of motion capture can be seen across many disciplines, as it allows for investigation into the methodology of human motion. Applications range from improving animated motion to identifying the effect of cerebral palsy on human gait.

Recent advancements in sensor technology have simplified the motion capture process, allowing for direct application of sensors to the subjects body. Previously motion capture was primarily accomplished via cameras, necessitating large initial expense as well as a dedicated motion capture environment. In Inertial motion capture, motion is captured by sensors placed on the subject's body. It is comparatively cheap, and can be performed in a wide variety of environments. For these reasons inertial motion capture is often preferred to optical motion capture.

For all its benefits, a weakness of inertial motion capture is error accumulation due to magnetic interference. Magnetic interference is caused primarily by metal objects close to the subject. As inertial sensors measure the ambient magnetic field such interference can generate error. The focus of this thesis is the use of inertial sensors in magnetically interfered environments. The developed algorithm operates by filtering the magnetic field data and discarding data compromised by interference. The resulting algorithm allows for motion capture resistant to magnetic interference, and shows the potential for inertial motion capture to be used effectively in a broader range of environments than is currently feasible.

TABLE OF CONTENTS

LIST OF TABLES	x
LIST OF FIGURES	xi
LIST OF EQUATIONS	xviii
LIST OF ABBREVIATIONS.....	xix
PREFACE.....	xx
CHAPTER 1 : INTRODUCTION AND MOTIVATION.....	1
1.1 Introduction	1
1.1.1 MARG and Magnetometer Introduction.....	1
1.1.2 Cause of Magnetic Interference	1
1.1.3 Why Magnetic Interference is a Problem in Inertial Motion Capture	3
1.2 Literature Review of Magnetic Interference Mitigation Methods	3
1.3 Approach of This Investigation.....	6
1.3.1 Investigation of Magnetic Shielding and Magnetic Field Characteristics	7
1.3.2 Development and Testing of Mitigation Algorithm	7
CHAPTER 2 : MAGNETIC SHIELDING INVESTIGATION.....	7
2.1 Introduction	7
2.2 Experimental Methodology.....	7
2.3 Experimentation	8
2.3.1 Drift Comparison of the Roll, Pitch, and Yaw Orientations.....	8
2.3.2 Stationary Drift Comparison – Undisturbed vs Disturbed.....	10
2.3.3 Oscillatory Drift Comparison – Undisturbed vs Disturbed	13
2.3.4 Conclusions.....	16
2.3.5 Magnetic Norm and Frequency Filtering.....	16
2.4 Magnetic Shielding to Alter the Norm.....	20
2.4.2 Conclusions.....	25
2.5 Applications in Inertial Motion Capture	25

2.5.1	Field Mapping	26
2.5.2	Shield Application	26
2.6	Conclusion.....	27
CHAPTER 3 : ALGORITHM METHODOLOGY		28
3.1	Introduction	28
3.2	Methodology	29
3.2.1	Magnetic Interference Mitigation Algorithm (MIMA) Overview	29
3.2.2	Sensor Fusion.....	34
3.2.3	Heading Inclusion Filter (HIF)	41
3.2.4	Selection of Tolerance Values	53
3.3	Experimental Setup	57
3.3.1	Introduction.....	57
3.3.2	Experimental Set-Up.....	58
3.4	Conclusion.....	62
CHAPTER 4 : EXPERIMENTATION, RESULTS, AND DISCUSSION.....		63
4.1	Effect of Initialization in a Disturbed Environment.....	63
4.1.1	Objective.....	63
4.1.2	Methods.....	63
4.1.3	Results.....	64
4.1.4	Discussion.....	64
4.2	Effect of Moving into and out of a Disturbed Environment	65
4.2.1	Objective	65
4.2.2	Methods.....	65
4.2.3	Results.....	66
4.2.4	Discussion.....	67
4.3	Effect of Extended Time on MIMA Reparation Ability	67
4.3.1	Objective.....	67
4.3.2	Methods.....	68
4.3.3	Results.....	68
4.3.4	Discussion.....	69
4.4	Sensitivity Analysis.....	70

4.4.1	Objective.....	70
4.4.2	Methods.....	71
4.4.3	Results.....	71
4.4.4	Discussion.....	76
CHAPTER 5 : FUTURE RESEARCH.....		80

APPENDIX

A.1	Basics of Magnetism	82
A.1.1	Magnetic Properties	82
A.1.2	Magnetic Permeability	83
A.1.3	Magnetic Saturation	83
A.1.4	Geomagnetic Field.....	83
A.2	Characterizing Magnetic Disturbances	86
A.2.1	Hard Iron Disturbances	86
A.2.2	Soft Iron Disturbances	88
A.2.3	Frequency-Based Disturbances.....	90
A.2.4	Magnetic Norm	90
A.3	Mitigating Magnetic Distortion.....	90
A.3.1	Hard Iron Disturbances	91
A.3.2	Soft Iron Disturbances	92
A.3.3	Combination Iron Disturbances	93
A.3.4	Time-Varying Disturbances.....	94
A.4	Software/Hardware Enhancements	95
A.4.1	Disturbance Modeling.....	95
A.4.2	Sensor Enhancements	96
A.5	Magnetic Shielding	97
A.5.1	Methodology.....	97
A.5.2	Factors Affecting Shielding Capacity.....	99
A.5.3	Shielding Placement.....	101
A.6	Other Considered Mitigation Methods.....	102
A.6.1	Sensor Shielding	102
A.6.2	High-Strength Artificial Heading	102

A.7	Shielding Experimentation.....	104
A.7.1	Variability in Yaw Drift.....	104
A.7.2	Effects of Magnetic Shield Layering.....	107
REFERENCES.....		111

LIST OF TABLES

Table

1. Terms used in the filter by Madgwick et al. and their meanings..... 36
2. Here the tolerance values chosen for the MIMA's operation are listed, as well as their meaning. 53
3. The ranges of each characteristic determined to be undisturbed for each testing session. 56
4. Information pertaining to the HIF test parameters investigated in the sensitivity analysis. The fixed values represent the values used for the experiment in Section 4.3, as well as the value that parameter assumed while not being varied..... 71

LIST OF FIGURES

Figure

1. A source of magnetic interference was placed next to a sensor in an otherwise magnetically undisturbed field. The disturbance altered the X, Y, and Z components of the magnetic field, in turn causing alteration of the field's magnitude, as seen in the behavior of the norm. 2
2. Two XSENS MTws (MARG sensor arrays) were placed on a six-degree-of-freedom vibratory table. One sensor was mounted on the magnetically undisturbed aluminum seat, and the other was mounted on the magnetically disturbed steel baseplate. 9
3. This plot shows relative magnitude of roll, pitch, and heading variability as calculated by the MTw. After 150 seconds, the absolute error in roll and pitch are less than .25 degrees, whereas heading is more than 5.5 degrees off from its initial orientation..... 10
4. The drift error in the undistorted field sensor is greater than that of the distorted field sensor. The implication of these results is that a distorted field can work to stabilize the heading calculations, even if about the wrong value. Here the red represents the disturbed yaw values and the blue represents the undisturbed yaw values..... 12
5. Two MTs were mounted on a vibratory shaker table, one magnetically disturbed and the other magnetically undisturbed. In this test the shaker table was made to move in an oscillatory motion about the vertical axis (Figure 6) at a frequency of one hertz..... 13
6. Explanation of the oscillatory motion of the shaker table. The table pivoted about a central, vertical axis in a repetitive, alternating motion. This is taken from the MOOG website (MOOG is the name of the vibratory table)..... 14
7. There was a drift error of approximately five degrees in the undisturbed sensor, taking over two minutes to stabilize, and exhibiting minor drift after that..... 15
8. There was a drift error of approximately two degrees in the disturbed sensor, taking approximately thirty seconds to stabilize..... 15

9.	MTw and its body coordinate system. Image from the XSENS MTw User Manual (XSENS, 2013).....	17
10.	After filtering, the plots for both the disturbed and undisturbed cases became clearer, and the differences in the disturbed and undisturbed norms were more readily identified.....	18
11.	In the undisturbed case, the X and Y components were oscillatory and the Z component was not. This is to be expected given the rotation about the Z axis.....	19
12.	In the disturbed case, the X and Z components were oscillatory and the Y component was not. This deviation from the undisturbed behavior showcases the effect magnetic distortion has on the components of the magnetic field.....	19
13.	An overhead view of the robotic arm and shielding experimental setup.....	21
14.	A side view of the robotic arm and shielding experimental setup.....	21
15.	An overhead view of the robotic arm and shielding experimental setup with the magnetic shield removed to show the disturbance source.....	22
16.	The magnetic norm as a function of shield type and angular location. It should be noted that regardless of whether the magnetic shield lowered or increased the norm, the local field was invariably altered.....	23
17.	Here the labels mean the following: LP – Low Permeability, MP – Medium Permeability, HP – High Permeability. The numbers following HP denote the thickness of the shield in inches. The effects of each of the five different shielding approaches are shown alongside the unshielded case, quantified by comparing their relative overall difference from the idealized case of norm equal to one.....	24
18.	The overall structure of the algorithm developed in this work. The traditional approach is to feed the complete MARG data into a KF. In order to make the process more resistant to magnetic interference, here the inertial and magnetic data are separated such that only the inertial data is Kalman filtered, and the magnetic data is accounted for afterwards via the Heading Inclusion Filter.....	31

19.	Block Diagram of the orientation filter from Madgwick et al. Reprinted from “Estimation of IMU and MARG orientation using a gradient descent algorithm”, by S. Madgwick, A. Harrison, R. Vaidyanathan, 2011 IEEE International Conference on Rehabilitation Robotics, Rehab Week Zurich, ETH Zurich Science City, Switzerland, June 29-July 1, 2011.	37
20.	The error range of gyroscope heading estimates as compared with the magnetometer heading estimates. In the undisturbed case the gyroscope error range and the magnetometer heading overlap, and can be combined for a more accurate overall estimate. In the disturbed case the estimates do not overlap, and to reach agreement the gyroscope assumes a value that will increase its agreement with the magnetometer	39
21.	Flowchart explanation of the Heading Inclusion Filter (HIF) process for including or excluding magnetic data points.	43
22.	Flowchart of the Heading Estimate Comparison Test (HECT)	44
23.	Flowchart representation of the Magnetic Characteristics Test (MCT)	47
24.	Flowchart of Change in Acceleration Test (CIAT)	49
25.	An explanatory subset of the considered linear compensation functions.	52
26.	An example of the effect of unequal point distribution on data fitting.....	52
27.	This graph shows the accelerometer readings (blue) over a 60 second stationary period. The red bars bracket a range of 0.08 m/s^2 , equivalent to the chosen tolerance value for this test.	55
28.	This plot shows the XY norm of a test in which the only time spent in a MUA was the first and last five seconds of the test. Here the disturbed and undisturbed periods are readily identified, and the range can be set to contain the undisturbed data, shown by the pair of red lines. This user-selected range is then used to determine tolerances within the Heading Inclusion Filter (HIF).....	57
29.	This figure shows the wooden staircase used to sensor apparatus. This provided adequate separation from the ground so as to avoid the disturbance it produced.	58

30.	Here the testing apparatus in its two-link formulation can be seen. The pivoting linkage is mounted onto an aluminum base. The MARG sensors are placed on one end of the linkage, while the other end is used to manually pivot the linkage.....	60
31.	Here the testing apparatus in its single link formulation can be seen, as well as the pipe and bumper magnetic disturbance set up.....	60
32.	Here the four markers mounted onto the MARG can be seen, labeled for post-processing determination of orientation. The red and orange vectors were calculated, then averaged into a single vector (blue arrow).....	61
33.	Graphical Comparison of the XKF and MIMA heading angles under motion conditions beginning in a disturbed environment and moving into an undisturbed environment. The red bars mark the periods in which a disturbance was present. These graphs highlight the drift error the XKF must incur while adjusting to the undisturbed environment, as well as the MIMA's ability to avoid such drift and provide accurate data even in areas of magnetic interference.	64
34.	Graphical comparison of the XKF and MIMA (red and blue respectively) calculated heading angles. The motion began and ended in an undisturbed area, with the interim spent in a disturbed area. The red bars are marked the periods in which a disturbance was present.....	66
35.	The MIMA and XKF were graphically compared over a period 190 seconds. The first and last five seconds were spent in a MUA, and the other 180 seconds in a MDA.....	69
36.	The pre- bias compensation heading data (blue line), juxtaposed with the magnetic data that passed the filter (black stars). In order to apply bias compensation a linear function will be added to the yaw so as to achieve optimal agreement with the magnetic data.	73
37.	A zoomed in replica of the data shown in Figure 36, with the left hand side showing the initial period in the undisturbed area and the right hand side showing the final period in the undisturbed area. Together they represent the entirety of the time spent outside the disturbed area.....	73
38.	This graph compares resulting MIMA angles when the time spent in the undisturbed area was cropped such that the minimum time spent in it without compromising accuracy was achieved.....	74

39.	Here the crucial magnetic characteristics involved in filtering, the XY Norm and the Z Norm are shown. Here the measurement difference between disturbed and undisturbed periods can be seen.	74
40.	This figure shows the effect altering the tolerance values has on how many magnetic data point are excluded. The graph on the left shows how varying the allowable difference between gyroscope and magnetometer heading (HECT) affects the number of passing points. Here an epoch is 0.05 seconds. The graph on the right shows how varying the required change in acceleration between time points, in this case .0083 seconds, affects the number of passing points.	75
41.	This figure shows the effect of varying the acceptable range of magnetic characteristics on the number of passing points. The left graphs shows the effect related to the XY norm, while the right graph is relative to the Z norm.	75
A-1.	Magnetic field lines emanating from the northern magnetic pole and returning at the southern magnetic pole. 82	
A-2.	The magnetic declination angle is determined by the angle between the Earth's Magnetic and geographic Norths.	84
A-3.	The left image shows the three components (X Y Z) that make up a complete vector. The right image exhibits the magnetic dip angle, which is the angle at which the geomagnetic field penetrates the Earth's horizontal surface.	85
A-4.	The magnitude of the geomagnetic field as a function of global location. The contour interval is 1000 nanoTeslas (1 milliGauss = 100 nanoTeslas). Reprinted from "The US/UK World Magnetic Model for 2015-2020: Technical Report", by A. Chulliat, S. Macmillan, P. Alken, C. Beggan, M. Nair, B. Hamilton, A. Woods, V. Ridley, S. Maus, and A. Thomson, 2015, National Geophysical Data Center, NOAA.	86

A-5.	Hard iron distortion are caused by a material possessing its own magnetic field combining with the ambient, geomagnetic one. In this example, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the hard disturbance is applied, signified by the central rectangle, the magnetic field near it is drastically bent, and the field is concentrated on one side and largely removed from the other. This disturbance is only an example of how hard iron can distort a magnetic field; there are many other possible forms and resulting distortions. This image was developed in the program QuickField (Tera Analysis Ltd., 2013).	88
A-6.	Soft iron distortion acts as a conductor of magnetic field lines, distorting the ambient field by providing a path of less resistance relative to the surroundings. As before, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the soft iron is applied, the field lines are funneled through it, reducing field intensity on the sides of the disturbance and concentrating them at its ends. This image was developed in the program QuickField (Tera Analysis Ltd., 2013).	89
A-7.	Hard iron distortion adds a constant magnitude to every orientation, and by subtracting said constant the distortion can be removed.	92
A-8.	Soft iron distortion concentrations the field at some angles and reduces it at others, producing an elliptical shape. Calibration can be applied by mapping the ellipse into a circle.	93
A-9.	Disturbances generally have components of hard and soft iron distortion. The disturbance can be removed via hard iron calibration followed by soft iron calibration.	94
A-10.	Periodic calibration can be used to stabilize a data stream. This image is a simulation of how periodic calibration can be used to stabilize drift from a noisy signal. The red line is the noisy data that is periodically calibrated. The blue line is what the signal would look like without any noise, i.e., it is the true signal.	95
A-11.	A hard iron disturbance can have a large area of effect, disturbing the surrounding area in unpredictable ways. The chosen disturbance is simple for demonstration purposes, real-life disturbances can be more complicated. Image was developed in the program QuickField (Tera Analysis Ltd., 2013).	98

A-12. Introducing a magnetic shield into the environment can absorb and redirect the magnetic field lines emanating from the disturbance, leaving the surrounding area unaffected. Notice, however, that the field lines are now concentrated heavily at the ends of the shield. This area is now more heavily distorted than before, so it is imperative that this redirection occurs into the correct area. This image was developed in the program QuickField (Tera Analysis Ltd., 2013).....	99
A-13. Permeability of an object is proportional to its effect on the field. In this figure, the object on the left has a lower permeability value than the object on the right. The effect of this is evident in the increased number of field lines passing through the object, as well as greater curving of the field lines, in response to the higher permeability object. These images were developed in the program QuickField (Tera Analysis Ltd., 2013).	100
A-14. A Helmholtz coil is a made up of two solenoids and generates a uniform magnetic field within a portion of the area enclosed by the solenoid loops.	103
A-15. The environment chosen for static drift test was outside in a small field. It was chosen to avoid the magnetic disturbances found inside buildings, as well as to minimize the disturbance from any other ferromagnetic objects.	105
A-16. Two sensors were placed on the grass, and three were placed on the concrete.	105
A-17. The undisturbed static drift test shows that while all five sensors remained within an approximately two-degree range, their drift was random and chaotic.	106
A-18. This plot shows the magnetic norm as a function of shield layering and angular location of the MARG. Only high-permeability shields were used, and the numbers in the legends refer to thousandths of an inch, i.e., 4 under 6 means that the .004” shield was placed underneath (closer to the disturbance) the .006” shield in the layering configuration.	108

LIST OF EQUATIONS

Equations

1.	Magnetic Norm	17
2.	Output Orientation Estimate	37
3.	Objective Function.....	37
4.	Objective Function Gradient.....	37
5.	Local Leveling	44
6.	Magnetic Norm	46
7.	Magnetic XY Norm	46
8.	Magnetic Z Norm.....	47
9.	Magnetic Dip Angle.....	47
10.	Global Heading Model.....	49
11.	Least Squares Optimization	50
12.	Bias Error Model.....	50
13.	Magnetic Heading.....	62

LIST OF ABBREVIATIONS

Abbreviations

1. CIAT:.....Change in Acceleration Test
2. HIF:Heading Inclusion Filter
3. HECT:Heading Estimate Comparison Test
4. IMC:Inertial Motion Capture
5. KF:.....Kalman Filter
6. MIMA:.....Magnetic Interference Mitigation Algorithm
7. MOF:.....Madgwick Orientation Filter
8. MMOF:.....MARG Madgwick Orientation Filter
9. XKF:.....XSENS Kalman Filter
10. deg:.....degrees
11. mG:.....milliGauss
12. m/s²:.....meters-per-second-squared
13. sec:.....seconds

PREFACE

The layout of this document is such that it mirrors the path underwent to understand, investigate, and ultimately address the problem confronted by this thesis. It begins in Chapter 1 with an explanation of the problem of magnetic interference, and details the motivation for solving said problem. Chapter 2 explores the use of magnetic shielding as a mitigation method, and develops a process for shield application, which is subjected to experimental testing. The strengths and limitations of the shielding approach are discussed, and the insight gained is highlighted. Chapter 3 builds on the findings of Chapter 2, explaining the methodology behind the developed magnetic interference mitigation algorithm. The various factors, data streams, and relationships that make up the algorithm, as well as their interactions, are explained, discussed, and set up for experimentation. Chapter 4 presents the results from the experiments laid out in Chapter 3, and discusses the resulting insights, limitations, and future improvements. Following Chapter 4 is the Appendices Section, where the central concepts of inertial motion capture are discussed, and can be referenced should more background information be desired.

CHAPTER 1 : INTRODUCTION AND MOTIVATION

1.1 Introduction

In the world of inertial motion capture (IMC), error due to magnetic field interference is a persistent and widespread problem (Cuesta-Vargas, 2010; Lutters, 30-31 March, 2009; Fong & Chan, 2010). A major determinant in the results of IMC is the ambient magnetic field, and interference in this field causes errors in its interpretation. A real world example is the attempt of IMC inside an automobile. Certain characteristics inherent in the car, such as its metal frame and engine block, cause magnetic interference in and around the car. When IMC is attempted in this environment, this interference causes errors, generally in the form of inaccurate joint angles and inaccurate report of extremity location. Depending on the severity of the interference and the length of capture time, these errors can grow to levels that severely restrict the reliability and usefulness of the motion capture data. Any environment containing metal, such as inside a building with a metal frame or outside on an ATV, is likely affected by some degree of magnetic interference. Therefore decreasing the error caused by magnetic interference is a vital step in making IMC more useful and reliable in a variety of applications, ranging from medical to military.

1.1.1 MARG and Magnetometer Introduction

The cornerstone of IMC is the Magnetic, Angular Rate, and Gravity sensor array (MARG), which utilizes a variety of sensors to determine orientation. All MARGs by definition contain an accelerometer, a gyroscope, and a magnetometer, and other sensors, such as barometers or thermometers, can also be included. This thesis focuses on the magnetometer, which is used to stabilize the orientation about the vertical axis by calculating Magnetic North from components of the Earth's magnetic field. The quality of this stabilization depends upon correct interpretation of the magnetic field components, which is problematic when MARGs are used in areas affected by magnetic interference.

1.1.2 Cause of Magnetic Interference

Certain objects, known as sources of magnetic interference, can cause areas of distorted or disturbed magnetic field around them. This field distortion alters the components of the magnetic field, causing the magnetometer to incorrectly calculate the direction of Magnetic

North. An object is considered a source of magnetic interference because it satisfies at least one of the following two criteria. First, the object conducts magnetic field lines differently than air (soft iron disturbance). Second, the object possesses its own inherent magnetic field, which combines with the ambient magnetic field (hard iron disturbance). If either or both of these criteria are satisfied, the object will alter the properties of the magnetic field around it and cause errors in the magnetometer's calculation of Magnetic North (Figure 1).

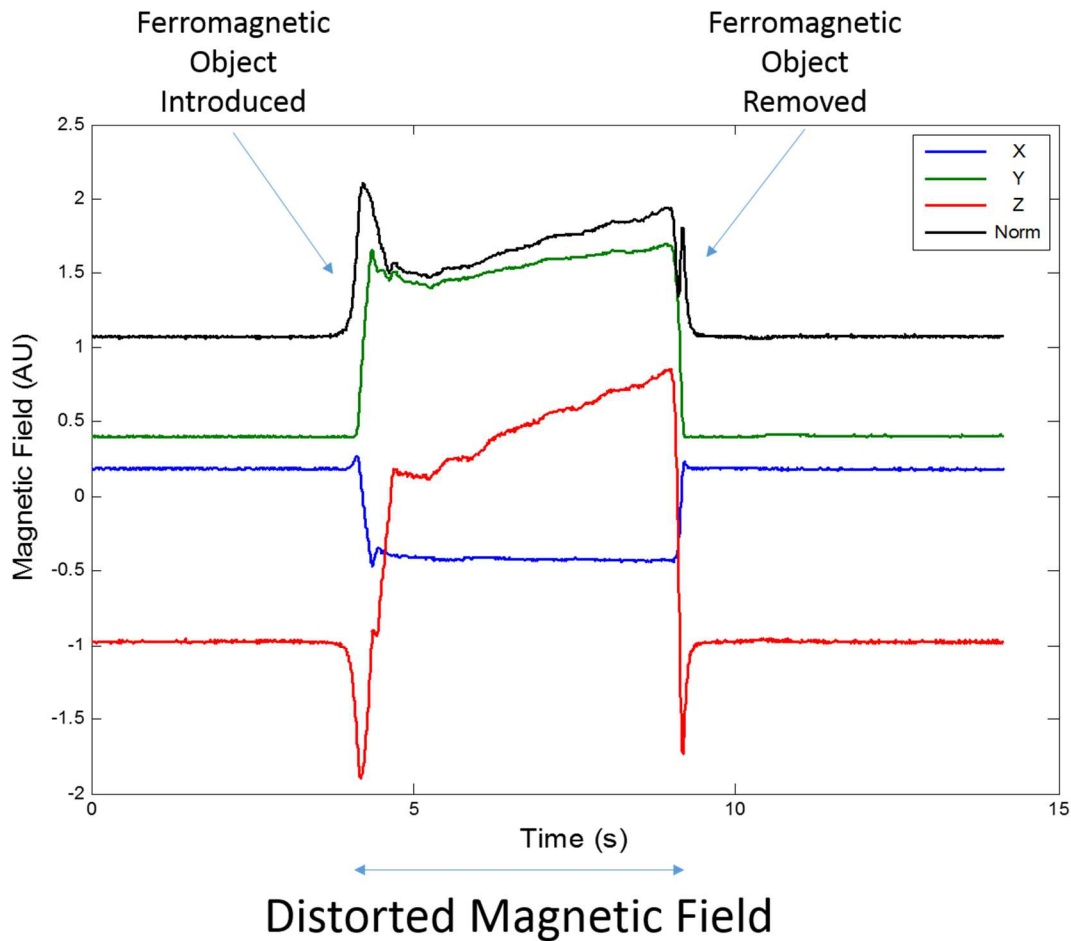


Figure 1: A source of magnetic interference was placed next to a sensor in an otherwise magnetically undisturbed field. The disturbance altered the X, Y, and Z components of the magnetic field, in turn causing alteration of the field's magnitude, as seen in the behavior of the norm.

The most common material that satisfies these criteria is iron, including in the form of steel. Ferromagnetic by nature, iron materials always satisfy the first criterion, and due to their potential for magnetization, often satisfy the second. In IMC environments iron is ubiquitous,

appearing in testing equipment, flooring, and wall supports.

The magnitude of magnetic interference caused by interference-producing materials increases with the amount of the object present, but it is also affected by other factors such as orientation or degree of magnetization. For more information on what causes magnetic interference, see the Appendices Section A.2.

1.1.3 Why Magnetic Interference is a Problem in Inertial Motion Capture

It has been stated that magnetic interference causes inaccurate measurement of Magnetic North, but the effect of such interference has yet to be related to the process of IMC. IMC combines MARG information with mathematical constraints and biomechanical modeling to generate motion data. MARG sensors are placed on the various body segments, and the data they record is processed to calculate the orientation of each body part. The orientation of each body part is then combined and processed to generate the overall posture of the subject. Magnetic interference adversely affects the calculated orientations, which in turn adversely affects the posture data. The applied mathematical constraints and biomechanical model work to mitigate the effect of magnetic interference, but non-negligible errors can remain, often at a level that makes the motion capture data unreliable.

1.2 Literature Review of Magnetic Interference Mitigation Methods

There have been many attempts to mitigate the problems associated with magnetic disturbance via filters, algorithms, and calibrations (Rotenberg, 2005; Roetenberg, 2007; Kok, 2012; Cutler, 2012; El-Gohary, 2013; Caruso, 2000; Konvalin, 2008b; Sabatini, 2006; Fong & Chan, 2010; Cuesta-Vargas, 2010). These attempts center on identifying the limitations of each sensor within the MARG and then intelligently pairing sensors together to overcome those limitations. Statistical analysis is applied to identify and compensate for noisy or erroneous sensor readings, and provide an objective and reliable method for combining sensor data.

Regarding filters and algorithms, Luinge utilized a Kalman Filter (KF) to combine gyroscope and accelerometer data. The accelerometer was able to detect the roll and pitch angles via the gravity vector, which was used to stabilize the integration drift of the gyroscope (Luinge, 2002). Expanding upon this, Rotenberg et al., 2005 and Rotenberg et al 2007, developed a specialized KF that fused magnetometer data with accelerometer data and gyroscope data, providing

stabilization about the yaw axis via the vector associated with Magnetic North. In order to deal with the problem of magnetic disturbance, the KF was formulated to apply less weight to the magnetometer data in the event of disturbance detection. This approach took advantage of the short-term stability about the yaw axis when only a gyroscope and accelerometer were fused. Integration-drift accumulates over time (on the order of a fraction of a degree per second), and ignoring the magnetometer data for short time periods when a magnetic disturbance is detected leads to only a very small drift accumulation. This sensor fusion formulation mitigates the effect of magnetic disturbances by accepting a small integration drift instead of the far larger error associated with inaccurate heading readings.

This approach worked well to eliminate drift, though its effectiveness declined with longer periods of disturbance (testing disturbances last between 5 and 15s), and the filter is less effective if the disturbance is introduced slowly (the disturbance is introduced to the sensor's environment by slowly increasing its proximity to the sensor rather than immediately being placed next to the sensor). (Roetenberg, 2007; Rotenberg, 2005). This is because drift accumulates the longer yaw stabilization is unavailable, and it is difficult to identify magnetic disturbances if the change in magnetic field components is small, as in the case of slowly introducing the disturbance.

Sabatini developed a quaternion based extended KF, including the effects of sensor bias and noise covariance (Sabatini, 2006). Prediction and modeling of inherent sensor errors such as bias and noise allowed for better integration drift mitigation, and this improvement was achieved without the use of magnetometer data.

In a similar approach, El-Gohary avoided magnetic interference by foregoing the use of magnetometers entirely, instead relying on implementation of Newton-Euler equations, noise modeling, and zero velocity updates (El-Gohary, 2013). This approach replaces the magnetometer with analytic and statistical relationship between sensors, and uses this information to stabilize the integration drift. Furthermore, noise modeling and zero velocity updates work to reduce the effect of sensor errors and mitigate position drift.

In the field of pedestrian navigation, Bird and Arden used zero velocity updates and a Kalman Filter that selectively excluded magnetic data producing headings that did not agree with the

predictions from the inertial sensors (Bird & Arden, 2011). In this formulation, zero velocity updates worked to mitigate drift from gyroscope bias, and the magnetic filter decreased error due to magnetic interference. The overall effect was to increase the accuracy for pedestrian position data in GPS-denied environments. Faulkner et. al developed a similar approach using an Extended Kalman Filter, using a magnetic filter and zero velocity updates to decrease drift in horizontal position estimates (Faulkner, Alwood, Taylor, & Bohlin, 2010). The results showed decreased horizontal drift error in pedestrian tracking in multi-story buildings and near large vehicles.

Regarding the calibration methods, the company Honeywell noted the problems that magnetic interference caused and published a paper detailing a method for two-dimensional compass calibration via ellipse-to-circle mapping for use in a low-cost navigation system (Caruso, 2000). The application of this paper was an e-compass, which contains a magnetometer and accelerometer. This publication was targeted toward aircraft navigation and included disturbance calibration and a model to mitigate the effect of heading error introduced due to pitch and roll angles. The calibration process utilized the magnetometer's constant position relative to the disturbance sources (it is mounted to the aircraft) in order to characterize the present disturbances. Once characterized, a mapping function could be applied to all incoming data to remove the data distortion due to the disturbance. The company MEMSense published two similar papers expounding upon the same method for use with its microelectromechanical sensors (Konvalin, 2008a; Konvalin, 2008b). This paper was published several years after the Honeywell paper, and MEMS technology was more advanced. As such, the focus of this paper was the application of calibration methods to such sensors.

Kok et al., 2012 expanded this mapping method to three dimensions. It is difficult and impractical to rotate an object through every three-dimensional orientation, a problem Kok et al. addressed by developing an optimization scheme that could predict a mapping function based on only a small number of orientations (Kok, 2012). However, as with the previous calibrations, this method is only valid at the exact location of calibration, or if the IMU's position is fixed relative to all sources of magnetic disturbance one is attempting to remove.

Afzal et al. developed a method of disturbance mitigation that relied more on hardware advancements than software (Afzal, Renaudin, & Lachapelle, September 2010). Twelve

magnetometers were combined into a single array composed of two orthogonal planar arrays. This configuration was designed to gather magnetometer data from the same point at a large variety of different orientations. The incoming data from all 12 magnetometers is surveyed, and the least disturbed data stream is chosen to be fed into a KF for heading determination (Afzal, Renaudin, & Lachapelle, September 2010). This method offers a form of real-time, adaptive calibration, but it is targeted more towards the less precise requirements of pedestrian navigation. Furthermore, the size of the sensor array in its current form would be prohibitive for mounting onto the human body.

The results of this literature survey show that magnetic interference is a persistent problem in IMC that many attempts have been made to address it. Systems capable of dealing with ephemeral disturbances have been developed, but their effectiveness degrades when the disturbance is prolonged or consistent. Other systems that do not rely on any magnetic data have been developed, but the concern with these systems is that without the magnetometer there is limited ability to correct for drift error. These investigations have decreased the error associated with magnetic interference, but a method for reliable full-body IMC in a magnetically disturbed environment has not yet been achieved.

1.3 Approach of This Investigation

A review of the literature and testing of the XSENS MVN Biomech motion capture system revealed that an effective and reliable method for magnetic interference mitigation has not yet been developed, and that such a development would lead to significant improvement in IMC accuracy. This thesis explores two approaches to mitigation method, one that physically alters the nature of the interference and one that digitally alters the effect of the interference. First, mitigation through the application of magnetic shielding is investigated. Second, an algorithm for identifying and correcting MARG data gathered in a magnetically disturbed environment is developed and tested. The shielding investigation revealed three key insights, which constitute the foundation upon which the digital approach is based.

1. What causes magnetic interference and how it can be identified.
2. How MARG sensors interpret and are affected by magnetic interference.
3. Why magnetic information is required for orientation calculation, and how it is utilized.

CHAPTER 2 : MAGNETIC SHIELDING INVESTIGATION

2.1 Introduction

The concept of magnetic shielding is similar to any other concept involving shielding, in that the goal is introduce an object into an environment such that portions of the environment become separated. In the case of magnetic shielding, the goal is to introduce an object, a magnetic shield, into an environment containing magnetic interference and a magnetic field sensor in such a way that the interference is blocked from interacting with the sensor. However, magnetic field lines cannot be blocked, only redirected, meaning that for a magnetic shield to be successful, it must redirect the interference around or away from the area it is intended to shield.

Magnetic shields function via redirection of magnetic interference, meaning that the shield itself is warping the magnetic field and is therefore a form of magnetic interference. However, it is a controlled form of interference, and can be used to counteract the effects of other, less controlled interference sources. Therefore, magnetic shielding applied to inertial motion capture (IMC) is a method of strategically placing interference in a way that other interference is mitigated. The efficacy of a magnetic shield is governed by a variety of parameters, such as permeability, size, shape, and placement, and each situation requires an evaluation of each of these parameters to achieve optimum shielding. For more information on magnetic shielding and the parameters that affect it, see the Appendices Section A.5.

2.2 Experimental Methodology

To investigate the effect of shielding on magnetic interference, several experiments were performed with the purpose of exploring how the MARG sensors calculated heading and the effect of shielding on magnetic fields. Of the five shielding experiments discussed in this chapter, the first three explored the properties of the MARG sensors regarding heading calculation, and the last two explored the effect of shielding of the magnetic field. Heading refers to the yaw angle measured with respect to the direction of Magnetic North, as determined by the measured components of the magnetic field. Each experiment was performed using XSENS' state-of-the-art wireless MARG called the Motion Tracker Wireless (MTw).

The MARG-focused experiments were performed in the following manner. Experiment one investigated MARG heading calculation and compared it to the orientation calculation

unaffected by magnetic interference (roll and pitch) to determine a relative magnitude of heading instability. Experiment two investigated the effects of an introduced magnetic interference on the heading calculations of a stationary MARG, while experiment three investigated the effects of magnetic interference on the heading calculations of a non-stationary MARG. Together these experiments led to a better understanding of heading calculation and instability, and how that instability manifests in the IMC environment.

The shielding-focused were performed in the following manner. Experiment one (four overall) investigated a key characteristic of magnetic fields (the norm) and the differences in its behavior while in undisturbed and disturbed environments. Experiment two (five overall) investigated the effect of shield type and material on the norm. Together these experiments led to a better understanding of magnetic field characteristics and how different types of shielding affected those characteristics.

In addition to the five experiments listed here, two other experiments were performed that investigated inter-sensor variability and the effects of magnetic shield layering. These experiments can be found under the Appendices Section A.9.

2.3 Baseline Testing through MARG Experimentation

2.3.1 Drift Comparison of the Roll, Pitch, and Yaw Axes

2.3.1.1 Objective

The purpose of this experiment was to compare the error and variability about the vertical axis (heading/yaw) with the error and variability about the two horizontal axes (roll and pitch). Roll and pitch can be obtained without the use of magnetic data, whereas heading requires magnetic data. By comparing the difference in error and variability of these different axes, the relative instability of the magnetically derived components could be determined.

2.3.1.2 Experimental Setup

A wireless MARG (Xsens MTw) was placed in a location several feet from any known sources of magnetic disturbance, initialized, and then used to record data in the form of roll, pitch, and yaw. The chosen location was in the center of the lab room, on top of an aluminum chair (non-ferromagnetic) mounted on top of the lab's vibratory shaker table (Figure 2).




Figure 2: Two XSENS MTws (MARG sensor arrays) were placed on a six-degree-of-freedom vibratory table. One sensor was mounted on the magnetically undisturbed aluminum seat, and the other was mounted on the magnetically disturbed steel baseplate.

2.3.1.3 Data Collection and Processing

Data was collected for approximately three minutes and then exported from the XSENS software MT Manager 4.2.1 as a text file containing the raw inertial and magnetic data and the calculated orientation in the form of roll, pitch, and heading. In order to compare the relative instabilities, each set of angles were shifted such that they all began at zero.

2.3.1.4 Results

The roll and pitch data remained very close to zero, oscillating between ± 0.25 degrees (Figure 3). The heading data was more variable, steadily drifting for the first 50 seconds at a rate of approximately 0.1 deg/s. After 65 seconds, the drift became less uniform, though continued to slowly increase, growing from five degrees to 5.66 degrees 85 seconds later at the test

conclusion.

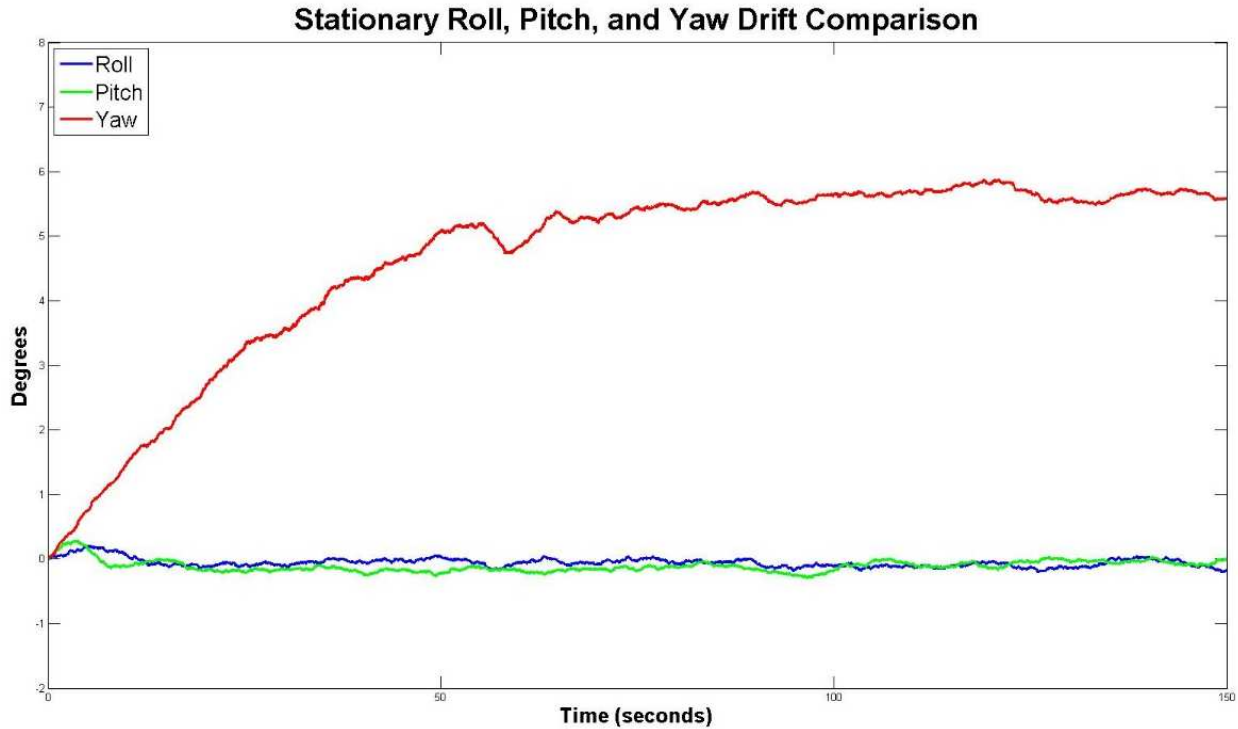


Figure 3: This plot shows relative magnitude of roll, pitch, and heading variability as calculated by the MTw. After 150 seconds, the absolute error in roll and pitch are less than .25 degrees, whereas heading is more than 5.5 degrees off from its initial orientation.

2.3.1.5 Discussion

The results of this experiment show that MARG-calculated heading is less stable than either roll or pitch. Though the MTw was removed from known sources of interference, the magnetic field of the test environment is not ideal. The magnetic field inside buildings can be volatile and prone to fluctuations. These uncertainties make it difficult for the MTw to settle on a heading value, producing the variability seen in Figure 3.

2.3.2 Stationary Drift Comparison – Undisturbed vs Disturbed

2.3.2.1 Objective

The purpose of this experiment was to extend the previous test by examining the effect of introduced interference on the heading variability. The MTw is known to combine the data from its various sensors to achieve an optimal orientation estimate (XSENS, 2013), and seeing how it

behaves under controlled introduction of interference will offer insight into its governing methodology.

2.3.2.2 Experimental Setup

Within the lab, two locations were chosen for MARG placement such that one sensed a distorted magnetic field and the other an undistorted field. One sensor was placed in an aluminum seat in the room's center, in an undistorted field, and the second sensor placed on the steel base below the chair, in a distorted field (Figure 2). Hence forth magnetic interference, distorted field, and disturbance will be used interchangeably to mean the same thing. After placement the sensors were initialized and set to record orientation data.

2.3.2.3 Data Collection and Processing

Data was collected for approximately 70 seconds and then exported as in Section 2.3.1.3. The roll, pitch, and heading values were again shifted to a common starting position and then plotted (Figure 4).

2.3.2.4 Results

The heading determined by the undisturbed MTw drifted 1.25 degrees in the first 18 seconds before settling into an oscillatory pattern with an amplitude of 0.25 degrees that decayed over time. The heading determined by the disturbed MTw followed a more linear pattern and drifted down 0.25 degrees over the duration of the test (75 seconds).

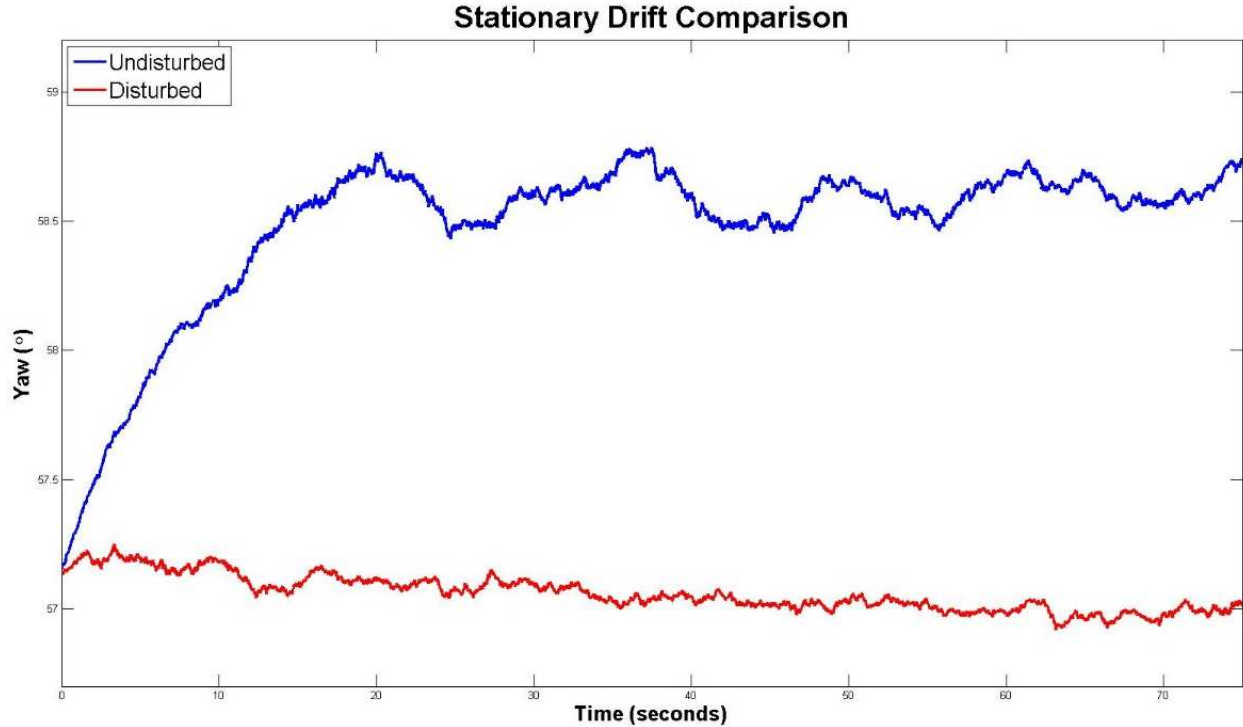


Figure 4: The drift error in the undistorted field sensor is greater than that of the distorted field sensor. The implication of these results is that a distorted field can work to stabilize the heading calculations, even if about the wrong value. Here the red represents the disturbed yaw values and the blue represents the undisturbed yaw values.

2.3.2.5 Discussion

The results of this experiment were at first counterintuitive. It was expected that the disturbed MTw would be more prone to drift, not less prone, than the undisturbed MTw. Understanding these results requires some knowledge of how the MTw calculates heading, based on which an explanation is offered. The MTw determines heading by combining data from its several sensors based on the confidence it has associated with each data stream, and proceeds in this fashion until equilibrium is reached. The steel baseplate constituted a source of magnetic interference, but it also presented a strong, consistent magnetic field. In contrast the undisturbed MTw was forced to interpret a weaker and more variable magnetic field. The MTw stabilized around the strong, disturbed field, but required time to reach equilibrium in the undisturbed field. It is important to note that a faster rate of stabilization does not necessarily indicate a more accurate heading. The distorted field sensor stabilized more quickly due to the interference, but not necessarily to the correct heading value.

2.3.3 Oscillatory Drift Comparison – Undisturbed vs Disturbed

2.3.3.1 Objective

The purpose of this experiment was to see if the results of the previous experiment applied to MARGs in motion as well as stationary ones, as they are often not stationary during IMC.

2.3.3.2 Experimental Setup

The MTws were placed as in Section 2.3.1.2, on a six-degree-of-freedom shaker table, one (undisturbed) in an aluminum seat in the room's center, and one (disturbed) on the steel base below the chair. The MTws were initialized, and the table was set to oscillate about the vertical axis (heading/yaw) at a frequency of one hertz for 200 seconds (Figure 5 and Figure 6).




Figure 5: Two MTs were mounted on a vibratory shaker table, one magnetically disturbed and the other magnetically undisturbed. In this test the shaker table was made to move in an oscillatory motion about the vertical axis (Figure 6) at a frequency of one hertz.




Figure 6: Explanation of the oscillatory motion of the shaker table. The table pivoted about a central, vertical axis in a repetitive, alternating motion. This is taken from the MOOG website (MOOG is the name of the vibratory table).

2.3.3.3 Data Collection and Processing

The shaker table was made to oscillate at one hertz and minimal displacement for 8 cycles of about 20 seconds each, with approximately 5 seconds in between cycles. Data was collected for approximately 200 seconds and then exported as before. The MTw-calculated heading values were shifted to a common starting position and then plotted (Figure 7 and Figure 8).

2.3.3.4 Results

The oscillatory motion was captured well in both situations, but the drift was present in both trials. As with the static test, the disturbed sensor stabilized more quickly, reaching a steady state after one cycle (~20 seconds). Its overall drift was two degrees. The undisturbed sensor drifted over the first two cycles by approximately 0.6 degrees, and then tapered off in the subsequent cycles for a total drift of 5.7 degrees.

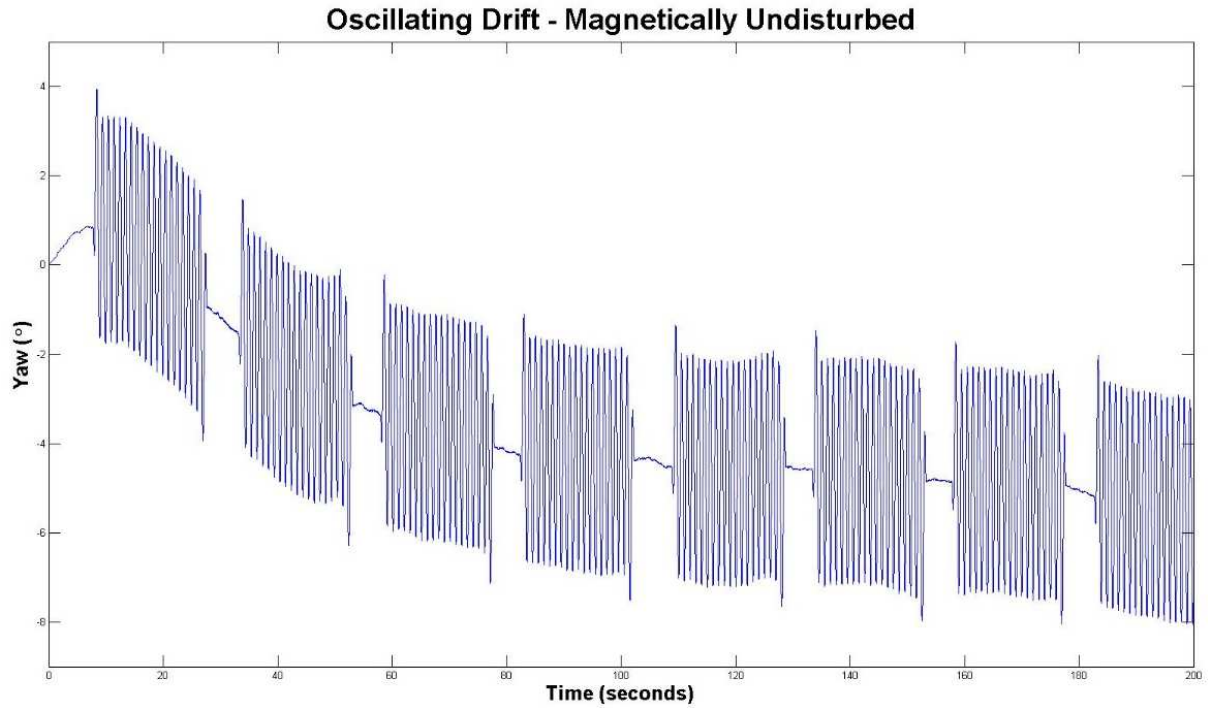


Figure 7: There was a drift error of approximately five degrees in the undisturbed sensor, taking over two minutes to stabilize, and exhibiting minor drift after that.

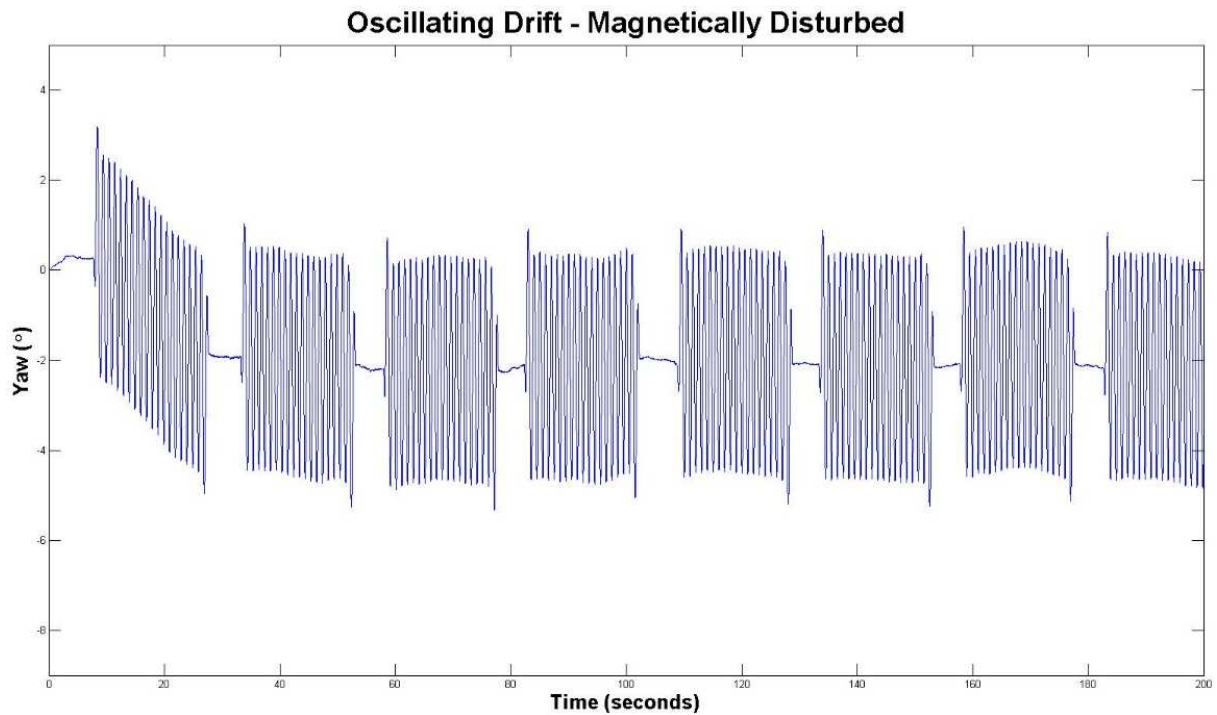


Figure 8: There was a drift error of approximately two degrees in the disturbed sensor, taking approximately thirty seconds to stabilize.

2.3.3.5 Discussion

As with the results of the previous experiment in Section 2.3.2.4, the disturbed sensor stabilized more quickly. The problems in heading detection for the undisturbed sensor were similar compared with the stationary test. It was concluded that due to the algorithm by which the MTw combined its various sensors data the oscillatory motion had little effect relative to the stationary case. The magnetometer is used to stabilize the heading value, not track it with high accuracy in the short term. The other sensors of the MTw needed to alter their processing to account for the motion, but the processing of the magnetic field remained largely unchanged.

2.3.4 Conclusions

- Heading instability is roughly an order of magnitude greater than that about the roll and pitch axes, though this can be affected by the properties of the present magnetic field.
- MARGs in disturbed areas can stabilize faster than those in undisturbed areas, deemed a result of the strong, consistent magnetic field offered by the disturbances.
- The previous point holds for stationary and oscillatory scenarios; motion seemed to have little to no effect on the stabilization.
- Stabilization, especially in a disturbed area, does not necessarily mean the calculated heading/yaw is accurate with respect to Magnetic North

2.3.5 Magnetic Norm and Frequency Filtering

2.3.5.1 Objective

The purpose of this experiment was to explore the use of the magnetic norm as a method for disturbance detection. The norm is an easily calculated metric of a magnetic field's magnitude and it would be insightful to determine its response to magnetic interference.

2.3.5.2 Experimental Setup

There was no setup for this experiment. The data required for analysis was drawn from the oscillatory drift test discussed in Section 2.3.3 and 2.3.3.4.

2.3.5.3 Data Collection and Processing

The raw unaltered magnetometer data, recorded in X, Y, and Z components, was extracted from the oscillatory drift test data. The first step in calculation of the norm accomplished by summing

the squares of each magnetometer component and taking the square root of the sum (Equation 1).

Equation 1: Magnetic Norm

$$M_N = \sqrt{M_x^2 + M_y^2 + M_z^2}$$

After calculating the norm, the component data was run through a lowpass filter that removed all frequencies larger than the motion frequency (one hertz). This filtering was done to improve visualization, but it is not necessarily useful for improving orientation accuracy because the MTw already accounts for this. The filtered norms (undisturbed and disturbed) were compared, and the filtered X, Y, and Z components (undisturbed and disturbed) were compared. The XYZ components were measured and reported relative to the coordinate system of the MTw (Figure 9). However, due to the orientation of the MTws, their coordinate systems were such that the Z axis was vertical and the X and Y axes were horizontal.




Figure 9: MTw and its body coordinate system. Image from the XSENS MTw User Manual (XSENS, 2013).

2.3.5.4 Results

The differences in filtered norms can be seen in Figure 10. Of the two differences, the first was the average magnitude. The undisturbed norm had an average of 1.01 AU whereas the disturbed

norm had an average of .91 AU. AU stands for arbitrary unit such that one AU is equivalent to the magnetic field's magnitude at calibration (XSENS, 2013). The second major difference was that disturbed norm exhibited oscillatory behavior, and the undisturbed norm did not. This signifies that the norm around the steel baseplate varied with motion whereas the norm around the shaker seat did not.

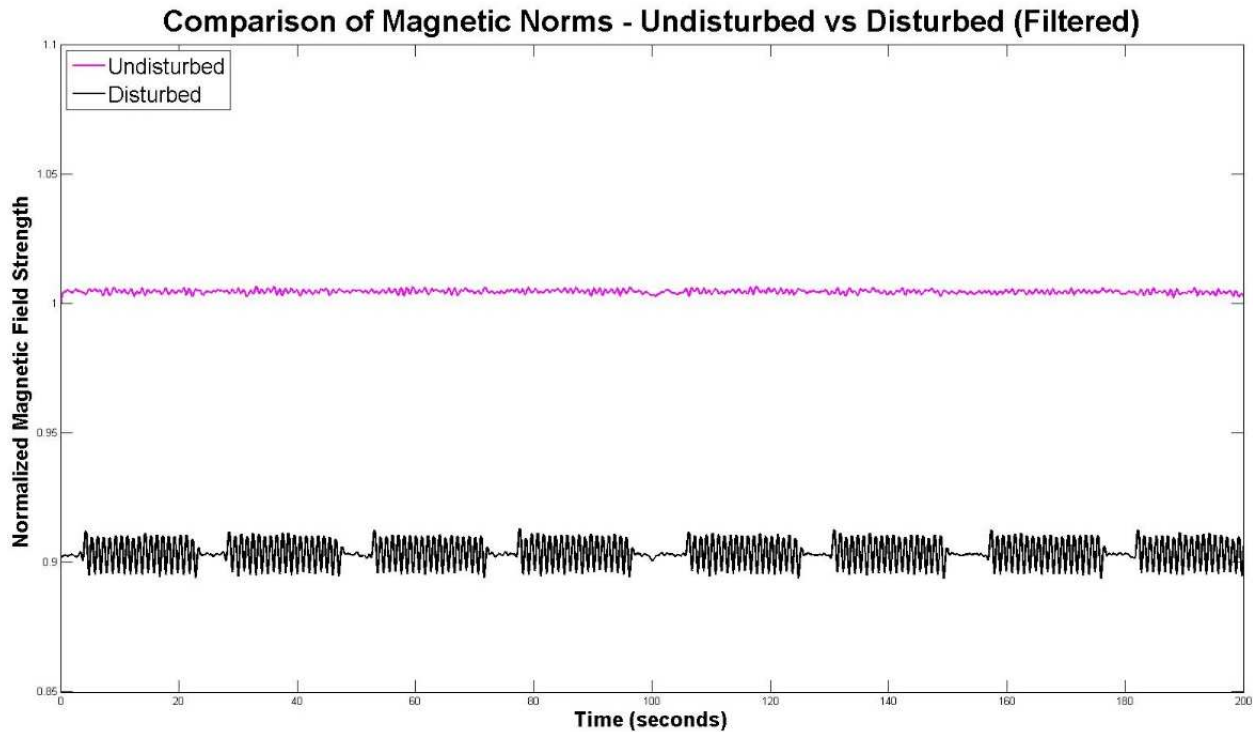


Figure 10: After filtering, the plots for both the disturbed and undisturbed cases became clearer, and the differences in the disturbed and undisturbed norms were more readily identified.

Application of disturbance had a clear effect on the magnetometer component data, shifting the relative component magnitudes, as well as changing which component was affected by the oscillatory motion. In the undisturbed data, the X and Y components exhibited oscillatory behavior, whereas in the disturbed data the X and Z components exhibited oscillatory behavior (Figure 11 and Figure 12).

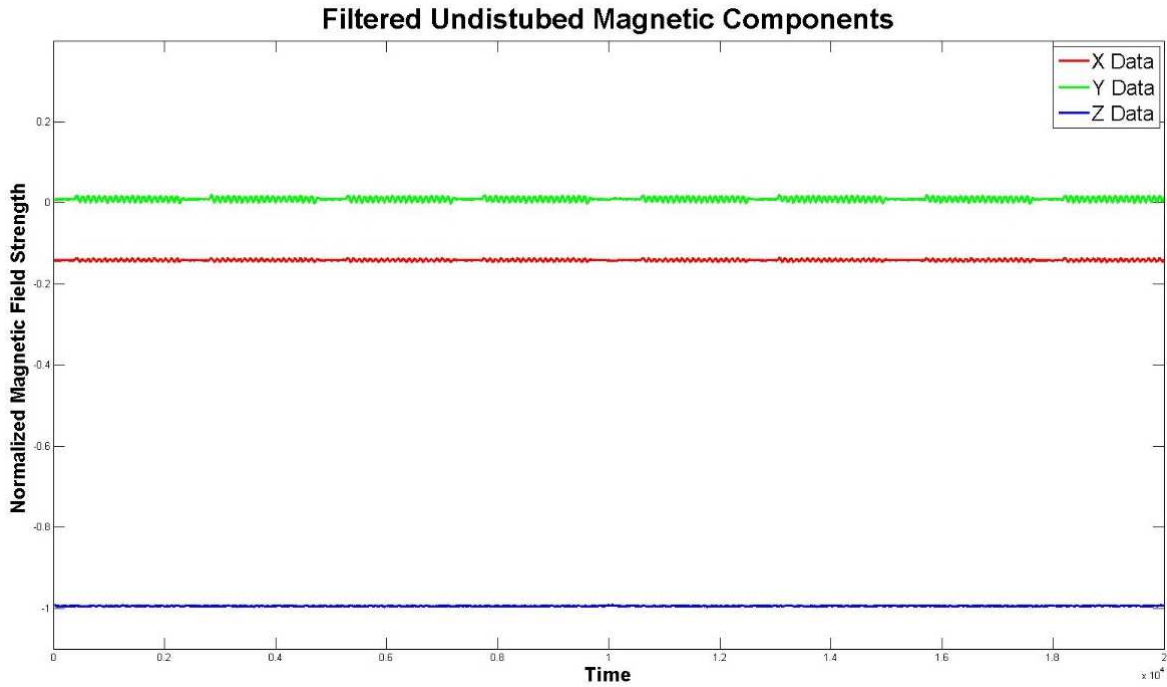


Figure 11: In the undisturbed case, the X and Y components were oscillatory and the Z component was not. This is to be expected given the rotation about the Z axis.

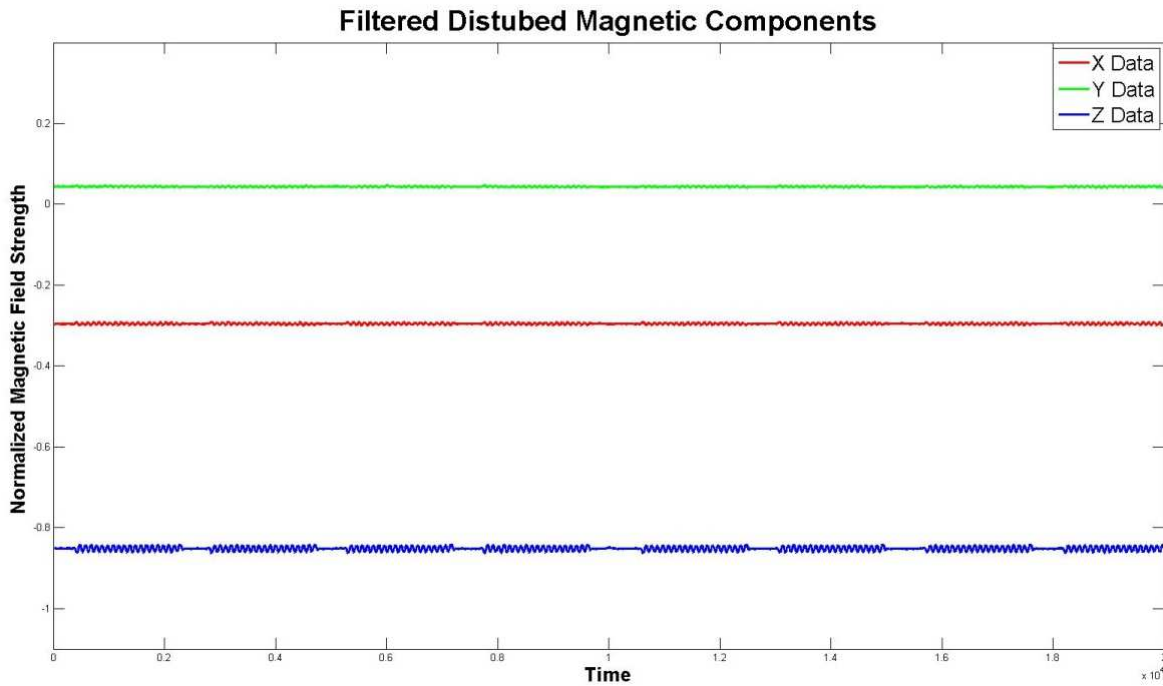


Figure 12: In the disturbed case, the X and Z components were oscillatory and the Y component was not. This deviation from the undisturbed behavior showcases the effect magnetic distortion has on the components of the magnetic field.

2.3.5.5 Discussion

The conclusion of this study was that the norm can be used as a method of disturbance detection by examining its average value and its deviation from that average value during motion. Sources of interference warp the magnetic field around them, altering the norm. This is supported by the differing value of the average norm in the undisturbed and disturbed cases for this experiment. Furthermore the field warping occurs in an irregular manner, such that the magnitude of the surrounding field changes with the measurement location. Therefore if the norm varies with motion the area is can be deemed disturbed. This is supported by the current experiment, as the motion of the disturbed MTw resulted in norms of varying value but the motion of the undisturbed MTw did not.

One could attempt to determine the existence of magnetic interference based on the norm's average value, but this is not as straight forward as it seems. A magnetic field can be considered undisturbed for the purposes of IMC if it is homogenous, as field homogeneity allows for consistent heading calculation regardless of the field's magnitude. However, due to the variability in environments the magnetic field can be homogenous while presenting a different magnitude than another undisturbed environment. It would be difficult to determine an ideal value of the norm with which to compare measurements, making consideration of the norm's value alone a difficult method of detecting interference. It is more reliable to ascertain the existence of magnetic interference based on variation of the norm.

2.4 Magnetic Shielding to Alter the Norm

2.4.1 Objective

The purpose of this experiment was to explore the effect of varying types of magnetic shielding on magnetic interference mitigation by examining how the norm around a disturbance source is altered by the application of shielding material.

2.4.2 Experimental Setup

This experiment involved the passing of a MARG over a shielded source of magnetic interference. The source was a magnetized ferromagnetic metal bar, comprising both a hard and soft iron (combination) disturbance. In order to ensure reproducible motion over the disturbed area, a pivoting robotic arm was employed, and the MARG was mounted on its end effector.

The robotic arm followed a circular motion, starting at -150 degrees and ending at 150 degrees, with 0 degrees corresponding to the location directly above the disturbance (Figure 13 and Figure 14). This allowed for smooth, repeatable motion over the disturbance, and for observation of effects near the shield edges. In order that shield lay flatter over the disturbance source, the ferromagnetic bar interference source was placed within a block of foam, creating a flat area for the shield to lay on (Figure 15).




Figure 13: An overhead view of the robotic arm and shielding experimental setup.




Figure 14: A side view of the robotic arm and shielding experimental setup.




Figure 15: An overhead view of the robotic arm and shielding experimental setup with the magnetic shield removed to show the disturbance source.

The shielding was accomplished by laying a magnetic shield over the disturbance, such that the shield was located between the MARG and the disturbance source. Five different kinds of magnetic shields were tested: three high-permeability shields of thickness .04", .06", and .01", a .01" medium-permeability shield, and a .007" low-permeability shield. The exact permeability of each shield type varies with factors such as magnetic field strength and shield geometry, but the listed maximum permeability for each of the shielding alloys was 7000, 150,000, and 300,000 for the low, medium, and high permeability alloys respectively (MuShield Magnetic Shielding, 1996-2015). Permeability is the measure of a material's ability to form a magnetic field within itself, and is reported here as a unit-less value describing the material's permeability relative to that of a vacuum. In order to more easily discern trends, double layers of each shield type were used in testing so as to produce exaggerated responses.

2.4.3 Data Collection and Processing

The double-layered shields were placed over the disturbance source, which the sensor was swung over by the robotic arm for a minimum MTw-to-disturbance distance of approximately two inches. The test was performed with no shielding material present, as well as with each of the five shield types. Data was recorded for all angles through which the robotic arm passed (-150 to 150 degrees). The raw data was extracted and the magnetometer components identified so as to compute the norm. The norm for each case was then plotted.

2.4.4 Results

Each shield altered the magnetic norm in a different way (Figure 16). With the exception of the medium-permeability shield, which inverted and exaggerated the distortion, each shield had a smoothing effect on the magnetic distortion, reducing the overall difference from the chosen ideal case of norm equal to one (Figure 17). It is not possible to know that one is truly the ideal case, as it varies from location to location, but XSENS uses norm equal to one as their ideal case, and the same was done here. The more important characteristic to notice is the variability in the norm, as an ideal case would have a norm that is completely static.

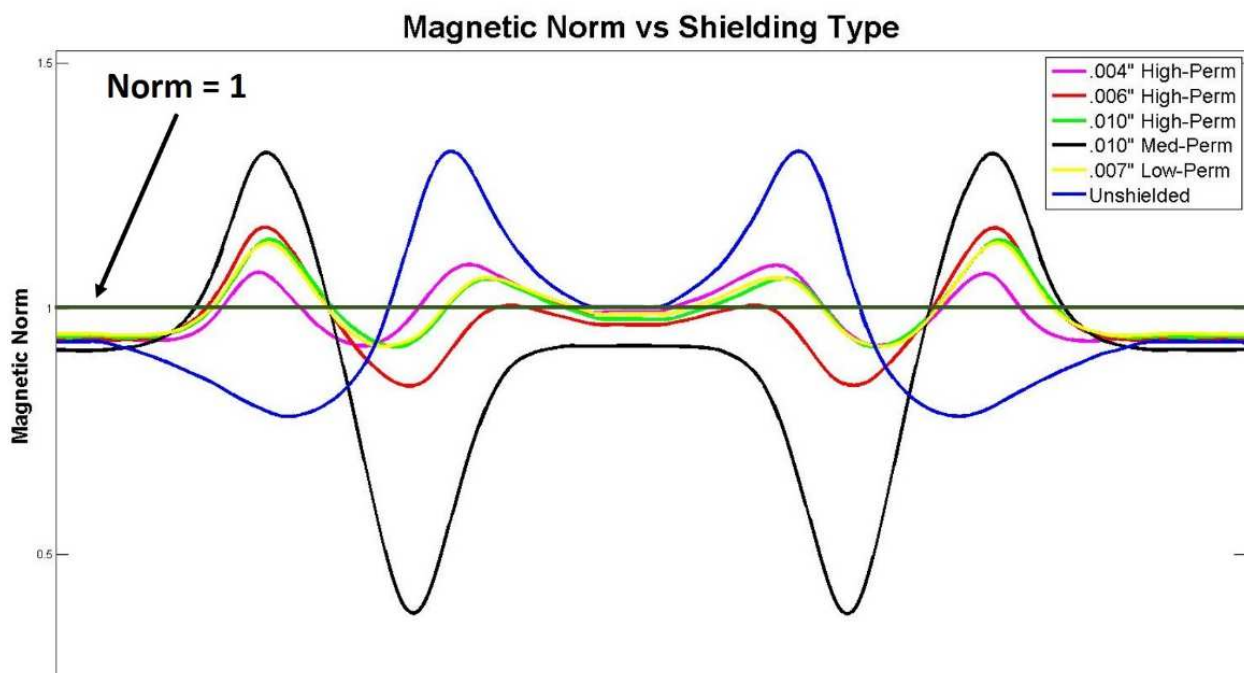


Figure 16: The magnetic norm as a function of shield type and angular location. It should be noted that regardless of whether the magnetic shield lowered or increased the norm, the local field was invariably altered.

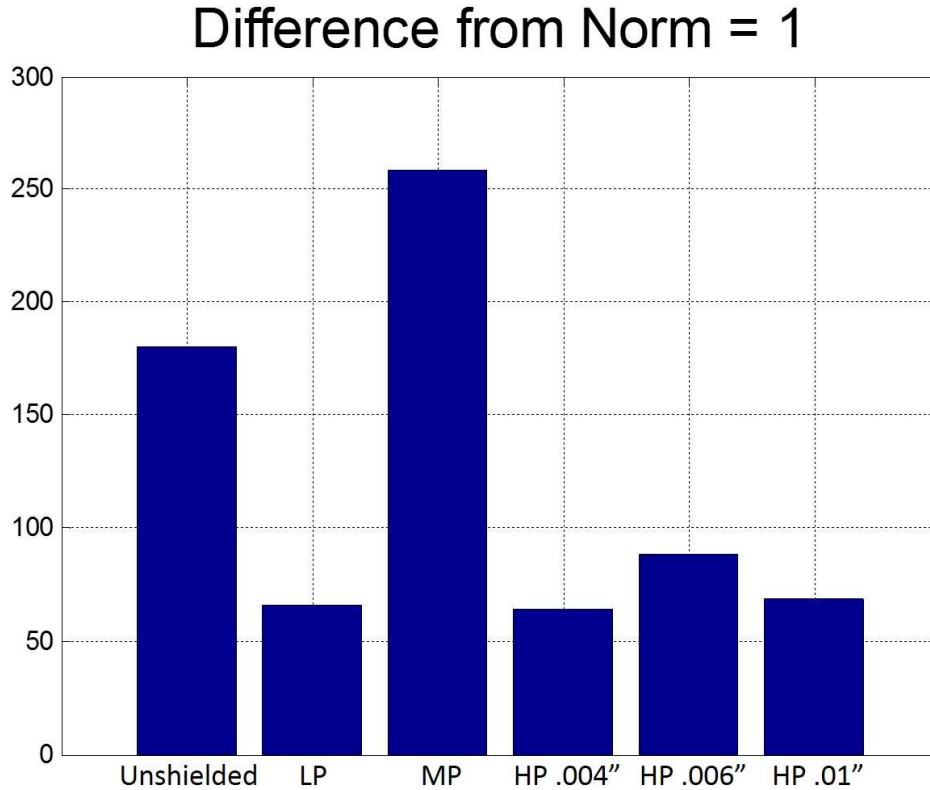


Figure 17: Here the labels mean the following: LP – Low Permeability, MP – Medium Permeability, HP – High Permeability. The numbers following HP denote the thickness of the shield in inches. The effects of each of the five different shielding approaches are shown alongside the unshielded case, quantified by comparing their relative overall difference from the idealized case of norm equal to one.

2.4.5 Discussion

The clearest conclusion from the results is that magnetic shielding has a non-negligible effect on magnetic field, evidenced by the threefold decrease in difference from the norm equal to one case when comparing the unshielded case to the .004" high-permeability shield case.

Furthermore, the said effect is strongly dependent on the permeability level and thickness of the shield. The low-permeability and high-permeability shields lowered the difference from the ideal case by an average of 60.3%, whereas the medium-permeability case increased it by 43.3%, an example of magnetic shielding exacerbating rather than attenuating the interference.

When evaluating the effect of shield thickness, it can be seen that increased thickness does not correlate with increased efficacy; in fact, the thinnest shield was the most effective. This is reflective of the inherent methodology behind magnetic shielding. In order to mitigate an

uncontrolled disturbance, a controlled one, in the form of a magnetic shield, is introduced. The introduced disturbance must alter the magnetic field enough to mitigate the uncontrolled disturbance, but not so much as to introduce a new one of its own. These results lead to the recommendation of using magnetic shielding only when necessary, and in the smallest amount possible.

When evaluating the effect of shield permeability, a conclusion similar to that for thickness is followed. The low-permeability and high-permeability shields improved the situation, but the medium-permeability shield worsened it. This again has to do with hitting the sweet spot in disturbance introduction. Permeability is a measure of the shield's magnetic "conductivity". It would stand to reason that relative to each other the low-permeability shield removed less distortion while introducing less, and the high-permeability shield removed more distortion but also introduced more. Conversely, the medium-permeability shield introduced more distortion and removed less of it, resulting in an overall increase. Overall, these results show that both shield type and thickness are important when identifying the most effective shielding material, and that each case must be done on a custom basis.

2.5 Experiment-based Conclusions

- The norm can be used as an indicator of magnetic interference based on both its magnitude and its variation during motion.
- The amount of change the norm undergoes is related to the amount of interference, with greater changes indicating more interference.
- Magnetic shielding can be used to mitigate magnetic interference by reducing the amount of change the norm undergoes during motion.
- Choosing the correct shield type can lead to a reduction in interference, but choosing the incorrect shield type can increase the interference.
- Magnetic field data are used to determine global orientation about the yaw axis and to prevent error from accumulating about that axis due to sensor error.

2.6 Applications in Inertial Motion Capture

Below is the recommended method of applying the findings and conclusions of this investigation to improve the accuracy of an IMC test. The experiments of this chapter were done using

XSENS hardware and software, but these recommendations can be applied to any IMC System.

2.6.1 Field Mapping

Within buildings, particularly ones with steel reinforcements, the geomagnetic field is highly disturbed (Cermakova, 2005), so it is likely there is some interference in most testing areas. This interference is specific to each location, so the magnetic field must be investigated. Once the location has been chosen, a tri-axial gaussmeter or magnetometer can be used to characterize the field. Before performing IMC, characterize the chosen location's magnetic field by surveying the entire location with a gaussmeter or magnetometer, noting areas in which the norm shows a lot of variability. In areas where variability was noted, a more thorough survey must be conducted so as to determine the cause of the disturbance. The magnitude of variability will grow in size as distance from the disturbances decreases.

2.6.2 Shield Application

Once the magnetic field has been characterized, shielding can be applied. Any areas where deviations were detected are candidates for shield placement. Magnetic interference can cause two separate types of disturbance, though most disturbances are a combination of the two types. The first type is a soft iron disturbance, in which the disturbance-causing item alters the path of the magnetic field lines around it, but does not produce its own magnetic field. An example of this disturbance type is a magnetic shield. The second type is a hard iron disturbance, in which the disturbance-causing item produces its own magnetic field that is superimposed on the ambient field. An example of this disturbance type is a conventional magnet, such as one would place on a refrigerator. Further information on disturbance types and magnetic fields can be found in the Appendices Sections A.2 and A.1 respectively.

If the disturbance is located in a wall or floor, a small, flat, high-permeability shield should be placed over it. This will concentrate the disturbance within the shield, reducing its range of effect. The shield is recommended to be flat so as to not impede the subject's motion. Application of the shield will not remove the magnetic interference entirely, rather it reduce its range of effect.

If the disturbance is due to a fixture in the environment, such as a treadmill, a more custom shielding approach is required. In the case of a treadmill (and most such fixtures), the

recommended shielding approach would be to cover the fixture in shielding material so as to contain the disturbance as much as possible. In this case a high permeability shielding material would likely introduce more interference than it would mitigate, so a low permeability material such as steel foil is recommended. Steel foil has the added benefit of a higher saturation value than high permeability shielding, meaning it can absorb more field lines.

In the case of the treadmill, wrap as much of it as possible in steel foil. The steel foil's low permeability will ensure it introduces little interference, and its high saturation level will allow it to contain more of the interference produced by the treadmill. Due to the large number of variables in this situation, a variety of shielding methods must be tried to obtain the best results for the specific motion capture scenario desired. A combination of shield perm abilities and geometries may be required to attain optimal mitigation.

2.7 Conclusion

Magnetic interference in the IMC environment is a recognized and as-of-yet unsolved problem. Methods of mitigating this interference focus mainly on filtering and processing of MARG data, instead of addressing the problem at its source via magnetic shielding. As such the focus of this investigation was to explore the efficacy of magnetic shielding in reducing the effect of magnetic interference on the orientation calculations of MARG sensor arrays.

Based on the results of this investigation, magnetic shielding can mitigate interference. The experiments presented in this chapter revealed that the magnetic norm can be a powerful tool in identifying magnetic interference, and that magnetic shielding can be used to lower said interference. This knowledge allows users to quantifiably find undisturbed locations for testing, identify an interference source as well as its magnitude, and a novel method for mitigating the interference associated with unavoidable disturbed locations. However, magnetic shielding can also exacerbate interference, and the effects of a shielding solution must be investigated with the norm to ensure their efficacy. Shielding material complicates the magnetic field, and without testing it is impossible to ensure it is reducing rather than increasing interference.

CHAPTER 3 : ALGORITHM METHODOLOGY

3.1 Introduction

The magnetic shielding investigation explored how MARGs calculate orientation, how they are affected by magnetic interference, and methods for identifying such interference. MARGs calculate orientation primarily based on integration of angular rate readings from the gyroscope, and stabilization relative to a global coordinate system is provided by the accelerometer and magnetometer. Magnetic interference, identified by any alteration to the ambient, homogenous magnetic field in either direction or magnitude, changes the global heading calculated by the magnetic field, leading to stabilization about the wrong direction. The stabilization is required to prevent the accumulation of gyroscopic integration error, which occurs in a manner that can be predicted and modeled. This insight constitutes the foundation of the digital approach to interference mitigation developed in this thesis.

There are existing methods of identifying magnetic interference, and the orientation filters (typically a Kalman Filter formulation) utilized by most MARGs are limited in their ability to compensate for errors due to distorted magnetic fields. Based on the results of the shielding investigation, two clear avenues of improvement emerge. The first avenue involves identifying and correcting magnetometer readings before they are processed by the orientation filter. The second avenue involves altering or replacing the orientation filter so that it is more capable of compensating for disturbance-related errors. Either of these approaches, or some combination of the two, would mitigate the error caused by distorted magnetic fields.

Successful application of the first approach requires a quantifiable and objective way to reliably identify the nature of the sensed magnetic field as disturbed or undisturbed. Magnetic fields tend to exist in a spectrum between disturbed and undisturbed rather than clearly falling under a single category, complicating such identification. Furthermore, each environment affects the magnetic field differently, and what would be considered disturbed in one environment may not be in another environment. This requires that the identification criteria be adaptable. An example of inter-environment variability can be seen when comparing indoor and outdoor IMC. Outdoor IMC will likely encounter a magnetic field very similar in nature and magnitude to the geomagnetic field. Indoor IMC will likely encounter a field of different nature due to the

warping effect of building structures. The geomagnetic field lines can be altered by their passage through the building exterior, but remain homogenous inside the building. Relative to the geomagnetic field, the indoor field is distorted, but because it remains homogenous it is not distorted in a way that will cause errors in IMC.

Successful application of the second approach requires a re-evaluation of the method by which the MARG's various sensor data is combined, known as the sensor fusion process. In an undisturbed environment the optimal balance between short-term and long-term stabilization offered by a recursive orientation filter such as the KF is ideal. However, these real-time orientation filters are unable to account for things not incorporated in their model, and their recursive nature limits adaptability and ability to correct itself. An example of this limitation occurs when IMC is initialized in a distorted environment. The orientation filter bases and stabilizes its heading relative to that distortion. After moving to an environment free of distortion, the filter cannot automatically correct itself. The calculated heading must continuously drift until it reaches the correct heading. Despite this limitation, it is difficult to avoid using these real-time orientation filters, specifically the Kalman Filter, due to the need to fuse noisy data from multiple sensors.

Examining the pros and cons of each avenue, it was originally intended to pursue the first approach. It was reasoned that filtering the magnetometer data would be a more practical and rewarding approach than modifying the orientation filter. However, initial testing showed that such filtering was not able to overcome the problems caused by the orientation filter's lack of adaptability. Based on this an algorithm was developed that processed the magnetic and inertial data separately, using a real-time orientation filter to fuse the inertial data and a separate filter for the magnetometer to exclude disturbed magnetic data, before combining the results into a single output. This approach maintained much of the stability and accuracy offered by real-time sensor fusion, and added an adaptive method of including magnetometer data.

3.2 Methodology

3.2.1 Magnetic Interference Mitigation Algorithm (MIMA) Overview

The magnetic interference mitigation algorithm (MIMA) was designed to mitigate MARG error in disturbed magnetic fields, and it differs from other such algorithms mainly in how it processes

magnetometer data. Examination of the traditional KF algorithm revealed that while the magnetometer was introduced to reduce error, it also has the potential to increase error. Furthermore, the error it reduces is primarily due to sensor bias in the gyroscope. This error can be modeled, as despite its nonlinear and unpredictable nature it can be approximated as a constant linear drift with a magnitude falling within a known range. In contrast the error introduced by magnetic interference cannot, at least with any practicality, be modeled as it is more erratic and unpredictable, and is not constrained the same as sensor error is. Moreover, the error introduced by magnetic interference can be greater than the error that would accumulate had the magnetometer been ignored.

Based on this examination, it was made the driving goal of the MIMA that magnetic field data be processed and included in a way that would either reduce error or have no effect. This was achieved by removing the magnetic data from the initial processing of the MARG data, leaving only inertial data. The general structure of the algorithm reflecting this can be seen in Figure 18. The error in this inertial data was modeled, and the magnetic data was filtered to remove all potential disturbed points. The initial inertial estimate, an error model, and the filtered magnetic data were then combined to correct the inertial data. This formulation prevents the magnetometer from introducing error, as all disturbed points are ignored in the correction process. In the case of no available undisturbed magnetic data, no correction is made. This resulting data will inevitably have drift errors, but the error will not be increased because the measurement was made in a disturbed field.

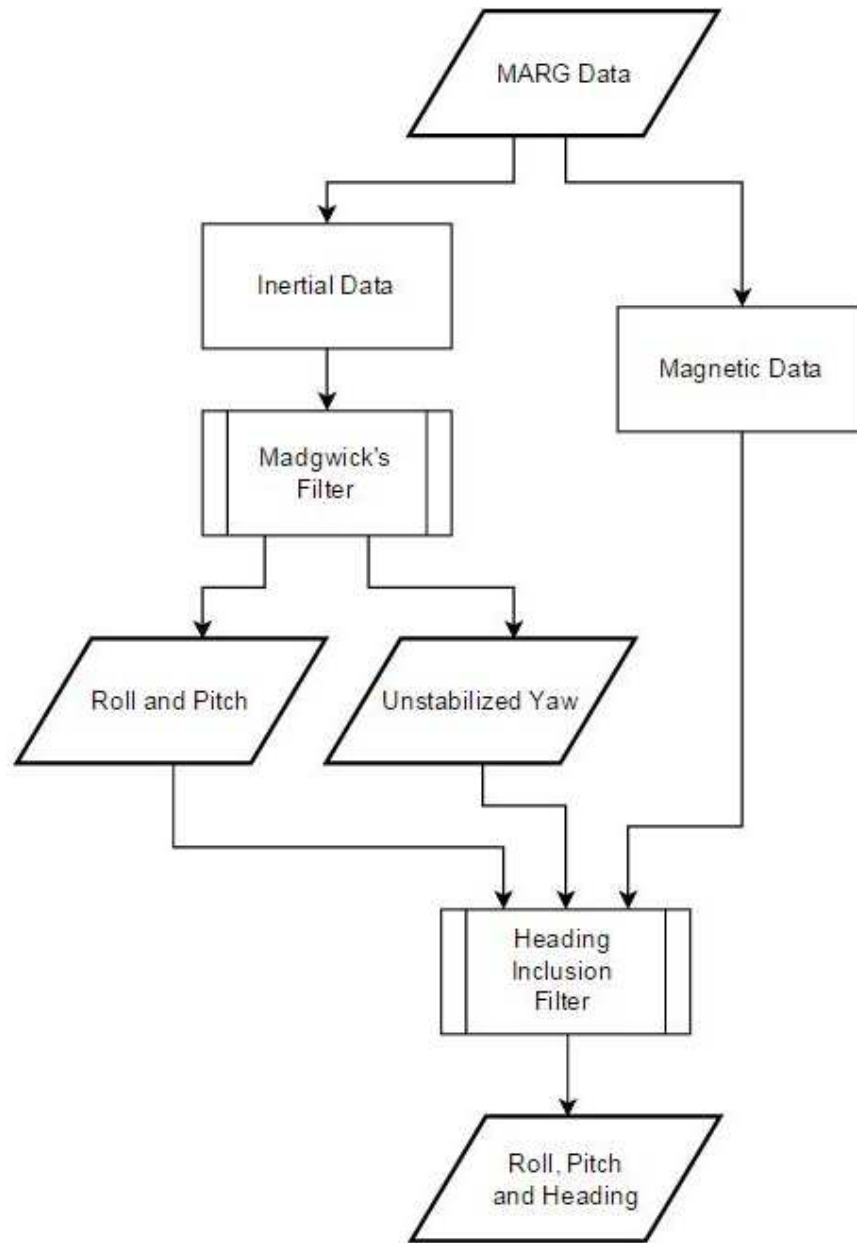


Figure 18: The overall structure of the algorithm developed in this work. The traditional approach is to feed the complete MARG data into a KF. In order to make the process more resistant to magnetic interference, here the inertial and magnetic data are separated such that only the inertial data is Kalman filtered, and the magnetic data is accounted for afterwards via the Heading Inclusion Filter.

A major difference between the developed MIMA and traditional KFs is that the MIMA is not in its current formulation recursive. Each time step processed by a KF is a function of the previous estimated data point and the current measured data point. The KF combines these values to generate an estimate of the current data point, making each data point a function of every

previous data point. This approach allows for real-time data processing, but does not allow future data points to retroactively affect current data points. The current formulation of the MIMA is not recursive, nor is it real-time, allowing estimates to be affected by measurements made both before and after it. In a disturbed environment this ability can be exploited to drastically improve data estimates, as will be shown in the following sections.

The transition from real-time and recursive to post-processing has the obvious drawback of not allowing the digitized motion to be viewed as it is recorded. However, the inertial portion of the MIMA is recursive, allowing that an estimate of the digitized motion could be viewed in real-time. This estimate would not be affected by magnetic interference, but would also lack any stabilization from the magnetometer, and could therefore accumulate error. Another drawback of not being recursive is increased computational effort. In a recursive formulation only the previous and current time steps are processed, whereas all time steps are processed in a non-recursive approach. This can lead to substantial computational requirement depending on the number of sensors used. However, the MIMA processes and incorporates only the magnetic data non-recursively, so the increase in computational effort is substantially reduced from if the entire algorithm was non-recursive. Finally, it should be noted that future development of a recursive MIMA formulation is possible, though its method of future data inclusion will need to be updated to allow this.

The MIMA evolved to include three distinct sub-processes:

- Determination of roll and pitch via sensor fusion
- Heading determination from magnetic and fused data
- Drift correction through compensation of gyroscope bias

Traditional sensor fusion of MARG data records highly accurate orientation via integration of the gyroscope readings. Accelerometer and magnetometer data is used to determine orientation relative to the global and mitigate the errors that grow over time due to error caused by the inherent bias of the gyroscope. The accelerometer-gyroscope pair alone is capable of accurate and stable determination of both roll and pitch angles, constituting an inertial measurement unit (IMU). The IMU, however, suffers from an inability to determine heading (yaw measured with respect to Magnetic North), or to stabilize the gyroscopic integration error associated with

heading. However, because the integration error is time-dependent, examination over short time periods diminishes this error. By this reasoning, the shorter the time period considered, the smaller the accumulated error. It is difficult to correctly state how quickly gyroscopic integration error grows due to the many affecting variables, but generally it is on the order of degrees per minute and somewhat linear in nature.

3.2.1.1 Determination of Roll and Pitch by Sensor Fusion

The foundation of the MIMA is an orientation filter applied to inertial data. Magnetic data is concurrently recorded but ignored in the sensor fusion process. The orientation filter fuses the accelerometer and gyroscope data to generate roll, pitch and yaw orientations. Of these orientations, the roll and pitch angles are in the global coordinate system, stabilized by the accelerometer's measurement of the gravity vector. The yaw angle is the result gyroscopic integration without correction of the related error, and is not in the global coordinate system. This follows the typical process for calculating orientation from inertial data.

3.2.1.2 Heading Determination via Inclusion of Magnetometer Data

The next step is to convert the yaw into heading by transforming it into the global coordinate system. Such a transformation requires heading information, which the magnetometer can provide. However, the magnetometer does not always offer correct heading data due to magnetic interference. In order to transform the yaw angles to global they must be combined with the undisturbed magnetic heading calculations, requiring a method of identifying disturbance. This method is termed the Heading Inclusion Filter (HIF) and operates by developing several criteria magnetic field data must pass in order to be considered undisturbed. The criteria are based upon inertial, environmental, and geomagnetic relationships so as to make use of all available data and thoroughly test the magnetic measurements.

Once the undisturbed portions of the magnetic data are identified, they are used to convert the yaw into heading by determining the optimal fit of the undisturbed magnetic data with the yaw angles. After optimal fitting, the yaw angles have been converted into heading, but lack stabilization of the gyroscopic integration error.

3.2.1.3 Bias Correction

Gyroscopic integration error is caused by the gyroscope's inherent bias, a systematic form of

error. The bias manifests by introducing a consistent error into the gyroscope's readings of angular velocity, and with each instance of integration the error in predicted orientation grows by an amount equal to the bias error multiplied by the time step of integration.

Gyroscopic bias is complex and difficult to model. Factors including temperature, pressure, and force on the sensor are capable of altering the bias value, leading MARG sensors such as the XSENS MTw to include thermometers and barometers for bias compensation. However, experimentation showed that under simple motion the bias could be modeled as a constant with little error. A constant bias will cause a linear drift in the integrated gyroscope data, meaning that addition of the correct linear function would effectively remove the error due to the bias. The general process of this bias compensation involves solving for the linear function that when added to the heading data generates the least amount of error with respect to the undisturbed magnetic data points.

3.2.1.4 Conclusion

In sum, the MIMA removes the magnetometer from the sensor fusion process in order to utilize the strengths of a real-time orientation filter to generate roll, pitch, and yaw angles. The removal of the magnetometer makes the resulting orientations immune to magnetic disturbance, but removes the ability for heading determination or stabilization. To compensate for these shortcomings, a magnetic filter is used to identify undisturbed magnetic data, which is combined with the yaw angle via a least-squares optimization to convert the yaw angles into heading. The next compensatory step is to stabilize the heading, accomplished by determining the linear function that when added to heading generates the least amount of error relative to the undisturbed magnetic data points. The end result is an algorithm with much of the stability of a traditional sensor fusion approach that can adaptively compensate for magnetic interference.

3.2.2 Sensor Fusion

Determination of MARG orientation via sensor fusion produces accurate results given undisturbed conditions. Of the three main sensors, the gyroscope, accelerometer, and magnetometer, the magnetometer is the most susceptible to disturbance. The accelerometer is used to identify the gravity vector, therefore it is susceptible to disturbance from dynamic motion, but such disturbance is more controlled and predictable than that encountered by the

magnetometer.

In the MARG array the purpose of the magnetometer is to stabilize the yaw estimate made by the gyroscope. Removing the magnetometer downgrades the MARG to an inertial measurement unit (IMU), and immunizes the fused data against magnetic disturbance. An IMU with sensor fusion processing is capable of producing stabilized roll, stabilized pitch, and un-stabilized, relative yaw. The MIMA begins by extracting the gyroscope and accelerometer data from the MARG and processing it in this fashion. The focus of the MIMA is on magnetic interference, therefore existing, open source code was used to complete the fusion of the inertial data (Madgwick, MADGWICKAHRS Implementation of Madgwick's IMU and AHRS algorithms, 2010; Madgwick, Harison, & Vaidyanathan, Estimation of IMU and MARG orientation using a gradient descent algorithm, 2011). The filter developed by Madgwick et al. constitutes the MIMA's starting point, with all other processes functioning by combining the calculated Euler Angles with raw MARG data. The basis for this filter will be discussed first, followed by a more conceptual examinations of its performance and limitations.

3.2.2.1 Madgwick's Orientation Filter

It has become standard practice for orientation calculations to be completed by a KF, and for good reason. KFs offer a statistically optimal fusion of the sensors data. However, the KF approach also has its draw backs, primarily its complexity of implementation, high sample rate requirements, and large computational burden. In comparison, the orientation filter developed by Madgwick et. al does not require high sample rates, is more computationally efficient, and is more straightforward in its implementation (Madgwick, Harison, & Vaidyanathan, Estimation of IMU and MARG orientation using a gradient descent algorithm, 2011).

Madgwick's Orientation Filter (MOF) uses a quaternion representation of orientation, and a gradient descent optimization approach to calculate the direction of the gyroscope's error as a quaternion derivative. The use of quaternions, as opposed to Euler Angles, avoids the singularities associated with gimbal lock, as well as reducing computational requirements. Furthermore, the gradient descent algorithm is formulated with a step size greater than or equal to the rate of orientation change so that only one iteration is required per time step.

For clarity of explanation the meanings of each to the terms used in MOF are listed in Table 1.

The leading superscripts carry the following meaning: ${}^S_E q$ describes the orientation of frame E (earth) relative to frame S (sensor), and ${}^E d$ describes a vector in frame E (earth).

Table 1: Terms used in the filter by Madgwick et al. and their meanings.

Term	Meaning
q	Orientation (Quaternion)
t	Time Step
${}^S \hat{a}_t$	Normalized Accelerometer Measurement
${}^S \omega$	Gyroscope Measurement
${}^E \hat{d}$	Gravity Vector
${}^S_E \hat{q}_{est}$	Orientation Estimate
${}^S_E \dot{q}_\omega$	Gyroscope Calculated Rate of Orientation Change
${}^S_E \dot{q}_{est}$	Estimated Rate of Orientation Change
f_g	Objective Function
J_g	Objective Function Jacobian
∇f	Objective Function Gradient
β	Magnitude of Gyroscope Error
\otimes	Quaternion Product

The final output of the filter is calculated by Equation 2, where the objective function and its gradient are described by Equation 3 and Equation 4 respectively. The objective function of Equation 3 uses the quaternion conjugate to convert the gravity vector from the earth frame to the sensor frame, from which the accelerometer readings are then subtracted. In an ideal situation the results of this subtraction would be zero, but in reality error ensures a nonzero difference. Minimization of this error is thus the goal, which is accomplished by calculating the gradient per Equation 4. The orientation error is a result of gyroscope inaccuracy, thus normalizing the gradient (Equation 2), provides the direction of said gyroscope error.

Equation 2 shows that the final orientation estimate is equal to the previous orientation estimate, plus the numerical integration of the gyroscope readings, and minus the numerical integration of the normalized gradient multiplied by some constant β . As the gradient is normalized, this constant is responsible for determining the magnitude with which gyroscope is corrected by the accelerometer. The choice of β value conceptually represents how much trust is placed in the gyroscope; a small β value makes only small corrections to the gyroscope estimate, whereas a

large β value makes only large corrections. A β value of 0.1 was used for the experiments of this thesis. A block diagram of this process can be seen in Figure 19, taken from the paper by Madgwick et al.

Equation 2: Output Orientation Estimate

$${}^S\hat{q}_{est,t} = {}^S\hat{q}_{est,t-1} + {}^S\dot{q}_{\omega,t}\Delta t - \beta \frac{\nabla f}{\|\nabla f\|} \Delta t$$

Equation 3: Objective Function

$$f({}^S\hat{q}_k, {}^E\hat{g}, {}^S\hat{a}) = {}^S\hat{q}_k^* \otimes {}^E\hat{g} \otimes {}^S\hat{q}_k - {}^S\hat{a}$$

Equation 4: Objective Function Gradient

$$\nabla f({}^S\hat{q}_k, {}^E\hat{g}, {}^S\hat{a}) = J^T ({}^S\hat{q}_k, {}^E\hat{g} + {}^S\hat{q}_k, {}^E\hat{d}) * f({}^S\hat{q}_k, {}^E\hat{g}, {}^S\hat{a})$$




Figure 19: Block Diagram of the orientation filter from Madgwick et al. Reprinted from “Estimation of IMU and MARG orientation using a gradient descent algorithm”, by S. Madgwick, A. Harrison, R. Vaidyanathan, 2011 IEEE International Conference on Rehabilitation Robotics, Rehab Week Zurich, ETH Zurich Science City, Switzerland, June 29-July 1, 2011.

3.2.2.2 *Effect of Magnetic Interference on Sensor Fusion*

Traditional approaches to sensor fusion such as the Kalman Filter and Madgwick's Filter described in Section 3.2.2.1 work well when fusing data with predictable error sources like sensor or process noise, but encounter when unpredictable error sources are introduced, such as magnetic interference. Upon encountering magnetic interference, the magnetometer's heading estimate changes in a manner different from what is expected based on the gyroscope's measurements, which causes erratic filter behavior if encountered for an extended period of time.

The traditional sensor fusion process begins with orientation estimation by the gyroscope. The gyroscope is known to be susceptible to error generating a range of acceptable values the true yaw can fall under. The error is determined based on the sensor's innate mechanical and electrical characteristics, such as noise or bias, which can be modeled and iteratively predicted (as in the KF), or assigned a constant corrective value (as in MOF). In the case of magnetic interference, the heading as calculated by the magnetometer falls outside the acceptable range determined by the gyroscope (Figure 20). In order to correct this inconsistency, each time step the orientation filter will assume the error associated with the gyroscope's heading estimate is in the direction opposite the magnetic heading, and the heading estimate will iteratively drift until it reaches agreement with the magnetic heading.

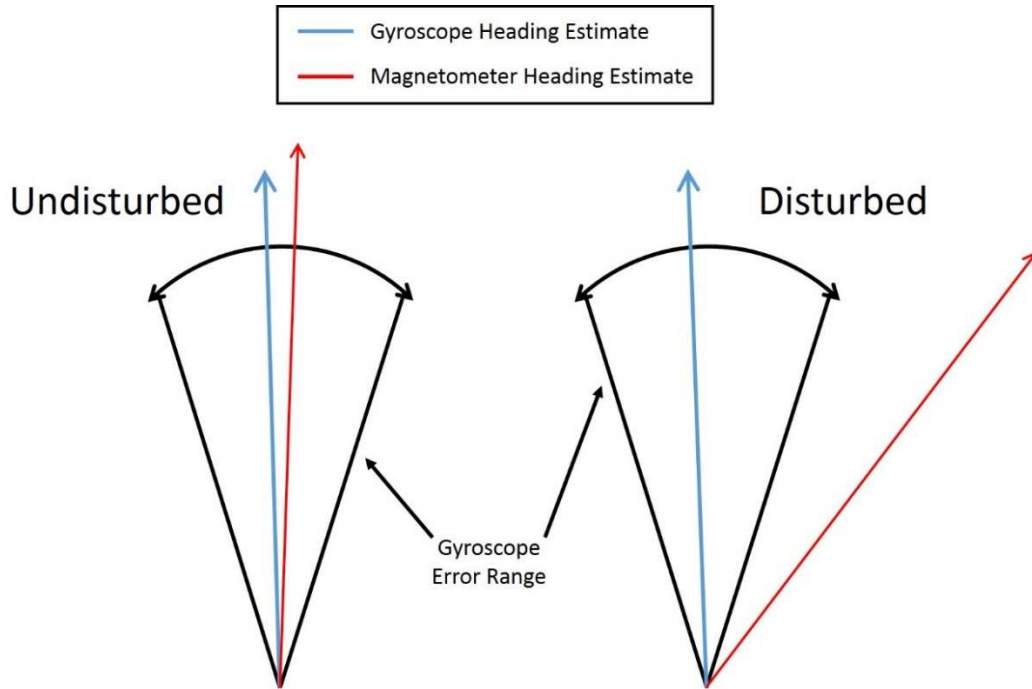


Figure 20: The error range of gyroscope heading estimates as compared with the magnetometer heading estimates. In the undisturbed case the gyroscope error range and the magnetometer heading overlap, and can be combined for a more accurate overall estimate. In the disturbed case the estimates do not overlap, and to reach agreement the gyroscope assumes a value that will increase its agreement with the magnetometer

In IMC ephemeral disturbances are often encountered, leading to a sensor fusion formulation in which magnetometer data is ignored during the disturbed period. In the case of ephemeral disturbances, generally on the order of a few seconds, this approach produces favorable results. When the magnetometer is ignored, the heading estimate may deviate from gyroscope's heading, the potential error increasing the longer the magnetometer is ignored.

In the case of a prolonged period in a disturbed field, the potential gyroscope error can grow beyond a preset threshold determined by the filter's architects. Once this threshold is crossed, the magnetometer data is re-included in the heading estimate, even if it is still disturbed. The heading estimate will then drift until it reaches agreement with the magnetically determined heading. Such an approach may at first appear to be an invitation to error, and it is. As sensor fusion schemes are typically recursive, incoming data must be processed before future data can be processed. Magnetometer data must be either ignored or included and the decision is made without knowledge of the data that is to come after. From this perspective it is logical to set a

pre-determined threshold for acceptable error, which once exceeded triggers restabilization of heading via magnetometer inclusion.

Magnetic interference is problematic recursive sensor fusion because it is difficult to model accurately or practically. It is not a property the sensor itself, rather of the environment. In a magnetically undisturbed environment, the KF approach is, in the opinion of this work, a highly accurate method for sensor fusion and orientation calculation. However, MOF achieves accuracy similar to the KF approach, is computationally more efficient, does not require high sampling rates, and more straightforward to implement and adapt. It is because of this opinion that MOF was the chosen method of sensor fusion, and that a recursive sensor fusion approach constitutes the core of the MIMA.

3.2.2.3 *Sensor Fusion of Inertial Data*

It is clear from Figure 19 that MOF operates on only the accelerometer and gyroscope, ignoring magnetometer data. As such MOF cannot determine heading, only relative yaw, but this does not make it useless in the calculation of heading. Without the stabilizing influence of the magnetometer, gyroscope integration error grows due to sensor bias, though the effect of random, zero-mean sensor errors are nullified. The result is yaw orientation, off by some constant offset with respect to Magnetic North, confounded only by bias error.

The power of this attribution should not be overlooked. Even without the magnetometer, MOF can calculate yaw, though not heading, and the only source of error is the gyroscope bias. Gyroscope bias causes error during integration, which accumulates over time. In the experiments of this thesis the rate of accumulation was on the order of 0.2 deg/s. As the considered time period of integration shrinks, the accumulated error decreases, allowing for great accuracy over very short time periods. Furthermore, if an accurate method of bias compensation can be developed, the only remaining error source will have been mitigated.

For the reasons expounded in this section, a recursive sensor fusion approach makes up the core of the MIMA, specifically a quaternion based gradient descent optimization approach developed by Madgwick et. al. Despite removal of the magnetometer from the fusion process, roll and pitch angles can be determined, as well as much of the information required to accurately determine heading. Developing a separate process to utilize the magnetic data allows for the

strengths of recursive sensor to be fully exploited, whilst the weakness associated with magnetic interference are nullified.

3.2.3 Heading Inclusion Filter (HIF)

3.2.3.1 Introduction

The Heading Inclusion Filter (HIF) is a novel method of thoroughly assessing the nature of magnetic field readings via a series of tests, designed by the author. The reason multiple tests are required is due to the complex and difficult-to-identify nature of magnetic fields, as well as a requirement of adaptability. There are many aspects of the magnetic field, and all must be considered to develop a reliable filter. Magnetic interference is difficult to determine because the nature of a magnetic field cannot be gauged simply by comparing it to a series of idealized parameters. Each environment presents its own unique values for those parameters, and accurate identification of interference requires taking this uniqueness into account.

A flowchart representation of the HIF's steps of operation pertaining to the inclusion and exclusion of magnetic data can be seen in Figure 21 . The inputs are the roll and pitch angles as determined by MOF, as well as the raw magnetometer data. In order to assess the magnetic data in a consistent manner it must be converted to a coordinate system independent of the MARG's orientation. The known roll and pitch angles are used to complete this transformation via a process known as local leveling. A degree of freedom about the vertical axis will remain as no yaw/heading information is available, but the magnetic data can be resolved into an XY component and a Z component.

After local leveling, the magnetic data is simultaneously subjected to six tests (Figure 21), classified under inertial agreement testing or magnetic characteristics testing. All six tests must be passed for the data to be deemed undisturbed. The inertial agreement testing consists of two tests, each using gyroscope or accelerometer data. The first of these tests, the Heading Estimate Comparison Test (HECT), compares the estimated heading changes of the gyroscope and magnetometer, requiring a certain level of agreement to pass. The second of these tests, the Change in Acceleration Test (CIAT), requires a certain magnitude of acceleration to be present for the data point to pass.

The next set of tests, the Magnetic Characteristics Tests (MCT), work solely on the locally-leveled magnetometer data, and require each characteristic tested to fall within a pre-determined acceptable range specific to each environment. This range is identified by examining the gathered magnetic data and determining the undisturbed periods based on knowledge of the environment and the data's behavior. The magnetic characteristics tested are the norm, the XY norm, the Z norm, and the dip angle. These characteristics are inter-related and testing all four could be viewed as redundant. The two main characteristics are the XY norm and the Z norm, from which the norm and dip angle are directly calculable. However, testing the norm and dip angle imposes constraints on the interrelation of the XY norm and Z norm, allowing for more overall specificity. The acceptable test ranges can thus be relaxed slightly to better accommodate sensor noise and slight field variations, increasing the potential for finding undisturbed data.

Explanation of the HIF will begin with a section detailing local leveling. Next the Heading Estimate Comparison Test will be explained, as it is basis of the HIF. The Change in Acceleration Test was introduced as a secondary test to address a weakness in the Heading Estimate Comparison Test, so the Magnetic Characteristics Testing will be discussed first. After all the tests have been considered, the process of heading determination and bias compensation will be examined, where the passing magnetic data is combined with the M-determined Euler Angles to determine accurate heading.




Figure 21: Flowchart explanation of the Heading Inclusion Filter (HIF) process for including or excluding magnetic data points.

3.2.3.2 Local Leveling

The first step of the HIF is to convert the measured magnetometer values into the global coordinate systems as much as possible. Using these roll and pitch angles determined by MOF,

the raw magnetometer data can be partially converted from the MARG coordinate system to the global. The conversion is not complete, but constrained to a single degree of freedom about the vertical (Z) axis. This partial transformation is known as local leveling, as seen in Equation 5.

Equation 5: Local Leveling

$$\begin{bmatrix} \mu_X \\ \mu_Y \\ \mu_Z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \eta & \sin \eta \\ 0 & \sin \eta & \cos \eta \end{bmatrix} \begin{bmatrix} \cos \xi & 0 & \sin \xi \\ 0 & 1 & 0 \\ \sin \xi & 0 & \cos \xi \end{bmatrix} \begin{bmatrix} \mu_x \\ \mu_y \\ \mu_z \end{bmatrix}$$

Here, the capitalized subscripts of μ denote the transformed coordinate system, and the lowercase subscripts denote the raw magnetometer readings. The MOF-calculated roll angle is symbolized by η , and the MOF-calculated pitch angle is symbolized by ξ .

3.2.3.3 Heading Estimate Comparison Test (HECT)

The first test in the HIF involves a comparison of the gyroscope and magnetometer estimates, and is termed the Heading Estimate Comparison Test (HECT). Gyroscope integration error decreases as the time period considered decreases. Shrinking said time period to fractions of a second minimizes the error, and an accurate estimate of heading change over that time period can be made by the gyroscope alone. The magnetometer data can be discretized into the same small time periods and the heading change over each time period can be calculated. This process of this comparison is shown in Figure 22.




Figure 22: Flowchart of the Heading Estimate Comparison Test (HECT)

Here φ is the sensor sample rate, γ is the yaw angle as estimated by the gyroscope, and μ is the heading as estimated by the magnetometer. The value of i denotes the time step and Φ_{tol} is the maximum allowed difference in the heading change estimates. Here the yaw is not measured with respect to Magnetic North as the heading is. As they describe rotation about the same axis, their values will differ by some constant, but their relative changes will be the same. Two estimates of heading change have been generated, one by the gyroscope and one by the magnetometer. The short time period considered ensures the gyroscope's estimation is highly accurate, and in the absence of magnetic distortion the magnetometer's estimation will also be accurate. Given this, the two estimations are compared and if the magnetic heading change is not the same as the gyroscope's yaw change, that magnetometer data is deemed disturbed. If the difference in the estimations falls within the allowed tolerance, the magnetic data is not deemed undisturbed, it is instead subjected to the next test.

3.2.3.4 Magnetic Characteristics Test (MCT)

In an ideal situation the ambient magnetic field would consist only and exactly of the geomagnetic field. Such an ideal situation is rare, and many IMC tests must be conducted inside buildings. Passage through the building exterior alters the geomagnetic field, but the field properties and relationships are generally changed, rather than destroyed. For example, a property of the geomagnetic field is the dip angle, the angle at which magnetic field lines contact the Earth's surface. Outside this angle is of a certain value and is constant. Inside, this angle is of a value different than it was outside, but it can still be considered a constant for practical purposes. The constant value implies field homogeneity, which constitutes an undisturbed field for the purposes of IMC. In order to develop a useful magnetic filter, the altered, inside value for each characteristic must also be known, and this can be determined via field mapping. In order to ascertain the nature of the tested magnetic field, a series of test were developed pertaining to the ambient magnetic field and its properties. These tests were grouped under the Magnetic Characteristics Test (MCT).

Tested Characteristics

The magnetic characteristics tested are as follows:

- Norm

- XY Norm
- Z Norm
- Dip Angle

Each of these characteristics have a specific value when applied to the geomagnetic field, and as such will have a certain value inside a building, making them ideal candidates for testing for disturbance in a magnetic field. Together these characteristics test both components of the magnetic field and how they relate to each other, making for an effective method of disturbance detection.

As previously mentioned, the norm is magnitude of the magnetic field, and is calculated via Equation 1, reproduced here.

Equation 6: Magnetic Norm

$$M_N = \sqrt{M_X^2 + M_Y^2 + M_Z^2}$$

M_N is the magnetic norm, M_x is the X component of the magnetic field, M_y is the Y component of the magnetic field, and M_z is the Z component of the magnetic field. The Z component is vertical, normal to the Earth's surface and the X and Y components are parallel to the plane of the Earth's surface. The combination of the X and Y components results in a vector that points in the direction of Magnetic North.

The XY Norm is a similar characteristic to the norm, calculated via

Equation 7: Magnetic XY Norm

$$M_{XY} = \sqrt{M_X^2 + M_Y^2}$$

where M_{XY} is the magnetic norm defined by the X and Y components of the field, M_x is the X component, M_y is the Y component. The Z Norm is a similar characteristic, calculated via

Equation 8: Magnetic Z Norm

$$M_Z = \sqrt{M_Z^2}$$

where M_Z is the magnetic norm, and M_Z is the Z component of the magnetic field. Finally, the dip angle is defined by the relation of the XY Norm and the Z Norm, via

Equation 9: Magnetic Dip Angle

$$\theta_{dip} = \tan^{-1} \left(\frac{M_{XY}}{M_Z} \right)$$

where θ_{dip} is the dip angle.

The local leveling performed in the initial step of the HIF (3.2.3.2) allows each of these magnetic field characteristics to be meaningfully calculated within the same coordinate system. The testing proceeds as seen in Figure 23, where M_i denotes a calculated magnetic characteristic (M_N , M_{XY} , M_Z , or θ_{dip}) at time step i , M is a value equal to the ideal value of the magnetic characteristic, and v is the acceptable deviation from the ideal value. The process shown in Figure 23 is repeated for each of the four characteristics, and only the values of μ that fall within the acceptable deviation for all four characteristics are passed.

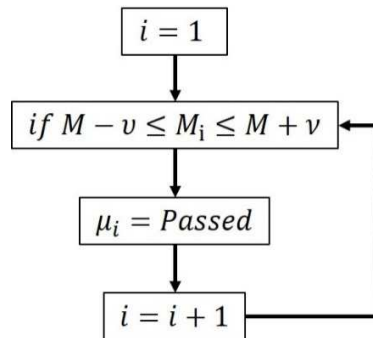


Figure 23: Flowchart representation of the Magnetic Characteristics Test (MCT)

Field Mapping

After local leveling of the magnetometer readings and computation of the four magnetic characteristics, the values of those characteristics must be compared with an ideal value and an associated error tolerance. As previously mentioned, building walls warp the geomagnetic field

as they pass through, but leave the inside magnetic field homogenous. Therefore values for what constitute an undisturbed field must be determined for each characteristic via the process of field mapping. Field mapping requires that the testing environment be sampled to determine the acceptable ranges of each characteristic.

To perform such field mapping, a MARG must be moved around within the testing environment avoiding areas of known disturbance. The purpose is to sample the area of undisturbed field and generate a range for each characteristic that is acceptable. Once the area has been surveyed, the MARG data must be processed so as to calculate each of the four characteristics. Assuming the MARG was kept within undisturbed regions, each characteristic should be largely constant, only deviating within a small range from the average. Once ranges have been generated for each of the characteristics, the test can be performed.

Testing

Once the MARG data has been recorded, the magnetometer components locally leveled, the characteristics calculated, and the acceptable range of values determined via field mapped, the testing can be completed. For each tested characteristic, any magnetic data point that does not fall within the specified range is rejected, and any data point that does not pass all four characteristic tests is rejected.

3.2.3.5 Change in Acceleration Test (CIAT)

Early in its development, the HIF included only the gyroscope comparison test and magnetic characteristics tests. In most situations these tests were sufficient, the HECT performing the majority of the filtering and the magnetic characteristic tests ensuring that no anomalies passed. However, heavy reliance on the HECT caused problems during stationary periods. Without motion the gyroscope and magnetometer will each predict zero change in heading, rendering the HECT useless, as it will pass every point. Generally the magnetic tests could compensate for this, but occasionally magnetic data would pass that would have normally been rejected had the HECT been operational.

In order to address this, another test based on the magnitude of acceleration was added, termed the Change in Acceleration Test (CIAT). In order to pass the CIAT, acceleration other than gravity must be present. What constitutes acceleration must be defined via a minimum value,

and it was found that this value must be defined by the user based on the nature of the test. For example if the motion consists predominantly of stationary periods, requiring a high minimum acceleration would reject many potentially good data points. Conversely if the motion consisted of few stationary periods, a high minimum acceleration would be acceptable. The CIAT process can be seen in flowchart form in Figure 24.




Figure 24: Flowchart of Change in Acceleration Test (CIAT)

Here i denotes the time step, a is the acceleration measured, and A_{min} is the minimum required acceleration value.

3.2.3.6 Conversion of Yaw to the Global

Once all the tests have been performed, the magnetometer data has been purged of all identifiable disturbed data points, and all that remains is a scattering of undisturbed points. It is this data alone that will be used to convert the MOF-calculated yaw into heading. The first step in said conversion is to identify the time steps paired with each of the undisturbed data points, and pull the values for each corresponding time point in the yaw data.

Now that the gyroscope readings temporally paired with the undisturbed heading reading have been identified, a least squares best fit between the two data sets is performed. This process can be modeled as seen in Equation 10 and Equation 11.

Equation 10: Global Heading Model

$$\gamma(t) = \alpha(t) + \delta + \varepsilon(t)$$

Here γ is the yaw as calculated by MOF, t denotes time, α is the true, error-free yaw, δ is the

offset value between the yaw and heading, and ε is the error designed to account for the gyroscope bias. Initially the error term is ignored by setting it equal to zero, and the value of δ is determined via minimization of the least squares error between γ and the magnetic data points, denoted via μ , as expressed in Equation 11. The values for i are set such that only the time points that passed the HIF are compared.

Equation 11: Least Squares Optimization

$$S = \sum_{i=1}^n (\gamma_i - \mu_i)^2$$

Here S is the sum of the passing points' error, and n is the number of data points that passed the HIF. Conceptually this represents shifting the gyroscope data up and down until an optimal fit is determined with respect to the undisturbed heading points. The end result of the HIF is heading calculation with a high immunity to magnetic disturbance.

3.2.3.7 Bias Compensation

In the HIF the yaw was transformed to the global (heading), but the error term was not accounted for. There was no mechanism for heading stabilization in MOF, thus time integration of the gyroscope bias leads to a drift error that increases with time. It was experimentally determined that under simple motion the bias drift could be estimated as linear with a high degree of accuracy. As such, stabilization of the heading data could be accomplished by designating the error term, $\varepsilon(t)$ from Equation 10, as a linear function of some optimal slope and intercept (Equation 12).

Equation 12: Bias Error Model

$$\varepsilon(t) = m * (t - b)$$

Here m is the gyroscope's bias value approximated as a constant, t is time, and b is the X intercept of this linear function. The intercept term is included to account for skewed distribution of magnetic data points. Values for m and b are determined via minimization of the least squares error between γ and μ (Equation 11), where in contrast with Section 3.2.3.6, the error term is included.

In the optimization the tested values of m were chosen cover a range larger than the most extreme experimental tests, considering up to 80 degrees of drift over the test's duration. The values of b were chosen to vary between its two possible extremes, intercepting at the beginning and end of the functions. A graphical example of the tested functions can be seen in Figure 25. This figure relates to a hypothetical 100 second test, and shows a subset of the linear functions considered for optimal bias compensation. This graphic is intended to clarify the purpose of m and b in Equation 12. Each red circle highlights a different b value (X intercept), and each blue line represents an m value paired with the b value. It can be seen that for each b value, a range of m values are tested, and a range of b values are tested.

The result of this approach is that every plausible linear function over the duration of the test is considered, and the function that leads to the greatest agreement with the magnetometer data is selected. The reason for varying the slope (m) is clear in that it is a direct estimator of the bias, but the reason for the intercept is less clear. The inclusion of the intercept value (b) is required because the distribution of the undisturbed magnetic data points can be heavily imbalanced, i.e. the first five seconds of a test have 27 undisturbed data points and the last five seconds have only four undisturbed data points as in Figure 26. Such imbalances affect the HIF conversion of yaw to heading, and varying the intercept allows the HIF to compensate for this.




Figure 25: An explanatory subset of the considered linear compensation functions.

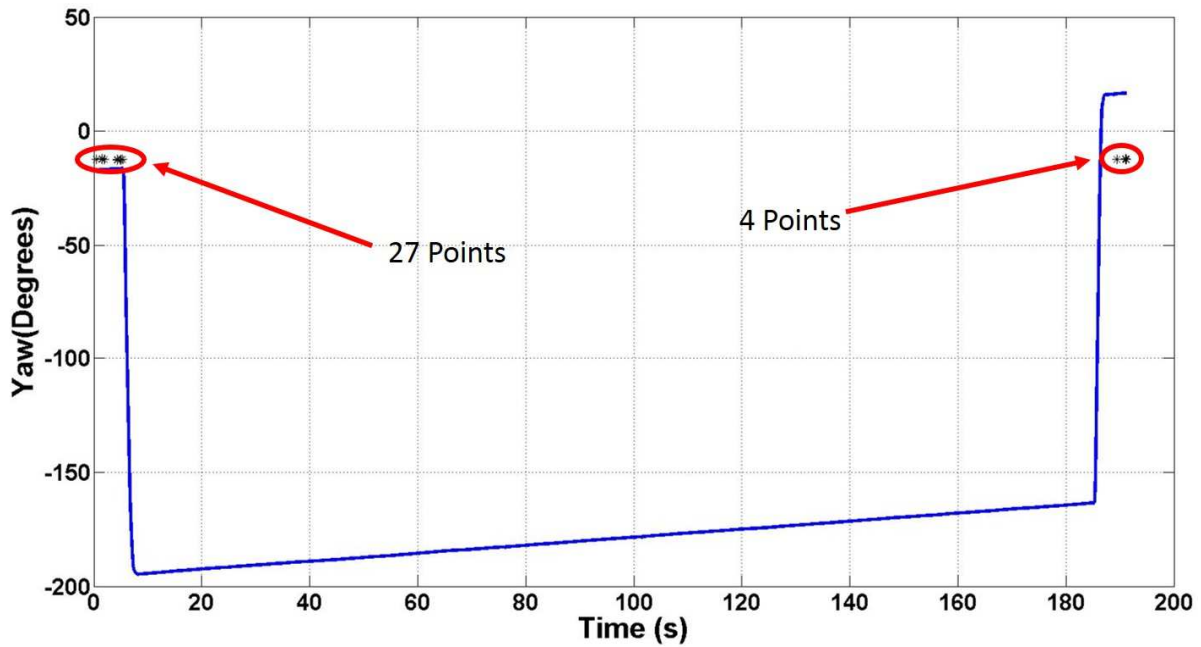


Figure 26: An example of the effect of unequal point distribution on data fitting.

Figure 26 shows how the distribution of points to be fit with affects the optimized location. In this test there are more than six times as many points on the left cluster as the right cluster. When Equation 11 is applied to determine the optimal fit, this imbalanced distribution causes the left cluster to be fit much closer than the right. The bias drift remains linear, however, and can still be modeled with a linear function. Figure 26: is an example of why optimization of the linear function must include the X intercept as a design variable in addition to the slope.

The end result of applying bias compensation is to increase the accuracy of the estimated global heading, stabilizing heading with magnetic data even in a magnetically disturbed environment. This process models the drift error caused by the bias, and combines that information with filtered magnetic data, including the magnetometer data such that undisturbed data will improve the heading estimate, but disturbed data has no effect.

3.2.4 Selection of Tolerance Values

In this section the methodology of the HIF has been expounded, and with that knowledge the choice of tolerance values can be explained. These values equate to the values used in the MIMA's general operation, as well as in the experiments of Chapter 4. The tolerance value for each test is listed in Table 2.

Table 2: Here the tolerance values chosen for the MIMA's operation are listed, as well as their meaning.

Test	Chosen Value	Meaning
HECT	0.1 (deg/epoch)	The maximum allowed difference in the estimates of heading change of the gyroscope and magnetometer over the epoch (0.05 sec)
CIAT	0.008 (m/s ²)	The minimum required change in acceleration over each time point (equals .01 at 100 Hz)
XY Norm	Test Dependent (milliGauss)	The absolute value of acceptable deviation from the designated undisturbed XY norm value
Z Norm	Test Dependent (milliGauss)	The absolute value of acceptable deviation from the designated undisturbed Z norm value

3.2.4.1 Heading Estimate Comparison Test (HECT) Tolerance

This test's tolerance was chosen to be 0.1 deg/epoch, where an epoch represents 0.05 seconds. This equates to 2 deg/s. Of all the tolerances, this was the most experimentally determined. The initial inclination was to make this tolerance extremely tight, as it seems simple and effective. If the magnetometer data does not agree with the gyroscope, interference must be present. In reality, even a magnetically undisturbed area will contain slight imperfections. Further, at low tolerances the noise characteristics of the sensors more prominently affect the test results.

Making this tolerance very stringent excluded too much data, as will later be seen in the Sensitivity Analysis of Section 4.4. An epoch value of 0.05 seconds at 100 Hz calculates the change in heading over five time steps. This time span is long enough reduce the effect of noise, and short enough that it was capable of detecting small disturbances. A maximum variation of 0.1 deg over each epoch, is still stringent, as any significant disturbance would be detected at this level.

3.2.4.2 Change in Acceleration Test (CIAT) Tolerance

Determining the tolerance value for this test was simpler than for the previous test, but with greater ramifications. This test was designed to address the HECT's weakness to stationary periods, therefore this tolerance value was largely based on the accelerometer's noise characteristics. During stationary periods, the noise characteristics of the accelerometer dictate that it will record a range of values, and the width of that range is the logical choice for the tolerance value.

Figure 27 shows accelerometer data from a MARG that was stationary for the plotted 60 seconds. The red error bars bracket a range of 0.08 m/s² (9.7 ± 0.04 m/s²) and exclude the vast majority of readings. A value of 0.008, ten times smaller than this, was chosen as tolerance value. Requiring a change in acceleration greater than 0.008 m/s² between each time step greatly reduces the amount of data from stationary periods that can pass the HIF. However, setting this tolerance so as to include the vast majority of stationary data, such as in Figure 27, is not necessarily desired, especially if the sensor is stationary in the magnetically undisturbed areas. As a result, this test became an adjustable tool within the HIF's infrastructure, a topic furthered discussed in the limit testing and Sensitivity Analysis of Sections 4.3 and 4.4

respectively.

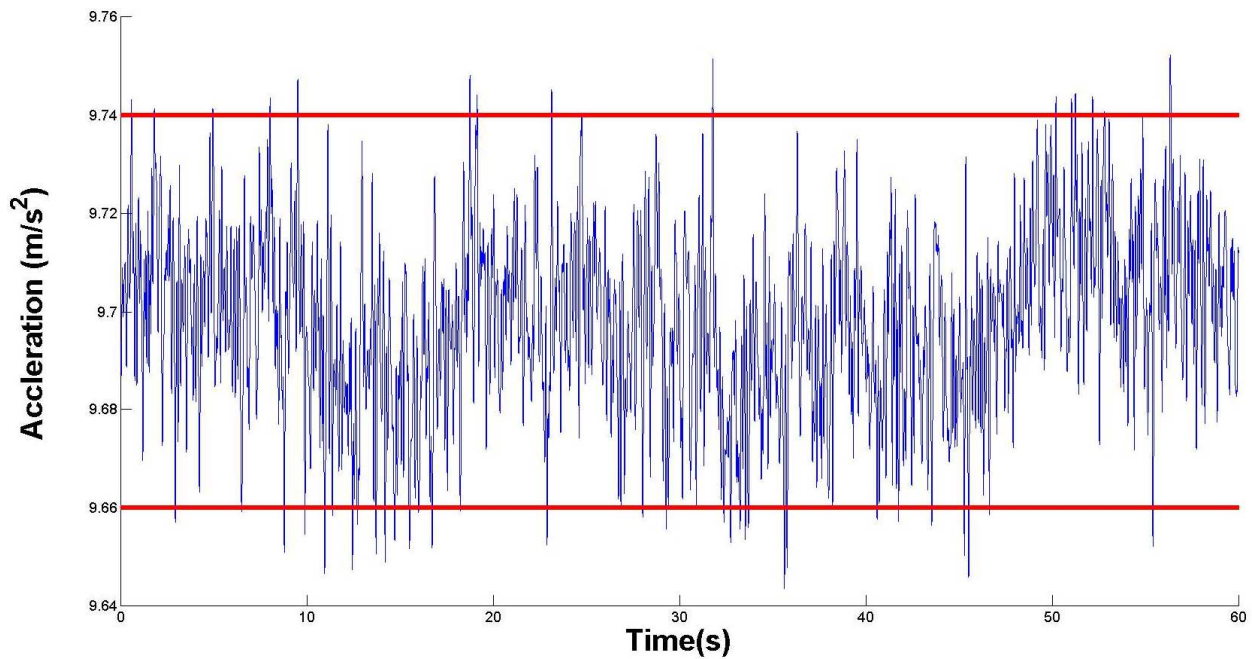


Figure 27: This graph shows the accelerometer readings (blue) over a 60 second stationary period. The red bars bracket a range of 0.08 m/s^2 , equivalent to the chosen tolerance value for this test.

3.2.4.3 Magnetic Characteristics Test (MCT) Tolerance

The tolerances for this test were initially based on the geomagnetic field, under the reasoning that a magnetically undisturbed field should closely resemble the Earth's magnetic field. As explained in Section 3.2.3.4, this is rarely the case. Furthermore, the magnetic field in a particular location can vary with time. Table 3 explores this fact by comparing the four testing sessions performed. The listed range corresponds to the maximum and minimum values of that characteristic determined to be undisturbed for each testing session.

Table 3: The ranges of each characteristic determined to be undisturbed for each testing session.

Testing Session	Dip Angle (deg)	Norm (mG)	XY Norm (mG)	Z Norm (mG)
1	64 – 72	42.3 – 43.7	13 – 18	38 – 41
2	67 – 74	44 – 46	13 – 17.5	41.5 – 42.7
3	67 – 71	43.5 – 45.5	14.5 – 17	40.5 – 42.5
4	63.7 – 65	43 – 44	18.2 – 19.2	38.6 – 39.6

The values in Table 3 show that the range for each characteristic varies between testing sessions, even though the testing location remained the same. Motion type did have an effect on the characteristics, evidenced by the wide range associated with the XY Norm in Session 1. In this session the two-linkage testing apparatus was used and the time in the undisturbed area was spent in motion. Despite this, the Norm and Z Norm were influenced less by the motion, as their ranges are similar to those of the other testing sessions.

The implication of the ranges listed in Table 3 is that it would be more advantageous to set the tolerances on a per case basis rather than using the same range for all tests, and this was the approach taken in the operation of the MIMA. This may seem like a laborious and impractical task, and while it would be easier to use a standard value for every test, this approach is relatively simple and affords the MIMA a great deal of adaptability, as it can maintain specificity across all environments.

To determine the acceptable range for a test, the portion of the test deemed to be in magnetically undisturbed area is noted, and each characteristic is plotted. The values on the plot corresponding with the magnetically undisturbed area are identified and used to set the range. Consider the example shown in Figure 28. In this instance knowing that the test was set to spend its first and last five seconds in the undisturbed area and the rest in a disturbed area make range determination an easy task. The range should be set as tightly as is possible without excluding the undisturbed data.

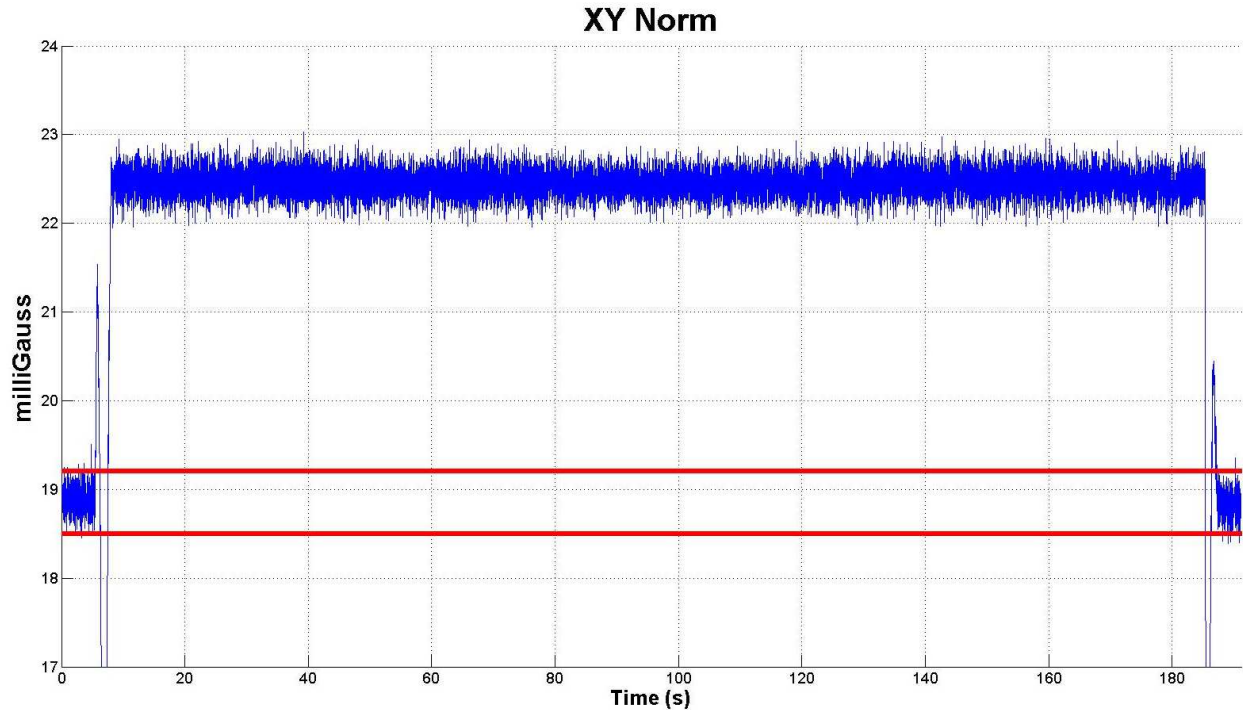


Figure 28: This plot shows the XY norm of a test in which the only time spent in a MUA was the first and last five seconds of the test. Here the disturbed and undisturbed periods are readily identified, and the range can be set to contain the undisturbed data, shown by the pair of red lines. This user-selected range is then used to determine tolerances within the Heading Inclusion Filter (HIF).

3.3 Experimental Setup

3.3.1 Introduction

The MIMA was designed to increase the stability of MARG orientation calculations in magnetically disturbed environments. As such the experimental protocol of this research was set up so as to test and showcase this increase in stability. In each experiment a known source of magnetic interference was placed in an otherwise undisturbed environment and the MARG was moved between the disturbed and undisturbed areas via a manually-operated, pivoting aluminum linkage. The chosen MARG, an XSENS MTw, was capable of exporting both fused and raw sensor data, with sensor fusion achieved via the XSENS' proprietary Kalman Filter (XKF). The raw data was processed by the MIMA, and the heading values as calculated by the MIMA and XKF were compared with optical data. This comparison highlighted the MIMA's ability to adaptively use data MARG data to mitigate error due magnetic interference.

3.3.2 Experimental Set-Up

The chosen testing environment was the 3D Bio-Motion Lab at the Center for Computer Aided Design at the University of Iowa. The room is designated for motion capture, inertial and optical, making it the logical choice for testing an inertial motion capture algorithm. A previously constructed wooden staircase was used as the mounting platform for the testing apparatus (Figure 29). The purpose of the staircase was to elevate the testing apparatus, as the floor was known to disturb the magnetic field in its proximity. The magnitude of the interference decreased exponentially with distance from the floor, such that distances greater than approximately 18 inches could be considered undisturbed.



Figure 29: This figure shows the wooden staircase used to sensor apparatus. This provided adequate separation from the ground so as to avoid the disturbance it produced.

Classifying the Bio-Motion Lab as an undisturbed area does not imply that the magnetic field inside it was undistorted relative to the geomagnetic field. Upon comparison with the known geomagnetic field values (Maus, 2015), the field was found to be distorted. Furthermore, magnetic field variability was also found between the different testing dates. However, in all

cases the field was found to be homogenous, which constitutes an undisturbed magnetic field for the purposes of inertial motion capture.

One could argue that the lack of agreement with the geomagnetic field and variation between test times makes the field disturbed, and this may be true based on how one defines magnetic interference. However, the magnetic field can vary slightly even in seemingly ideal areas, and heading can be accurately calculated in any homogenous magnetic field, regardless of magnitude. Furthermore, expecting less disturbance than encountered in the Bio-Motion Lab is not realistic. The Lab is a designated motion capture room, and any disturbances present represent a realistic testing environment.

The testing apparatus consisted of an aluminum base with a pivoting aluminum linkage mounted on top of it (Figure 29). Aluminum is non-ferromagnetic and does not introduce magnetic interference. Two linkage formulations were used, one comprised of two links and one of only one link, seen respectively in Figure 29 and Figure 30. The apparatus was used by grasping the side designated as the handle and manipulating it to produce the desired motion. In the case of the two-link formulation, the second link was also pivoted by hand. In order to introduce a source of disturbance, a magnet or a metal pipe was mounted on the staircase (Figure 31). Here the pipe constitutes a practically encountered interference source, and the bumper ensures that a constant, realistic distance is maintained between the pipe and MARG. The magnet placement area denotes where the donut magnet was placed when being used as the interference source instead of the pipe.




Figure 30: Here the testing apparatus in its two-link formulation can be seen. The pivoting linkage is mounted onto an aluminum base. The MARG sensors are placed on one end of the linkage, while the other end is used to manually pivot the linkage.




Figure 31: Here the testing apparatus in its single link formulation can be seen, as well as the pipe and bumper magnetic disturbance set up.

Optical motion capture was used as the standard of comparison between the MIMA and XKF. The Bio-Motion Lab is outfitted with 16 MotionAnalysis Eagle-4 Digital cameras; optical motion capture is considered the gold standard for the industry, and is immune to magnetic interference. Four markers were placed on the MARG sensors as seen in (Figure 32).




Figure 32: Here the four markers mounted onto the MARG can be seen, labeled for post-processing determination of orientation. The red and orange vectors were calculated, then averaged into a single vector (blue arrow).

The captured data was processed with the software Cortex, and lowpass filtered at 8 Hz. Orientation was calculated from the markers by calculating two vectors, defined by A and B and C and D (Figure 32). In order to minimize error, these two vectors were then averaged into a single vector, whose direction represented the orientation of MARG within the marker-based coordinate system (Figure 32). Though it would be customary to use just two markers to determine the vector, four markers were used to increase accuracy as well as to simplify processing. If four markers are used, their relative geometry can be defined and exploited to fill in if markers is temporarily not captured by the cameras. In order to achieve global orientation,

the marker coordinate system had to be calibrated to coincide with Magnetic North. The linkage was planar and required only the vector parallel to the floor and directed at Magnetic North for calibration. This vector was determined by using a MARG to measure the magnetic field components at the pivot point of the testing apparatus. The components were then converted to magnetic heading via Equation 13.

Equation 13: Magnetic Heading

$$\mu = \tan^{-1}\left(-\frac{M_y}{M_x}\right)$$

After calibration to Magnetic North, the marker data provided accurate, global orientation data on the MARG, and could be used to objectively compare MIMA and XKF data.

3.4 Conclusion

This chapter has elucidated the methodology of the developed algorithm's operation. It was explained that the MIMA differs from traditional methods of processing MARG data in that it excludes the magnetic data from the KF process, reintroducing it post-fusion in a way that can only improve the orientation estimate. In environments of consistent magnetic interference, this affords the MIMA several advantages over the traditional KF approach, mainly defined by its ability to repair data affected by magnetic interference based on its relation to the surrounding information.

After the methodological explanation, an experimental setup was developed to test the abilities and limits of the MIMA, as well as identify its governing factors and the methodology of their interactions. These factors are what govern the MIMA's reparation ability, and thorough analysis will provide insight into the MIMA's strengths and limitations as well as how the MIMA can be tailored to each situation for optimal results. In Chapter 4 the experimental setup developed here is utilized to conduct said analysis via a series of experiments.

CHAPTER 4 : EXPERIMENTATION, RESULTS, AND DISCUSSION

4.1 Effect of Initialization in a Disturbed Environment

4.1.1 Objective

The purpose of this experiment was to compare the resulting MIMA and XKF yaw angles when testing began in an area of known magnetic interference and then moved into an undisturbed area. The expectation of this experiment was that the XKF would initialize at an incorrect heading, become confused by the unexpected change in magnetic field, and require time to compensate before producing accurate heading estimates. Conversely the MIMA was expected to produce accurate yaw angles for the entire test, as it could repair the initial, disturbed data with the later, undisturbed data.

4.1.2 Methods

The motion of this experiment was performed on the two-linkage version of the testing apparatus (Figure 30). The linkage was manually manipulated to begin the test in an area of magnetic disturbance stationary for ten seconds. Though the test only captured ten seconds of disturbed data, the sensor was placed in the disturbed area approximately one minute before testing. This period of time allows the XKF to stabilize itself, and such stabilization is recommended by XSENS (XSENS, 2013). XSENS also recommends avoiding initialization in disturbed environments, but this is not always possible and this test was meant to address this.

After ten seconds, the linkage was pivoted 180 degrees, into an undisturbed area and oscillated over a range of 60 degrees repeatedly for 70 seconds. Generation of the magnetic disturbance was accomplished by placement of a magnet in close proximity to the sensor, as shown in Figure 31. In order to ensure the magnet's disturbing effects were completely removed after the disturbed portion of the test was completed, the magnet was removed from the testing area once the allotted ten seconds had passed.

After the data was collected, the magnetic characteristics were inspected to determine the range of acceptable values for each. This inspection was also used to determine the tolerance value of the accelerometer test. Once these values were determined, they were input into the MIMA and the data was processed. The resulting heading angles were then graphically compared with the

XKF and optical data.

4.1.3 Results

As expected, beginning in a disturbed environment before moving to an undisturbed one caused error in the XKF in the form of an adjustment drift, whereas the MIMA's accuracy remained the same for the duration of the test. The initial error of the XKF was -34 degrees, and the MIMA's initial error was -2.7 degrees (Figure 33). At the completion of the test the XKF had an error of 4 degrees and the MIMA error was -.6 degrees (Figure 33).

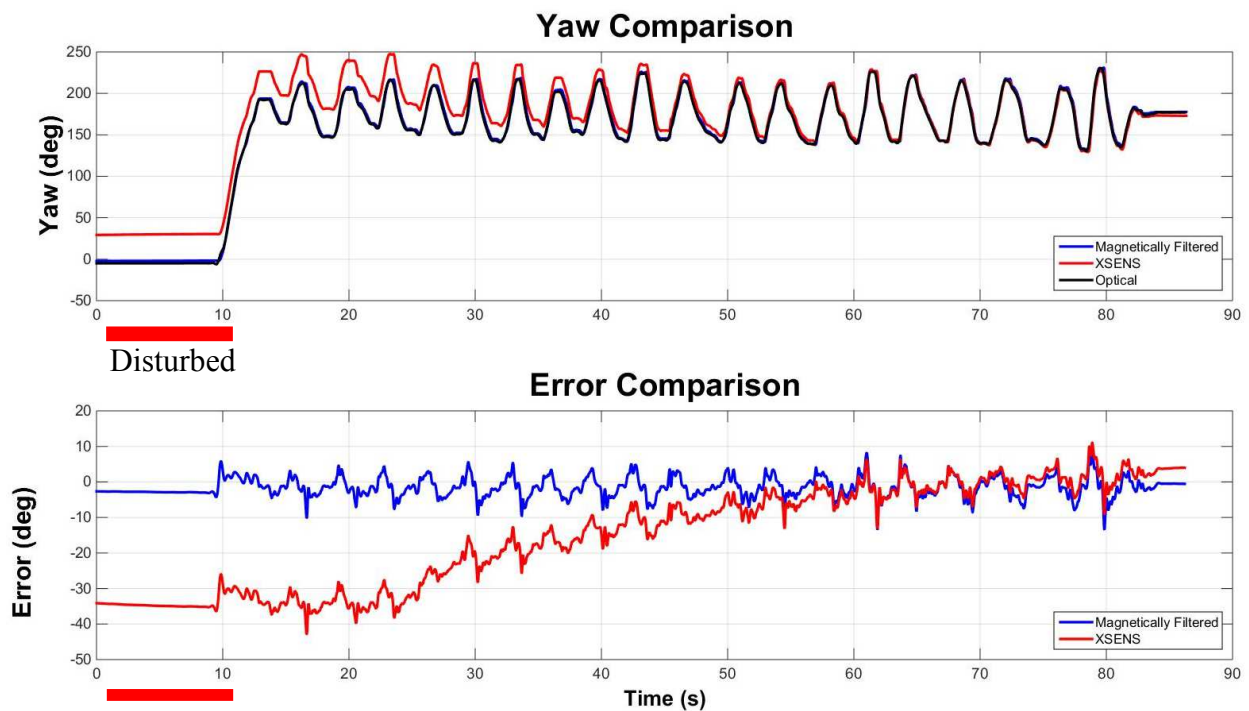


Figure 33: Graphical Comparison of the XKF and MIMA heading angles under motion conditions beginning in a disturbed environment and moving into an undisturbed environment. The red bars mark the periods in which a disturbance was present. These graphs highlight the drift error the XKF must incur while adjusting to the undisturbed environment, as well as the MIMA's ability to avoid such drift and provide accurate data even in areas of magnetic interference.

4.1.4 Discussion

This experiment is a quintessential example of the problems encountered by traditional Kalman Filters (KF). It is expected that the XKF was aware it was beginning in a magnetically disturbed

environment, but it still based its heading on the magnetic field because this was the only available information. After initializing and building up a data history in a disturbed area, the XKF was unable adjust to an undisturbed area without requiring an adjustment period (Figure 33). For the first 15 seconds in the undisturbed area, the MIMA and XKF calculate very similar heading yaw values relative to their respective starting positions. During this time period, both algorithms are relying heavily on the gyroscope to calculate heading. After 15 seconds have passed, the XKF begins to drift towards the correct heading because it can no longer ignore the magnetometer readings. Overall it takes 50 seconds before the XKF reaches the same level of accuracy as the MIMA.

The MIMA maintains the same level of accuracy for the duration of the test, as was expected. The MIMA relies heavily on the gyroscope, using the magnetometer readings only to determine global yaw information and estimate a linear bias compensation function. In this test the main factor that made the MIMA more accurate than the XKF is that the MIMA was not required to initialize based on data it knew to be disturbed. The MIMA ignored the magnetometer data during the disturbed period, relying entirely on the gyroscope to capture any motion. Once in the undisturbed period the MIMA was able to use that data to convert the disturbed gyroscope readings to the global, as well as compensate for any gyroscopic bias drift error that had occurred.

4.2 Effect of Moving into and out of a Disturbed Environment

4.2.1 Objective

The purpose of this experiment was to compare the resulting MIMA and XKF yaw angles when testing began and ended in an undisturbed area, and spent the rest of the time in a disturbed area. The expectation of this experiment was the same as before, with the XKF expected to become confused by the disturbance while the MIMA remained accurate. However, the manifestation of the XKF error was expected to occur differently. This time the XKF would initialize at the correct heading and incur error as time passed.

4.2.2 Methods

The motion of this experiment was performed on the two-linkage version of the testing apparatus (Figure 30). The linkage was manually manipulated to begin the test in an undisturbed area,

remaining stationary for three seconds, after which the linkage was pivoted into the disturbed area. One minute later the linkage was pivoted back into the undisturbed area and one oscillation was completed.

Total time for the second undisturbed period was five seconds, and the sensor was allowed to stabilize for a minute before the test began. The disturbance was generated via a magnet that was removed from the testing area during the undisturbed periods. After the data was collected, the data was inspected to determine values for the magnetic characteristics and the acceleration test tolerance. The data was then processed and the resulting heading angles were then graphically compared with the XKF and optical data.

4.2.3 Results

The results of this experiment were as expected with both the MIMA and XKF initializing at similar values in the undisturbed field, but with the XKF drifting when exposed to the disturbed field. The initial XKF and MIMA errors were -2.5 and -1.8 degrees respectively, while the final errors were 28.1 and 1.2 degrees respectively.

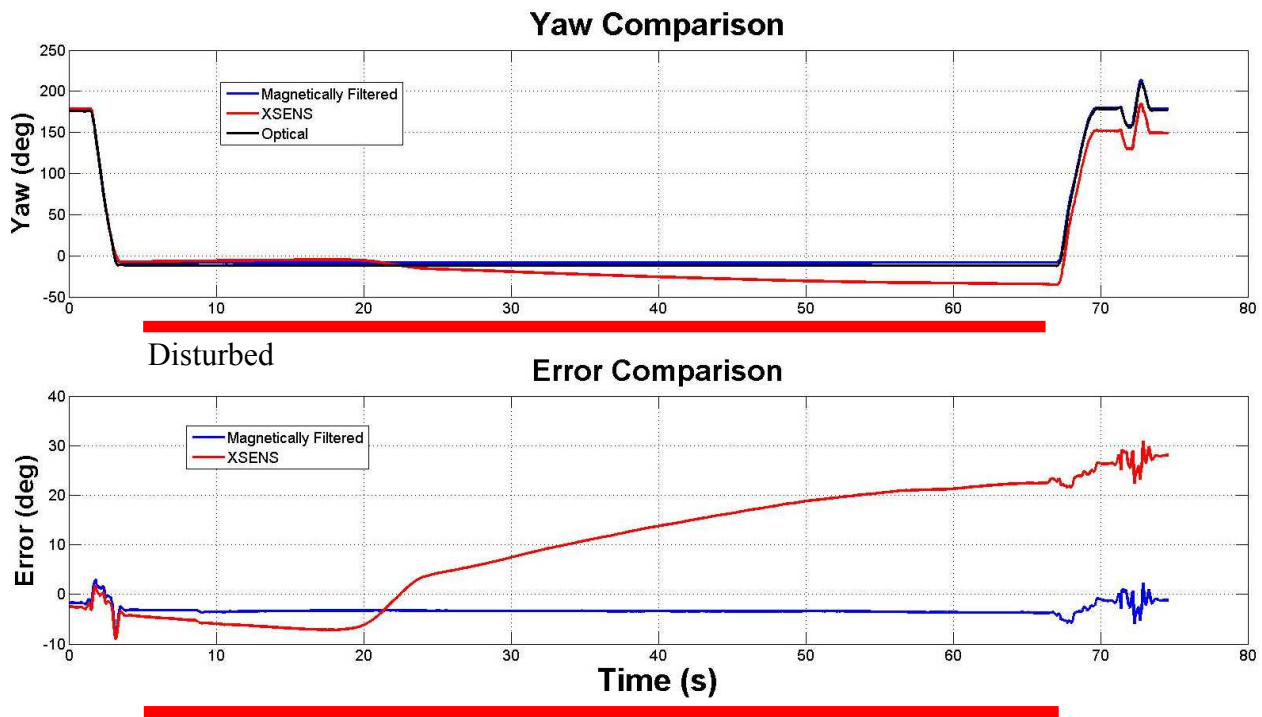


Figure 34: Graphical comparison of the XKF and MIMA (red and blue respectively) calculated heading angles. The motion began and ended in an

undisturbed area, with the interim spent in a disturbed area. The red bars are mark the periods in which a disturbance was present.

4.2.4 Discussion

The experiment showcases the problem traditional KFs experience when moving into a disturbed area. Even when initialized correctly, the KF is unable to determine accurate heading in a disturbed area and accumulates error as times passes. Furthermore, the KF is unable to take swift corrective measures when data capture ends in an undisturbed area. The XKF initialized correctly, and its error increased linearly for the first 15 seconds in the disturbed area, overestimating the correct heading. The direction of accumulation then switched and grew nonlinearly, such that it began to underestimate the correct heading. As in the previous experiment, this drift pattern occurred because the KF at first relied on the gyroscope for heading estimates, but eventually attempted to incorporate the magnetometer data. The initial drift period was due to gyroscope bias, and the second drift period was due to magnetic interference, explaining the change in drift direction.

In contrast the MIMA was relatively unaffected by the disturbance, drifting only three degrees between start and finish, compared to the 30.6 degree drift of the XKF heading. During the undisturbed periods the MIMA was able to gather enough magnetometer data to accurately determine the initial heading, as well as determine accurate bias compensation. During the disturbed period the magnetometer data was completely ignored by the MIMA, relying instead on the gyroscope. Though the gyroscope had drift error it was compensated for by the bias compensation portion of the algorithm.

4.3 Effect of Extended Time on MIMA Reparation Ability

4.3.1 Objective

The purpose of this experiment was to test the MIMA's ability to repair disturbed data, specifically to determine how it would respond to a longer testing period with a greater percentage of the gathered data being disturbed. Testing this ability would determine the practicality of the MIMA in more relatable terms. For example if IMC is to be done in a disturbed area, how often and for how long must the sensor be exposed to undisturbed areas in order to maintain accuracy in the disturbed areas.

4.3.2 Methods

The motion for this experiment was performed on the single linkage version of the testing apparatus, beginning and ending in an undisturbed area for five seconds, with the 180 intermediate seconds spent in a disturbed area. In previous experiments where a magnet had been used to generate the disturbance, here a metal pipe was used. The magnet was ideal for highlighting the abilities of the MIMA, but constituted an extreme form of disturbance rarely encountered in true motion capture. The metal pipe was a more direct analogue to forms of interference encountered in IMC.

As proximity to the pipe has a direct effect on the disturbance magnitude, a separation distance of 4 cm was chosen and enforced via a foam bumper attached to the pipe. Said bumper ensured that the sensor-disturbance distance would remain consistent at a total distance of 5 cm when the marker width was included. The distance was chosen because it was an approximation of the disturbance proximity that would be encountered if an MARG were placed on a hand that was grasping a disturbance source. Specifically this was meant to emulate the situation in which a person riding an ATV grasps the handlebars.

After collection, the data was inspected and ranges for the magnetic characteristics and acceleration test were defined. The data was processed, and the resulting heading angles were graphically compared with the XKF and optical data. A difference in this test's processing relative to the others was that the optical data was not determined by the calibration process used to align the marker coordinate system with magnetic north. That calibration process was completed based on magnetic heading, which given the insights of this thesis is known to be potentially inaccurate. Therefore the marker data was shifted to align with the heading determined by XKF during the initial period in the undisturbed area. Specifically the XKF-calculated heading and optically calculated heading from 1.5 to 4 seconds was extracted and averaged. The optical data was then shift by an amount equal to the difference between the two averages, such that its average optical heading from 1.5 to 4 seconds was the same as that of the XKF in the same time span.

4.3.3 Results

As the purpose of this test was to test the limits of the MIMA's ability to maintain accuracy, the

expectation was that if the testing time was increased three-fold, the accuracy of the MIMA would be compromised. The results were unexpected in that the accuracy actually improved relative to the previous test, despite increasing the disturbed period length by over 275%, from 65 seconds to 180, while only increasing the undisturbed period length 25%, from 8 seconds to 10. The MIMA began at .7 degrees of error and ended with .4 degrees of error. For comparison the same values for the XKF were .3 and -78.8 degrees. These results are presented graphically in Figure 35.




Figure 35: The MIMA and XKF were graphically compared over a period 190 seconds. The first and last five seconds were spent in a MUA, and the other 180 seconds in a MDA.

4.3.4 Discussion

The results of this experiment suggest that under the tested motion of stationary periods punctuated by planar motion, the MIMA can potentially generate accurate data over an indefinite time period in a disturbed area, so long as the beginning and end are punctuated with undisturbed periods. In a situation such as this, where all of the data was captured by a MARG either stationary or slowly moving in a plane, the bias can be well approximated as a constant. The conditions on the sensor were largely constant, meaning the bias will also be largely constant.

When integrated a constant bias manifests as a linear drift, which can be well estimated by two data points or data clusters. The MIMA models the bias as a linear function, leading to the suggestion that under the presented motion conditions the MIMA would remain accurate for an indefinitely long period of disturbance.

Two questions are raised by these results. The first question is how does a linear model for bias compensation perform under conditions of varying and dynamic motion? The answer to this question lies in the relationship between gyroscope bias and varying motion conditions. Even under stationary conditions, the bias is slightly nonlinear so it should not be surprising that the nonlinearity increases with varying motion conditions. Sensor acceleration, ambient temperature, and ambient pressure can all effect the bias (Shiau, Huang, & Chang, 2012). These relationships make modeling the bias under varying motion conditions a complicated endeavor. It is expected that as motion increases in variability, modeling the bias as linear will generate greater and greater error, though testing would need to be done to determine exactly how much. Regardless, it is clear that the bias compensation must include some form of modeling other than linear, and as such this issue is further discussed in Chapter 5.

And second, under the present motion conditions, how long must the undisturbed periods be to achieve accurate results? This question is more easily addressed than the previous one, and its answer can be estimated with data from this experiment. To address this question, in the next experiment the magnetic data was cropped to shorten the length of the undisturbed periods, and the accuracy was calculated.

4.4 Sensitivity Analysis

4.4.1 Objective

The results of the experiment discussed in Section 4.3 showed the MIMA was able to maintain accuracy despite spending only a small percentage of the overall time in an undisturbed area. This sparked an investigation into the MIMA's limiting factor for accuracy in the face of magnetic disturbance. The methodology of the MIMA's bias correction suggest that for the tested motion factors other than time in the undisturbed area affect its reparation ability. The purpose of this experiment was to identify those factors as well as their method of interaction.

4.4.2 Methods

In order to examine the factors affecting the MIMA's accuracy, the data from the previous experiment (Section 4.3) was examined, and the algorithm was stopped before bias compensation was applied. The uncompensated data was then plotted in relation to the magnetic data that had passed the filtering process, as this was the data on which both the conversion from yaw to heading and bias compensation was based. After examining the plotted data, a sensitivity analysis was done to assess the effects of altering the MIMA's test tolerances. The sensitivity of the HECT, CIAT, XY norm test, and Z norm test were investigated by varying their respective tolerance values as shown in Table 4.

Table 4: Information pertaining to the HIF test parameters investigated in the sensitivity analysis. The fixed values represent the values used for the experiment in Section 4.3, as well as the value that parameter assumed while not being varied.

Test	Fixed	Range	Parameter Meaning
HECT	0.1 deg/epoch	0 – 3 deg/epoch	Equates to the maximum allowed difference in the estimates of heading change of the gyroscope and magnetometer over the epoch (0.05 sec)
CIAT	.008 m/s ²	0 – .02 m/s ²	Equates to the minimum required change in acceleration over each time point (.0083 sec)
XY Norm	0.5 milliGauss	0 – 1 milliGauss	Equates to the absolute value of acceptable deviation from the designated undisturbed XY norm value
Z Norm	1.0 milliGauss	0 – 1 milliGauss	Equates to the absolute value of acceptable deviation from the designated undisturbed Z norm value

4.4.3 Results

Pertaining to the initial portion of the experiment, Figure 36 and Figure 37 detail when and how often magnetic data passed in this test. Here each black stars represent a single passing magnetic data point, and the blue line represents the pre-bias compensation heading. It was shown that in this test the magnetic data could be altered such that the initial 5.2 seconds and final 1.7 seconds of magnetic data could be removed without affecting the accuracy of the MIMA. This cropping resulted in .2 seconds of magnetic in the initial UA period, and 1.5 seconds in the final UA

period. The resulting error was .6 degrees at the start of the test, and 0.5 degrees at the end of test (Figure 38).

Pertaining to the sensitivity analysis, the results can be seen in Figure 39, Figure 40, and Figure 41. Each of the four tolerance values analyzed exhibited the same trend, in which the number of passing points initially increase quickly with increased tolerance range, but then reach a plateau and further increases result in the inclusion of few additional points.

Note that due to the formulation of acceleration test, the test parameter represents the minimum required change in acceleration between time points, and increasing its value leads to more stringent requirements. This is in contrast with the other three tests, where increasing the value of the test parameter results in less stringent requirements.

The CIAT was shown to affect the greatest number of points, excluding up to 464 points, compared to the 109, 35, and 32 points potentially excluded by the HECT, XY Norm Test, and Z Norm Test respectively.

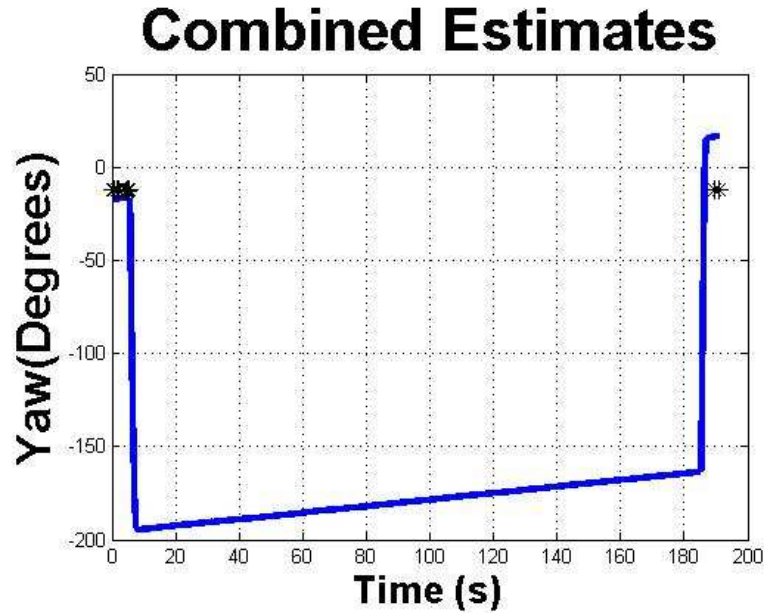


Figure 36: The pre- bias compensation heading data (blue line), juxtaposed with the magnetic data that passed the filter (black stars). In order to apply bias compensation a linear function will be added to the yaw so as to achieve optimal agreement with the magnetic data.




Figure 37: A zoomed in replica of the data shown in Figure 36, with the left hand side showing the initial period in the undisturbed area and the right hand side showing the final period in the undisturbed area. Together they represent the entirety of the time spent outside the disturbed area.

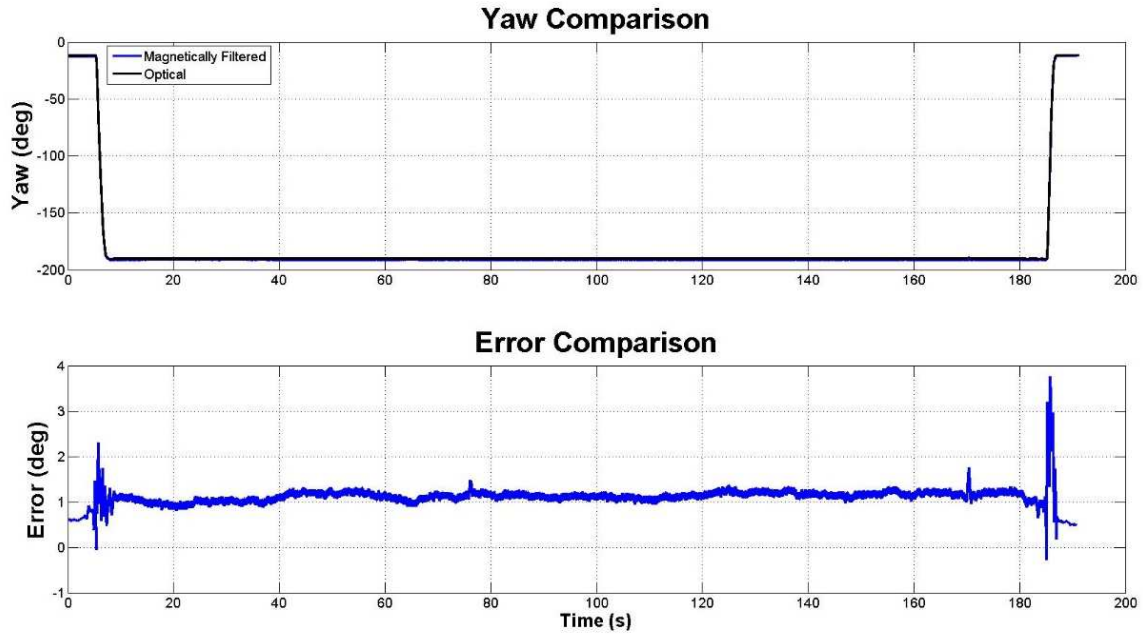


Figure 38: This graph compares resulting MIMA angles when the time spent in the undisturbed area was cropped such that the minimum time spent in it without compromising accuracy was achieved.

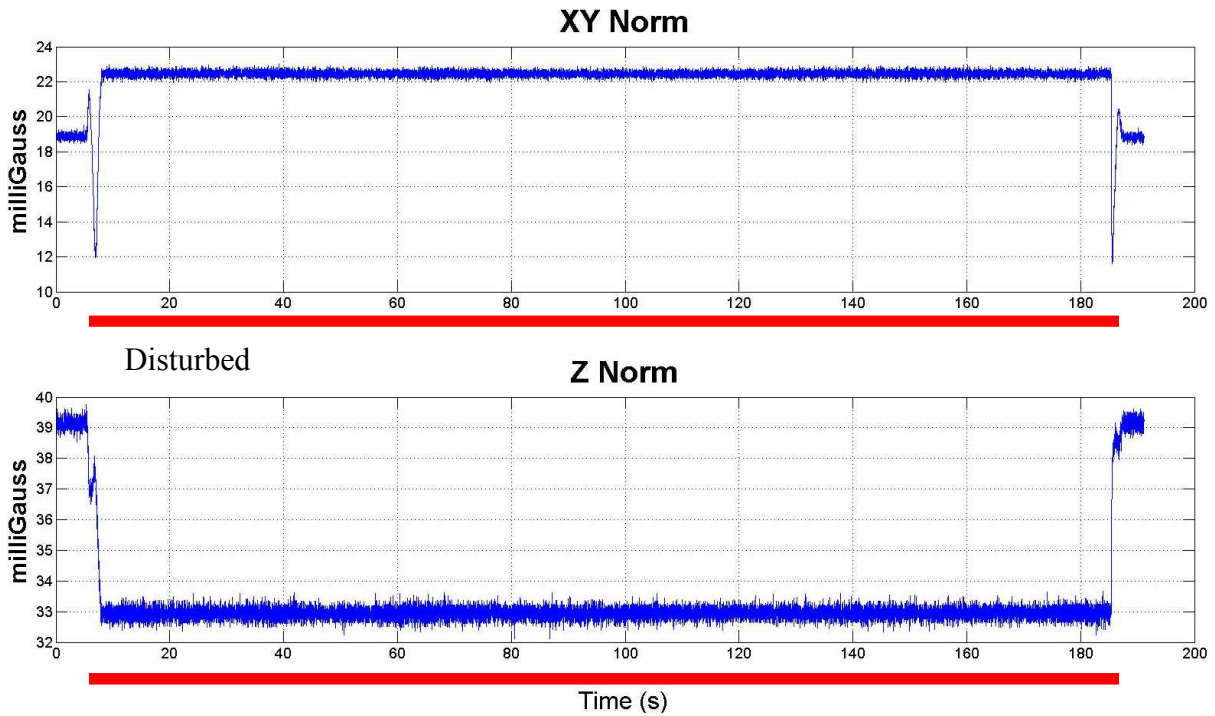


Figure 39: Here the crucial magnetic characteristics involved in filtering, the XY Norm and the Z Norm are shown. Here the measurement difference between disturbed and undisturbed periods can be seen.




Figure 40: This figure shows the effect altering the tolerance values has on how many magnetic data point are excluded. The graph on the left shows how varying the allowable difference between gyroscope and magnetometer heading (HECT) affects the number of passing points. Here an epoch is 0.05 seconds. The graph on the right shows how varying the required change in acceleration between time points, in this case .0083 seconds, affects the number of passing points.

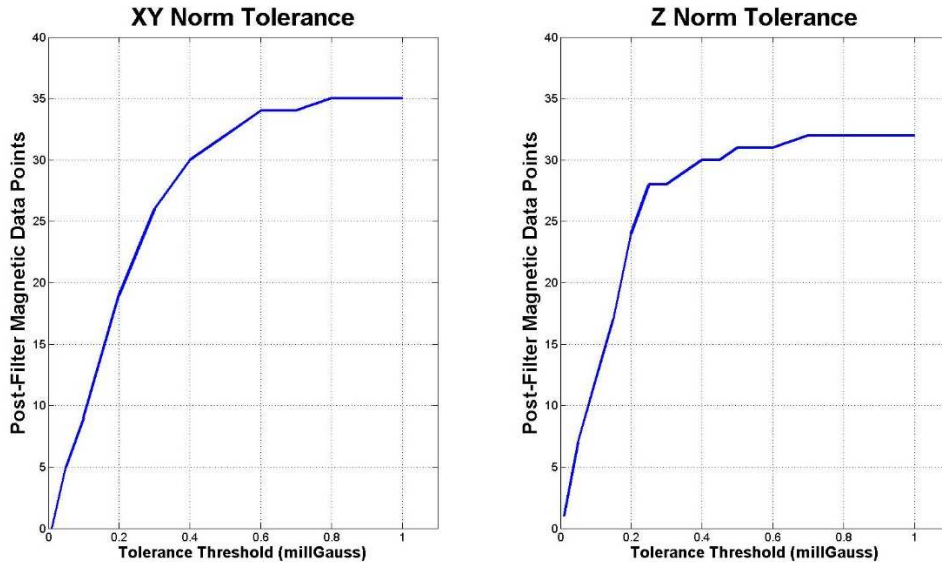


Figure 41: This figure shows the effect of varying the acceptable range of magnetic characteristics on the number of passing points. The left graphs shows the effect related to the XY norm, while the right graph is relative to the Z norm.

4.4.4 Discussion

The initial portion of this experiment revealed that in the given test, the time spent in the undisturbed area was not necessarily the factor limiting the MIMA's accuracy. Inspection of the magnetic data that passed the filtering test revealed that passing values tend to come in groups, such that several would be accumulate over a fraction of a second. The MIMA estimates bias error as a linear function, requiring only two data points to determine the slope and intercept. Though two points in close proximity could be used to determine compensation parameters, accuracy would tend to increase as the distance between the two points increased. Further, performing a linear fit between two groups of points would tend to be more accurate than a fit between two individual points.

Based on the logic just presented, the MIMA could accurately compensate for bias drift given two clusters of points of at least several seconds apart. Applying this to the data shown in Figure 36 and Figure 37, the initial period in the undisturbed area resulted in several clusters of filtered magnetic data, but the final period produced only one. The expectation would then be that time spent in the undisturbed area could be cropped such that only one point cluster from each period remained. This was found to be true, as excluding all other magnetic data outside of two such point clusters did not adversely affect accuracy.

Pertaining to the sensitivity analysis, the results indicate that each test parameter has an acceptable range that garners the majority of the magnetic data, and increases beyond this range have little effect. This plateau effect can be explained by the multi-dimensionality of the HIF. Each parameter was responsible only for excluding a portion of the data points, as the other parameters could be relied on to reject much of the magnetic data. This results in each parameter having a limit that once crossed has no effect on the included data. Above this limit, the test is not selective enough to operate on the portion of the data it is intended to, and the other tests have already excluded the remainder of the data. The sensitivity analysis revealed that for the given values of the other parameters, the acceleration parameter for the CIAT was the most sensitive to variation, and thus responsible for excluding the greatest number of points. This sensitivity can be exploited to increase accuracy of the MIMA, given knowledge about the testing environment.

The CIAT was introduced to address a weakness in the HECT, as during stationary periods both the gyroscope and magnetometer would predict no change heading, regardless of magnetic interference. The HECT was designed to identify disturbed data that was not excluded by the magnetic characteristic tests. This identification is only required if disturbed data is difficult to identify by its magnetic characteristics. However, if the majority of disturbed data can be excluded based on magnetic characteristics alone, the HECT need not be as stringent. In this case the CIAT can be relaxed, resulting in the inclusion of many more data points.

The sensitivity analysis also showed the effect of sensor noise on the HIF. Sensor noise results in a range of recorded values even while the sensor is stationary (Figure 39). This range must be included when determining the acceptable test ranges. Knowledge of noise behavior allows for the acceptable ranges to be set as tightly as possible without risking that any undisturbed data be unintentionally excluded. Under testing conditions where disturbed and undisturbed periods cannot be easily distinguished, the separate tests must work in a more interdependent manner and judicious choice of test ranges is imperative to the success of the HIF.

4.5 Comparison with Madgwick's Complete Orientation Filter

4.5.1 Objective

The previous experiments have compared the MIMA with what in the opinion of this author is the state-of-the-art in inertial motion capture, the KF formulation used by the XSENS MTw. For the sake of completeness, however, it was determined that the MIMA also be compared with the complete version of the orientation filter upon which it is based, the MOF. The MIMA utilizes the IMU filter developed by Madgwick et al., but a version of the filter adapted for magnetic data was also developed for MARG use that incorporates the magnetometer (Madgwick, Harison, & Vaidyanathan, 2011). The filter methodology is the same, except that the error gradient is calculated with magnetometer as well as accelerometer data. The purpose of this experiment is to compare the MIMA with this complete version of the MOF, henceforth referred to as the MARG Madgwick Orientation Filter (MMOF).

4.5.2 Methods

For this comparison the same methodology as in Section 4.3.2 was used, as well as the same raw data. However, instead of comparing the MIMA results with the XKF, the MMOF was used. As

for the other experiments in this thesis, the MMOF filter parameter was set equal to 0.1 (seen in Figure 19 as β). Furthermore, the starting value of the optical data was set to agree with the initial heading as determined by the XKF. This was done because the MMOF requires a short start-up period to stabilize, and its initial heading values are unreliable. It can be seen in Figure 42 that this stabilization period for this experiments was less than 0.5 seconds.

4.5.3 Results

The MMOF behaved similarly to the XKF, but drifted more quickly to achieve agreement between the gyroscope and magnetometer, and did so in a linear manner in contrast with the XKF's nonlinear manner. This expedited drift caused the MMOF to converge to the correct heading more quickly than the XKF in undisturbed environments, but also to diverge from the correct heading more quickly than the XKF in disturbed environments. When compared with the MIMA, the results were very similar to the comparison with the XKF. The calculated headings agreed when testing began in an undisturbed environment, but the MMOF accumulated error while in the disturbed area until its calculated heading agreed with that of the magnetometer. Once returned to the undisturbed area, the MMOF attempted to correct itself, but was only able to reduce the error accumulated error by approximately half, ending the test with an absolute error of 33.5 deg compared to the MIMA's 0.4 deg.




Figure 42: Comparison of the MIMA and MMOF, using the same data as in Figure 35.

4.5.4 Discussion

The MMOF behaved as expected, and generated heading values similar to those of the XKF. The major difference was the manner by which error accumulated. The XKF was able to recognize magnetic interference, and delayed its inclusion for a short time period. Once included the drift was nonlinear, a result of the complex KF calculations. The MMOF has no protocol for dealing with magnetic interference, and therefore when exposed to it the response was immediate. The gradient descent optimization method caused the heading to drift linearly so as to agree with the new heading. Here altering the value of β would have altered the slope of the drift, as it would represent a change in how prominently the magnetometer data was included. A smaller value would have decreased the slope, whereas a larger value would have increased it.

The comparison of the MIMA and the MMOF shows that MMOF is just as vulnerable to magnetic interference as the XKF, and lends credence to the generalized statement that any orientation filter incorporating magnetic data will be negatively affected by magnetic interference. If accurate data is to be gathered by a MARG in an area of magnetic interference, there must be a protocol for dealing with the resulting erroneous heading values, such as is exhibited in the MIMA.

CHAPTER 5 : FUTURE RESEARCH

The ultimate goal of the research presented in this thesis is to achieve accurate inertial motion capture data in magnetically disturbed environments. The algorithm developed in this work shows how data from the sensors within a MARG can be intelligently combined such that the effect magnetic interference is mitigated. In a similar fashion, the future vision for this research is the development of a full-body motion capture system capable of mitigating magnetic interference. Said mitigation would be accomplished via synergistic combination of data from a network of multiple, interconnected MARG sensors.

Just as the MIMA repairs disturbed data based on its relation to undisturbed data, a synergistic sensor network could repair data from a disturbed sensor based on its relation to the surrounding sensors it is interconnected with. This ability would take advantage of the local nature of magnetic interference, as a motion capture subject may encounter interference at his or her extremities, but not at their trunk. The undisturbed trunk data would stabilize the disturbed extremities.

As mentioned in Section 4.3.4, more advanced modelling of the gyroscope bias is a likely requirement for expansion of the MIMA to motion more complicated than the controlled, planar type experimented with in this work. A synergistic sensor approach could fulfill this requirement. Incorporating a biomechanical model into the sensor network would impose constraints on sensor inter-relation, which could be applied to improve the estimates of every sensor in the network. Furthermore, this approach would address not only error due to gyroscope bias, but also the sources of error accumulated in the expansion from single sensor measurement of a mechanical linkage to multi-sensor measurement of the human body. Such error sources include skin artifact, in which the sensor data is degraded by skin movement, and position errors due to double integration of accelerometer data, in which accelerometer bias is included in the double integration, causing an exponential increase in position estimate error.

The potential of this approach would be realized by an increase of acceptable motion capture environments, so as to include magnetically interfered areas. A mundane example would be motion capture on a treadmill, currently made problematic by the steel in the support bars. In that same line of thinking, motion capture on ATVs and bicycles would be improved. Applied to

the military, testing within vehicles and aircraft, areas of extreme magnetic interference, would be more feasible. IMC has grown into an attractive alternative to optical motion capture, and immunity to magnetic interference would make it even more practical and widely-applicable.

In sum the results of this work show that inertial motion capture can be improved by intelligent processing and combination of existing data. These encouraging results suggest that this methodology can be expanded to include more data streams in order to achieve the ultimate goal of inertial motion capture in environments of consistent and unavoidable magnetic interference.

APPENDIX

A.1 Basics of Magnetism

A.1.1 Magnetic Properties

Every object, from the largest star to the smallest molecule, is affected by magnetism and has its own magnetic properties. The most well-known of these properties is magnetic field, which can be defined as the space surrounding a magnet in which a magnetic force is experienced. A magnetic field is defined by a vector, meaning it has both a direction and a magnitude, and it is traditionally represented as a series of arrowed directional lines, known as field lines. The arrows on the lines show the direction of the field, pointing away from the northern magnetic pole and towards the southern magnetic pole (Figure A-1: Magnetic field lines emanating from the northern magnetic pole and returning at the southern magnetic pole.).

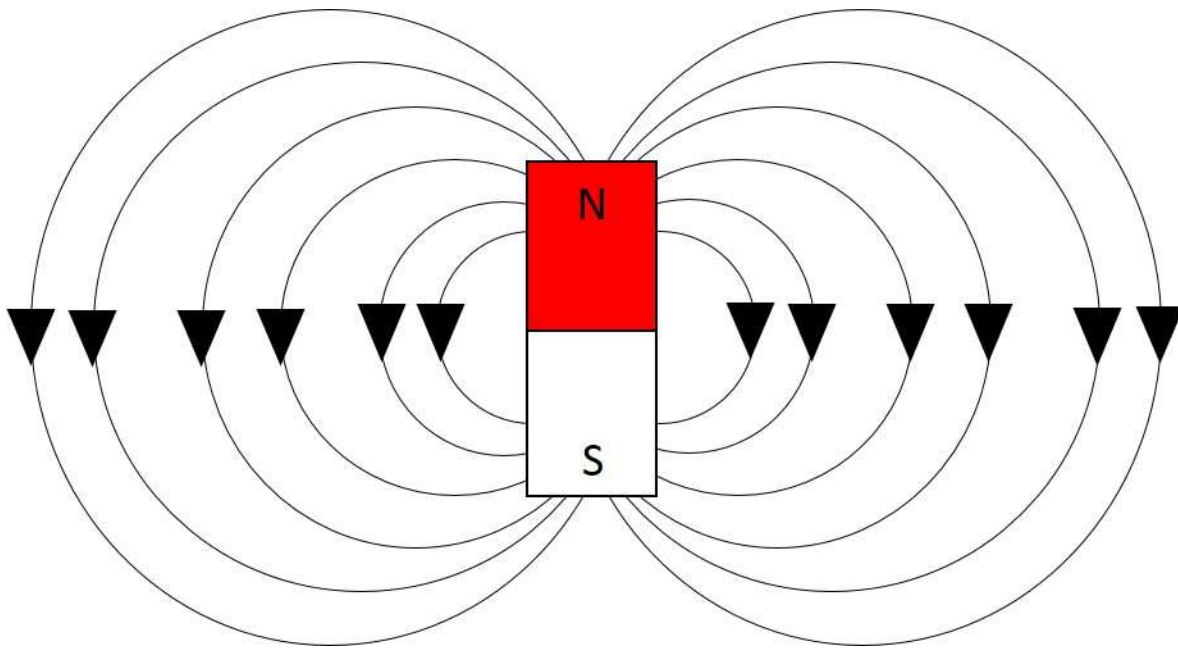


Figure A-1: Magnetic field lines emanating from the northern magnetic pole and returning at the southern magnetic pole.

Gauss's law for magnetism states that an individual magnetic field has divergence equal to zero, meaning that the magnetic flux out of any enclosed surface is and must be equal to the magnetic flux into said enclosed surface. More succinctly, the net magnetic flux over any enclosed surface

is equal to zero. This means that every magnetic pole must exist with an equal and opposite pole; a northern magnetic pole cannot exist without a southern magnetic pole to balance it out. If you were to take a bar magnet and split it in half, two new poles would form at the new ends of the broken pieces, and you would have two pairs of equal and opposite poles.

While every object does create its own magnetic field, most are orders of magnitude lower in strength than the fields created by dedicated magnets. How a material interacts with the existing magnetic field is generally more important than the field it creates. The interaction of a material with a magnetic field is in general governed by two properties: permeability and saturation level.

A.1.2 Magnetic Permeability

Magnetic permeability is the measure of a material's ability to support the formation of a magnetic field within itself. A good analogue is electrical conductivity. Magnetic permeability can be thought of as a measure of how well a material "conducts" magnetic fields through itself. In practice, we are concerned with relative permeability rather than absolute permeability. Relative permeability is how the material's permeability compares to that of a vacuum. Note that air has an absolute permeability that is nearly the same as a vacuum's. Therefore, air has a relative permeability of almost exactly one.

A.1.3 Magnetic Saturation

Magnetic saturation describes the point at which applying a stronger magnetic field to an object no longer causes a change in the object's magnetic properties. This is very similar to bottoming out the shocks on a car; the shocks can compress in order to absorb the shock of an impact, but once they are fully compressed, no more impact reduction can be achieved. Pertaining to shielding, saturation level of the shield must be considered; if the disturbance fully saturates a shield, any disturbance above the saturation level will not be mitigated by the shield.

A.1.4 Geomagnetic Field

Just as every object generates a magnetic field, so does the Earth. Though there are many theories and explanations, the prevailing one is dynamo theory (Stacey, 2008). According to this theory, the geomagnetic field is generated by the Earth's core via complex interactions of electrical currents, heat, and conductive materials. As such the planet can be modelled as a gigantic magnet with field lines leaving from the South Pole and entering in the North Pole.

A good way to conceptualize the geomagnetic field (Earth's magnetic field) is to think of Earth acting as if it has a gigantic bar magnet placed inside of it, roughly aligned with its geographic north-south axis, though offset by an angle of about 11.5 degrees (Basavaiah, 2011). The angle the magnetic axis makes with the geographic one is termed the declination angle, and this angle varies both with geographic location and time (Figure A-2: The magnetic declination angle is determined by the angle between the Earth's Magnetic and geographic Norths). Despite the variability in the declination angle, its calculation and prediction of the geomagnetic field has been made trivial by a variety of models and databases, such as those offered by Natural Resources Canada and the British Geological Survey.

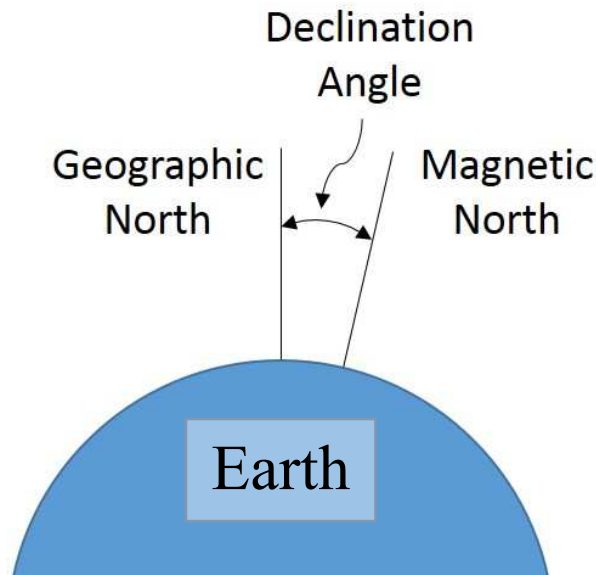


Figure A-2: The magnetic declination angle is determined by the angle between the Earth's Magnetic and geographic Norths.

It was previously mentioned that magnetic fields are defined by vectors, and this is no different for the geomagnetic field. Furthermore, the geomagnetic field is three-dimensional in nature, not simply pointing in the north-south direction as generally thought. At any point on Earth, the magnetic field can be represented with three mutually perpendicular components, which on the horizontal plane that makes up the Earth's surface can be represented via an up-down component, a left-right component, and a forward-backward component. These three components can be combined into a single direction and magnitude, and this direction is characterized by the dip or inclination angle (Basavaiah, 2011). The dip angle is the angle made by the horizontal plane of the Earth's surface and the Earth's magnetic field lines (Figure A-3:

The left image shows the three components (X Y Z) that make up a complete vector. The right image exhibits the magnetic dip angle, which is the angle at which the geomagnetic field penetrates the Earth's horizontal surface. As with the declination angle, the dip angle varies with both time and location. Furthermore, the magnitude of the geomagnetic field varies with time and location (Figure A-4: The magnitude of the geomagnetic field as a function of global location. The contour interval is 1000 nanoTeslas (1 milliGauss = 100 nanoTeslas). Reprinted from “*The US/UK World Magnetic Model for 2015-2020: Technical Report*”, by A. Chulliat, S. Macmillan, P. Alken, C. Beggan, M. Nair, B. Hamilton, A. Woods, V. Ridley, S. Maus, and A. Thomson, 2015, National Geophysical Data Center, NOAA.), ranging approximately from 65 microTesla at the Poles to 25 microTesla at the Equator (Stacey, 2008) (Basavaiah, 2011).




Figure A-3: The left image shows the three components (X Y Z) that make up a complete vector. The right image exhibits the magnetic dip angle, which is the angle at which the geomagnetic field penetrates the Earth's horizontal surface.




Figure A-4: The magnitude of the geomagnetic field as a function of global location. The contour interval is 1000 nanoTeslas (1 milliGauss = 100 nanoTeslas). Reprinted from “*The US/UK World Magnetic Model for 2015-2020: Technical Report*”, by A. Chulliat, S. Macmillan, P. Alken, C. Beggan, M. Nair, B. Hamilton, A. Woods, V. Ridley, S. Maus, and A. Thomson, 2015, National Geophysical Data Center, NOAA.

A.2 Characterizing Magnetic Disturbances

There are two main categories of magnetic disturbances that pertain to IMC. The first are hard iron disturbances, and the second are soft iron disturbances. In actuality, every disturbance is a combination of these two types, though one type usually dominates.

A.2.1 Hard Iron Disturbances

Hard iron disturbances are caused when a magnetized object is introduced to the environment, combining with the preexisting field (Figure A-5: Hard iron distortion are caused by a material possessing its own magnetic field combining with the ambient, geomagnetic one. In this example, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the hard disturbance is applied, signified by the central rectangle, the magnetic field near it is drastically bent, and the field is concentrated on one side and largely removed from the other. This disturbance is only an example of how hard iron can distort a magnetic field; there are many other possible forms and resulting distortions. This image was developed in the program QuickField . (Gebre-Egziabher, 2001). As with all magnetic fields, the field strength increases

with proximity to the magnetic object. If a hard iron disturbance is placed in a room, the magnitude of disturbance it produces varies with location around the disturbance, as well as with proximity. The location matters because at some places the hard iron magnetic field may be perpendicular with the preexisting field, causing maximum distortion, or the hard iron magnetic field could be parallel with the preexisting field, causing no distortion. Most likely, the actual effects will be somewhere in between. The magnitude of a hard iron disturbances is determined by and increases with the degree of magnetization. The characterization of hard iron disturbances is complicated by the fact that a ferromagnetic material placed within the geomagnetic field will be slowly magnetized, which with the exception of a select few instances is true of every ferromagnetic material (Basavaiah, 2011) (Cermakova, 2005). This means that every ferromagnetic material on Earth constitutes a hard iron disturbance, though their magnetic fields are generally negligible. Furthermore, this magnetization affect applies to all magnetic fields, not just the geomagnetic field. A ferromagnetic material placed in the vicinity of a magnet will also become magnetized.




Figure A-5: Hard iron distortion are caused by a material possessing its own magnetic field combining with the ambient, geomagnetic one. In this example, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the hard disturbance is applied, signified by the central rectangle, the magnetic field near it is drastically bent, and the field is concentrated on one side and largely removed from the other. This disturbance is only an example of how hard iron can distort a magnetic field; there are many other possible forms and resulting distortions. This image was developed in the program QuickField (Tera Analysis Ltd., 2013).

A.2.2 Soft Iron Disturbances

Unlike hard iron disturbances, soft iron disturbances do not have their own magnetic field. Instead soft iron disturbances provide a path of less resistance for magnetic field lines to travel through, distorting the preexisting field (Figure A-6: Soft iron distortion acts as a conductor of magnetic field lines, distorting the ambient field by providing a path of less resistance relative to the surroundings. As before, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the soft iron is applied, the field lines are funneled through it, reducing field

intensity on the sides of the disturbance and concentrating them at its ends. This image was developed in the program QuickField .) (Gebre-Egziabher, 2001). As with hard iron disturbances, the effects of the distortion vary with location and proximity. Factors that determine the magnitude of a soft iron disturbance are magnetic permeability and magnetic saturation. The more permeable a material is, the better it conducts magnetic field lines, and the greater the distortion it causes. The higher the saturation level of an object, the more magnetic field it can absorb, and the greater the distortion it causes. For most materials, the magnitude of the soft iron response is proportional to the external magnetic field (Gebre-Egziabher, 2001). More information on magnetic permeability and magnetic saturation can be found in the Appendices Sections A.1.2 and A.1.3 respectively.




Figure A-6: Soft iron distortion acts as a conductor of magnetic field lines, distorting the ambient field by providing a path of less resistance relative to the surroundings. As before, the undisturbed field consists of a series of parallel, equidistant, vertical lines. When the soft iron is applied, the field lines are funneled through it, reducing field intensity on the sides of the disturbance and concentrating

them at its ends. This image was developed in the program QuickField (Tera Analysis Ltd., 2013).

A.2.3 Frequency-Based Disturbances

Many electronic devices are powered by alternating current electricity, generally operating at 60 Hz. Because the movement of electrical charges induces a magnetic field, these electronics create magnetic fields of equal and opposite magnitude many times a second during operation. In general, the distortion caused by these is far smaller than those caused by hard or soft iron, though there are some exceptions, the most notable of which are power lines, which can create very large magnetic distortions (Dimcev, 29-31 May, 2000). However, power lines are generally easy to avoid, and if they cannot be avoided, a lowpass filter can be used on the data to remove any frequency-dependent magnetic field.

A.2.4 Magnetic Norm

Calculation of the magnetic norm is a simple way to detect the presence of any magnetic interference (Xsens, 2013). In an undisturbed geomagnetic field, the norm of the magnetic field is constant regardless of orientation, as it is a measure of the field magnitude only. Once the value of the norm has been established in an area free from magnetic disturbances, any norm measurement that deviates from this value signifies distortion. The norm is an effective predictor of disturbed data, both by its value and its variability. If the undisturbed norm value is known, this can be used to identify disturbances, but often this information is not available, or available only with limited certainty. A more robust method is to track the variation in the norm rather than its value. In an undisturbed field the norm is constant at all locations, therefore a varying norm signifies disturbance. The more often and greater said norm varies, the greater the magnitude of the disturbance.

A.3 Mitigating Magnetic Distortion

This section discusses distortion mitigation accomplished in two ways, calibration and disturbance modeling. The first method requires that the magnetometer remain constant in distance and orientation relative to the disturbance source, and the second method requires a magnetically controlled environment and rigorous modeling. In IMC none of these requirements are satisfied. These methods can reduce error due to magnetic interference in IMC, but they are far from being a complete solution. Regardless, their explanation offers insight into the

characteristics of magnetic interference, and this is the reason for their inclusion.

A.3.1 Hard Iron Disturbances

Hard iron disturbances are common and problem-causing, but under the right circumstances calibration can remove the effects of hard iron. Take the example of a magnetometer mounted on a magnetized metal block. Without calibration, accurate measurement would be impossible due to the distorting nature of the magnetized block. However, because the block produces its own independent field, and the magnetometer is mounted to the block, remaining in the same relative position regardless of the orientation of the block, calibration can be applied to effectively remove any distorting effects of the block. Say you were to take an undisturbed magnetometer and rotate it 360 degrees whilst remaining in the same location. Because the norm is the same for every orientation, a plot of any two components as a function of degrees rotated would result in a perfect circle with its center at the origin. Furthermore, because the magnetometer is mounted onto the block, it receives the same magnetic field from the block regardless of the orientation. Therefore, rotating the block and magnetometer in an otherwise undisturbed field would result in a perfect circle with its center some distance from the origin. The hard iron adds a constant bias to the magnetometers results, resulting in a shifted plot. Applying an equal and opposite bias will return the circle's center to the origin, and all future measurements would be distortion-free (Figure A-7: Hard iron distortion adds a constant magnitude to every orientation, and by subtracting said constant the distortion can be removed.).




Figure A-7: Hard iron distortion adds a constant magnitude to every orientation, and by subtracting said constant the distortion can be removed.

A.3.2 Soft Iron Disturbances

Soft iron disturbances are more common than hard iron disturbances. As with hard iron, under the right circumstances, soft iron disturbances can be removed. Take the example of a magnetometer mounted onto a high-permeability, high-saturation, metal block that is completely free of magnetization. If we were to repeat the rotation experiment done with the soft iron, the plot of any two magnetic components vs degrees rotated would be a perfect ellipse with its center located at the origin. An ellipse results because the distortion redirects field lines so as to concentrate them in some locations and reduce the, in others. Thus, the major axis of the ellipse will represent norms greater than the undisturbed, and the minor axis will represent norms smaller than the undisturbed. By mapping the ellipse into a circle and applying the mapping function to all incoming data, the soft iron disturbance is removed (Figure A-8: Soft iron distortion concentrates the field at some angles and reduces it at others, producing an elliptical shape. Calibration can be applied by mapping the ellipse into a circle.).




Figure A-8: Soft iron distortion concentrates the field at some angles and reduces it at others, producing an elliptical shape. Calibration can be applied by mapping the ellipse into a circle.

A.3.3 Combination Iron Disturbances

In practice there is no such thing as a purely hard iron disturbance or a purely soft iron disturbance. Disturbances are always encountered as combination of the two types, though usually the majority of the disturbance is due to one or the other. An example of a hard iron disturbance is a permanent magnet. The magnet generates its own magnetic field, but because it is made of a material with a high magnetic permeability, it will also distort the preexisting field around it. Conversely, a soft iron disturbance could be caused by a steel plate, but due to its ferromagnetic nature the plate will invariably become somewhat magnetized over time, either as a result of exposure to local magnetic fields or the geomagnetic field.

In order to deal with combination disturbances, a combination calibration procedure must be used. By combining the process for hard and soft calibration, combination disturbances can be mitigated as well. Rotation of a magnetometer mounted on a combination disturbance will result

in a norm-vs-rotation angle plot that is an ellipse with its center displaced some distance from the origin. As with hard distortion, the offset from the origin must first be determined and corrected, and as with soft distortion, the elliptical shape must be mapped into a circle (Figure A-9: Disturbances generally have components of hard and soft iron distortion. The disturbance can be removed via hard iron calibration followed by soft iron calibration.). Once this is completed, the combination disturbance has been mitigated.

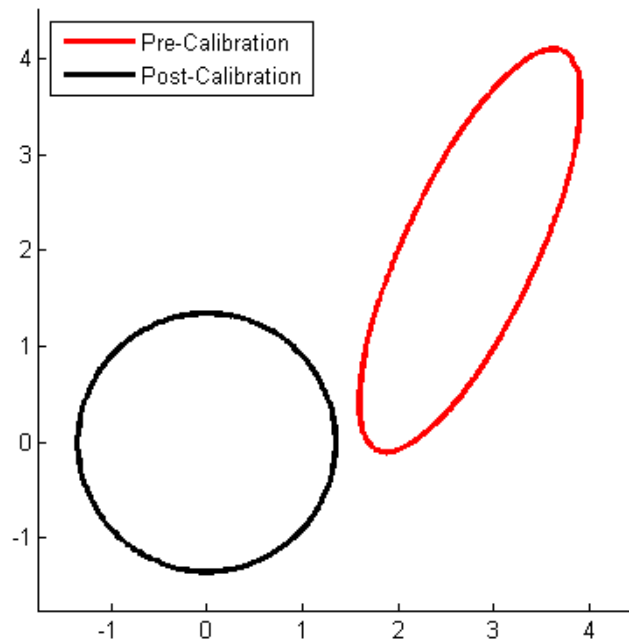


Figure A-9: Disturbances generally have components of hard and soft iron distortion. The disturbance can be removed via hard iron calibration followed by soft iron calibration.

A.3.4 Time-Varying Disturbances

Time-varying disturbances are magnetic disturbances that change as time passes. One strategy for dealing with time-varying disturbances is to periodically perform calibration via the method previously explained. The more often the calibration procedure is performed, the more accurate the overall data will be because the bias and scaling of the magnetometer data are being constantly adjusted. In IMC this method can reduce error, but the number of calibrations required is often prohibitive, as it requires a break in the motion to perform the calibration. Often a balance between accuracy and number of calibrations is found, and the remaining error is tolerated (Figure A- 10: Periodic calibration can be used to stabilize a data stream. This image is

a simulation of how periodic calibration can be used to stabilize drift from a noisy signal. The red line is the noisy data that is periodically calibrated. The blue line is what the signal would look like without any noise, i.e., it is the true signal.).

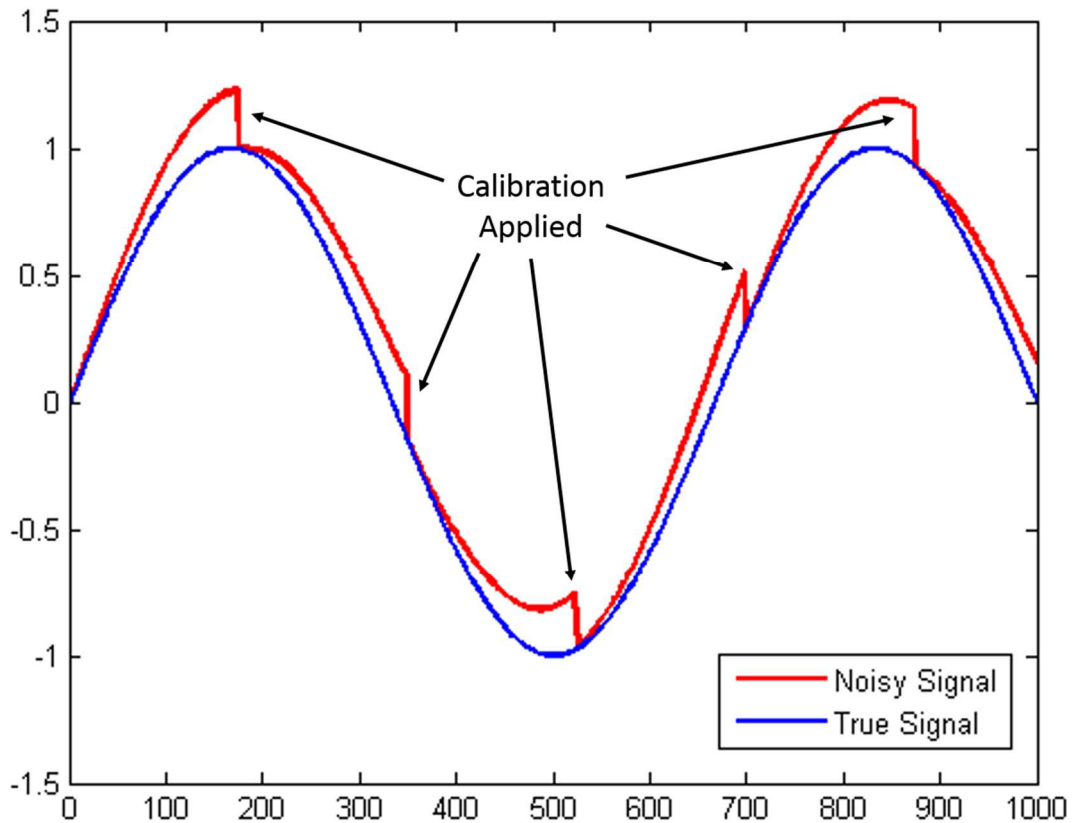


Figure A- 10: Periodic calibration can be used to stabilize a data stream. This image is a simulation of how periodic calibration can be used to stabilize drift from a noisy signal. The red line is the noisy data that is periodically calibrated. The blue line is what the signal would look like without any noise, i.e., it is the true signal.

A.4 Software/Hardware Enhancements

A.4.1 Disturbance Modeling

Another strategy for mitigating time-varying magnetic disturbances is to develop a model of the disturbance that can act like continuously updating calibration based on the predicted disturbance. If all the factors in the data-capture environment that affect the local magnetic field,

can be identified, even if an exact relationship between the disturbance and magnetic field distortion can be established, a model can be created. By taking data in a multitude of situations and gathering data on each factor, the problem can be made into one of optimization, and a relationship can be determined. An example of this is the magnetic calibration model developed by Springmann and Cutler for the electronics onboard a spacecraft (Cutler, 2012). By modeling the time-varying factors, namely the magnetic field due to the solar panels, they were able to determine a relationship that predicted the effect the solar panels would have on the magnetometer readings. Once this was determined, they could alter the calibration to fit the changing magnetic disturbances (Cutler, 2012).

This method has several drawbacks, the most problematic of which is identification of all relevant factors. Given a closed and highly controlled environment, such as a spacecraft, this method is feasible. However, in a real-world testing environment it becomes far more complicated; dozens of relationships corresponding to every feasible source of interference would need to be established. Furthermore, the model is restricted to use strictly in the environment it was designed for – it is not adaptive.

To gauge the complexity of such a modeling task, consider that Johnston and Stacey found that a vehicle passing within 10 meters of a magnetometer could register a .1 milligauss change in the sensed field. Given that magnetic field strength grows exponentially with decreasing distance from the source, something significantly smaller than a car that is only a few meters away can have a far greater effect than the car (Stacey, 1968). Applying the knowledge to the motion capture environments reveals the scope of the modeling process. In addition to the disturbances inside the room, wall supports, floor reinforcements, and electrical wiring would need to be accounted for. For simple, controlled systems this method may be sufficient, but in a complicated environment such as that encountered in IMC it is impractical.

A.4.2 Sensor Enhancements

Instead of relying solely on software methods to mitigate disturbances, hardware alterations can also be made. One approach is to increase the number of magnetometers gathering data at a specific location. Afzal et al. developed a method of detecting and compensating for magnetic perturbations by combining data from several magnetometers placed very close to each other but

at different orientations (Afzal, Renaudin, & Lachapelle, September 2010). The magnetometer with the least detected distortion was chosen and used in conjunction with a Kalman filter to calculate a corrected heading, providing a near real-time calibration method for heading determination (Afzal, Renaudin, & Lachapelle, September 2010). This method is effective at disturbance mitigation, but as it currently stands, the required hardware is too bulky for inertial motion capture.

A.5 Magnetic Shielding

A.5.1 Methodology

The term “magnetic shielding” can be misleading, given that it is impossible to truly block magnetic field lines (MuShield Magnetic Shielding, 1996-2015). Per Gauss’s law for magnetism, the magnetic flux into and out of an object must be the same, so the magnetic field lines leaving one pole must eventually reach the other; however, the path that the lines take can be altered. Magnetic shielding is accomplished by introducing a high-permeability, high-saturation material into the environment so as to redirect the magnetic disturbances away from the shielded area (Gabrielson, 1993; MuShield Magnetic Shielding, 1996-2015). The shielding material introduces a path of least resistance for the field lines to travel through, thereby increasing the magnetic flux through the shield and decreasing the magnetic flux in the area surrounding the shield (Figure A-11: A hard iron disturbance can have a large area of effect, disturbing the surrounding area in unpredictable ways. The chosen disturbance is simple for demonstration purposes, real-life disturbances can be more complicated. Image was developed in the program QuickField . and Figure A-12: Introducing a magnetic shield into the environment can absorb and redirect the magnetic field lines emanating from the disturbance, leaving the surrounding area unaffected. Notice, however, that the field lines are now concentrated heavily at the ends of the shield. This area is now more heavily distorted than before, so it is imperative that this redirection occurs into the correct area. This image was developed in the program QuickField). The similarities between magnetic shields and soft iron disturbances are striking, and this is to be expected because a magnetic shield is simply a soft iron disturbance with properties that we can control. By varying several properties of the shielding material, such as shape, size, permeability, saturation, or disturbance proximity, it is possible to favorably alter the path of magnetic field lines. In a controlled environment with a

highly specific goal, such shielding is generally a straightforward task, but in an uncontrolled environment with several different goals, shielding methodology can get very complicated.




Figure A-11: A hard iron disturbance can have a large area of effect, disturbing the surrounding area in unpredictable ways. The chosen disturbance is simple for demonstration purposes, real-life disturbances can be more complicated. Image was developed in the program QuickField (Tera Analysis Ltd., 2013).




Figure A-12: Introducing a magnetic shield into the environment can absorb and redirect the magnetic field lines emanating from the disturbance, leaving the surrounding area unaffected. Notice, however, that the field lines are now concentrated heavily at the ends of the shield. This area is now more heavily distorted than before, so it is imperative that this redirection occurs into the correct area. This image was developed in the program QuickField (Tera Analysis Ltd., 2013)

A.6 Factors Affecting Shielding Capacity

It is difficult to develop a definitive relationship between shield parameters and shielding capacity, mainly because shielding capacity is a broad term. Shields can be employed to accomplish a variety of tasks, from shielding sensitive instruments from high-frequency electromagnetic interference to shielding the construction environment of a satellite to protect it from magnetic forces in space. For the purposes of this document, shielding capacity will be defined as the ability of a shield to mitigate sources of magnetic disturbance that are commonly found in inertial motion capture labs. With shielding capacity narrowed in scope, permeability and size/shape are the major determinants of shielding capacity. If a shield does not have a high enough permeability, magnetic field lines will not be drawn into strongly enough to appreciably alter the area surrounding the field. However, if the permeability is too high, the area affected by

the shield can grow too large, actually introducing magnetic distortion in the form of a strong soft iron disturbance (Figure A-13: Permeability of an object is proportional to its effect on the field. In this figure, the object on the left has a lower permeability value than the object on the right. The effect of this is evident in the increased number of field lines passing through the object, as well as greater curving of the field lines, in response to the higher permeability object. These images were developed in the program QuickField .).




Figure A-13: Permeability of an object is proportional to its effect on the field. In this figure, the object on the left has a lower permeability value than the object on the right. The effect of this is evident in the increased number of field lines passing through the object, as well as greater curving of the field lines, in response to the higher permeability object. These images were developed in the program QuickField (Tera Analysis Ltd., 2013).

Volume has a similar effect as permeability: the larger the volume, the more the surrounding magnetic field lines are drawn into the shield. Conversely, shape has a far more complicated relationship with shielding capacity. The shape with the highest potential for shielding capacity is a sphere, followed by a cylinder, but for the purposes of inertial motion capture, these shapes are generally of little use, as will be discussed in the next section. The general underlying theme in shield shapes is that the smaller the radius of curvature of a surface, the less effective it is at shielding. The more angular and cornered a shield is, the lower its shielding capacity. Magnetic field does not turn corners well, so abrupt changes in the direction the shield is conducting can lead to flux leakage (MuShield Magnetic Shielding, 1996-2015). Flat shields can circumvent

this problem and are often the best choice of shape for practical use in inertial motion capture. However, magnetic shields have no effect on magnetic field lines that enter perpendicularly to the shield surface, and flat shields are particularly vulnerable to this. The more perpendicular the field's lines are to the flat shield, the less effective it is (MuShield Magnetic Shielding, 1996-2015).

Another factor that can have a large effect is the saturation capacity of the shielding material. Regardless of the other qualities of the shield, if the material saturates too quickly it will not be effective. Once a material has reached its saturation limit and can absorb no more magnetic field lines, it can have no further effect on the surrounding field, regardless of increasing disturbance magnitude. It is difficult to practically obtain a material that has both a high permeability and a high saturation point, and a balance between the two properties must be found that optimizes shielding capacity.

A.7 Shielding Placement

Although the preceding factors are the main determinants in a material's shielding capacity, the effectiveness of a shield, in relation to inertial motion capture, is highly dependent on its placement in the environment relative to both the source disturbance and the magnetic sensors. As previously mentioned, magnetic field can only be redirected, not extinguished. Therefore, magnetic shielding is simply a process of altering the magnetic properties of an environment in a way that removes the magnetic distortion from an area of interest by moving and condensing it into a designated area. In order to apply this process to the inertial motion capture environment, one must possess a detailed knowledge of both the magnetic properties of the environment and the inertial system's ability to process and mitigate sensed disturbances.

As previously mentioned, the highest shielding capacity is achieved via a spherical or cylindrical shield, but this shape is highly impractical, if not impossible, to apply. A common source of disturbance is metal support bars in the floor of the testing room; there is no way to fully enclose these bars in a shield, and applying even a half cylinder over them will render that area unusable for testing as it will provide a physical obstacle that the testing subject must avoid. Generally, the only practical method of applying shielding to the floor is with a flat sheet, and this shape has a low shielding capacity. Other sources of disturbances, such as computer equipment, can be

completely enclosed in a spherical or cylindrical shield, but this would require a custom shield for each piece of equipment, and the amount of shielding material would likely be prohibitively expensive. A more practical approach in this situation is to consolidate the equipment and shield only the pieces that are the causing the greatest disturbance.

A.8 Other Considered Mitigation Methods

A.8.1 Sensor Shielding

The most obvious solution to the problem of magnetic interference is to prevent the effects of the disturbance from reaching the magnetometer. This would reason that the sensors themselves should be shielded, not the disturbance source. If the magnetometer is encased in shielding material, would it not be immune to disturbances? Despite this reasoning, sensor shielding is not feasible for two reasons. First, magnetic field lines cannot be blocked, only redirected, so placing a shield around the sensor would actually increase the amount of disturbance the sensor is subjected to. It may help to negate the effects of surrounding disturbances, but it would introduce a new disturbance that would be almost invariably worse. Second, there is no way to separate magnetic field lines emanating from a disturbance from the geomagnetic field lines. Tracking the norm allows for detection of disturbances, but provides us with no information as to the nature of the disturbance. Therefore, even if magnetic field lines could be blocked, shielding the sensor would block the desired geomagnetic field along with the disturbance.

A.8.2 High-Strength Artificial Heading

A major source of the interference problem is that the geomagnetic field is weak in comparison with the magnetic disturbances. Introducing an artificial magnetic heading of strength an order of magnitude greater than the Earth's could feasibly remove this problem. However, generation of such an artificial heading would be highly impractical. Measuring the geomagnetic field provides a constant heading when measured over a very large area, which is what makes it so useful. However, this is only the case because the Earth's magnetic field spans such a vast distance. When traveling across significant portions of the Earth, on the order of hundreds of miles, the measured heading will change perceptibly. Because motion capture is generally performed over much smaller distances, it can be assumed that the geomagnetic heading remains constant regardless of location. Reproduction of this constant heading direction is where the

difficulty lies, as placing a super-strong magnet at one end of the room would fail to do this. Instead we must consider a device such as the Helmholtz coil. By symmetrically placing a pair of solenoids along a common axis, a region of uniform magnetic field is created in the area surrounding said axis (Figure A-14: A Helmholtz coil is a made up of two solenoids and generates a uniform magnetic field within a portion of the area enclosed by the solenoid loops.). By manipulating the properties of the solenoids, the volume of uniform field can be manipulated, as well as its strength.




Figure A-14: A Helmholtz coil is a made up of two solenoids and generates a uniform magnetic field within a portion of the area enclosed by the solenoid loops.

Theoretically, this is a promising concept, but in practice it is not feasible for practical use with motion capture. The size and power requirements of such a coil would be enormous, and implementation would be very expensive. Furthermore, this would only mitigate hard iron disturbances. Soft iron disturbances would continue to absorb and redirect the field, and would do so at a magnitude proportional to that of the artificial heading.

A.9 Shielding Experimentation

A.9.1 Variability in Yaw Drift

A.9.1.1 Objective

The purpose of this experiment was to compare yaw drift across a number of sensors in order to gain insight into how magnetic data was converted to heading data by the MARG. Placing five sensors in the same, undisturbed area and comparing the calculated headings would reveal how the filtering algorithm behaved.

A.9.1.2 Experimental Setup

The purpose of this experiment was to compare yaw drift across a number of sensors. Five sensors were taken to an open space, over a dozen meters from any building, and used to record the yaw orientation when the sensors were completely stationary (Figure A-15: The environment chosen for static drift test was outside in a small field. It was chosen to avoid the magnetic disturbances found inside buildings, as well as to minimize the disturbance from any other ferromagnetic objects.). Sensors one and four were placed in the grass, and the rest were placed on the sidewalk a few feet away (Figure A-16: Two sensors were placed on the grass, and three were placed on the concrete.). This was done in case one of the locations was affected by an unknown magnetic disturbance, though the small difference in location seemed to have negligible effect.



Figure A-15: The environment chosen for static drift test was outside in a small field. It was chosen to avoid the magnetic disturbances found inside buildings, as well as to minimize the disturbance from any other ferromagnetic objects.




Figure A-16: Two sensors were placed on the grass, and three were placed on the concrete.

A.9.1.3 Data Collection and Processing

The only steps involved in collection and processing were sensor initialization and recording and exportation of the recorded orientation data. Once done, the orientation data from each sensor was set to begin at the same location for easy comparison, and the results were plotted.

A.9.1.4 Results

The yaw orientation of every sensor remained within approximately two degrees of each other. No difference between the sensors placed on the concrete and those placed in the grass was found. Furthermore, all five sensors calculated erratic motion (Figure A-17: The undisturbed static drift test shows that while all five sensors remained within an approximately two-degree range, their drift was random and chaotic.).

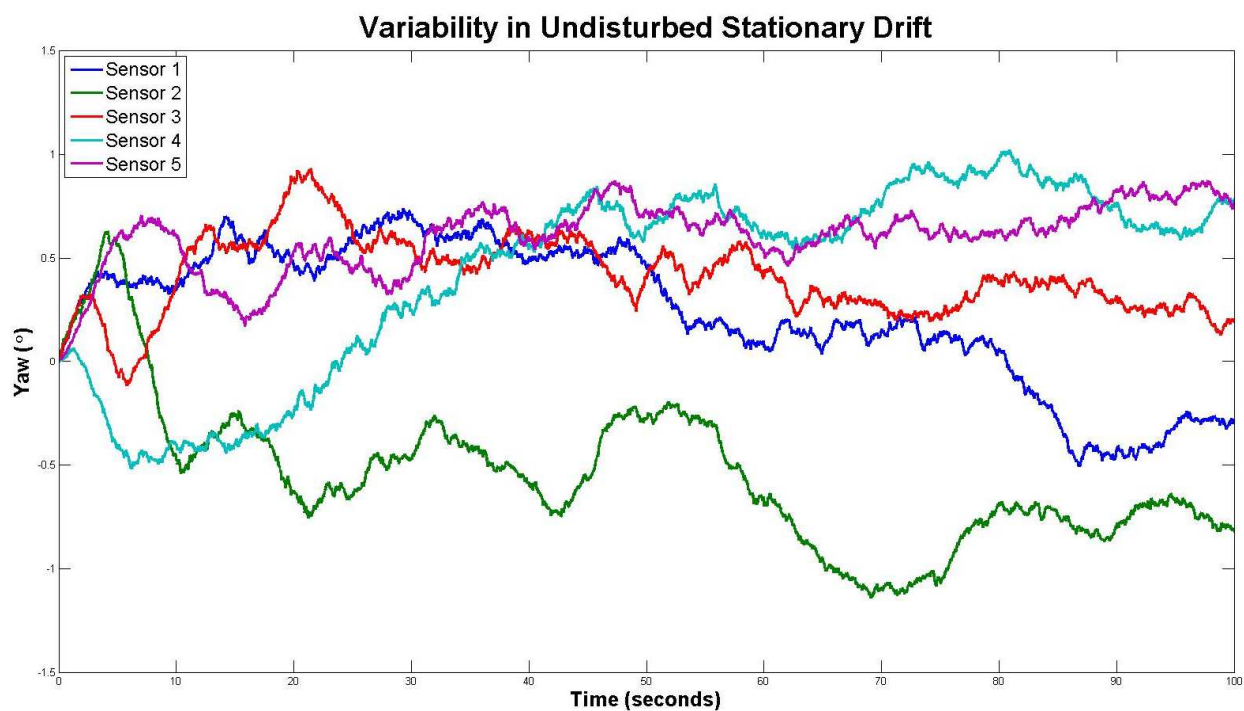


Figure A-17: The undisturbed static drift test shows that while all five sensors remained within an approximately two-degree range, their drift was random and chaotic.

A.9.1.5 Discussion

The main conclusion that can be drawn from these results is that orientation drift can cause significant errors in orientation. While these sensors all calculated data that was within one degree of its starting orientation, they were in a nearly ideal environment. It is easy to see how linking several of these sensors together can easily lead to errors of a few degrees, not even including the effects of magnetic interference. Furthermore, the drift is at least somewhat random and showed little sign of stabilization over the 100 seconds of data collection.

A.9.2 Effects of Magnetic Shield Layering

A.9.2.1 Objective

The purpose of this test was to determine the effect of layering on a magnetic shield's effectiveness. Each shield is of different thickness and/or material and therefore has a unique effect on the magnetic field. It logically follows that combining two shields would shield differently than each shield on its own.

A.9.2.2 Experimental Setup

The setup from the previous shielding experiment was replicated, the exception being which shields were used. Only the three high-permeability shields (of thickness .004", .006", and .010") were tested, leaving six possible permutation pairs – all of which were tested.

A.9.2.3 Data Collection and Processing

This process was identical to that of the other, non-layered shielding experiment.

A.9.2.4 Results

For clarity in explanation of these results, 4 under 6 means that the .004" shield was placed underneath (closer to the disturbance source) the .006" shield in the layering configuration, and the same pattern was followed for the other combinations (Figure A-18: This plot shows the magnetic norm as a function of shield layering and angular location of the MARG. Only high-permeability shields were used, and the numbers in the legends refer to thousandths of an inch, i.e., 4 under 6 means that the .004" shield was placed underneath (closer to the disturbance) the .006" shield in the layering configuration.). Layering the shields was found to affect their shielding capacity, as measured by alteration of the norm. As in the previous shielding experiment, the goal was to minimize variability in the norm. The results showed that the 10

under 4 case produced the least variable norm, and the 10 under 6 case produced the most variable norm.

A.9.2.5 Discussion

The results show that layering has an effect on shielding capacity. When comparing the order in which the shields are placed, there is an optimal choice. The most normalizing combination was 10 under 4, whereas the least normalizing combination was 10 under 6, despite having a shield thickness difference of only two thousandths of an inch. The conclusion is that small differences in shield parameters can have a large impact on the shield's efficacy, and that layering is another factor that must be considered when designing a magnetic shield.

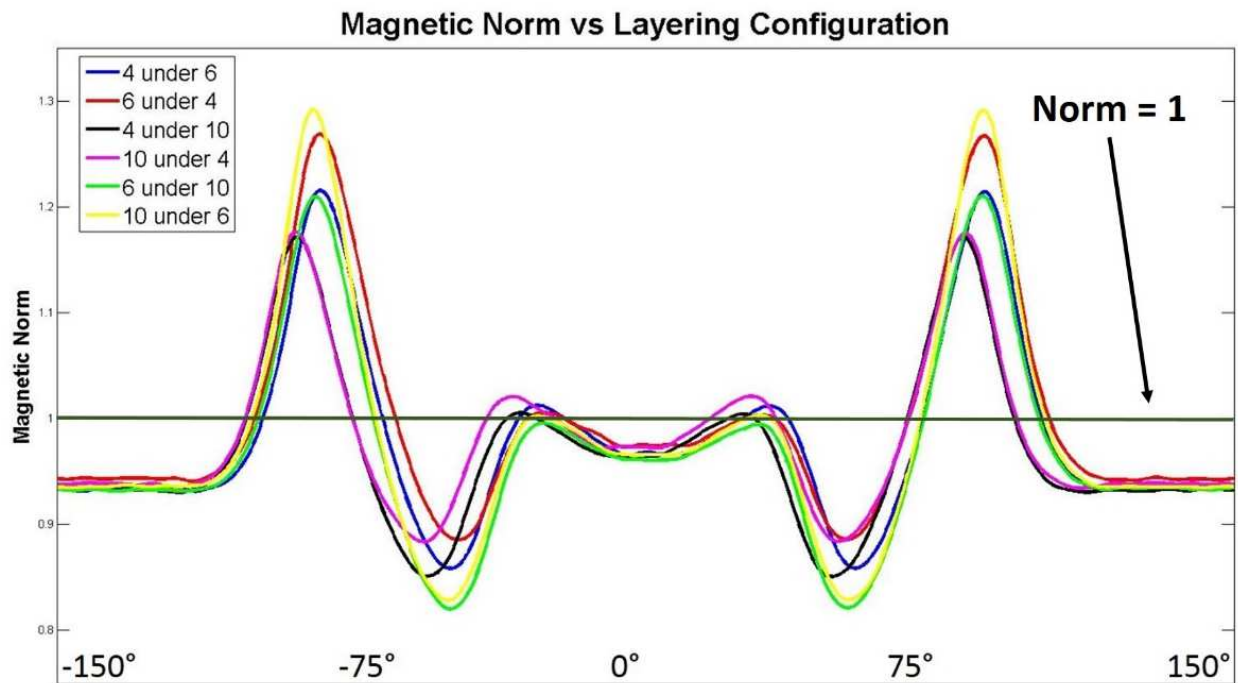


Figure A-18: This plot shows the magnetic norm as a function of shield layering and angular location of the MARG. Only high-permeability shields were used, and the numbers in the legends refer to thousandths of an inch, i.e., 4 under 6 means that the .004" shield was placed underneath (closer to the disturbance) the .006" shield in the layering configuration.

REFERENCES

- Afzal, M. H., Renaudin, V., & Lachapelle, G. (September 2010). Assessment of Indoor Magnetic Field Anomalies using Multiple Magnetometers. *Proceedings of the 23rd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2010)*, (pp. 525-533). Portland, Oregon.
- Basavaiah, N. (2011). *Geomagnetism - Solid Earth and Upper Atmosphere Perspectives*. Springer and Capital Publishing Company.
- Bird, J., & Arden, J. (2011, April). Indoor Navigation with Foot-Mounted Strapdown Inertial Navigation and Magnetic Sensors [Emerging Opportunities for Localization and Tracking]. *IEEE Wireless Communications*, 18(2), 28-35.
- Caruso, M. J. (2000). *Application of Magnetic Sensors for Low Cost Compass Systems*. Technical Report, Honeywell, SSEC.
- Cermakova, E. (2005). Magnetization of Steel Building Materials and Structures in the Natural Geomagnetic Field. *Acta Polytechnica*, 45(6), 47-52.
- Cuesta-Vargas, A. I. (2010). The use of inertial sensor system for human motion analysis. *Physical Therapy Reviews*, 15(6), 462-473.
- Cutler, J. C. (2012, July-August). Attitude-Independent Magnetometer Calibration With Time-Varying Bias. *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 4, 1080-1088.
- Dimcev, R. S. (29-31 May, 2000). Measurement of Power Frequency Magnetic Fields in Different Environments. *Electrotechnical Conference, 2000. MELECON 2000. 10th Mediterranean*. 3, pp. 1107-1110. IEEE.
- El-Gohary, M. A. (2013). *Joint Angle Tracking with Inertial Sensors*. PhD Thesis, Portland State University.
- Faulkner, T. W., Alwood, R., Taylor, D. W., & Bohlin, J. (2010). GPS-denied Pedestrian Tracking in Indoor Environments Using an IMU and Magnetic Compass. (pp. 198-204). San Diego, CA: Proceedings of the 2010 International Technical Meeting of The Institute of Navigation.
- Fong, D. T.-P., & Chan, Y.-Y. (2010). The Use of Wearable Inertial Motion Sensors in Human Lower Limb Biomechanics Studies: A Systematic Review. *Sensors*, 11556-11565.
- Gabrielson, B. C. (1993). *Non-permeable Cable Magnetic Shielding Properties*. Chesapeake Beach, MD: Security Engineering Services.

- Gebre-Egziabher, D. G. (2001). *A Non-Linear, Two-Step Estimation Algorithm for Calibrating Solid-State Strapdown Magnetometers*. Department of Aeronautics and Astronautics at Stanford University.
- Kok, M. J. (2012). Calibration of a magnetometer in combination with inertial sensors. *Information Fusion (FUSION), 2012 15th International Conference on* (pp. 787-793). Singapore: IEEE.
- Konvalin, C. (2008a). *Calculating Heading, Elevation, and Bank Angle*. MEMSense Technical Document, MEMSense.
- Konvalin, C. (2008b). *Compensating for Tilt, Hard Iron and Soft Iron Effects*. MEMSense Technical Document, MEMSense.
- Luinge, H. (2002). *Inertial Sensing of Human Movement*. PhD Thesis, University of Twente.
- Lutters, R. D. (30-31 March, 2009). The Drift of the Xsens Moven Motion Capturing Suit during Common Movements in a Working Environment. *Proceedings of the 19th CIRP Design Conference - Competitive Design* (p. 338). Cranfield University.
- Madgwick, S. O. (2010, April 30). MADGWICKAHRS Implementation of Madgwick's IMU and AHRS algorithms. (Matlab, Compiler) Retrieved October 14, 2104, from <http://www.x-io.co.uk/open-source-imu-and-ahrs-algorithms/>
- Madgwick, S. O., Harison, A. J., & Vaidyanathan, R. (2011, June 29 - July 1). Estimation of IMU and MARG orientation using a gradient descent algorithm. *Rehabilitation Robotics (ICORR), 2011 IEEE International Conference on Rehabilitation Robotics* (pp. 1-7). Zurich, Switzerland: IEEE.
- Maus, S. S. (2015). *The US/UK World Magnetic Model for 2015-2020*. National Geophysical Data Center, NOAA. Retrieved from www.ngdc.noaa.gov.
- MuShield Magnetic Shielding. (1996-2015). *MuShield Magnetic Shield Design Kit and Applications Guide*. Design Kit and Applications Guide, The MuShield Company, Inc., Manchester.
- MuShield Magnetic Shielding. (1996-2015). *MuShield Magnetic Shield Design Kit and Applications Guide*. Design Kit and Applications Guide, The MuShield Company, Inc., Manchester.
- Roetenberg, D. C. (2007, September). Estimating Body Segment Orientation by Applying Inertial and Magnetic Sensing Near Ferromagnetic Materials. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(3), 469-471.

- Rotenberg, D. H. (2005, September). Compensation of Magnetic Disturbances Improves Inertial and Magnetic Sensing of Human Body Segment Orientation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 13(3), 395-405.
- Sabatini, A. M. (2006, July). Quaternion-Based Extended Kalman Filter for Determining Orientation by Inertial and Magnetic Sensing. *IEEE Transactions on Biomedical Engineering*, 53(7), 1346-1356.
- Shiau, J.-K., Huang, C.-X., & Chang, M.-Y. (2012). Noise Characteristics of MEMS Gyro's Null Drift and Temperature Compensation. *Journal of Applied Science and Engineering*, 15(3), 239-246.
- Tera Analysis Ltd. (2013). *QuickField Finite Analysis System Version 6.0 User's Guide (Version 6.0)*. Svendborg, Denmark: Tera Analysis Ltd.
- XSENS. (2013). MTw User Manual. *MTw User Manual - MTw Hardware, MT Manager, Awinda Protocol*. Enschede, The Netherlands.
- Xsens. (2013). *User Guide MVN, MVN Biomech, MVN Awinda*. Xsens.