



**ABSOLUTE POSITIONING USING
THE EARTH'S MAGNETIC ANOMALY
FIELD**

DISSERTATION

Aaron J. Canciani, Capt, USAF
AFIT-ENG-DS-16-S-074

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENG-DS-16-S-074

ABSOLUTE POSTIONING USING
THE EARTH'S MAGNETIC ANOMALY FIELD

DISSERTATION

Presented to the Faculty
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

Aaron J. Canciani, B.S.E.E., M.S.E.E.
Capt, USAF

September 2016

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENG-DS-16-S-074

ABSOLUTE POSTIONING USING
THE EARTH'S MAGNETIC ANOMALY FIELD
DISSERTATION

Aaron J. Canciani, B.S.E.E., M.S.E.E.
Capt, USAF

Committee Membership:

John Raquet, PhD
Chairman

Matthew Fickus, PhD
Member

Maj Scott Pierce, PhD
Member

ADEDEJI B. BADIRU, PhD
Dean, Graduate School of Engineering and Management

Abstract

Achieving worldwide dependable alternatives to GPS is a challenging engineering problem. Current GPS alternatives often suffer from limitations such as where and when the systems can operate. Navigation using the Earth's magnetic anomaly field, which is globally available at all times, shows promise to overcome many of these limitations. We present a navigation filter which uses the Earth's magnetic anomaly field as a navigation signal to aid an inertial navigation system (INS) in an aircraft. The filter utilizes highly-accurate optically pumped cesium (OPC) magnetometers to make scalar intensity measurements of the Earth's magnetic field and compare them to a map using a marginalized particle filter approach. The accuracy of these measurements allows observability of not only the INS errors, but also long-wavelength errors in the measurements. We demonstrate navigation accuracy of 13 meters DRMS with a high quality magnetic anomaly map at low altitudes with real flight data. We identify altitude and map quality as the two largest variables which effect navigation accuracy. We further demonstrate navigation accuracies of several kilometers over a cross-country flight at 3000 meters altitude with a continental-sized magnetic anomaly map. We demonstrate that the majority of this error is caused by poor map quality. We predict navigation accuracies of a high altitude cross-country flight with an improved continental-sized map through simulation and show a range of accuracies from tens of meters to hundreds of meters, depending on altitude. We then conduct a simulation over the continental United States to predict accuracies with respect to variables like location, altitude, and velocity. Finally, we address the problem of map availability by presenting a method for a self-building world magnetic anomaly map which uses Gaussian process regression to model the error in existing large-scale mag-

netic anomaly maps. We use real data to demonstrate the benefit in map accuracy that a few flight lines can provide to a large area.

Table of Contents

	Page
Abstract	iv
List of Figures	x
List of Tables	xvii
List of Abbreviations	xix
I. Introduction	1
1.1 Technical Motivation	2
1.2 Claimed Contributions	2
1.3 Literature Review	4
Aerial Magnetic Navigation	4
Magnetic Ground Navigation	7
Magnetic Space Navigation	8
Magnetic Underwater Navigation	8
Magnetic Indoor Navigation	9
Gravity Navigation	10
Cramer-Rao Lower Bound For Map-Matching Algorithms	10
Literature Review Conclusion	11
1.4 Dissertation Organization	11
II. BACKGROUND	13
2.1 Components of the Earth's Magnetic Field	13
Core Field	13
Crustal Field	15
Induced Field	16
External Fields	16
Ionospheric Effects on the Magnetic Field	17
Magnetospheric Effects on the Magnetic Field	19
Coupling Currents	21
2.2 Earth's Magnetic Anomaly Field	21
International Geomagnetic Reference Field	22
Anomaly Definition and Assumptions	25
2.3 Magnetic Anomaly Modeling and Transforms	29
Upward Continuation	29
Magnetic Map Time Projections	34
Modeling Magnetic Anomaly Fields	38
Map Pre-Conditioning	43

	Page
2.4 Magnetic Anomaly Maps	46
Flying Over a Modern Map At Map Altitude - The	
Ideal Case	48
North American Magnetic Anomaly Map	51
World Digital Magnetic Anomaly Map	51
2.5 Types of Magnetic Measurements	54
Scalar Intensity Measurements	54
Intensity Gradient Measurements	56
Vector Measurements	59
Tensor Measurements	60
2.6 Magnetic Sensors	60
Optically Pumped / Alkali-Vapor Magnetometers	61
Fluxgate Sensors	67
2.7 Obtaining Accurate Magnetic Measurements	68
Aircraft Sources	69
Aircraft Magnetic Compensation Systems	71
Temporal Variations	74
2.8 Creating Magnetic Anomaly Maps	75
Flight Path	75
Data Processing	79
2.9 Rao-Blackwellized Particle Filtering	82
2.10 Background Conclusion	89
III. Filter Design	90
3.1 Temporal Variation Modeling	90
Types of Temporal Variations	90
Characterizing Temporal Variations as a Random	
Process	92
Temporal Variation Variables	100
3.2 Temporal Variation Observability Analysis	104
Temporal Variation Frequencies	104
Anomaly Field Frequencies	107
3.3 Measuring the Magnetic Anomaly	109
Map Quality	110
Altitude Dependent Variations	110
Corrupting Sources	111
Measurement Equation	112
3.4 Filter Design	115
3.5 Magnetic Navigation Algorithm	123
3.6 Filter Design Conclusions	126

	Page
IV. Flight Test Results	128
4.1 Ideal-Case Results	129
Flight Test Data	129
Map Stability	132
Navigation Results	133
Ideal-Case Conclusions	136
4.2 Cross Country Results	137
NAMAD Map Errors	137
Map Gradient vs. Altitude	139
Cross-Country Flight Test Data	141
Cross-Country Results	144
Cramer-Rao Lower Bound Analysis	148
Cross-Country Navigation Conclusions	150
V. Continental Simulation Results	153
5.1 Magnetometer Errors	153
5.2 Generating Magnetometer Measurements	155
5.3 Altitude	157
5.4 Core Field Effects	160
5.5 Aircraft Effects	161
5.6 Temporal Variations	163
5.7 Sensor Errors	164
5.8 Map Quality	165
5.9 Generating INS Measurements	167
Dynamics Model	169
Noise Model	171
5.10 Simulation Validation	172
5.11 Simulation Tradespace Predictions	174
5.12 Conclusions	179
VI. Self-Building Magnetic Field World Model	181
6.1 Characterizing Errors in Existing Magnetic Anomaly Maps	185
6.2 Gaussian Process Regression Models	187
6.3 Determining GPR Hyper-parameters	191
6.4 Upward and Downward Continuation	194
6.5 Results	199
6.6 Practical Implementation of the Self-Building World Model	203
6.7 Navigation Simulation Using GPR	209
Incorporating Map Covariance	209
Simulation Setup	211

	Page
Line Spacing Trade Space Study	214
Full GPR Simulation	216
6.8 Self Building World Model Conclusions	228
VII. Conclusion	229
7.1 Future Work	234
Aircraft Field Modeling	234
Measurement Types	235
Attitude Determination	236
Bibliography	237

List of Figures

Figure	Page
1	Magnetic Sources of the Earth [55] 14
2	Frequency Content vs. Amplitude of Earth's Magnetic Field [42] 17
3	Solar Quiet (Sq) Currents on Day Side of Earth [77] 18
4	Solar Quiet (Sq) Currents on Day Side of Earth [78] 19
5	Magnetospheric Currents [33]..... 20
6	Contribution of Individual Harmonics [45]..... 23
7	Spherical Harmonic Power [2] 24
8	Perpendicular Crustal and Core Field Vectors 25
9	Projection of Magnetic Anomaly Onto Core Field 26
10	Example of Projection Assumption Accuracies 27
11	Induced and Remnant Magnetization [27] 28
12	Upward Continuation Filter Spectrums 31
13	Edge Effects From Upward Continuation..... 32
14	Contribution of Ring of Area as a Function of Horizontal Distance 33
15	Magnitude of Sample Time Reduction Filter 37
16	Phase of Sample Time Reduction Filter..... 37
17	Synthetic Map Created at 1km Altitude and Subset Map 40
18	True and Calculated Synthetic Map Upward Continued 500m from 1km 41
19	True and Calculated Synthetic Map Upward Continued 50m from 1km 42
20	Time Reduction Filtering on 5 Year Old Map..... 44

Figure		Page
21	Time Reduction Filtering on 15 Year Old Map.....	45
22	Removing a 2D Trend	46
23	Example of a Smoothed Periodic Map Tile [20]	47
24	Freely Available Maps from the USGS (Post-1990 Era Maps Outlined in Black)	50
25	North American Magnetic Anomaly Map – 1km Grid Spacing at 305 meter Altitude [3]	52
26	North American Magnetic Anomaly Map Contributing Surveys and Line Spacings [3]	53
27	World Digital Magnetic Anomaly Map [43]	55
28	Sander Geophysics Magnetic Gradient System [22]	57
29	All 9 Components of a Magnetic Tensor (6 unique) [69]	61
30	Geometrics 823A Magnetometer [21]	62
31	Specification Sheet for Geometrics 823A [21]	63
32	Sensor Dead-Zones for Geometrics 823A [20]	65
33	Specifications from Barrington SpaceMag Catalog [4]	68
34	Aircraft Stinger Holding Magntometer [22]	70
35	Before and After Magnetic Compensation [52]	72
36	Temporal Variations Recorded at Magnetic Base Stations Separated by 10's of Kilometers [52]	74
37	Basic Aeromagnetic Survey Lines	76
38	Example Diurnal Variations [58]	91
39	Example Micro-Pulsations [58]	91
40	Example Magnetic Storms [58]	92
41	Daily Recorded Temporal Variations at Boulder CO Observatory in 2010.....	93

Figure		Page
42	Montly Average Temporal Variations Showing Secular Variations	93
43	Zero-Mean Daily Variations to Emphasize Hourly Variations	94
44	5 Quiet Days From Single Month to Determine Hourly Average	95
45	Year of Data with Hourly Averages Determined for Each Month Removed	96
46	Removing the Secular Variation From 1 Month of Data in Order to Perform Valid Auto-Correlation	98
47	Estimating the FOGM time constant	99
48	Verifying the FOGM Model	100
49	Time-Aligned Temporal Variations at Mid and Equatorial Latitudes	101
50	Time-Aligned Temporal Variations at Polar Latitudes	102
51	Temporal Variations Under Various Magnetic Storm Conditions	103
52	Amplitude vs. Time-Frequeuncy of the Temporal Variations [42].....	105
53	Temporal Variation Spatial Frequencies [13]	106
54	Temporal Variations at Magnetic Base Stations 10's of km Apart [52]	106
55	Comparison of Temporal Variation and Crustal Field Frequencies at Aircraft Velocities and Altitudes [42]	109
56	Geo-Survey Aircraft Taking Off for Test Flight - Magnetometers Located on Tail Boom	129
57	2012 Magnetic Anomaly Map Over Louisa VA and 2015 Flight Profile Over Mapped Area	130

Figure		Page
58	Difference Between the Expected Measurements From Interpolation Function (Using True GPS Positions) and the Raw Measurements (made zero-mean)	132
59	Difference Between the 2012 Magnetic Anomaly Map and the 2015 Magnetic Anomaly Map	133
60	North and East Error Over 1-Hour Segment of Flight Profile	135
61	Filter Estimation of Temporal Variations	137
62	North American Magnetic Anomaly Database [3]	138
63	Cross-country flight path	141
64	Actual Recorded Measurement and Expected Measurement From NAMAD Map	142
65	Measurement Errors and Map Gradient	143
66	Navigation Grade INS Drifts and Tactical Grade INS Drifts Over Full Flight	145
67	Filter Results Using Navigation Grade INS, Real Magnetometer Measurements, and NAMAD Map	146
68	Filter Results Using Simulated Tactical Grade INS, Real Magnetometer Measurements, and NAMAD Map	146
69	Navigation Accuracy Using Hypothetical Improved-Quality Magnetic Anomaly Map	151
70	Comparison of CRLB and Filter Predicted Standard Deviations	152
71	CRLB Predicted Standard Deviation at 10 altitudes	152
72	Magnetic Anomaly Gradient Over United States at 1km (Contrast Stretched)	152
73	Upward Continuation Filter Spectrums for 100x100 km map	159
74	Upward Continuation of North American Magnetic Anomaly Database	161

Figure		Page
75	Compensation of Aircraft Effects on Geo-Survey Aircraft [23]	162
76	Non-removable aircraft effects for simulation measurement corruption	163
77	Temporal Variations From 10 Different Flight Segments, Taken From Real Temporal Variations	164
78	Line Spacing of North American Magnetic Anomaly Database	167
79	Comparison of Simulated and Actual Flight Test Error Standard Deviations over Virginia	174
80	Comparison of Simulated and Actual Flight Test Error Standard Deviations over Texas	175
81	Simulated Trajectories Over the United States	177
82	Expected DRMS Accuracy Over the United States at 200 m/s at 300 m Altitude	178
83	Navigation Error Cumulative Distributions	180
84	Expected DRMS Accuracy Over the United States at 200 m/s at 3000 m Altitude	180
85	Texas Regional Magnetic Anomaly Map at 900 meters altitude	192
86	Virginia Regional Magnetic Anomaly Map at 300 meters altitude	193
87	NAMAD Region of Virginia Survey at 300 meters altitude	193
88	Calculated Covariance Function for the VA Survey Area	194
89	Calculated Covariance Function for the TX Survey Area	195
90	Upward Continuation Transform at Several Altitudes	197
91	Downward Continuation Transform at Several Altitudes	198
92	Calculation of α Parameter for Downward Continuation	199

Figure		Page
93	Comparison of Stabilized and Unstabilized Downward Continuation	200
94	Original Absolute Value of Error in Texas NAMAD Map and Error After GPR From 5 Flight Lines	201
95	Original Absolute Value of Error in Texas NAMAD Map and Error After GPR From 10 Flight Lines	202
96	Original Absolute Value of Error in Virginia NAMAD Map and Error After GPR From 5 Flight Lines	203
97	Original Absolute Value of Error in Virginia NAMAD Map and Error After GPR From 10 Flight Lines	204
98	Standard Deviation of GPR Map Prediction	204
99	Validity of GPR Map Prediction Over Rectangular Flight Path	205
100	Multiple Map Layers Illustration	208
101	One-Dimensional Particle Weighting Using Map-Based Covariances	211
102	Self Building World Model Navigation Simulation Setup	212
103	Sparse Flight Line Spacing: 100 km Line Spacing 10 km Line Width	213
104	Sparse Flight Line Spacing: 200 km Line Spacing 10 km Line Width	213
105	Simple Binary Covariance Method (No GPR)	214
106	Flight Line Masked Map Errors	215
107	Horizontal Errors Using Binary Method	216
108	Comparison of Long and Short Wavelength Map Errors	218
109	Grid of Regional GPR Models	219
110	GPR Calculated Map Covariance	219

Figure		Page
111	GPR Corrected Map Values and STD Along Flight Line (Short Wavelength Errors)	220
112	GPR Corrected Map Values and STD Along Flight Line (Long Wavelength Error)	220
113	East Navigation Error Using GPR Method	224
114	North Navigation Error Using GPR Method	225
115	Measurement Residuals Using GPR-Derived Particle Covariances	226

List of Tables

Table	Page
1	Height/Line Spacing Ratio vs. Aliased Power [53] 108
2	Max Crustal Field Temporal Frequencies When Sampled From an Aircraft 108
3	Measurment Equation Validation 131
4	“Best Case” Flight Test Navigation Accuracy Results..... 134
5	Magnetic Anomaly Gradients of NAMAD Map at Various Altitudes and Predicted Navigation Accuracy 141
6	Navigation Accuracy Results with Tactical-Grade INS 148
7	Navigation Accuracy Results with Hypothetical Improved Map 148
8	Flight Test 1 Simulation Parameters 173
9	California Navigation Accuracy Tradespace Statistics 176
10	US Navigation Accuracy Tradespace Statistics 177
11	Information Comparison for True High Resolution Maps and Errors Between High Resolution Maps and NAMAD..... 186
12	Gaussian Hyper-Parameters Calculated From Observations Along 5 Flight Lines 195
13	Results of Gaussian Process Regression Modeling on 5 and 10 Flight Lines Over Texas and Virginia Regions 201
14	Comparison of Bootstrapping and Standard GPR Approaches 206
15	Line Spacing and Line Width Navigation Accuracy Tradespace 216
16	Comparison of Corrective Ability of GPR Method for Short and Long Wavelength Map Errors 221
17	Navigation Results for 100 km Line Spacing and 10 km Line Width Map Corrections (Short Wavelength Errors) 222

Table		Page
18	Navigation Results for 100 km Line Spacing and 10 km Line Width Map Corrections (Long Wavelength Errors)	223
19	Navigation Results for 200 km Line Spacing and 10 km Line Width Map Corrections (Long Wavelength Errors)	227

List of Abbreviations

Abbreviation	Page
GPS	Global Positioning System 1
INS	Inertial Navigation System 3
USGS	United States Geological Survey 5
EKF	Extended Kalman Filter 6
MAD	mean absolute difference 6
ICP	Iterative Closest Point 6
UPF	unscented particle filter 6
SITAN	Sandia inertial terrain-aided navigation 6
SLAM	Simultaneous Location and Mapping 9
IGRF	International Geomagnetic Reference Field 22
IAGA	International Association of Geomagnetism and Aeronomy 22
RTP	Reduction to the Pole 35
RTE	Reduction to the Equator 35
ESDI	Equivalent Source Dipole Inversion 38
NAMAD	North American Magnetic Anomaly Database 51
WDMAM	World Digital Magnetic Anomaly Map 51
FOM	Figure of Merit 73
MPF	marginalized particle filter 82
FOGM	First Order Gauss Markov 90
FOGM	First-Order Gauss-Markov 95
DST	Disturbance Storm Time 103
CDF	Cumulative Distribution Function 178

Abbreviation		Page
GPR	Gaussian process regression	183
SLAM	Simultaneous Localization and Mapping	183
AGL	above ground level	186

ABSOLUTE POSTIONING USING THE EARTH'S MAGNETIC ANOMALY FIELD

I. Introduction

There is currently a desire to develop a backup aerial navigation system when the Global Positioning System (GPS) is unavailable. These backup navigation systems are often referred to as alternative navigation systems. In the context of aerial navigation, these alternative navigation systems do not typically attempt to replace GPS, but rather augment it. If GPS is being disrupted an alternative navigation system can provide positioning information. There are currently a wide variety of alternative navigation systems, ranging from older radio-based navigation techniques [62] to newer computer-vision based approaches [71]. Additional navigation techniques include star-trackers [35], terrain height matching [28], and gravity gradiometry [54]. In a testament to the usefulness of GPS, none of these technologies are adequate as an overall replacements for GPS. Each of these alternative technologies tends to work best in a certain environment or under specific conditions. For example, star trackers do not work well during the day or in cloudy conditions. Vision-aided navigation requires unique features, making oceans, deserts, and forests problematic. Radio-based communications can be limited by altitude. While GPS is exceptionally accurate, its world-wide availability at all times may be the most difficult attribute to replace. Developing alternative navigation systems which better match the world-wide availability of GPS is an important step in addressing the vulnerabilities of GPS.

1.1 Technical Motivation

Aerial navigation using magnetic anomaly fields can overcome many of the drawbacks of current alternative navigation systems. The Earth’s magnetic anomaly field is available world-wide, including over oceans, forests, and deserts. This availability is an attribute that vision-based and terrain height alternative navigation systems do not have. Furthermore, the magnetic anomaly signal is available at all times of the day and under all weather conditions. The general availability of the magnetic anomaly field is better than most of the previously discussed navigation systems. The magnetic anomaly field is also very difficult to jam. The power from the magnetic field of a dipole decays with distance at the rate of d^{-5} while radio signals decay with distance at the rate of d^{-2} . This indicates that it would take much more energy to disrupt a magnetic dipole source field at distance d than it would to disrupt a radio-based signal. In addition to being difficult to disrupt, the magnetic anomaly field is measured with a completely passive instrument, emitting no signal of its own. In contrast, terrain-height navigation must emit radar or laser signals to measure terrain height, making it an active system. It is clear that magnetic anomaly navigation is able to address many of the *availability* obstacles of a robust alternative navigation system. This research attempts to determine if the *accuracy* of magnetic anomaly navigation is sufficient for practical use onboard aircraft. An alternative navigation system which was both widely-available and sufficiently accurate for a given situation would overcome many of the drawbacks of current alternative navigation systems.

1.2 Claimed Contributions

The main contribution of this research is a proven navigation system which uses the Earth’s magnetic anomaly field to navigate. The navigation system functions by matching measurements from a scalar magnetometer to a magnetic anomaly map to

correct the drift in a navigation grade Inertial Navigation System (INS). This navigation system was tested on real flight data with real magnetic anomaly maps. Low altitude magnetic anomaly navigation with a high quality map provided positioning information to tens of meters accuracy. High altitude navigation with high quality maps is shown in simulation to provide positioning information on the order of 10's to 100's meters of accuracy, depending on spatial frequency content of the magnetic field as well as altitude. Achieving this level of accuracy is a two order of magnitude improvement over current published navigation systems using the magnetic anomaly field for airborne navigation [79].

A detailed analysis of the measurement equation for magnetic anomaly navigation was performed in order to remove the various corrupting sources in a magnetometer measurement. The current literature on magnetic anomaly navigation does not provide a detailed analysis of these corrupting components. These corrupting sources, in the context of magnetic anomaly navigation, include the aircraft field, temporal variations, and the Earth's core field. By fully developing methods to remove each of these components, a magnetic anomaly navigation system can reach its highest potential accuracy.

Finally, a method for creating a self-building world map is presented. This method allows a small number of aircraft flights making measurements of the magnetic anomaly field to make large regional corrections to the poor-quality magnetic anomaly maps which are widely available. A self-building world model addresses the important concern of map availability.

1.3 Literature Review

Absolute navigation using geophysical fields is an active research area. The prospect of passive navigation systems using naturally occurring environmental signals is understandably appealing. There are two major categories to consider when discussing navigation using geophysical fields: the platform environment, and the specific geophysical fields or signals being used for navigation. There are five major platform environments seen in the literature for geophysical field navigation: ground [56], air [79], sea [67], space based platforms [59], and indoor platforms [61]. Each of these platforms presents unique challenges to geophysical navigation. There are three main types of geophysical fields or signals commonly explored for navigation: magnetic fields [79], gravity fields [17], and terrain elevation [28]. Perhaps the most extensively explored topic in the overall field of absolute navigation using geophysical fields is terrain-elevation navigation for aircraft. Terrain-elevation navigation is a mature technology, with fielded military systems being deployed decades ago [28]. The use of such systems is restricted to aerial use over land. Current research into geophysical field navigation shows promise to expand the availability and accuracy of geophysical field navigation past the proven success of terrain-elevation navigation systems.

Aerial Magnetic Navigation.

Magnetic navigation has been implemented on a large variety of platforms. This research is focused on aerial navigation, but can benefit from research in other environments such as ground, indoor, underwater, and space. Research on magnetic navigation using these environments is especially important considering there is limited experimental results for aerial magnetic navigation, while a larger body of research exists on indoor navigation.

Absolute positioning for aircraft using magnetic fields has been experimentally demonstrated. A fair amount of research has been conducted on the theoretical aspects of aerial magnetic navigation systems, including computer simulations, which will be discussed shortly. Limited experimental results are available in the literature. The most applicable experimental results found were in a paper by Wilson and Kline-Schoder [79]. Wilson and Kline-Schoder designed a navigation filter which compared measurements from a 3-axis magnetometer to a magnetic anomaly map obtained from the United States Geological Survey (USGS) over Vermont. The navigation filter used the magnetic measurements to aid a simple airspeed dead-reckoning system. The magnetic measurements were matched to the map through a batch process which matched a sequence of measurements to a location on the map as well as an estimate of wind speed—the primary source of drift in the dead-reckoning system. Wilson and Kline-Schoder demonstrated accuracies of around 1–2 kilometers with post-processed data from a flight test.

Goldenberg discusses absolute positioning on aircraft using magnetic fields in [24]. This paper presents useful discussion on many aspects of a magnetic navigation system. Goldenberg identifies mapping as one of the primary limitations of magnetic anomaly navigation. There currently does not exist high accuracy magnetic maps over the entire world, but rather a patchwork of maps. He suggests the World Magnetic Model may address this concern. Goldenberg addresses the mathematics of upward continuation for map data. Because magnetic maps exist in three dimensions, maps need to be continued upward when flying at higher altitudes. Goldenberg states that above 30 kilometer scalar measurements are more useful for magnetic navigation, and below 30 kilometer gradient measurements would be most useful. Finally, he gives a background of the corrupting effects of an aircraft’s magnetic field and addresses

placement and calibration of magnetometers.

There are many papers on geomagnetic aerial navigation in the literature which approach the problem from a theoretical standpoint, with results obtained via simulation. The majority of these papers present various algorithms for matching a sequence of measurements to a magnetic map. The algorithms can be placed into two major groups: batch algorithms and sequential algorithms. The focus of listing these algorithms is to demonstrate the wide availability of possible approaches to the map-matching problem, not to choose a best approach.

Caifa et al. compare a sequential Extended Kalman Filter (EKF) approach with a batch mean absolute difference (MAD) approach [12]. They conclude the sequential filter approach is more robust in areas of limited magnetic variability. Liu et al. present a batch Iterative Closest Point (ICP) algorithm as a method to choose an update location for a Kalman Filter, with the complete algorithm being called a Nearest Point Kalman Filter (NPKF) [38]. They conclude that the algorithm successfully constrains the drift of an INS. Ming-ming et al. presented an unscented particle filter (UPF) approach which used the position solution as an update to a Kalman Filter to obtain the remaining kinematic states [47]. They conclude that the algorithm constrains the drift of an INS. Xu et al. present the results of a batch map-matching differential evolution algorithm [80]. The algorithm uses the differential evolution technique to quickly search through a map grid for a likely location based on a sequence of measurements. They conclude that the algorithm is a faster way to search through a large grid. Dai and Kang present the results of using the Sandia inertial terrain-aided navigation (SITAN) algorithm with magnetic field measurements [16]. The SITAN algorithm is the experimentally proven terrain following

system developed by Sandia National Laboratories. They conclude that the algorithm, developed for terrain maps, also works with magnetic data to constrain the drift of an INS. Wang uses principal component analysis to implement a batch map-matching algorithm [76]. He describes sequences of measurements by several factors such as standard deviation, roughness, gradient deviation, information entropy, and Fisher information content to describe a sequence of measurements. He then uses this unique label to find possible matching map locations to predict position. He concludes that the technique improves matching probability over more simple batch map-matching techniques. Liu explores the use of a particle filter for geomagnetic aerial navigation [37]. He concludes that the particle filter works better than the EKF for geomagnetic navigation. Feng analyzes the use of the Iterative Closest Contour Point algorithm and concludes that it has better accuracy than the traditional contour matching algorithm [18].

Magnetic Ground Navigation.

Shockley successfully implemented ground vehicle navigation using only magnetic field measurements [56]. He experimentally demonstrated the ability to get near meter-level position fixes on an intermittent basis. He found that areas like highways had less magnetic variation and led to degraded filter performance. Shockley created his own maps and demonstrated that the measurements had enough repeatability to implement a correlation type filter with a sequence of magnetic measurements. Kauffman and Raquet followed Shockley's work by integrating an INS with the magnetic measurement filter [32]. Kauffman showed improved performance especially over areas of low magnetic features when the INS allowed the filter to coast to an area of better magnetic feature availability.

Magnetic Space Navigation.

Shorshi and Bar-Itzhack successfully demonstrated satellite navigation using magnetic measurements [59]. Shorshi and Bar-Itzhack used real satellite magnetometer measurements and related them to the International Geomagnetic Reference Field (IGRF) using an Extended Kalman Filter. The algorithm achieved performance on the order of 1–10 kilometers. Shorshi and Bar-Itzhack also compared the performance of using just the magnetic intensity sensor with the magnetic vector sensor. Using the magnetic vector sensor required knowledge of the satellite attitude and improved performance due to the filter bringing in three separate measurements instead of one. Psiaki et al. designed a similar system which also matched measurements from real satellite data to the IGRF [50]. They found performance to be on the order of 4–8 kilometers, which agrees with Shorshi and Bar-Itzhack’s work.

Magnetic Underwater Navigation.

Tyren proposed one of the first magnetic navigation systems, specifically for use on submarines [67]. He presented actual recorded data from two separate magnetometers separated by a fixed distance to calculate ground speed. While his proposed method was a velocity aiding system and not an absolute positioning system, his work is still widely cited as one of the first papers to discuss using magnetic anomalies for navigation. Tyren also identified the problem of separating vehicle and Earth magnetic fields and discussed methods for separating the two signals. Jie discusses underwater geomagnetic navigation using a Kalman Filter [30]. Jie’s paper discusses removing the diurnal variations that are inherent in magnetometer measurements. He concludes that the navigation solution is more accurate when removing the mean diurnal variation. Zhang discusses using the previously mentioned ICP algorithm for underwater navigation with magnetometer measurements [81]. He concludes that the

algorithm is able to restrict the drift of an INS. May and Meisinger discuss naval testing of a magnetic anomaly navigation system aboard naval platforms [46]. May and Meisinger describe repeatability tests in which the magnetic measurements over repeat ground tracks have standard deviations of just 1.8 nano-Teslas. They describe this as impressive with respect to the fact that magnetic field gradients existed in the testing area as high as 350 nano-Teslas/mile. They also note that the standard deviation of the differences between measurements and a magnetic anomaly map vary between 5 and 30 nano-Teslas. These experimental results help motivate the potential accuracy of magnetic anomaly navigation, indicating that a large signal to noise ratio potentially exists in a magnetic anomaly field measurement.

Magnetic Indoor Navigation.

Storms et al. experimentally demonstrated magnetic positioning indoors [61]. They showed obtainable accuracies at the decimeter level when navigating through previously mapped areas as well as successful navigation in a leader/follower scenario without a map. Storms et al. used a Kalman filter based on principles from a terrain height navigation system developed in [48]. They identify magnetometer calibration as one of the biggest difficulties in high accuracy navigation. They also discuss the stability of the magnetic map. The stability of magnetic maps is also of concern in aerial navigation. A static map assumption is seen almost universally across the literature, when this in fact is not a correct assumption. Haverinen and Kemppainen used a particle filter to implement one dimensional navigation inside building corridors [26]. They specifically demonstrated the stability of the magnetic field inside buildings over time (40 days) and conclude the fields are relatively stable. Havernin and Kemppainen also achieved performance on the order of 1 meter. Vallivaara et al. developed a Simultaneous Location and Mapping (SLAM) robot which navigated

and mapped the three dimensional magnetic field inside a building to an accuracy of less than a meter [70]. Judd and Vu demonstrate using the magnetic vector field as a fingerprint which can aid a dead-reckoning system when revisiting a location [31]. Gozick et al. used the magnetometers available on mobile phones to identify magnetic landmarks and create magnetic maps used for navigation [25].

Gravity Navigation.

Navigation using the Earth’s gravity field has a fundamental difference from magnetic and terrain navigation in that the gravity vector is not directly measurable with current instruments. Accelerometers can measure specific force, but cannot separate the effects of aircraft movement from gravity. This necessitates the need for gravity gradient measurements to implement a gravity map-matching navigation system. Richeson describes key aspects of a gravity gradiometry map-matching navigation system [54]. Richeson concludes that current gravity gradiometry systems of the required accuracy either do not exist or are too large for most applications. DeGregoria also found that current technology is not at the appropriate level to implement gravity gradiometry navigation [17]. Both authors predict promising navigation accuracies under the assumption of much more accurate future instruments.

Cramer-Rao Lower Bound For Map-Matching Algorithms.

The Cramer-Rao lower bound (CRLB) is a lower limit on the covariance of an unbiased estimator. The CRLB allows the performance of a specific filter to be compared to the best possible performance with the given information. Niclas Bergman derives the CRLB for a terrain aided navigation system which aids an INS with terrain height measurements matched to a map in [8]. The derivation of this CRLB is not trivial, although the final result states that the CRLB is equivalent to the

error covariance in the discrete time Extended Kalman Filter where the non-linear measurement function h has been replaced by its gradient evaluated at the true state value at time k [8]. This is an extremely useful result, since the EKF equations are well understood and easily implemented. The extension of this derivation to magnetic anomaly maps is not difficult. Bergman further validates this result in a separate paper in which he shows a Terrain-Aided navigation system achieving the CRLB after convergence [9].

Literature Review Conclusion.

This section has provided a literature review for magnetic anomaly navigation. It is clear that minimal research has been done in the specific area of aerial magnetic field navigation. The only experimental results in the literature achieved accuracies of approximately 2 kilometers. There is a large deal of theoretical papers on the topic, however. These papers almost all concern themselves with the method of matching measurements to a spatial map, and include limited discussion of the unique problems encountered in aerial magnetic field navigation, such as navigating in a three dimensional field. The navigation techniques for indoor, ground, and space magnetic navigation provide useful insights to aerial magnetic navigation, but are a very different problem than aerial magnetic anomaly navigation. Finally, gravity and terrain navigation provide additional analogous approaches to aerial magnetic field navigation, but have major differences from the type of navigation system this research plans to implement.

1.4 Dissertation Organization

The remainder of this dissertation is organized as follows. Chapter II presents a thorough geo-physical background on magnetic anomaly fields and Rao-Blackwellized

(marginalized) particle filtering. Chapter III details the design of the navigation filter used for magnetic anomaly navigation. Chapter IV provides the results and analysis from our flight tests. Chapter V presents a continental simulation over the United States to explore magnetic anomaly navigation accuracy with respect to location, altitude, and velocity. Chapter VI presents a method for a self-building world magnetic anomaly map. Chapter VII suggests future work and draws final conclusions on the research which was conducted.

II. BACKGROUND

The purpose of this background section is to provide readers with a limited geophysics background enough information to apply navigation techniques to the Earth's magnetic field. The chapter attempts to provide no more information than is pertinent to the design of a navigation system. Topics covered include the major components of a magnetic field measurement as well as the definition of magnetic anomaly fields. A background on magnetic anomaly map transforms is then provided as well as ways to model magnetic anomaly fields. The current availability of magnetic anomaly maps is presented and a standard method to calibrate out aircraft fields is discussed. Finally, a basic tutorial on implementing the Rao-Blackwellized particle filter is provided.

2.1 Components of the Earth's Magnetic Field

The Earth's magnetic field is the superposition of a large variety of magnetic sources. The origin and interaction of each of these sources is complex. At a high level, the Earth's magnetic field can be thought of as consisting of internal and external sources. The internal sources consist of the core, crustal, and induced fields. The external sources consist of the ionosphere, the magnetosphere, and coupling currents [29]. Fig. 1 shows these components of the Earth's magnetic field.

Core Field.

The first component of the internal magnetic field is the core, or main Earth field. The main Earth field is what makes a compass point north. As a first order approximation, it is a dipole. It accounts for the vast majority of the measured magnetic field near the surface of the Earth. It is caused by the motion of conductive fluids deep within the Earth. Heat from the Earth's inner core causes fluid movement in the




Figure 1. Magnetic Sources of the Earth [55]

liquid outer core, and the Earth's rotation causes this fluid to rotate. The rotating conductive materials generate electrical fields, which in turn generate magnetic fields. The field caused by these deep sources accounts for around 95-99 percent of the total measured field at any point in the vicinity of the Earth's surface [29]. The magnitude of the main Earth field varies widely from around 30–70 micro-Teslas. The main Earth field undergoes secular variations. These are variations in the field over time. The changes to the main field are not insignificant—the field is remodeled every five years to account for observed changes [29]. The wavelengths associated with the core field are long. Spherical harmonic models of the core field generally include the first 13 wave numbers of the magnetic field [29]. This indicates the shortest spatial wavelength that exists as a result of the core field is around 4000 kilometers. Intuitively, such a low frequency signal would not allow sub kilometer level navigation.

Crustal Field.

The second component of the internal sources is the crustal, or lithosphere field. The crustal sources are caused by the permanent or induced magnetization of rocks in the Earth’s crust. Due to the lower temperatures of the Earth’s crust, many magnetic materials are at a temperature below the Curie point—the point at which materials change from having induced magnetization to permanent magnetization [55]. There are two main causes for the mineral’s magnetization. The first is remnant magnetization. This is caused by a past induced magnetic field changing to a permanent magnetic field when the mineral cooled below the Curie temperature [55]. The second type of magnetization is induced magnetization. This occurs when the present day magnetic field induces a magnetic field in a magnetically susceptible mineral. The magnetic field generated by these sources is small compared to the deep sources. The crustal field accounts for around 1–5 percent of the total measured magnetic field in the vicinity of the Earth’s surface [13]. This amounts to several hundred nano-Teslas. An important aspect of the crustal field is that it changes so slowly that it may be considered static. It also includes high spatial-frequency information. This makes it an ideal candidate for map based navigation. An approximation for the highest expected crustal field wavelengths is simply the measurement altitude above the surface of the Earth. Flying at 1 kilometers altitude, spatial magnetic field wavelengths of 1 kilometers are expected to exist [42]. As altitude increases, individual magnetic features tend to “blur” together. Increasing altitude thus acts like a low pass filter on the crustal magnetic field. Intuitively, this indicates altitude will be a major factor in navigation performance.

Induced Field.

External sources will induce currents in the conductive mantle of the Earth. These currents in turn will create their own magnetic fields. Although these currents exist deep within the Earth, it is important to differentiate them from the main Earth field. Separating these components can often be difficult. These fields are time varying due to the time varying nature of the external sources which induce them. Models of a region's ground conductivity can help in predicting the effect of induced magnetic fields [13].

External Fields.

The combined effect on the measured magnetic field of all external sources is often referred to as the temporal variations. Temporal variations are caused by many distinct sources with varying characteristics. The amplitude and frequency content of the variations can vary dramatically. As a general rule, higher frequency variations tend to have lower amplitudes as shown in Fig. 2. From a navigation standpoint, this is beneficial, as the larger temporal variations will vary more slowly with time. Slowly changing variations tend to look like constant biases over short duration flights. Constant bias type errors are often easier to estimate and correct than time varying errors. On the opposite end of the spectrum, very high frequency variations are also not as problematic for a navigation system. High frequency errors will tend to average out, as they appear similar to white noise. When flying over a spatial map, the spatial-frequency information of the map is transformed to a temporal frequency that depends on flight velocity. Intuitively, temporal variations that occur at similar frequencies to those that exist in the map data will be the most difficult errors to estimate and remove, and will likely be un-observable in a navigation system.




Figure 2. Frequency Content vs. Amplitude of Earth's Magnetic Field [42]

Ionospheric Effects on the Magnetic Field.

The ionosphere is a major contributor to the temporal variations. The ionosphere is ionized by solar radiation, which creates an electrically conducting plasma where electric currents can flow. The flow of these currents is primarily driven by the heating of the atmosphere by the sun. The day-night cycle creates differential solar heating which causes atmospheric tidal winds. The gravitational attraction of the moon can also create these tidal winds. As the electrically conducting plasma moves relative to the main Earth magnetic field, electric currents are created. These currents in turn create magnetic fields. Fig. 3 shows these currents flowing on the day side of the Earth. The magnetic field induced as a result of the differential solar heating is called the Sq, or solar quiet variations. Sometimes the currents which create these variations are referred to as solar currents. The currents caused by the gravitational attraction of the moon are called lunar currents. The magnetic field caused by the solar and lunar currents is relatively smooth and periodic, because these systems exist




Figure 3. Solar Quiet (Sq) Currents on Day Side of Earth [77]

directly between the sun and Earth, and the Earth is rotating. The period of the solar currents is 24 hours. The strength of the magnetic field induced by the solar currents depends on latitude, season, and time of day. At the middle latitudes, the solar currents can cause variations of 20 to 40 nano-Teslas. Near the magnetic equator these variations can be as high as 100- 200 nano-Teslas as a result of the equatorial electrojet (described below). The lunar currents have a period of 12 hours. They cause magnetic field variations of around 1 nano-Tesla - much smaller than the solar currents [5].

The equatorial electrojet (EEJ) is an eastward current on the day-time side of the Earth flowing along the magnetic dip equator. The previously described solar currents cause an east-west directed electrostatic field which interacts with the magnetic dip equator. The magnetic dip equator is the point at which the main Earth geomagnetic field is horizontal. The resultant current is a narrow band (about 6 degrees of latitude) in the equatorial regions. The EEJ can induce magnetic fields 5–10 times stronger than the mid-latitude solar currents. Similar to the solar and lunar currents, the EEJ manifests itself as a 24 hour period with respect to stationary measurements




Figure 4. Solar Quiet (Sq) Currents on Day Side of Earth [78]

taken near the magnetic equator. Like the solar currents, the EEJ is driven by the time of day, season, and solar activity. This indicates that magnetic measurements near the equator will have stronger Sq variations than at mid latitudes. Fig. 4 shows the relative strengths of these solar quiet fields. The Sq arrows are pointing to approximately noon local time [5].

The final major ionospheric current systems are the auroral currents. These currents are caused by polar electrojets and other transient currents. They are usually located above 67 degrees latitude but may occur lower. The auroral currents have a periodicity which is related to the 27 day rotation period of the sun. They are strongly influenced by the surface activity of the sun. Maximum activity occurs during the Spring and Fall when the Earth is at the equinoxes. The magnetic field induced by the auroral currents can be much stronger than the other ionospheric current systems. Typical variations are on the order of 1500 nano-Teslas. This indicates that navigation near the Earth's poles may potentially be difficult [5].

Magnetospheric Effects on the Magnetic Field.

The Earth's magnetosphere is the area of space surrounding the Earth in which charged particles from the sun interact with the Earth's magnetic field. The cur-




Figure 5. Magnetospheric Currents [33]

rents of the magnetosphere are driven primarily by the solar wind coming from the sun. The solar wind is a stream of plasma ejected from the sun at high velocities. This stream of particles is composed of electrons and protons. When these charged particles reach Earth they begin to interact with the Earth's magnetic field. These interactions can cause currents which create their own magnetic fields. The magnetosphere field is primarily driven by ring currents and the currents of the magnetopause and magnetotail. Fig. 5 shows a depiction of the major current systems of the magnetosphere [55].

Ring currents are the most prominent current systems in the inner magnetosphere [44]. Ring currents are caused by currents which flow along the magnetopause. These currents cancel the Earth's field outside the magnetopause boundary and stretch the field outward in the characteristic tail shape [55]. The resulting cavity contains sheet currents which are aligned with the equatorial plane. These currents interact with radiation belts near the Earth and create ring currents which partially circle the Earth [55]. The ring currents achieve full loop closure by coupling with the ionosphere. These coupling currents are described shortly. The combined effect of

the ring currents and the magnetopause and magnetotail currents is approximately 20–30 nano-Teslas during quiet times but can be in the hundreds of nano-Teslas during geomagnetic storms [55].

Coupling Currents.

The previously described magnetosphere currents are roughly solenoidal [55]. This means they tend to flow in closed paths. Because of the variable conductive structure of the near Earth region, this closure is sometimes achieved through the coupling of the magnetosphere sources and the ionosphere sources [55]. The coupling of sources to close a current loop forms another distinct magnetic source. When this coupling takes place near the poles of the Earth it is called a Field Aligned Current (FAC). This name comes from the fact that the Earth magnetic field lines are vertical near the Earth’s poles, so the coupling currents are flowing along the main Earth magnetic field lines. During quiet periods at high latitudes, field aligned currents create magnetic fields on the order of 30–100 nano-Teslas [55]. There are also coupling currents which flow between the magnetosphere and the Sq currents. These induce a magnetic field with a magnitude of around 10 nano-Teslas or less. Finally, the equatorial electrojet can also couple with the magnetosphere with currents causing a magnetic field on the order of 15-40 nano-Teslas [55]. The FAC are another contributing factor to the highly varying fields near the Earth’s poles and another reason why magnetic navigation near the poles would be difficult.

2.2 Earth’s Magnetic Anomaly Field

Due to their wide availability, we are attempting to use standard magnetic anomaly maps for navigation. Understanding how these maps are made and what they represent is important for developing a navigation system which utilizes these maps. A

magnetic anomaly, by definition, is the vector deviation from a reference field. The reference field in this context is the main Earth field, and the anomalies of interest are crustal field sources. Geophysicists and industry have been making magnetic anomaly maps for decades to study the Earth’s subsurface. These maps give valuable insight into the locations and types of minerals buried underground and are used commonly in industry to discover resources such as oil and diamonds. Understanding what exactly a magnetic anomaly map represents involves understanding a few subtleties that will be presented in this section. What is typically referred to as a “magnetic anomaly map” is *not* a map of the magnetic vector deviation from a reference field. It is instead an intensity measurement which is an *approximation* to the projection of the vector magnetic anomaly in the direction of a reference field. While this may sound somewhat arbitrary at first, it stems from practical issues which make this type of map far easier to generate than a true magnetic anomaly map. A magnetic anomaly map is created by differencing the measured magnetic intensity over a survey area from a particular reference field. Without a specific set of assumptions, this scalar subtraction of two vector field intensities would hold little physical meaning. These assumptions will be detailed in the following sections and the physical significance of the magnetic anomaly maps will be described.

International Geomagnetic Reference Field.

A magnetic anomaly map is created by differencing the measured magnetic intensity from a reference field. The most commonly used reference field is the International Geomagnetic Reference Field (IGRF). This model is put out every 5 years by the International Association of Geomagnetism and Aeronomy (IAGA) [63]. The IGRF is created to provide a standard world-wide core field model. This allows for a common reference field to be used in the processing of surface and aeromagnetic anomaly field

Single Harmonics




Figure 6. Contribution of Individual Harmonics [45]

data [34]. The IGRF primarily captures the core field and includes time derivative terms to account for secular variations, which are changes in the core field over time.

The IGRF is a spherical harmonic model. A spherical harmonic model attempts to fit a periodic model onto a sphere with a set of coefficients. A spherical harmonic model can contain any number of harmonics of degree n . These harmonics are also often called wave-numbers. Each wave-number corresponds to a given spatial wavelength in the Earth's magnetic field. Because the Earth's magnetic field can be roughly approximated by a tilted dipole, it is no surprise that the majority of the power in the Earth's magnetic field is at low wave-numbers. Fig. 6 shows the contribution from the first few wave-numbers of a spherical harmonic model.

In the IGRF model, each of these wave-numbers have weighted coefficients to best fit the observed magnetic field around the world. Wave-numbers are also useful for




Figure 7. Spherical Harmonic Power [2]

separating out components of the Earth’s magnetic field. Fig. 7 shows the power in the Earth’s magnetic field at increasing wave-numbers. It is clear that there is a change in order of magnitude around wave-number 15. Geophysicists have interpreted this break to indicate that for wave-numbers 1–12 the field comes primarily from the Earth’s core, and for wave-numbers greater than 16 the field comes primarily from the Earth’s crust [34]. Wave-numbers 13–15 include contributions from both the core field and the crustal field. The problem of separating magnetic sources is difficult: because of the spatial frequency overlap, it is not possible to independently measure the low-degree crustal or the high degree core field. The IGRF model usually is computed up to wave-number 10. This means the model primarily captures the core field. It is important to note, however, that very small contributions from the crustal field are included in the model and small contributions from the core field are ignored in the model. The intent of subtracting the IGRF from raw measurements is to remove the regional effects of the core field. It is important to realize that removing the core field can never be achieved perfectly, and that calling the residual from this difference the crustal field is an approximation.




Figure 8. Perpendicular Crustal and Core Field Vectors

Anomaly Definition and Assumptions.

We have established that magnetic anomaly maps are primarily capturing the Earth's crustal field. While a magnetic anomaly could be caused by man-made sources or space weather effects, most magnetic surveys take steps to avoid man-made sources corrupting the measurements, and temporal variations are recorded at a base station and removed. Mitigating man-made effects and temporal variations helps ensure the magnetic anomaly is due to crustal sources alone. Previously, it was stated that magnetic anomaly maps are created by the scalar difference between the measured field and a reference field. The fact that this scalar subtraction has physical significance involves several important assumptions about the vector fields being subtracted. As an exaggerated counter-example, consider two vectors which are orthogonal, B_{crustal} and B_{core} , and their vector summation B_{total} , as shown in Figure 8. It is clear that subtracting the length of B_{core} from B_{total} does NOT yield B_{anomaly} , yet this scalar subtraction is exactly what is done to create magnetic anomaly maps. It is clear that in the given example this subtraction would have no physical meaning.

The magnetic surveys instruments used to make magnetic anomaly maps are inherently limited in that they are scalar intensity instruments. Chapter 3.5 explains why these scalar instruments are used over vector instruments. In short, they are far more accurate. Fortunately, this scalar subtraction does have physical meaning




Figure 9. Projection of Magnetic Anomaly Onto Core Field

due to properties of the Earth’s core and crustal fields. The correct interpretation of a typical scalar magnetic anomaly map is that it represents an approximation to the projection of the magnetic anomaly along the reference field direction. This is only a good approximation if $B_{core} \gg B_{anomaly}$. The fact that this represents the projection along the core field makes sense when the relative sizes of the two vectors are considered. The core field is on the order of 50,000 nano-Teslas while the crustal field variations are on the order of 100’s of nano-Teslas. Because the core field is so much stronger than the crustal field, only crustal sources which “stretch” or “shrink” the total field vector in the core field direction are observable. Fig. 9 shows an example of this projection. To illustrate how accurate the scalar subtraction of $B_{total} - B_{core}$ approximates the projection of $B_{anomaly}$, consider Fig. 10 (not to scale). We approximate the projection of the anomaly vector onto the core field vector using $B_{total} - B_{core}$. In the top part of the figure, realistic values for the Earth’s anomaly field and core field are used. The error is seen to be small. In the bottom part of the figure, we assume a very large anomaly which breaks the assumption that $B_{core} \gg B_{anomaly}$. In this case, the projection of the anomaly field onto the core field cannot be given by the simple scalar subtraction $B_{total} - B_{core}$. It is important to note the non-linearity in this assumption. The bottom example shows an anomaly which is 200 times greater than the top example, yet the error is over 40,000 times greater.




Figure 10. Example of Projection Assumption Accuracies




Figure 11. Induced and Remnant Magnetization [27]

A second property of the crustal field also increases the validity of the scalar subtraction $B_{\text{total}} - B_{\text{core}}$. In the examples in Fig. 10, the vector direction of the anomaly was 45 degrees. The core and anomaly field vectors, however, are often more aligned in reality. This is because of the two types of magnetization in the Earth's crust: remnant magnetization and induced magnetization. If a mineral with a high magnetic susceptibility and no remnant magnetization is not in the presence of a magnetic field, it has no magnetic field. When an external field is applied to it, such as the Earth's core field, it has a magnetic field which is aligned with the inducing field. When a mineral has remnant magnetization it has a magnetic field even when there is no external field present. If a mineral has both remnant and induced magnetization and is placed within an external field, its total magnetic field will be the vector sum of the induced and remnant magnetic fields. Fig. 11 shows an example of both remnant and induced magnetization. It is clear then that in the absence of remnant magnetization, where a mineral only has induced magnetization, the B_{core} and B_{anomaly} vectors would be aligned. In this case, it is clear that it is perfectly acceptable to subtract the scalar B_{core} from B_{total} to get B_{anomaly} . In reality, induced magnetization is indeed the dominant form of mineral magnetization within the Earth's crust, leading to closer alignment between the anomaly and core

fields [34].

2.3 Magnetic Anomaly Modeling and Transforms

Understanding the physics of magnetic anomaly fields is essential to implementing a navigation filter. This section describes how to perform the necessary operations of upward continuation and time transformations, as well as how to create an accurate model of a magnetic anomaly field.

Upward Continuation.

Magnetic anomaly fields are potential fields. Potential fields, by definition, satisfy Laplace’s equation. Laplace’s equation states that the sum of the curvatures of a function in each direction of the function equals zero [10]:

$$\frac{\delta^2 \phi}{\delta x^2} + \frac{\delta^2 \phi}{\delta y^2} + \frac{\delta^2 \phi}{\delta z^2} = 0. \quad (1)$$

Laplace’s equation directly leads to a set of identities known as Green’s Identities. Green’s first and second identities can be used to derive the upward continuation integral [10]. The upward continuation integral allows the calculation of magnetic intensity at any value above an infinite 2-dimensional plane and is given by

$$U(x, y, z_0 - \Delta z) = \frac{\Delta z}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{U(x', y', z_0)}{[(x - x')^2 + (y - y')^2 + \Delta z^2]^{3/2}} \delta x' \delta y'. \quad (2)$$

From an algorithmic viewpoint, generating a new map at a higher altitude consists of four nested for-loops (Every value on the 2-dimensional grid contributes to a *single* value at a higher altitude). Fortunately, a much more computationally efficient Fourier domain equivalent expression can be derived [10]. The upward-continued

altitude in the Fourier domain can be given by

$$\mathcal{F}[U_{z+\Delta z}] = \mathcal{F}[U_z]F[\psi], \quad (3)$$

where $F[\psi]$ is the upward continuation filter. The upward continuation filter attenuates shorter wavelengths more than longer wavelengths and is given by

$$F[\psi] = e^{-\Delta z|k|}. \quad (4)$$

$|k|$, in this context, is equal to the absolute value of the wavenumber in the frequency domain and is simply

$$|k| = \sqrt{k_x^2 + k_y^2}. \quad (5)$$

The upward continuation filter is real valued and therefore does not affect the phase of the map data. Fig. 12 shows the magnitude of the upward continuation filter in the Fourier domain for a upward continuation height of 1 kilometer with spatial frequencies ranging from 0–1 Hz (cycles/kilometer). Taking the inverse Fourier transform of $\mathcal{F}[U_{z+\Delta z}]$ is the final step needed to generate a map at a higher altitude. There are two key assumptions in the upward continuation that are not ever met in practice. The first is that an infinite horizontal map tile is available. Obviously, real map tiles will be finite and violate this assumption. The second assumption which is violated in practice is that all magnetic sources lie below the map tile. The total magnetic field measured in the vicinity of the Earth is a superposition of many magnetic sources. While the largest of these sources are internal to the Earth, there are external sources which will be above a map tile as well. More sophisticated methods of upward continuation may be utilized which use both aerial survey data and satellite data to perform a more accurate upward continuation and better reflect the external sources [27]. It is important to note that the upward continuation filter is essentially a low pass filter.




Figure 12. Upward Continuation Filter Spectrums




Figure 13. Edge Effects From Upward Continuation

This means that the higher we upward continue data, the more the high frequency features will be attenuated. This is important from a navigation standpoint, because it is these high frequency features which contribute to high accuracy navigation. Intuitively, then, navigation accuracy will decrease at higher altitudes.

We wish to understand how errors manifest themselves in the upward continuation process. We can begin by considering the assumption that we have an infinite sized map tile in the horizontal direction. From the upward continuation integral, we can see that every point on the infinite horizontal map contributes to the intensity at a given higher altitude. Because we have a finite sized grid, this will cause errors in the upward continuation. Clearly this effect is amplified at the edges of the map. As shown in Fig. 13, upward continuing near the edge of the map immediately depends on values outside the grid. Fig. 13 also demonstrates how higher altitude upward continuations depends on values which are further away horizontally. We can use the upward continuation integral to approximate the contribution from a torus shaped ring of area as a function of horizontal distance r as shown in Fig. 14 and given by

$$\text{Weight} = \frac{2r}{[r^2 + \Delta z^2]^{3/2}}. \quad (6)$$




Figure 14. Contribution of Ring of Area as a Function of Horizontal Distance

We considered the weight of a section of area, rather than an individual map point, because as you move horizontally away from the upward continued point being calculated an increasing amount of grid points contribute to the calculated value. Notice the breaks in the log-log plot occur at the upward continuation altitude. All points on a map which are at a lateral distance less than or equal to the upward continuation height have a significant effect on the upward continued value. Moving further away laterally than the upward continuation height, the weight of the points drops off with a slope of -1 for the first order of magnitude past the upward continuation altitude, settling on a final slope of -2. The slope of the weight function can give intuition as to the size map needed to perform an upward continuation without introducing larger errors. It is apparent that the dominant variable in determining needed map size is the upward continuation altitude—the point at which the weight function begins to decrease exponentially. From the slope of the loglog weight function a rough approximation can be made. Moving ten times further away than the upward continuation altitude decreases the weights of a torus shaped ring of area by a factor of

10. Moving 100 times further away than the upward continuation altitude decreases the weights by a factor of 1000. Therefore, a map must contain data which is 1–2 orders of magnitude larger than the upward continuation altitude around the horizontal location of the upward continued point being calculated. This approximation lets us quantify the edge effects of a map. If a 100 kilometer $\times$ 100 kilometer map is being upward continued 1 kilometer, at a minimum the outside 10 kilometers of the upward continued map should be considered inaccurate.

An additional source of error is introduced in upward continuation comes from the Fourier domain implementation of the transform. The two dimensional Fourier transform assumes a periodic map. The sharp discontinuities on the edge of the map will introduce errors. It is recommended to apply a tapering of the edges of the map to “wrap-around” and make the map periodic [10]. This process will be further examined in the Preconditioning section within Chapter 3.3.

Magnetic Map Time Projections.

The magnetic anomaly field is primarily an induced field [34]. There are two main causes for a mineral’s magnetization. The first is remnant magnetization. This is caused by a past induced magnetic field changing to a permanent magnetic field when the mineral cooled below the Curie temperature at its initial formation [55]. The second type of magnetization is induced magnetization. This occurs when the present day magnetic field induces a magnetic field in the minerals. An induced magnetic field will have the same orientation as the inducing field. Magnetic anomalies are induced primarily by the Earth’s core field. The core field is well modeled by the International Geomagnetic Reference Field. The IGRF has been modeled for many years and significant changes have been observed. Over the past 30 years over Ohio,

the declination of the core field has changed by over 6 degrees. This change in the inducing field will clearly change the magnetic anomalies which are being induced. This could make an older map inaccurate for use in navigation. Fortunately, there is a method to account for the change the magnetic field will undergo if the previous and current inducing fields are known. This is possible because, fundamentally, the field is only a function of the inducing field and the distribution of magnetic minerals in the ground. These magnetic mineral distributions only change on geological time-scales, and can be considered static for navigation purposes. The process of transforming magnetic anomaly maps to reflect different inducing fields is primarily known as Reduction to the Pole (RTP) or Reduction to the Equator (RTE) in the field of geophysics. Under an inducing field, a symmetric distribution of minerals in the ground can appear on a magnetic anomaly map to be asymmetric. Furthermore, because of the inducing field, the magnetic anomaly may be shifted laterally on map as well. When studying magnetic anomaly maps for geophysical exploration purposes, these asymmetries and lateral shifts are not desirable. Geoscientists perform the RTP transforms to make a magnetic anomaly map appear how it would at the magnetic pole, removing asymmetries and lateral shifts. This technique can also be used when stitching together different map tiles from surveys flown at different times. The map tiles can all be transformed to a common year before being stitched together. To achieve the highest levels of navigation accuracy, all magnetic maps should be transformed to reflect the present day IGRF. The Fourier domain time reduction transform is derived in Blakely [10] and given by

$$\mathcal{F}[\psi_t] = \frac{\Theta'_m \Theta'_f}{\Theta_m \Theta_f}, \quad (7)$$

where Θ'_m and Θ'_f are derived from the new inclination and declination of the source field (m) and ambient field(f) and Θ_m and Θ_f are derived from the old inclination

and declination and of the source and ambient fields. When making the assumption that the magnetic fields are completely induced, $m = f$, and a simplified expression can be given as

$$\mathcal{F}[\psi_t] = \left(\frac{\Theta'_f}{\Theta_f} \right)^2, \quad (8)$$

where both the old and new Θ_f are given by the complex expression

$$\Theta_f = f_z + i \frac{f_x k_x + f_y k_y}{|k|}, \quad (9)$$

where:

$$\begin{aligned} f_x &= \cos D \cos I, \\ f_y &= \cos D \sin I, \\ f_z &= \sin I, \end{aligned} \quad (10)$$

and k_x k_y are the x and y wavenumbers with $|k| = \sqrt{k_x^2 + k_y^2}$. The declination and inclination of the core field are give by D and I .

In practice the $k_x = k_y = 0$ wavenumber must be set equal to 1, otherwise a divide-by-zero error will occur. The zero-centered magnitude spectrum of the filter is shown in Fig. 15. The spectrum will always consist of rays from the origin with a constant value. The phase spectrum is shown in Fig. 16. It is important to note that these filters assume a constant inclination and declination of the Earth's core field over the mapped area. This is a good assumption for smaller regional maps but may not be adequate for larger maps. Time reduction transforms do not suffer from edge effect errors as much as upward continuation transforms. In upward continuation, the values outside the map are needed at higher altitudes. This is not the case in time reduction, but the errors will still primarily manifest themselves at the edges of the map due to breaking assumptions such as having a periodic map and Gibbs




Figure 15. Magnitude of Sample Time Reduction Filter




Figure 16. Phase of Sample Time Reduction Filter

phenomenon errors. The preconditioning section of Chapter 2.3 will examine ways to minimize these types of errors.

Modeling Magnetic Anomaly Fields.

When analysing and testing the upward continuation transform and time transform it is useful to have an accurate mathematical model of a magnetic anomaly field. Analyzing synthetic map data has the benefit of having a true solution. As mentioned previously, aeromagnetic surveys are only flown at one altitude. To test upward continuation with real data, you would need maps at multiple altitudes. The creation of accurate synthetic magnetic anomaly maps is vital to an accurate analysis. To create the maps, we were motivated by Equivalent Source Dipole Inversion (ESDI) techniques [72]. ESDI is a way to represent a magnetic anomaly map as an equivalent layer of magnetic dipoles with varying susceptibility. An assumption is made that the magnetic sources are completely induced. The magnetic dipoles will therefore have the same orientation as the inducing field—in this case the IGRF. A least squares method is used to solve for the needed susceptibility of each dipole to match a given magnetic anomaly map. If performed correctly, the resulting model is an accurate representation of the magnetic anomaly field not only at the grid point, but also at higher altitudes and under different inducing fields. In light of this technique, we created synthetic magnetic anomaly maps by placing these dipoles on a grid with the same orientation as the IGRF, and with varying but correlated magnetic susceptibilities. The dipoles are placed on the ground at spacing equal to the desired map altitude. This simple technique gives the synthetic maps similar frequency content to a real map, in which the shortest wavelengths are approximately equal to the altitude [53]. The equations for the magnetic anomaly intensity from a single dipole are:

$$B_T = \cos(I) \cos(D) B_x + \cos(I) \sin(D) B_y + \sin(I) B_z, \quad (11)$$

where D and I are the IGRF magnetic declination and inclination, respectively, and B_x , B_y , and B_z are given by

$$B_x = \frac{3C_m}{r^5} \left[j_x \left(\Delta x^2 - \frac{r^2}{3} \right) + j_y (\Delta y \Delta x) + j_z (\Delta z \Delta x) \right], \quad (12)$$

$$B_y = \frac{3C_m}{r^5} \left[j_x (\Delta x \Delta y) + j_y \left(\Delta y^2 - \frac{r^2}{3} \right) + j_z (\Delta z \Delta y) \right], \quad (13)$$

$$B_z = \frac{3C_m}{r^5} \left[j_x (\Delta x \Delta z) + j_y (\Delta z \Delta y) + j_z \left(\Delta z^2 - \frac{r^2}{3} \right) \right], \quad (14)$$

where j_x , j_y , and j_z are given by

$$j_x = J \cos(I') \cos(D'), \quad (15)$$

$$j_y = J \cos(I') \sin(D'), \quad (16)$$

$$j_z = J \sin(I'), \quad (17)$$

where J is the magnetic susceptibility of the dipole and I' and D' are the inclination and declination of the dipole vector. The location with respect to the magnetic dipole at which the magnetic field is being calculated is given by Δx , Δy and Δz , and r is the three dimensional distance from the observation point to the dipole. C_m is the magnetic constant equal to 10^{-7} (Newtons per square meter). We assume the dipoles are completely induced, with no remnant magnetization, so $I = I'$ and $D = D'$. Under these assumptions, the entire field for a single dipole is determined only by 3 parameters—the IGRF inclination and declination and the magnetic susceptibility. After placing dipoles throughout a grid at spacing equal to the altitude, the final map can be created by summing up the contributions of each dipole.

Figure 17 shows the synthetic map created using the previously described process.




Figure 17. Synthetic Map Created at 1km Altitude and Subset Map

A box is drawn around a subset of the map. This is the actual data which will be passed to the upward continuation filter. This ensures we have an accurate representation of the edge effects. When we analytically calculate the higher altitude truth map it will be a function of data outside of the boxed area. This is of course the situation that always exists in reality because our maps are finite. Because we have a complete analytical description of the map we can easily compute the true solution at any altitude and compare it with what the upward continuation function provides. We already motivated that the size of a map being upward continued should be 1-2 orders of magnitude larger than the upward continuation altitude to achieve negligible error at the center of the map. Another way of stating this is that all finite maps will have errors on their edges with widths equal to 1–2 orders of magnitude larger than the upward continuation altitude. We demonstrate this effect on the synthetic maps, which are 10×10 kilometers. From our rule of thumb, we would not expect a 500 meter upward continuation to be valid on a map this small. Fig. 18 shows the results of this upward continuation.

It is clear that the map has large errors covering a significant portion of the map.




Figure 18. True and Calculated Synthetic Map Upward Continued 500m from 1km




Figure 19. True and Calculated Synthetic Map Upward Continued 50m from 1km

Again, this result was expected given that a map size of 50×50 kilometers would be more appropriate for an upward continuation of this altitude. Fig. 19 shows the results of upward continuing the map 50 meters. This altitude meets the requirement of being 1–2 orders of magnitude less than the size of our map. As can be seen in Fig. 19, the errors are only significant on the very edges of the map. This is unavoidable and will always be an artifact of upward continuation.


As discussed previously, the magnetic anomaly field changes slightly over time as the external inducing field changes. We wish to quantify the magnitude of this change using our synthetic maps as well as actual observed changes in the Earth’s core field. The core field over southwest Ohio presently (2015) has a declination of

-5.57 degrees and an inclination of 67.313 degrees. Five years ago (2010) these values for declination and inclination were -5.037 and 68.135, respectively. We computed the synthetic maps for both of these values of the inducing field and compared the error between the two maps. Fig. 20a shows how much the map changed over the given five year period. We then took the old map from 2010 and performed a time reduction on it to the current year inducing field values for declination and inclination. Fig. 20b shows the error between the map transformed to the year 2015 and the true 2015 map. As shown, the errors are greatly reduced from the case of not performing any time reduction. As with upward continuation, the majority of the errors lie on the edge of the map. These negative artifacts can be reduced by pre-conditioning techniques, described shortly. Fig. 21 shows similar plots for a map being projected from the year 2000 to the year 2015. Again, a great deal of improvement can be seen in the error between transformed and non-transformed maps. Any map being used for navigation should be passed through a time-reduction filter to bring it up to date with the present day IGRF values.


Map Pre-Conditioning.

Preconditioning the map tiles can lead to greater accuracy in both upward continuation and time reduction filters. The Fourier transform assumes all signals are periodic. The abrupt edges at the end of a map tile violate this assumption and can lead to edge effects from Gibbs Phenomenon. The following procedures can help mitigate these effects:

1. Make the map data have zero mean
2. Remove linear trends from map data
3. Make the map data periodic




(a) Error over 5 Years with No Time Reduction Filtering




(b) Error over 5 Years with Time Reduction Filtering

Figure 20. Time Reduction Filtering on 5 Year Old Map



(a) Error over 15 Years with No Time Reduction Filtering



(b) Error over 15 Years with Time Reduction Filtering

Figure 21. Time Reduction Filtering on 15 Year Old Map




Figure 22. Removing a 2D Trend

Removing the mean from the dataset is self-explanatory. For procedure 2, remove any general upward or downward trends in the map data from east to west or north to south. This does not have a large effect on procedures such as upward continuation because of the low-pass nature of the filter. Making the map data periodic can be done by expanding copies of the map in each direction and ensuring smooth transitions between map tiles through a cosine roll-off function. Fig. 22 illustrates removing a one dimensional trend and Fig. 23 shows an example of this procedure.

2.4 Magnetic Anomaly Maps

When navigating with a map-based system, map availability and map accuracy are two essential factors in determining the feasibility and accuracy of the navigation system. Determining the accuracy of a magnetic anomaly map is a difficult problem. Magnetic anomaly maps have been created all over the world over the last 70 years. Before the widespread use of GPS around 1990, the magnetic sensor readings were geo-located using relatively crude aerial photography methods. Although there have been improvements in actual sensor accuracy over time, the pre-GPS era maps are almost certain to be filled with inaccuracies due to poor geo-location. It is difficult




Figure 23. Example of a Smoothed Periodic Map Tile [20]

to distill the accuracy of these maps down to a simple number, i.e. “The maps are accurate to within 10 nano-Teslas”. When reading the description of old map tiles it is not uncommon to find accuracy claims around 10 nano-Teslas. It may be helpful to keep a level of skepticism regarding these claims for pre-GPS era maps, especially with regard to the exact positioning of the map. Maps which were created with the aid of GPS and differential GPS have the added benefit of also having more modern accurate magnetic sensors. Most magnetic intensity surveys are flown with optically pumped cesium magnetometers with an absolute accuracy of about 1–3 nano-Teslas. With the use of differential GPS to geo-locate the measurements, these maps are quite accurate. The following sections discuss map types and flight conditions with respect to navigation accuracy.

Flying Over a Modern Map At Map Altitude - The Ideal Case.

As alluded to previously, there are several major variables which will effect navigation accuracy in a magnetic navigation system. Using a map with GPS era position accuracies is a clear requirement to achieve the best navigation accuracies. Using a more recently created magnetic map can also improve navigation accuracy. This is due to the previously described minor changes in the magnetic anomaly field over time. These changes can be mitigated with the time-transform techniques presented, but map error will generally be smallest with a more recent map. Flying at the mapped altitude is also necessary to achieve the best possible navigation accuracies. When flying above the mapped altitudes, the map data must undergo the upward continuation transform. This transform makes several assumptions that do not hold in practice, such as having an infinite sized map in the horizontal direction, and having all magnetic sources originate under the map tile. Breaking these assumptions will increase map error when flying at altitudes which require upward continuation (Note: downward continuation is not normally required for aircraft navigation because most map surveys are already flown at very low altitudes). Flying at higher altitudes is also inherently less accurate due to the fact that much of the high-frequency/high-gradient parts of the crustal signal get filtered out as altitude increases.

The line spacing that was flown to create the magnetic surveys is also important to navigation accuracy. A magnetic navigation system is given a grid of magnetic values and must interpolate between values on the grid to estimate the magnetic intensity at any location on the map. As previously described, surveying at a line-spacing less than or equal to survey altitude ensures the field has been fully sampled according to the Nyquist theorem. If an aircraft flew with 8 kilometer line spacing at 1 kilometer altitude, the true signal cannot be reconstructed through any interpolation meth-

ods. When navigating using an under-sampled map, high frequency content in the measurements which was not captured in the under-sampled map will corrupt the signal and decrease navigation accuracy. It is important to note that this problem is mitigated by flying at higher altitudes. If a magnetic survey was flown at 8 kilometer line spacing at an altitude of 1 kilometer, and the data is upward continued to 8 kilometers, the field will be fully sampled. There is of course negative aspects previously described with flying at higher altitudes with respect to navigation accuracy, so a tradeoff must be made at some point.

Meeting these requirements requires the use of a regional map tile. There are currently no large scale maps created with all post-GPS era magnetic surveys. This indicates best performance could not be achieved on a cross country or other long distance flight. New magnetic maps are being created all the time and many countries do have large scale magnetic maps created with modern instruments including GPS. The United States does not have any continental sized maps created with modern instruments. What exists instead is a patchwork of modern accurate surveys as well as older, less accurate surveys. The ownership of the maps is another potential issue. Surveys which were financed by governments are generally available to the public but private maps created by industry are generally unavailable. Fig. 24 shows the locations of post-GPS era high accuracy maps available over the United States freely available through the United States Geological Survey. Maps which were created after 1990 are outlined in black. A navigation system in practice could use these high accuracy maps when available and default to one of the larger-scale maps described in the following sections when unavailable.




Figure 24. Freely Available Maps from the USGS (Post-1990 Era Maps Outlined in Black)

North American Magnetic Anomaly Map.

The North American Magnetic Anomaly Database (NAMAD) was created as a joint effort between the United States Geological Survey, the Geological Survey of Canada, and the Consejo de Recursos Minerales of Mexico [3]. It is freely available online through the USGS [68]. Fig. 25 shows the magnetic terrain of the entire North American continent. Fig. 26 shows the individual component surveys which were used to create the map over the United States. It can be seen in Fig. 26 that a large majority of the surveys on the east and west sides of the country were flown at line spacing of 2 kilometers or less, and the central part of the country was flown at 4–8 kilometers or less. The line spacings from the individual survey components of the North American Magnetic Anomaly Map indicate that this single map could be used for navigation at any altitudes greater than 4.8 kilometers almost anywhere over the continental US, with decreased accuracy expected over the few survey areas which were flown at 8 kilometer line spacing. The map is technically at an altitude of 305 meters but it very under-sampled at this altitude. The feasibility of using the North American Magnetic Anomaly Map for navigation at lower altitudes will likely require empirical testing.

World Digital Magnetic Anomaly Map.

The World Digital Magnetic Anomaly Map (WDMAM) is the result of an international effort to create a world magnetic anomaly map [43]. The map represents the Earth’s magnetic anomaly field at 5 kilometer altitude with a 3 arc-minute resolution (approximately 5.5 kilometers at the equator). The World Digital Magnetic Anomaly Map is freely available online and can be seen in Fig. 27. The WDMAM is subject to the same issues as the NAMAD—it was created from data with varying accuracies. Unlike the NAMAD there are areas where aeromagnetic data is completely absent,




Figure 25. North American Magnetic Anomaly Map – 1km Grid Spacing at 305 meter Altitude [3]




Figure 26. North American Magnetic Anomaly Map Contributing Surveys and Line Spacings [3]

especially over the oceans. Aircraft flying at altitudes of 5 kilometers or greater could potentially use the WDMAM to navigate. Determining the feasibility of using the WDMAM for navigation will likely require empirical testing.

2.5 Types of Magnetic Measurements

There are several types of magnetic measurements which could potentially be used in a navigation system. Each measurement type has advantages and disadvantages on both a theoretical level as well as stemming from practical engineering issues. Understanding these advantages and disadvantages is important for choosing the best measurement type for a given navigation system.

Scalar Intensity Measurements.

The magnetic field surrounding the Earth is a vector field, and consequently has both magnitude and direction. A magnetic intensity sensor can only sense the magnitude of the vector field at a single point in space. The magnitude, B_T is given by

$$B_T = \sqrt{B_N^2 + B_E^2 + B_D^2}, \quad (18)$$

where B_N , B_E , and B_D are the north, east, and down components of the vector field respectively. On a theoretical level, an intensity measurement contains the least amount of information of any of the possible measurement types. If all other variables were equal, a navigation system using only intensity measurements should be outperformed by the other measurement types of equal accuracy. Furthermore, a scalar intensity measurement measures the total magnetic field, when in practice we need to measure the anomaly field. Therefore a scalar intensity measurement will always be corrupted by non-anomaly field sources. These sources are primarily the aircraft field and temporal variations caused by space and ionospheric weather conditions. From




Figure 27. World Digital Magnetic Anomaly Map [43]

a practical standpoint, scalar intensity measurements have many benefits. The first benefit relates to their accuracy. Scalar intensity instruments are far more accurate than the current state of the art vector sensors. Modern magnetic intensity instruments have absolute accuracies on the order of 1 nano-Teslas and sensitivities on the order of pico-Teslas [27]. These instruments are so accurate that they are usually not the dominant factor in determining the accuracy of crustal field measurements. The ability of an aircraft compensation system to remove aircraft fields contributes more to the absolute accuracy of anomaly field measurements than the actual sensor absolute accuracy. Another benefit of magnetic intensity measurements is their simplicity. As a stand-alone instrument there is no need to be concerned with orientation to other sensors. Finally, the majority of existing magnetic maps are made with scalar sensors, lessening the need for magnetic mapping.

Intensity Gradient Measurements.

Two or more magnetic intensity sensors separated by a known baseline may be used to form an intensity gradient measurement. Because magnetic intensity exists as a three dimensional field, the true magnetic gradient is a vector with both a magnitude and a direction in three dimensions. A gradient measurement fundamentally has more information than a scalar intensity measurement due to the two directional components of the measurement. Measuring the magnetic intensity gradient vector in a world frame has several difficulties. Two magnetic intensity sensors are required to measure just one component of this vector. Survey aircraft often have instruments configured in such a way to compute two orthogonal horizontal gradients, as well as a vertical gradient. Fig. 28 shows the magnetic gradient system used by Sander Geophysics. The two wing-tip instruments measure a body frame transverse component of the gradient. The wing-tip gradients are then averaged and differenced with the



Figure 28. Sander Geophysics Magnetic Gradient System [22]

top tail instrument to form a body frame longitudinal component of the gradient. Finally, the two tail instruments are differenced to form a body frame vertical gradient measurement. It is clear that a system such as this can measure the complete gradient vector, including its magnitude and direction. However, this vector would exist in an aircraft body frame. If a navigation system had a stored gradient map, it would exist in a world frame. The navigation system would then have to complete a coordinate transform of the body frame gradient vector into a world frame. This transform would require the aircraft attitude. A navigation-grade INS may have attitude solutions with accuracies on the order of 0.01 degrees after an hour of flight, and could potentially be used to complete this coordinate transform. Even with a navigation grade INS, however, this transform would introduce measurement errors. The body-frame gradient could potentially be used to determine attitude information itself when compared to a stored world-frame gradient map; however this would require a known position.

Despite the previously mentioned difficulties, there are still many benefits to a gradient measurement. If a map of just the magnitude of the total gradient is stored in a navigation system, the coordinate transform is no longer an issue. This type of

measurement has the benefit of removing temporal variations—the ionospheric and space weather effects described previously. The temporal variations vary over time but also vary spatially. The wavelengths of these spatial variations are far greater than the size of an aircraft. This indicates that the temporal variations at any moment in time will be common to each instrument and will be subtracted out when a gradient measurement is formed. Removing these corrupting effects could potentially increase navigation accuracy.

A practical issue with gradient measurements is the baseline, or separation, of the sensors. A true theoretical gradient measurement could only be formed by two sensors separated by an infinitesimally small distance. In reality, due to the limited sensitivity and accuracy of the sensors, the instruments tend to give more accurate gradient measurements the further apart they are on the aircraft. This measurement is therefore the average gradient between the two sensor locations, and not the gradient at a single point. The gradient changes by a very small amount between each sensor, so calling the average gradient between the sensors the true gradient is not a bad assumption. It is important to note that a gradient measurement does *not* help remove aircraft field corrupting effects. Temporal variations are removed because the spatial gradient is almost zero—temporal variations are similar up to 50 kilometers apart. The aircraft effects are not removed even if two sensors are placed infinitely close together, because the aircraft gradient itself is not zero. A gradient-based system measures a gradient, and the aircraft field has a non-negligible gradient, unlike the temporal variations.

Vector Measurements.

Vector measurements can directly measure the components of the total magnetic field vector. On a theoretical level, a vector measurement contains more information than a magnetic intensity measurement. A navigation system could store either three maps for each vector component in a world frame, or three maps consisting of magnitude, inclination, and declination. These maps would contain the same information overall, although the later would likely be more helpful for comparing performance to a intensity-only system. The greatest issue with using a vector system lies in the accuracy of current vector instruments. These instruments are typically 1–2 orders of magnitude less accurate than current systems. Whatever is gained by bringing in more information to the navigation filter is likely lost by the poor quality of the measurements with current sensors.

A vector system clearly would suffer from some of the same limitations described for gradient-vector systems. The largest issue is the need for aircraft attitude to transform the body frame measurements into a world frame. As before, a navigation grade INS may suffice for short duration flights. If the magnetic navigation system performs well enough, a boot-strapping effect may occur where the magnetic measurements keep the navigation grade INS attitude accurate enough for continuous use of vector measurements. The vector measurements are also subject to some of the previously described limitations of a intensity measurement—the vector sensor is measuring the total magnetic field, not the crustal field. As with the intensity measurements the vector measurements will be corrupted with temporal variations and aircraft fields. It is also important to note that even small errors in aircraft attitude could lead to large errors when resolving vector measurements. For example, a 0.01 degree error when resolving a 50,000 nano-Tesla magnetic field measurement

(approximate magnitude of core field) could lead to nearly 10 nano-Teslas of mis-resolved core field projected in the wrong direction. Finally, magnetic vector maps have not been widely produced around the world, unlike magnetic intensity maps. Even if a high accuracy magnetic vector instrument was created, there would still be a large mapping effort needed to take advantage of this sensor.

Tensor Measurements.

A tensor measurement is the vector equivalent of the magnetic intensity gradient measurements. A tensor measurement gives the derivative of the three vector components, B_{xx} , B_{yy} , and B_{zz} , as well as the three unique cross derivative terms B_{xy} , B_{xz} , and B_{yz} , as shown in Fig. 29. Tensor measurements are clearly limited by the same inaccuracies of the vector instruments required to compute the tensor. If accurate vector magnetometers were available, a tensor measurement could potentially bring the most information into a navigation filter while simultaneously subtracting corrupting fields such as the temporal variations and aircraft fields. Like vector measurements, tensor measurements would require re-mapping of the magnetic field to capture the tensor value.

2.6 Magnetic Sensors

There are a large variety of magnetic sensors which exist to measure magnetic fields. This section describes the two magnetic instruments which are used most often in airborne applications. It does not attempt to give an in-depth description of the physics of the instruments, but rather an analysis on their use for airborne applications, including specifications of common instruments. Many other types of magnetometers exist for laboratory use which are not described here.




Figure 29. All 9 Components of a Magnetic Tensor (6 unique) [69]

The main instrument used in most aeromagnetic surveys is an optically pumped or alkali-vapor magnetometer [27]. These instruments measure magnetic intensity and are absolute instruments, using the atomic properties to measure the absolute magnetic field. The optically pumped magnetometers provide the measurements which actually make the magnetic anomaly maps. The second type of magnetometer used is the flux-gate magnetometer. These are vector instruments and measure the *relative* magnetic field. They can only measure the magnetic field with respect to some un-calibrated baseline. The flux-gate magnetometers are used to help estimate and remove the corrupting aircraft magnetic fields [27].

Optically Pumped / Alkali-Vapor Magnetometers.

Almost all modern magnetic surveys use a nuclear resonance magnetometer as their primary survey instrument. Nuclear resonance magnetometers use the atomic properties of gasses which are sensitive to an external magnetic field [27]. There are several types of nuclear resonance magnetometers including proton-precession,



Figure 30. Geometrics 823A Magnetometer [21]

alkali-vapor, and Overhauser magnetometers [27]. Alkali-vapor magnetometers, also known as optically-pumped magnetometers, are the instruments of choice in industry for aero-magnetic surveys due to their superior sensitivities and faster sampling rates [27]. Optically pumped magnetometers can use several different types of alkali vapors including cesium and potassium [40]. See [19] for a description of the physics of an optically pumped magnetometer which uses cesium gas.

Optically pumped magnetometers are a mature technology. They are small, lightweight, stand-alone instruments which can digitally provide an extremely sensitive and accurate magnetic field measurement. Fig. 30 shows a common airborne magnetometer, the Geometrics 823A. There are several important sensor specifications for an optically pumped magnetometer. A specification sheet from the Geometrics 823A is provided in Fig. 31.

MODEL G-823A AIRBORNE CESIUM MAGNETOMETER SENSOR SPECIFICATIONS

OPERATING PRINCIPLE:	Self-oscillating split-beam Cesium Vapor (non-radioactive)
OPERATING RANGE:	20,000 to 100,000 nT
OPERATING ZONES:	The earth's field vector should be at an angle greater than 10° from the sensor's equator and greater than 10° from the sensor's long axis. Automatic hemisphere switching.
SENSITIVITY WITH CM-201:	< 0.004 nT/√Hz rms. Typically 0.02 nT P-P at a 0.1 second sample rate (90% of all readings falling within the P-P envelope) using CM-201 Mini-Counter
HEADING ERROR:	± 0.15 nT over entire 360° equatorial and polar spins Not specified on 823B
ABSOLUTE ACCURACY:	< 3 nT throughout range
OUTPUT:	Cycle of Larmor frequency = 3.498572 Hz/nT, RS-232 data at 9600 baud, concatenated data streams from up to 6 sensors
MECHANICAL:	
Sensor:	2.375" (60.32 mm) diameter, 5.75" (146 mm) long, 12 oz. (339 g) – any orientation in 7" (177.8 mm) diameter stinger
Sensor Electronics:	2.5" (63.5 mm) diameter, 11" (279.4 mm) long, 22 oz. (623 g)
CABLES: Sensor to electronics: Electronics to Junction Box:	Standard 109 in. (9 ft. or 2.77 m) with connector on electronics end. Other lengths available from 2.4 ft. (0.75m) to 12.75ft (3.75m) at 40 inch (1 m) increments. Lengths approx. due to cable variations. RS-232 to Computer, standard 8 m, 60 m max. Larmor to external counter with coupler over Coax, standard 10m, 50m max
OPERATING TEMPERATURE:	-30°F to +122°F (-35°C to +50°C)
STORAGE TEMPERATURE:	-48°F to +158°F (-45°C to +70°C)
ALTITUDE:	Up to 30,000 ft. (9,000 m)
WEATHERPROOF:	O-Ring sealed for operation in the rain and/or 100% humidity
POWER:	24 to 32 VDC, 0.75 amp at turn-on and 0.5 amp thereafter

Figure 31. Specification Sheet for Geometrics 823A [21]

Operating Range.

Optically pumped magnetometers have a limited range of magnetic field magnitudes they can measure. These instruments are designed to measure Earth magnetic fields. If in the presence of strong man-made fields, the instruments may fail. Fields larger than the max of 100,000 nano-Teslas are not likely to be encountered when flying on an aircraft.

Operating Zones.

Optically pumped magnetometers are subject to dead zones. The instruments must be orientated at a certain angle with respect to the total Earth field vector being measured. This angle is usually small—the Geometrics 823A experiences a dead-zone when the earth’s vector is less than 10 degrees from either the sensor equator or the sensor’s long axis. Fortunately, entering a dead-zone only temporarily disrupts measurements. Accuracy of future measurements is not affected by temporarily flying with an orientation that leads to a sensor dead-zone. Normally, when flying an airborne survey, survey lines are all flown in the same direction and the sensor is simply mounted in such a way that the sensor dead-zones will never be an issue. In a non-survey flight, this limitation manifests itself as an engineering problem—it can be addressed by using two instruments mounted at different orientations or having a sensor platform which can rotate. Both of these techniques have been implemented successfully to mitigate the problem of sensor dead-zones. It is important to note that it is the orientation of the total field vector which causes the dead-zones, not magnetic anomalies like crustal sources, aircraft fields, or temporal variations. Magnetic anomalies are so small compared to the main Earth field that they have little influence on the total vector direction. This means it is easy to predict the needed sensor orientation along any given flight using a core field model such as the IGRF.




Figure 32. Sensor Dead-Zones for Geometrics 823A [20]

Sensitivity and Peak-to-Peak.

Sensor sensitivity is not defined uniformly across all manufacturers, although the most acceptable definition describes sensitivity as the larger value of either sensor noise or sensor resolution [60]. Sensor resolution is the smallest increase or decrease in the true value being measured which the sensor can detect. For instance, the sensor may report the magnetic field to five significant figures but only be able to detect changes to four significant figures. Sensor resolution often depends on sensor rate. Sensor noise also depends on sensor rate, and the units for all three of these specifications (resolution, noise, and sensitivity) are given in units of nano – Teslas/ $\sqrt{\text{Hertz}}$ RMS. The units of $\sqrt{\text{Hertz}}$ indicate that if the frequency is increased by a factor of four, the sensor noise will only increase by a factor of two. This frequency dependency arises from the fact that at lower sample rates the white noise begins to average out and therefore decreases. The peak-to-peak value can be defined differently by different manufacturers [60]. The Geometrics 823-A lists their peak-to-peak value as 0.02 nano-Teslas at 10 Hertz, and state that 90% of all measurements will fall within the peak-to-peak envelope. Sensitivity may be measured by placing a magnetometer in a carefully controlled magnetic environment where the true field value is constant. In

a perfect instrument the measured value would be constant. The sensitivity, or peak-to-peak value, describe how much the measurements would vary under this constant field. It is important to distinguish sensitivity from sensor accuracy. An instrument could have an extremely high sensitivity and not actually be reporting the correct magnetic field value—it could be off by some constant, or slowly varying bias.

Absolute Accuracy.

Optically pumped magnetometers are absolute instruments. They use known atomic properties of alkali gasses to give an absolute reading of an external field. Absolute accuracy is a tricky specification to describe in the context of measuring the Earth’s magnetic anomaly field. The actual instruments have a certain absolute accuracy, but the absolute accuracy we care about is the accuracy of measuring the Earth’s magnetic anomaly field. The magnetometer measures the total field, including the aircraft magnetic field, which can never be totally removed. The actual sensor absolute accuracy may be well under 1 nano-Teslas for an optically-pumped magnetometer but if the aircraft field can never be removed to that extent, it will be the aircraft compensation system which is driving the absolute accuracy of the sensor with respect to measuring the Earth’s magnetic anomaly field. Sensor specification sheets can be confusing with respect to absolute accuracy specifications. Oftentimes, they describe the absolute accuracy of compensated measurements on board an aircraft, taking into account factors like heading errors, aircraft fields, and temporal variations. The Geometrics 823-A lists the absolute accuracy as less than 3 nano-Teslas. The actual sensor drift, given in [60] as 0.1 nano-Teslas, is stated to be much less than the absolute error. In general, the sensor accuracy itself is better than 1 nano-Teslas but the actual ability to observe the Earth’s magnetic anomaly field is only accurate to 1–3 nano-Teslas. A good way to measure sensor accuracy is to

watch how a gradient measurement drifts in a constant field. Because the gradient measurement is subtracting out the common field between each instrument, any drift is likely due to sensor inaccuracies.

Heading Errors.

Optically pumped magnetometers are subject to heading errors. It is important to note that this is a separate error from the dead zones encountered based on sensor orientation. Heading error is caused by a non-uniform permeability of the sensor materials [60]. As the sensor rotates, the magnetic field vector interacts with different parts of the sensor, and the non-uniform permeability of the sensor materials can create a heading-dependent error. Many manufacturers will provide plots of the heading error as a function of orientation angle, allowing some of the heading error to be estimated and removed [60]. It is important to note that while these errors are normally referred to as “heading errors” they actually are affected by rotation about all three axes, not just heading.

Fluxgate Sensors.

Fluxgate magnetometers are vector instruments which measure the relative magnetic field. They are much less accurate than optically pumped magnetometers which measure only magnetic intensity. Airborne surveys use fluxgate magnetometers to estimate and remove aircraft corrupting fields [27]. The Earth’s magnetic field will induce a secondary magnetic field in an aircraft which depends on the aircraft’s orientation with respect to the Earth’s magnetic field. A survey aircraft will fly a calibration flight over a known magnetic intensity and observe the how the magnetic field changes based on the aircraft orientation. It estimates a set of coefficients which can then be used to remove the majority of the aircraft magnetic field for the rest of

Performance	
Number of axes	Three (right hand XYZ coordinate system)
Polarity	+ve non-inverting output when pointing North
Full-scale measuring range	$\pm 100 \mu\text{T}$
Bandwidth at -3dB	$> 100\text{Hz}$
Measurement noise floor	10 to $\leq 20 \text{pTrms}/\sqrt{\text{Hz}}$ at 1Hz
Scaling	$100 \text{mV}/\mu\text{T}$
Offset error	$\pm 100 \text{nT}$ (max)
Offset temperature coefficient	$\pm 1 \text{nT}/^\circ\text{C}$
Scaling temperature coefficient	$\pm 100 \text{ppm}/^\circ\text{C}$
Calibration error	$\pm 0.5\%$ (at DC)
Orthogonality error	$< 1^\circ$
Alignment error (to mounting face)	$< 1^\circ$
Linearity error	0.0033% (least squares fit)
Frequency response	Maximal flat response ($\pm 5\%$) from DC to 10Hz
Hysteresis	$< 2 \text{nT}$ for exposure to up to 2 x full scale
Excitation breakthrough	$< 20 \text{mV}$ pk-pk at 16kHz

Figure 33. Specifications from Barrington SpaceMag Catalog [4]

the flight.

The most accurate fluxgate magnetometers are accurate to within about ± 100 nano-Teslas. Fig. 33 shows the specifications for a Bartington Space-Mag vector magnetometer. It can be seen that the instruments have very low noise compared to the magnitude of magnetic anomalies but the absolute errors can be as large as 100 nano-Teslas. These absolute errors are the cumulative effect of many sources of error including calibration error, temperature errors, orthogonality and alignment errors. Fluxgate magnetometers will be essential to a magnetic navigation system in order to remove aircraft magnetic fields but at the current level of accuracy are likely not sufficient as the primary navigation measurement.

2.7 Obtaining Accurate Magnetic Measurements

Obtaining accurate measurements of the Earth's magnetic anomaly field is a challenging engineering problem. The raw sensor data must be compensated to remove

the effect of corrupting sources. The magnitude of these corrupting sources can be greater than the sensor accuracy. This indicates that in the context of magnetic navigation, the quality of compensation can be what drives measurement accuracy from a filtering standpoint. Magnetic measurements are corrupted from two main sources—the aircraft, and temporal variations. Compensation of these two effects will be presented in this section, which follows mainly from [52]. It is important to note that not all of the presented procedures can be applied in real-time.

Aircraft Sources.

The aircraft is a source of corruption in magnetometer measurements. There are 4 main sources of corruption from the aircraft [52]:

Permanent Magnetization.

Aircraft permanent magnetization fields are magnetic fields caused by actual magnetic components in the aircraft. These fields stay relatively constant and would appear as a constant bias in the magnetometer measurements. Any part of the aircraft or piece of equipment inside the aircraft which has a permanent magnetic field can contribute to the overall permanent magnetic field. Modern survey aircraft are periodically de-gaussed—a procedure which helps to remove any aircraft permanent magnetic fields [52].

Induced Magnetization.

Induced aircraft magnetic fields are caused by the aircraft flying in an external magnetic field—the Earth’s main field. When a material with a given magnetic susceptibility is placed within a magnetic field, a second magnetic field may be induced. These magnetic fields are heavily dependent on the orientation of the aircraft within



Figure 34. Aircraft Stinger Holding Magnetometer [22]

the Earth's main field. Many aircraft are primarily constructed from aluminum alloys which are non-magnetic [52]. This means the body of the aircraft will not have induced magnetic fields (although it may have eddy current, explained below). The aircraft engines are the largest source of induced magnetic fields [52]. Therefore, one of the primary mitigation strategies for magnetometers is to place them as far away from the engines as possible. This is often accomplished through the use of stingers. Stingers are rigid structures extending from an aircraft to add physical separation between the magnetometers and sources of measurement corruption. Fig. 34 shows a stinger on a geo-survey aircraft. Stingers only can remove a portion of the magnetic field, and therefore additional techniques are needed to remove corrupting magnetic fields. The main technique is the use of modern aeromagnetic compensation systems to estimate and remove aircraft magnetic fields based on the orientation of the aircraft with respect to the main earth field. The details of these systems are presented below.

Aircraft Electronics.

Any steady or changing flow of electrical current within an aircraft can create corrupting magnetic fields. Steady currents create magnetic fields following the Bio-Savart Law and changing electric fields may induce magnetic fields in nearby conductors according to Faradays law of induction. The corrupting effect of these fields is

primarily mitigated by the aircraft magnetic compensation system.

Eddy Currents.

Eddy currents are electrical currents that can run through any conductor in the aircraft. They are caused by the aircraft maneuvering in a magnetic field. These eddy currents create their own magnetic fields which obey Lenz’s Law—that is, they oppose the magnetic field which created them. Eddy currents are also primarily removed by the aircraft magnetic compensation system. Compensating eddy currents requires knowledge of the rotation rate within the external field.

Aircraft Magnetic Compensation Systems.

The corrupting magnetic sources of an aircraft as it moves through an external magnetic field can be quite complicated to measure and model. In the past, attempts were made to identify these individual fields and cancel them out by actively creating opposing magnetic fields with coils of wire near the magnetic sensors [52]. Modern compensation systems are far more elegant and use a pre-flight calibration process to estimate and remove the effects of these fields using a flux-gate magnetometer, which gives the orientation of the aircraft within the Earth’s magnetic field [52]. The calibration procedure works by flying the aircraft in a square pattern (4 different headings) and performing a series of maneuvers in the roll, pitch, and yaw axis by about 5-10 degrees [52]. The compensation flight is flown away from cultural anomalies and over an area of low crustal field variations. The flight is also flown at high altitudes which further diminish the effect of the crustal sources. These flight conditions ensures that any variations seen in the measurements are due to the aircraft changing its heading, roll, and pitch within the Earth’s main field, and not a changing Earth magnetic field. A calibration routine then minimizes the observed variation in




Figure 35. Before and After Magnetic Compensation [52]

the measurements by estimating different coefficients in an aircraft magnetic model [52]. These coefficients, as well as the orientation within the Earth’s main field given by the flux-gate magnetometer, are used to remove errors from the measurements for the remainder of the aircraft flight. Fig. 35 shows an example of the observed variations of a magnetometer before and after magnetometer compensation. Note that the actual value of the magnetic field does not need to be known—it is the variations about an arbitrary level which are minimized.

The 18 coefficients to be estimated include three permanent magnetization coefficients, 6 induced magnetization coefficients, and 9 eddy current magnetization coefficients [41]. The complete model for the disturbance field, given by [41], is

$$B_{dist} = B_{perm} + B_{ind} + B_{eddy}, \quad (19)$$

where

$$B_{perm} = a_1 \cos X + a_2 \cos Y + a_3 \cos Z, \quad (20)$$

$$B_{ind} = B_t (a_4 + a_5 \cos X \cos Y + a_6 \cos X \cos Z + a_7 \cos^2 Y + a_8 \cos Y \cos Z + a_9 \cos^2 X), \quad (21)$$

where B_t is the total measured magnetic intensity (scalar value). The eddy current term is given by

$$\begin{aligned} B_{eddy} = B_t & (a_{10} \cos X \cos \dot{X} + a_{11} \cos X \cos \dot{Y} + a_{12} \cos X \cos \dot{Z} \\ & + a_{13} \cos Y \cos \dot{X} + a_{14} \cos Y \cos \dot{Y} + a_{15} \cos Y \cos \dot{Z} \\ & + a_{16} \cos Z \cos \dot{X} + a_{17} \cos Z \cos \dot{Y} + a_{18} \cos Z \cos \dot{Z}). \end{aligned} \quad (22)$$

The direction cosine terms are computed from the fluxgate sensor readings as

$$\begin{aligned} \cos X &= \frac{T}{B_t}, \\ \cos Y &= \frac{L}{B_t}, \\ \cos Z &= \frac{V}{B_t}, \end{aligned} \quad (23)$$

where T, L , and V are the components of the total field along the transverse, longitudinal, and vertical directions as measured by the flux-gate magnetometer. The derivatives are with respect to time. Geo-survey companies will often give a Figure of Merit (FOM) to describe the quality of the magnetic measurement compensation. They compute the FOM by flying a clover-leaf pattern which overflies the same point in space at four different headings while performing the same roll, pitch, and yaw maneuvers [52]. The FOM is computed by differencing the magnetometer readings for each pair of maneuvers, i.e. banking ± 5 degrees, and summing the differences. After removing the temporal variations as well as applying the compensation coefficient corrections, the FOM can often be less than 1 nano-Tesla.




Figure 36. Temporal Variations Recorded at Magnetic Base Stations Separated by 10's of Kilometers [52]

Temporal Variations.

Temporal variations, also known as diurnal variations, corrupt magnetic measurements when measuring the Earth's magnetic anomaly field. As described in the Magnetic Sources section, they originate primarily in the Earth's ionosphere and magnetosphere. Temporal variations change both spatially and temporally. The primary way to remove temporal variations in an aeromagnetic survey is with a magnetic base station. Any variations recorded in the stationary magnetic base station (located away from cultural effects) can be attributed to the temporal variations. These recorded variations can be subtracted from the aircraft measurements. Temporal variations are often consistent on a regional level [52]. Fig. 36 shows recorded temporal variations at four magnetic base stations separated by tens of kilometers. The low frequency components are very consistent on a regional level, whereas the high frequency components can vary between base stations. This indicates that using a magnetic base station to remove temporal variations will not fully remove high frequency variations [52]. High quality magnetic anomaly maps also go through processes known as tie-line leveling and micro-leveling to attempt to further remove the temporal variations. These procedures are described in Chapter 3.8: Creating Magnetic Anomaly

Maps. Removing temporal variations with a magnetic base station generally happens during processing of the magnetic data. In the context of a navigation system, it is not infeasible to transmit corrections in real-time; however, this may be undesirable. Furthermore, transmitting magnetic base station corrections would still not remove all temporal variations, especially in areas that do not have a nearby magnetic base station. A method to remove temporal variations in real-time is one of the major contributions of this research and is presented in Chapter 3.

Temporal variations originating from sources other than the ionosphere and the magnetosphere are also routinely removed from magnetometer data. Examples of these sources include lightning strikes or cultural anomalies. These are often removed “by hand” from a magnetic data set. It is clear that this type of error would also be difficult to remove in real-time.

2.8 Creating Magnetic Anomaly Maps

Magnetic anomaly navigation is a map-based navigation system. Consequently, a thorough understanding how magnetic anomaly maps are made is necessary for correctly designing a magnetic navigation system. This section presents the basic flight path considerations of a standard magnetic survey as well as the data processing which takes place to turn the measurements into a finished product—the magnetic anomaly map. This section follows from [39] and [52].

Flight Path.

There are important decisions to be made when planning a magnetic survey:




Figure 37. Basic Aeromagnetic Survey Lines

Location and Orientation of Survey Lines.

A standard magnetic survey is flown as a series of survey lines as well as tie lines. Fig. 37 shows the basic flight path for an aeromagnetic survey. Survey lines generally lie in the same orientation and are at an angle determined to give the best interpretive results given the specific geology of the survey area. This normally constitutes flying perpendicular to any prominent geological strike in the survey area [52]. Tie lines, or control lines, are flown perpendicular to the survey lines. Tie lines serve two main purposes. The first is to aid in interpolating the data. The tie lines give better observability of how the magnetic field changes in the direction perpendicular to the survey lines. The second purpose is to level the magnetic data. As stated

previously, the temporal variations can only be partially removed with a base station. The intersection of the tie lines and the survey lines give two measurements at the same physical location. After magnetic compensation, any difference between these two measurements can be attributed to the temporal variations. Tie-line leveling is explained in more detail in the next section.

Altitude and Drape Surface.

The altitude of a survey is primarily determined by the depth of the buried anomalies which are being measured, as well as the desired survey resolution. Flying low will lead to higher resolution surveys. This is because increasing altitude essentially acts like a low pass filter. Short wavelength anomalies can be missed if the survey altitude is too high. Oftentimes the surveys are flown at the lowest safe altitude. This usually necessitates flying a drape-surface. A drape surface is a varying altitude that loosely follows the terrain height. This means that most magnetic anomaly maps are not flown at a constant barometric altitude, but rather a constant height above terrain. This is required to keep terrain height changes from appearing as magnetic anomalies. A magnetic map which was created at a drape surface may need to be upward continued to a constant altitude for use in a magnetic navigation system. This type of upward continuation can be more complicated than the standard level-surface to level-surface upward continuation presented previously. One technique is to perform many upward continuations at discrete levels thereby creating a stack of uneven map surfaces. This stack of uneven map surfaces can then be used to interpolate to one constant altitude.

Line Spacing and Sensor Sampling Rate.

The line spacing is the spacing between the survey lines. The line spacing is driven by the desired resolution, as well as the survey altitude. The survey altitude determines the upper limit on the frequency content of measured magnetic field. Magnetic fields are said to be Laplacian, which ensures that they are smooth and devoid of abrupt changes and discontinuities [52]. Flying at a line spacing equal to the survey height is sufficient to fully sample the field with respect to the Nyquist criteria [53]. While flying at a line spacing equal to the altitude is ideal, cost constraints normally require a larger line spacing. For magnetic navigation, it is important to know whether a map being used for navigation is fully sampled, or close to fully sampled. A fully sampled map ensures correct reconstruction of the signal when interpolating. The ability to fully reconstruct the signal from a gridded map is a feature of magnetic navigation not seen in similar navigation systems such as terrain height navigation.

Magnetic anomaly maps are not sampled uniformly. Modern magnetic sensors are capable of sensor readings of 10 Hertz or more. At an aircraft velocity of 50 meters/second, this would give samples along the flight lines separated by five meters. If flying at an altitude of 100 meters and attempting to fully sample the field, the survey line spacing would be 100 meters. This indicates that the map is sampled every 5 meters in the survey line direction, but sampled only every 100 meters in the perpendicular direction. Depending on the line spacing and the sampling rate, it should be understood that a magnetic anomaly map is potentially under-sampled in one direction and fully sampled in the other direction. Sometimes an anti-aliasing filter will be applied to magnetic survey data to prevent errors from under-sampling a magnetic anomaly map.

Data Processing.

Once a dataset has been recorded from an aeromagnetic survey, the raw data must be turned into a finished map product. Understanding what is done to the raw data to create the magnetic anomaly map is essential for magnetic navigation. When an aircraft is navigating with a magnetometer it is recording raw data. This raw data must be related to a processed map in real time. It is clear that this comparison will be imperfect, especially with respect to types of processing which could never take place in real time, such as a human operator noticing and then removing a blip in the data caused by a lightning strike. The basic steps for creating a finished magnetic anomaly map were obtained from [39], and are given below.

1. Verifying and editing raw data

The first step in processing the data is to view the data for any irregularities caused by instrument error, cultural anomalies, or natural events such as lightning strikes. Removing these types of errors may be done “by-hand” in specific situations like an instrument failure. In cases where there exists measurement noise that is clearly not caused by magnetic anomalies in the ground, low pass filtering may be used.

2. Geo-locating data

Once the magnetic measurements and GPS data have been examined and determined to be reasonable, the magnetic measurements, which are often GPS time-stamped, can be geo-located.

3. Lever-arm correction

A lever-arm often exists between the magnetometer instruments and the GPS. A lever arm correction must be applied to spatially align the magnetometer measurements with the recorded GPS locations. A standard lever-arm correction

can be applied if aircraft attitude is available. If unavailable, the GPS position may be interpolated using the aircraft velocity to adjust the GPS position measurements.

4. Temporal variation corrections

Temporal variations are recorded at a base station and time-stamped with GPS time. The base station measurements are aligned in time with the aircraft measurements and then subtracted from the aircraft measurements. This removes the majority of the temporal variations, especially longer wavelength variations. Base station measurements are considered representative of the survey area to a distance of about 50 kilometers [39]. It is important to note that this process will never fully remove the temporal variations. Tie-line leveling, described below, attempts to remove the remaining temporal variations. Subtracting the temporal variations in this way will also add an arbitrary DC shift in the data, which is not considered an issue as magnetic anomaly maps do not attempt to represent the absolute value of the field [39].

5. Removing the Earth's main field

Magnetic anomaly maps are primarily concerned with the short wavelength features of the measured magnetic field caused by crustal sources. Consequently, the main Earth field is routinely subtracted from the measurements using a model such as the IGRF. The IGRF changes slowly enough that position does not need to be known to high accuracy to successfully subtract out the field. A magnetic navigation system with a kilometer of error could still easily look-up and subtract out with IGRF to a high degree of accuracy.

6. Tie-line leveling

Magnetic surveys are flown in one principle direction, with perpendicular tie-

lines flown to create a series of intersecting points. Any difference in the measured magnetic field at these intersections is an error. Tie-line leveling seeks to systematically reduce the error of these intersections by adjusting the magnetic data. Various methods exist to accomplish tie-line leveling, and are not discussed in detail here. Tie-line leveling is often a very manual process and requires a skilled operator [39].

7. Micro-leveling

Micro-leveling is the final adjustment to the raw sensor data. Micro-leveling methods are often proprietary but generally involve some kind of filtering [39]. The primary purpose of micro-leveling is to remove any apparent errors in the visualized map product. Oftentimes this process is called de-corrugation because of the appearance of corrugated lines in a magnetic anomaly map which are clearly in error.

8. Gridding the data

The goal of gridding the magnetic measurement data is to create a smooth uniform grid of magnetic data which honors the original measured values. Various types of interpolation are used to grid the data. Linear interpolation is often sufficient in the survey-line direction while cubic interpolation is recommended in the tie-line direction [52]. Because the map data is sampled so much higher in the survey line-direction, care must be taken not to introduce aliasing when down-sampling to the grid spacing [39]. Oftentimes an anti-aliasing low-pass filter may be used to remove aliasing effects [52].

2.9 Rao-Blackwellized Particle Filtering

Now that we have described the magnetic field and how it is measured and compensated, we will shift our attention to estimation methods that are useful for magnetic field navigation. The marginalized particle filter (MPF), also known as the Rao-Blackwellized particle filter, is a useful extension to the basic particle filter for high dimension non-linear filtering problems [57]. As the number of states in a filtering problem increases, the number of particles needed to accurately represent the probability distribution for each state increases exponentially. Fully modeling an inertial navigation system can easily exceed the feasible number of states a basic particle filter can handle. A basic inertial navigation system model may have position, velocity, and attitude states, giving a 9-dimensional filtering problem. More complex models can have over 20 states with the addition of altitude aiding, sensor bias, and sensor scale factor states. The MPF addresses the problem of high dimension filtering problems by marginalizing out the states appearing linearly in the dynamics [57]. Particles are only needed to represent the non-linear states, with each particle having an associated Kalman Filter to represent the remaining states. This section aims to provide the basic equations needed to implement an MPF in code. The following section follows from [57], which should be referenced for a more complete discussion of the MPF as well as the full derivation.

Assume a general non-linear filtering problem with a discrete state-space model and a set of measurements which may be related to the filter states. The filter aims to estimate the posterior probability density function after an observed measurement.

The non-linear system may be expressed as

$$\begin{aligned}x_{t+1} &= f(x_t, w_t), \\ y_t &= h(x_t, e_t).\end{aligned}\tag{24}$$

where x_t is the state variable at time t , and y_t is the measurement at time t . The time x_{t+1} represents the state variable propagated forward one discrete time step. The terms f and h are non-linear functions representing the state dynamics and measurement model, respectively. The w_t and e_t terms are the dynamics and measurement noises, respectively. Now assume that some of the states are non-linear and some of the states are linear. We wish to partition the states into a group of linear states x_t^n and a group of non-linear states x_t^l :

$$x_t = \begin{bmatrix} x_t^n \\ x_t^l \end{bmatrix}.\tag{25}$$

Now the linear and non-linear states are expressed as a sum of a non-linear part as well as a linear part:

$$\begin{aligned}x_{t+1}^n &= f_t^n(x_t^n) + A_t^n(x_t^n)x_t^l + w_t^n, \\ x_{t+1}^l &= f_t^l(x_t^n) + A_t^l(x_t^n)x_t^l + w_t^l, \\ y_t &= h_t(x_t^n) + C_t(x_t^n)x_t^l + e_t,\end{aligned}\tag{26}$$

where x_t^n represents the non-linear states and x_t^l represents the conditionally linear states. A state being “linear” simply means its posterior PDF can be expressed as a Gaussian random variable—it may still have non-linear dynamics, although conditioned on the non-linear states the dynamics become linear. Consequently, both the non-linear and linear states potentially propagate forward in time based on some function f^n or f^l , of the non-linear states. They may also propagate forward in time based on some function of the linear states—represented by the matrices A^n and A^l .

These matrices may also be conditioned on the non-linear states x^n . Finally, the measurement may be both a function of the non-linear states, represented by h_t , as well as a function of the linear states, represented by the matrix C_t .

Clearly identifying the above functions and matrices in the system model is essential to transition to the main equations for the MPF. It is worth noting that a common special case arises with many navigation-related filtering problems, where the state dynamics are linear and only the measurement equation has non-linearities. In this case the functions f^n and f^l become simple matrices and the model may be rewritten as

$$\begin{aligned}x_{t+1}^n &= A_{n,t}^n(x_t^n) + A_{l,t}^n(x_t^n)x_t^l + w_t^n, \\x_{t+1}^l &= A_{l,t}^n(x_t^n) + A_{l,t}^l(x_t^n)x_t^l + w_t^l, \\y_t &= h_t(x_t^n) + C_t(x_t^n)x_t^l + e_t.\end{aligned}\tag{27}$$

The steps for the MPF are

1. Initialization
2. Calculate and normalize importance weights
3. Apply Kalman filter measurement update
4. Apply particle filter time update
5. Apply Kalman filter time update
6. Iterate from step 2.

In the section below each of the above steps are outlined in detail to a level in which they can be implemented in code. We simplify the presentation by assuming the initial noise matrix for the system dynamics is diagonal with no cross-covariance terms. The

Kalman filter time update equations are more complex when this assumption is not met and [57] should be referenced if the more complete equation set is needed.

1. Initialization:

The MPF is initialized similar to the standard particle filter. Recall each particle has an associated Kalman filter. For each particle the non-linear states are initialized with particles drawn from an initial distribution and the linear states are given an initial covariance P_0 . For $i = 1, \dots, N$ initialize the particles, $x_{0|-1}^{n,(i)} \sim p_{x_0^n}(x_0^n)$ and set $\{x_{0|-1}^{l,(i)}, P_{0|-1}^{(i)}\} = \{\bar{x}_0^l, \bar{P}_0\}$. Frequently the particles will be drawn from a Gaussian distribution with a given mean and covariance.

2. Calculate and normalize the importance weights:

The recursive part of the algorithm begins by assigning importance weights to each particle based on the current observation. The importance weight for each particle is determined by calculating the probability of obtaining the observed measurement if the measurements are Gaussian distributed with a mean μ given by

$$\mu = h_t(x_t^{n,(i)}) + C_t \hat{x}_{t|t-1}^{l,(i)}, \quad (28)$$

and a covariance M given by

$$M = C_t P_{t|t-1} C_t^T + R_t. \quad (29)$$

Using the observation y_t we can denote the residual $e = y_t - \mu$. We can evaluate the probability from the Gaussian distribution for the i th particle with

$$q_t^{(i)} = \exp\left(-\frac{1}{2}(e M^{-1} e^T)\right). \quad (30)$$

The weights then must be normalized according to

$$\tilde{q}_t^{(i)} = \frac{q_t^{(i)}}{\sum_{j=1}^N q_t^{(j)}}. \quad (31)$$

3. Resampling:

At this point in the algorithm the particles may be resampled according to whichever resampling strategy is chosen. See the Particle Filtering section for more information on resampling strategies.

4. Kalman Filter Measurement Update: The next step is to update the linear states of the Kalman Filter with the measurement. At this point we will drop the (i) superscripts—it should be understood that a Kalman Filter is associated with each particle and calculations may be different for each particle. We first define the Kalman gain matrix

$$K_t = P_{t|t-1} C_t^T M_t^{-1}. \quad (32)$$

We next calculate a measurement residual and apply the Kalman gain for each particle to determine the Kalman filter updated linear state:

$$\hat{x}_{t|t}^l = \hat{x}_{t|t-1}^l + K \left(y_t - h_t(x_t^{n,(i)}) - C_t \hat{x}_{t|t-1}^{l,(i)} \right). \quad (33)$$

We next calculate the updated Kalman filter covariance for each particle according to

$$P_{t|t} = P_{t|t-1} - K_t M_t K_t^T. \quad (34)$$

5. Particle Filter Time Update:

The particles are propagated in time according the system dynamics model. Each individual particle also has random noise added to the appropriate states. When the non-linear states are conditioned on their value at time t , their PDF's are conditionally Gaussian and given by

$$p(x_{t+1}^n | X_t^n, Y_t) = \mathcal{N}\left(f_t^n + A_t^n \hat{x}_{t|t}^l, A_t^n P_{t|t} (A_t^n)^T + Q_t^n\right). \quad (35)$$

Implementing this in code isn't directly obvious from the expression because each particle must have added noise that is likely not a diagonal matrix due to the $A_t^n P_{t|t} (A_t^n)^T + Q_t^n$ term. In a MATLAB like high-level programming language the *randn()* or similar method is often used to add noise to a model assuming a known noise covariance σ^2 for a given state. The noise is then simply added to a particle state by calling $\sigma \cdot \text{randn}()$. This simplification is only possible when the noise matrix Q is diagonal. The method needed here is to compute the Cholesky decomposition of $A_t^n P_{t|t} (A_t^n)^T + Q_t^n$ when adding the noise component of the dynamics model. For MATLAB this would give

$$x_{t+1}^{n,(i)} = f_t^n(x_{t|t}^{n,(i)}) + A_t^n x_{t|t}^{l,(i)} + \text{chol}\left(A_t^n P_{t|t}^{(i)} (A_t^n)^T + Q_t^n\right) \text{randn}(m, 1), \quad (36)$$

where *chol* is MATLAB's Cholesky decomposition function and m is the number of non-linear states.

6. Kalman filter time update

The final step is to propagate the linear Kalman filter states forward in time. This includes propagating their means as well as covariances. First define a few intermediate matrices (again, these can be different for each particle):

$$N_t = A_t^n P_{t|t} (A_t^n)^T + Q_t^n, \quad (37)$$

$$L_t = A_t^l P_{t|t} (A_t^n)^{-1}, \quad (38)$$

$$z_t = x_{t+1|t}^n - f_t^n(x_{t|t}^n). \quad (39)$$

The Kalman filter time update for the state means is then given by

$$\hat{x}_{t+1|t}^l = A_t^l \hat{x}_{t|t}^l + f_t^l(x_{t|t}^n) + L_t(z_t - A_t^n \hat{x}_{t|t}^l). \quad (40)$$

Finally, the covariance is propagated forward in time according to

$$P_{t+1|t} = A_t^l P_{t|t} (A_t^l)^T + Q_t^l - L_t N_t L_t^T. \quad (41)$$

At each time step the particles may be used to determine the expected mean and covariance of the linear states. The expected mean is given by

$$\hat{x}_{t|t}^l = \sum_{i=1}^N \hat{q}_t^{(i)} \hat{x}_{t|t}^{l,(i)}, \quad (42)$$

and the expected covariance of the linear states is given by

$$\hat{P}_{t|t}^l = \sum_{i=1}^N \hat{q}_t^{(i)} \left(P_{t|t}^{(i)} + (\hat{x}_{t|t}^{l,(i)} - \hat{x}_{t|t}^l)(\hat{x}_{t|t}^{l,(i)} - \hat{x}_{t|t}^l)^T \right). \quad (43)$$

The expected means of the non-linear states is given by

$$\hat{x}_{t|t}^n = \sum_{i=1}^N \hat{q}_t^{(i)} \hat{x}_{t|t}^{n,(i)}, \quad (44)$$

and the expected covariance of the non-linear states is given by

$$\hat{P}_{t|t}^n = \sum_{i=1}^N \hat{q}_t^{(i)} \left((\hat{x}_{t|t}^{n,(i)} - \hat{x}_{t|t}^n)(\hat{x}_{t|t}^{n,(i)} - \hat{x}_{t|t}^n)^T \right). \quad (45)$$

2.10 Background Conclusion

This chapter has provided a geophysics focused background on the magnetic anomaly field. The intent of the chapter was to allow a person with limited geophysics knowledge to understand the subtleties of magnetic anomaly navigation. First, each of the major components of the total measured magnetic field was discussed. Magnetic anomaly navigation uses a subset of the total measured field to navigate, so understanding the corrupting components of a measurement is necessary. The definition of a magnetic anomaly was then clearly defined as well as methods to perform transformations on magnetic anomaly maps. These transformations would be necessary to implement in a magnetic anomaly navigation filter. The current availability of magnetic anomaly maps was discussed along with the methods used to create these maps. An understanding of how magnetic anomaly maps are created is essential to understanding how to relate a total field measurement back to a magnetic anomaly map. Different types of magnetic sensors were also discussed along with how vector sensors are used to provide aircraft field calibration. Removal of the aircraft field is especially important when magnetic anomaly navigation is being implemented on a real air force aircraft without the clean environment of a boom. Finally a brief discussion of Rao-Blackwellized particle filtering was discussed.

III. Filter Design

This chapter details the design of the navigation filter for magnetic anomaly navigation. The first sections detail the analysis of the temporal variations which corrupt the magnetic anomaly measurements. The temporal variations will be estimated and removed by the navigation filter and consequently a model must be developed which properly describes them. Magnetometer data from magnetic observatories is used to gain insight into modeling the temporal variations. Analysis indicates that a First Order Gauss Markov (FOGM) process is a useful way to model the temporal variations. Choosing the parameters of the FOGM process is discussed. Next, the observability of the temporal variations is discussed. The navigation filter will attempt to simultaneously estimate the aircraft's position as well as the temporal variations, and this requires observability of the temporal variations. Finally, the full navigation filter design is presented. The filter is a 13 state Rao-Blackwellized particle filter which primarily models the errors in an INS, as well as the temporal variations.

3.1 Temporal Variation Modeling

This section details the modeling of the temporal variations which corrupt a measurement of the magnetic anomaly field. This modeling was accomplished using real temporal variation data recorded from magnetometer base stations.

Types of Temporal Variations.

There are three main types of temporal variations. The first type is the diurnal variations. The diurnal variations have a period of 24 hours and are caused by the solar quiet currents in the ionosphere as described in Chapter 2. The diurnal variations are generally smooth and vary by around 20 nano-Teslas within a single




Figure 38. Example Diurnal Variations [58]

day. At mid latitudes the diurnal variations drop and at equatorial latitudes the diurnal variations rise. Fig. 38 shows an example of the diurnal variations. The next type of temporal variation is called a micro-pulsation. Micro-pulsations have shorter wavelengths and shorter amplitudes than the diurnal variations. Unlike the diurnal variations, they are not predictable. Fig. 39 shows an example of micro-pulsations. The final type of temporal variation is a magnetic storms. Magnetic storms are caused by the interaction of the solar wind with the Earth's magnetosphere. The intensity




Figure 39. Example Micro-Pulsations [58]




Figure 40. Example Magnetic Storms [58]

of magnetic storms can vary widely. The overall intensity of a magnetic storm can be predicted but the actual variations seen on a magnetometer are very unpredictable. Fig. 40 shows an example magnetic storm.

Characterizing Temporal Variations as a Random Process.

We wish to model the temporal variations which will be corrupting our measurements of the magnetic anomaly field. We will attempt to model them as a random process. To accomplish this we start by examining a year's worth of data from a magnetic observatory. Magnetic observatory data is published daily at observatories all over the world. We downloaded data from a magnetic observatory in Boulder, CO. The raw data is shown in Fig. 41. It is difficult to draw any conclusions from this data. We know, however, that the magnetic field undergoes secular variations over time. The secular variations are a slow drift in the Earth's core field. We can compute the average monthly variations to estimate the magnitude of the secular variations. Fig. 42 shows the monthly averages for the temporal variations. It can be seen from the plot that the secular variations over one year are larger than the daily variations. The secular variations that would be observed during a single flight, however, are definitely much smaller than the daily variations. We now wish to examine the daily variations. We know that the temporal variations have a strong 24 hour periodic




Figure 41. Daily Recorded Temporal Variations at Boulder CO Observatory in 2010




Figure 42. Montly Average Temporal Variations Showing Secular Variations




Figure 43. Zero-Mean Daily Variations to Emphasize Hourly Variations

component called the diurnal variation. This is caused by the rotation of the Earth, as explained in Chapter 2. We can simply plot the raw data from Fig. 41 with the data from each day made to be zero-mean. The result of this operation is shown in Fig. 43. It is apparent there is a large periodic swing in the data, as expected. During the duration of a flight an aircraft could certainly encounter the daily variations seen here. Finally, we wish to examine the hourly variations caused by micro-pulsations. We can examine the micro-pulsations by removing the secular and daily variations from the data. To accomplish this we utilize a simple technique used by magnetic observatories. A baseline daily variation can be determined by taking five “quiet” days of a month and averaging them together. We do not wish to simply average the entire month, because stormy days can be very different than the quiet day average. An




Figure 44. 5 Quiet Days From Single Month to Determine Hourly Average

example of selecting five quiet days and averaging them together is shown in Fig. 44. We computed a daily average temporal variation for each month of our dataset from five quiet days during the month. We then subtracted these monthly averages from the raw data. This subtraction not only removes the diurnal, or daily variations, but also the secular variations. The result of this operation is shown in Fig. 45. At this point the data appears quite random, without any obvious trends. It appears to resemble a moving bias, which can be modeled as a First-Order Gauss-Markov (FOGM) process. Removing the daily variations is a necessary condition to apply this model. A FOGM process is by definition zero-mean. If the daily variations were not removed there would be a large periodic swing in the data which would not fit a FOGM model well. A FOGM model is characterized by two parameters—a variance




Figure 45. Year of Data with Hourly Averages Determined for Each Month Removed

and a time constant. We can get a rough estimate of the variance simply by observing Fig. 45. The standard deviation of the data is about 10 nano-Teslas, giving a FOGM variance of 100 nano-Teslas. In general, the most appropriate way to estimate the variance is likely to just use the variance of the data from the previous few hours. If a navigation filter had access to base station data while it was being initialized, this would be easy to compute. If an aircraft was flying with GPS and a magnetic anomaly map, it would also have observability of the temporal variations and could compute this parameter. If it did not have access to base station data or GPS, it could use a hard-coded number for the variance. The second parameter needed to characterize a FOGM model is the time constant. The time constant describes how long the data takes to de-correlate with itself. The time constant can be determined by observing when the signal's auto-correlation decreases to 36.8% of its starting value. To estimate the time constant we took a month of raw base station data and removed the secular variation. This is necessary to obtain a valid auto-correlation. The month of data as well as the secular variation are shown in Fig. 46. A normalized auto-correlation of this data is shown in Fig. 47. As expected, there are peaks at both 12 and 24 hours. Because we plan to remove these longer periodic trends, we are only concerned with how long it takes the data to decorrelate with itself for small time lags. Fig. 47 shows that the signal decorrelates with itself down to 36.8% after approximately 140 minutes.

In summary, we are choosing to model the temporal variations in the navigation filter as a first-order Gauss-Markov process. We have shown that this model is applicable when an average daily variation is removed. This daily variation is computed by simply averaging a few quiet days over the past month. We need to determine both a variance and time constant for the FOGM model. We can use hard-coded




Figure 46. Removing the Secular Variation From 1 Month of Data in Order to Perform Valid Auto-Correlation




Figure 47. Estimating the FOGM time constant




Figure 48. Verifying the FOGM Model

numbers such as 100 nano-Teslas for the variance and 140 minutes for the time constant for a approximate fit to the data. For a more accurate fit, we can compute these parameters from the previous few hours of recorded temporal variations. The flight test results shown later in Chapter V show little sensitivity to these parameters with respect to filter performance. Fig. 48 shows three true daily temporal variations as well as three generated daily temporal variations using the derived model.

Temporal Variation Variables.

The raw data for generating the temporal variation model came from one base station in Colorado. It is important to have an understanding of how the temporal variations change spatially as well as how they change with respect to other variables such as magnetic storms. Fig. 49 shows recorded temporal variations at 6 different mid and equatorial latitudes. The temporal variations are from the same 24 hour period but are shifted to be aligned in local time. The most noticeable aspect of the plot is that French Guiana appears to have its noon-day swing in the opposite




Figure 49. Time-Aligned Temporal Variations at Mid and Equatorial Latitudes




Figure 50. Time-Aligned Temporal Variations at Polar Latitudes

direction. This is expected at equatorial latitudes, and is due to the EEJ described in Chapter 2. While there is a fair deal of variation in the plot, from a random-process standpoint the temporal variations appear consistent among the different mid and equatorial latitudes. Contrast this consistency with Fig. 50, which shows polar latitudes. These temporal variations are from the same 24 hour period as Fig. 49. It is clear that these temporal variations are much less smooth. It does not appear appropriate to use the same random process model at polar latitudes when modeling the temporal variations. The greater variation in the polar latitudes is expected, and due to several factors explained in Chapter 2.

Finally, we wish to explore what the temporal variations look like during high




Figure 51. Temporal Variations Under Various Magnetic Storm Conditions

magnetic storm conditions. Fig. 51 shows recorded temporal variations at five different Disturbance Storm Time (DST) indices at Boulder, CO. The DST index is a measure of the magnetosphere ring current and is often used to characterize magnetic storm conditions. Higher absolute values correspond with stronger magnetic storm conditions. The magenta curve corresponds with one of the strongest magnetic storms on record. For very strong magnetic storms it is clear that the parameters for the FOGM process would need to be modified. Hard-coded values based on an average may not suffice. This is likely best answered empirically, and any filter framework using magnetic measurements should be tested during magnetic storm conditions.

3.2 Temporal Variation Observability Analysis

In the previous section we developed a model for the temporal variations to be used in a navigation filter. Even a perfectly modeled state does not guarantee observability of that state in a given filter. We wish to explore the observability of the temporal variations. There are straightforward methods to determine observability of linear systems; however, a map-based navigation system is highly non-linear. Non-linear methods for determining observability do exist, but can be quite complex. We will utilize a more ad hoc approach. Assume we have magnetometer measurements in which we have already removed the core field and applied aircraft compensation. Our remaining measurement is approximately the sum of the Earth's crustal field as well as the temporal variations. On one extreme, if these two components existed at identical time-frequencies in our measurements it is clear there would be no way to separate the two components. On the other extreme, if the two components were at completely different frequencies which were spread far enough apart on the spectrum, a simple filtering operation could resolve the components. We wish to analyze the frequency components of both the temporal and crustal frequencies in order to get a general idea of signal separability. If the signals have limited frequency overlap we would expect some level of temporal variation observability in our filter model.

Temporal Variation Frequencies.

The temporal variations exist as a multi-dimensional signal, changing over time as well as space. When discussing the frequency of the temporal variations, we need to account for the fact that the temporal variations have both a spatial-frequency as well as a time-frequency. A magnetic base station observes only the time-frequencies of the temporal variations at a single location. A real aircraft moving through space will have measurements which capture both of these frequency components.




Figure 52. Amplitude vs. Time-Frequency of the Temporal Variations [42]

We first consider the time-frequencies of the temporal variations. Fig. 52 shows a plot of the amplitude of the temporal variations vs frequency. There is a general trend of decreasing amplitude as frequency increases. It is important to note that the approximate accuracy of a survey aircraft to measure the magnetic anomaly field is 1 nano-Tesla. This indicates that magnetic anomaly maps have an approximate uncertainty of 1 nano-Tesla. Therefore, temporal variations with a frequency of around 0.02 Hertz or higher have amplitudes which are already smaller than the accuracy of magnetic anomaly maps.




Figure 53. Temporal Variation Spatial Frequencies [13]




Figure 54. Temporal Variations at Magnetic Base Stations 10's of km Apart [52]

We next wish to determine the spatial frequency content of the temporal variations. Fig. 53 shows the spatial frequencies of various components of the temporal variations. It is clear from the plot that the largest spatial-period of any of the components is approximately 100 kilometers. Fig. 54 shows the recorded temporal variations during the same time interval at magnetic observatories which are separated tens of kilometers. It is clear from the plot that the low frequency components of the signal are all quite similar. The high frequency components seem to generally line up as well; however, there is certainly variation. It is apparent that the

higher time-frequencies manifest themselves as higher spatial-frequencies as well. This is beneficial because we know from Fig. 52 that higher time-frequency components have smaller amplitudes. Given the large separation of the magnetic observatories in Fig. 54, the consistency observed between base stations, and the approximate speeds at which aircraft fly, it appears reasonable to conclude that the time-frequencies are the dominant temporal variation component which would be observed in an aircraft magnetometer.

Anomaly Field Frequencies.

When flying over a spatial map at a given velocity, the spatial frequencies from the anomaly field are manifested as time-frequencies in the magnetometer measurements. The magnitude of these time-frequencies is clearly related to aircraft velocity. If an aircraft flies over the crustal field at twice the velocity, the measurements in the magnetometer will have double the frequency. The frequency content of the crustal field is well understood in the geo-physical survey industry. When flying a aeromagnetic survey, the line spacing must be such to fully sample the field, otherwise aliasing in the map can occur. Reid showed that the aliased power expected from an aeromagnetic survey at height h and line spacing Δx is an exponential function [53], given by

$$F_{\text{aliased}} = e^{\frac{-2\pi h}{\Delta x}}. \quad (46)$$

Table 1 shows the aliased power given by Equation 46 for several ratios of height and sample spacing. As shown in Table 1, when sampling the magnetic field during an aeromagnetic survey, nearly 100 percent of the signal is captured if the sample spacing is one half that of the height. This fact, along with the Nyquist frequency theorem, indicates the shortest expected wavelength at a given height h is itself h . We can use this fact to estimate what the max time-frequency would be in an

Table 1. Height/Line Spacing Ratio vs. Aliased Power [53]

Height / Line Spacing	Percent Aliased Power
0.25	21%
0.5	4.3%
1	0.19%
2	0.0003%

Table 2. Max Crustal Field Temporal Frequencies When Sampled From an Aircraft

	50 m/s	100 m/s	500 m/s
500 m	0.1 Hz	0.2 Hz	1 Hz
1 km	0.05 Hz	0.1 Hz	0.5 Hz
5 km	0.01 Hz	0.02 Hz	0.1 Hz
10 km	0.005 Hz	0.01 Hz	0.05 Hz

aircraft’s magnetometer measurements as it flew over the crustal field at height h and velocity V_t .

$$F_{time}^{max} = \frac{1}{h} V_t \quad (47)$$

We can use this equation to estimate the crustal field time-frequencies which would appear in our magnetometer measurements at several various aircraft altitudes and velocities. As seen in Table 2, the max frequencies range from 1 Hertz when flying low and fast, to 0.005 Hertz when flying high and slow. Comparing these time-frequencies expected when flying over the anomaly field to Fig. 2 indicates that the temporal frequencies of the external magnetic sources do overlap with the temporal frequencies of the anomaly field seen at aircraft altitudes and velocities. Fig. 55 shows the temporal variation amplitude vs. frequency plot shaded to coincide with the anomaly frequencies seen in an aircraft magnetometer at normal altitudes and velocities. The green regions are frequencies a filter could “know” are not possible from the Earth’s crustal field. These frequencies should be observable. The red shaded region shows frequencies in which the temporal variations and anomaly field time-frequencies are overlapped. These frequencies should not be observable. The orange




Figure 55. Comparison of Temporal Variation and Crustal Field Frequencies at Aircraft Velocities and Altitudes [42]

shaded region shows where observability is not possible, but the temporal variations are less than the accuracy of the magnetic anomaly map. Although the signal is being corrupted at these frequencies, the absolute accuracy of the measurements is relatively unchanged. This analysis of spatial and temporal frequencies of the anomaly field compared to other magnetic sources indicates that separating these sources while flying at aircraft velocities and speeds should be possible to some extent.

3.3 Measuring the Magnetic Anomaly

Defining the measurement equation for magnetic anomaly navigation is an important first step in designing an estimation filter. The measurement equation describes how the raw measurements from a magnetometer relate back to the filter states. An intermediate goal is to determine how the raw measurements relate back to a magnetic anomaly map. The individual errors present in the raw magnetometer measurement

are presented below.

Map Quality.

Map quality is a significant factor in magnetic anomaly navigation. Many magnetic anomaly maps were created before the use of GPS. The magnetometer data therefore may not be geo-located as accurately as modern magnetic surveys. A map created without the use of GPS will possibly lead to degraded navigation performance. If using a magnetic anomaly map that was made without the use of GPS, a navigation filter would likely need to have a greater measurement error covariance. In general, it is suspected that the highest accuracy navigation will be obtained with recent magnetic anomaly maps. Map resolution is also an important concern. As stated previously, a magnetic survey must be flown with sample spacing no greater than the height above the terrain to fully sample the field. Not all magnetic anomaly maps meet this criteria. These maps will therefore have un-captured high frequency components which will lead to degraded navigation performance. Finally, many magnetic anomaly maps are created at a “drape” altitude. These maps are flown at a constant height above terrain rather than a constant barometric height. Fortunately, these maps can be upward continued to a constant barometric height. This process requires the original drape altitude which was flown, however, and this may not always be available. Treating a draped map as a constant elevation map will lead to decreased navigation performance.

Altitude Dependent Variations.

Navigation accuracy is dependent on altitude. The Earth’s magnetic anomaly field exists in three dimensions. As altitude increases, the spatial-frequency content of the signal decreases. This decreased spatial-variation in the field degrades navigation

performance. Navigation at aircraft altitudes has the benefit of being far closer to the magnetic sources than navigation at satellite altitudes, and consequently can achieve much better accuracy. Magnetic anomaly maps are only flown at one altitude—usually very low. As described in Chapter 2, we can upward continue this data. This process is error prone, however, especially with a small map, due to edge effects. The need for upward continuation can be ignored if flying at the survey altitude; however, this is certainly not feasible for all or even most aircraft missions.

Corrupting Sources.

A magnetometer mounted on an aircraft measures the total magnetic field. We are attempting to use a distinct component of this field for navigation, the crustal field. The other three components of the total field can be thought of as measurement errors, or corrupting sources. The total measurement includes four main components.

1. Main Earth Field
2. Aircraft Field
3. Temporal Variations
4. Crustal Field

The measurement equation must include all four of these terms. The main earth field is the easiest of the three corrupting fields to remove. The main earth field is well modeled by the IGRF, a freely available model which is easy to evaluate for a given latitude, longitude, altitude, and time [27]. The IGRF field changes so slowly with respect to location that only a rough position estimate must be known to subtract out this field. The aircraft field is mostly removable. Modern aircraft compensation systems can remove the aircraft field to fractions of a nano-Tesla [27]. After

the removal of the main Earth field using the IGRF and applying aircraft field compensation, the temporal variations are the largest remaining error. If unaccounted for, temporal variations can affect navigation performance. The effect of temporal variations can be mitigated, however. In [14], the observability of real-world temporal variations by a magnetic anomaly navigation filter was shown in a simulation environment. When the filter estimated the temporal variations, the strength of the magnetic storm conditions did not have a large effect on navigation performance. Transmitting the temporal variations from a nearby base station is a feasible but potentially undesirable way to remove the temporal variations.

Measurement Equation.

In light of the given error sources, we can present the full measurement equation for magnetic anomaly navigation. First, we present the errors in the 3D map function to be used by the navigation system:

$$M_3(lat, lon, h) = f_I^3(f_U(M_{h_0} + \delta M_{h_0}) + \delta U), \quad (48)$$

where:

M_{h_0} is the two dimensional grid of magnetic intensity values at height h_0

δM_{h_0} are the errors in the original map grid

f_U is the upward continuation function which transforms M_{h_0} to several discrete increasing altitudes, giving a three dimensional grid of values from the original two dimensional grid

δU is the error in the upward continuation transform

f_I^3 is an operation which returns a three dimensional interpolation function given a three dimensional grid of values

lat, lon, h are the latitude, longitude, and height, respectively, at which the magnetic

intensity is being evaluated

$M_3(lat, lon, h)$ is a 3D interpolation function which returns expected magnetic intensity at a given latitude, longitude, and altitude

The magnetic anomaly measurement equation for a *post-compensated* measurement can be given by

$$Z_t = M_3(lat, lon, h) + I(lat, lon, alt, t) + \delta I(lat, lon, alt, t) + \delta C(\theta, \phi, \psi) + V(lat, lon, h, t) + H + b + w \quad (49)$$

where

Z_t is the raw measurement from the magnetometer at time t

$M_3(lat, lon, h)$ is the pre-computed 3D interpolation function which returns expected magnetic intensity

I is the IGRF model, which depends on both position and time

δI is the error in the IGRF model, which depends on both position and time

θ, ϕ, ψ are aircraft Euler angles

δC is the residual aircraft field after compensation as a function of attitude (given by Euler angles)

V is the temporal variation, which is a function of both position and time

$H(\theta, \phi, \psi)$ is the magnetometer heading error, which is a function of aircraft attitude [11]

b is the time-correlated magnetometer bias

w is the magnetometer white noise

This measurement equation is intentionally as general as possible—simplifications can likely be made depending on the specific circumstances of the flight. We will

end this section by presenting the ideal situation for a magnetic anomaly navigation system. Assume a high-quality magnetic anomaly map exists over the area in which we want to fly. Assume we wish to navigate over this map at the survey altitude h_0 . These two assumptions allow us to avoid using upward continuation. We therefore create an interpolation function directly out of the two dimensional grid and call it M_2 . Furthermore, assume we are flying in an aircraft ideally suited for magnetic anomaly navigation, such as a geo-survey aircraft. The compensation systems of these aircraft routinely remove the aircraft field to a fraction of a nano-Tesla, so we will drop the δC term. The δI term, which represents the errors in the IGRF field, will be assumed constant, and will be grouped in with the temporal variations, V . This is a reasonable assumption, as the core field is a long wave-length model. Finally, assume we group the heading error H , sensor bias b , and white noise w , into a single white noise source $\hat{w}$ with a noise strength of q . Under the ideal case presented, our measurement equation simplifies to

$$Z_t = M_2(lat, lon) + I(lat, lon, alt, t) + V(lat, lon, h, t) + \hat{w} \quad (50)$$

We wish to relate the measurement to an aircraft's latitude and longitude. Removing the I term is trivial—the IGRF model is simply evaluated at the estimated position. As stated previously, this is a long wave length model and an approximate location will suffice (any errors will likely be constant offsets). This leaves the V term as the only remaining non-white error in the measurement equation. We could make another assumption that we have a base station transmitting the temporal variations to the aircraft, but this may be impractical from an operational standpoint. We could also group the temporal variations (V term) in with the white noise term, although this would greatly increase the noise strength of $\hat{w}$, and calling the resulting noise source Gaussian is a poor assumption. In [14], an alternative method to handle V was

presented—it is added as a observable filter state. With the simplifications made, the measurement relates to an aircraft’s latitude and longitude primarily through a highly non-linear function $M_2(lat, lon)$ plus a linear term V . We will apply a marginalized particle filter in the next section to estimate the aircraft’s latitude and longitude.

3.4 Filter Design

We will now present our filter design for magnetic anomaly navigation. Map based navigation systems are highly non-linear. We chose to use a particle filter to handle these non-linearities. Specifically, we are applying a marginalized particle filter (MPF), also known as a Rao-Blackwellized particle filter [57]. We wish to model an INS, which will require a high number of filter states. Particle filters are known to handle high dimensional problems poorly [57]. The MPF uses particles to represent only the non-linear states. Each particle has an associated Kalman filter to estimate the linear states. For our purposes, the only non-linear states are the latitude and longitude states. Our filter will consist of a 9-state Pinson INS error model from [65] which contains estimates of the INS position, velocity, and tilt errors, as well as two barometer states, a temporal variation state as described in [14], and a constant offset state. An inertial navigation system is unstable in the vertical channel and will quickly diverge without altitude aiding. We apply altitude aiding directly to the mechanization equations to constrain the vertical channel. This barometer aiding in the mechanization equations is modeled with the two barometer states. For the constant offset state c , recall that part of forming the measurement equation involves subtracting the IGRF—if an accurate location is not known when the magnetic navigation system is initialized, subtracting the IGRF with an approximate location will lead to a bias. The IGRF is a long-wavelength model so we can assume this bias is constant. It is important to distinguish this constant bias from the temporal variation

bias. The temporal variation bias is designed to capture small drifts in the field of around 10–20 nano-Teslas. If the errors in subtracting the IGRF field led to a 50 nano-Tesla offset, the temporal variation state would not model this well. This gives a total of 13 states which are given as

$$\mathbf{x} = [\delta lat \quad \delta lon \quad \delta alt \quad \delta v_n \quad \delta v_e \quad \delta v_d \quad \epsilon_x \quad \epsilon_y \quad \epsilon_z \quad \delta h_a \quad \delta \hat{a} \quad V \quad c]^T, \quad (51)$$

where:

δlat , δlon and δalt are the INS position errors

δv_n , δv_e , and δv_d are the INS velocity errors

ϵ_x , ϵ_y , and ϵ_z are the INS tilt errors

δh_a is the aiding altitude error

$\delta \hat{a}$ is the vertical acceleration error

V is the filter estimated temporal variation

c is the filter estimated constant bias error

To implement a MPF, we next have to partition the states into non-linear and linear components. See [57] for a complete description of the MPF. The two horizontal position states are non-linear and the remaining states are linear:

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{x}_t^n \\ \mathbf{x}_t^l \end{bmatrix}. \quad (52)$$

Eq. 51 is already partitioned correctly. We now express the linear and non-linear states as a sum of a non-linear part as well as a linear part, as shown in [57]. These are given by

$$\begin{aligned} \mathbf{x}_{t+1}^n &= f_t^n(\mathbf{x}_t^n) + \mathbf{A}_t^n(\mathbf{x}_t^n)\mathbf{x}_t^l + \mathbf{w}^n, \\ \mathbf{x}_{t+1}^l &= f_t^l(\mathbf{x}_t^n) + \mathbf{A}_t^l(\mathbf{x}_t^n)\mathbf{x}_t^l + \mathbf{w}^l. \end{aligned} \quad (53)$$

The horizontal position states are considered non-linear states because their posterior distributions (distributions after a map-matching update) are potentially multi-modal, non-Gaussian distributions. The horizontal position state dynamics, however, are still linear. This means there are no non-linear functions in Eq. 53. We can therefore simplify Eq. 53 by replacing the non-linear functions f^n and f^l with matrices:

$$\begin{aligned}\mathbf{x}_{t+1}^n &= \mathbf{A}_{n,t}^n(\mathbf{x}_t^n) + \mathbf{A}_{l,t}^n(\mathbf{x}_t^n)\mathbf{x}_t^l + \mathbf{w}^n, \\ \mathbf{x}_{t+1}^l &= \mathbf{A}_{n,t}^l(\mathbf{x}_t^n) + \mathbf{A}_{l,t}^l(\mathbf{x}_t^n)\mathbf{x}_t^l + \mathbf{w}^l.\end{aligned}\tag{54}$$

It is helpful to think of these matrices as follows.

$\mathbf{A}_n^n$ represents how the non-linear states propagate with respect to themselves

$\mathbf{A}_l^n$ represents how the non-linear states propagate with respect to the linear states

$\mathbf{A}_n^l$ represents how the linear states propagate with respect to the non-linear states

$\mathbf{A}_l^l$ represents how the linear states propagate with respect to themselves

Each of these four matrices is needed in the equations for the marginalized particle filter. These four matrices are simply the INS Pinson error model (see [65]) partitioned into four components, with the addition of the barometer, temporal variation, and constant offset states. These four matrices are given by

$$\mathbf{A}_n^n = \begin{bmatrix} 0 & 0 \\ \frac{v_e \tan L}{r_e \cos L} & 0 \end{bmatrix}_{2 \times 2}, \tag{55}$$

$$\mathbf{A}_l^n = \begin{bmatrix} \frac{-v_n}{r_e^2} & \frac{1}{r_e} & 0 & \mathbf{0}_{1 \times 8} \\ \frac{-v_e}{r_e^2 \cos L} & 0 & \frac{1}{r_e \cos L} & \mathbf{0}_{1 \times 8} \end{bmatrix}_{2 \times 11}, \tag{56}$$

$$\mathbf{A}_n^l = \begin{bmatrix} 0 & 0 \\ -v_e \left(2\omega \cos L + \frac{v_e}{r_e \cos L^2} \right) & 0 \\ 2\omega (v_n \cos L - v_d \sin L) + \frac{v_n v_e}{r_e \cos L^2} & 0 \\ 2\omega v_e \sin L & 0 \\ -\omega \sin L & 0 \\ 0 & 0 \\ -\omega \cos L - \frac{v_e}{r_e \cos L^2} & 0 \\ \mathbf{0}_{4 \times 1} & \mathbf{0}_{4 \times 1} \end{bmatrix}_{11 \times 2}, \quad (57)$$

$$\mathbf{A}_l^l = \begin{bmatrix} \mathbf{N}_{7 \times 7}^l & \mathbf{D}_{7 \times 2} & \mathbf{0}_{7 \times 2} \\ \mathbf{0}_{2 \times 7} & \mathbf{B}_{2 \times 2} & \mathbf{0}_{2 \times 2} \\ \mathbf{0}_{2 \times 7} & \mathbf{0}_{2 \times 2} & \mathbf{C}_{2 \times 2} \end{bmatrix}_{11 \times 11}, \quad (58)$$

where

$$\mathbf{N}^l = \begin{bmatrix} 0 & 0 & 0 \\ \frac{v_e^2 \tan L - v_n v_d}{r_e^2} & \frac{v_d}{r_e} & -2 \left(\omega \sin L + \frac{v_e \tan L}{r_e} \right) \\ -v_e \frac{v_n \tan L + v_d}{r_e^2} & 2\omega \sin L + \frac{v_e \tan L}{r_e} & \frac{v_n \tan L + v_d}{r_e} \\ \frac{v_n^2 + v_e^2}{r_e^2} & -2 \frac{v_n}{r_e} & -2 \left(\omega \cos L - \frac{v_e}{r_e} \right) \\ \frac{-v_e^2}{r_e^2} & 0 & \frac{1}{r_e} \\ \frac{v_n}{r_e^2} & \frac{-1}{r_e} & 0 \\ \frac{v_e \tan L}{r_e^2} & 0 & \frac{-\tan L}{r_e} \end{bmatrix}$$

$$\begin{bmatrix}
-1 & 0 & 0 & 0 \\
\frac{v_n}{r_e} & 0 & -f_d & f_e \\
2\omega \cos L + \frac{v_e}{r_e} & f_d & 0 & -f_n \\
0 & -f_e & f_n & 0 \\
0 & 0 & -\omega \sin L - \frac{v_e \tan L}{r_e} & \frac{v_n}{r_e} \\
0 & \omega \sin L + \frac{v_e \tan L}{r_e} & 0 & \omega \cos L + \frac{v_e}{r_e} \\
0 & \frac{-v_n}{r_e} & -\omega \cos L - \frac{v_e}{r_e} & 0
\end{bmatrix}_{7 \times 7}, \quad (59)$$

where

L is latitude in radians

r_e is the Earth's radius equal to 6378135 meters

ω is the Earth's rotation rate equal to $7.2921151467 \times 10^{-5}$ rad/sec

f_n , f_e , and f_d are the north, east, and down specific forces

v_n , v_e , v_d are the north, east, and down velocities

and

$$\mathbf{B} = \begin{bmatrix} \frac{-1}{\tau_b} & 0 \\ -k_3 & 0 \end{bmatrix}_{2 \times 2}, \quad (60)$$

where

τ_b is the barometer error time constant

k_3 is a barometer aiding constant used in the third order altitude aiding feedback loop

Additionally,

$$\mathbf{D} = \begin{bmatrix} k_1 & 0 \\ \mathbf{0}_{2 \times 1} & \mathbf{0}_{2 \times 1} \\ -k_2 & 1 \\ \mathbf{0}_{3 \times 1} & \mathbf{0}_{3 \times 1} \end{bmatrix}_{7 \times 2}, \quad (61)$$

where k_1 , and k_2 are barometer aiding constants used in the third order altitude aiding feedback loop

and

$$\mathbf{C} = \begin{bmatrix} \frac{-1}{\tau_{TV}} & 0 \\ 0 & 0 \end{bmatrix}_{2 \times 2}, \quad (62)$$

where

τ_{TV} is the temporal variation bias time constant

We now address the measurement equation for magnetic anomaly navigation. We will use the simplified measurement equation given by Eq. 50. We wish to put this measurement equation into a valid form for the marginalized particle filter. From [57], the general form for the measurement equation of a marginalized particle filter expresses the measurement as a sum of a non-linear function of the non-linear states, a linear function of the linear states, and a measurement noise w :

$$\mathbf{y}_t = h_t(\mathbf{x}_t^n) + \mathbf{C}_t \mathbf{x}_t^l + w. \quad (63)$$

We slightly rearrange the measurement equation from Eq. 50 by bringing the IGRF term to the left side of the equation. Our measurement is therefore the raw magnetometer measurement minus the predicted IGRF field. We also break the measurement equation into a non-linear part represented by the map interpolation function,

and a linear portion which captures the temporal variation and constant bias states. The measurement equation is now

$$Z_t^m - I(\hat{lat}, \hat{lon}, \hat{alt}) = M_2(lat_{ins} + \delta lat, lon_{ins} + \delta lon) + \mathbf{C}\mathbf{x}_t^l + w, \quad (64)$$

where

Z_t^m is the magnetometer raw measurement at time t

I is the IGRF intensity computed with an approximate position

M_2 is the non-linear 2D interpolation function created from the 2D map grid.

$\mathbf{C}$ is the matrix representing the linear portion of the measurement equation

w is the measurement noise

The $\mathbf{C}$ matrix shows that the measurement is also the sum of the temporal variation and constant bias states:

$$\mathbf{C} = \begin{bmatrix} \mathbf{0}_{1 \times 9} & 1 & 1 \end{bmatrix}_{1 \times 11}. \quad (65)$$

Finally, we must define the measurement noise matrix and the system noise matrix.

The measurement noise is modeled a zero-mean Gaussian with a variance σ_{mag}^2 :

$$R = E[w^2] = \sigma_{mag}^2, \quad (66)$$

where σ_{mag}^2 is the measurement accuracy for measuring the magnetic anomaly field.

This is not the same thing as the magnetometer absolute accuracy, which is accurate to less than 1 nano-Tesla. While the magnetometer can measure the total field to a very high accuracy, this measurement noise captures the accuracy of measuring the magnetic anomaly field, which is a single component of the total field, and therefore

has a higher noise variance due to the various errors presented in the measurement equation section. The non-linear states do not have any dynamic noise terms:

$$E[\mathbf{w}_n^2] = \mathbf{Q}^n = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}_{2 \times 2}. \quad (67)$$

The linear-state dynamics noise captures the velocity and angular random walk of the IMU and the driving noise for the temporal variation bias:

$$E[\mathbf{w}_1^2] = \mathbf{Q}^l = \text{diag} \left(\begin{bmatrix} 0 & \mathbf{VRW}_{1 \times 3} & \mathbf{ARW}_{1 \times 3} & B & 0 & T & 0 \end{bmatrix} \right)_{11 \times 11}, \quad (68)$$

$$B = \frac{2\sigma_b^2}{\tau_b}, \quad (69)$$

$$T = \frac{2\sigma_{tv}^2}{\tau_{tv}}, \quad (70)$$

where

$\mathbf{VRW}$ is the zero-mean Gaussian noise covariance of the velocity random walk

$\mathbf{ARW}$ is the zero-mean Gaussian noise covariance of the angular random walk

B and T are the driving noise strengths for the barometer and temporal variation moving biases

σ_b^2 is the barometer error variance

τ_b is the barometer error time constant

σ_{tv}^2 is the temporal variation variance

τ_{tv} is the temporal variation time constant

The filter approach can now be implemented in the MPF algorithm given in [57].

3.5 Magnetic Navigation Algorithm

In this section we present the filter algorithm which is needed to implement the previously described magnetic anomaly navigation system. This algorithm follows closely from the MPF algorithm given in [57], with added detail on the map pre-computation steps as well as how the INS reference trajectory is used. Recall that each particle has an associated Kalman filter. Because there are only non-linearities in the measurement equation, and not in the state dynamics, many of the Kalman filter equations can be evaluated once for all particles [57].

Algorithm Input: Magnetic anomaly map, raw magnetometer measurements, INS ΔV 's and $\Delta\Theta$'s, barometer-derived altitude, IGRF model

Algorithm Output: Errors in an INS navigation solution including position errors, velocity errors, and tilt errors

Notation: (i) denotes an individual particle, **bold** denotes vectors and matrices, superscripts l and n denote the linear and non-linear states respectively, and t denotes current time step, and $t - \Delta t$ denotes the previous time step.

1. Initialization:

- (a) Create interpolation function $M_2 = f(lat, lon)$ from magnetic anomaly grid. Choose number of particles equal to N .
- (b) For $i = 1, \dots, N$ initialize the non-linear particles by drawing from a Gaussian distribution with mean μ_0^n and covariance P_0^n : $\mathbf{x}_0^{n,(i)} \sim \mathcal{N}(\mu_0^n, P_0^n)$. Next, initialize all the linear state particles with the same initial conditions: $\mathbf{x}_0^{l,(i)} = \bar{\mathbf{x}}_0^l$. Finally, initialize the linear particle covariance, which is a single common matrix for all of the linear particles, as $\mathbf{P}_0 = \bar{\mathbf{P}}_0$

2. INS Mechanization: Use the new ΔV 's and $\Delta\Theta$ measurements to compute the

INS solution at time t . Barometer aiding occurs at this step.

3. Compute Dynamics Matrix: Use the INS solution to compute $\mathbf{A}_n^n$, $\mathbf{A}_n^l$, $\mathbf{A}_l^n$, and $\mathbf{A}_l^l$ at time t .
4. Evaluate IGRF Model: Use the filter estimated position, given as $INS + \hat{\mathbf{x}}$, to evaluate the IGRF model for the filter estimated latitude, longitude, and altitude.
5. Particle Filter Measurement Update: For $i = 1, \dots, N$

- (a) Evaluate the expected magnetic intensity for each particle hypothesis. The expected intensity is based on both the particle's location hypothesis (non-linear states 1 and 2) and the temporal variation hypothesis (linear state 8).

$$I_t^{(i)} = M_2(lat_t^{ins} + x_t^{n,(i)}(1), lon_t^{ins} + x_t^{n,(i)}(2)) + x_t^{l,(i)}(8) \quad (71)$$

Using the raw magnetometer measurement y_t , denote the residual particle intensity by $e_t^{(i)} = y_t - (I_t^{(i)} - IGRF)$.

- (b) The measurements are assumed Gaussian with mean $e_t^{(i)}$ and covariance $V_t = \mathbf{C}\mathbf{P}_{t|t-\Delta t}\mathbf{C}^T + R$. $\mathbf{C}$ is the linear measurement matrix, in this case $\begin{bmatrix} \mathbf{0}_{1 \times 9}, 1, 1 \end{bmatrix}_{1 \times 11}$, and $\mathbf{P}_{t|t-\Delta t}$ is the linear state's Kalman filter covariance. We can evaluate the probability for each particle from a Gaussian distribution given by

$$q_t^{(i)} = \exp\left(-\frac{1}{2}(e_t^{(i)}V_t^{-1}e_t^{(i)T})\right) \quad (72)$$

The weights then must be normalized according to

$$\tilde{q}_t^{(i)} = \frac{q_t^{(i)}}{\sum_{j=1}^N q_t^{(j)}} \quad (73)$$

6. Resampling: Resample particles with a chosen resampling strategy. We resampled at each step with sequential importance sampling [36].
7. Kalman Filter Measurement Update: The next step is to update the linear states of the Kalman Filter with the measurement. We first define the Kalman gain matrix and covariance matrix, which are the same for all particles:

$$\mathbf{K}_t = \mathbf{P}_{t|t-\Delta t} \mathbf{C}^T V_t^{-1}, \quad (74)$$

$$\mathbf{P}_{t|t} = \mathbf{P}_{t|t-\Delta t} - \mathbf{K}_t V_t \mathbf{K}_t^T. \quad (75)$$

We next calculate a measurement residual $J_t^{(i)}$ and apply the Kalman gain for each individual particle to determine the Kalman filter updated linear state. For $i = 1, \dots, N$

$$J_t^{(i)} = M_2(lat_t^{ins} + \mathbf{x}_t^{n,(i)}(1), lon_t^{ins} + \mathbf{x}_t^{n,(i)}(2)) + \mathbf{x}_t^{l,(i)}(8), \quad (76)$$

$$\hat{\mathbf{x}}_{t|t}^{l,(i)} = \hat{\mathbf{x}}_{t|t-\Delta t}^{l,(i)} + \mathbf{K}_t \left(y_t - J_t^{(i)} \right). \quad (77)$$

8. Particle Filter Time Update: The particles are propagated in time according the system dynamics model. First, compute the Cholesky decomposition $\mathbf{q}_{chol} = chol \left(\mathbf{A}_l^n \mathbf{P}_{t|t} (\mathbf{A}_l^n)^T + \mathbf{Q}^n \right)$. Random noise must be added to the non-linear states, which is denoted by a function *randn* to represent adding white Gaussian noise. The particles are then propagated according to

$$\mathbf{x}_{t+1}^{n,(i)} = \mathbf{A}_n^n \mathbf{x}_{t|t}^{n,(i)} + \mathbf{A}_l^n \mathbf{x}_{t|t}^{l,(i)} + \mathbf{q}_{chol} \cdot randn()_{2 \times 1}. \quad (78)$$

9. Kalman filter Time Update: Define two intermediate matrices, N_t and L_t , which are the same for each particle, and compute the propagated Kalman filter co-

variance, which is also the same for each particle:

$$\mathbf{N}_t = \mathbf{A}_l^n \mathbf{P}_{t|t} (\mathbf{A}_l^n)^T + \mathbf{Q}^n, \quad (79)$$

$$\mathbf{L}_t = \mathbf{A}_l^l \mathbf{P}_{t|t} \mathbf{A}_l^n \mathbf{N}_t^{-1}, \quad (80)$$

$$\mathbf{P}_{t+1|t} = \mathbf{A}_l^l \mathbf{P}_{t|t} (\mathbf{A}_l^l)^T + \mathbf{Q}^l - \mathbf{L}_t \mathbf{N}_t \mathbf{L}_t^T. \quad (81)$$

Finally, define a time update residual for each particle and calculate the time update. For $i = 1, \dots, N$

$$\mathbf{z}_t^{(i)} = \mathbf{x}_{t+1|t}^{n,(i)} - \mathbf{A}_n^n \mathbf{x}_{t|t}^{n,(i)}, \quad (82)$$

$$\mathbf{x}_{t+1|t}^{l,(i)} = \mathbf{A}_l^l \mathbf{x}_{t|t}^{l,(i)} + \mathbf{A}_l^n \mathbf{x}_{t|t}^{n,(i)} + \mathbf{L}_t (\mathbf{z}_t^{(i)} - \mathbf{A}_l^n \mathbf{x}_{t|t}^{l,(i)}). \quad (83)$$

10. Determine the expected means and covariances of the non-linear and linear states. The covariance of the linear states is already computed as $\mathbf{P}_{t+1|t}$

$$\hat{\mathbf{x}}_{t|t}^l = \sum_{i=1}^N \hat{q}_t^{(i)} \hat{\mathbf{x}}_{t|t}^{l,(i)}, \quad (84)$$

$$\hat{\mathbf{x}}_{t|t}^n = \sum_{i=1}^N \hat{q}_t^{(i)} \hat{\mathbf{x}}_{t|t}^{n,(i)}, \quad (85)$$

$$\hat{\mathbf{P}}_{t|t}^n = \sum_{i=1}^N \hat{q}_t^{(i)} \left((\hat{\mathbf{x}}_{t|t}^{n,(i)} - \hat{\mathbf{x}}_{t|t}^n) (\hat{\mathbf{x}}_{t|t}^{n,(i)} - \hat{\mathbf{x}}_{t|t}^n)^T \right). \quad (86)$$

11. Move to the next measurement time t and iterate from step 2.

3.6 Filter Design Conclusions

This chapter has provided a detailed explanation of the design of a 13 state Rao-Blackwellized particle filter which estimates the errors in a navigation grade INS. A complete measurement equation for magnetic anomaly field navigation is provided.

This measurement equation is used to design a navigation filter. The filter uses aircraft field compensated magnetometer measurements which have had the IGRF field removed. These measurements provide an estimate of both the aircraft's position as well as an estimate of the temporal variations. The temporal variation state must be modeled correctly and a discussion on the observability of this state was provided.

IV. Flight Test Results

This chapter details the results of the magnetic anomaly navigation flight tests which were conducted to test the feasibility of magnetic anomaly navigation. The first flight test involved both a mapping mission as well as a navigation mission. The mapping mission created a magnetic anomaly map over a previously mapped area which was created three years ago. The two maps matched up to better than a nano-Tesla, showing the stability of magnetic anomaly maps. The navigation mission flew through the previously mapped area recording magnetometer measurements as well as INS and barometer data, and GPS data for truthing. The filter was implemented using this recorded data and obtained 13 meters DRMS error. These results were considered “ideal case” results because the flight occurred at a low altitude over a high quality map. The obtained results are far more accurate than anything presented previously in the literature. The second flight test details the results of a “cross country” flight test. This flight test flew from Virginia to Iowa at an altitude of approximately 3000 meters. Unlike the previous flight test over a regional high quality magnetic anomaly map, no high quality map was available over such a large region. This necessitated the use of the North American Magnetic Anomaly Database (NAMAD). This map is of poor quality compared to the high accuracy regional map used over Louisa, VA. The navigation accuracy obtained with this flight test was not nearly as good as the ideal-case results achieved in the first flight test. The decrease in accuracy is explained by two factors—flying at a higher altitude and using a much lower quality map. Using the low quality map a tactical grade INS achieved a DRMS error of approximately 3 kilometers over a 5 hour flight. This was a significant improvement over the unaided tactical grade INS, which drifted over 50 kilometers. When using a high quality *simulated* map, navigation accuracy was shown to be approximately 150 meters DRMS over a 5 hour flight at 3000 meters

altitude. This is still less than the accuracy obtained in the first flight test due to the much higher altitude of the cross-country flight. Finally, Cramer-Rao Lower Bound (CRLB) analysis is used to provide an altitude vs. accuracy trade-space analysis.

4.1 Ideal-Case Results

Flight Test Data.

We contracted out a flight test in 2015 to a geo-survey company called Sander Geophysics to test the accuracy of the developed navigation system. Fig. 56 shows the geo-survey aircraft taking off. The magnetometer used to generate the results in this chapter is located in the lower tail boom. A magnetic anomaly map was created by Sander Geophysics over Louisa, VA in 2012. We re-mapped a subset of this previously mapped area to evaluate map stability. We flew over the mapped area collecting data with a Geometrics 823A optically pumped cesium magnetometer, a barometer, and a navigation-grade inertial navigation system. We input this collected data and the 2012 magnetic anomaly map into our navigation filter. The magnetic anomaly map as well as the flight path is shown in Fig. 63. The flight begins with a U shaped trajectory covering a large area of the 2012 map followed by figure eights over the area we re-mapped in 2015. The gradients in the magnetic field at this altitude



Figure 56. Geo-Survey Aircraft Taking Off for Test Flight - Magnetometers Located on Tail Boom




Figure 57. 2012 Magnetic Anomaly Map Over Louisa VA and 2015 Flight Profile Over Mapped Area

and at this location can be considered on the moderate to high end of what generally exists around the United States at similar altitudes.

The magnetic anomaly map can be used to create an interpolation function. The expected measurement for the flight can be obtained by passing the true GPS flight position through the interpolation function. This expected measurement is not used by the filter, as it requires the use of GPS, but is only used to gain insight into how well the actual raw measurements match the magnetic anomaly map. We validate our measurement equation in Table 3. Each row of the table removes an additional source of error from the raw measurements. The standard deviation of the difference between the expected measurements and the actual measurements is then given. The standard deviation decreases sharply with the removal of the IGRF field. The temporal variations are removed by subtracting the measured variations of a nearby base station. This base station data is not assumed available while flying, but is shown here to help validate the measurement equation. A slight decrease in the standard deviation of

Table 3. Measurement Equation Validation

Error Term	STD of Difference Between Pred. and Meas. Magnetic Field (nT)
No Corrections	21.21 nT
IGRF Correction	3.35 nT
Plus Temporal Variation Correction	2.09 nT
Plus Aircraft Field Correction	1.55 nT

the expected and actual measurements is observed when the temporal variations are removed. Finally, the effects of the aircraft field are removed with an aircraft-field compensation system. These compensated measurements were provided by Sander Geophysics and are computed by applying compensation coefficients obtained during a calibration flight to the remaining flight data. The lowest standard deviation of the difference between the expected measurements and the actual measurements is obtained after applying this final correction. The final standard deviation of the error between the expected and actual measurements is 1.55 nano-Teslas, roughly 1% of the total variation seen in the magnetic intensity along the flight profile. This is analogous to a high signal to noise ratio, and provides valuable insight into expected navigation performance. Fig. 58 shows a six minute segment of the errors between the expected measurement and the raw magnetometer measurements with all of the previously described errors removed. As can be seen in the figure, the errors are time correlated with a time constant on the order of less than a minute. We believe these remaining errors are the result of using an under-sampled map, although a test flight using a fully sampled map would be necessary to determine if this is truly the case. If this is the case, then using a fully sampled map should lead to higher navigation accuracy, as we believe these errors are the dominant errors remaining in the measurements.




Figure 58. Difference Between the Expected Measurements From Interpolation Function (Using True GPS Positions) and the Raw Measurements (made zero-mean)

Map Stability.

In addition to collecting data to test the navigation system, we flew a magnetic survey over a subset of the original 2012 survey area. The resulting map was almost identical to the previous map. We computed the difference between the two maps, as shown in Fig. 59. Both maps, with the exception of the edges, have a max difference of around 2 nano-Teslas. The errors on the edges of the map are an artifact of how the maps were processed—a larger map would push these errors out farther. When ignoring the edge effects, the standard deviation of the error between the two maps is 1.2 nano-Teslas. This gives further evidence that the magnetic anomaly maps are stable over time, or at least will not be a driving error term in a navigation filter. These errors would appear as slowly varying biases on the order of a single nano-Tesla—the previously shown errors vary much quicker than this, and by a larger amount, so map stability does not appear to be a driving error term for navigation accuracy.




Figure 59. Difference Between the 2012 Magnetic Anomaly Map and the 2015 Magnetic Anomaly Map

Navigation Results.

We ran the designed navigation filter using the collected flight data combined with the 2012 survey map. The INS mechanization was run independently of the rest of the filter, with barometer aiding occurring in the mechanization equations. No feedback was provided by the filter, although feedback may be required with a less accurate INS. The INS provided the reference trajectory for the filter, which estimated the errors in the INS. The INS solution is used while evaluating the map interpolation function, as shown in Eq. 64, as well as when computing the dynamics matrix, which depends on latitude, velocity, and specific forces. We ran the filter against an hour long segment of the flight profile using only INS, barometer, and magnetometer measurements. We then used the true GPS solution to evaluate the performance of the filter. Fig. 60 shows the north and east errors over the hour long flight as well as the predicted filter standard deviations. The statistics for the flight are given in Table 13. Overall, the system obtained a DRMS error of 13 meters. The filter clearly constrained the drift of the INS. Furthermore, the filter's covariance

Table 4. “Best Case” Flight Test Navigation Accuracy Results

	North Channel	East Channel
Mean	-2.2 meters	2.7 meters
Standard Deviation	9.0 meters	8.9 meters
DRMS	13.1 meters	
Unaided INS Error After 1 Hour	344 meters	

bounds appear reasonable. It is important to note that the filter’s covariance bounds are strongly correlated with the gradient of the magnetic field at a given location. Steeper magnetic gradients lead to a more accurate position estimate, and shallower gradients lead to a degraded position estimate. The navigation-grade INS provides an excellent short term solution to help the navigation solution pass through these areas of low magnetic gradient without losing too much accuracy. A high quality (navigation grade) INS is essential to achieving the accuracies demonstrated here. It is also apparent that the filter standard deviation has reached a steady-state that varies around 10–20 meters. The mean error in the north and east channels are both less than 3 meters, and the plots indicate it is reasonable to consider the errors zero-mean, with no clear bias existing in the north or east channels. The standard deviations of the north and east channels are very similar. This is expected due to the fact that the steep gradients which exist in the map are in a north-east to south-west direction. Flying over a different area may have better or worse errors in one channel due to consistently steeper north or east gradients. Due to the accurate modeling of the INS, the corrected latitude and longitude errors also lead to corrected velocity errors of approximately 0.1 meters/second.

The filter also estimated the temporal variations. The temporal variations were rather calm on the day of the flight, so in order to better determine the filter observability of these variations we further artificially corrupted the measurements with




Figure 60. North and East Error Over 1-Hour Segment of Flight Profile

actual recorded data from a different magnetically stormy day in a second filter run. Previously shown in Fig. 58, even with the temporal variations removed, the measurements have remaining high frequency errors with a standard deviation of around 2 nano-Teslas. Consequently, the filter is only able to estimate the larger, longer wavelength (therefore, lower frequency) components of the temporal variations. Fig. 61 shows the filter's estimation of the temporal variations on the stormy day. When the magnetometer measurements were corrupted to simulate the stormy day, the filter achieved a DRMS error of 15 meters—only a mild increase in errors compared to a DRMS of 13.1 meters on the quiet day. It is important to note that the filter temporal variation state is estimating *any* residual error in the measurements, so it will not match the actual temporal variations exactly. The temporal variation state is also capturing any residual aircraft fields, variations from the map altitude, or un-captured high frequency crustal components (the 2012 map was not quite fully sampled). The apparent lag in the estimation of the temporal variations is caused by

the tuning parameters related to the temporal variation state. Better filter performance was achieved when using a long time constant for the temporal variation state. Restricting the temporal variation state to slow changes increases filter stability. If a short time constant was used for the temporal variations, the filter began to assume most of the changes in the measurements were due to changing temporal variations, which degraded filter performance. A bad magnetic storm in the context of magnetic anomaly navigation does not so much correspond to large temporal variations as it does to large higher frequency variations. A magnetic storm could cause a large rise in the magnetic field, but if this change is slow relative to how the crustal field is changing throughout the flight, the filter will likely retain observability of the temporal variations. Temporal variations or aircraft fields with frequencies which overlap with the crustal field frequencies are much more problematic for a navigation filter than slowly varying errors. Conversely, temporal variations and aircraft field frequencies which are much higher than the crustal field frequencies tend to average out, and do not greatly degrade navigation performance.

Ideal-Case Conclusions.

The presented navigation filter achieved DRMS errors of around 13 meters under favorable conditions with actual flight test data. These conditions included a relatively low altitude, a high quality magnetic anomaly map, and a clean sensor environment. The availability of high quality magnetic anomaly maps is a large obstacle for magnetic anomaly navigation at the accuracies demonstrated here. Older less accurate maps are certainly useable; however, they would likely not obtain the accuracies demonstrated here. Finally, a clean sensor environment such as the one located on the geo-survey aircraft may not be feasible for all applications. Using a magnetometer inside the aircraft will also lead to decreased performance, and further




Figure 61. Filter Estimation of Temporal Variations

experimentation with real data is likely necessary to characterize performance for this case. The accuracies demonstrated using only passive sensors demonstrate that navigation using the Earth's magnetic anomaly field may be a viable approach for future GPS backup and alternative systems for aircraft in flight.

4.2 Cross Country Results

NAMAD Map Errors.

Map accuracy is a large concern for magnetic anomaly navigation. Magnetic mapping is costly and time-consuming. Because the Earth's magnetic anomaly field changes in three dimensions, with decreasing frequency content at higher altitudes, high resolution maps cannot be made from space. High resolution magnetic anomaly maps must be made with aerial surveys, which can be expensive in the context of continental or world-sized maps. The North American Magnetic Anomaly Database (NAMAD) is a large-scale compilation map of North America at 305 meters above




Figure 62. North American Magnetic Anomaly Database [3]

ground level. The NAMAD is shown in Fig. 62. It was created by merging hundreds of smaller magnetic surveys into one continuous map. To navigate using the NAMAD, it is important to understand the types of errors present in the map. The first major error in the NAMAD results from poorly geo-located data. Many of the magnetic surveys used to make the NAMAD were created before the use of GPS. The magnetic measurements from these surveys were often geo-located using crude methods with potentially kilometer level errors. A correctly geo-located map is clearly a basic requirement for a successful map-based navigation system. The second major type of error in the NAMAD results from under-sampling the magnetic field. The surveys used in the NAMAD have varying resolutions, with some grid spacing as large as 8 kilometers [3]. To fully sample a magnetic anomaly field, grid spacing must be approximately equal to height [53]. This indicates that there will be high frequency errors when flying over these low-resolution maps. It is important to note, however, that these under-sampled maps become fully sampled when upward continuing the data. For example, a magnetic survey flown at 300 meters with 1 kilometer grid spacing will be under-sampled at 300 meters but fully sampled when the map is upward continued to 1 kilometer altitude. This is due to the fact that the upward continuation operation is a low-pass filter. In practice the accuracy of this upward continuation

depends on many factors such as the original under-sampled map not having aliased components and the mitigation of edge effects. While accurate upward continuation depends on many factors, it is important to note that the high frequency errors present at low map altitudes will be mostly eliminated when flying at higher altitudes.

There are several other common errors in NAMAD. NAMAD is a patchwork of small magnetic surveys. Consequently, it is unable to capture long-wavelength magnetic anomalies. Furthermore, many magnetic surveys were leveled to an arbitrary constant value. We previously demonstrated the observability of temporal variations in a magnetic anomaly navigation system. These variations were observable due to their long-wavelengths compared to the short-wavelength magnetic anomalies. The long-wavelength errors in the NAMAD will likely also be observable when sufficient variation exists in the magnetic anomaly map. The final type of error in the NAMAD pertains to some older surveys used to make the map. Many older magnetic anomaly maps were hand contoured. In the creation of the NAMAD these hand-drawn maps were digitized. These types of errors are difficult to categorize but clearly digitizing older hand-counterred maps will yield less accurate results than a modern magnetic survey.

Map Gradient vs. Altitude.

Altitude is a major variable which affects the performance of a magnetic anomaly navigation system. The lower an aircraft flies, the more spatial variation will exist in the magnetic anomaly map. As stated previously, the shortest horizontal wavelength in a magnetic anomaly map is approximately equal to the height above ground level. It is important to note, however, that spatial variation is an indirect contributor to navigation accuracy. Fundamentally, it is the map gradient which determines

navigation accuracy when using the proposed navigation system. For real magnetic anomaly data, high spatial variation and steep gradients are closely correlated. The steepness of magnetic gradients decreases quickly with altitude. We used the NAMAD map to evaluate how magnetic anomaly gradients change with altitude. We upward continued the NAMAD map to several altitudes and examined the map gradients. Table 5 shows the means and standard deviations of the NAMAD gradients at various altitudes, along with a predicted navigation accuracy assuming the ability to measure the magnetic anomaly field to an accuracy of 1 nano-Tesla. The predicted navigation accuracies are obtained by multiplying measurement accuracy by the inverse of the magnetic gradient. Note that this method of predicted navigation accuracy is simply a first order approximation. The predicted accuracy at 300 meters matches well with our actual flight test data in [15]. In this chapter we will use real and simulated data to verify the predicted performance at higher altitudes. It is important to note that while we have no control over the steepness of the magnetic anomaly gradients, this is only one of two related variables, the other being sensor accuracy. If future magnetic sensors allow for more accurate measurement of the magnetic anomaly field, the accuracies at all altitudes will be increased. Finally, it is important to stress the difference between sensor accuracy and measurement accuracy. Measurement accuracy is the ability to measure the magnetic *anomaly* field while sensor accuracy is the ability to measure the *total* field. Measuring the magnetic anomaly field requires accurate removal of aircraft effects, temporal variations, and the core field. A more accurate sensor alone does not increase measurement accuracy without addressing these other issues.

Table 5. Magnetic Anomaly Gradients of NAMAD Map at Various Altitudes and Predicted Navigation Accuracy

	Mean (nT/km)	STD (nT/km)	Pred. Nav. Accuracy (m)
300m AGL	34.0	76.0	29
1km AGL	19.1	29.3	52
10km AGL	3.9	3.4	256



Figure 63. Cross-country flight path

Cross-Country Flight Test Data.

The cross-country flight path from the flight test data is shown in Fig. 63. The flight originated in Louisa, VA and ended in Iowa. The flight took a total of 5 hours and was around 1500 kilometers long. The aircraft flew with a navigation grade INS, a Geometrics 823 optically pumped cesium magnetometer, a barometer, and a GPS to collect truth data. The magnetometer was out on a boom at the back of the aircraft to minimize aircraft interference. Using the GPS solution from the flight test, we can obtain the expected measurement profile along the flight path. This expected measurement is not used by the filter, but is shown here to give intuition as to how well the true measurements match the expected measurements from the map. Fig. 64 shows the actual recorded measurements from the flight as well as the expected measurements from the NAMAD map using GPS. It is clear from Fig. 64 that the actual measurement is tracking the expected measurement throughout the 5 hour flight. The same data is shown in Fig. 65a as an error plot. Included in Fig. 65a






Figure 64. Actual Recorded Measurement and Expected Measurement From NAMAD Map

is the standard deviation of the observed measurement errors from the previous regional results when using a high quality magnetic anomaly map. When using a high quality magnetic anomaly map the standard deviation of the difference between the actual and expected measurements was only 1.3 nano-Teslas. In this case, the standard deviation of the difference between the actual and expected measurements is 40 nano-Teslas. It is apparent that one of the driving variables for navigation accuracy when using the NAMAD will be the map accuracy. Fig. 65b shows the magnetic gradient along the flight profile in nano-Teslas/kilometer. The map gradient can give insight into the approximate accuracy of the navigation system. If flying over an area of the map with a map gradient of 40 nano-Teslas/kilometer, a 40 nano-Tesla error in the measurements will cause approximately kilometer level errors. With the errors shown between the expected magnetic measurements and the NAMAD map, as well as the computed gradient, we would expect navigation accuracies of a few kilometers. This is much worse than what was previously demonstrated with the regional results, but we can confidently state map errors lead to most of this decreased accuracy.



(a) Error Between Actual Recorded Measurement and Expected Measurement From NAMAD Map



(b) Magnetic Gradient Along Flight Profile

Figure 65. Measurement Errors and Map Gradient

Finally, we wish to show the INS data from the flight. On-board the aircraft we used a navigation-grade INS to compute an estimated position solution. As we will show shortly, this navigation-grade INS solution, which only drifted a few kilometers throughout the 5 hour flight, was accurate enough that no apparent benefit is gained by the magnetic anomaly navigation system. To estimate the actual navigation potential of the NAMAD map, we needed to further degrade the INS measurements. Degraded INS measurements allowed the INS solution to drift far enough that a noticeable benefit from using the NAMAD map was observed. The navigation grade INS measurements were corrupted by adding a moving bias term into the accelerometer measurements. Fig. 66 shows the drifts of the real navigation grade INS as well as the tactical grade INS, which was simulated by corrupting the navigation grade INS measurements.

Cross-Country Results.

We first ran the navigation filter using all real data from our test flight. Fig. 67 shows the errors in the filter computed solution, as well as the errors in the INS. As can be seen in the figure, the INS out-performs the filter. In this particular case nothing was gained by bringing magnetometer measurements into the filter. This was an expected result. As we stated earlier, the expected accuracy using the NAMAD map is on the order of several kilometers. Because the INS only drifts several kilometers, we would not expect the filter solution to be much better than the INS.

Because of the high accuracy of the navigation-grade INS we were using, we were not able to demonstrate the navigation potential of using the NAMAD map. To estimate this expected accuracy, we instead use tactical-grade INS measurements. These tactical-grade measurements were made by corrupting the real accelerome-




Figure 66. Navigation Grade INS Drifts and Tactical Grade INS Drifts Over Full Flight

ter measurements with moving bias terms. We ran the filter using these degraded tactical-grade INS measurements and observed a noticeable increase in filter performance over the unaided tactical-grade INS. Fig. 68 shows the north and east errors of the filter solution. Table 6 shows the statistics for the filter run. The filter solution had a DRMS error of 3.2 kilometers while the tactical grade INS drifted over 50 kilometers. This demonstrates that using the NAMAD map to navigate does have navigation potential, just not at a very accurate level. Furthermore, a large portion of the navigation error can be attributed to NAMAD map error. We would expect far better results with a more accurate map. This indicates that better magnetic anomaly navigation accuracy is a very solvable problem, in the sense that it has a known solution (not in the sense of cost, time etc.). While the accuracy results of this navigation system are less than impressive, we successfully demonstrated constraining a tactical grade INS to a few kilometers over a five hour flight using only a freely available continental-sized magnetic anomaly map and all real flight test data.

As previously stated, a large majority of the positioning errors when using the real magnetometer measurements and the NAMAD map arise from poor map qual-




Figure 67. Filter Results Using Navigation Grade INS, Real Magnetometer Measurements, and NAMAD Map




Figure 68. Filter Results Using Simulated Tactical Grade INS, Real Magnetometer Measurements, and NAMAD Map

ity. Using a high quality magnetic anomaly map we were able to match expected measurements using the true GPS solution to actual measurements to within a few nano-Teslas. With the NAMAD map, the expected measurements and actual measurements were much worse, with a standard deviation of 40 nano-Teslas. We wish to predict navigation accuracy on this higher altitude cross country flight if a more accurate magnetic anomaly map were to become available. To simulate this more accurate map, we decided to call the NAMAD map “truth” and generated magnetometer measurements by corrupting the expected measurements from the NAMAD map with errors similar to what we observed with the regional high accuracy map. We did this by corrupting the expected measurements with a FOGM random process with a variance of 4 nano-Teslas and a time constant of 30 seconds. This random process creates errors which are statistically similar to what we observed with real data. It is important to note that it should have actually been the magnetic map which was corrected, not the measurements, but correcting the map wasn’t a feasible approach. We then ran the filter using these more accurate measurements, the navigation grade INS, and the NAMAD map. Fig. 69 shows the north and east errors using a hypothetical improved-quality map. Table 7 shows the statistics of the positioning errors over the flight. With a more accurate magnetic anomaly map the filter achieved DRMS errors of about 160 meters – a large improvement over the navigation grade INS. Previously, we showed DRMS errors of 13 meters over a 1-hour flight using a high-accuracy magnetic anomaly map. The large difference in accuracy is caused by the cross-country flight flying at an altitude nearly 10 times higher than the flight test over the regional high accuracy map. As described previously, altitude is a major driving variable in magnetic anomaly navigation accuracy. Higher altitudes have shallower magnetic anomaly gradients than lower altitudes. The cross-country flight held an altitude of around 2.5–2.7 kilometers while the previous test flight flew at 300

Table 6. Navigation Accuracy Results with Tactical-Grade INS

	North Channel	East Channel
Mean	1.3 km	0.8 km
Standard Deviation	2.5 km	1.4 km
DRMS	3.2 km	
Unaided INS DRMS	34.3 km	

Table 7. Navigation Accuracy Results with Hypothetical Improved Map

	North Channel	East Channel
Mean	23.5 meters	-0.5 meters
Standard Deviation	128.9 meters	91.1 meters
DRMS	159.6 meters	
Unaided INS DRMS	1534.7 meters	

meters.

Cramer-Rao Lower Bound Analysis.

The Cramer-Rao lower bound (CRLB) is the minimum theoretical error covariance of an unbiased estimator [7], given by

$$\mathbf{P}_t \leq E \left[(\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1})(\mathbf{x}_t - \hat{\mathbf{x}}_{t|t-1})^T \right]. \quad (87)$$

For a derivation of the CRLB, see [6]. In the context of navigation filters, the CRLB is a useful tool to evaluate the performance of a navigation filter, and to further predict navigation accuracy under various conditions. In [6], the CRLB was derived for terrain aided navigation – a very useful result due to the similarities in magnetic anomaly navigation. The derived CRLB applied to the following system model:

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k + \mu_k + \mathbf{v}_k, \\ \mathbf{y}_k &= h(\mathbf{x}_k) + \mathbf{e}_k, \end{aligned} \quad (88)$$

where:

$\mathbf{x}_k$ is the position state of an aircraft

μ_k represents the dynamics of the aircraft as measured by an INS

$\mathbf{v}_k$ is dynamic noise

$\mathbf{y}_k$ is a terrain-height measurement

h_k represents interpolation of a grid of terrain height measurements using the filter's estimated position state

$\mathbf{e}_k$ is measurement noise

Starting with the definition of the CRLB, [6] derived the final CRLB for this specific navigation filtering problem as simply equal to the familiar Kalman filter covariance propagate and update equations, with the Kalman filter measurement matrix $\mathbf{H}$ equal to the gradient of $h(\mathbf{x}_k)$ evaluated at its true state:

$$\begin{aligned}\mathbf{P}_{k+1} &= \mathbf{P}_k - \mathbf{P}_k \mathbf{H}_k (\mathbf{H}_k^T \mathbf{P}_k \mathbf{H}_k + \mathbf{R}_k)^{-1} \mathbf{H}_k^T \mathbf{P}_k + \mathbf{Q}_k, \\ \mathbf{H}_k &= \frac{\delta h^T(\mathbf{x}_k)}{\delta \mathbf{x}_k}.\end{aligned}\tag{89}$$

It is simple to extend this result to the magnetic anomaly navigation system presented in [15]. Both navigation systems have linear dynamics based on an INS model and a non-linear measurement update based on a spatial grid of values. The only difference is the addition of the linear temporal variation state to the measurement, which in this context can just be brought inside the non-linear function h .

We evaluated the CRLB for magnetic anomaly navigation using the test flight data previously discussed. This required the gradient of the NAMAD map, which we computed using finite differencing. A navigation grade INS model was used along with the previously described model for the magnetic measurement error—a first order

Guass Markov process with $\sigma = 4$ nano-Teslas and $\tau = 30$ seconds. Fig. 70 shows the CRLB covariance along the flight path compared to the simulated filter covariance. The CRLB clearly follows the trend of the filter simulation, and is slightly better, as expected. This indicates that the filter is near ideal. It is clear that the CRLB is a useful tool for analyzing the accuracy of magnetic anomaly navigation. We used the CRLB to examine the trade-space that exists between accuracy and altitude. We computed the CRLB along the same flight path at 10 different altitudes ranging from the NAMAD map altitude of 305 meters to a max altitude of 9 kilometers. Fig. 71 shows the CRLB at these 10 different altitudes along the cross-country flight path. Accuracy decreases with altitude as expected. The CRLB also validates the accuracy of the first order accuracy approximation we gave previously of simply taking the measurement accuracy and dividing it by the map gradient. This simple accuracy approximation can be used to suggest how performance changes with respect to location over the United States. Fig. 72 shows the magnetic anomaly gradient over the United States at 1 kilometer altitude. There is no color scale on the plot because the contrast was stretched to provide a better illustration of the gradient. The brightest areas on the map would correspond with sub-kilometer level accuracy while the darkest areas on the plot would correspond with accuracies of several kilometers.

Cross-Country Navigation Conclusions.

In this section we successfully used real magnetometer data, as well as a widely available magnetic anomaly map of North America, to constrain the drift of a tactical grade INS to a few kilometers over a five hour flight. Furthermore, we showed that the majority of position errors in the magnetic anomaly navigation system can be attributed to poor map quality. We showed that with high quality maps navigation accuracies of around 150 meters DRMS are possible at altitudes of 3 kilometers,




Figure 69. Navigation Accuracy Using Hypothetical Improved-Quality Magnetic Anomaly Map

with accuracy increasing to as good as 10 meters DRMS when flying a few hundred meters above ground level. Our results motivate that practical implementation of high accuracy magnetic anomaly navigation can be made possible by a modern magnetic survey of the United States. Finally, we presented the Cramer-Rao lower bound for magnetic anomaly navigation and used it to predict navigation accuracy at a wide range of altitudes.




Figure 70. Comparison of CRLB and Filter Predicted Standard Deviations




Figure 71. CRLB Predicted Standard Deviation at 10 altitudes




Figure 72. Magnetic Anomaly Gradient Over United States at 1km (Contrast Stretched)

V. Continental Simulation Results

This chapter details the results of a continental simulation on magnetic anomaly navigation. Having conducted successful flight tests using magnetic anomaly navigation techniques, we now wish to predict performance in other areas we have not yet flown. To accomplish this we need to create a realistic simulation environment. The simulation is validated by comparing the predicted simulation performance to the actual flight test performance. The simulation environment works by generating a realistic INS solution as well as realistic magnetometer measurements and passing these measurements into the previously developed navigation filter. The only inputs required to run the simulation are a magnetic anomaly map over the area an aircraft will be flying as well as a trajectory. We first compared flight test performance to simulation predicted performance over Louisiana, VA and Texas. Once the simulation environment was validated we created a grid of flight lines across the continental United States to test navigation accuracy with respect to location. We also ran several simulations at varying altitudes and velocity to explore these other key variables. We first look at the performance over the state of California, as this is one of the few areas of the country with widespread accurate magnetic anomaly maps which are fully sampled and publically available. We end by presenting performance over the United States as a whole. These results are likely limited by under-sampled maps over large portions of the country which lack high spatial frequency magnetic features.

5.1 Magnetometer Errors

There are many variables which affect the accuracy of magnetic anomaly navigation. Thus far, flight test performance has only been shown over limited cases—at low altitudes over a area with high spatial variability, and at high altitudes over an

area of low spatial variability. To assist in predicting magnetic anomaly navigation anywhere in the world, the authors have developed an accurate simulation model in which magnetic anomaly navigation performance can be predicted anywhere in which a magnetic anomaly map exists. This simulation model takes into consideration all of the major variables which affect magnetic anomaly navigation accuracy. These variables include:

1. Altitude
2. Temporal Variations
3. Aircraft Effects
4. Sensor Errors
5. Map Quality
6. INS Quality
7. Magnetic Anomaly Spatial Variation

Each of these errors will be described in more detail below. The simulation model is tested against real flight test data in order to validate its predictive abilities.

The magnetic anomaly navigation simulation model's primary purpose is to generate realistic sensor measurements. Once these measurements are generated, the data is simply fed into the previously developed navigation system. Generating accurate magnetic anomaly measurements is not trivial due to the large variety of error sources which exist in the measurements. In order for the simulation model to have useful

predictive abilities, the generated measurements must truly match what would be observed in a real flight test. Although generating magnetic anomaly measurements is the focus of this chapter, we also generate an accurate INS solution which reflects the quality of the INS being utilized.

5.2 Generating Magnetometer Measurements

The simulation model developed in this paper primarily generates realistic magnetometer and INS measurements. These measurements are fed into the magnetic anomaly navigation system described in Chapter 3. The simulation requires two major inputs—a magnetic anomaly map over which the aircraft will be flying and a given trajectory. The magnetometer measurement simply starts out as the expected magnetometer measurement from the given trajectory over the magnetic anomaly map. This is achieved by a three dimensional interpolation over a stack of gridded map tiles. As explained below, this stack of map tiles can be created by a single map tile at a lower altitude through upward continuation. It is important to note that the physical properties of a magnetic anomaly field make this interpolation very accurate. Magnetic anomaly fields are continuous fields, and if fully sampled according to the Nyquist rate, the field between grid points on a map can be accurately reconstructed [10]. This is a subtle but important difference between similar navigation techniques such as terrain following, in which the signal cannot be said with any certainty to be fully sampled, meaning interpolation can lead to unknown errors. The generation of the magnetometer measurements starts with this expected magnetic anomaly from the three dimensional interpolated map tiles, and is further corrupted with the core field, temporal variations, aircraft effects, and sensor noise. We use the term “corrupted” in our context to indicate these are fields which we need to remove in order to achieve the highest navigation accuracies, as our maps are only capturing

the magnetic anomaly. Taking into account the above described error sources, we can generate magnetometer measurements using a measurement equation similar to what was given in Chapter 3. Each of the below terms is used to create artificial magnetometer measurements along a flight path. The first equation describes how a two dimensional map tile is turned into a three dimensional interpolation function which gives expected magnetic anomaly intensity at any location:

$$z_e = f_I(f_U(map), lat_t, lon_t, h_t), \quad (90)$$

where:

lat_t, lon_t, h_t are the true latitude, longitude, and height of an aircraft

z_e is the expected magnetic anomaly at (lat_t, lon_t, h_t)

map is a two dimensional map tile at height h_0

f_U is the upward continuation function which creates a vertical stack of map tiles

f_I performs a 3D interpolation on a stack of map tiles to return the expected value at (lat_t, lon_t, h_t)

The next equation shows how the expected magnetic anomaly intensity measurement is corrupted with the core field, the aircraft effects, temporal variations, and white noise. Recall real recorded temporal variations from a base station are used and that the aircraft effects are modeled by a tunable FOGM process:

$$z_c = z_e + C + T_V + A_C + w, \quad (91)$$

where:

z_m is the measured value at (lat_t, lon_t, h_t)

z_e is the expected magnetic anomaly at (lat_t, lon_t, h_t)

C is the expected core/main field measurement at (lat_t, lon_t, h_t)

T_V are the temporal variations

A_C are the aircraft effects

w is the sensor noise

The next few sections describe each of the terms from the above equations in detail.

5.3 Altitude

Unlike other map-based navigation systems such as terrain following, the magnetic anomaly field changes in three dimensions. Although it is not correct to call altitude an “error source” failing to account for changing altitudes will lead to errors. Aeromagnetic surveys are almost never completed at more than one altitude. Fortunately, magnetic anomaly fields can be upward continued [10]. In the context of the simulation model, the upward continuation process described below is used to create a three dimensional grid of points to interpolate over. The uncorrupted measurements for the simulation start as the expected values from this three dimensional interpolation.

The first step in implementing both a magnetic anomaly simulation as well as the main navigation filter is to create a stack of multiple two dimensional map tiles from upward continuing a single tile at a lower altitude. This stack of map tiles can then be interpolated in three dimensions to get an expected magnetic measurement anywhere over the map at varying altitudes. The vertical spacing for the map tiles depends on altitude but can generally be much larger than the map’s horizontal spacing. It is important to note that downward continuation appears in the geoscience literature (e.g.—[10]) but has limited application in the context of magnetic anomaly navigation because it leads to a non-unique solution [27]. Aeromagnetic surveys are

normally flown at a low altitude, so the need for downward continuation is limited.

From an algorithmic viewpoint, generating a new map at a higher altitude consists of four nested for-loops (Every value on the two dimensional grid contributes to a *single* value at a higher altitude). Fortunately, a much more computationally efficient Fourier domain equivalent expression can be derived (see [10]). The upward-continued altitude in the Fourier domain can be given by

$$\mathcal{F}[U_{z+\Delta z}] = \mathcal{F}[U_z]F[\psi]. \quad (92)$$

Where $F[\psi]$ is the upward continuation filter. The upward continuation filter attenuates shorter wavelengths more than longer wavelengths and is given by

$$F[\psi] = e^{-\Delta z|k|}, \quad (93)$$

where $|k|$, in this context, is equal to the absolute value of the wavenumber in the frequency domain and is simply

$$|k| = \sqrt{k_x^2 + k_y^2}, \quad (94)$$

where k_x and k_y are the x and y direction wavenumbers, respectively. A wavenumber is simply the spatial domain analog of a time frequency.

The upward continuation filter is real valued and therefore does not affect the phase of the map data. Fig. 73 shows the magnitude of the upward continuation filter in the Fourier domain for a upward continuation height of both 1 kilometer and 100 meters with spatial frequencies ranging from 0–1 Hertz (cycles/kilometer). Taking the




Figure 73. Upward Continuation Filter Spectrums for 100x100 km map

inverse Fourier transform of $\mathcal{F}[U_{z+\Delta z}]$ is the final step needed to generate a map at a higher altitude. There are two key assumptions in the upward continuation that are not ever met in practice. The first is that an infinite horizontal map tile is available. Obviously, real map tiles will be finite and violate this assumption. The second assumption which is violated in practice is that all magnetic sources lie below the map tile. The total magnetic field measured in the vicinity of the Earth is a superposition of many magnetic sources. While the largest of these sources are internal to the Earth, there are external sources which will be above a map tile as well. It is important to note that the upward continuation filter is essentially a low pass filter. This means that the higher we upward continue data, the more the high frequency features will be attenuated. This is important from a navigation standpoint, because it is these high frequency features which contribute to high accuracy navigation. Intuitively, then, navigation accuracy will decrease at higher altitudes. Fig. 74 shows the NAMAD at its base altitude of 300 meters above terrain as well as at five kilometers. Due to the errors in upward continuation, it is important to ensure a map tile is approximately 50 times larger in the horizontal direction than the desired upward continuation height. This empirically derived ratio will ensure the errors caused by upward continuation are limited to approximately the outside 10% of the map. In general, the outermost edges of an upward continued map tile will not be accurate.

5.4 Core Field Effects

Magnetic anomaly maps, by definition, are the deviation from a model of the Earth's main magnetic field, also known as the core field. This is the overall approximately dipole magnetic field which makes a compass point north. This magnetic field has little spatial variation and changes over time, making it a poor candidate for a map-based navigation system. When generating realistic magnetometer mea-




Figure 74. Upward Continuation of North American Magnetic Anomaly Database

surements, the core field must be included. Because a magnetic anomaly map shows the deviation from a core field model, the same core field model used to make the map should be used by the navigation system. Many magnetic anomaly maps are made using the International Geomagnetic Reference Field model [63]. This model is updated every five years by an international team of scientists and includes time derivative terms to account for how the core field is drifting over time. If the IGRF was used to make a magnetic anomaly map, the IGRF should either be subtracted from a magnetometer measurement or added back to the magnetic anomaly map to allow magnetic measurements to match a magnetic anomaly map. For our simulation we simply evaluated the most current IGRF model along our trajectories. The inputs to this model are latitude, longitude, altitude, and time.

5.5 Aircraft Effects

The magnetometer measurements must also be corrupted with aircraft effects. Currently, there exists many compensation algorithms to remove aircraft magnetic




Figure 75. Compensation of Aircraft Effects on Geo-Survey Aircraft [23]

fields, such as the Tolles-Lawson model [66] described in Chapter 2. For the results shown in Chapters 4.1-4.2, the Tolles-Lawson model was used to provide real time compensation of the magnetometer measurements. We are primarily concerned with what these algorithms *cannot* remove, because this is what we will use to corrupt our measurements. Our corrupted measurements, therefore, represent *post-compensated* measurements. The magnitude of the aircraft effects is best derived empirically for a given aircraft. The flight tests described in Chapter 4 utilized a geo-survey aircraft which was modified to minimize aircraft effects, including placing the magnetometer out on a boom. In this type of environment, aircraft effects are quite small. Fig. 75 shows the magnitude of these aircraft effects on a geo-survey aircraft. It can be seen the aircraft effects were reduced to a fraction of a nano-Tesla from the original oscillations of nearly 10 nano-Tesla's. It is important to realize that these errors could be much larger on a less ideal platform.

We chose a simple model to represent the non-removable components of the aircraft effects. These errors are modeled as a stochastic first-order Gauss Markov (FOGM) process. A FOGM process is characterized by an exponential autocorrela-




Figure 76. Non-removable aircraft effects for simulation measurement corruption

tion. It can be fully described by two parameters - a variance σ and a time constant τ . The process autocorrelation is then

$$\mathbf{R}_x = \sigma^2 e^{-\beta|\tau|}. \quad (95)$$

The two parameters σ and τ can be varied to represent the non-removable aircraft effects for a given platform. For the trade-space simulations conducted below, we chose a σ value of 1 nano-Tesla and a τ value of 1000 seconds. Fig. 76 shows an example of generated aircraft effect errors using these specific parameters.

5.6 Temporal Variations

Magnetic fields from the Earth's ionosphere and magnetosphere, referred to as temporal variations, will corrupt magnetometer measurements. To create a realistic simulation, magnetometer measurements must be corrupted with these temporal variations. Rather than model these errors, we instead corrupt the measurements with actual recorded temporal variation data. All over the world there exist stationary magnetometers monitoring these temporal variations. This data is easily accessible




Figure 77. Temporal Variations From 10 Different Flight Segments, Taken From Real Temporal Variations

from organizations such as the USGS [1]. We took a year’s worth of temporal variation data from a Colorado base station and corrupted our measurements with random segments of this data. This ensures that our measurements are corrupted as realistically as possible with the temporal variation data. Fig. 77 shows temporal variation data from 10 different flight segments.

5.7 Sensor Errors

Sensor errors are not currently a dominant error source with magnetic anomaly navigation. It is important to note the distinction between measurement errors and sensor errors. Because we are navigating with a subset of the total measured magnetic field, our ability to remove the unwanted components such as the aircraft fields and the temporal variations are what drive our measurement errors. Even in an ideal-case scenario we can only remove these effects to approximately 1 nano-Tesla. The sensor errors, however, are much smaller than this. The sensitivities of the optically pumped cesium magnetometers we used are on the order of pico-Teslas. The bias stabilities are

also well below one nano-Tesla [60]. This indicates that the sensor errors will likely not be a driving term. Nonetheless, for completeness we corrupt the measurements with a white noise error source. This white noise $\mathbf{w}$ is zero-mean random process with a variance σ^2 .

$$E[\mathbf{w}] = 0, \quad (96)$$

$$E[\mathbf{w}^2] = \sigma^2. \quad (97)$$

We choose a σ of 100 pico-Teslas for our simulations.

5.8 Map Quality

Map quality is a major variable affecting navigation accuracy in a magnetic anomaly navigation system. Determining the accuracy of a magnetic anomaly map is a difficult problem. Magnetic anomaly maps have been created all over the world over the last 70 years. Before the widespread use of GPS around 1990, the magnetic sensor readings were geo-located using relatively crude aerial photography methods. Although there have been improvements in actual sensor accuracy over time, the pre-GPS era maps are almost certain to be filled with inaccuracies due to poor geo-location. It is difficult to distill the accuracy of these maps down to a simple number, e.g., “The maps are accurate to within 10 nano-Teslas”. When reading the description of old map tiles it is not uncommon to find accuracy claims around 10 nano-Teslas. It may be helpful to keep a level of skepticism regarding these claims for pre-GPS era maps especially with regard to the exact positioning of the map. We did not include the modeling of map geolocation errors in this simulation. These types of errors are very difficult to characterize and any modern survey would not be subject to them. A map with large geolocation errors is not appropriate for high accuracy navigation. This is fairly intuitive—if the map itself is not geo-located to better than

one kilometer, for example, the 10's of meter accuracy we have demonstrated with a high-quality map is likely not possible.

In many of the large trade-space simulations in this chapter we use the North American Magnetic Anomaly Database map [3] to make navigation accuracy predictions. This is a compilation map made from many separate aeromagnetic surveys. We use this map to help show regional trends in expected magnetic anomaly navigation accuracy but it is important to point out that this map itself would not be very useful for navigation as many of the contributing individual surveys are decades old. The map is useful for representing the overall spatial variation in a given area but leads to poor quality navigation results as demonstrated in Chapter 4.2. The geolocation errors in certain areas are simply too large to achieve high accuracy navigation. Consequently, the simulation model provided will not reflect real navigation performance if a poor quality map is used.

It is important to note the varying line spacings which contributed to the NAMAD map. As mentioned previously, this map was used to determine expected navigation accuracy on a regional level. Fig. 78 shows the individual component surveys which were used to create the map over the United States, as well as the line spacing of each survey. It can be seen in Fig. 78 that a large majority of the surveys on the east and west sides of the country were flown at line spacing of 2 kilometers or less, and the central part of the country was flown at 4–8 kilometers or less. The map is technically at an altitude of 305 meters but it very under-sampled at this altitude due to the large line spacings (see Chapter 2). This indicates that the actual navigation accuracy will be better than predicted, as the current map used to make the regional accuracy predictions is missing many of the high frequency wavelengths which contribute to




Figure 78. Line Spacing of North American Magnetic Anomaly Database

high navigation accuracy.

5.9 Generating INS Measurements

The developed magnetic anomaly navigation system requires the use of an inertial navigation system. Consequently, the simulation environment must generate realistic INS measurements. The magnetic anomaly navigation filter is loosely coupled with the INS. No corrections are ever made to the INS, and only the position of the INS is used by the filter during map-matching updates—the actual INS mechanization

happens entirely independently from the filter. Because of this, instead of generating delta-V and delta-theta measurements to model an INS, we simply directly generate INS errors using a 15 state INS error model which includes position, velocity, attitude, and accelerometer and gyroscope bias error states. These errors are added to a true trajectory to obtain the simulated INS trajectory. We can change the quality of the INS by simply changing the velocity and angular random walks of the INS model as well as the magnitude of the accelerometer and gyroscope biases. While not truly simulating an INS, this simple method creates realistic INS trajectories, which are the only outputs from the INS which are actually needed by the magnetic anomaly navigation filter.

The INS model used is commonly known as the Pinson error model. A full derivation of the model is given in [65]. The Pinson error model includes two main components—a first-order dynamics model describes how INS errors propagate over time, and a noise model describes how noise enters the system, causing the INS to drift over time. The continuous model is given by

$$\dot{\mathbf{x}} = \mathbf{F}\mathbf{x} + \mathbf{w}, \quad (98)$$

where $\mathbf{F}$ describes the continuous dynamics and $\mathbf{w}$ described the platform noise.

The first 9 states of the model describe the errors in INS computed position, velocity, and attitude. The attitude errors are captured by INS orientation tilt errors. The final 6 states describe the IMU sensor biases in both the accelerometers and the

gyroscopes. The full state vector is

$$\mathbf{x} = \begin{bmatrix} \Delta lat \\ \Delta lon \\ \Delta alt \\ \Delta v_n \\ \Delta v_e \\ \Delta v_d \\ \Delta \epsilon_x \\ \Delta \epsilon_y \\ \Delta \epsilon_z \\ bias_a^x \\ bias_a^y \\ bias_a^z \\ bias_g^x \\ bias_g^y \\ bias_g^z \end{bmatrix}. \quad (99)$$

Dynamics Model.

The INS dynamics model is given by

$$F = \begin{bmatrix} \mathbf{N}_{9 \times 9} & \mathbf{A}_{9 \times 6} \\ 0_{6 \times 9} & \mathbf{B}_{6 \times 6} \end{bmatrix} \quad (100)$$

There are 3 major sub-matrices defined. $\mathbf{N}_{9 \times 9}$ describes how the INS errors in position, velocity, and attitude propagate. $\mathbf{B}_{6 \times 6}$ describes how the IMU biases propagate. Finally, $\mathbf{A}_{9 \times 6}$ describes how the accelerometer and gyroscope bias errors propagate into the velocity and attitude states respectively [65].

$$\mathbf{N}_{9 \times 9} = \begin{bmatrix}
0 & 0 & 0 \\
\frac{v_e^2 \tan L - v_n v_d}{r_e^2} & \frac{v_d}{r_e} & -2 \left(\omega \sin L + \frac{v_e \tan L}{r_e} \right) \\
-v_e \frac{v_n \tan L + v_d}{r_e^2} & 2\omega \sin L + \frac{v_e \tan L}{r_e} & \frac{v_n \tan L + v_d}{r_e} \\
\frac{v_n^2 + v_e^2}{r_e^2} & -2 \frac{v_n}{r_e} & -2 \left(\omega \cos L - \frac{v_e}{r_e} \right) \\
\frac{-v_e^2}{r_e^2} & 0 & \frac{1}{r_e} \\
\frac{v_n}{r_e^2} & \frac{-1}{r_e} & 0 \\
\frac{v_e \tan L}{r_e^2} & 0 & \frac{-\tan L}{r_e} \\
-1 & 0 & 0 & 0 \\
\frac{v_n}{r_e} & 0 & -f_d & f_e \\
2\omega \cos L + \frac{v_e}{r_e} & f_d & 0 & -f_n \\
0 & -f_e & f_n & 0 \\
0 & 0 & -\omega \sin L - \frac{v_e \tan L}{r_e} & \frac{v_n}{r_e} \\
0 & \omega \sin L + \frac{v_e \tan L}{r_e} & 0 & \omega \cos L + \frac{v_e}{r_e} \\
0 & \frac{-v_n}{r_e} & -\omega \cos L - \frac{v_e}{r_e} & 0
\end{bmatrix}_{7 \times 7} \quad (101)$$

where:

L is latitude in radians

r_e is the Earth's radius equal to 6378135 meters

ω is the Earth's rotation rate equal to $7.2921151467 \times 10^{-5}$ rad/sec

f_n , f_e , and f_d are the north, east, and down specific forces

v_n , v_e , v_d are the north, east, and down velocities

The sub-matrix $\mathbf{B}$ describes the model for the sensor biases. In this case, the sensor

biases are modeled as first order Gauss Markov processes. The matrix is given by

$$\mathbf{B} = \begin{bmatrix} \frac{-1}{\tau_a} & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{-1}{\tau_a} & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{-1}{\tau_a} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{-1}{\tau_g} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{-1}{\tau_g} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{-1}{\tau_g} \end{bmatrix}, \quad (102)$$

where:

τ_a is the accelerometer bias time constant

τ_b is the gyroscope bias time constant

The sub-matrix $\mathbf{A}$ is given by

$$\mathbf{A} = \begin{bmatrix} 0_{3 \times 3} & 0_{3 \times 3} \\ \mathbf{C}_{\mathbf{b}}^{\mathbf{n}} & 0_{3 \times 3} \\ 0_{3 \times 3} & \mathbf{C}_{\mathbf{b}}^{\mathbf{n}} \end{bmatrix}, \quad (103)$$

where $\mathbf{C}_{\mathbf{b}}^{\mathbf{n}}$ is the direction cosine matrix which converts from the body frame to the navigation frame.

Noise Model.

The INS noise model includes four main noise sources. The first two are velocity and angular random walks. These are white noise terms that are simply added to the velocity and attitude states and consequently integrated into larger errors. The second two noise terms are the bias driving noise terms. These are the first order Gauss Markov white driving noises needed in the sensor bias models. For simplicity

we will assume that each accelerometer axis and each gyroscope axis has the same magnitude VRW and ARW terms. The same simplification is made for the sensor biases—the accelerometer biases in each axis have the same noise strength and the gyroscope biases in each axis have the same noise strength. The noise model is given by

$$E[\mathbf{w}(t)\mathbf{w}^T(t+\tau)]\delta(\tau) = \mathbf{Q} = \begin{bmatrix} 0_{3 \times 1} \\ VRW_{3 \times 1} \\ ARW_{3 \times 1} \\ a_{3 \times 1}^{bias} \\ g_{3 \times 1}^{bias} \end{bmatrix}. \quad (104)$$

Where the driving noise terms are computed from the first order gauss Markov variance σ and time constant τ according to $q = \frac{2\sigma^2}{\tau}$.

5.10 Simulation Validation

In the following sections we will compare simulation predicted performance with the actual achieved performance using real collected data. The first flight test occurred over Louisa, VA. The flight test lasted about an hour and consisted of flying large figure-8's over a high quality magnetic anomaly map. When running the navigation filter with real magnetometer and INS data the filter achieved DRMS errors of approximately 10 meters over an hour long flight. As stated before, the only inputs to the simulation will be the true trajectory the aircraft flew as well as the magnetic anomaly map. There are two major sets of parameters that need to be defined. The first set of parameters are the statistics for the FOGM process which describe the aircraft induced magnetometer measurement errors. The second set of parameters describe the performance of the inertial navigation system. Recall that the measurements are corrupted with a random segment of real recorded temporal variation data.

Table 8. Flight Test 1 Simulation Parameters

Parameter	Value
Aircraft Effects Variance	1 nT
Aircraft Effects Time Constant	300 seconds
INS Velocity Random Walk	$(5 \times 10^{-5} m/s^2 \sqrt{Hz})^2$
INS Angular Random Walk	$(6 \times 10^{-2} deg/h \sqrt{Hz})^2$
INS Accelerometer Bias Variance	$(2 \times 10^{-4} m/s^2)^2$
INS Gyro Bias Variance	$(3 \times 10^{-3} deg/h)^2$
INS Bias Time Constants	3600 seconds

Table 8 shows the simulation parameters for the first flight test.

As can be seen from Table 8, the aircraft effects are modeled as being around 1 nano-Tesla in magnitude. This is worse than a high quality geo-survey aircraft can achieve but also helps to capture any un-modeled effects. The INS model represents a navigation grade INS. The modeled INS will drift a few kilometers on average in an hour. In this case, the temporal variations dominate the long wavelength errors, and the aircraft effects dominate the short wavelength errors. Fig. 79 shows a comparison of the horizontal position DRMS errors for the simulation results and the actual flight test results. It is clear the results match up very closely. This lends credence to the simulation model being useful for predicting performance for other trajectories over different magnetic anomaly maps. It is also worth noting that this results in general should be considered “best case” for magnetic anomaly navigation. Flying at low altitudes is a major contributor to the high accuracy for this flight test. The magnetic anomaly field has high spatial variation at low altitudes. We will use the exact same parameters for the second flight test scenario to demonstrate the simulation model is not being “tuned” to match the actual results.

The second flight test occurred over central Texas. This test flight encompassed




Figure 79. Comparison of Simulated and Actual Flight Test Error Standard Deviations over Virginia

three separate flights at various increasing altitudes. It is important to note that the spatial variation in the magnetic anomaly field over this area of Texas is much smaller than the spatial variation from the previous test flight over Louisa, VA. This test flight also flew at altitudes up to three kilometers. This further reduces the spatial variation in the map. Consequently, lower navigation accuracies are expected for these flight tests. The same simulation parameters from Table 8 are used to model the aircraft effects and INS quality. Fig. 80 shows the results of the second flight test at the various altitudes. As expected, the navigation accuracies are lower than the previous flight test, but fortunately the simulation accurately predicts this lower performance. While there are some differences between the simulated and the real data results, the overall magnitudes and trends appear to be largely correct.

5.11 Simulation Tradespace Predictions

The NAMAD was used to perform several trade-space studies demonstrating magnetic anomaly navigation performance over the United States with respect to key




Figure 80. Comparison of Simulated and Actual Flight Test Error Standard Deviations over Texas

variables. It is important to understand that this does *not* mean actually using the NAMAD map to navigate will achieve these performances. In the simulations we assume the NAMAD map has no map errors. This is far from accurate in certain areas of the country due to poor geolocation and under-sampling errors. The simulation results are still useful, however, to show regional trends in expected magnetic anomaly navigation performance if a good map were to exist. The key variables explored in the simulation are velocity, altitude, and location. A large simulation was conducted in which simulated aircraft trajectories create a grid over the United States, as shown by the red lines in Fig. 81. The DRMS accuracies along the flight path were then used to calculate navigation error probability distributions for flying over California and the continental United States. One issue with this type of analysis is the varying sample spacing of the contributing magnetic anomaly maps to the NAMAD map. Based on the line spacing map provided in Fig. 78, California has been surveyed with much better line spacing than the majority of the country. This is beneficial, because high spatial frequency variations improve navigation performance. Many of the con-

Table 9. California Navigation Accuracy Tradespace Statistics

Parameters (altitude/velocity)	DRMS	95% CDF
300 m / 200 m/s	10 m	54 m
300 m / 100 m/s	16 m	87 m
3 km / 200 m/s	32 m	146 m
3 km / 100 m/s	59 m	289 m
10 km / 200 m/s	250 m	919 m

tributing maps from the Midwest have large sample spacings, indicating these areas are potentially under-sampled. This indicates that the spatial frequency content of these maps are fundamentally limited by this large sample spacing. This limited frequency information will degrade navigation accuracy in these areas, but this is a function of an under-sampled map, not the true magnetic field which exists in these areas. Due to this limitation, results are first presented over the state of California, where line spacing is not a large issue. Given the large size and varying geological terrain of California, these statistics are likely a better measure of predicted navigation performance than a simulation over the entire US map, which includes large under-sampled areas. Using our simulation framework we explored navigation accuracy at several different velocities and altitudes. The error statistics from these simulations are shown in Table 9. This table shows the error statistics for five different simulation scenarios. The first two represent flying at 300 meters above terrain at 100 meters/second and 200 meters/second. The second two show flying 3 kilometers above terrain at 100 meters/second and 200 meters/second. The final scenario represents flying at 10 kilometers altitude at 200 meters/second. The DRMS of the data is given as well as a 95% column, which can be read as “95% of errors are less than x ”.

In addition to calculating the navigation error statistics over CA, we also wished to conduct an analysis over the entire United States as well. Again, it is important




Figure 81. Simulated Trajectories Over the United States

Table 10. US Navigation Accuracy Tradespace Statistics

Parameters (altitude/velocity)	DRMS	95% CDF
300 m / 200 m/s	19 m	159 m
300 m / 100 m/s	28 m	243 m
3 km / 200 m/s	52 m	349 m
3 km / 100 m/s	82 m	560 m
10 km / 200 m/s	276 m	1342 m

to mention that many of the large errors encountered are likely a function of low map spatial frequency information caused by under-sampling. Using the data from the flight lines over the United States we created navigation accuracy maps which show navigation errors as a function of position. The navigation accuracy maps are low pass filtered to better show the regional trends throughout the United States. Fig. 82 shows the simulation scenario which led to the highest accuracies and serves as a baseline to compare the other scenarios. The simulation scenario used to make this map had an aircraft fly at a velocity of 200 meters/second (450 miles/hour) at an altitude of 300 meters, using random temporal variations. Generally, this would not be a practical velocity at such a low altitude but we present it to show the “best-case” results and as a reference for the other scenarios. Finally, Fig. 84 shows another US map plot for the simulation scenario of flying at 3 kilometers altitude at 200 meter-s/second for comparison to the previously shown baseline map in Fig. 82. Note that the color scale for this map is different between Fig. 82 and 84.




Figure 82. Expected DRMS Accuracy Over the United States at 200 m/s at 300 m Altitude

There are several conclusions which can be drawn from the simulation results. The first is the importance of INS quality. There is a large difference between the DRMS and 95% values for the Cumulative Distribution Function (CDF). This indicates that occasionally the navigation solution drifts much further than the DRMS solution. This likely happens while flying through areas of low spatial variability. This indicates that INS quality is a very important variable for this type of navigation system. A better quality INS allows the navigation system to “coast” through the areas of low spatial variability until an accurate position fix can be obtained by flying over an area of high spatial variability. We predict a more accurate INS would reduce the difference between the median and 95% CDF values. The next observation relates to the NAMAD map line spacing. It can be seen from the navigation accuracy maps over the United States that there are clear areas over the country in which navigation accuracy is degraded. The degraded areas of the map appear correlated with the line spacings of the individual surveys which contributed to the NAMAD map, shown in Fig. 78. This indicates that if a better quality map were available which was fully sampled, the large errors over these areas may diminish. The errors could also be caused, however, by a true lack of spatial variation in the maps over

these parts of the United States. If these areas of the map are truly under-sampled, navigation accuracy would be expected to increase significantly if a fully sampled map was obtained. It is clear the navigation accuracies are much better over the state of California. Again, this could be caused by a true increase in overall spatial variability in CA. The alternative is that it is simply a more fully sampled map, and areas of the country which are currently under-sampled in the NAMAD map will actually have higher accuracies than our simulations are currently predicting.

5.12 Conclusions

This chapter presented a simulation tool to predict navigation accuracy anywhere a magnetic anomaly map is available. The method presented involves generating a realistic INS solution based on a given trajectory and realistic magnetometer measurements corrupted with temporal variations, aircraft fields, and sensor errors. These realistic measurements are then provided to the author's previously designed magnetic anomaly navigation filter to provide predicted navigation performance along the given trajectory. The simulation solutions were evaluated by comparing predicted performance to actual flight test results which provided real magnetometer measurements and INS data. After the simulation results were shown to be in agreement with the test flight results, the simulation was used to perform a trade-space study over the United States using the North American Magnetic Anomaly Database. Navigation performance was predicted at various altitudes and velocities over the United States in order to gain insight into expected magnetic anomaly navigation performance.




Figure 83. Navigation Error Cumulative Distributions




Figure 84. Expected DRMS Accuracy Over the United States at 200 m/s at 3000 m Altitude

VI. Self-Building Magnetic Field World Model

There are several goals for a self-building world magnetic anomaly model. The first goal pertains to the collection of data for the model. To be “self-building” the model must be able to incorporate random flight data from aircraft equipped with magnetometers. Magnetic anomaly maps are normally created by careful survey procedures, including flying in a grid-pattern at a set altitude to fully sample the field. A self-building model must be able to incorporate the addition of random flight lines as they are flown. The second goal pertains to the optimal use of the sparse data being collected. As a counter-example, consider the trivial approach of simply updating a large three dimensional grid with data as it is being collected, perhaps using a nearest neighbors approach to assign the collected data to the three dimensional grid points. It would take a tremendous number of random flights through this three dimensional grid to begin to populate it enough for useful interpolation. The reason this approach is trivial is because it does not take advantage of the properties of the magnetic field. The final goal of a self-building model is to include an uncertainty term. Consider an aircraft flying the same route as a previous aircraft, and attempting to navigate using the magnetic anomaly model. The navigation system in this case should have a high degree of trust in the model. Alternatively, if flying somewhere for the first time, the magnetic anomaly map should be trusted to a lesser extent.

There is a great deal of existing literature regarding the modeling of magnetic fields [34][10]. Specifically, the modeling of magnetic anomaly fields (caused by crustal sources) is a well understood science in the field of geophysics. Magnetic anomaly fields are routinely modeled for geological studies as well as for commercial purposes such as oil and mineral exploration. This type of modeling can be seen as an inverse problem, in which a set of observations is used to solve for a magnetic source distri-

bution. There are two separate goals in this type of modeling. The first is to model the correct distributions of the source materials, and the second is to model the resulting field [72]. The magnetic distributions which can lead to an observed magnetic anomaly field are not unique—there are an infinite number of source distributions which can lead to an observed field. Solving for the correct distributions is therefore a heavily user-involved process in which an experienced modeler uses geological information from other sources to aid in solving for the distribution [27]. This type of modeling is clearly not applicable to a self-building model.

The second type of modeling only attempts to model the observed field. In this case, solving for the correct distribution is not important, as any of the infinite distributions that cause the field are an equally valid representation of the field. This type of modeling is used when transformations of a magnetic anomaly field are needed, such as predicting the field at higher altitudes or calculating what the field would look like with a zero degree inclination inducing field (useful in geological studies). ESDI techniques are a common way to accomplish this type of modeling [72]. ESDI seeks to represent the observed field as the combined effect of a set of magnetic dipoles placed at or below ground level. The ESDI inversion problem solves for the magnetic susceptibilities of these magnetic dipoles, which are often simply placed on a grid. Furthermore, due to the superposition principle of these magnetic dipoles, additional dipoles can always be added to improve the fit of a model, a process called bootstrapping [74]. At first glance, this process sounds appealing to the problem at hand, and was in fact the first method the author pursued to solve the problem of a self-building world model. An ESDI model would be able to incorporate measurements from any altitude and be computationally efficient when using a bootstrap approach. In practice however, this turned out to be very difficult. Inversion of potential field

data is notoriously unstable. Common ways to address this instability involve using a damping factor with ridge regression techniques [74]. Choosing a damping factor involves making a tradeoff between stability and how closely the solution matches the observations. Determining these damping factors can end up being a heavily user-involved approach. The instability of this type of inverse modeling, which is actually the best physical representation of the field, was abandoned due the instability issues encountered.

The existing literature of magnetic anomaly modeling suggests another approach to the self-building world model, a process called least-squares collocation in geodetic applications and kriging in mining and geological applications [27]. Least-squares collocation is often used in the gridding of magnetic anomaly data, in which the magnetic field value at a given point is the weighted average of nearby observations. In [73], three dimensional collocation was used to grid magnetic data from various altitudes to a common altitude, and shown to be much faster and computationally efficient than ESDI approaches.

More generally, kriging and least-squares collocation can be described as a Gaussian process regression (GPR) technique. GPR is a non-parametric regression approach useful when there exists no obvious parameter based model (linear, polynomial, etc). GPR was used in [70] to model the two dimensional magnetic vector field measured inside a building for a Simultaneous Localization and Mapping (SLAM) robot. The authors used a squared exponential covariance model, which guarantees smoothness, and used pre-computed hyper-parameters to describe the squared exponential covariance model. The GPR techniques lead to accurate two dimensional models of the magnetic field in a SLAM based setting. As pointed out in [75], these

type of models do not take advantage of the physics of magnetic fields, namely that they are curl-free and divergence free. [75] derives a Gaussian process model which obeys Maxwell’s equations, and shows the ability for a set of observations to model both a full magnetic vector field as well as the location of the actual magnetic sources. This model is appealing because a self-building world model will have sparse measurements, so the more representative the model is of reality, the more information can be obtained from these sparse measurements. Unfortunately, this model requires observations of the magnetic *vector* field, whereas the high accuracy magnetic anomaly navigation we have previously demonstrated relied on magnetic *intensity* fields. The magnetic intensity measurements are used because the instruments are 1–2 orders of magnitude more accurate than vector sensors. The one property of the magnetic fields we can capture with a Gaussian process is their guaranteed smoothness, assured because the magnetization at aircraft altitudes is assumed zero, and consequently the magnetic field is infinitely continuously differentiable [75].

Based on the proven success of GPR in both magnetic anomaly gridding as well as indoor SLAM applications, we choose a GPR method as the basis for creating a self-building world model. Magnetic anomaly mapping has been undertaken all over the world for decades. There exist several large-scale compilation maps, including the NAMAD at 300 meters altitude as well as the WMM at 5 kilometers altitude [3][43]. Because of the sparse nature of the observations available in a self-building world magnetic anomaly map, we choose to not throw out decades of information from these surveys, but rather correct the errors in the existing surveys. We demonstrate the usefulness of this approach in Table 11. The standard deviation of two high resolution surveys is compared to the standard deviation of the errors between these same high resolution surveys and their NAMAD counterparts. It appears there

is less information to capture in the error maps than in the maps themselves, so these error maps are what we will actually try to model.

The remainder of this chapter is organized as follows. In Chapter 6.1 we characterize the errors in existing magnetic anomaly maps. In Chapter 6.2 we present a GPR model which uses the sum of two anisotropic squared-exponential covariance functions to correct errors in existing magnetic anomaly maps, and discuss the significance of the model’s hyper-parameters. In Chapter 6.3 we discuss the selection of GPR hyper-parameters. In Chapter 6.4 we discuss the challenges of mapping in 3 dimensions and present line-filters for upward and downward continuation. Chapter 6.5 presents the results of using a small series of flight lines to correct errors in the NAMAD over two areas in which we also have high resolution surveys as truth data. Chapter 6.6 discusses the practical issues of implementing a self-building world magnetic anomaly model. Finally, Chapter 6.7 shows an example of how a cross-country navigation flight can obtain a more accurate navigation solution using only a small number of flight lines and the GPR method.

6.1 Characterizing Errors in Existing Magnetic Anomaly Maps

The first step to fitting a Gaussian process regression model to the NAMAD error is to understand the types of errors present in the NAMAD, and any other large compilation type magnetic anomaly map. Interestingly, sensor accuracy is one of the smallest contributors of error in the NAMAD map. Most of the magnetic surveys in the NAMAD were created after 1960 [3]. Before 1960, flux-gate magnetometers were primarily used for magnetic surveys. These instruments were indeed much less accurate than current magnetometers. However, starting in 1960 proton precession magnetometers started being used for aero-magnetic surveys and were soon replaced

Table 11. Information Comparison for True High Resolution Maps and Errors Between High Resolution Maps and NAMAD

	STD (nT)
Virginia High Resolution Survey	106.9
Error between HiRes VA map and NAMAD	88.4
Texas High Resolution Survey	121.58
Error between HiRes TX Map and NAMAD	3.2

by cesium vapor magnetometers [27]. Both proton precession and cesium vapor magnetometers have absolute errors which are trivial compared to the other errors sources in older magnetic anomaly maps.

The most dominant error in old magnetic anomaly maps is likely poor geo-location. Before the GPS came online, aero-magnetic survey data were geo-located using crude navigation methods. While the magnetic measurements themselves may have been very accurate, they could be geo-located with errors on the order of a kilometer. It is easy to see that these types of errors are likely to be spatially correlated. This is advantageous because GPR models use spatial correlation to help predict values at locations without observations. The other main type of error seen in the NAMAD is caused by under sampling, as described in Chapter 2. The NAMAD map exists at 300 meters above ground level (AGL) but many of the surveys used line spacings far greater than 300 meters, sometimes up to 5 kilometers or more [3]. This means there will be many errors in the NAMAD maps which have spatial-correlations approximately equal to the map altitude. This also indicates that when flying at higher altitudes, for example 5 kilometers, these type of errors will be mostly filtered out. At 5 kilometers altitude the original under sampled surveys become fully sampled. This indicates correcting errors in the NAMAD will be far easier at high altitudes, where under-sampling errors are not present

There are two other common errors sources in the NAMAD. Because the NAMAD is a compilation map made from many smaller maps, there isnt any way to capture long-wavelength crustal sources. Furthermore, these individual surveys were often adjusted to some arbitrary constant level. These type of errors will clearly have large spatial correlations, which means they will be relatively easy to correct. There are many remaining types of errors in the NAMAD map, but these errors are likely far outweighed by the poor geolocation and under sampling which is present in the map. Some of these errors include poor aircraft compensation, poor removal of diurnal variations, and hand-contouring of magnetic anomaly data from very old maps.

6.2 Gaussian Process Regression Models

We choose to use a Gaussian process regression model for the implementation of a self-building world magnetic anomaly model. Gaussian processes are a family of statistical distributions which extend multivariate Gaussian distributions to infinite dimensionality. A set of N observations, in this case a set of N measurements along multiple flight lines, can be thought of as a single draw from a N -variate Gaussian distribution. Formally, the Gaussian process can be described as a distribution over functions such that

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), K(\mathbf{x}, \mathbf{x}')) \quad (105)$$

where the Gaussian process is defined by a mean function $\mu(\mathbf{x})$ and by its covariance function $K(\mathbf{x}, \mathbf{x}')$. The variables $\mathbf{x}$ and $\mathbf{x}'$ are used to denote different observations, which are also two different dimensions of the N -variate Gaussian distribution, and therefore have a correlation described by K . Gaussian processes are well suited to modeling spatially correlated measurements, because each of the N observations is related to each other through the covariance matrix $K(\mathbf{x}, \mathbf{x}')$. Often Gaussian pro-

cesses are assumed to be zero mean to simplify calculations. In this case, the Gaussian process is fully defined by its covariance matrix K . Different choices of the covariance matrix K can be used to model vastly different types of problems. As stated previously, we wish to take advantage of the guaranteed infinite differentiability (i.e., smoothness) of a magnetic anomaly field and therefore use the squared exponential covariance function

$$K(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left[\frac{-|\mathbf{x} - \mathbf{x}'|_2^2}{2l^2} \right], \quad (106)$$

where:

σ_f is the expected amplitude of the signal

$\mathbf{x}$ is a single pair of points, $[x_1, x_2]'$, which could be the North and East directions in a local navigation frame

$\mathbf{x}'$ is a second pair of points

$|\mathbf{x} - \mathbf{x}'|_2$ denotes the L_2 norm, or Euclidian distance between two points

l is the length scale parameter which captures the effects of the correlation distance

There are several important properties of the standard squared exponential function. The first property, already mentioned, is its smoothness. The second is the fact that it is isotropic. This means that the covariance depends only on the distance between points, not on the direction. On a two dimensional contour map, this indicates the covariance function would look like a circle. We wish to modify the covariance function so that the north and east directions on our map can have different correlation lengths. We can write the squared exponential in its anisotropic form as

$$K(\mathbf{x}, \mathbf{x}^T) = \sigma_f^2 \exp \left[- \left(\frac{(x_1 - x_1^T)^2}{2l_1^2} + \frac{(x_2 - x_2^T)^2}{2l_2^2} \right) \right], \quad (107)$$

where:

$$\mathbf{x} = [x_1, x_2]^T \text{ and } \mathbf{x}^T = [x_1^T, x_2^T]^T$$

l_1 is the x_1 length scale parameter (i.e., north direction)

l_2 is the x_2 length scale parameter (i.e., east direction)

σ_f is the square root of the variance parameter

Covariance functions may be added together to capture additional behavior of a modeled process. We know from the analysis of errors in magnetic anomaly maps that we would expect two major correlation lengths in the data. The first would capture the effects of having an under sampled map, and would be approximately equal to the altitude above terrain. The second would capture the long wavelength, poor geo-location, or constant bias type errors expected in the NAMAD maps. We therefore make our final covariance function equal to the sum of two anisotropic squared exponential functions. Additionally, with a slight abuse of notation for brevity, we add an additional noise source to capture our measurement uncertainty and stabilize the covariance function:

$$K(\mathbf{x}, \mathbf{x}^T) = K_1(\mathbf{x}, \mathbf{x}^T) + K_2(\mathbf{x}, \mathbf{x}^T) + \sigma_n^2 \delta(x, x^T), \quad (108)$$

where K_1 and K_2 are equal to Equation 107 with different σ and length scale parameters.

Once a covariance function is chosen, the next step is determining the hyper-parameters of the Gaussian process. The covariance function we have chosen has 6 hyper-parameters consisting of anisotropic squared exponential functions each with a sigma parameter and two length scale parameters. The calculation of these hyper-parameters will be discussed in the following section. We now describe how to actually

perform the regression given a set of N observations. We wish to predict the magnetic anomaly map value at a location $\mathbf{x}_*$. We begin by forming two matrices which calculate the covariance between all sets of observations, $\mathbf{x}$, and between all prediction points $\mathbf{x}_*$. The matrix $\mathbf{K}$ computes the covariance between all combinations of observations and is given by

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \cdots & k(\mathbf{x}_1, \mathbf{x}_N) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) & \cdots & k(\mathbf{x}_2, \mathbf{x}_N) \\ \cdots & \cdots & \ddots & \vdots \\ k(\mathbf{x}_N, \mathbf{x}_1) & k(\mathbf{x}_N, \mathbf{x}_2) & \cdots & k(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix}. \quad (109)$$

The next matrix, $\mathbf{K}_*$ computes the covariance between each observation and the prediction point $\mathbf{x}_*$:

$$\mathbf{K}_* = \begin{bmatrix} k(\mathbf{x}_*, \mathbf{x}_1) & k(\mathbf{x}_*, \mathbf{x}_2) & \cdots & k(\mathbf{x}_*, \mathbf{x}_N) \end{bmatrix}. \quad (110)$$

Finally, the scalar K_{**} is given by

$$K_{**} = k(\mathbf{x}_*, \mathbf{x}_*) = \sigma_{f1}^2 + \sigma_{f2}^2 + \sigma_n^2. \quad (111)$$

With GPR we seek to determine the conditional probability of observing any $\mathbf{y}_*$ at location $\mathbf{x}_*$ given the observations $\mathbf{y}$ located at $\mathbf{x}$. It can be shown [51] that this conditional probability is given by

$$\mathbf{y}_* | \mathbf{y} \sim \mathcal{N}(\mathbf{K}_* \mathbf{K}^{-1} \mathbf{y}, \mathbf{K}_{**} - \mathbf{K}_* \mathbf{K}^{-1} \mathbf{K}_*^T). \quad (112)$$

Since the above distribution is a Gaussian we can directly read off the mean and covariance of the estimate $\mathbf{y}_*$ as

$$\bar{\mathbf{y}}_* = \mathbf{K}_* \mathbf{K}^{-1} \mathbf{y}, \quad (113)$$

$$\text{var}(\mathbf{y}_*) = \mathbf{K}_{**} - \mathbf{K}_* \mathbf{K}^{-1} \mathbf{K}_*^T. \quad (114)$$

6.3 Determining GPR Hyper-parameters

Accurate determination of GPR hyper-parameters is essential to obtaining a useful prediction. The response of a given covariance function can lead to very different results with different hyper-parameters. As shown in [51], Bayes theorem tells us that the hyper-parameters $\boldsymbol{\theta}$ can be estimated by maximizing $\log p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ given by

$$\log p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) = -\frac{1}{2} \mathbf{y}^T \mathbf{K}^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K}| - \frac{n}{2} \log 2\pi. \quad (115)$$

The maximization of the log likelihood equation above can be found using common conjugate-gradient methods. It does not take many flight lines to solve for useful hyper-parameters. We used a series of five common flight lines over high resolution aero-magnetic surveys in Virginia and Texas to estimate the hyper-parameters for each region. These high resolution surveys were treated as “truth” for our simulations, and could only be measured along flight lines. Recall our observations are actually the errors in the NAMAD map. We compute these errors, or residuals, by simply differencing our measurements along the flight lines with the predicted measurements from the NAMAD map. The two chosen high-resolution surveys reflect very different errors in the NAMAD map. The Virginia survey was conducted at an altitude of 300 meters. The NAMAD map in this area is under sampled at 300 meters altitude and consequently the residuals contain high frequency errors. This indicates we should have a length scale parameter approximately equal to the altitude. There are likely




Figure 85. Texas Regional Magnetic Anomaly Map at 900 meters altitude

also long wavelength errors, so there should also be a mid to long range length scale parameter. The high resolution survey in Texas was conducted at 900 meters altitude. The NAMAD map, when upward continued to 900 meters is not under sampled. The dominant errors in the Texas residuals are likely long wavelength errors due to poor geo-location. We do not expect a short length-scale parameter to be estimated for the Texas map. The hyper-parameters that are estimated from each of these flights should be very different, and reflect the discussed spatial correlations. Fig. 85 shows the true high resolution Texas survey as well as the error in the NAMAD map. Fig. 86–87 show mesh versions of the Virginia high resolution survey and NAMAD map, respectively. It is clear that the Virginia NAMAD map is missing many of the high frequency components in the high resolution survey, and is therefore very under sampled. The lines on the maps denote the flight lines along which observations were collected, and are the only information about the true field that are used to estimate the hyper-parameters. Fig. 88 shows the covariance functions evaluated with the calculated hyper-parameters for the Virginia survey. It is clear that the search for optimal hyper-parameters found two distinct length scales for each of the squared




Figure 86. Virginia Regional Magnetic Anomaly Map at 300 meters altitude




Figure 87. NAMAD Region of Virginia Survey at 300 meters altitude




Figure 88. Calculated Covariance Function for the VA Survey Area

exponential functions, just as expected. This is due to the high frequency errors from under sampling which are in the NAMAD map. The shape of the covariance function is circular for both the short length-scale and long length-scale portions of the function. This indicates the spatial correlation in one direction was not much greater than the other direction. Fig. 89 shows the covariance function calculated for the Texas survey. This function clearly has only one squared exponential term. Unlike the Virginia function, however, this covariance function detected more spatial correlation in one direction than the other, leading to an elliptical shaped covariance function. As demonstrated with these two distinct covariance functions, the chosen GPR model shows remarkable ability to structure its own model to match the data being observed.

6.4 Upward and Downward Continuation

So far we have only considered the two dimensional case for the chosen GPR model. It may seem obvious to extend the two dimensional covariance function into three dimensions. This is certainly possible, but a better method exists. The process




Figure 89. Calculated Covariance Function for the TX Survey Area

Table 12. Gaussian Hyper-Parameters Calculated From Observations Along 5 Flight Lines

Hyper-Parameter	Virginia Value	Texas Value
σ_{f1}	69.2 nT	0 nT
L_1^1	140 m	-
L_1^2	125 m	-
σ_{f2}	72.2 nT	3 nT
L_1^1	5 km	5.72 km
L_1^2	5 km	3.45 km

of upward and downward continuing magnetic anomaly data is a well understood concept in geological studies. For upward continuation, under strict enough conditions, there is an exact analytical solution to the magnetic anomaly field at higher altitudes. These conditions include knowledge of the horizontal field on an infinite two dimensional plane as well as the requirement that all magnetic sources are located below the plane (See Chapter 2 for more details). While these conditions obviously are not met in practice, accurate upward continuation of magnetic anomaly data is still possible, and is routinely implemented in geological studies. As shown in [10], the process of upward and downward continuing magnetic anomaly data can be implemented as a simple filtering operation in the Fourier domain. Upward continuation attenuates high frequencies while leaving low frequencies unchanged while downward continuation does the reverse, increasing high frequency and leaving low frequencies unchanged. The accuracy of these filtering operations relates back to the initial assumptions required for upward continuation. Because we do not have an infinite two dimensional map, our solution will never be perfect. A reasonable approximation can be obtained, however, and will likely be more accurate than simply using spatial correlations, as the GPR model uses. The two dimensional filtering operation which implements upward and downward continuation can also be applied as a line filter. The upward continuation line filter is given as

$$\mathcal{F}\{L_{h+\Delta h}\} = \mathcal{F}\{L_h\}e^{-|k|\Delta h}, \quad (116)$$

and the downward continuation line filter is given as

$$\mathcal{F}\{L_{h+\Delta h}\} = \mathcal{F}\{L_h\}e^{|k|\Delta h}, \quad (117)$$




Figure 90. Upward Continuation Transform at Several Altitudes

where

L denotes magnetic measurements along a flight line

h and Δh denote altitude and change in altitude

$\mathcal{F}\{\}$ denotes a Fourier transform

k denotes the wavenumber

Fig. 90 and Fig. 91 shows the Fourier domain response of these two filters for a 10 kilometer long flight line with sample spacing equal to 100 meters at several altitudes. It is clear that for upward continuation the decaying exponential simply decays steeper and steeper in a bounded fashion. For downward continuation, however, the filter is an increasing exponential and grows without bound. This leads to serious stability issues with downward continuation. Any noise in a flight line being downward continued will soon begin to completely dominate all other frequencies, and the computed signal at a lower altitude will be completely irrelevant. One method to address these stability issues is to low-pass filter the data as well, to stop the noise from blowing up the downward continued signal. This involves a trade-off;




Figure 91. Downward Continuation Transform at Several Altitudes

however, as applying too much low pass filtering will give a stable, but completely incorrect solution. The process of selecting the low-pass filter type and cutoff is discussed extensively in geological studies. We used a classical method called Tikhonov regularization [64]. The Tikhonov regularization form of the downward continuation transform is

$$\mathcal{F}\{L_{h+\Delta h}\} = \mathcal{F}\{L_h\} \frac{e^{|k|\Delta h}}{1 + \alpha k^2 e^{-k\Delta h}}, \quad (118)$$

where the new parameter α is the regularization parameter. This parameter provides the tradeoff between low pass filtering and downward continuation. To determine the α parameter we used the L-curve approach as described in [64]. The L-curve approach involves actually performing the full downward continuation on the signal multiple times over a geometric sequence of alpha values. The sequence of alpha values should span a large range and have a common ratio of 1.1 [49]. The L_∞ norm of the difference between successive downward continuations (in the spatial domain) gives a characteristic curve when plotted verse α as shown in Fig. 92. The optimal α value lies on a local minimum of this curve, and can be found by calculating the point




Figure 92. Calculation of α Parameter for Downward Continuation

of max curvature on the curve. The top plot of Fig. 93 shows the Tikhonov regularization downward continuation transform as well as the unregularized transform for a 200 meter downward continuation. The bottom plot shows the two corresponding downward continued solutions for a random flight line, as well as the original flight line. The instabilities without using Tikhonov regularization are clear in the plot. It is important to point out that there are inherent inaccuracies in using both the upward and downward continuation transforms. We use them because they provide a better estimate of the signal at varying altitudes than extending the GPR model to 3 dimensions would provide, as the GPR just uses spatial correlations. It seems that simply using spatial correlations would not be able to predict the increased frequency content at lower altitudes. In Section VII we discuss how to handle the practical limitations of these transforms.

6.5 Results

To simulate the process of a self-building world model we used two main sets of real data. The first was a high resolution survey over Louisa, Virginia at 300 meters




Figure 93. Comparison of Stabilized and Unstabilized Downward Continuation

altitude. The second data set was a high resolution survey over Texas. For each of these data sets we also created additional maps by taking a subset of the NAMAD over the respective survey areas. Recall that we are actually modeling the error in the NAMAD, not the total field. The NAMAD map was upward continued for the Texas survey to bring it to a common altitude. The NAMAD and Virginia surveys were already at similar altitudes. The starting error of the two map tiles is simply the difference between the true map and the NAMAD map. We wish to use a small set of flight lines to correct this error. The predicted map is the NAMAD map plus the predicted NAMAD error, as calculated by the GPR. As a performance metric we used the mean and standard deviation of the errors between the true and predicted maps. We used a common set of five and then ten flight lines over the Virginia and Texas survey areas and measured the performance of the GPR. Fig. 94 shows the absolute value of the original error in the NAMAD and the error after the GPR model is applied on five random flight lines over the Texas survey area. Fig. 95 shows the same plot for ten random flight lines. A large improvement is clearly obtained with just these few flight lines. This is due to the fact that there were few high frequency errors

Table 13. Results of Gaussian Process Regression Modeling on 5 and 10 Flight Lines Over Texas and Virginia Regions

	Mean (nT)	Standard Deviation (nT)
TX Starting NAMAD Error	1.40	3.28
TX 5 Flight Lines	-0.33	1.28
Tx 10 Flight Lines	0.04	0.60
VA Starting NAMAD Error	122.24	88.44
VA 5 Flight Lines	0.73	87.01
VA 10 Flight Lines	-4.45	78.14




Figure 94. Original Absolute Value of Error in Texas NAMAD Map and Error After GPR From 5 Flight Lines

in the Texas NAMAD map, and therefore the errors had longer spatial correlations. Fig. 96 shows the absolute value of the original error in the NAMAD and the error after the GPR is applied on five random flight lines over the Virginia survey area. Fig. 97 shows the same plots after ten flight lines. It is clear that the performance improvement in this case is smaller. This is due to the high frequency errors in the Virginia map that exist due to the under sampling of the NAMAD map.

The means and standard deviations of these results are shown in Table 13. It is clear that the GPR for the Texas map worked very well. The Virginia map shows




Figure 95. Original Absolute Value of Error in Texas NAMAD Map and Error After GPR From 10 Flight Lines

less improvement, but the mean of the map is immediately corrected with a few flight lines, fixing the long wave-length errors that exist in the NAMAD map. It is clear the GPR is not able to substantially decrease the standard deviation of the Virginia map. This is not surprising—fundamentally these are high frequency errors and the only way to detect them is to have observations in the area. Fortunately this problem would not arise very much in practice. The Virginia map is very low with an altitude of 300 meters AGL. Most aircraft obviously fly much higher than this, where errors are likely to be spatially correlated to a much greater extent. It is worth mentioning again that an under sampled map at a low altitude becomes fully sampled after a sufficient increase in altitude. Fig. 98 shows the predicted standard deviation of the Texas map using the ten random flight lines. The black planes in the figure show the initial uncertainty estimate, which is calculated from the Gaussian process hyper-parameters. The pinched in areas of the mesh plot reveal where the flight lines were located. Due to the high spatial correlation in the Texas NAMAD error, the variance over the entire map area is much smaller than the initial estimated uncertainty. Finally, Fig. 99 evaluates the performance of the GPR by comparing the




Figure 96. Original Absolute Value of Error in Virginia NAMAD Map and Error After GPR From 5 Flight Lines

error in a rectangular flight path over the Texas map before and after the ten flight line GPR. It is clear that not only is the error greatly reduced from just ten random flights over a 1000 square kilometer area, but the GPR also “knows” how accurate the map is, and increases the variance at the beginning and end of the flight path, which happened to be long distances from any of the previous observations.

6.6 Practical Implementation of the Self-Building World Model

The examples we have demonstrated so far consist of correcting the errors in a limited region of the NAMAD. In practice the corrections that are made to the NAMAD will always need to be applied to subsets of the full map. GPR has a complexity of $\mathcal{O}(N^3)$, where N is the number of observations being input to the GPR. It would not be feasible to apply the GPR to the entire map at once. These map tiles can actually be made quite large with the addition of an extra step during the data collection process. We know that we only have to sample a magnetic anomaly field at a horizontal distance equal to its height to fully sample the field. If we throw out oversampled data points we can apply GPR to fairly large tiles. Consider a




Figure 97. Original Absolute Value of Error in Virginia NAMAD Map and Error After GPR From 10 Flight Lines




Figure 98. Standard Deviation of GPR Map Prediction




Figure 99. Validity of GPR Map Prediction Over Rectangular Flight Path

map being constructed at 1 kilometer altitude as a 100×100 kilometer grid with 1 kilometer grid spacing. 100 random flights through this map tile would consist of roughly $100 \times 100 = 10,000$ observations when limiting the observations to a single measurement per square kilometer. Solving this GPR would require inverting a $10,000 \times 10,000$ matrix, an operation which takes only several seconds on a modern computer. Another approach to limit computational complexity would be to perform a “bootstrapping” approach. Once a given map tile has reached approximately 10,000 observations, or whatever computational limit applies, the result of the GPR may be used to define a new predicted map, replacing the role of the NAMAD map. Further observations would then attempt to correct errors in this new map. The downside of this approach is the loss of optimal behavior for a GPR model and the difficulty of interpreting the variance of this new map. Table 14 shows a performance comparison of a bootstrapping approach to the previously given Texas results. For this specific example, the bootstrapping approach works surprisingly well, although further study would be required to evaluate its performance in general, as it is certainly not optimal in any way.

Table 14. Comparison of Bootstrapping and Standard GPR Approaches

	STD (nT)	Computation Time (sec)
TX GPR 10 Flights at Once	0.60	171
TX GPR 2x5 Flights	0.70	43
TX GPR 10x1 Flights	0.77	15

To make the presented approach “self-building”, a fair amount of work would need to be completed to initially process the data. As stated previously, at a minimum the NAMAD map must be split into tiles, and software would need to correctly discard data points based on altitude to avoid causing excessive computation time. Each time a tile receives a new flight line the NAMAD error can be recomputed. Computation time in this case isn’t terribly important because this would be an offline process. The map would likely not update itself on the fly, because this would require all participating aircraft to be sharing information, which isn’t reasonable. Additionally, the NAMAD map must be upward continued to a level altitude. This is perfectly feasible but would likely require the assistance of the scientists who created the map in the first place. When the NAMAD map claims to be 300 meters above terrain, this is a general statement. The map is likely defined on a somewhat arbitrary surface which is draped over the terrain. To upward continue the NAMAD map to level altitude knowledge of this specific drape surface is required.

The final practical consideration for real implementation of such a system is a decision on what level or levels to calculate the map. As stated previously, we can upward and downward continue data, but this process is error prone. Upward continuation is far more accurate than downward continuation. In practice, computing the map at several altitudes is likely the best approach. Fig. 100 illustrates this layered approach. Assume three map layers exist, one at 500 meters, one at 1 kilometer, and one at 10

kilometers. This corresponds to four zones: below 500 meters, between 500 meters and 1 kilometer, between 1 kilometer and 10 kilometers, and above 10 kilometers. A given flight through any of these zones can always be upward continued to the level directly above it. Upward continuing multiple levels could work, but accuracy would decrease. In a parallel operation, if (and only if) the aircraft is within a prescribed vertical height of a layer below it, the flight line can be downward continued to this map. This is due to the large stability issues encountered with downward continuation. Large downward continuations could potentially introduce huge errors in a map tile. At this point the multiple map tiles can be used to compute a new map tile at any altitude with a simple Fourier domain grid filter, called interval continuation, as described in [27]. Using two map tiles for an upward or downward continuation adds additional constraints that can greatly improve the accuracy of these operations. This interval continuation filter is given in matrix notation as

$$\mathcal{F}\{M_z\} = \frac{1}{q} \begin{bmatrix} e^{+(z-H)k} - e^{-(z-H)k} \\ e^{-zk} - e^{+zk} \end{bmatrix}^T \begin{bmatrix} \mathcal{F}\{M_0\} \\ \mathcal{F}\{M_H\} \end{bmatrix}, \quad (119)$$

$$q = e^{-Hk} - e^{+Hk}, \quad (120)$$

where:

z is the altitude above the lower grid

$\mathcal{F}\{M_z\}$ if the Fourier transform of the grid at altitude z

$\mathcal{F}\{M_0\}$ if the Fourier transform of the lower altitude grid z

H is the height between the upper and lower grids

$\mathcal{F}\{M_H\}$ if the Fourier transform of the higher altitude grid

k is equal to the grid wavenumber $k = \sqrt{k_x^2 + k_y^2}$




Figure 100. Multiple Map Layers Illustration

6.7 Navigation Simulation Using GPR

We now wish to test the previously developed GPR method in a full navigation simulation. In this simulation, navigation accuracy will be examined over mapped areas which have been sparsely surveyed—the maps are surveyed on a grid but this grid is very under-sampled. These simulations are not attempting to show the effect of completely random flight lines in a self-building world model. Instead, they aim to show the effect which minimal surveying has on the accuracy of a navigation filter. There are several important goals of this simulation. The first is to develop a method to incorporate map covariance into the previously developed navigation filter. The second is to examine the relationship between navigation accuracy and the number of corrective survey lines through a mapped area. Finally, we wish to predict realistic navigation performances over a magnetic anomaly map which has had minimal corrections from a sparse survey.

Incorporating Map Covariance.

The first design decision that must be made when incorporating a GPR-derived map into the previously designed navigation system is how to incorporate the map covariance into the navigation filter. We propose including the map covariance into the filter measurement covariance matrix R . This indicates that as our aircraft moves about the map, the measurement covariance will be changing with respect to location. The exact location of the aircraft is not known, however, so a simple covariance lookup based on location will not suffice. Instead, we propose to make each individual particle have its own measurement covariance based upon its *hypothesized* location. In this way there is never an “incorrect” measurement covariance—even if a particle is in the wrong location, its covariance is still accurate because it is the covariance of the particle’s *assumed* location. To understand why this works consider, the following

four cases.

1. Large Particle Position Error, Large Particle Covariance: The particle will either be down-weighted with respect to other particles or share a similar weight
2. Large Particle Position Error, Small Particle Covariance: The particle will be down-weighted with respect to other particles, possibly eliminated via re-sampling
3. Small Particle Position Error, Large Particle Covariance: The particle will either be weighted high with respect to other particles or share a similar weight
4. Small Particle Position Error, Small Particle Covariance: The particle will be weighted high with respect to other particles

To further illustrate this idea of particle covariance weighting, consider the one-dimensional case shown in Fig. 101. Several hundred particles are placed along the x axis, denoted by blue circles. The solid red line shows the true map value with respect to position along the x axis. The solid blue line shows the map's covariance with respect to position along the x axis. The green line denotes a single measurement. The red circles denote the calculated particle weight at each location. From Fig. 101, it is clear that some particles are heavily down-weighted. These particles coincide with areas of the map with large measurement errors. However, not all particles with large measurement errors are down-weighted. Areas of the map with large map uncertainty can still be weighted equal to even the true location particle because we cannot trust the map at certain locations. In this way the filter acts pessimistically, only killing off particles which have a relative large measurement error *and* a small map covariance.




Figure 101. One-Dimensional Particle Weighting Using Map-Based Covariances

Simulation Setup.

Using the previously described method to incorporate map covariance into the navigation filter, we designed several navigation simulations to test the GPR method. We used the data from the cross-country test flight as the basis for the simulation. The flight flew from Virginia to Iowa at an altitude of approximately 3000 meters ASL. The map we used was the NAMAD. In these simulations we assumed NAMAD was truth, and then further corrupted the map with known errors. Fig. 102 shows the overall simulation setup, with the red line denoting the airplane flight path over the NAMAD map. The black lines denote sparse flight line corrections that were made to the NAMAD map at grid spacing of 100 kilometers. As explained shortly, the NAMAD map is corrupted with errors and these errors can then be measured along these flight lines. The flight lines are considered sparse because they are not sufficient to fully sample the magnetic field.

The first goal of these navigation simulations was to test how grid spacing for these sparse flight lines affected navigation accuracy. To test grid spacing, several scenarios




Figure 102. Self Building World Model Navigation Simulation Setup

were created. The first scenario, shown in Fig. 103, shows a flight line spacing of 100 kilometers. To allow the filter time to get an adequate fix, we want the flight line corrections to have a sufficient width. The width of the flight lines in Fig. 103 is 10 kilometers. This indicates that when performing an actual survey, several flight lines next to each other may be needed to achieve the specified width—10 kilometers in this case. Based on the magnetic field sampling discussion in Chapter 2, for the 10 kilometer line width this may be a single flight line at 10 kilometers altitude, or two flight lines at five kilometers altitude. Note that this is simulating sparse, but intentional, surveying of the magnetic field. In a truly “self-building” model, where no intentional surveying is done, the effective width of the flight lines would be approximately equal to the altitude at which they were flown. Fig. 104 shows an increased grid spacing of 200 kilometers.




Figure 103. Sparse Flight Line Spacing: 100 km Line Spacing 10 km Line Width




Figure 104. Sparse Flight Line Spacing: 200 km Line Spacing 10 km Line Width




Figure 105. Simple Binary Covariance Method (No GPR)

Line Spacing Trade Space Study.

Before we evaluate the GPR model, we wish to test a simpler method in order to conduct a large trade-space study over grid spacing and line width. For this trade-space study we develop what we call the binary map correction method. In this case, map corrections are only made inside the width of flight lines and the map covariance is simply “large” or “small” depending on if it is over a flight line or not. Fig. 105 shows an example covariance map for this simple binary covariance method. As can be seen in the figure, most of the time the map covariance has a large and constant value, but over the flight lines the covariance clamps down to a much smaller value.

Fig. 106 shows a zoomed in view of the map errors used to corrupt the NAMAD model. As can be seen in the figure, areas over a flight line have no errors because these errors are assumed to have been corrected. Fig. 107 shows an example run for the filter computed horizontal standard deviation. The green shading in the figure shows




Figure 106. Flight Line Masked Map Errors

areas in which the aircraft flew over a flight line. The filter error clearly decreases sharply when flying over flight line corrections. Between flight line corrections the filter’s drift is dominated by the drift of the INS.

Using the binary covariance method, a trade space study was performed over the grid spacing and line width parameters which could be used to sparsely sample a map. Table 15 shows the results of this trade space study. The overall trends are obvious and expected—smaller line spacing and larger line width both lead to better navigation accuracy. We have highlighted two cells of Table 15 to show the two sets of parameters in which we will test the GPR method. These highlighted cells have line spacings of 100 kilometers and 200 kilometers at a line width of 10 kilometers. At 10 kilometers altitude, this type of survey would take 10–20 times fewer flight lines to complete than a survey which fully sampled the map. Interestingly, the navigation accuracy does not decrease as much as one might expect. At a line spacing of 100




Figure 107. Horizontal Errors Using Binary Method

Table 15. Line Spacing and Line Width Navigation Accuracy Tradespace

Line Width	5 km	10 km	25 km	50 km
25 km Line Spacing	66 m	60 m	-	-
50 km Line Spacing	80 m	80 m	57 m	-
100 km Line Spacing	148 m	112 m	81 m	55 m
200 km Line Spacing	286 m	256 m	172 m	145 m

kilometers the navigation accuracy was 112 meters DRMS and at 200 kilometers line spacing the navigation accuracy was 256 meters DRMS.

Full GPR Simulation.

We next tested a full GPR navigation simulation. These simulations consisted of four main steps. The first was to generate map errors. We tried generating both long wavelength and short wavelength map errors. These map errors were created by applying the upward continuing filter described in Chapter 2 to two-dimensional white noise. The extent of the upward continuation determined the lengths of the spatial correlations of the map errors. Fig. 108 shows both a long and short wavelength error

map. Recall that the GPR method will attempt to model this error map, not the actual magnetic field.

The second step of the simulation consisted of generating regional GPR models. Recall that we use the observations themselves to determine the hyper-parameters of the Gaussian process covariance function. The covariance function we have chosen has six hyper-parameters consisting of anisotropic squared exponential functions each with a sigma parameter and two length scale parameters. The GPR covariance function describes how the observations from the survey are correlated with each other, and is useful for interpolating outside of the flight lines. We do not want to solve for a single GPR covariance function because the GPR covariance length scales and sigma parameter are likely different in different regions. It is best to have regional GPR models based on the observations within that region. For our simulation, we created a grid of GPR models over the area the aircraft flew. Fig. 109 shows the covariance function which was solved for at various locations within the grid. When interpolating outside of a flight line, it is most beneficial to use the GPR model which is closest to the location at which the data is to be interpolated.

The third step of the simulation consisted of using the GPR model to compute a corrected NAMAD map as well as to compute a map covariance. The GPR models shown in Fig. 109 were used to generate the covariance map shown in Fig. 110. The corrected map and covariance are generated by simply evaluating the GPR model at the map location, using the closest GPR model. Fig. 111 and 112 show the errors in both the GPR-corrected NAMAD map as well as the uncorrected NAMAD map for both the short and long wavelength errors.





Figure 108. Comparison of Long and Short Wavelength Map Errors




Figure 109. Grid of Regional GPR Models




Figure 110. GPR Calculated Map Covariance




Figure 111. GPR Corrected Map Values and STD Along Flight Line (Short Wavelength Errors)




Figure 112. GPR Corrected Map Values and STD Along Flight Line (Long Wavelength Error)

Table 16. Comparison of Corrective Ability of GPR Method for Short and Long Wavelength Map Errors

	Mean	Standard Deviation
Map Error: Short Wavelength	2.6 nT	20.6 nT
GPR Corrected Map Error: Short Wavelength	0.6 nT	15.1 nT
Map Error: Long Wavelength	-1.7 nT	52.7 nT
GPR Corrected Map Error: Long Wavelength	1.7 nT	15.3 nT

It is clear that the GPR method has a more apparent benefit when the map errors have long spatial wavelengths. This is fairly intuitive—large spatial correlations in the map error will allow interpolations outside the flight lines to be more accurate than when the map error has short spatial correlations. Table 16 shows how well the GPR method was able to correct the NAMAD map with both the short correlation map errors as well as the long correlation map errors. As expected, it is clear that there is a much larger potential correction to the map when the spatial correlations are large. It is interesting to note that for very short spatial correlations, the GPR method effectively collapses down to the simple binary method presented previously, in which map corrections only exist directly on flight lines. When outside of a flight line the covariance is very large, and no map corrections are possible.

The final step in the simulation is to use the corrected NAMAD map as well as the associated covariance map within the previously discussed navigation filter. Note that each of the three previous steps can be completed offline before a flight. The GPR calculations do not directly interact with the navigation filter at all; they simply provide a corrected NAMAD map as well as a covariance map. The first scenario we tested assumed a 100 kilometer grid spacing survey was conducted which could be used to correct errors in the NAMAD. The flight line width from this survey was 10 kilometers. Again, this indicates that a single line on the grid could actually be

Table 17. Navigation Results for 100 km Line Spacing and 10 km Line Width Map Corrections (Short Wavelength Errors)

Scenario: 100x10	DRMS Error
Uncorrupted Map (Base Case)	47 m
INS Only	5400 m
Corrupted Map with No Flight Line Corrections	445 m
Corrupted Map with Binary Method Flight Line Corrections	106 m
Corrupted Map with GPR Method Flight Line Corrections	96 m

several lines depending on altitude. The simulation altitude used was 1 kilometers so each line on the grid would actually represent ten flight lines to achieve the flight line width of 10 kilometers. The first scenario also corrupted the NAMAD map with short wavelength errors. Table 17 shows the DRMS accuracies for the first scenario. DRMS accuracy is given for five cases. The first is the base case. This case runs the filter with no map errors at all. This is used to show the navigation potential for a perfect map, assuming 1 nano-Tesla level measurement errors. The second case shows the navigation accuracy for the unaided INS. The third case shows the navigation accuracy when the map is corrupted with errors and no map corrections. A large constant measurement uncertainty is needed in this case to ensure the filter does not diverge. This shows that even though the map is corrupted, it still has potential to correct the drift of the INS. The fourth case shows the binary method for correcting the NAMAD map. Again, this method corrects the NAMAD map only within the width of the flight lines. It also uses either a very large covariance outside the flight lines or a very small covariance when inside the flight line width. Each particle is weighted according to this binary covariance. Finally, the fifth case shows the full GPR method results. This case uses the GPR corrected NAMAD map as well as the computed covariance map.

For the short wavelength errors it is clear there is not a large difference between

Table 18. Navigation Results for 100 km Line Spacing and 10 km Line Width Map Corrections (Long Wavelength Errors)

Scenario: 100x10	DRMS Error
Uncorrupted Map (Base Case)	47 m
INS Only	5400 m
Corrupted Map with No Flight Line Corrections	717 m
Corrupted Map with Binary Method Flight Line Corrections	135 m
Corrupted Map with GPR Method Flight Line Corrections	106 m

the binary method and the GPR method. This is expected as the GPR method cannot interpolate outside of the flight lines very well. There is still a small improvement, however. This is likely due more to the accurate covariance map than the corrected map itself. Even with short wavelength errors on a map surveyed at 100 kilometers line spacing, sub 100 meter DRMS errors were achieved. Sampling a map at 100 kilometers line spacing is far more realistic than sampling the same map at 1 kilometer line spacing and the sub 100 meter performance shows promise for navigating with sparse map corrections.

The second scenario used the same 100 kilometer grid spacing and 10 kilometer line width, but corrupted the map with long spatial correlation errors instead of short correlation errors. For the long wavelength errors it is clear the GPR method outperformed the binary method. This was the expected result for the long wavelength errors. The GPR method was able to make corrections not only inside the width of the flight lines, but also well outside of the flight lines through the GPR interpolation. The overall accuracy of this scenario is close to the previous scenario. While the errors are easier to model in this case, the un-modeled errors in the short wavelength scenario looked far more “white” and therefore did not have as great of an effect on the navigation solution as the long wavelength errors. With 100 kilometer grid spacing this scenario still achieved approximately 100 meters DRMS errors while




Figure 113. East Navigation Error Using GPR Method

utilizing only a small number of flight lines over the survey area. The east and north errors for scenario two are shown in Fig. 113–114. It is clear that the errors are corrected not only inside the actual flight lines but also before and after the flight lines as well. This is due to the GPR model accurately interpolating outside of the flight lines. Finally, Fig. 115 shows the average particle residuals and the average particle one sigma standard deviation throughout the flight. The calculated particle residuals appropriately fall inside of the particle standard deviation lines throughout the flight, indicating stable filter behavior and showing the particle covariance method working as expected.

The third scenario used 200 kilometer grid spacing with a 10 kilometer line width.




Figure 114. North Navigation Error Using GPR Method




Figure 115. Measurement Residuals Using GPR-Derived Particle Covariances

Table 19. Navigation Results for 200 km Line Spacing and 10 km Line Width Map Corrections (Long Wavelength Errors)

Scenario: 200x10	DRMS Error
Uncorrupted Map (Base Case)	47 m
INS Only	5400 m
Corrupted Map with No Flight Line Corrections	717 m
Corrupted Map with Binary Method Flight Line Corrections	280 m
Corrupted Map with GPR Method Flight Line Corrections	192 m

This resulted in a map area with far less corrective flight lines than the previous two scenarios (see Fig. 103-104). This scenario used the NAMAD map corrupted with the long wavelength errors. The overall error for both the binary method and the GPR method are worse than the previous two scenarios. This is expected as there were only half as many corrective flight lines flown over the NAMAD map. The navigation error for the binary method was 280 meters and the navigation error for the GPR method was 192 meters. The GPR method clearly outperformed the binary method, and the usefulness of the GPR method is better shown in the 200 kilometer line spacing scenario than in the 100 kilometer line spacing scenarios. The GPR method constrained the drift of the INS from 5400 meters DRMS down to 192 meters DRMS with a very limited amount of corrective flight lines.

The potential for a self-building world model is made especially clear in this scenario. Flying over even a few corrected map areas can significantly improve navigation performance. Furthermore, the usefulness of the GPR method is shown by the success of these simulations. The GPR method does all of the needed model computation offline and simply provides a navigation filter a corrected map as well as a covariance map. The navigation filter is then able to use these corrected maps and covariance information to provide a far better navigation solution than an uncorrected map would allow.

6.8 Self Building World Model Conclusions

Aerial magnetic anomaly navigation is a promising new alternative navigation system. The biggest weakness of this type of map based navigation system is map availability and quality. A self-building world magnetic anomaly model would allow the gradual creation of an accurate magnetic anomaly map with covariance bounds. Multiple-sensor navigation filters could assign the needed level of trust to magnetic anomaly derived position measurements by using the map covariance estimates. In this chapter we showed the usefulness of a Gaussian process regression model with a covariance function equal to the sum of two anisotropic squared exponentials functions. The Gaussian process was able to distinguish between an under sampled map and a fully sampled map, and interpolate to new map points accordingly. The large spatial correlations in the NAMAD dataset over a region in Texas were used to correct nearly the entire map to better than 1 nano-Tesla using only ten flight lines through a 1000 square kilometer area. The flexibility of this model will allow large corrections to the NAMAD model to be made with sparse measurements, continually increasing the usefulness of magnetic anomaly navigation. The GPR method was tested in several simulation scenarios using real flight test data. A trade space over map flight line spacing showed the navigation potential of sparse surveys with line spacings as great as 200 kilometers providing significant aiding to an INS. Finally, three navigation simulations showed how the GPR method can provide navigation errors on the order of 100 meters using corrective flight lines spaced 100 kilometers apart and navigation errors on the order of 200 meters using corrective flight lines spaced 200 kilometers apart. The benefits of the GPR method over simpler methods are made very clear when large spatial correlation errors exist in a magnetic anomaly map.

VII. Conclusion

This dissertation has detailed a large volume of research which has been focused on the topic of airborne magnetic anomaly navigation. The major contributions of each chapter are listed below.

Chapter 1

Chapter 1 presented the motivation for developing a magnetic anomaly navigation system. Many current alternative navigation systems suffer from common limitations such as when and where they can operate. The magnetic anomaly field holds promise as a global navigation signal available at all times, which is almost impossible to jam. These traits make it stand out among other alternative navigation techniques. The main contribution of Chapter 1 was a thorough literature review of magnetic anomaly navigation systems. The majority of the existing literature on airborne magnetic anomaly navigation presents simulation-only results. Among actual experimental results, it is clear that the flight test results presented in this research achieved higher accuracies than anything else found in the literature. The literature review also helped reveal the major differences between airborne magnetic navigation and other types of platforms such as indoor, ground, sea, and space. Airborne magnetic anomaly navigation exists in a “sweet spot” in which a large amount of signal is available, unlike space platforms, but most man-made sources have diminished, unlike ground and indoor platforms. The relatively fast velocity of aircraft also give observability of the temporal variations—this technique would likely not work on slow moving ships at sea.

Chapter 2

The main contribution of Chapter 2 was a thorough background on geophysical principles useful for magnetic anomaly navigation, as well an overview of the marginalized particle filter. Chapter 2 begins with a detailed description of the components of the Earth’s magnetic field, including the core field, crustal field, and space weather effects. Special attention is given to the Earth’s crustal field (or anomaly field), which serves as the primary navigation signal. Details on modeling magnetic anomaly fields are presented, as well as the necessary transforms to upward continue magnetic anomaly fields and perform time transformations. The creation of magnetic anomaly maps is discussed, including the removal of aircraft effects using compensation systems. An overview of large-scale maps is then presented. Next, a discussion on the various types of magnetic instruments was given, with a focus on their usefulness for magnetic anomaly navigation. Finally, the marginalized particle filter was presented.

Chapter 3

The main contribution of Chapter 3 was the detailed design of a magnetic anomaly navigation system. This navigation system consisted of an 18 state marginalized particle filter which used a magnetic anomaly map, a magnetometer, an INS, and a barometer to perform absolute positioning. This navigation system is far more detailed than anything found in the literature review of Chapter 1. The navigation system takes into account the fact that the magnetic anomaly field changes in three dimensions by upward continuing two dimensional map tiles. It also utilizes a navigation grade INS and estimates the errors of the INS using the magnetometer measurements. This allows the navigation system to coast between periods of high

and low magnetic anomaly variability. The magnetic anomaly navigation system also presented a proven method to remove measurement-corrupting temporal variations. An analysis on the observability of temporal variations indicated partial observability of the temporal variations at aircraft altitudes and velocities. The temporal variations were then added as a state to the marginalized particle filter. In this way the filter simultaneously estimated its position and the temporal variations, increasing navigation accuracy.

Chapter 4

Chapter 4 presented the flight test results in which the navigation system of Chapter 3 was applied to real data. The main contributions are navigation results which are far more accurate than any other aerial magnetic anomaly navigation systems in the literature. The “best-case” results used a high quality magnetic anomaly map and a geo-survey aircraft flying at low altitudes over Louisa, Virginia. Over an hour long flight test the navigation system achieved horizontal DRMS accuracies of 13 meters. This accuracy was considered “best-case” from both a phenomenology perspective as well as an engineering perspective. From a phenomenology perspective, the aircraft was flying at a low altitude of 300 meters AGL. There are greater variations in the magnetic anomaly field at low altitudes than there are at high altitudes. Fundamentally, higher navigation accuracies will always be achieved at low altitudes when using a magnetic anomaly navigation system, and there is no direct way to mitigate this effect, unlike the following engineering issues. From an engineering perspective, high map quality and a magnetically quiet aircraft environment were important factors in achieving the “best-case” results. The magnetic anomaly map which was used was fully sampled and created with the use of GPS. This type of map is not globally avail-

able. In areas of poor map quality, navigation accuracy will decrease. The use of a clean geo-survey aircraft also led to higher navigation accuracies. The magnetometer was out on a boom far from the aircraft engines. Placing the magnetometer inside the aircraft creates a more difficult compensation problem. Both of these engineering perspective issues are likely solvable to some extent. Further mapping can be conducted to address map quality and availability concerns. More advanced aircraft effect compensation may also be possible allowing less ideal aircraft environments.

Chapter 4 also presented the results of a cross-country flight test which flew from Virginia to Iowa. This flight test occurred at 3000 meters AGL, ten times higher than the previous flight test. There was also no high quality magnetic map available along the entire flight trajectory. These two factors led to decreased navigation performance. The navigation errors, which were primarily attributed to poor map quality, were on the order of several kilometers. A simulation was then conducted to predict the expected navigation accuracies had a high quality magnetic anomaly map been available. This simulation predicted navigation accuracies on the order of 150 meters. Another contribution from this flight test was the fact that existing magnetic anomaly maps are accurate enough for kilometer level navigation. While this positioning accuracy is relatively poor, it may be useful in certain applications such as over the ocean where other alternative navigation systems may fail or be jammed. Another contribution is the demonstrated need for high accuracy magnetic anomaly maps.

Chapter 5

Chapter 5 presented the results of a continental magnetic anomaly navigation sim-

ulation. A simulation framework was presented which can be used to predict magnetic anomaly navigation along any trajectory over a magnetic anomaly map. The simulation created magnetometer measurements by corrupting the measurements with real temporal variations and aircraft effects modeled by a first-order Gauss Markov random process. Predicted accuracies over the continental United States were created by running the simulation on a grid of flight paths over the United States using the North American Magnetic Anomaly Database. A strong correlation was observed between areas of poor map line-spacing and navigation errors. This indicated the simulation results over the United States are likely more accurate over areas of the country which are fully sampled in the North American Magnetic Anomaly Database. The state of California was one such area and the navigation results at several altitudes and velocities are presented here. Navigation accuracies over CA range from tens of meters when flying low and fast to hundreds of meters when flying high and slow. The main contribution of this Chapter was a tool which can be used to accurately predict magnetic anomaly navigation anywhere a magnetic anomaly map exists.

Chapter 6

Chapter 6 presented the results of a self-building world model method for magnetic anomaly maps. This self-building world model attempted to incrementally build a magnetic anomaly map as individual flight lines were collected by magnetometer-equipped aircraft. The main contribution of this section was the observation that correcting existing magnetic anomaly maps is a better strategy than creating the maps with no prior information. This is due to the large spatial correlations of the errors which exist in many magnetic anomaly maps. These errors are due to the poor data geolocation of maps made prior to the use of GPS. Real data was used to show

how even a few flight lines were able to correct large areas of the North American Magnetic Anomaly Database. Map under-sampling was shown to be a more difficult error to correct, however this type of error diminishes with increasing altitude.

7.1 Future Work

Future work which can follow on this dissertation will focus primarily on aircraft field calibration and modeling, as well as navigation using gradient, vector, and tensor measurements. Research into these areas will help overcome current obstacles preventing more accurate magnetic anomaly navigation. Measurement accuracy in the context of magnetic anomaly navigation describes how well the magnetic anomaly field can be measured, as opposed to the total field. To measure the magnetic anomaly field the other corrupting sources must be removed. These corrupting sources include aircraft effects, temporal variations, and core field effects. We wish to further study each of these three corrupting sources in the hopes of improving magnetic anomaly navigation. Measurement type, which includes scalar, gradient, vector, and tensor measurements, plays an important role in these studies. The various measurement types each bring different advantages and disadvantages to goal of increasing measurement accuracy.

Aircraft Field Modeling.

The removal of the aircraft field is an existing obstacle in achieving higher accuracy magnetic anomaly navigation. The removal of aircraft fields may require a more robust calibration procedure than what is currently implemented for survey aircraft. Development of new calibration procedures would benefit greatly from a specific aircraft magnetic field model. We hope to model an aircraft's magnetic field as it flies

through an external field. We will primarily accomplish this by simulation, but may also take actual measurements of an aircraft's magnetic field while on the ground. An aircraft-specific model may lead to insights into better aircraft effects removal while in flight. We hope to research the benefit of placing multiple magnetometers throughout an aircraft as well as the benefits of more accurate vector magnetometers. We will focus specifically on the removal of the aircraft's magnetic gradient, as gradient based approaches for magnetic anomaly navigation are likely necessary to improve accuracy, as explained below.

Measurement Types.

While there seems to be potential for more robust aircraft calibration routines, there is little possibility to better remove temporal variations when using scalar measurements. Currently we can estimate the long-wavelength components of the temporal variations. The overlapping frequencies cannot be removed and corrupt measurements. These un-removable frequencies can be as high as a nano-Tesla. Considering magnetic anomaly gradients at high altitudes are only several nano-Teslas/kilometer, this equates to hundreds of meters of error. Switching to magnetic anomaly gradient measurements has the potential to remove these temporal variations. The spatial gradient of the temporal variations is very small relative to the magnetic anomaly gradients. Base stations hundreds of kilometers apart often record temporal variations that agree to a few nano-Teslas, indicating the spatial gradients of the temporal variations are 1-2 orders of magnitude less than the spatial gradients of the magnetic anomaly field at high altitudes.

Current state of the art magnetometers are not accurate enough to form usable gradient measurements. The absolute errors of the individual magnetometers drift

too much relative to each other to provide better performance than scalar intensity measurements. Although the temporal variations are removed, new errors are introduced that have a larger negative impact than the temporal variations. More accurate magnetometers, however, are likely to become a reality in the near future. A more accurate magnetometer would allow the use of gradient measurements, which could potentially increase navigation accuracy by allowing removal of the temporal variations. We hope to research how the current approach for magnetic anomaly navigation using scalar measurements translates to gradient measurements. Expanding on this idea, we also hope to research how vector and tensor measurements, which are currently very inaccurate, could potentially increase navigation accuracy with future improved sensors.

Attitude Determination.

Spacecraft routinely determine coarse attitude using two vector measurements known in a both a world and body frame. We wish to research how future high accuracy vector and gradient measurements could be used for attitude determination of aircraft flying with vector and gradient maps. This opens up interesting possibilities of tightly-coupled navigation filters which determine both position and attitude. This type of navigation may be especially promising in indoor environments.

Bibliography

1. USGS Data Series 321. Illinois, Indiana, and Ohio magnetic and gravity maps and data: A website for distribution of data. United States Geological Survey, 2014.
2. J.L. Ahern. Spherical Harmonics. <http://principles.ou.edu/mag/earth.html>, 2009. [Online; accessed Nov-2014].
3. Viki Bankey, Alejandro Cuevas, David Daniels, Carol Finn, Israel Hernandez, Patricia Hill, Robert Kucks, Warner Miles, Mark Pilkington, Carter Roberts, Walter Roest, Victoria Rystrom, Sarah Shearer, Stephen Snyder, Ronald Sweeney, Julio Velez, J.D Phillips, and D. Ravat. Digital data grids for the magnetic anomaly map of north America: U.S. geological survey open-file report 02-414. U.S. Geological Survey, 2002.
4. Bartington. *Barrington Space Mag Brochure*. Rev. DS2714.
5. Andrew J. Beck. Magnetic fields - earth and extraterrestrial. Nasa space vehicle design criteria [environment], NASA, 1969.
6. N. Bergman. On the Cramer-Rao bound for terrain-aided navigation. Automatic control group, Linköping, 1997.
7. N. Bergman, L. Ljung, and F. Gustafsson. Point-mass filter and Cramer-Rao bound for terrain-aided navigation. In *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, pages 565–570, 1997.
8. Niclas Bergman. Point-mass filter and Cramer-Rao bound for terrain-aided navigation. *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, 1997.
9. Niclas Bergman. Point-mass filter and Cramer-Rao bound for terrain-aided navigation. In *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, pages 565–705. IEEE, December 1997.
10. Richard J. Blakely. *Potential Theory in Gravity and Magnetic Applications*. Cambridge, 1996.
11. Dmitry Budker. Optical magnetometry. Cambridge University Press, 2013.
12. Guo Caifa, Anlian Li, Hong Cai, and Yang Huabo. Algorithm for geomagnetic navigation and its validity evaluation. *College of Aerospace and Materials Engineering*, National University of Defence Technology.
13. Joseph Cain and Richard Blakely. The magnetic field of the earth's lithosphere: The satellite perspective. *Eos, Transactions American Geophysical Union*, 80.14(156), 1999.

14. Aaron Canciani and John Raquet. Absolute positioning using the earth's magnetic anomaly field. *Navigation*, 2015.
15. Aaron Canciani and John Raquet. Airborne magnetic anomaly navigation. *IEEE AES*, 2015.
16. Zhiguo Dai. Geomagnetic field aided inertial navigation using the sitan algorithm. In *Systems and Informatics (ICSAI), 2014 2nd International Conference on*, pages 79–83. IEEE, March 1990.
17. Anthony DeGregoria. *Gravity gradiometry and map matching: an aid to aircraft inertial navigation systems*. Thesis, Air Force Institute of Technology, 2010.
18. Haonan Feng. Simulation design of geomagnetic aided inertial navigation system. In *Systems and Control in Aerospace and Astronautics, 2008. ISSCAA 2008. 2nd International Symposium on*, pages 1023–1026. IEEE, December 2008.
19. Geometrics. Cesium optically pumped magnetometers. In *Technical Report M-TR91*. jan 200.
20. Geometrics. *G-822A and G-823A and B Cesium Magnetometer Manual*, 2004. Rev. B.
21. Geometrics. *G-823A Cesium Magnetometer Data Sheet*, 4 2013.
22. Sander Geophysics. High Resolution Magnetic Gradiometer Surveys. <http://www.sgl.com/MagGradient.html>. [Online; accessed Nov-2014].
23. Sander Geophysics. Aeromagnetic survey procedures. Presentation, Sander Geophysics, 2015.
24. Felix Goldenberg. Geomagnetic navigation beyond the magnetic compass. In IEEE, editor, *Proc. of Position, Location, and Navigation Symposium*, pages 684–694, 2006.
25. B. Gozick, K.P. Subbu, R. Dantu, and T. Maeshiro. Magnetic maps for indoor navigation. *Instrumentation and Measurement, IEEE Transactions on*.
26. Janne Haverinen and Anssi Kemppainen. A global self-localization technique utilizing local anomalies of the ambient magnetic field. In *Proc. of IEEE International Conference on Robotics and Automation*, Kobe International Conference Center, Kobe, Japan, 2009.
27. William J. Hinze, R. Von Frese, and Afif H. Saad. *Gravity and Magnetic Exploration: Principles, Practices, and Applications*. Cambridge, U.K.: Cambridge UP, 1998.

28. Jeff Hollowell. Heli/sitan: A terrain referenced navigation algorithm for helicopters. In *Position Location and Navigation Symposium*, pages 616–625. IEEE, March 1990.
29. G. Hulot. *Terrestrial Magnetism*. Space Science Series of ISSI. Springer, 2011.
30. Feng Jie. Kalman filter underwater passive geomagnetic navigation technology research considering the impact of diurnal variation. In *Image and Signal Processing, 2009. CISP '09. 2nd International Congress on*, pages 1–4. IEEE, October 2009.
31. T. Judd and T. Vu. Use of a new pedometric dead reckoning module in gps denied environments. In *Position, Location and Navigation Symposium, 2008 IEEE/ION*, pages 120–128, 2008.
32. Kyle Kauffman and John Raquet. Navigation via h-field signature map correlation and ins integration. In *Radar Conference, 2014 IEEE*, pages 1390–1395. IEEE, May 2014.
33. M.G. Kivelson and C.T. Russell. *A Brief History of Solar Terrestrial Physics*, chapter Introduction to Space Physics. Cambridge: Cambridge UP, 1995.
34. R. A. Langel and William J. Hinze. *The Magnetic Field of the Earth's Lithosphere: The Satellite Perspective*. Cambridge, U.K.: Cambridge UP, 1998.
35. C. Liebe. Accuracy performance of star trackers - a tutorial.
36. Jun S Liu, Rong Chen, and Tanya Logvinenko. A theoretical framework for sequential importance sampling with resampling. In *Sequential Monte Carlo methods in practice*, pages 225–246. Springer, 2001.
37. Ming Liu. Research on particle filter based geomagnetic aided inertial navigation algorithm. In *Systems and Control in Aeronautics and Astronautics (ISSCAA), 2010 3rd International Symposium on*, pages 1023–1026. IEEE, June 2010.
38. Ying Liu, Meiping Wu, Hu Xiaoping, and Hongwei Xie. Geomagnetism aided inertial navigation system. In *Systems and Control in Aerospace and Astronautics*, pages 1–5. IEEE, December 2008.
39. A.P.J. Luyendyk. Processing of airborne magnetic data. *Journal of Australian Geology and Geophysics*, 1997.
40. Stevan A. Macintyre. Magnetic field measurement. In John G. Webster, editor, *The Measurement, Instrumentation and Sensors Handbook*. 1988.
41. Mioara Manda and Monika Korte. *Geomagnetic Observations and Models*. Springer Science and Business Media, 2010.

42. Richard Marshall. Geomagnetic pulsations in aeromagnetic surveys. In *Proc. of NAECON*, Magnetic Anomaly Maps and Data for North America, nov 2014.
43. S. Maus. Emag2: Earth magnetic anomaly grid (2-arc-minute resolution). National Geophysical Data Center, NOAA, 2009.
44. S. Maus, M. Rother, C. Stolle, W. Mai, S. Choi, D. Luhr, and C. Cooke, D. Roth. Third generation of the potsdam magnetic model of the earth (pomme). *Geochemistry Journal International*, 16(1):206–14, 2005.
45. Stefan Maus. POMME-5 Magnetic Model of the Earth. <http://www.geomag.org/models/pomme5.html>, 2008. [Online; accessed Nov-2014].
46. Marvin May and Paul Meisinger. Testing of the Geomagnetic Navigation Concept. Naval Command, Control and Ocean Surveillance Center RDTE Div., Det Warminster, 1992.
47. Li Ming-Ming, Hong-Qian Lu, Hang Yin, and Xian-Lin Luang. Novel algorithm for geomagnetic navigation. *Journal of Central South University of Technology*, 18(791-99):791–99, 2011.
48. Ingemar Nygren. Robust and efficient terrain navigation of underwater vehicles. In *Position, Location and Navigation Symposium, 2008 IEEE/ION*, pages 1390–1395. IEEE, May 2008.
49. R. Pasteka, R. Karcol, D. Kusnirak, and A.. Mojzes. Regcont: A matlab based program fo rstable downward continuation of geophysica lpotential fields using tikhonov regularization). *Computers and Geoscience*.
50. Mark L. Psiaki, Stephen M. Fox, and Lejin Huang. Ground tests of magnetometer-based autonomous navigation (magnav) for low-earth-orbiting spacecraft. *Journal of Guidance, Control, and Dynamics*, 16(1):206–14, 1993.
51. C. Rasmussen and K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
52. Colin Reeves. Aeromagnetic surveys; principles, practice and interpretation. Geosoft, 2005.
53. A. B. Reid. Aeromagnetic survey design. *Geophysics*, 45(5):937–76, 1980.
54. Justin Richeson. *Gravity Gradiometer Aided Inertial Navigation Within Non-GNSS Environments*. Dissertation, University of Maryland, 2008.
55. Terence J. Sabaka, Nils Olsen, and Robert A. Langel. A comprehensive model of the quiet-time, near-earth magnetic field: Phase 3. *Geophysical Journal International*, 151(1):32–68, 2002.

56. Jeremiah A Schockley. *Ground Vehicle Navigation Using Magnetic Field Variation*. Dissertation, Air Force Institute of Technology, 2012.
57. Thomas Schon, Fredrik Gustafsson, and Per-Johan Nordland. Marginalized particle filters for mixed linear/nonlinear state-space models. *Signal Processing, IEEE Transactions on*, 53.
58. Scintrex. Magnetic application guide. Manual, Scintrex Limited, 1996.
59. Gil Shorshi and Itzhack Bar-Itzhack. Satellite autonomous navigation based on magnetic field measurements. *Journal of Guidance, Control, and Dynamics*, 18.
60. Kenneth Smith. Total field magnetometer performance published specifications and what they mean. In *Technical Report TR-120*. feb 1997.
61. William Storms, Jeremiah Shockley, and John Raquet. Magnetic field navigation in an indoor environment. In *Ubiquitous Positioning Indoor Navigation and Location Based Service (UPINLBS), 2010*, pages 1–10. IEEE, 2010.
62. Kayton M. Tazartes, D. A. and J. G. Mark. *Avionics Navigation Systems*. John Wiley and Sons, 1997.
63. E. Thebault et al. International geomagnetic reference field: the 12th generation. *Earth, Planets and Space*, 67.
64. A. Tikhonov and B. Arsenin. *Solutions of Ill-Posed Problems*. John Wiley and Sons, 1977.
65. D.H. Titterton and J. L. Weston. *Strapdown Inertial Navigation Technology*. London, UK, 1997.
66. W.E. Tolles and J.D. Lawson. Magnetic compensation of mad equipped aircraft. Airborne Instruments Lab, INC., 1950.
67. Carl Tyren. Magnetic terrain navigation. In *Unmanned Untethered Submersible Technology, Proceedings of the 1987 5th International Symposium on*, pages 245–256. IEEE, June 1987.
68. USGS. North American Magnetic Anomaly Database. http://crustal.usgs.gov/projects/namad/the_project.html, 2002. [Online; accessed Jun-2014].
69. USGS. Earthquake Glossary - moment tensor. <http://earthquake.usgs.gov/learn/glossary/?term=moment%20tensor>, 2012. [Online; accessed Nov-2014].
70. I. Vallivaara, J. Haverinen, A. Kemppainen, and J. Roning. Simultaneous localization and mapping using ambient magnetic field. In *IEEE Conference on Multisensor Fusion and Integration*, 2010.

71. Michael Veth. *Fusion of Imaging and Inertial Navigation Sensors*. Dissertation, Air Force Institute of Tecnology, 2009.
72. R. Von Frese and W. Hinze. Spherical earth gravity and magnetic anomaly analysis by equivalent point source inversion. *Earth and Planetary Science Letters*.
73. R. Von Frese, W. Hinze, D. Ravatt, and H. Goyal. Statistical prediction of satellite magnetic anomalies. *Geophysics Journal International*.
74. R. Von Frese, W. Hinze, D. Ravatt, and C. McGue. Improved inversion of geopotential field anomalies for lithospheric investigations. *Geophysics*.
75. N. Wahlstrom, M. Kok, T. Schon, and F. Gustafsson. Modeling magnetic fields using gaussian processes. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, 2013.
76. Peng Wang. Geomagnetic aided navigation suitability evaluation based on principal component analysis. In *Industrial Control and Electronics Engineering (ICICEE), 2012 International Conference on*, pages 324–329. IEEE, August 2012.
77. Wikipedia. Ionosphere. https://commons.wikimedia.org/wiki/File:Diurnal_ionospheric_current.jpg, 2005. [Online; accessed Nov-2014].
78. Wikipedia. Equatorial. https://en.wikipedia.org/wiki/File:Eej_cm4.png, 2008. [Online; accessed Nov-2014].
79. John M. Wilson and Robert J. Kline-Schoder. Passive navigation using local magnetic field variations. *ION NTM*, 2006.
80. Zunyi Xu, Yi Liu, and Lei Yan. A new correlation matching algorithm based on differential evolution for aircraft geomagnetic aid navigation. *Applied Mechanics and Meterials*, 94-96:2032–2038, 2011.
81. Yue Zhang. Geomagnetism-aided navigation based on matching algorithm for underwater vehicles. In *Selected and Revised Results of the 2011 International Conference on Mechanical Engineering and Technology*, pages 59–65. IEEE, November 2012.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From — To)		
15-09-2016		Doctoral Dissertation		Sept 2014 — Sept 2016		
4. TITLE AND SUBTITLE Absolute Positioning Using the Earth's Magnetic Anomaly Field				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Canciani, Aaron J, Capt, USAF				5d. PROJECT NUMBER 16G150		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENG-DS-16-S-074		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Dr. Robert Lutwak Defense Advanced Research Projects Agency 675 N Randolph St, Arlington, VA 22203 (703) 526-2863 robert.lutwak@darpa.mil				10. SPONSOR/MONITOR'S ACRONYM(S) DARPA		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES This work is declared a work of the U.S. Government and is not subject to copyright protection in the United States.						
14. ABSTRACT Achieving worldwide alternatives to GPS is a challenging engineering problem. Current GPS alternatives often suffer from limitations such as where and when the systems can operate. Navigation using the Earth's magnetic anomaly field, which is globally available at all times, shows promise to overcome many of these limitations. We present a navigation filter which uses the Earth's magnetic anomaly field as a navigation signal to aid an inertial navigation system (INS) in an aircraft. The filter utilizes highly-accurate optically pumped cesium (OPC) magnetometers to make scalar measurements of the Earth's magnetic field and compare them to a map using a marginalized particle filter approach. We demonstrate navigation accuracy of 13 meters DRMS with a high quality magnetic anomaly map at low altitudes with real flight data. We conduct a simulation over the continental United States to predict accuracies with respect to variables like location and altitude. Finally, we address the problem of map availability by presenting a method for a self-building magnetic anomaly model.						
15. SUBJECT TERMS Magnetic Anomaly, Navigation, Alternative Navigation, GPS-Denied, Magnetic Fields						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. John F. Raquet, AFIT/ENG	
U	U	U	UU	264	19b. TELEPHONE NUMBER (include area code) (937) 255-3636, x4580; john.raquet@afit.edu	