

AFIT/DS/ENG/04-04



TARGET RECOGNITION USING LATE-TIME RETURNS FROM
ULTRA-WIDEBAND, SHORT-PULSE RADAR

DISSERTATION
Kenneth J. Pascoe
Major, USAF

AFIT/DS/ENG/04-04

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

Approved for public release; distribution unlimited.

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT/DS/ENG/04-04

TARGET RECOGNITION USING LATE-TIME RETURNS FROM
ULTRA-WIDEBAND, SHORT-PULSE RADAR

DISSERTATION

Presented to the Faculty of the
Graduate School of Engineering and Management
of the Air Force Institute of Technology
Air University

In Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

Kenneth J. Pascoe, B.S.E.E., M.S.
Major, USAF

June 2004

Approved for public release; distribution unlimited.

AFIT/DS/ENG/04-04

TARGET RECOGNITION USING LATE-TIME RETURNS FROM
ULTRA-WIDEBAND, SHORT-PULSE RADAR

DISSERTATION

Kenneth J. Pascoe, B.S.E.E., M.S.

Major, USAF

Approved:

Dr. Michael J. Havrilla
Dissertation Advisor

Date

Dr. Milton E. Franke
Dean's Representative

Date

Dr. Peter S. Maybeck
Committee Member

Date

Dr. Aihua W. Wood
Committee Member

Date

Dr. Robert A. Calico
Dean

Acknowledgments

Thanks to all the people who helped to make this research happen. Thanks to my committee: Dr. Michael Havrilla, the late Major Bill Wood, Dr. Peter Maybeck, Dr. Aihua Wood, and Dr. Milton Franke and my original advisor, Lieutenant Colonel Pete Collins.

Particular thanks to my advisor through much of this work, Major Bill Wood. He guided the whole effort, helped me through some rough periods, kept me out of trouble, and conducted most of the scattering simulations used in this research. The world is diminished by his untimely loss. May he rest in peace.

Particular thanks also to Dr. Maybeck for several extensive critiques of this document, and to Dr. Wood for an extensive critique which managed to find a different set of errors. If the reader finds any problems with the document, the fault lies with the author. That there are not many more errors and omissions must be credited to Dr. Maybeck, Dr. Wood, Major Wood, and Dr. Franke. This document improved in style and substance due to their patient efforts. Their questions forced me to think and learn.

Thanks to Dr. Carl Baum and Mr. Bill Prather of the Air Force Research Laboratory High Power Microwaves Division at Kirtland Air Force Base in Albuquerque, New Mexico, for advice, guidance, and funding. Thanks also to Dr. Tapan Sarkar of Syracuse University for advice.

Most importantly, special thanks and eternal love to my wife, who waited patiently while I worked. She handled the little things, the big things, the crises, the disappointments, the good news, the frequent bad news, and the checkbook. She raised the children with minimum assistance and maximum success. I owe everything to her. And thanks to the children for waiting. Dad is back to play with you now.

Table of Contents

	Page
Acknowledgments	iii
List of Figures	viii
List of Tables	xx
List of Abbreviations	xxi
Abstract	xxii
 I. Introduction	 1-1
1.1 Background and Purpose	1-1
1.2 Impulse Radar Technology	1-2
1.3 Problem Statement	1-3
1.4 Overview of Solution	1-4
1.5 Contributions	1-7
1.6 Organization of Report	1-10
 II. Background	 2-1
2.1 Notation, Units and Some Definitions	2-1
2.1.1 Complex Variables	2-3
2.1.2 Bandwidth	2-4
2.1.3 Linear Algebra Terms	2-5
2.1.4 Dyadic Operations	2-6
2.2 Scattering Theory	2-7
2.2.1 The Singularity Expansion Method	2-8
2.2.2 Late Time	2-21

	Page
2.3 Signal Processing	2-23
2.3.1 The Matrix Pencil Method	2-23
2.3.2 The Total Least-Squares Matrix Pencil Method	2-30
2.3.3 Modified Matrix Pencil Methods	2-31
2.4 Stochastic Estimation	2-33
2.4.1 The Kalman Filter	2-34
2.4.2 Multiple Model Adaptive Estimation	2-38
2.5 Previous Target Recognition Techniques and the E-Pulse	2-43
2.6 Chapter Summary	2-47
III. Building the Library: Signal Processing Research	3-1
3.1 Overview of Target Signature Development	3-1
3.2 The Modified Total Least-Squares Matrix Pencil Method	3-3
3.2.1 Evaluation of M-TLS-MPM with Kumaresan-Tufts Data	3-3
3.2.2 Screening Invalid Poles	3-19
3.3 Evaluation of M-TLS-MPM with Synthetic Data	3-22
3.3.1 Synthetic Data	3-22
3.3.2 Evaluating Noise Suppression by the LRHA	3-26
3.3.3 Evaluating the Significance Criterion P	3-30
3.3.4 Signal Reconstruction and Error	3-42
3.3.5 Application to Theoretical Sphere Data	3-49
3.4 Combining Data from Various Azimuths	3-56
3.5 Developing Target Signatures with M-TLS-MPM	3-58
3.6 Chapter Summary	3-83
IV. Reading from the Library: Stochastic Estimation Research	4-1
4.1 Overview	4-1

	Page
4.2 Applying the Kalman Filter	4-4
4.2.1 Filter Tuning	4-17
4.2.2 Filter Tests with Synthetic Data	4-21
4.3 Applying Multiple Model Adaptive Estimation	4-45
4.3.1 Selecting a Target	4-49
4.3.2 Confusion Matrices	4-52
4.3.3 Verifying Target Identification via Chi-Squared Test	4-54
4.4 Target Recognition Using Radar Data	4-57
4.5 Comparison to the E-Pulse	4-70
4.6 Chapter Summary	4-78
V. Conclusion and Recommendations	5-1
5.1 Contributions	5-2
5.2 Recommendations for Future Research	5-4
5.2.1 Recommendations for Future Electromagnetics Research	5-4
5.2.2 Recommendations for Future Signal Processing Research	5-7
5.2.3 Recommendations for Future Stochastic Estima- tion Research	5-9
Appendix A. Additional Discussion on Electromagnetic Theory . . .	A-1
A.1 Mie Series Convergence	A-1
A.2 Sphere Theoretical Poles	A-3
A.3 Proof of Differentiability for Sphere SEM Poles	A-11
Appendix B. Measurement Techniques	B-1
B.1 Target Descriptions	B-3

	Page
Appendix C. Computational Techniques	C-1
C.1 Computational Results in the Literature	C-1
C.1.1 Cylinders and Cones	C-1
C.1.2 Thin Wire Dipoles	C-2
C.1.3 Rectangular Plates	C-3
C.2 Target Simulation	C-4
Appendix D. Data Processing	D-1
D.1 The Inverse Discrete Fourier Transform	D-1
D.2 Negative Frequency Data and the IDFT	D-3
D.3 Frequency Data Windowing	D-4
D.4 Data Smoothing	D-8
Bibliography	BIB-1
Vita	VITA-1

List of Figures

Figure		Page
1.1.	Target Recognition System Diagram	1-5
1.2.	Down Converter Diagram	1-7
2.1.	Simulated time-domain scattering from a cylinder, six inches long and two inches in diameter, axially-polarized illumination with bandwidth 0.1 to 18 GHz	2-9
2.2.	Simulated time-domain scattering from a cylinder, six inches long and two inches in diameter, axially-polarized illumination with bandwidth 0.1 to 18 GHz, emphasizing late-time scattering	2-9
2.3.	Two Dirac Delta Functions and constant, Freq. Domain . . .	2-11
2.4.	Dirac Delta Function (solid) and cosine (dotted), Time Domain	2-11
2.5.	Scatterer Illuminated by Plane Wave	2-15
2.6.	MMAE Diagram	2-43
3.1.	Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM	3-7
3.2.	Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM	3-7
3.3.	Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM	3-8
3.4.	Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM	3-8
3.5.	Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA	3-11
3.6.	Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA	3-11
3.7.	Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA	3-12

Figure		Page
3.8.	Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA	3-12
3.9.	Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data	3-14
3.10.	Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data	3-14
3.11.	Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data	3-15
3.12.	Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data	3-15
3.13.	Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data	3-16
3.14.	Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data	3-16
3.15.	Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data	3-17
3.16.	Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data	3-17
3.17.	Pole locations before removing high-damping-rate poles. Line near top indicates $\omega = \pi/\Delta t$	3-20
3.18.	Synthetic data Version 1, no noise added, vs range in inches .	3-23
3.19.	Synthetic data poles	3-25
3.20.	MNTSE (dB) of $y_H(t)$, normalized by mean $x(t)^2$, vs SNR, found via M-TLS-MPM with 0 to 2 iterations of LRHA . . .	3-28
3.21.	Synthetic Data Version 1, $y(t)$ with 13 dB SNR, dashed, and uncorrupted signal $x(t)$, solid	3-28

Figure		Page
3.22.	Synthetic Data Version 1, $y_H(t)$ 13 dB SNR and 1 iteration of LRHA, dashed, and uncorrupted signal $x(t)$, solid	3-29
3.23.	Synthetic Data Version 1, $y_H(t)$ 13 dB SNR and 2 iterations of LRHA, dashed, and uncorrupted signal $x(t)$, solid	3-29
3.24.	Estimated ($\cdot$) and actual ($\diamond$) poles, 25 dB SNR, $P = 1.3$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-32
3.25.	Estimated ($\cdot$) and actual ($\diamond$) poles, 17 dB SNR, $P = 0.9$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-32
3.26.	Estimated ($\cdot$) and actual ($\diamond$) poles, 9 dB SNR, $P = 0.5$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-33
3.27.	Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.3$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-33
3.28.	Estimated ($\cdot$) and actual ($\diamond$) poles, 7 dB SNR, $P = 0.5$ (offset +0.1), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-35
3.29.	Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.4$ (offset +0.1), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-35
3.30.	Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.5$ (offset +0.2), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-36
3.31.	Estimated ($\cdot$) and actual ($\diamond$) poles, 25 dB SNR, $P = 1.6$ (offset +0.3), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits	3-37
3.32.	MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20$ (no offset)	3-38
3.33.	MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20$ (no offset)	3-38

Figure		Page
3.34.	MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 + 0.2$	3-39
3.35.	MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 + 0.2$	3-39
3.36.	MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 - 0.2$	3-40
3.37.	MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 - 0.2$	3-40
3.38.	Overlay of measured and reconstructed signals	3-47
3.39.	Overlay of measured and reconstructed signals	3-48
3.40.	Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 15$, and theoretical poles ($\circ$)	3-51
3.41.	Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 6$, and theoretical poles ($\circ$)	3-52
3.42.	Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 2$, and theoretical poles ($\circ$)	3-53
3.43.	Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 1$, and theoretical poles ($\circ$)	3-53
3.44.	Estimated poles ($\times$) from 3-inch radius sphere, unwindowed Mie Series data, using M-TLS-MPM, no noise added, $P = 2$, and theoretical poles ($\circ$)	3-55
3.45.	Late time signal and reconstructed signal from 3-inch radius sphere, unwindowed Mie Series data, using M-TLS-MPM, no noise added	3-55
3.46.	Poles estimated ($\cdot$) from 5" by 1" diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-62

Figure		Page
3.47.	Poles estimated ($\cdot$) from 5" by 1" diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)	3-62
3.48.	Poles estimated ($\cdot$) from 3" by 1" diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-63
3.49.	Poles estimated ($\cdot$) from 3" by 1" diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)	3-63
3.50.	Poles estimated ($\cdot$) from 6" by 2" diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-65
3.51.	Poles estimated ($\cdot$) from 6" by 2" diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-65
3.52.	Poles estimated ($\cdot$) from 4.5" square plate, ϕ -pol, using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)	3-66
3.53.	Poles estimated ($\cdot$) from 4.5" square plate, θ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-66
3.54.	Poles estimated ($\cdot$) from Block 112, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-69
3.55.	Poles estimated ($\cdot$) from Block 112, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-69
3.56.	Signal reconstruction example for Block 112, ϕ -pol, error 1.6%	3-70
3.57.	Signal reconstruction example for Block 112, θ -pol, error 2.9%	3-70
3.58.	Poles estimated ($\cdot$) from Block 123, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)	3-71

Figure		Page
3.59.	Poles estimated ($\cdot$) from Block 123, θ -pol, using M-TLS-MPM, no noise added, $P = 1.6$, combined poles ($\square$), and published poles ($\circ$)	3-71
3.60.	Signal reconstruction example for Block 123, ϕ -pol, error 2.9%	3-74
3.61.	Signal reconstruction example for Block 123, θ -pol, error 4.8%	3-74
3.62.	Poles estimated ($\cdot$) from cube, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-75
3.63.	Poles estimated ($\cdot$) from cube, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)	3-75
3.64.	Signal reconstruction example for 3" cube, ϕ -pol, error 2.2% .	3-76
3.65.	Signal reconstruction example for 3" cube, θ -pol, error 4.1% .	3-76
3.66.	Poles estimated ($\cdot$) from Combo A, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and combined poles of components ($\circ$)	3-77
3.67.	Poles estimated ($\cdot$) from Combo A, θ -pol, using M-TLS-MPM, no noise added, $P = 0.7$, combined poles ($\square$), and combined poles of components ($\circ$)	3-77
3.68.	Poles estimated ($\cdot$) from Combo B, ϕ -pol using M-TLS-MPM, no noise added, $P = 0.9$, combined poles ($\square$), and combined poles of components ($\circ$)	3-79
3.69.	Poles estimated ($\cdot$) from Combo B, θ -pol, using M-TLS-MPM, no noise added, $P = 1.2$, combined poles ($\square$), and combined poles of components ($\circ$)	3-79
4.1.	Target Recognition System Diagram	4-2
4.2.	Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$	4-23
4.3.	Mean and standard deviation of Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$	4-23

Figure		Page
4.4.	Scaled residuals for Synthetic Data Version 1	4-24
4.5.	Mean and standard deviation of scaled residuals for Synthetic Data Version 1	4-24
4.6.	Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode	4-26
4.7.	Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode	4-26
4.8.	Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode	4-27
4.9.	Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode	4-27
4.10.	Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode	4-28
4.11.	Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode	4-28
4.12.	Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode	4-29
4.13.	Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode	4-29
4.14.	Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, third mode suppressed	4-31
4.15.	Mean and standard deviation of Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, third mode suppressed	4-31
4.16.	Scaled residuals for Synthetic Data Version 1, 3 dB SNR, third mode suppressed	4-32
4.17.	Mean and standard deviation of Scaled residuals for Synthetic Data Version 1, 3 dB SNR, third mode suppressed	4-32
4.18.	Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-33

Figure		Page
4.19.	Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-33
4.20.	Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-34
4.21.	Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-34
4.22.	Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode, third mode (this mode) suppressed	4-36
4.23.	Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode, third mode (this mode) suppressed	4-36
4.24.	Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-37
4.25.	Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed	4-37
4.26.	Synthetic Data Version 2, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\hat{x}}$, filter based on Version 1 frequencies	4-38
4.27.	Mean and standard deviation of Synthetic Data Version 2, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\hat{x}}$, filter based on Version 1 frequencies	4-38
4.28.	Scaled residuals for Synthetic Data Version 2, 3 dB SNR, filter based on Version 1 frequencies	4-40
4.29.	Mean and standard deviation of Scaled residuals for Synthetic Data Version 2, 3 dB SNR, filter based on Version 1 frequencies	4-40
4.30.	Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-41
4.31.	Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-41

Figure		Page
4.32.	Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-42
4.33.	Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-42
4.34.	Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-43
4.35.	Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-43
4.36.	Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-44
4.37.	Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies	4-44
4.38.	MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 25 dB data	4-48
4.39.	MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data	4-48
4.40.	MMAE-MAP with $\beta_k = 1$, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data	4-51
4.41.	Percent of failures of χ^2 test, filters based on synthetic data with 13dB SNR	4-57
4.42.	Output of Filter 1 with signal model for simulated 1" $\times$ 1" $\times$ 2" block target; measurement (solid), reconstruction (dashed) and 1- σ bounds (dash-dot)	4-58
4.43.	Output of Filter 2 with signal model for simulated 1" $\times$ 2" $\times$ 3" block target; measurement (solid), reconstruction (dashed) and 1- σ bounds (dash-dot)	4-58

Figure	Page
4.44. Output of Filter 3 with signal model for simulated $1.1'' \times 1.9'' \times 2.8''$ block target; measurement (solid), reconstruction (dashed) and $1-\sigma$ bounds (dash-dot)	4-59
4.45. Output of Filter 4 with signal model for simulated $4.5''$ square plate target; measurement (solid), reconstruction (dashed) and $1-\sigma$ bounds (dash-dot)	4-59
4.46. HCP of Filter 1 (correct answer, solid line), Filter 2 (dashed), Filter 3 (dotted), Filter 4 (dash-dot). Correct filter has highest HCP at filter equilibrium.	4-61
4.47. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, ϕ -pol	4-62
4.48. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, θ -pol	4-64
4.49. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, θ -pol	4-67
4.50. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated cylinders and cone, θ -pol	4-67
4.51. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated Cone, Cylinder A, Cylinder B, Cylinder C, ϕ -pol	4-68
4.52. MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated Cube, Combo A, Cone, Cylinder C, θ -pol	4-69
4.53. Percent of failures of χ^2 test vs SNR (dB), Signatures: Simulated Cube, Combo A, Cone, Cylinder C, θ -pol	4-69
4.54. E-Pulse for Synthetic Data Version 1, estimated from 25 dB data	4-72
4.55. E-Pulse for Synthetic Data Version 2, estimated from 25 dB data	4-72
4.56. E-Pulse for Synthetic Data Version 3, estimated from 25 dB data	4-73
4.57. E-Pulse for Synthetic Data Version 4, estimated from 25 dB data	4-73
4.58. Eight convolutions of E-Pulse for Synthetic Data Version 1 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start	4-75

Figure		Page
4.59.	Eight convolutions of E-Pulse for Synthetic Data Version 2 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start	4-75
4.60.	Eight convolutions of E-Pulse for Synthetic Data Version 3 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start	4-76
4.61.	Eight convolutions of E-Pulse for Synthetic Data Version 4 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start	4-76
4.62.	E-Pulse, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 25 dB data . .	4-77
4.63.	E-Pulse, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data . .	4-77
A.1.	Theoretical poles for 3" radius sphere for n=1 through n=27, q=1 (·) and q=2 (+)	A-7
A.2.	Theoretical poles for 3" radius sphere for n=1 through n=35, q=1 (·) and q=2 (+)	A-7
A.3.	Theoretical poles for 3" radius sphere for n=1 through n=85, q=1 (·) and q=2 (+)	A-8
A.4.	Theoretical poles with $\Omega > -8.4$ for 3" radius sphere	A-9
A.5.	Estimated coefficient magnitudes for theoretical poles with lowest $ \Omega $ for 3" radius sphere	A-10
A.6.	Theoretical late-time return from a 3" radius sphere, incident illumination 0.1 to 18 GHz	A-10
B.1.	Stepped-Frequency Radar System	B-2
B.2.	Target measurement geometry illustrating ϕ polarization . . .	B-5
C.1.	Coordinates and geometry of illumination	C-6
C.2.	Geometry of 'Tophat' Target	C-8
C.3.	Geometry of Cone Target	C-8

Figure		Page
C.4.	Phase difference, JRMBOR (no phase correction) minus Mie, degrees	C-9
C.5.	Sphere backscattered fields, time domain, Mie vs. JRMBOR .	C-10
C.6.	Geometry of ‘Combo A’ Target	C-12
C.7.	Geometry of ‘Combo B’ Target	C-12
C.8.	Real and imaginary components of frequency domain data from 3” cube, computed with grid densities of 12 ppw (solid line) and 20 ppw (dashed), focused on lower frequencies	C-13
C.9.	Time domain data from 3” cube, computed with grid densities of 12 ppw (solid line) and 20 ppw (dashed), focusing on late time	C-13
D.1.	Sphere time domain response, twin Hann windows, V vs ns .	D-7
D.2.	Sphere time domain response, single Gaussian window, V vs ns	D-7
D.3.	Sphere poles, data windowed by Hanning, Gaussian, Blackman windows	D-8
D.4.	Smoothed and raw measured data, 2-18 GHz	D-10
D.5.	Smoothed and raw measured data, 15-18 GHz	D-10
D.6.	Smoothed and raw measured data, Time Domain	D-12
D.7.	Poles estimated from smoothed and raw data	D-12

List of Tables

Table		Page
3.1.	Pole locations for synthetic data, giganepers/sec and gigaradians/sec	3-25
3.2.	Coefficients A_m , B_m for synthetic data, millivolts	3-26
4.1.	Comparison of Resonance and Response Concepts	4-5
4.2.	Routh Array	4-11
4.3.	Confusion matrix, synth data, SNR = -9 dB	4-54
4.4.	Confusion matrix, sim. radar data, SNR = -3 dB, θ -Pol. Signals	4-64
4.5.	Confusion matrix, sim. radar data, SNR = +3 dB, θ -Pol. Signals	4-65
4.6.	Confusion matrix, sim. radar data, SNR = 11 dB, θ -Pol. Signals	4-65
A.1.	Number of Mie series terms required for 15-digit convergence	A-3
A.2.	Sphere Theoretical Poles	A-6
B.1.	Targets measured with bandwidth 2-18 GHz, dimensions in inches	B-4
C.1.	First Four Pole Pairs of Wire Dipole, sd/c , ...	C-3
C.2.	First Four Computed Pole Pairs For Square Plate	C-4
C.3.	Dimensions (inches) of targets simulated with CARLOS code	C-7
C.4.	Dimensions (inches) of targets simulated with CARLOS code	C-11

List of Abbreviations

Abbreviation	Page
SEM Singularity Expansion Method	xxii
MAP Maximum <i>A Posteriori</i>	xxii
UWB ultra-wideband	1-2
AFRL/DEH Air Force Research Laboratory, High Power Microwaves Division	1-3
dB decibels	1-3
MMAE-MAP Multiple Model Adaptive Estimation - Maximum <i>A Pos- teriori</i>	1-4
MPM Matrix Pencil Method	1-6
M-TLS-MPM Modified Total-Least-Squares Matrix Pencil Method .	1-6
MoM method of moments	2-8
HCP hypothesis conditional probability	2-40
M-TLS-MPM Modified Total-Least-Squares Matrix Pencil Method .	3-1
CNL complex natural logarithm	3-4
MSE mean squared error	3-5
CRB Cramer-Rao bound	3-5
NTSE normalized temporal-average squared error	3-27
MNTSE mean normalized temporal-average squared error	3-27
PA&A Pole Association and Averaging	3-56
ppw points per wavelength	C-10

Abstract

In theory, targets may be identified by the natural resonances present in the late-time electromagnetic signals they scatter after illumination with transient waveforms. For perfectly conducting targets, the resonances in the electric and magnetic fields take the form of first-order poles in the Laplace domain, and exponentially damped sinusoids in the time domain. Although in theory a target's late-time scattering consists of an infinite number of poles, only a few of these are associated with significant scattering terms. The frequencies and damping rates of these terms are characteristic of the object's size and shape. The goal of this research is to develop an algorithm capable of distinguishing targets based on the late-time scattering they exhibit.

Target identification is pursued in two stages called library building and library reading. To build the library, signatures of simple-shape targets (blocks, cylinders, etc.) are developed based on estimates of the targets' natural resonances. The complex frequencies associated with the natural resonance modes are aspect-independent, according to Singularity Expansion Method (SEM) theory. Poles are identified for simulated and measured scattering from various targets using a novel variant of the Matrix Pencil Method, called the Modified Total Least-Squares Matrix Pencil Method. An all-aspect signature is produced by selectively grouping and averaging the poles estimated from single-aspect measurements. The library is read with a technique based on Kalman filtering, called Maximum *A Posteriori* (MAP) Multiple-Model Adaptive Estimation. A linear, time-domain, state-space system model is developed for the late time scattering. Several target filters, each representing different target types, are operated simultaneously. Residuals produced by each filter are compared, and the filter based on the system model that best matches the scattered signal represents the correct target.

The resulting algorithm is capable of accurately recognizing signals composed of sums of damped sinusoids, even with signal-to-noise ratios approaching 0 dB. It is able to distinguish between late-time signals but is dependent on both proper filter tuning and accurate target signatures for effective discrimination. The technique is suited to identification problems involving short ranges to the target, such as classification of land mines and unexploded ordnance that may be shallowly buried or concealed in foliage. This technique may also be employed on targets which are concealed in baggage or behind walls.

TARGET RECOGNITION USING LATE-TIME RETURNS FROM ULTRA-WIDEBAND, SHORT-PULSE RADAR

I. Introduction

1.1 Background and Purpose

In this research a target recognition technique is developed that recognizes the radar signatures of targets by comparison with stored models. The radar signatures of interest are late-time, resonant or ‘ringing’ returns excited by an ultra-wideband, short-pulse radar. An advantage of the technique developed in this research is that the resonant frequencies found by the new technique are, in theory, independent of the aspect, or direction from which a target is viewed [9, 13, 18, 90]. This allows for a target description which takes up little memory, on the order of one hundred parameters per object. Aspect-dependent techniques based on high-resolution radar, by comparison, require on the order of one billion parameters to describe a single complicated target [25]. Techniques based on synthetic aperture radar images or other images require a similar amount of storage. The aspect-independent nature of the signatures makes it desirable to excite the target’s natural ringing modes and use them to distinguish it from other targets.

This work presents and characterizes a technique which can be a useful component of a target recognition system. Such a system must provide acceptable confidence of a target’s identity, with minimum dwell time and computation time required for a decision. The research provides a working algorithm that quickly decides which signature in a library best fits a set of sensor data. Such an algorithm can be used in parallel with other algorithms, which make use of early-time radar scattering or

other information. However, characterizing other algorithms and integrating them with this one were outside the scope of this research.

A potential application of this target recognition technique that is of interest to the Air Force is detection and identification of vehicles or other targets on the ground and obscured by foliage or camouflage. Unmanned aircraft could employ the technique to enable the precision engagement of hidden enemy ground forces and ground-based air defenses early in a conflict. Other military applications may include detection and identification of obscured targets such as weapons concealed in clothing, in baggage, or behind walls [93]. The technique may also prove useful for aerial detection of mines and other unexploded ordnance.

1.2 Impulse Radar Technology

This work is possible due to the convergence of new technology with existing techniques and theory. The key technological advance is ultra-wideband, short pulse radar. Ultra-wideband transmission and reception make an ‘almost ideal’ impulse radar possible. In the past, impulse response was simulated and measured using long frequency sweeps and post processing via Fourier Transform methods. Now, ultra-wideband generators can excite near-impulse short pulses, illuminate the target instantaneously and measure the near-impulse response directly [1, 108]. Existing techniques of adaptive state estimation should make it relatively easy to extract frequency components from the returned signal. As explained further in Section 2.1, the wider the bandwidth of energy involved in a pulse, the shorter the pulse’s time duration.

An ideal impulse, denoted in mathematics by the Dirac delta function $\delta(t)$, has infinite frequency content. It contains constant amplitude power over all frequencies. A realizable approximation of an impulse is a short, sharp pulse having near-constant power per unit frequency over a very wide bandwidth. Transmitters producing such short pulses are called ultra-wideband (UWB) sources. The

High Power Microwaves Division of the Directed Energy Directorate of Air Force Research Laboratory (AFRL/DEH) has years of experience developing UWB transmitters for radar and other applications. AFRL/DEH, located at Kirtland Air Force Base, New Mexico, has produced UWB pulses with constant power spectral density (within 3 decibels (dB)) across the bandwidth from 35 megahertz to 4 gigahertz, corresponding to a rise time of 85 picoseconds [108].

1.3 Problem Statement

Targets are assumed to be detected prior to the start of the identification process by an unspecified detection sensor. Target locations are known. In this research, targets are assumed to be stationary, but extending the technique to moving targets should not be difficult. It only involves correcting the radar data to account for relative velocity between the radar and the target using the approximate formula $f_A = f_T(1 + 2v/c)$ [128], where f_A is the apparent frequency seen by the radar, f_T is the true radar frequency, v is target velocity, and c is the speed of light. The targets are assumed highly conducting and will be modelled as perfect electrical conductors (PEC). The PEC assumption simplifies the problem somewhat by neglecting terms associated with surface resistance.

The recognition sensor illuminates the target with a series of ultra-wideband, short radar pulses. The target recognition technique receives the backscattered illumination from the target. It processes the late-time portion of the backscattered signal via adaptive estimation techniques [76] based on the Kalman filter [65, 77]. The output will be a perceived target classification based on a match to a known late-time signature. The output could be ‘unknown’ if there is insufficient confidence to declare a target. The target recognition algorithm should work quickly, without requiring extraordinary computational power.

The technique is only expected to work at relatively short ranges due to power limitations. The late-time portion of the scattered signal is much weaker than the

early-time or middle-time portions of the signal. It is typically 20 dB or more down from early time, although the actual difference between early and late time depend on the target being illuminated, the frequency of illumination, and the aspect from which the target is viewed. Since a 12 dB reduction in target cross-section reduces radar detection range by half [59], the effective range of any late-time technique is limited to one quarter or less of the effective range of a technique making use of early time radar scattering from a target.

Development of the algorithm requires good estimates of the complex natural resonance frequencies, or Singularity Expansion Method (SEM) poles, of each known target. Target scattering must be measured or simulated, and the late-time results analyzed to provide estimates of the frequencies. Once radar data is generated, poles must be estimated. An effective estimation code will be resistant to noise. It will provide accurate estimates of poles, and will not provide additional, spurious results. The pole estimation code is not required to work quickly because it is intended to be used well in advance of the operation of the target recognition code. The pole estimation code will be used to populate a library of target signatures that will be utilized in real-time by the target recognition code. Scattering is much weaker in late time than in early time, resulting in lower signal-to-noise ratios and increased difficulty in accurately assessing the SEM poles from radar data.

1.4 *Overview of Solution*

A target recognition algorithm is developed based on a novel application of the the Maximum *A Posteriori* version of the Multiple Model Adaptive Estimation algorithm (MMAE-MAP) [76]. The algorithm inputs radar scattering information into a bank of Kalman filters. Each filter contains a state-space model which represents the expected late-time scattering from a known target. The filter models are based on the complex resonant frequencies of the terms making up the late-time response. Each target is represented by a different set of complex frequencies, in the form of




Figure 1.1 Target Recognition System Diagram

damping and oscillating rates. The number of frequencies in the filter model also varies from target to target. The scattered signal received by the radar is input into each filter in the bank. Within each filter the signal is compared to the model. The filter reports the residual, which is the difference between the actual measurement and the best prediction of it before its arrival based on the assumed model. Residuals are scaled and compared, and the filter generating each residual stream is assigned a probability of being correct. When all filters reach steady-state operation, or at the end of the data set, the algorithm is terminated. The filter with the highest probability at termination is reported as the correct match. A χ^2 test [44, 101] on the residuals of the best-matching filter is intended to provide additional indication that the filter adequately matches the signal. A diagram of the recognition algorithm is shown in Figure 1.1.

For comparison, a target recognition algorithm based on extinction pulse convolution [113] was implemented. The E-pulse technique had much less tolerance to

white noise than the MMAE-MAP technique. The signal-to-noise ratio required for successful target recognition in 90% of attempts using the E-pulse was about 20 dB higher than for the MMAE-MAP method.

The target recognition algorithms were tested with measured and simulated data. A variety of simple geometric shapes (cylinders, spheres, blocks) were measured using swept frequencies over as wide a bandwidth as possible. Target responses were predicted for simple shapes, and for two simple shapes in close proximity, using computational electromagnetics codes. Computational results were compared to measured and theoretical results when available. Frequency-domain results from measurements and simulations were transformed into the time domain using an Inverse Fast Fourier Transform. The resulting equivalent impulse response was the source of late-time data. This data was used both as a means to estimate natural frequencies of targets and as input data for the target recognition algorithm.

To estimate the natural frequencies of the targets from their late-time scattering, a novel variant of the Matrix Pencil Method (MPM) [119], denoted the Modified Total Least-Squares Matrix Pencil method (M-TLS-MPM) was developed, tested, and implemented. Existing variants of MPM have a published history of application to late-time scattering problems [48, 49, 51]. M-TLS-MPM provides some additional noise resistance over published versions of MPM at lower signal-to-noise ratios. Pole estimates from various aspect angles of the same target were averaged to produce all-aspect signatures. In theory, the same poles are present at all aspect angles, but pole estimates varied somewhat from one angle to the next. If a target can be adequately described from many or all aspect angles by a single set of poles, a target recognition system can easily store signatures for many targets in its library.

The most significant limiting factor in the success of this algorithm is the quality of the signatures associated with each target type. These signatures are based on estimates of the natural resonance frequencies of each target. The estimates varied little in oscillating rate ω , but somewhat more in damping rate Ω , as aspect angle




Figure 1.2 Down Converter Diagram

was varied. Resonant frequencies were found which supported target recognition and provided a degree of aspect independence, but they did not in all cases match previously published poles found by other techniques.

Given that late time returns last on the order of nanoseconds, the identification algorithm is unable to process radar data in real time. In a realistic application of this system, the transmitted and received radar signals $\bar{y}_R(t)$ and $\bar{u}_R(t)$ can be reduced in frequency via a down converter, shown schematically in Figure 1.2. The down converter reduces the frequency of the returned radar signal by mixing it with the output of a local oscillator operating at a frequency between the radar signal and the center frequency of a band-pass filter known as the intermediate filter (IF). The output of the IF is then sampled. This receiver architecture is called super-heterodyne [40]. An alternative to a down converter is simply to sample the returned signal at its actual frequency, record it and play it back at a lower speed, effectively reducing the frequency. The reduced-frequency radar signals $\bar{y}(t_i)$ and $\bar{u}(t_i)$ from the down converter or recorder are fed into the target recognition algorithm.

1.5 Contributions

The major contribution reported herein is the novel application of estimation techniques based on Kalman filters to the problem of target recognition via Singul-

arity Expansion Method (SEM) parameters. The application of the MMAE-MAP technique to scattering predicted by the SEM is developed in detail. The algorithm is tested against sets of synthetic data which fit the theoretical description of late-time scattering according to SEM. It is also tested with measured and simulated radar data, for which the SEM frequencies are estimated. The new target recognition algorithm is shown to be resistant to additive, white, Gaussian noise.

To date, all efforts at target identification via late-time scattering involve resonance annihilation (RA) techniques [14, 18, 52, 58, 92, 118], a generalized likelihood ratio test (GLRT), [90, 91] an autoregressive filter with exogenous input (ARX), optimal instrumental variables (IV) or an autoregressive moving average with exogenous input (ARMAX) [24, 25]. The RA techniques, E-pulse, ξ -pulse, K-pulse, S-pulse, and resonance annihilation filter, rely on convolution of a predetermined pulse with the incoming data to eliminate specific frequency components. This research applies techniques based on Kalman filtering to the problem of resonance-based target recognition for the first time. This is also the first comparison of Kalman-based techniques and resonance annihilation techniques. Existing RA methods require signal-to-noise ratios (SNR) of 20 to 42 dB, depending on target aspect and number of targets, to achieve a 90 percent likelihood of correctly identifying the target [92]. For a GLRT method, the required SNR for a 90% chance of correct ID was 2 to 7 dB [90, 91]. For both RA and GLRT techniques, tests involved simple wire models. The ARX, IV and ARMAX methods have been applied only to cones and cylinders, with an eye toward cleanup of unexploded ordnance (UXO) [24, 25]. The research involving ARX, IV and ARMAX used a 25 dB SNR signal and tried to estimate frequencies as continuous variables instead of simply selecting a target type. Correct identifications ranged from 40% to 100% for ARX and 60% to 80% for IV. ARMAX results were not included, but were characterized as poor.

In a comparison between the MMAE-MAP target recognition technique and the E-Pulse, the former identified signals in more than 90% of cases when the SNR

was above -7 dB, while the E-Pulse required an SNR above 12 dB for similar results. This comparison used data that was not representative of actual targets, as explained in Sections 3.3.1, 4.3 and 4.5, because the E-Pulse had poor results with simulated scattering data.

The MMAE-MAP technique successfully identified targets based on their simulated scattering. Targets examined in this research included cones, cylinders, rectangular blocks, and a square plate, all of which were modeled as perfect electrical conductors. The MMAE-MAP technique was often able to identify targets correctly with SNR of 3 dB, although in some cases an SNR of 7 dB was required. In a few cases, a higher SNR was required to achieve a 90% or greater successful rate of identification, although the results may be improved by additional filter tuning experiments.

Another contribution was developing an improved version of the Matrix Pencil Method (MPM) which combines the Low-Rank Hankel Approximation (LRHA) used in the Modified Matrix Pencil Method 3 (MMP3) [66,69] with the Total-Least-Squares (TLS) version of MPM [51]. The new version, Modified TLS-MPM, decreases the estimation error in some low-SNR cases by 10 to 15 dB compared to the TLS-MPM, at the cost of a significantly increased number of computations. It shows a slight improvement in estimation error compared to the MMP3 and reduces the computational load slightly. It was found that four iterations are sufficient to gain essentially all of the improvement that can be gained from the LRHA.

The LRHA is also able to suppress additive, white, Gaussian noise in signals consisting of a sum of damped sines, with a reduction of 11.7 dB observed at many noise levels. The LRHA has been characterized more completely here than in previous literature, and recommendations have been made for its use.

1.6 Organization of Report

Chapter II presents background information on the areas of research presented here. Chapter III discusses signal processing research applied toward developing a library of target signatures. In Chapter IV, stochastic estimation techniques used to compare signature data to the library of signatures are presented. Conclusions are summarized at the end of each chapter and in Chapter V. Additional work performed during this research which was not central to the problem of target recognition is presented in appendices.

II. Background

2.1 Notation, Units and Some Definitions

Mathematical notations used in this document include: $\bar{x}$ is a column vector, $\bar{x}^*$ is its complex conjugate, $\bar{x}^T$ is its transpose (always non-conjugate unless followed by an asterisk), and $\bar{x}^{T*}$ is its conjugate or Hermitian transpose. An estimate of the correct value of $\bar{x}$ is denoted $\hat{\bar{x}}$. Vector $\tilde{r}$ is a unit column vector in the direction of $\bar{r}$, and $\mathbf{A}$ is a matrix. Superscripts i , r , and p indicate that a vector is an incident or reflected (scattered) field, or a field's polarization, respectively.

A primed variable, such as $\bar{r}'$, has the same domain as its unprimed counterpart, $\bar{r}$ in this case. The primed and unprimed variables are independent. Primed and unprimed variables are used in Green's functions to take into account two variables, typically position vectors of an energy source and a point of observation. Typically either the primed or unprimed variable will be a variable of integration while the other will not. $\mathbf{A}_{m \times n}$ indicates the matrix $\mathbf{A}$ has m rows and n columns.

A variable with superscript L (e.g., $\Phi^L(s)$) is the Laplace transform of the indicated variable (otherwise $\Phi(t)$). The Laplace transform is taken to be two-sided except where stated; in other words, it involves the integral from $-\infty$ to ∞ , while the inverse Laplace transform is taken along the Bromwich contour from $\Omega_B - j\infty$ to $\Omega_B + j\infty$ [10, 64]. The Laplace transform is defined as

$$\mathcal{L}\{\Phi(t)\} = \Phi^L(s) = \int_{-\infty}^{\infty} \Phi(t) \exp(-st) dt, \quad (2.1)$$

where s is a complex number. The inverse Laplace transform is

$$\mathcal{L}^{-1}\{\Phi^L(s)\} = \Phi(t) = \frac{1}{j2\pi} \int_{\Omega_B - \infty}^{\Omega_B + \infty} \Phi^L(s) \exp(st) ds \quad (2.2)$$

A variable with superscript F ($\bar{E}^F(\omega)$) is the Fourier transform of the indicated variable ($\bar{E}(t)$). The Fourier transform is also two-sided, integrated from $-\infty$ to ∞ ,

while the inverse Fourier transform is integrated along the $j\omega$ axis from $-j\infty$ to $+j\infty$ [62, 100]. The Fourier transform is defined as

$$\mathcal{F}\{\bar{E}(t)\} = \bar{E}^F(\omega) = \int_{-\infty}^{\infty} \bar{E}(t) \exp(-j\omega t) dt \quad (2.3)$$

where ω is taken to be real unless otherwise specified. The inverse Fourier transform is

$$\mathcal{F}^{-1}\{\bar{E}^F(\omega)\} = \bar{E}(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \bar{E}^F(\omega) \exp(j\omega t) d\omega \quad (2.4)$$

The Fourier transform and its inverse can be defined with different limits of integration. If the infinite limits are replaced with a finite time or frequency window, i.e. replace ∞ and $-\infty$ with w and $-w$, the Fourier transform and its inverse can be applied to limited sets of data. The transform assumes that whatever was presented in the band $(-w, w)$ is repeated outside the band. That leads to the notion of zero padding to extend the bandwidth of data, discussed further in Appendix D, Section D.2.

When discussing measured radar data, the term ‘broadside’ is used. Broadside incidence occurs when a major feature such as the flat face of a plate or the length of a long cylinder is perpendicular to the direction of propagation of the incident radar wave. Picture yourself looking into a mirror. If you can see your own face, you see the broadside of the mirror. Broadside is used as an azimuth reference. Broadside illumination occurs when the incident wave approaches the object’s surface in a direction parallel to the surface normal, but in the opposite direction. In broadside illumination, much (but not all) of the energy is scattered from the object in the direction of the surface normal. If you see another object in the room reflected in the mirror, then the aspect of the mirror you are seeing is off-broadside at some angle away from the surface normal. Coordinate systems for all targets used in this research are discussed in Appendices B and C.

There are several synonymous terms for the function of the algorithm developed in this research. The process could be termed identification, classification, recognition, characterization, etc. Identification, in United States military usage, means deciding whether a target is friendly, hostile, or neutral. That decision is beyond the scope of this research. The algorithm developed in this research categorizes a target according to type, which relates to the military term ‘recognition’ [55]. The terms recognize, identify, classify, and characterize will be used interchangeably in this document. The word ‘target’ is used interchangeably with the word ‘object’ to signify a physical object being deliberately illuminated by a radar.

2.1.1 Complex Variables. The complex frequency $s = \Omega + j\omega$ has units of nepers per second for the real (damping) component Ω , and radians per second for the imaginary (oscillating) component $j\omega$ [4, 60]. The oscillating component of frequency, ω , has a ‘cyclic’ companion quantity $f = \omega/(2\pi)$ defined in terms of cycles per second or hertz. In the literature, the damping component of frequency Ω has no comparable ‘cyclic’ quantity $F = \Omega/(2\pi)$. Since damping and oscillating are joined in the complex frequency $s = \Omega + j\omega$, a consistent description for them based on either radians or cycles is needed. Since it is often more convenient to describe the oscillating component of frequency in terms of hertz, gigahertz, etc., a ‘cyclic’ unit for damping is needed as well. That then allows the convenient relation $s = 2\pi(F + jf) = 2\pi\Xi$. While there is no recognized unit for Ξ , some authors apply hertz to it as well as to f . That convention is followed in this document.

The term pole location, or pole, refers to a complex frequency where the unknown (field) quantity blows up to infinity. Poles of complex-valued functions appear in the form $(s - s_n)^{-p}$.

Resonant frequency and natural frequency are the same thing. Each term refers to poles associated with an object. Energy on the surface of the object resonates according to the object’s resonant or natural frequency.

To find the effect of the pole in the time domain, the inverse Laplace transform, Equation (2.2), is applied, requiring the integration of a complex function with respect to a complex variable. Each pole is transformed into a time-dependent term, such as a scattering or current mode, in a series. If integration is performed around the pole by Cauchy Residue Theorem [64,125], the resulting value is the residue. In the time domain, the residue is the coefficient associated with the pole.

For electromagnetic quantities in the Laplace domain, a time variation term $\exp(st)$ is assumed, equivalent to the $\exp(j\omega t)$ typically assumed for narrowband electromagnetic problems. For quantities in the Fourier domain, $\exp(j\omega t)$ is assumed.

2.1.2 Bandwidth. The terms narrowband, wideband, and ultra-wideband (UWB) are a source of confusion in scientific and technical literature. Throughout this work, the term ultra-wideband refers to signals with frequency content covering a decade or more of nominal bandwidth. In other words, the nominal highest frequency is ten or more times the nominal lowest frequency. Here, the 3 dB (half-power) level is used to define the nominal bandwidth, so the signal does have some power outside the nominal band. Absolute bandwidth can be considered, such as 2 gigahertz, or fractional bandwidth, in which the absolute bandwidth is divided by the center frequency and displayed as a percentage. A system with a decade of bandwidth will have a fractional bandwidth of approximately 163%. The United States Federal Communications Commission (FCC) defines a UWB system as one with greater than 500 MHz of bandwidth or 20% fractional bandwidth, using the -10 dB power limits to define high and low frequencies [32]. The FCC ‘ultra-wide’ definition in Section 15.503 of the FCC Rules and Regulations is in fact quite narrow when compared to the short-pulse generators in the literature [108]. A third measure of bandwidth, more appropriate to this research, is band ratio. Band ratio is the highest nominal frequency divided by the lowest. Therefore a system with a decade of bandwidth has a band ratio of 10.

In this report a narrowband signal is defined as one with less than 1% fractional bandwidth, measured from the -20 dB points. Any signal with fractional bandwidth greater than 1% and less than 163% will herein be considered wideband. Although FCC defines UWB, it does not define either narrowband or wideband. The 1% definition used here is comparable to FCC rules 15.231, 15.237 through 15.241, and 15.247, which allow devices with similar but slightly narrower bandwidths to use frequencies in various bands [32].

Short-pulse UWB signals are not swept through a band one frequency at a time. The bandwidth in the UWB case refers to instantaneous bandwidth, in which all frequencies in the band are present in a single pulse. The shorter in time a pulse is, the wider its bandwidth in the frequency domain.

This work builds on research in electromagnetic scattering, signal processing, and stochastic estimation and control. Background information in each of these areas is presented in Sections 2.2, 2.3 and 2.4 respectively.

2.1.3 Linear Algebra Terms. If a matrix has as many rows as columns, it is square. If it has more rows, it is called thin or overdetermined. If it has more columns than rows, it is called fat or underdetermined. A square matrix $\mathbf{A}$ is unitary if the transpose of its conjugate, $\mathbf{A}^{T*}$, equals its inverse, $\mathbf{A}^{-1}$. A unitary matrix has a determinant with a magnitude of 1 [45]. Separate symbols for conjugate ($\mathbf{A}^*$) and transpose ($\mathbf{A}^T$) are used, and both are employed when the conjugate transpose is desired ($\mathbf{A}^{T*}$). The conjugate transpose is also called the Hermitian transpose as well as the adjoint matrix. A real square matrix $\mathbf{A}$ is called normal if $\mathbf{A}\mathbf{A}^{T*} = \mathbf{A}^{T*}\mathbf{A}$.

A square matrix $\mathbf{A}$ composed of column vectors $\bar{A}_i$ is orthogonal if $\bar{A}_i^T \bar{A}_j = 0$ for all $i \neq j$. It is orthonormal if $\bar{A}_i^T \bar{A}_j = 0$ when $i \neq j$ and $= 1$ when $i = j$, or, using the Kronecker delta function, $\bar{A}_i^T \bar{A}_j = \delta_{ij}$. Put differently, the matrix is orthonormal if $\mathbf{A}^T \mathbf{A} = \mathbf{A} \mathbf{A}^T = \mathbf{I}$, where $\mathbf{I}$ is the identity matrix.

Eigenvalues and eigenvectors are solutions of the eigenvalue problem $\mathbf{A}\bar{x} = \lambda\bar{x}$, where $\mathbf{A}$ is a known, real, square matrix, λ are the complex eigenvalues and $\bar{x}$ are the complex eigenvectors. Generalized eigenvalues and eigenvectors are solutions of the generalized eigenvalue problem $\mathbf{A}\bar{x} = \mathbf{B}\lambda\bar{x}$, where $\mathbf{B}$ is a known, real, square matrix of the same size as $\mathbf{A}$.

Square matrices have eigenvalues, while rectangular matrices do not. For a rectangular matrix, singular values are found instead. They are related this way: the singular values of a complex matrix $\mathbf{C}$ are the real, nonnegative square roots of the eigenvalues of $\mathbf{C}^{T*}\mathbf{C}$. Both $\mathbf{C}^{T*}\mathbf{C}$ and $\mathbf{C}\mathbf{C}^{T*}$ are positive semidefinite matrices [38, 46], meaning that all of their eigenvalues are positive or zero. Right singular vectors of a complex, rectangular matrix $\mathbf{C}$ are the same as the eigenvectors of the matrix $\mathbf{C}^{T*}\mathbf{C}$, while left singular vectors of $\mathbf{C}$ are the eigenvectors of $\mathbf{C}\mathbf{C}^{T*}$.

If a matrix is square and positive definite, a square root can be computed for it. Positive definite means the eigenvalues of the matrix are all greater than zero. While a real scalar has only one square root of ambiguous sign, a matrix has many possible formulations of the square root. In this research the Cholesky factorization [77] is used to compute the square root $\sqrt[\mathcal{C}]{\mathbf{A}}$ of the positive definite matrix $\mathbf{A}$ such that

$$\sqrt[\mathcal{C}]{\mathbf{A}} \sqrt[\mathcal{C}]{\mathbf{A}}^T = \mathbf{A} \quad (2.5)$$

where $\sqrt[\mathcal{C}]{\mathbf{A}}$ has the same number of rows and columns as $\mathbf{A}$. Note that Equation (2.5) does not uniquely define the Cholesky square root. The Cholesky square root is a lower triangular matrix. If a matrix has a single element, the Cholesky square root is the same as the square root of a scalar.

2.1.4 Dyadic Operations. Vectors can be multiplied in three ways that produce either a scalar, vector, or dyadic answer. It is common in electromagnetics literature to display matrix multiplications in terms of the dot, cross, and dyadic

products. However, in signal processing and stochastic estimation literature, it is more common to use linear algebra notation.

The dyadic product, $\mathbf{E} = \bar{a}\bar{b}$, or outer product of two vectors $\bar{a}$ and $\bar{b}$ is a matrix formed by transposing the second vector and multiplying.

$$\mathbf{E} = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} \begin{pmatrix} b_1 & b_2 & b_3 \end{pmatrix} = \begin{pmatrix} a_1 b_1 & a_1 b_2 & a_1 b_3 \\ a_2 b_1 & a_2 b_2 & a_2 b_3 \\ a_3 b_1 & a_3 b_2 & a_3 b_3 \end{pmatrix} \quad (2.6)$$

If two vectors are shown side-by-side with no dot or cross between them, as in $\bar{a}\bar{b}$, dyadic multiplication is implied. Linear algebra notation for the dyadic product of vectors $\bar{a}$ and $\bar{b}$ is $\mathbf{E} = \bar{a}\bar{b}^T$.

2.2 Scattering Theory

Any object illuminated by an electromagnetic wave will scatter some of the energy in, generally speaking, all directions. Actual prediction of the scattering magnitude in each direction is difficult, but can be accomplished exactly for a limited set of scatterer geometries through Green's functions [129], plane wave spectra [27] and modal solutions to boundary value problems [4, 23, 116]. However, for scattering from real, three-dimensional, objects of interest, approximate techniques are typically used, and these fall into the asymptotic and computational categories.

Asymptotic techniques improve in accuracy as frequency increases. They include geometric optics (GO) and its extensions: the geometric theory of diffraction (GTD) and the uniform geometric theory of diffraction (UTD) [83]. Other asymptotic techniques are physical optics (PO) and its extension, the physical theory of diffraction (PTD) [4, 83]. Asymptotic models have been used to gain insight into measured data by identifying scattering centers, which are features on an object which reflect significant energy to the radar. At sufficiently high frequencies, an object may be approximated as a collection of scattering centers [34, 107].

Computational techniques solve Maxwell’s equations numerically, over small segments of the scatterer, along with the interaction, or mutual impedance, between each segment. The equations may be solved in differential or integral form and may be formulated as either time domain or frequency domain equations. Scattering from a target can be regarded as a time-invariant transfer function, so the time-domain and frequency-domain solutions can be related to each other via the Fourier transform. Computational techniques include the method of moments (MoM), in which integral equations are discretized, and techniques which discretize differential equations, such as finite difference (FD), finite volume (FV) and the finite element method (FEM), which can be formulated in either the frequency or time domain [105]. Computational techniques simulate electromagnetic interaction (propagation, scattering, surface and volume currents, etc.) to a high degree of fidelity, so they are also known as scattering simulations.

The results from one analytic technique, the Singularity Expansion Method (SEM) [6,9], are the theoretical basis for the target recognition algorithm presented herein. Various authors [110, 110, 114, 124, 131] have used SEM in conjunction with computational methods to identify natural resonances of targets.

2.2.1 The Singularity Expansion Method. Measurements involving illumination of simple shapes by ultra-wideband, short-pulse radar have led to an interesting observation. The strong initial return from an object was followed by a weaker signal that took the form of a damped sinusoid or sum of damped sinusoids. Simulated scattering from a cylinder is shown in Figure 2.1. The late-time portion of the signal is shown in a different scale in Figure 2.2, where the domination of the return by a single damped sinusoid is obvious. The frequency of oscillation of the damped sine shown in Figures 2.1 and 2.2 is 665 MHz. With the same target and azimuthally-polarized illumination, a more complicated late-time return occurs, indicating the presence of different significant, damped-sine terms in the signal.




Figure 2.1 Simulated time-domain scattering from a cylinder, six inches long and two inches in diameter, axially-polarized illumination with bandwidth 0.1 to 18 GHz




Figure 2.2 Simulated time-domain scattering from a cylinder, six inches long and two inches in diameter, axially-polarized illumination with bandwidth 0.1 to 18 GHz, emphasizing late-time scattering

This observation led to attempts to formulate an analytic expression for this phenomenon. The Singularity Expansion Method (SEM), developed by Baum in 1971 [6, 9], extends scattering theory to account for ultra-wideband illumination.

SEM has been applied to a limited number of targets so far. There are theoretical results for a sphere [6] and wire loops [20, 134], plus results based on MoM or other integral equations for a variety of other targets: polygonal wire loops [124], elliptical wire loops [114], wire dipoles [110, 110, 131], cylinders, rectangular plates, rectangular boxes, etc. Modal solutions have been found for bodies of revolution [86]. Most of these targets were assumed to be isolated in free space. Some solutions exist for wires in the presence of ground planes [133], objects in or near lossy, dielectric media [12, 37, 42], coupled system consisting of two wires or two loops [43], etc.

Electromagnetic scattering due to narrowband illumination is usually approximated as single-frequency, or zero-bandwidth, illumination which must have infinite time duration according to Fourier transform theory [100]. Conversely, in SEM, scattering from ultra-wideband, short-pulse sources is approximated by scattering from an impulse with infinite bandwidth and zero time duration. A target's response (surface current, scattered field, current in a circuit or waveguide, etc.) to an impulsive forcing function (incident wave, impressed voltage or current, etc.) is treated as a series of isolated singularities in the complex frequency $s = \Omega + j\omega$ plane. Isolated singularities include poles of integer order, essential singularities, and branch points [64]. Branch points may be of integer order, or may be algebraic, logarithmic, or transcendental [125].

To clarify time-domain and frequency-domain concepts, representations of the Dirac delta function in frequency and time domains are shown in Figures 2.3 and 2.4. The dotted lines in the two figures are Fourier Transform pairs, as are the solid lines. Figure 2.4 shows a time domain plot in which the solid line represents a delta function $\delta(t)$. The line extends upward at $t = 0$ to the top of the graph, while the function actually extends upward to infinity; the area under the function is unity.




Figure 2.3 Two Dirac Delta Functions (dotted) and constant (solid), Frequency Domain




Figure 2.4 Dirac Delta Function (solid) and cosine (dotted), Time Domain

The dotted line is a sinusoid, which is the time domain transform of the pair of delta functions in the frequency domain, shown in Figure 2.3. In the frequency domain plot, the sum of two delta functions, $\delta(\omega - 25) + \delta(\omega + 25)$, is shown with dotted lines extending toward infinity at $\omega = -25$ and $\omega = 25$. These two delta functions represent power concentrated exactly at single frequencies. The resulting time domain signal in Figure 2.4 is a single frequency sinusoidal signal that extends for all time, from $-\infty$ to ∞ . The solid line in Figure 2.3 is a constant. It is the frequency domain transform of the delta function, solid line, in Figure 2.4.

To develop SEM, it is useful to start with a description of forced response. The forced response of an object of volume V to electromagnetic stimulation $\bar{F}^L$, which may be an incident wave, an impressed current, etc., is [7]:

$$\int_V \mathbf{K}^L(\bar{r}, \bar{r}', s) \cdot \bar{J}^L(\bar{r}', s) dV' = \bar{F}^L(\bar{r}, s) \quad (2.7)$$

where $\bar{r}$ and $\bar{r}'$ are position vectors to any point in the volume V of the object, s is the Laplace-domain frequency, and the equation kernel $\mathbf{K}^L$ is based on the Green's function of the target and its surroundings. The forcing function $\bar{F}^L$ and the object's physical configuration determine the response $\bar{J}^L$, which may be a current or a scattered field. An equation of this form is called a Fredholm integral equation of the first kind [127]. If the object is a perfect electrical conductor (PEC), the integral is taken over the surface area D rather than the volume V , and the position vectors trace the surface rather than the volume of the object.

Now consider an unforced response. A PEC object which is not illuminated obeys this homogeneous integral equation:

$$\int_D \mathbf{K}^L(\bar{r}, \bar{r}', s) \cdot \bar{J}^L(\bar{r}', s) dD' = \bar{0} \quad (2.8)$$

in which the integral is taken over the surface D instead of the volume V . Every nontrivial solution to this equation is an unforced response. Unforced responses are any responses which may exist ($\bar{J}^L \neq \bar{0}$) in the absence of excitation. The kernel is

dependent on the frequency s . The unforced response may exist for some values of s and not others. The values for which an unforced response exists are called natural frequencies. Because excitation may be absent, neither the polarization nor angle of incidence of an exciting source has an effect on these natural frequencies. The total unforced scattering may be found from the sum of the scattering due to each natural-frequency term.

Unforced response includes response after excitation has ceased. Natural modes that make up the unforced response represent energy stored, briefly, on the surface of the object. The energy was deposited there by an incident wave or other excitation. An object's energy-storing ability is determined by the object itself, although the amount of energy stored is related to the amount of energy added to the system by an incident field. That means that some quantities in a mathematical representation of unforced scattering are functions of the incident wave, but the natural frequencies themselves are not [6].

Scattering from a finite-sized, perfectly conducting object in free space due to an impulse is found to be an infinite series of first order poles plus an entire function [6, 9]. The first-order poles, $(s - s_n)^{-1}$ correspond to damped sinusoids ($\exp(s_n t)u(t)$) in the time domain via the Laplace transform, as demonstrated by:

$$\begin{aligned}
(s - s_n)^{-1} &= \int_0^\infty \exp[-(s - s_n)t]dt \\
&= \int_0^\infty \exp(s_n t) \exp(-st)dt \\
&= \int_0^\infty \exp(-\Omega_n t) \exp(-j\omega_n t) \exp(-st)dt \\
&= \int_{-\infty}^\infty \exp(-\Omega_n t) [\cos(j\omega_n t) - j \sin(\omega_n t)]u(t) \exp(-st)dt \\
&= \mathcal{L}\{\exp(-\Omega_n t) [\cos(j\omega_n t) - j \sin(\omega_n t)]u(t)\} \tag{2.9}
\end{aligned}$$

An entire function is a function with no singularities in the finite s plane, but which may have one at infinity [64, 123]. If the object is on or near a ground plane or embedded in a lossy dielectric volume, two branch points of order 1 are introduced in the s -plane, with a branch cut between them [37]. A branch point is a point at which a complex-valued function may have more than one value [64]. The branch points can cause discontinuity in the impedance or admittance matrices developed in integral-equation electromagnetics simulations. Giri [37] analytically predicted branch points in the kernel function $\mathbf{K}(r, r', s)$ for a dipole antenna in a lossy medium, but found their effect was negligible for the cases studied.

If a PEC object is resistively loaded, the SEM series includes second order poles, $(s - s_n)^{-2}$, corresponding to damped sinusoids multiplied by ramps ($t \exp(s_n t)$) in the time domain. Given targets constructed of metal, the effect of the second order poles should be negligible compared to the effect of the first-order poles [132].

A plane wave in space can be defined by its electric field. In the frequency domain, the field is

$$\bar{E}^{iL}(\bar{r}, s) = E_0 f^L(s) e^{-\gamma \tilde{r}^i \cdot \bar{r}} \tilde{r}^p \quad (2.10)$$

while in the time domain, it is

$$\bar{E}^i(\bar{r}, t) = E_0 f\left(t - \frac{\tilde{r}^i \cdot \bar{r}}{c}\right) \tilde{r}^p \quad (2.11)$$

where $f(t)$ is the waveform, and E_0 is the magnitude of the plane wave. The position vector $\bar{r}$ refers to any point in space at which the plane wave is observed at a given time t . The unit vector $\tilde{r}^i$ indicates the direction of propagation of the incident waves. The polarization direction, or direction of the electric field, is given by $\tilde{r}^p$. The propagation constant γ is equal to s/c in free space, where c is the speed of light. An impulse is defined as the Dirac delta function $f(t) = \delta(t)$ in the time domain, which has the transform $f^L(s) = 1$ in the Laplace frequency domain. Current density $\bar{J}$ on the object's surface D is related to the incident field via

$$\bar{E}^{iL}(\bar{r}_D, s) = \int_D \mathbf{K}^L(\bar{r}_D, \bar{r}'_D, s) \cdot \bar{J}^L(\bar{r}'_D, s) d\bar{r}'_D \quad (2.12)$$

in which $\mathbf{K}^L(\bar{r}_D, \bar{r}'_D, s)$ is a dyadic kernel related to the Green's function for the object [8,18], and $\bar{r}'_D$ is the position vector to any point on the object's surface. This is a restatement of Equation (2.7) in which the forcing function is an electric field and the response is a surface current. Since $\bar{E}^{iL}(\bar{r}_D, s)$ is time limited, there is a period when equation (2.12) is unforced.




Figure 2.5 Scatterer Illuminated by Plane Wave

Now consider a time domain representation of the incident signal, an impulse. The duration of the pulse is zero, so the position of the wave relative to the target can be defined precisely. The time period can be distinguished during which the incident pulse passes over the object from the time after the pulse has passed. After the forcing function ceases, i.e., after the incident wave has passed, the object exhibits unforced response, or ringing, for a brief period called “late time”, explained in Section 2.2.2. In an actual radar application, an ultra-wideband, short-duration

pulse would be used in place of the theoretical impulse. Short pulses produced in the laboratory have rise times on the order of picoseconds, and durations of a few nanoseconds. If a long radar pulse were used, the ringing in late-time could be confounded with early-time and the information in late-time ringing would be unusable.

Assuming that $f^L(s) = 1$, the solution to the surface currents induced on the object by the field in Equation (2.12) is given by [9, 18]:

$$\bar{J}^L(s) = E_0 \sum_n \eta_n(\tilde{r}^i, \tilde{r}^p) \bar{v}_n(\bar{r}_D)(s - s_n)^{-p_n} \exp(st_0 - s_n t_0) + \bar{F}^L(\tilde{r}^i, \tilde{r}^p, \bar{r}, s) \quad (2.13)$$

where η_n is the coupling coefficient and $\bar{v}_n$ is the natural mode, normalized to a maximum value of one. The coefficient contains magnitude and phase information about the natural frequency response, while the mode gives its direction and distribution on the object's surface. The turn-on time t_0 defines the point in the time domain after which the SEM representation becomes valid. The integer p_n is the order of the pole, and is equal to one for finite, perfectly electrically conducting objects. $\bar{F}$ is an entire function which may have a value of zero. It is the sum of the entire functions associated with each pole term according to Mittag-Leffler's Theorem [125]. This theorem states that a partial fraction expansion of a meromorphic function consists of an entire function plus a series of pole terms, each accompanied by a polynomial [18, 125]. The vector $\bar{r}_D$ is the position of any point on the object's surface (if conducting) or volume (if dielectric). The origin is at the center of a notional sphere of minimum radius to enclose the target completely, as shown in Figure 2.5. The direction of propagation of the incident fields, $\tilde{r}^i$, and the scattered fields, $\tilde{r}^r$, are also illustrated in Figure 2.5.

The coupling coefficient is [18]

$$\eta_n(\tilde{r}^i, \tilde{r}^p) = \frac{\tilde{r}^p \cdot \int_D \exp(-\gamma \tilde{r}^i \cdot \bar{r}_D) \bar{\nu}_n(\bar{r}_D) d\bar{r}_D}{\int_D \int_D \bar{\nu}_n(\bar{r}_D) \cdot \frac{\partial}{\partial s} \mathbf{K}(\bar{r}_D, \bar{r}'_D, s) |_{s=s_n} \cdot \bar{\nu}_n(\bar{r}'_D) d\bar{r}'_D d\bar{r}_D} \quad (2.14)$$

and the mode is defined by:

$$\int_D \mathbf{K}^L(\bar{r}_D, \bar{r}'_D, s) \cdot \bar{\nu}_n(\bar{r}'_D) d\bar{r}'_D = \bar{0} \quad (2.15)$$

The surface currents on the scatterer give rise to a scattered electric field given by:

$$\bar{E}^{rL}(\bar{r}, s) = \frac{E_0 \exp(-\gamma \tilde{r}^r \cdot \bar{r})}{4\pi r} \sum_n f^L(s_n) \Lambda_n \mathbf{M}_n \cdot \tilde{r}^p (s - s_n)^{-p_n} e^{(s-s_n)t_0} + \bar{G}^L(\tilde{r}^i, \tilde{r}^p, \bar{r}, s) \quad (2.16)$$

Here $\bar{G}^L$ is the scattering due to the entire function $\bar{F}^L$. In Equation (2.16), the scattering coefficient Λ_n , and the scattering matrix $\mathbf{M}_n$ are defined as

$$\Lambda_n = -s_n \mu_0 \left[\int_D \int_D \bar{\nu}_n(\bar{r}_D) \cdot \frac{\partial}{\partial s} \mathbf{K}(\bar{r}_D, \bar{r}'_D, s) |_{s=s_n} \cdot \bar{\nu}_n(\bar{r}'_D) d\bar{r}'_D d\bar{r}_D \right]^{-1} \quad (2.17)$$

$$\mathbf{M}_n = \left[\int_D \exp(-\gamma_n \tilde{r}^i \cdot \bar{r}_D) \mathbf{T}^i \cdot \bar{\nu}_n(\bar{r}_D) d\bar{r}_D \right] \left[\int_D \exp(\gamma_n \tilde{r}^r \cdot \bar{r}_D) \mathbf{T}^r \cdot \bar{\nu}_n(\bar{r}_D) d\bar{r}_D \right] \quad (2.18)$$

Integrals are taken over the scatterer's surface, D . The transverse matrices found in $\mathbf{M}_n$ are

$$\mathbf{T}^i = \mathbf{I} - \tilde{r}^i \tilde{r}^{iT} \quad (2.19)$$

$$\mathbf{T}^r = \mathbf{I} - \tilde{r}^r \tilde{r}^{rT} \quad (2.20)$$

where $\mathbf{I}$ is a 3×3 identity matrix. The two unit vectors are joined in a dyadic product, $(\tilde{r}^i \tilde{r}^{iT})$ or $(\tilde{r}^r \tilde{r}^{rT})$, also called an outer product, which produces a rank-one matrix. Vectors are assumed to be column vectors, so the incident direction $\tilde{r}^i$ can be defined as

$$\tilde{r}^i = \begin{pmatrix} r_x^i \\ r_y^i \\ r_z^i \end{pmatrix} \quad (2.21)$$

Here Cartesian (x,y,z) coordinates are shown, but any orthogonal system such as spherical, cylindrical, etc., could be used. The outer product of the two vectors is a dyad:

$$\tilde{r}^i \tilde{r}^{iT} = \begin{pmatrix} r_x^i r_x^i & r_x^i r_y^i & r_x^i r_z^i \\ r_y^i r_x^i & r_y^i r_y^i & r_y^i r_z^i \\ r_z^i r_x^i & r_z^i r_y^i & r_z^i r_z^i \end{pmatrix} \quad (2.22)$$

The transverse $\mathbf{T}$ matrix is similar to an orthogonal matrix transformation called the Householder Reflection [38]:

$$\mathbf{P} = \mathbf{I} - 2 \frac{\bar{v} \bar{v}^T}{\bar{v}^T \bar{v}} \quad (2.23)$$

Using the inverse Laplace transform, the time domain scattered field is

$$\bar{E}^r(\bar{r}, t) = \sum_n \bar{C}_n e^{s_n(t-r/c+t_0)} u(t-r/c+t_0) + \bar{G}(\tilde{r}^i, \tilde{r}^p, \bar{r}, t) \quad (2.24)$$

where

$$\bar{C}_n = \frac{E_0}{4\pi r} \Lambda_n \mathbf{M}_n \cdot \tilde{r}^p \exp(-s_n t_0) \quad (2.25)$$

and the incident field is assumed to be an impulse, i.e., $f^L(s) = 1$.

The time axis ($t_l = t - r/c + t_0$) is now shifted so the origin is the start of retarded late time; that is, the late-time portion of the scattered signal as seen by

a radar receiver. During late time, the entire function $\bar{G}$ is identically equal to zero [15]. The scattered field is then

$$\bar{E}^r(\bar{r}, t_l) = \sum_n \bar{C}_n \exp(s_n t_l) u(t_l) \quad (2.26)$$

The complex natural frequency is $s = \Omega + j\omega$, where Ω and ω are both real. Ω is the decay rate, or growth rate if positive, and ω is the frequency of oscillation. Thus

$$\exp(st_l) = \exp(\Omega t_l) \exp(j\omega t_l) = \exp(\Omega t_l) [\cos(\omega t_l) + j \sin(\omega t_l)] \quad (2.27)$$

In this document, complex-valued fields are used to account for phase shifts in the frequency domain. In the time domain, all fields are considered real, so all imaginary terms must sum to zero. Thus every pole with nonzero ω is part of a conjugate pair; another pole k has the complex conjugate of both the frequency ($s_k = s_n^* = \Omega_n - j\omega_n$) and the coefficient ($\bar{C}_k = \bar{C}_n^*$). If a pole has $\omega_n = 0$, it can be treated as the sum of two identical poles, each with a real coefficient. The scattering term associated with each conjugate pair or real-axis pole is called a mode. Because negative oscillating frequencies are treated explicitly as $-j\omega_n$, unsigned oscillating frequencies ω_n can be considered to be positive. So Equation (2.24) can be put in the following form, *now summing over modes m instead of poles n* :

$$\bar{E}^r(t_l) = \sum_m [\bar{C}_m \exp(\Omega_m t_l + j\omega_m t_l) + \bar{C}_m^* \exp(\Omega_m t_l - j\omega_m t_l)] u(t_l) \quad (2.28)$$

The complex coefficients can be separated into rectangular components via $\bar{C}_m = \bar{A}_m + j\bar{B}_m$ and $\bar{C}_m^* = \bar{A}_m - j\bar{B}_m$, where $\bar{A}_m$ and $\bar{B}_m$ are both real, so Equation (2.28) can be rewritten as

$$\begin{aligned}
E^r(t_l) &= \sum_m \exp(\Omega_m t_l) u(t_l) [\bar{A}_m \exp(j\omega_m t_l) \\
&\quad + j\bar{B}_m \exp(j\omega_m t_l) + \bar{A}_m \exp(-j\omega_m t_l) - j\bar{B}_m \exp(-j\omega_m t_l)] \\
&= \sum_m \exp(\Omega_m t_l) u(t_l) \cdot \\
&\quad \{ \bar{A}_m [\cos(\omega_m t_l) + j \sin(\omega_m t_l)] + j\bar{B}_m [\cos(\omega_m t_l) + j \sin(\omega_m t_l)] \\
&\quad + \bar{A}_m [\cos(j\omega_m t_l) + \sin(-j\omega_m t_l)] - j\bar{B}_m [\cos(j\omega_m t_l) - j \sin(-j\omega_m t_l)] \} \\
&= 2 \sum_m \exp(\Omega_m t_l) [\bar{A}_m \cos(\omega_m t_l) - \bar{B}_m \sin(\omega_m t_l)] u(t_l)
\end{aligned} \tag{2.29}$$

In the complex frequency plane (the s -plane or Laplace domain), each term of Equation (2.29) corresponds to a conjugate pair of first order poles, and the weighting coefficients $\bar{A}_m$ and $\bar{B}_m$ are the real and imaginary parts of the residues of each pole. Transformed into the time domain, the effect of the two poles is the damped and phase-shifted sinusoid in each term of Equation (2.29). If second order poles are present, some terms in the time domain series would be of the form $t \exp(\Omega_m t_l) [\bar{A}_m \cos(\omega_m t_l) - \bar{B}_m \sin(\omega_m t_l)]$, or $\bar{A}_m t_l \exp(\Omega_m t_l)$ for poles on the $-\Omega$ axis.

The coefficients $\bar{C}_m$, $\bar{A}_m$, and $\bar{B}_m$ are shown as vectors because they have components corresponding to the polarization of the scattered field $\bar{E}^r(t_l)$. If the radar receiving the scattered signal has a linearly-polarized antenna, the signal apparent to the radar will have only a single polarization component. The received signal may then be treated as a scalar $E^r(t_l)$. The coefficients A_m and B_m of the signal will then be real scalars, while C_m and C_m^* will, of course, be complex scalars. In the remainder of this document, the coefficients of each scattering mode are treated as scalars, with the understanding that they are dependent on polarization, so that a

target measured vertically and horizontally will have different coefficients on each scattering mode in the vertical and horizontal cases.

2.2.2 Late Time. Fundamental to SEM is the notion of early, middle, and late time of the scattering return. Early time is the time required for the incident pulse to transit the target once from front to back. SEM predicts that during early time, the series of damped sinusoids will not converge. This makes it difficult to draw useful conclusions about early-time scattering from SEM, but some work has been done on defining the entire function shown in Equation (2.13) [15,17]. Middle time is the period during which the incident pulse transits back from the far end of the target to the near end. During middle time, the pole series will converge, but is not necessarily unique; the entire function will be zero [11,15]. Late time is after the return transit. In late time, the incident wave is no longer present, so the only scattering is due to the unforced, natural response of the target. During late time, the series converges and is unique. Convergence allows us to approximate the late-time scattering with a finite sum of terms [17,18]. A key result of SEM is that the natural frequency of each term in the pole series is independent of target aspect, although coefficients A_m and B_m in Equation (2.29) are not [9].

The incident wave travels along the surface of a perfect electrical conductor, or through a dielectric volume. In Figure 2.5, the incident wave travels along the ray path, which is a combination of creeping waves along the scatterer's surface, and diffracted waves from one peak to another on the surface. If the notional minimum circumscribing sphere in Figure 2.5 is 5 meters in radius, the total path around the target is on the order of 25 meters from front to back to front. About half of that is early time, the rest middle time. The first transit lasts approximately $12.5/c = 41.7$ nanoseconds. The middle time is from about 41.7 to 83.4 nanoseconds, and late time is more than 83.4 nanoseconds from the instant energy first arrived at the point on the target nearest the transmitter. At the receiver, late time is 83.4 nanoseconds

after the first energy reaches the receiver. For the different times to be distinct, the incident pulse must be short, ideally an impulse.

For physically realizable late-time signals, the damping time constant is negative, indicating a pole in the left half of the Laplace domain. This pole placement corresponds to a signal which loses energy over time in the absence of excitation. Clearly, the principle of conservation of energy prevents an unforced signal from gaining energy and growing in magnitude. In addition, it prevents such a signal from continuing indefinitely as energy is radiated. The pole search may be constrained further by noting that the poles in the lower half plane are conjugates, with frequency s_n^* and residue C_n^* , of poles in the upper half plane. Also, poles with a high-magnitude real component of frequency (low damping time constant) vanish too quickly to be adequately detected and identified. Therefore, only poles existing in a limited region of one quadrant of the complex plane are of concern. Even within this area, poles will be most strongly excited if the imaginary part of the natural frequency, $j\omega$, is found in the incident signal. Summarizing, the excited poles will be within the area bounded by $\Omega \in [\Omega_{min}, 0)$ and $\omega \in [0, \omega_{max}]$, and most will be in the smaller area bounded by $\Omega \in [\Omega_{min}, 0)$ and $\omega \in [\omega_{min}, \omega_{max}]$. Here ω_{min} and ω_{max} are determined by the incident signal bandwidth, and Ω_{min} is determined by the sampling rate and computational considerations.

Practical discussions of radar cross-section [59] cover several scattering components attributable to definable physical sources. These components include specular reflections, edge and tip diffraction, etc. On the other hand, analytical descriptions of scattering [72, 129] based on Green's functions do not distinguish between these sources. They only show the total scattering. Late-time scattering as described by SEM also gives only the total response in its time window. It does not distinguish between various types of ringing, although it is known as a practical matter that the various mechanisms active on a target in late time can include travelling waves, creeping waves, multiple diffractions, and cavity responses [59]. These same mecha-

nisms operate during early time as well, but may be dominated by other phenomena such as specular reflections.

2.3 *Signal Processing*

Given a set of late-time data from a target, the target's natural frequencies must be determined. Among the types of techniques that may be employed are linear predictors, maximum likelihood methods, and evolutionary algorithms. Linear predictor is a term encompassing Prony's Method [3], the Matrix Pencil Method [49], ESPRIT [115], and others, which use matrix manipulations to work backward from a data set to the signal that gave rise to it, and from there to the components that make up the signal [121]. Linear predictors could more appropriately be called linear estimators, since they estimate current, rather than future, values of variables. Maximum likelihood methods adjust parameters until errors are minimized, using search techniques such as gradient search. Maximum likelihood techniques iterate from a starting point supplied by some other analysis, possibly a linear predictor. They require far more computation than the linear predictors.

2.3.1 The Matrix Pencil Method. The Matrix Pencil Method (MPM) is an extension of Prony's method [3, 98] which can take in data on a given target and analyze it to extract the natural frequencies [49, 119]. A matrix pencil is a linear combination of two matrices, as in $\mathbf{X}_1 - \lambda \mathbf{X}_0$, in which λ is a scalar variable [51]. The term 'pencil' in MPM comes from projective geometry, in which a pencil is a set of lines which intersect at a point [28]. The term was used by Jain [53] and later by Jain, Sarkar, and Weiner [54] to group related functions.

An explanation of the theory behind MPM starts with the signal model. The signal is a sum of a finite number N of exponentials sampled at evenly spaced intervals, such as:

$$x(t_i) = \sum_{n=1}^N C_n \exp(s_n t_i) \quad (2.30)$$

where both C_n and s_n are, in general, complex numbers. The resulting signal $x(t_i)$ is real, so for every term $C_n \exp(s_n t_i)$ there is a corresponding term $C_p \exp(s_p t_i)$, $p \neq n$, for which $C_n = C_p^*$ and $s_n = s_p^*$. The Laplace-domain frequencies s_n can be transformed into z -domain frequencies using the sample period $\Delta t = t_i - t_{i-1}$ and the relations:

$$z_n = \exp(s_n \Delta t) \quad (2.31)$$

and:

$$z_n^i = \exp(s_n i \Delta t) = \exp(s_n t_i), \quad t_0 = 0 \quad (2.32)$$

where z_n is complex if and only if s_n is complex, so the signal may be defined as:

$$x(t_i) = \sum_n C_n z_n^i \quad (2.33)$$

Given a set of noise-free sampled data $x(t_i), i = 0, 1, \dots, K - 1$, the data is formed into a data matrix called a Hankel matrix [121]. A Hankel matrix has the same value in each element along each anti-diagonal, as shown in Equations (2.34) and (2.35). Given matrix elements $y_{r,c}$ at rows r and columns c , an anti-diagonal is a set of elements for which $r + c$ is a constant. The shape of the Hankel matrices is determined by the free-moving window length L [49], also known as the pencil parameter. For best estimation results, L is chosen between $K/3$ and $2K/3$ [120]. The matrices have $(K - L)$ rows and L columns. Let

$$\mathbf{X}_0 = \begin{pmatrix} x(t_0) & x(t_1) & x(t_2) & \cdots & x(t_L) \\ x(t_1) & x(t_2) & x(t_3) & \cdots & x(t_{L+1}) \\ x(t_2) & x(t_3) & x(t_4) & \cdots & x(t_{L+2}) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x(t_{K-L-1}) & x(t_{K-L}) & x(t_{K-L+1}) & \cdots & x(t_{K-2}) \end{pmatrix} \quad (2.34)$$

$$\mathbf{X}_1 = \begin{pmatrix} x(t_1) & x(t_2) & x(t_3) & \cdots & x(t_{L+1}) \\ x(t_2) & x(t_3) & x(t_4) & \cdots & x(t_{L+2}) \\ x(t_3) & x(t_4) & x(t_5) & \cdots & x(t_{L+3}) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x(t_{K-L}) & x(t_{K-L+1}) & x(t_{K-L+2}) & \cdots & x(t_{K-1}) \end{pmatrix} \quad (2.35)$$

Each element in the matrix $\mathbf{X}_1$ is from data sampled one period later than the data in the corresponding element of $\mathbf{X}_0$. It is therefore the result of the multiplication of each term in Equation (2.33) by its corresponding z_n . As an aside, most authors implement MPM with a Hankel matrix, but the technique can also be implemented with a Toeplitz matrix [49]. A matrix with Toeplitz symmetry has the same values across each diagonal, rather than each anti-diagonal. Only Hankel data matrices are used herein, to conform with the majority of the literature.

The data matrix $\mathbf{X}_0$ can be decomposed into three matrices, as:

$$\mathbf{X}_0 = \mathbf{Z}_1 \mathbf{C} \mathbf{Z}_2 \quad (2.36)$$

$$\mathbf{Z}_1 = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ z_1 & z_2 & \cdots & z_N \\ z_1^2 & z_2^2 & \cdots & z_N^2 \\ \vdots & \vdots & \ddots & \vdots \\ z_1^{N-L-1} & z_2^{N-L-1} & \cdots & z_N^{N-L-1} \end{pmatrix} \quad (2.37)$$

$$\mathbf{Z}_2 = \begin{pmatrix} 1 & z_1 & z_1^2 & \cdots & z_1^{L-1} \\ 1 & z_2 & z_2^2 & \cdots & z_2^{L-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & z_N & z_N^2 & \cdots & z_N^{L-1} \end{pmatrix} \quad (2.38)$$

and

$$\mathbf{C} = \begin{pmatrix} C_1 & 0 & \cdots & 0 \\ 0 & C_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & C_N \end{pmatrix} \quad (2.39)$$

where $\mathbf{Z}_1$ and $\mathbf{Z}_2$ are Vandermonde matrices [45]. The values C_n found on the main diagonal of $\mathbf{C}$ are the amplitude coefficients of the signal model in Equation (2.33).

In a similar fashion, $\mathbf{X}_1$ can be decomposed as:

$$\mathbf{X}_1 = \mathbf{Z}_1 \mathbf{C} \mathbf{Z}_0 \mathbf{Z}_2 \quad (2.40)$$

where $\mathbf{Z}_1$, $\mathbf{Z}_2$ and $\mathbf{C}$ are as defined in Equations (2.37), (2.38), and (2.39), respectively, and

$$\mathbf{Z}_0 = \begin{pmatrix} z_1 & 0 & \cdots & 0 \\ 0 & z_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & z_N \end{pmatrix} \quad (2.41)$$

Using Equations (2.36) and (2.40), the matrix pencil $\mathbf{X}_1 - \lambda \mathbf{X}_0$ can be rewritten as:

$$\begin{aligned} \mathbf{X}_1 - \lambda \mathbf{X}_0 &= \mathbf{Z}_1 \mathbf{C} \mathbf{Z}_0 \mathbf{Z}_2 - \lambda \mathbf{Z}_1 \mathbf{C} \mathbf{Z}_2 \\ &= \mathbf{Z}_1 \mathbf{C} [\mathbf{Z}_0 - \lambda \mathbf{I}] \mathbf{Z}_2 \end{aligned} \quad (2.42)$$

The rank of the matrix pencil is N , provided $N \leq L \leq (K - N)$, for an arbitrary value of λ [120]. However, if $\lambda = z_n$, a row of $[\mathbf{Z}_0 - \lambda \mathbf{I}]$ goes to zero, and the rank of the matrix pencil is reduced by one. The frequencies z_n may therefore be found by converting the matrix pencil $\mathbf{X}_1 - \lambda \mathbf{X}_0$ into a generalized eigenvalue problem as:

$$\mathbf{X}_0^\dagger \mathbf{X}_1 = \lambda \mathbf{I} \quad (2.43)$$

The matrix $\mathbf{X}_0^\dagger$ is the Moore-Penrose pseudo-inverse of $\mathbf{X}_0$ [38], defined via the singular value decomposition (SVD) [38, 121]. The SVD is

$$\mathbf{U}_0 \mathbf{\Sigma}_0 \mathbf{V}_0^{T*} = \mathbf{X}_0 \quad (2.44)$$

where $\mathbf{\Sigma}_0$ and $\mathbf{X}_0$ have L columns and $(K - L)$ rows. $\mathbf{U}_0$ and $\mathbf{V}_0$ are both square matrices. $\mathbf{U}_0$ has $(K - L)$ rows and columns, while $\mathbf{V}_0$ has L rows and columns.

As an example, consider

$$\mathbf{\Sigma}_0 = \begin{pmatrix} \sigma_1 & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2 & 0 & \cdots & 0 \\ 0 & 0 & \sigma_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \sigma_L \\ 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.45)$$

where $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_L$.

The matrices $\mathbf{U}_0$ and $\mathbf{V}_0$ contain the left and right singular vectors of $\mathbf{X}_0$. The pseudo-inverse of $\mathbf{X}_0$ is computed from the decomposed matrices as

$$\mathbf{X}_0^\dagger = \mathbf{V}_0 \mathbf{\Sigma}_0^{-1} \mathbf{U}_0^{T*} \quad (2.46)$$

Here it must be noted that Σ_0^{-1} is not a true inverse, except in the case of a square matrix $\mathbf{X}_0$. It is defined, using the scalar inverses of the nonzero elements of Σ_0^{-1} , as

$$\Sigma_0^{-1} = \begin{pmatrix} \sigma_1^{-1} & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2^{-1} & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & \sigma_3^{-1} & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & \sigma_L^{-1} & 0 & \cdots & 0 \end{pmatrix} \quad (2.47)$$

Thus, in the noise-free case, the z-domain frequencies z_n can be found and the Laplace-domain frequencies s_n calculated from them.

Instead of perfectly-known data $x(t_i)$, where $i = 1, \dots, K$, available data $y(t_i)$, $i = 1, \dots, K$, will in fact be corrupted with noise. The noisy data is arranged into Hankel data matrices $\mathbf{Y}_0$ and $\mathbf{Y}_1$ as was done for uncorrupted data in Equations (2.34) and (2.35). If the presence of additive noise is assumed in each element of the matrices $\mathbf{Y}_0$ and $\mathbf{Y}_1$, the noise contributions $\mathbf{W}_0$ and $\mathbf{W}_1$ can be separated, in theory, from the uncorrupted matrices $\mathbf{X}_0$ and $\mathbf{X}_1$, and the matrix pencil can be rewritten as

$$\begin{aligned} \mathbf{Y}_1 - \lambda \mathbf{Y}_0 &= (\mathbf{X}_1 + \mathbf{W}_1) - \lambda(\mathbf{X}_0 + \mathbf{W}_0) \\ &= (\mathbf{X}_1 - \lambda \mathbf{X}_0) + (\mathbf{W}_1 - \lambda \mathbf{W}_0) \end{aligned} \quad (2.48)$$

If λ is a generalized eigenvalue of the matrix pencil $\mathbf{X}_1 - \lambda \mathbf{X}_0$, in other words, if $\lambda_n = z_n, n = 1, \dots, N$, the rank of the matrix pencil is reduced by one. However, because of the presence of noise matrices $\mathbf{W}_0$ and $\mathbf{W}_1$ in Equation (2.48), only noisy estimates $\hat{z}_n$ of these generalized eigenvalues are available.

Once the eigenvalue estimates $\hat{z}_n$ are known, the natural frequency estimates may be found via the natural logarithm. Although not pointed out explicitly in the literature [48–51, 54, 119, 120], this must be a complex natural logarithm [64, 125],

$y = Ln(x)$. Here, it is defined with a complex range with real components $\Re(y) = \ln|x|$ in $(-\infty, \infty)$ and imaginary values $\Im(y) = \arg x$ in $(-\pi, \pi]$. The domain of the complex natural logarithm used here is a single Riemann sheet with a branch cut along the $-\Re(x)$ axis. The natural frequencies are estimated via

$$\hat{s}_n = Ln(\hat{z}_n)/(\Delta t) \quad (2.49)$$

To reduce noise, one data matrix is filtered prior to estimating eigenvalues. The truncated pseudo-inverse of the Hankel matrix $\mathbf{Y}_0$ [48] is found by modifying the SVD computation in Equation (2.46). The truncated pseudo-inverse is the pseudo-inverse of the low-rank approximation [121] described below. The data matrix is assumed to be of rank N , given the N terms summed in the underlying signal model. The largest N singular values, called the dominant or principal singular values [47], are assumed to represent the underlying signal $x(t_i)$ and are retained. The singular vectors associated with those singular values are assumed to span the signal subspace of the matrix [121]. The remaining $L - N$ singular values of the matrix are associated with noise, rather than signal, thus are set to zero. The $L - N$ singular vectors associated with the non-principal singular values span a subspace orthogonal to the signal subspace, which is assumed to be the noise subspace of the data matrix $\mathbf{Y}_0$ [121]. The low-rank approximation to the pseudo-inverse of $\mathbf{Y}_0$, also called the truncated pseudo-inverse, is given by:

$$\mathbf{Y}_{0T}^\dagger = \mathbf{V}_0 \boldsymbol{\Sigma}_{0T}^{-1} \mathbf{U}_0^{T*} \quad (2.50)$$

where $\boldsymbol{\Sigma}_{0T}^{-1}$ is the reduced-rank version of $\boldsymbol{\Sigma}_0^{-1}$ from Equation (2.47), in which σ_{N+1}^{-1} through σ_L^{-1} have been replaced with zeros. To employ the truncated pseudo-inverse effectively, some means must be found to determine how many singular values to retain. MPM in its original form provides no solution to this problem, but the Total Least-Squares Matrix Pencil Method does, as discussed in the next section.

2.3.2 The Total Least-Squares Matrix Pencil Method. An improvement to the Matrix Pencil Method is the Total Least-Squares (TLS) Matrix Pencil Method [47, 120]. In the basic version of MPM, one data matrix, $\mathbf{Y}_0$, is perturbed by SVD truncation. In TLS-MPM, both matrices, $\mathbf{Y}_0$ and $\mathbf{Y}_1$, are perturbed. To illustrate, let error matrices $\mathbf{E}_0$ and $\mathbf{E}_1$ be the differences $[\mathbf{Y}_0 - \mathbf{Y}_{0R}]$ and $[\mathbf{Y}_1 - \mathbf{Y}_{1R}]$. The ranks of $\mathbf{Y}_{0R}$ and $\mathbf{Y}_{1R}$ are both equal to N and the Frobenius norms [121] of the perturbation matrices $\|\mathbf{E}_0\|_F$, $\|\mathbf{E}_1\|_F$ are minima [47], as guaranteed by the SVD [121]. The Frobenius norm of the error matrix $\mathbf{E}_0$ is [121]:

$$\|\mathbf{E}_0\|_F = \text{trace}\{\mathbf{E}_0^{T*}\mathbf{E}_0\} = \text{trace}\{[\mathbf{Y}_0 - \mathbf{Y}_{0R}]^{T*}[\mathbf{Y}_0 - \mathbf{Y}_{0R}]\} \quad (2.51)$$

The data matrices $\mathbf{Y}_0$ and $\mathbf{Y}_1$ are not calculated directly. A single Hankel data matrix $\mathbf{Y}$ is computed from all data points. Its approximate matrix is found via SVD truncation. The right singular matrix $\mathbf{V}$ contains N principal and $L - N$ non-principal right singular vectors of $\mathbf{Y}$. The principal vectors form a new truncated matrix $\mathbf{V}_T$, of L rows and N columns. The vectors making up the columns of $\mathbf{V}_T$ are orthogonal to the remaining, non-principal vectors of $\mathbf{V}$. The first $L - 1$ rows of $\mathbf{V}_T$ form $\mathbf{V}_{0T}$, while the last $L - 1$ rows form $\mathbf{V}_{1T}$.

The z -domain frequencies z_n are then estimated from eigenvalues λ of the matrix pencil

$$\mathbf{V}_{1T} - \lambda \mathbf{V}_{0T} \quad (2.52)$$

Because $\mathbf{V}_{0T}$ and $\mathbf{V}_{1T}$ are orthogonal to the assumed noise subspace, TLS-MPM gains additional noise resistance over MPM.

As with the standard MPM, the estimated natural resonance frequencies $\hat{s}_n$ are found from the eigenvalue estimates $\hat{z}_n$ via the same complex natural logarithm defined previously as $Ln(z)$. $Ln(\hat{z}_n)$ is divided by the sample period Δt to find $\hat{s}_n$ as in Equation (2.49).

An additional benefit of TLS-MPM is a criterion to estimate the desired number of poles N based on knowledge of the accuracy of the data. While the basic MPM uses an arbitrarily assumed number of exponential terms, TLS-MPM determines a number of terms based on the size of the singular values of the Hankel matrix and the known or assumed level of noise in the data. Given data that is accurate to P digits, the singular values associated with the signal are found via

$$\sigma_n \geq \sigma_1 \times 10^{-P} \quad (2.53)$$

Singular values less than the criterion $\sigma_1 \times 10^{-P}$ are assumed to represent noise and are set to zero. The distance between a rank- $[N+1]$ matrix and its rank- N truncation is given by the singular value σ_{N+1} . If σ_{N+1} is less than the noise assumed to be present in the signal, the noise is sufficient to account for the difference between the rank- N truncated matrix and the rank- $[N+1]$ matrix. Therefore, the singular value σ_{N+1} represents noise rather than signal [97].

The significance level P is described as the number of noise-free digits, but in fact, need not be an integer. To represent noise variations in increments of less than 20 dB, fractional values of P are required. These values can be interpreted as narrow bounds of uncertainty on the final digit in the data. If the last digit of data is given as 4 but the true value is known to be between 3 and 5, it can be viewed as a fraction of a noise-free digit, because an amount of usable information is present that is less than that of a noise-free digit.

2.3.3 Modified Matrix Pencil Methods. A separate improvement to MPM is the addition of the low-rank Hankel approximation (LRHA) of Li, Liu, and Razavilar [66]. The LRHA is a pre-filtering technique applied to a data matrix before employing a linear predictor algorithm. Li, *et al* [66], proposed its use with the Kumaresan-Tufts (KT) algorithm [61] to create the Modified KT algorithm. This technique, under the name reduced-rank Hankel approximation (RRHA), was employed with

MPM by Lu, Wei, Evans, and Bovik [69]. Three variants were developed. The one discussed here was denoted Modified Matrix Pencil Method 3 (MMP3). The term LRHA will be used rather than RRHA from here onward for consistency with Li, *et al* [66]. LRHA has not been employed in conjunction with TLS-MPM, only with the basic MPM.

The LRHA is an iterative technique that alternates between reducing the rank of a matrix via SVD truncation and returning the matrix to Hankel symmetry by averaging across the elements of each anti-diagonal. Let $\mathbf{X}_0$ be the noise-free data matrix, with Hankel symmetry, used in MPM. Let $\mathbf{Y}_0$ be a noise-corrupted version of $\mathbf{X}_0$. Let $\mathbf{Y}_{0T}$ be an approximate version of $\mathbf{Y}_0$, filtered by SVD truncation as explained in Equation (2.50). A different approximation of $\mathbf{Y}_0$, denoted $\mathbf{Y}_{0H}$, is found by averaging the values across each anti-diagonal of $\mathbf{Y}_{0T}$. Li, *et al*, [66] show that, compared to $\mathbf{Y}_{0T}$, $\mathbf{Y}_{0H}$ is a superior approximation to $\mathbf{X}_0$ in that

$$\|\mathbf{Y}_{0H} - \mathbf{X}_0\|_F \leq \|\mathbf{Y}_{0T} - \mathbf{X}_0\|_F \quad (2.54)$$

where $\|\cdot\|_F$ indicates the Frobenius norm [121]. The proof given by Li, *et al*, is that, for the elements in any anti-diagonal n , that is, elements at row i and column j satisfying $i + j = n$, the errors between the noise-free data elements $x(n)$ and the approximate data elements $y_T(i, j)$ and $y_H(n)$, are related by

$$\begin{aligned} \frac{1}{N} \sum_{i+j=n, i \geq 0, j \leq n} |x(n) - y_H(i, j)|^2 &= |x(n) - y_H(n)|^2 \\ &+ \frac{1}{N} \sum_{i+j=n, i \geq 0, j \leq n} |y_H(n) - y_T(i, j)|^2 \\ &\geq |x(n) - y_H(n)|^2 \end{aligned} \quad (2.55)$$

where there are N elements in anti-diagonal n of matrix $\mathbf{X}_0$, $\mathbf{Y}_{0T}$, or $\mathbf{Y}_{0H}$.

Applying Equation (2.55) to every anti-diagonal in each matrix, it can be seen that

$$\|\mathbf{Y}_{0H} - \mathbf{X}_0\|_F^2 \leq \|\mathbf{Y}_{0T} - \mathbf{Y}_{0H}\|_F^2 + \|\mathbf{Y}_{0H} - \mathbf{X}_0\|_F^2 = \|\mathbf{Y}_{0T} - \mathbf{X}_0\|_F^2 \quad (2.56)$$

Thus the error in the reduced-rank Hankel approximate matrix is less than the error in the low-rank approximate matrix.

2.4 Stochastic Estimation

Here a variety of techniques are discussed which are fundamental to the target recognition algorithm. Each target in the library will be represented by a model in a Kalman filter. The basic theory of such filters is presented in Section 2.4.1 and the application of the filters to the target recognition problem is detailed in Section 4.2. Several filters will be combined in a multiple-model adaptive estimation (MMAE) [76] algorithm. MMAE is presented in Section 2.4.2 and developed into a target recognition algorithm in Section 4.3.

In stochastic control and estimation theory [76, 77], a frequently-seen problem is estimating the current values of the state variables of a given dynamic system. Often, the goal is to use the state variable estimates to control a system. In this case, however, there is no system to control. The state estimation process is simply used to gather information about a signal. An estimate of the states is formed from noise-corrupted measurement data and some knowledge of the system's dynamics.

The stochastic estimation techniques discussed in this research are based on noise of an assumed form. Noise in this case is assumed to be additive, white, Gaussian, and zero-mean. The additive assumption is simply that the noise adds or subtracts a random amount from the underlying signal. The white assumption is that the noise has infinite bandwidth. Realistically, noise cannot have infinite bandwidth because that implies infinite power involved in the noise process, which is impossible.

However, the noise can have nearly constant power per unit frequency (i.e., power spectral density value) over the entire bandwidth of the system, so that there is no distinguishable difference between realistic wideband noise and the idealized white noise. The assumption of a Gaussian probability distribution for the noise is justified by the Central Limit Theorem [44], which states that as more and more random variables are added together, the distribution of their sum becomes more nearly Gaussian, regardless of the probability distributions of the individual variables. The zero-mean assumption is that at any point in time, the sum of all possible values the noise could attain at that point, weighted by the probability of the noise achieving those values, is zero. In other words, the expected value of the noise at any point in time is zero, simply because the noise is equally likely to be above or below zero.

It is also assumed that the noise is an ergodic process, meaning a random process which has a time average that is statistically equivalent to the average over all possible realizations of the random variable corresponding to the random process at any one particular time. Since the noise is zero-mean, a time average over a sufficient length of time should tend toward zero.

2.4.1 The Kalman Filter. Given any system which is linear in its state variables, operates according to a known differential equation, and is subject to additive, white, Gaussian noise, the best estimate of the true states of the system is provided by a Kalman filter [65, 77]. Kalman filters are widely used in control applications in which measurements may be corrupted by noise, such as aircraft flight control [65, 78, 80] and navigation [94]. Nishiyama [96] applied a Kalman filter to a frequency estimation problem, but his signal model was different from the one presented here. The late-time scattered signal is a real one consisting of damped sinusoid terms, while Nishiyama assumed a complex signal with undamped sinusoids.

The Kalman filter generates an estimate of the states based on a probability-weighted average between the measurements and the expected behavior of the system. Weights are determined from the expected covariance of both measurements and states. In a discrete-time Kalman filter, every sample period, the most recently updated value of each state is propagated using the system dynamics equation and known control inputs. As measurements are taken, the updated states are determined using the measurements and the state values immediately prior to the measurements. Covariances are calculated for both sets of state estimates based on expected noise in the dynamic process and the measurement equipment.

The linear system model state and output equations for a continuous-time system with sampled data are [77, 106]

$$\frac{d\bar{x}(t)}{dt} = \mathbf{F}(t)\bar{x}(t) + \mathbf{B}(t)\bar{u}(t) + \mathbf{G}(t)\bar{w}(t) \quad (2.57)$$

$$\bar{y}(t_i) = \mathbf{H}(t_i)\bar{x}(t_i) + \mathbf{D}(t_i)\bar{u}(t_i) + \bar{v}(t_i) \quad (2.58)$$

where $\mathbf{F}$, $\mathbf{B}$, $\mathbf{H}$, $\mathbf{D}$, and $\mathbf{G}$ are matrices which in general are time-varying (although time-invariant matrices are used in the models for this research) and may be complex in a particular system. They represent system dynamics ($\mathbf{F}$), control inputs ($\mathbf{B}$), measurements or outputs ($\mathbf{H}$), direct coupling between inputs and outputs ($\mathbf{D}$), and noise inputs to the process ($\mathbf{G}$). The vectors $\bar{x}$, $\bar{y}$ and $\bar{u}$ are the state, measurement and control input vectors, respectively. The vectors $\bar{w}$ and $\bar{v}$ represent dynamics driving noise and measurement noise, respectively. These two noise sources are assumed independent of each other and Gaussian. The dynamics driving noise is assumed zero-mean, and has strength $\mathbf{Q}$. The measurement noise is also assumed zero-mean, and has covariance $\mathbf{R}$. Note that a state-space model is not unique. A given set of differential equations can be expressed in many different state space forms, giving us a great deal of flexibility. For a discrete-time (or sampled data)

system, $\mathbf{F}$ is not used directly. Instead, it is used to calculate the state transition matrix $\Phi(t_i, t_{i-1})$.

The linear system model in equivalent discrete-time form is [77]

$$\bar{x}(t_{i+1}) = \Phi(t_{i+1}, t_i)\bar{x}(t_i) + \mathbf{B}_d(t_i)\bar{u}(t_i) + \mathbf{G}_d(t_i)\bar{w}_d(t_i) \quad (2.59)$$

$$\bar{y}(t_i) = \mathbf{H}(t_i)\bar{x}(t_i) + \mathbf{D}(t_i)\bar{u}(t_i) + \bar{v}(t_i) \quad (2.60)$$

where the discrete-time dynamics driving noise $\bar{w}_d$ has covariance $\mathbf{Q}_d(t_i)$, defined below in Equation (2.65). The state transition matrix Φ is a solution of

$$\frac{d\Phi(t, t_0)}{dt} = \mathbf{F}(t)\Phi(t, t_0) \quad (2.61)$$

$$\Phi(t_0, t_0) = \mathbf{I}$$

in which $\mathbf{I}$ is a $p \times p$ identity matrix. If $\mathbf{F}$ is time-invariant, the equation can be solved in the Laplace domain, producing the resolvent matrix

$$\Phi^L(s) = [s\mathbf{I} - \mathbf{F}]^{-1} \quad (2.62)$$

The answer can be transformed back into the time domain, yielding

$$\Phi(t - t_0) = \mathcal{L}^{-1}\{[s\mathbf{I} - \mathbf{F}]^{-1}\}|_{(t-t_0)} \quad (2.63)$$

where $\mathcal{L}^{-1}$ indicates the inverse Laplace transform. [77, 106]

The filter state estimate $\hat{\hat{x}}(t_i)$ and state covariance $\mathbf{P}(t_i)$ are computed recursively from initial values $\hat{\hat{x}}(t_0) = \bar{x}_0$ and $\mathbf{P}(t_0) = \mathbf{P}_0$. From filter start, and after each measurement update, the filter propagates $\hat{\hat{x}}(t_{i-1})$ and $\mathbf{P}(t_{i-1})$ forward in time. The filter computes a predicted value for the state, $\hat{\hat{x}}(t_i^-)$, and covariance $\mathbf{P}(t_i^-)$ for a time t_i^- immediately prior to the next measurement update at t_i . The Kalman filter propagation equations corresponding to the discrete-time system are [77]

$$\begin{aligned}
\hat{\bar{x}}(t_i^-) &= \Phi(t_i, t_{i-1})\hat{\bar{x}}(t_{i-1}^+) + \int_{t_{i-1}}^{t_i} \Phi(t_i, \tau)\mathbf{B}(\tau)\bar{u}(\tau)d\tau \\
&= \Phi(t_i, t_{i-1})\hat{\bar{x}}(t_{i-1}^+) + \mathbf{B}_d(t_{i-1})\bar{u}(t_{i-1})
\end{aligned} \tag{2.64}$$

and

$$\begin{aligned}
\mathbf{P}(t_i^-) &= \Phi(t_i, t_{i-1})\mathbf{P}(t_{i-1}^+)\Phi^T(t_i, t_{i-1}) + \int_{t_{i-1}}^{t_i} \Phi(t_i, \tau)\mathbf{G}(\tau)\mathbf{Q}(\tau)\mathbf{G}^T(\tau)\Phi^T(t_i, \tau)d\tau \\
&= \Phi(t_i, t_{i-1})\mathbf{P}(t_{i-1}^+)\Phi^T(t_i, t_{i-1}) + \mathbf{G}_d(t_{i-1})\mathbf{Q}_d(t_{i-1})\mathbf{G}_d^T(t_{i-1})
\end{aligned} \tag{2.65}$$

The control input $\bar{u}$ is assumed to be held constant during a sample period. That allows $\mathbf{B}_d(t_{i-1})$ to be defined as

$$\mathbf{B}_d(t_{i-1}) = \int_{t_{i-1}}^{t_i} \Phi(t_i, \tau)\mathbf{B}(\tau)d\tau \tag{2.66}$$

The discrete-time system noise covariance is defined by

$$\mathbf{G}_d(t_{i-1})\mathbf{Q}_d(t_{i-1})\mathbf{G}_d^T(t_{i-1}) = \int_{t_{i-1}}^{t_i} \Phi(t_i, \tau)\mathbf{G}(\tau)\mathbf{Q}(\tau)\mathbf{G}^T(\tau)\Phi^T(t_i, \tau)d\tau \tag{2.67}$$

in which there is some flexibility in the choice of both $\mathbf{G}_d$ and $\mathbf{G}$. If the $\mathbf{Q}_d$ matrices have as many rows as the state vector, and are square, it is possible to choose $\mathbf{G}_d = \mathbf{I}$. For this choice of $\mathbf{G}_d$, the discrete-time dynamics noise covariance can be found via

$$\mathbf{Q}_d(t_{i-1}) = \int_{t_{i-1}}^{t_i} \Phi(t_i, \tau)\mathbf{G}(\tau)\mathbf{Q}(\tau)\mathbf{G}^T(\tau)\Phi^T(t_i, \tau)d\tau \tag{2.68}$$

For every sample time t_i , a new measurement vector $\bar{y}(t_i)$ is fed into the filter. The filter then computes a residual $\bar{r}(t_i)$ with filter-computed covariance $\mathbf{A}(t_i)$. The Kalman gain $\mathbf{K}(t_i)$ is computed and used to update both the state vector estimate $\hat{\bar{x}}(t_i^+)$ and its covariance $\mathbf{P}(t_i^+)$ to their values at a time t_i^+ immediately after the update at time t_i . The Kalman filter update equations for the system are [77]:

$$\bar{r}(t_i) = \bar{y}(t_i) - \mathbf{H}(t_i)\hat{\hat{x}}(t_i^-) \quad (2.69)$$

$$\mathbf{A}(t_i) = \mathbf{H}(t_i)\mathbf{P}(t_i^-)\mathbf{H}^T(t_i) + \mathbf{R}(t_i) \quad (2.70)$$

$$\mathbf{K}(t_i) = \mathbf{P}(t_i^-)\mathbf{H}^T(t_i)[\mathbf{A}(t_i)]^{-1} \quad (2.71)$$

$$\hat{\hat{x}}(t_i^+) = \hat{\hat{x}}(t_i^-) + \mathbf{K}(t_i)\bar{r}(t_i) \quad (2.72)$$

$$\mathbf{P}(t_i^+) = \mathbf{P}(t_i^-) - \mathbf{K}(t_i)\mathbf{H}(t_i)\mathbf{P}(t_i^-) \quad (2.73)$$

To improve numerical stability, Equation (2.73) can be replaced by the Joseph form [77]:

$$\mathbf{P}(t_i^+) = [\mathbf{I} - \mathbf{K}(t_i)\mathbf{H}(t_i)]\mathbf{P}(t_i^-)[\mathbf{I} - \mathbf{K}(t_i)\mathbf{H}(t_i)]^T + \mathbf{K}(t_i)\mathbf{R}(t_i)\mathbf{K}(t_i)^T \quad (2.74)$$

This form is equivalent, but avoids matrix subtraction, preserving both the positive definiteness and symmetry of the state covariance $\mathbf{P}(t_i^+)$. A different algorithm known as the U-D factored form provides numerical precision and stability superior to the Joseph form, but the Joseph form was easier to implement in the software package used in this research. The U-D factored form can be implemented in future research.

2.4.2 Multiple Model Adaptive Estimation. If the exact system dynamics are unknown, but can take any one of several known forms, the form which best describes the system can be found using Multiple Model Adaptive Estimation (MMAE) [71, 76]. MMAE algorithms have been applied to flight control systems which conduct failure detection and identification [89], target tracking algo-

rithms [31], control of spacecraft structures [39], and detection of jamming in navigation systems [140].

One widely used application of stochastic estimation and control techniques involves failure detection and identification (FDI) in a control system. FDI algorithms must first detect the occurrence of a failure, then attempt to diagnose it to allow corrective action. In target recognition, it can be assumed that a new target has been detected, leaving us with a problem parallel to the failure mode identification portion of FDI. As soon as data comes in, something must be identified, here a target signature rather than a failure mode. It is useful to review FDI techniques here. In addition to MMAE, there are the Chi-Square (χ^2) test [44] and the generalized likelihood ratio test (GLRT) [136]. The Chi-Square test is effective for detection, but weak in parameter identification. In typical implementations of GLRT [135,142,143], a single Kalman filter provides a residuals vector to a bank of matched filters. Each filter contains a model of the system affected by an anticipated error, and looks for the presence of an offset in the residuals which indicates the failure has occurred. In addition, a known GLRT limitation [135] is that windowing the estimate of the error onset time θ leads to a reduction in the accuracy of the estimate of the size of the failure, and may cause a delay in identifying failures. If a time window is not used, the set of hypotheses to be examined by a GLRT method will grow, leading to an unrealizable algorithm [135].

In MMAE, several Kalman filters run simultaneously using the same measured data. Each filter runs with a different system model representing the system as it may operate, or variations of the same model with different parameter values. Each filter (denoted with a subscript k) develops its own state estimate $\hat{\hat{x}}_k(t_i)$ and compares it to the measurement vector $\bar{y}(t_i)$, generating a residual $\bar{r}_k(t_i) = \bar{y}(t_i) - \mathbf{H}_k(t_i)\hat{\hat{x}}_k(t_i^-)$. Note that this form of the residual is sometimes called the “innovation” [56, 57]. Authors using the innovation term will define the residual as $\bar{r}_k(t_i) =$

$\bar{y}(t_i) - \mathbf{H}_k(t_i)\hat{x}_k(t_i^+)$ [77]. In the rest of this document the term residual will be applied to the quantity calculated immediately prior to measurement update.

The scaled residuals can then be calculated from the Cholesky square root of the inverse of the $\mathbf{A}_k(t_i)$ matrix, $\sqrt[{}^c]{\mathbf{A}_k^{-1}}^T \bar{r}_k$, to normalize each residual to its anticipated magnitude. Computationally, the scaled residuals are not explicitly calculated, but rather the quadratic form $\bar{r}_k^T \mathbf{A}_k^{-1} \bar{r}_k$ ($= \bar{r}_k^T \sqrt[{}^c]{\mathbf{A}_k^{-1}} \sqrt[{}^c]{\mathbf{A}_k^{-1}}^T \bar{r}_k$) is computed within MMAE. The matrix $\mathbf{A}_k(t_i) = \mathbf{H}_k(t_i)\mathbf{P}_k(t_i^-)\mathbf{H}_k^T(t_i) + \mathbf{R}_k(t_i)$ is the covariance computed by each filter for its own residuals.

The MMAE method involves monitoring the scaled residuals from each model and calculating the probability that any one model is the right one. This is called the hypothesis conditional probability (HCP) $p_k(t_i)$ because it is conditioned on the history of actual measurement values [76]. The component filter producing the smallest scaled residuals is the one which best represents the true system dynamics. MMAE systems can provide a weighted sum of the state estimates of each filter, in which the weights are the computed probabilities that each of the filters is correct. An alternative is to use the state estimate produced by the filter with the highest HCP. Using the single best estimate of the state instead of the weighted average results in the maximum *a posteriori* (MAP) form of the MMAE system.

The standard, Bayesian blended MMAE, and the MAP version of MMAE, are based upon very similar equations, but are very different conceptually. Both implement a bank of Kalman filters, and both compute an HCP for each filter based on its residual vector and residual covariance matrix. The difference is what is done with the information. The standard MMAE develops a single, blended state estimate based on a probability-weighted average of the states estimated by each filter. It is therefore a conditional mean estimator. The MMAE-MAP algorithm, in contrast, is a conditional mode estimator. Instead of forming a blended estimate from every filter's state vector, it uses only the state estimate from that filter which has the highest HCP at that time. The term conditional in the case of the two estimators

and the hypothesis conditional probability, refers to the fact that the probability assigned to each filter is conditioned on the history of all measurements seen until that point.

The hypothesis conditional probability is determined by the iteration

$$p_k(t_i) = \frac{f_{\bar{y}(t_i)|\bar{a},\bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_k, \bar{Y}_{i-1})p_k(t_{i-1})}{\sum_{j=1}^K f_{\bar{y}(t_i)|\bar{a},\bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_j, \bar{Y}_{i-1})p_j(t_{i-1})} \quad (2.75)$$

where the conditional probability density function is calculated from [136]

$$f_{\bar{y}(t_i)|\bar{a},\bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_k, \bar{Y}_{i-1}) = \beta_k \exp\left\{-\frac{1}{2}\bar{r}_k^T(t_i)\mathbf{A}_k^{-1}(t_i)\bar{r}_k(t_i)\right\} \quad (2.76)$$

where

$$\beta_k = \frac{1}{(2\pi)^{m/2}|\mathbf{A}_k(t_i)|^{1/2}} \quad (2.77)$$

and m is the number of elements in the measurement vector. The residual covariance matrix $\mathbf{A}_k$ is generated in the k th Kalman filter as:

$$\mathbf{A}_k(t_i) = \mathbf{H}_k(t_i)\mathbf{P}_k(t_i^-)\mathbf{H}_k^T(t_i) + \mathbf{R}_k(t_i) \quad (2.78)$$

Because the HCPs are calculated recursively, they must be initialized at a certain value. For this application, prior to employing the algorithm, the models in each of the filters are assumed equally likely to describe the system correctly. Therefore, for a bank of K filters, the probability of each filter is initialized to $1/K$.

Note that each filter's HCP is weighted by the HCP computed for that filter from the previous sample, as shown in Equation (2.75). In effect, probability flows from each filter that poorly matches the measurement to each filter that matches the measurement well. A problem in this formulation is that a filter's HCP may reach zero and be unable to regenerate. In control and estimation problems, the physical problem may change over time, so that a filter which was not needed early

on may become important later. If its probability falls to zero however, it cannot contribute to the overall MMAE system after that point. One solution is to bound the probability to some small finite value, such as 0.01 or less, to allow probability to flow back to any filter. To keep the mathematical definition of probability intact, an upper bound could be imposed on the HCP of each filter which is the sum of the lower bounds on all other filters. Another simple solution is to sum the bounded probabilities and divide each by the sum, in effect rescaling them to ensure they sum to one. The probabilities set to the lower bound would end up only slightly lower than the arbitrary lower bound, which is no problem at all. An alternative to lower bounds is to implement a transition probability model using the model that the random process resulting in the probability vector for all time, composed of components $p_k(t_i)$ for $k = 1, 2, \dots, K$, is a Markov process, that is, the present value of the probability vector depends only on the previous single value and not on the entire time history [76].

The weighted average state estimate is

$$\hat{x}_M(t_i^+) = \sum_{k=1}^K \hat{x}_k(t_i^+) p_k(t_i) \quad (2.79)$$

which has covariance

$$\mathbf{P}_M(t_i^+) = \sum_{k=1}^K p_k(t_i) \{ \mathbf{P}_k(t_i^+) + [\hat{x}_k(t_i^+) - \hat{x}_M(t_i^+)] [\hat{x}_k(t_i^+) - \hat{x}_M(t_i^+)]^T \} \quad (2.80)$$

where $\hat{x}_k$ is the state estimate produced by the k th filter given the assumption that parameter vector $\bar{a}_k$ is correct. In a control system which employs MMAE, the weighted average estimate $\hat{x}_M$ is critically important. It is the answer MMAE is designed to find. In the target recognition scheme, however, neither $\hat{x}_M$ nor its conditional covariance $\mathbf{P}_M$ are needed. The key in this research is the HCP p_k associated with each filter k . Equations (2.75) and (2.76) will cause the probability p_k associated with the best-matching filter to rise, and all others to drop. After




Figure 2.6 MMAE Diagram

several iterations, it is desired that the correct probability will approach one, so a target identification answer can be reported along with the associated probability that the choice is correct.

Figure 2.6 depicts a schematic diagram for an MMAE system. The system in question is the more typical Bayesian blended version of MMAE. A diagram of an MAP version of MMAE would have all the features seen in Figure 2.6 except those inside the dotted box.

2.5 Previous Target Recognition Techniques and the E-Pulse

In addition to resonance annihilation techniques, Mooney, Ding and Riggs [90,91] implemented a generalized likelihood ratio test for target identification using late-time transients. The technique has some theoretical similarity to the Multiple Model Adaptive Estimation approach, as shown by Hanlon [41]. The implementation of Mooney, *et al* did not use Kalman filters. Instead, it involved a likelihood

function that incorporated the signal data over the entire time window, as well as all unknown parameters. Other methods which have been used to distinguish cones from cylinders are an autoregressive filter with exogenous input (ARX), optimal instrumental variables (IV) and autoregressive moving average with exogenous input (ARMAX) [24, 25].

Given a library of natural frequency signatures, the state of the art in recognizing a target involves resonance annihilation via pulse convolution [92, 118]. Several authors have done work on resonance annihilation techniques called the kill pulse (K-pulse) [58], extinction pulse (E-pulse), [18] single-mode pulse (S-pulse) [52], ξ -pulse [14], and resonance annihilation filter [117]. The idea is to determine a pulse which will eliminate the modes associated with a known set of natural frequencies when convolved with a signal. Several pulses can be convolved with copies of the input signal in a parallel scheme to produce several resultant signals.

Rothwell developed the E-Pulse [112] as a generalized K-Pulse [58], and used it both to recognize targets and also to estimate the natural frequencies of a target [112]. Rothwell also developed the single-mode pulse as an E-pulse variant [112]. The various resonance annihilation techniques, such as E-pulse, K-Pulse, ξ -Pulse, and S-Pulse, were developed to identify a target based on known resonances [14, 18, 52, 58, 92, 118]. Other researchers working with resonance annihilation techniques include Ruiz, Gallego, and Carrion [36, 117, 118], Ilavarasan, Ross, Rothwell and Chen [52], and Mooney, Ding and Riggs [92].

The E-pulse is described here since it is compared to the MMAE-MAP method in Chapter IV. E-pulses exist in two forms, forced and natural, both of which are weighted series of subsectional basis functions. The basis functions $\beta(t)$ can be Dirac delta functions, rectangular pulses, trapezoids, polynomials defined on short intervals, etc. Subsectional (finite-duration) basis functions are used because the E-pulse must have a finite duration. The forced E-pulse designed to eliminate the effects of N poles will have $N+1$ basis functions, where the first function is considered

the forcing function for the resonance which is eliminated by the remaining pulses. The natural E-pulse consists of N basis functions, which are considered to eliminate their own resonances [112].

An E-pulse $e(t)$ is made up of forcing $e^f(t)$ and extinction $e^e(t)$ components:

$$e(t) = e^f(t) + e^e(t) = e^f(t) + \sum_{m=1}^M \alpha_m \beta_m(t) \quad (2.81)$$

The extinction component is the sum of basis functions weighted by coefficients α_m . The forcing function may be any finite-duration function, and is commonly chosen to be a weighted basis function for simplicity. In the case of the natural E-pulse, the forcing component $e^f(t)$ is set to zero. The E-pulse is convolved with the late-time portion of the measured signal.

To show the theory behind the E-pulse, a signal uncorrupted by noise is considered, although in practice, all signals are corrupted by noise to some extent. In the noise-free case, a portion of the resulting signal $c(t) = e(t) \odot y(t)$ will have zero amplitude. The symbol $\odot$ indicates convolution. To illustrate the concept, consider an E-pulse $e_i^e(t)$ composed of a single basis function $\beta_i(t)$. When convolved with a signal, it is intended to eliminate the effect of a single pole s_i .

$$e_i^e(t) = \delta(t) - \exp(s_i T) \delta(t - T) \quad (2.82)$$

In the Laplace domain, the pulse is

$$e_i^{eL}(s) = 1 - \exp[(s_i - s)T] \quad (2.83)$$

so that at the frequency of interest,

$$e_i^{eL}(s_i) = 1 - \exp(0) = 0 \quad (2.84)$$

When an expression $e_i^e(t) \odot y(t)$ involving convolution is transformed into the Laplace domain, it becomes a multiplication expression $e_i^{eL}(s)y^L(s)$. The act of

convolving $e_i^e(t)$ with a signal $y(t)$ in the time domain has an effect analogous to multiplying $e_i^{eL}(s)$ by $y^L(s)$ in the frequency domain: the s_i frequency component is annihilated.

To annihilate many frequencies, the signal $y(t)$ may be convolved successively with the E-pulse $e_i^e(t)$ of each of them. Alternatively, in the Laplace domain, the signal $y^L(s)$ may be multiplied by the E-pulse $e_i^{eL}(s)$ of each frequency. An equivalent technique is to produce a single E-pulse composed of several weighted basis functions, each of which eliminates the contribution of a single pole term. The weights α_m applied to each basis function β_m can be found via a matrix equation. For an E-Pulse developed in the Laplace domain [113], the equation is:

$$\begin{pmatrix} \beta_1^L(s_1) & \beta_2^L(s_1) & \cdots & \beta_N^L(s_1) \\ \beta_1^L(s_2) & \beta_2^L(s_2) & \cdots & \beta_N^L(s_2) \\ \vdots & \vdots & \ddots & \vdots \\ \beta_1^L(s_N) & \beta_2^L(s_N) & \cdots & \beta_N^L(s_N) \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{pmatrix} = \begin{pmatrix} -e^{fL}(s_1) \\ -e^{fL}(s_2) \\ \vdots \\ -e^{fL}(s_N) \end{pmatrix} \quad (2.85)$$

Equation (2.85) can be expressed more simply as:

$$\mathbf{B}^L \bar{\alpha} = -\bar{e}^{fL}(\bar{s}) \quad (2.86)$$

where $\mathbf{B}^L$ is the matrix of basis functions β^L .

The effectiveness of the pulse is very sensitive to the basis length $T = p\Delta t$, $p = 1, 2, \dots$. For continuous data, T could be any time, but given sampled data, an integer multiple of the sample period is the easiest to implement. Here, T is any positive integer p times the sample period Δt . The matrix $\mathbf{B}^L$ becomes singular when $T = p\pi/\omega_n$. Singularity is a useful characteristic for the natural E-pulse, but must be avoided for the forced E-pulse. In the case of the natural E-pulse, Equation (2.85) becomes

$$\mathbf{B}^L \bar{\alpha} = \bar{0} \quad (2.87)$$

For this equation to have a nonzero solution, the matrix $\mathbf{B}^L$ must be singular, i.e., must have two rows that are linearly dependent on each other. A way to guarantee linear dependence is to set the pulse length to be such that $T = p\pi/\omega_n$. The length must be chosen so that only those two rows are dependent.

Given a library of K targets, each with its own E-pulse, a decision must be made about which target best matches the signal. Several E-pulses $e_k(t)$, $k = 1, 2, \dots, K$ can be convolved with the measured signal $y(t)$ to produce several convolved signals $c_k(t)$, $k = 1, 2, \dots, K$. The convolved signal indicating the best target match is the one with the lowest energy discrimination number (EDN) [18, 92]:

$$EDN = \left[\int_{T_L}^{T_L+W} c^2(t) dt \right] \left[\int_0^{T_e} [e^f(t)]^2 dt \right]^{-1} \quad (2.88)$$

The EDN is found for each resultant signal c . T_e is the duration of the E-pulse, and T_L is the beginning of the late-time period plus T_e . W is the window width, chosen to be the same for all convolutions. The convolved signal with the lowest EDN is selected as the best match, indicating which target signature best fits the data [118].

2.6 Chapter Summary

In this chapter, background information has been provided on the electromagnetic theory which pointed the way to the new target recognition algorithm. The theory behind the pole estimation algorithm used to build the signature library has been presented. In addition, the procedures necessary to implement the Kalman filter and the multiple model adaptive estimation algorithm have been shown. In the next chapter, a novel variant of the Matrix Pencil Method, the Modified Total Least-Squares Matrix Pencil Method, is developed and applied to the problem of SEM pole estimation for late-time scattered signals.

III. Building the Library: Signal Processing Research

3.1 Overview of Target Signature Development

The target recognition scheme developed in this research involves comparing the target signature to a known library of signatures. To implement it, the library must first be filled. The signatures in the library are simply sets of resonant frequencies associated with each target. It is therefore necessary to estimate an adequate set of natural frequency estimates via late-time data from a target. An adequate set is a number of estimated poles sufficient to reconstruct late time signals from any azimuth. The targets used in this research are simple geometric shapes such as cylinders and blocks, or combinations thereof. Libraries contain signatures for four to ten targets, in the form of Kalman filters, as explained in Chapter IV.

In this chapter, variants of the MPM are applied to the problem of estimating poles from data. A novel algorithm, called the Modified Total-Least-Squares Matrix Pencil Method (M-TLS-MPM) [49, 51, 120], is developed and evaluated. M-TLS-MPM combines elements of both the Total Least-Squares Matrix Pencil Method (TLS-MPM) and the Modified Matrix Pencil Method 3 (MMP3) [69]. These three algorithms are applied to the problem of estimating poles from a noisy, complex signal in Section 3.2.1 and the results are compared.

M-TLS-MPM is used to estimate the poles of measured and computed time-domain, late-time scattering data. The algorithm is able to estimate poles from data with 9 dB of signal-to-noise ratio (SNR) and still achieve results close to the Cramer-Rao bound. The definition of SNR is presented in Section 3.3.1.

From SEM theory, it is expected that the late-time scattered signal contains poles that are aspect-independent. That is, under ideal conditions, the same natural resonance frequencies should be found in target measurements taken from different directions. However, while a pole pair may be present, the coefficients associated with the poles may be so small in magnitude that the pole pair's effect is insignificant.

Since there is no guarantee that any pole or pole pair will be significant, there is no guarantee that it will be detected in a given measurement. In this research, a subset of the poles was found to be significant over a wide range of azimuths, allowing the development of signatures that represent a target across a span of azimuths.

The techniques applied here estimate the poles from the available data. That is very different from theoretically developing the poles, or from analyzing the impedance matrix of a computational electromagnetics code [68, 137]. Estimating poles via any variant of MPM [49, 119], or its alternatives, provides a model of the scattered signal as presented, with all the noise and other distortions present in a realistic signal. The distortions, such as Gibbs phenomenon [62], pedestal interaction, etc., can give rise to apparent poles which are not truly in the underlying signal. Only a small number of scattering modes will be excited sufficiently that their poles can be estimated at all. The remaining poles are present in theory, but in practice are lost in noise. Given measured scattering data, the only course of action is to estimate poles from that data. The best available pole estimator is sought. The estimators implemented here use data from a single measurement or computation, at a single aspect angle. There are then several pole sets for a given target. These pole sets are combined into a single signature which is intended to represent the target from a variety of angles.

The overall architecture of the target recognition system involves building a library offline (prior to system operation), and reading it online (immediately after signal reception). Therefore, computational efficiency in the library building phase is of much less importance than accuracy of estimation. It is acceptable for a given technique to take hours to generate a best answer for a given target, as long as the resulting answer can be read in a fraction of a second. Thus, it is of no concern that the iterative M-TLS-MPM algorithm may take a few minutes per data set, as long as it provides the best available estimate of the poles.

3.2 The Modified Total Least-Squares Matrix Pencil Method

In this section a new variant of MPM is proposed. Elements of the Modified Matrix Pencil Method 3 (MMP3) and the Total Least-Squares Matrix Pencil Method (TLS-MPM) are combined into a new algorithm named the Modified Total Least-Squares Matrix Pencil Method (M-TLS-MPM). In addition, the performance of the new algorithm is compared to that of both TLS-MPM and MMP3. Both M-TLS-MPM and MMP3 make use of an iterative Low-Rank Hankel Approximation (LRHA). The number of iterations of the LRHA required, relative to noise strength, is investigated.

MMP3 uses most, but not all, of the features of the TLS-MPM. It is in fact a total least-squares technique because both matrices $(\mathbf{Y}_0, \mathbf{Y}_1)$ are perturbed by SVD truncation prior to forming the matrix pencil $(\mathbf{Y}_{0T}^\dagger \mathbf{Y}_{1T})$, as in TLS-MPM [47, 51]. This aspect of MMP3 was not noted by its developers [69, 109]. There are two TLS-MPM techniques not employed in MMP3. First, partitions $\mathbf{V}_{0T}$ and $\mathbf{V}_{1T}$ of the truncated, right-singular matrix $\mathbf{V}_T$ are used in the matrix pencil, as shown in Equation (2.52), instead of data matrices $\mathbf{Y}_{0T}$ and $\mathbf{Y}_{1T}$. Second, the significance criterion P from Equation (2.53) is used to determine the number of poles.

Here is a summary of the M-TLS-MPM algorithm. First the data is formed into a single Hankel matrix. Then SVD truncation and Hankel approximation are performed in an iterative fashion. Four iterations are recommended, as explained in Section 3.3.2. The matrix pencil is formed from partitions of right singular matrix of the Hankel matrix. The eigenvalues of the matrix pencil $(\mathbf{V}_{0T}^\dagger \mathbf{V}_{1T})$ are found and converted to natural frequencies. Frequencies that are not physically realizable or otherwise invalid are eliminated.

3.2.1 Evaluation of M-TLS-MPM with Kumaresan-Tufts Data. The algorithm was tested with the damped-exponential data used by Kumaresan and Tufts [61] in 1982 to test their linear predictor technique. The same data set has been

used by other authors [66, 69] to compare their techniques to the Kumaresan-Tufts Method. Since the data is complex, it does not represent a scattered time-domain radar signal. It was used only to make a consistent comparison between the new M-TLS-MPM algorithm and two existing variants of MPM. This complex data set is designated ‘Kumaresan-Tufts Data’ or ‘KT data’ in this document. KT data is determined by:

$$y(t_i) = \sum_{n=1}^2 C_n \exp[(\Omega_n + \omega_n)t_i] + r(t_i) \quad (3.1)$$

where each term in the series is a single, unpaired complex pole, rather than a pole pair or real-axis pole. There are two poles, with damping rates $\Omega_1 = -0.2$ and $\Omega_2 = -0.1$, and oscillating rates $\omega_1 = 2\pi(0.42)$ and $\omega_2 = 2\pi(0.52)$. The sample time $\Delta t = t_{i+1} - t_i$ is one second. The coefficients are set to $C_1 = 1$ and $C_2 = 1$. The noise $r(t_i)$ is Gaussian and complex. The duration of the data set is 25 seconds.

With the sample time given, the data is under-sampled, so MPM and its variants are unable to estimate the frequencies accurately. The signal is under-sampled because the sampling rate of one sample per second was less than the Nyquist rate [62], which in this case is ω_2/π , or 1.04 samples per second. The z -Domain frequencies may be estimated with no difficulty from the eigenvalues of the matrix pencil, but mapping them to Laplace-domain frequencies exposes the problem that $\omega_2\Delta t > \pi$. Thus, the estimate for the pole at $-0.1 + j2\pi(0.52)$ will occur around $-0.1 - j2\pi(0.48)$ in the s plane. One solution to this problem lies in redefining the complex natural logarithm (CNL) so that the higher frequency may be captured. The CNL, $y = Ln(x)$, was defined in Section 2.3.1. It is suitable for conjugate pairs of frequencies for which $|\omega|\Delta t < \pi$. This is the CNL definition which is most appropriate for any real data, such as measured or simulated time domain radar data.

The alternative CNL, suited for KT data, is denoted $y = Ln'(x)$ in this document. Its complex range has a real component $\Re(y) = \ln|x|$ in $(-\infty, \infty)$ and an

imaginary component $\Im(y) = \arg x$ in $[0, 2\pi)$. The domain of the complex natural logarithm $y = Ln'(x)$ is a single Riemann sheet [64] with a branch cut along the $+\Re(x)$ axis. Because KT data consists of two poles in the upper half-plane, this definition of the CNL can capture both of them. Lower-half-plane poles $\omega < 0$ cannot be estimated using this definition.

The other solution to the undersampled data was to resample it at a higher rate. The KT data in its original, undersampled form was used to allow a direct comparison to the results already seen in the literature. In addition, Equation (3.1) was used to generate a new data set, over a duration of 25 seconds with a sample rate of 0.5 seconds. The new data is denoted 'resampled KT data' in this document. The analysis in this section was done for both the original and the resampled KT data.

The performance of M-TLS-MPM was compared to both MMP3 and TLS-MPM using the original KT data. Each technique was evaluated by comparing the mean squared error (MSE) of the parameter estimates to the Cramer-Rao bound (CRB) [44, 61] for those parameters. The performance of the three algorithms is displayed in terms of mean squared error versus noise power, both reported in decibels. The noise power (in dB referenced to a watt) is thus

$$S_{noise} = 10 \log(2\sigma^2) \quad (3.2)$$

where σ^2 is the noise variance. Since the noise is complex, the variance applies to both the real and imaginary components, and the total noise power is twice the variance. This is the same formula given by Kumaresan and Tufts [61] as a definition for the SNR. It is only a proper definition of SNR if the mean signal power is equal to one, which is not true for KT data. The mean signal power for signal $y(t_i)$, where $i = 1, \dots, I$, is

$$S = \frac{1}{I} \sum_{i=0}^{I-1} y(t_i)y^*(t_i) \quad (3.3)$$

The mean signal power of the KT data is 0.4321 watts rather than 1 watt. The SNR can be found by subtracting the noise power (in dB) from -3.6442 dB.

To assess the error in parameter estimation, 500 noise-corrupted signals were formed at each desired noise power by summing KT data with complex Gaussian noise. The desired noise powers were -25 dB to -1 dB, every 2 dB. The estimated poles were associated with the actual poles according to their imaginary components. The algorithm was set to look for two poles, so the analysis could focus on the error in estimation of the damping and oscillating rates of those two poles. Estimated poles for which $\hat{\omega} \leq 2\pi(0.47)$ were associated with $s_1 = -0.2 + j2\pi(.42)$, while poles for which $\hat{\omega} > 2\pi(0.47)$ were associated with $s_2 = -0.1 + j2\pi(.52)$. The MSE was computed as

$$MSE(\alpha) = \frac{1}{K} \sum_{k=1}^K (\hat{\alpha}_k - \alpha)^2 \quad (3.4)$$

where α is one of these parameters: $\Omega_1, \Omega_2, \omega_1, \omega_2$. The average was based on K estimated poles that were associated with each actual pole. Ideally, $K = 500$, but for low SNR cases, sometimes there were fewer than 500 estimated poles associated with an actual pole.

At higher noise powers, it was possible that the error on an ω estimate was large enough that the estimate was associated with the wrong pole. Such spurious associations may have increased the MSE of a parameter estimate.

Figures 3.1 through 3.4 show the mean-square error found for both M-TLS-MPM and MMP3, each with one iteration of LRHA, and TLS-MPM. The CRB used by Kumaresan and Tufts [61] is also shown.

The MSE of M-TLS-MPM was lower than that of MMP3 for estimates of ω , particularly in the range of -9 to -11 dB noise power, corresponding to 5.4 to 7.4 dB SNR. For estimation of Ω , both MMP3 and TLS-MPM beat M-TLS-MPM when




Figure 3.1 Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM




Figure 3.2 Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM




Figure 3.3 Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM




Figure 3.4 Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM

noise power was less than -15 dB, corresponding to a SNR of 11.4 dB. The advantage was approximately 1 dB of reduction in the MSE of $\hat{\Omega}$.

Oddly, the MSE of MMP3 and TLS-MPM estimates of Ω_1 and Ω_2 also beat the CRB by approximately 1 dB at these noise levels. To have an MSE smaller than the CRB is simply impossible. However, this result is consistent with the literature on MMP3 [66, 69, 109] and Kumaresan-Tufts [61] for damping parameter estimates. The result was consistent over many noise power levels, each of which involved 500 realizations of a noisy signal. It is likely that an error existed in the original derivation of the CRB [61], but recomputing the CRB is beyond the scope of this research.

For estimation of ω , the M-TLS-MPM result was approximately equal to the MMP3 and TLS-MPM results for noise power less than -15 dB (11.4 dB SNR) and significantly better below that. The MSE of the parameter estimates $\hat{\omega}_1$ and $\hat{\omega}_2$ was always greater than the CRB. The same MSE formulation was used for all parameters.

The CRB is a lower limit on variance, assuming the estimate of the parameter is unbiased, so that the expected value of the estimate is the true value of the parameter. For MSE to be equal to variance, the mean of the set of estimates must be equal to the true parameter, making the mean of the estimation error equal to zero. If the mean of the estimation error is not zero, the MSE will be larger than the variance by the square of the mean of the error.

Bias was observed in some estimates produced by all variants of MPM, which is why the MSE of the estimates was used instead of the variance. Typical bias ranged from ± 0.001 for all parameters at -25 dB noise power to ± 0.05 for Ω_1 and Ω_2 , and ± 0.5 for ω_1 and ω_2 at -1 dB noise power. The bias decreased slightly, as the number of iterations of LRHA increased, but that benefit became less significant because the variance was the larger part of the error. The bias of M-TLS-MPM was

only slightly larger than that of MMP3, on the order of 0.001 % at 5 dB SNR, and less than that of TLS-MPM for one or more iterations of LRHA.

Figures 3.5 through 3.8 show the effect of the first few iterations of LRHA on the mean-square error. Reduction in the MSE of the parameter estimate improves little after the second iteration of LRHA, but some benefit can be seen until four iterations for the Ω estimates at the higher noise powers. There is little or no benefit to increasing the iterations of the LRHA past four. The benefit of M-TLS-MPM depends on the parameter being estimated, so the conclusions reached here can be taken as indicators of performance trends rather than specifications to be achieved under all circumstances.

The analysis done with KT data was repeated with resampled KT data to verify the conclusions drawn with the original, undersampled KT data. Equation (3.1) was used with the same values for the damping rates, oscillating rates, and coefficients. The resampled KT data consisted of samples taken every 0.5 seconds, twice the original rate. The samples were taken over the same time span, so resampled KT data consisted of twice as many time samples as KT data. In this case, the CNL defined in Section 2.3.1, with the branch cut on the $-\Re(z)$ axis, was used. Since $\omega_2\Delta t < \pi$, the data was no longer undersampled, so the alternate CNL, $Ln'(z)$, was no longer needed. That allowed the algorithm to find poles with negative oscillating rates, although both poles in the resampled (and original) KT data have $\omega > 0$.

For the resampled data, the noise power was computed via Equation (3.2), the signal power was computed via Equation (3.3), and the MSE was computed via Equation (3.4). The signal power was 0.3891 watts, slightly smaller than in the original version of the data. At each desired noise power level, 500 noise-corrupted signals were formed. Poles were estimated via TLS-MPM, MMP3, and the new M-TLS-MPM algorithm. To associate estimates with poles, any estimate for which $\omega < 2\pi(0.47)$ was considered to be an estimate of the first pole, $s_1 = -0.2 + j2\pi(.42)$, while the remaining estimates were associated with the second pole, $s_2 = -0.1 + j2\pi(.52)$.




Figure 3.5 Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA




Figure 3.6 Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA




Figure 3.7 Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA




Figure 3.8 Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA

Figures 3.9 through 3.12 show the mean-square error found for both M-TLS-MPM and MMP3, each with one iteration of LRHA, and TLS-MPM. The CRB used by Kumaresan and Tufts [61] is also shown. A frequent problem noted in these figures is that the MSE was less than the CRB at high SNR, which is not possible. Most likely there was an error in the original CRB derivation. Since the mean signal power for the resampled KT data was 0.3891 watts, the SNR can be computed by subtracting the noise power from -4.0991 dB.

Figure 3.9 shows the MSE achieved for the estimate of Ω_1 . The MSE of the M-TLS-MPM estimate was clearly superior to that of TLS-MPM at noise powers greater than -13 dB. The improvement seen by M-TLS-MPM was 6 to 10 dB. M-TLS-MPM also beat MMP3 when noise power was above -5 dB (0.9 dB SNR).

Figure 3.10 shows the MSE achieved for the estimate of Ω_2 . Again, M-TLS-MPM proved significantly better than TLS-MPM for noise powers above -11 dB. The improvement in MSE was about 6 to 8 dB. MMP3 produced estimates with smaller MSE than M-TLS-MPM in the range of -9 to -5 dB noise power, but M-TLS-MPM and MMP3 produced similar results for higher noise powers.

Figure 3.11 shows the MSE achieved for the estimate of ω_1 . The performances of M-TLS-MPM and MMP3 were very similar. Both algorithms beat TLS-MPM by about 2 dB when noise power was above -13 dB (8.9 dB SNR).

Figure 3.12 shows the MSE achieved for the estimate of ω_2 . Again, the performances of M-TLS-MPM and MMP3 were very similar. Both algorithms beat TLS-MPM by about 2 dB when noise power was above -11 dB (6.9 dB SNR).

Figures 3.13 through 3.16 show the mean-square error found for M-TLS-MPM with one to five iterations of the LRHA. The CRB used by Kumaresan and Tufts [61] is also shown. The benefit of M-TLS-MPM over TLS-MPM is due to the iterative LRHA. Much if the benefit comes from the first iteration, but further iterations provide some smaller benefit as well. The second iteration helps with damping rate




Figure 3.9 Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data




Figure 3.10 Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data




Figure 3.11 Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data




Figure 3.12 Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from MMP3 and M-TLS-MPM (1 iteration of LRHA each) and TLS-MPM using resampled KT data




Figure 3.13 Cramer-Rao bound for $\hat{\Omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data




Figure 3.14 Cramer-Rao bound for $\hat{\Omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data




Figure 3.15 Cramer-Rao bound for $\hat{\omega}_1$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data




Figure 3.16 Cramer-Rao bound for $\hat{\omega}_2$ and Mean-Square Error from M-TLS-MPM with 1 to 5 iterations of LRHA using resampled KT data

at high noise power, but further iterations achieve little benefit. The fourth and further iterations appear to provide no significant improvement to the estimates of the oscillating rates.

Figure3.13 shows the MSE achieved for the estimate of Ω_1 by varying the number of iterations of the LRHA. At noise powers below -9 dB, increased iterations actually appeared to increase the MSE slightly, on the order of 0.1 dB for each iteration. Above -9 dB noise power, the second iteration provided additional reduction in MSE, on the order of 3 dB.

Figure3.14 shows the MSE achieved for the estimate of Ω_2 by varying the number of iterations of the LRHA. At noise powers below -11 dB, increased iterations again appeared to increase the MSE slightly, under 0.1 dB for each iteration. Above -11 dB noise power, the second iteration provided additional reduction in MSE, on the order of 2 dB. Some slight reduction in MSE was noted for the third iteration at noise powers from -7 dB to -3 dB.

Figure3.15 shows the MSE achieved for the estimate of ω_1 by varying the number of iterations of the LRHA. At noise powers below -11 dB, increased iterations actually appeared to increase the MSE slightly, on the order of 0.1 dB for each iteration. Above -9 dB noise power, the second iteration provided about 0.2 dB of additional reduction in MSE.

Figure3.16 shows the MSE achieved for the estimate of ω_2 by varying the number of iterations of the LRHA. At noise powers below -11 dB, increased iterations actually appeared to increase the MSE slightly, on the order of 0.1 dB for each iteration. Above -7 dB noise power, the second iteration provided about 1 dB of additional reduction in MSE. Some additional reduction in MSE, about 0.5 dB, was noted for the third iteration at noise powers from -7 dB to -3 dB.

To conclude this section, it was shown that the new M-TLS-MPM algorithm provides parameter estimates with smaller mean squared error than those of TLS-

MPM at low SNR, and smaller than MMP3 at some low SNR levels. The benefit is more pronounced for damping rates than for oscillating rates. M-TLS-MPM should be used in preference to TLS-MPM, unless processing time is a key criterion or if the SNR is high. M-TLS-MPM uses slightly less computation than MMP3 for an equal number of iterations, since M-TLS-MPM uses only a single SVD on the data matrix $\mathbf{Y}$ rather than separate SVD computations for $\mathbf{Y}_0$ and $\mathbf{Y}_1$. However, the bulk of the computational load is in the LRHA employed by both algorithms. If the SNR of the data is unknown, either the M-TLS-MPM or MMP3 algorithm should be applied. No more than four iterations of the LRHA should be used, and the bulk of the benefit of LRHA will be achieved after two iterations.

3.2.2 Screening Invalid Poles. A few of the poles output by MPM are invalid for various reasons. They may be nonphysical due to growing magnitude or (rarely) being complex poles with their conjugates missing, as explained below. They may be insignificant due to rapidly decreasing magnitude or association with a very small coefficient. Successful target recognition requires the best possible rendering of target poles that are both valid and significant. All invalid poles must be eliminated. The procedures for eliminating them are discussed below.

Poles with $\Omega > 0$ correspond to growing-magnitude exponentials in the time domain. It is not physically possible for an unforced stable system's response to grow in magnitude. Growth can only result from adding energy to a system, i.e. forcing the system. Poles with $\Omega > 0$ may be due to noise and an attempt to identify too many poles. Sarkar [119] discusses this problem and recommends that the poles with $\Omega > 0$ be eliminated. Such poles were encountered rarely in this research, and were eliminated.

Some of the insignificant or non-physical pole estimates occurred because of the SVD truncation used in the matrix pencil method. Some small singular values were set equal to zero prior to computing the eigenvalues of the matrix pencil. The zeroed

singular values resulted in many eigenvalues z_n that were very small in magnitude but not exactly equal to zero. These small, insignificant eigenvalues were easily discriminated from significant eigenvalues by their much larger-magnitude damping rates, $|\Omega_n|$.

Figure 3.17 shows all poles found by TLS-MPM prior to screening out invalid poles. The poles fell into two widely separated groups based on the size of Ω : a ‘significant’ group on the right side (closer to $\Omega = 0$) and an ‘insignificant’ group on the left, with much lower (larger magnitude) Ω . Thus the poles on the left side of the chart, declared invalid, had an Ω about 100 times as large as valid poles. If these poles represented actual damped sine components of the signal, they would damp out so quickly that they would contribute essentially nothing to the return.




Figure 3.17 Pole locations before removing high-damping-rate poles. Line near top indicates $\omega = \pi/\Delta t$

It was possible to separate these poles automatically via a simple criterion. If the damping rate was such that a damped sine term lost 90% of its magnitude (20

dB of magnitude) in a single sample period, it vanished too quickly to be significant. Thus, a damping limit Ω_{max} was defined as

$$\exp(\Omega_{max}\Delta t) = 0.1, \quad (3.5)$$

or, isolating Ω_{max} , as

$$\Omega_{max} = \ln(0.1)/(\Delta t). \quad (3.6)$$

The results in Figure 3.17 are typical of unscreened poles found via TLS-MPM. The poles fell into two areas based on the size of $|\Omega|$: the possibly significant region ($-300 < \Omega < 0$ Giganepers/s) and the clearly insignificant region ($\Omega < -5 \times 10^3$ Giganepers/s). The insignificant poles were separated from the potentially significant poles by a vast empty region in the s plane. They tended to have the $|\Omega|$ two orders of magnitude larger than the potentially significant $|\Omega|$ values. The sample period of this data was $\Delta t = 6.1043 \times 10^3$ Giganepers/s, so given a 90% magnitude-drop criterion, Ω_{max} was 377 Giganepers per second. This choice of Ω_{max} fell within the empty region in Figure 3.17, closer to the valid poles. Poles with $\Omega < \ln(0.1)/(\Delta t)$ were eliminated.

Another problem was unpaired complex poles. Such poles were expected in the case of complex KT data, but for real-valued data, valid complex poles must occur in conjugate pairs. Invalid, unpaired complex poles appeared rarely at the exact upper limit of frequency resolution given by the Nyquist criterion [62], at $\omega = \pi/\Delta t$. These poles were due to the branch cut in the complex natural logarithm function $Ln(z)$. All of the eigenvalues that gave rise to these poles were on the negative real axis of the z -plane. Some unpaired complex poles are shown in the upper left and right corners of Figure 3.17, indicated by a horizontal line at $+\pi/\Delta t = 628.3185$ GHz. The data used to generate Figure 3.17 was real-valued. Poles for which $\omega\Delta t = \pi$ were eliminated. As a practical matter, most unpaired complex poles also had large-magnitude damping rates as well.

3.3 Evaluation of M-TLS-MPM with Synthetic Data

The complex KT data was used in the previous section to test the M-TLS-MPM algorithm because that data set has been used by several developers [61, 66, 69] of linear predictor algorithms. However, it is of no further interest in this research. Since radar data is treated as real in the time domain, the algorithm must be tested with real data. Also, while KT data consisted of two poles, the late-time signal from a target frequently had many significant poles. Additional data sets were developed, which consisted of a greater number of poles either in conjugate pairs or on the real axis of the s -plane, resulting in real signals. The additional data is referred to herein as synthetic data. It is explained in greater detail in Section 3.3.1.

For the real data, the definition of the complex natural logarithm (CNL) $\hat{s} = Ln(\hat{z})$ introduced in Section 2.3.1 was again used. The alternate definition $\hat{s} = Ln'(\hat{z})$ employed with KT data was not suitable because negative frequencies were present, and it was not needed because $|\omega_n|\Delta t < \pi$ for all n . To review, the range of the CNL has a real component $\Re(\hat{s}) = \ln|\hat{z}|$ in $(-\infty, \infty)$ and an imaginary component $\Im(\hat{s}) = \arg x$ in $(-\pi, \pi]$. The domain of the complex natural logarithm used in this section is a single Riemann sheet with a branch cut along the $-\Re(\hat{z})$ axis.

3.3.1 Synthetic Data. Three general sources of data were considered. They were measured, simulated, and synthetic data. Both measured and simulated data were collected in the frequency domain as discussed in Appendices B and C. It is possible to collect either measured or simulated data directly in the time domain, but the RCS ranges available for research employed swept-frequency measurement. The computational electromagnetics codes available for use reported results in the frequency domain. Measured and simulated data were transformed into the time domain via the IFFT as described in Appendix D. Since only the late-time portion of the scattered signal can be described by the SEM pole series, it is important that the matrix pencil method only be applied to late-time data. The beginning of the




Figure 3.18 Synthetic data Version 1, no noise added, vs range in inches

data set must be after the start of late time, and the end of the data set should be before the decaying signal is lost in the noise.

The synthetic data did not represent actual targets; it was simply a set of damped sine series signals. However, it was designed to be closer in form to an actual radar signal than the KT data used in Section 3.2.1. A plot of an example synthetic signal is shown in Figure 3.18. It is the weighted sum of three damped sinusoids (corresponding to pole pairs in the s domain via the Laplace transform), one non-oscillating exponential ($\omega_1 = 0, A_1 = 1, B_1 = 0$) (corresponding to a real axis pole), and white noise:

$$y(t_i) = 2W \sum_{n=1}^4 \exp(\Omega_m t_i) [A_m \cos(\omega_m t_i) - B_m \sin(\omega_m t_i)] + r(t_i) \quad (3.7)$$

Equation (3.7) is designed to match the theoretical late-time radar signal shown in Equation (2.29). The natural frequencies, or poles, are defined by the damping rate Ω_m and the oscillating rate ω_m . The coefficients A_m and B_m are the residues

associated with each pole. Discrete samples occur at evenly spaced times t_i . A weighting term $W = 0.001$ is used to drive the overall magnitude of the synthetic signal closer to that of the measured radar signals.

Seven poles were used: three pole pairs and a single pole on the real axis. The poles were similar in the size of the real and imaginary components to the poles expected for targets of interest within the measured bandwidth. A noise term, $r(t_i)$, was added to investigate noise resistance of the pole estimation and target recognition algorithms. The noise was additive, white and Gaussian, and was generated by a pseudorandom number algorithm that allowed the same noise signal to be repeated exactly, if desired. The noise variance was adjusted to achieve a desired signal to noise ratio (SNR) in the tests.

SNR was defined as the mean square signal over the time window of interest, divided by the mean square noise over the same window, expressed in decibels (dB). For the synthetic data, the time window was the entire 301-sample data set, because the entire set matched the theoretical form of late-time scattering. However, the signal strength decayed during that window. An alternate formulation was considered, which could be called peak SNR, which is the peak power in the time window divided by the average noise power. It was determined that peak SNR was no more useful than average SNR, because it only provided information about the early part of the time window.

There was one assumption used in the synthetic data that may not be present in actual radar data. The coefficients A_m and B_m were chosen within an order of magnitude of each other to allow unbiased comparisons between their estimates. An actual target's scattered signal is due to many significant modes, each with varied coefficient magnitudes. All synthetic data coefficients had real and imaginary components with magnitudes between 0.3 and 1. Several variants of synthetic data were developed with different pole locations and different coefficients. The four versions were used to test the target recognition algorithm discussed in Chapter IV.

For each version, poles were estimated via M-TLS-MPM and the estimates used to develop Kalman filters. The poles used to develop each version of synthetic data are listed in Table 3.1.

m	Version 1	Version 2	Version 3	Version 4
1	-2	-5	-10	-4
2	$-3 \pm j30$	$-3 \pm j20$	$-5 \pm j50$	$-1 \pm j60$
3	$-4 \pm j60$	$-4 \pm j50$	$-2 \pm j90$	$-1 \pm j90$
4	$-5 \pm j110$	$-3 \pm j100$	$-1 \pm j130$	$-2 \pm j140$

Table 3.1 Pole locations for synthetic data, giganepers/sec and gigaradians/sec




Figure 3.19 Synthetic data poles, Versions 1 ($\cdot$), 2 ($\times$), 3 ($+$), and 4 ($\diamond$)

The poles for each version of synthetic data are also shown in a plot in Figure 3.19. The plot has a logarithmic scale along the horizontal (Ω) axis, while the vertical (ω) axis has a linear scale. That is because Ω is known to have values only on one side of the origin, while ω can be positive, negative, or zero. It is shown later that the error in estimation of Ω is greater, relative to the size of Ω , than the estimation error on ω relative to its size.

Each version of the data was implemented eight times with different coefficients for each mode. This was intended to mimic actual late-time radar data, which in theory has the same poles but different coefficients when measured from different aspect angles. Each data set was then called an aspect. For example, the Version 1, Aspect 1 data set was used for much of the analysis below, because the distribution of its poles in the s plane was similar to that of cylinders. The coefficients were chosen to separate the phase of each mode into different quadrants of the phase plane. The coefficients used with each version of synthetic data are shown in Table 3.2.

m	Aspect 1	Aspect 2	Aspect 3	Aspect 4
1	0.30, 0	0.40, 0	0.50, 0	0.60, 0
2	1.00, -0.30	0.90, -0.40	0.80, -0.50	0.70, -0.60
3	1.00, 0.30	0.90, 0.40	0.80, 0.50	0.70, 0.60
4	-0.30, -1.00	-0.40, -0.90	-0.50, -0.80	-0.60, -0.70
m	Aspect 5	Aspect 6	Aspect 7	Aspect 8
1	0.70, 0	0.80, 0	0.90, 0	1.00, 0
2	0.60, -0.70	0.50, -0.80	0.40, -0.90	0.30, -1.00
3	0.60, 0.70	0.50, 0.80	0.40, 0.90	0.30, 1.00
4	-0.70, -0.60	-0.80, -0.50	-0.90, -0.40	-1.00, -0.30

Table 3.2 Coefficients A_m , B_m for synthetic data, millivolts

3.3.2 Evaluating Noise Suppression by the LRHA. The LRHA, used in the M-TLS-MPM algorithm, acts to suppress noise in the data matrix. After application of the LRHA, the filtered version of the signal was extracted from the data matrix and compared to the signal prior to LRHA application.

Noise was added to the Synthetic Data Version 1, Aspect 1. The noise strength was set so that the SNR ranged from 1 to 25 dB. The algorithm was set to look for the correct number of poles, seven in this case.

The evaluation criterion was an ensemble average across all samples of the normalized temporal average of the squared error. For each signal, the error, the difference between the noisy signal and the uncorrupted signal, was computed at each sample time. The error of that signal was squared and summed over each sam-

ple period. The temporal sum of squared error was then normalized by dividing it by the temporal sum of the square of the uncorrupted signal. This error is designated the normalized temporal-average squared error (NTSE) in this document. An ensemble mean was then formed from the NTSE of each of the 100 noisy signals. The ensemble average is designated the mean normalized temporal-average squared error (MNTSE). Averaging squared errors first temporally and then in an ensemble sense is equivalent to averaging first in an ensemble sense and then temporally. The MNTSE was computed for each signal before the application of the LRHA, and after one, two, three, or four iterations of the LRHA. For the signal prior to the application of the LRHA, the NTSE and MNTSE were equal to the variance of the noise added to the signal. MNTSE was computed from 1 dB to 25 dB SNR, every 2 dB.

The first iteration of LRHA reduced the MNTSE by approximately 11.7 dB. This result was consistent across initial SNR levels from 5 dB to 25 dB, as shown in Figure 3.20. Additional iterations had much less effect on the noise. The second iteration yielded an additional approximately 0.5 dB reduction in the SNR. After the second iteration there was a slight increase in the MNTSE. The third iteration actually increased MNTSE by less than 0.06 dB at SNR levels above 1 dB. The fourth iteration produced a further increase in MNTSE of about 0.12 dB. The fifth produced an additional 0.13 dB increase, the sixth an additional 0.11 dB, the seventh an additional 0.09 dB. A noisy signal prior to application of the LRHA is shown in Figure 3.21 along with the uncorrupted signal. The signal after one iteration of the LRHA is shown in Figure 3.22 along with the uncorrupted signal. The signal after two iterations of the LRHA is shown in Figure 3.23 along with the uncorrupted signal. As is expected from the small change in NTSE due to the second iteration, there is no apparent difference between Figures 3.22 and 3.23.

The theory behind the LRHA is that every iteration will improve the parameter estimates. This was shown to be the case for all parameters at most SNR levels and




Figure 3.20 MNTSE (dB) of $y_H(t)$, normalized by mean $x(t)^2$, vs SNR, found via M-TLS-MPM with 0 to 2 iterations of LRHA




Figure 3.21 Synthetic Data Version 1, $y(t)$ with 13 dB SNR, dashed, and uncorrupted signal $x(t)$, solid




Figure 3.22 Synthetic Data Version 1, $y_H(t)$ 13 dB SNR and 1 iteration of LRHA, dashed, and uncorrupted signal $x(t)$, solid




Figure 3.23 Synthetic Data Version 1, $y_H(t)$ 13 dB SNR and 2 iterations of LRHA, dashed, and uncorrupted signal $x(t)$, solid

up to 5 integrations of the LRHA in Section 3.2.1. No theory has been stated that the mean MSE of the signal itself will decrease at each new iteration of the LRHA.

The signal $y_H(t_i)$ resulting from the LRHA is the one for which the M-TLS-MPM computes eigenvalues. It is a different signal than either the noise-corrupted original signal $y(t_i)$ or the underlying, uncorrupted signal $x(t_i)$. However, $y_H(t_i)$ is a good approximation of $x(t_i)$. Across a wide range of SNR values for $y(t_i)$, the LRHA reduced noise by approximately 12 dB in the cases cited.

If the goal of using the LRHA is to suppress noise in a signal, then two iterations are appropriate for this application. If reducing the error of an estimate is the key criterion, as it was in this research, then a point of diminishing returns is reached at about four iterations, as shown in Figures 3.5 through 3.8 in Section 3.2.1. Based on these findings, four iterations of the LRHA were used in all further applications of M-TLS-MPM this research.

3.3.3 Evaluating the Significance Criterion P . In this section the effect of the singular value significance level P on pole estimation is considered. As discussed in Section 2.3.1, the number of significant poles underlying a measured signal is unknown. The number can be estimated by counting the singular values of the data matrix $\mathbf{Y}$ that exceed a significance limit $\sigma_{max} \times 10^{-P}$. An appropriate setting for P is SNR/20, where the SNR is expressed in decibels, given that P represents the number of noise-free digits on a signal [120]. To reiterate the discussion in Section 2.3.2, increasing the number of noise-free digits in a signal by one increases the SNR by 20 dB. Given that the correct SNR is not known with certainty, some criterion must be determined from the data which will allow an appropriate setting for P .

Setting the P level too high allows the algorithm to produce some pole estimates based on eigenvalues that may represent the noise in the Hankel data matrix. That may result in spurious estimates. Setting P too low means that some

eigenvalues representing the underlying signal may be excluded from consideration, preventing the algorithm from estimating poles that are significant. It is anticipated that offsetting P from the prescribed value has less effect when the SNR is higher. The experiments reported in this section determined how far P can vary from the SNR-based value while still providing adequate estimation of poles.

The M-TLS-MPM technique was used with four iterations of the LRHA. The SNR was varied from 1 to 25 dB. The significance factor P was set to 1/20th of SNR (in dB) plus an offset. SNR was rounded up to the next tenth, so 1.25 was treated as 1.3. The offset ranged from -0.2 to 0.5. There were 100 runs for each offset at each SNR. In each case, the synthetic data known as Version 1, Aspect 1 was used. The noise added to each run of the signal was actually a repeatable, pseudo-random data set. The same 100 noise samples were used at each level of SNR. Results for an offset of zero and an SNR of 25, 17, 9, and 5 dB are shown in Figures 3.24 through 3.27, respectively.

The quality of estimation was evaluated by associating estimated poles with actual poles. The s -plane was divided into zones based on the oscillating rates (ω) of the actual poles. Zone limits were at $\omega = \pm 15, 45, 85,$ and 135 Grad/s. The horizontal lines in Figures 3.24 through 3.27 are the zone limits around each actual pole. Estimated poles within the zone of an actual pole were associated with it. Only the upper left quadrant of the s plane is shown. For each pole shown above the real axis, the lower left quadrant contained a conjugate. The mean error, error variance, and mean-squared error were computed based on the number of estimated poles associated with each actual pole.

The M-TLS-MPM algorithm provided good pole estimates from signals as noisy as 9 dB SNR, as seen in Figure 3.26. Below that, results became unacceptably poor, as seen in Figure 3.27. Error grew, but in addition some of the actual poles had no estimate associated with them at all in some trials. All complex poles occurred in pairs, regardless of SNR, P setting, or number of LRHA iterations, as was expected




Figure 3.24 Estimated ($\cdot$) and actual ($\diamond$) poles, 25 dB SNR, $P = 1.3$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits




Figure 3.25 Estimated ($\cdot$) and actual ($\diamond$) poles, 17 dB SNR, $P = 0.9$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits




Figure 3.26 Estimated ($\cdot$) and actual ($\diamond$) poles, 9 dB SNR, $P = 0.5$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits




Figure 3.27 Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.3$ (no offset), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits

for a real signal. The only exceptions were poles for which $\omega\Delta t = \pi$, which were screened out, as discussed in Section 3.2.2. Since M-TLS-MPM is intended to be applied in a laboratory setting, far in advance of actual target recognition, it is assumed that a SNR of 9 dB or higher can be attained in late time.

For SNR equal to 9 dB or greater and $P = \text{SNR}/20$ (no offset), each measurement had seven estimated poles, with each estimate near an actual pole. Below 9 dB SNR, there was often no estimate for the highest pole. Some bias can be seen in the pole estimates shown in Figures 3.24 through 3.27. The estimates of Ω tend to be low (toward a larger $|\Omega|$) for Ω_1 at all SNR levels and for each Ω at SNR 17 dB and below. The estimates of ω_2 and ω_4 were a bit high at 9 dB. At 5 dB SNR, the algorithm was ineffective.

With P offset by +0.1 from the prescribed value, that is, $P = \text{SNR}/20+0.1$, similar results were achieved for the no-offset case. Seven poles were estimated in each case down to 15 dB SNR, but there were some instances of additional, spurious estimates below that. Even down to 7 dB SNR, results were recognizable, although errors grew, particularly on damping rates, as seen in Figure 3.28. Sometimes the real axis pole was missed at 7 dB. At 5 dB SNR, the results were still somewhat useful. Error had increased, but there tended to be estimates available for most of the poles, as seen in Figure 3.29. Below 5 dB, the results were worthless.

For P offset by +0.2, up to four extra poles were estimated in few cases at 25 dB, and more spurious poles were estimated at lower SNR levels. The spurious pole estimates varied in position from measurement to measurement. Figure 3.30 shows the poles estimated at 5 dB SNR with P offset by +0.2. Also, some spurious estimates were introduced at SNR of 19 dB and below. However, there were few of these close to the actual poles, and as with the 25 dB SNR case, they did not occur in the same position from one measurement to another. Above 9 dB, the computed mean-square error increased because of the spurious poles in the same zone as each actual pole. For 5 dB SNR, there were better estimates of the actual poles with P




Figure 3.28 Estimated ($\cdot$) and actual ($\diamond$) poles, 7 dB SNR, $P = 0.5$ (offset +0.1), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits




Figure 3.29 Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.4$ (offset +0.1), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits




Figure 3.30 Estimated ($\cdot$) and actual ($\diamond$) poles, 5 dB SNR, $P = 0.5$ (offset +0.2), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits

offset by +0.2, shown in Figure 3.30 than there were with no offset, as shown in Figure 3.27. Given the low SNR, it is possible that some of the singular values larger than $\sigma_1 \times 10^{-P}$ were due to noise, rather than signal in the no-offset case. As the LRHA algorithm was iterated, the low-rank data matrix was regenerated using an incomplete set of singular values representing poles, as well as some singular values due to noise. The data matrix would then have been reconstructed poorly. With a higher P setting, more singular values were used to reform the data matrix, giving a greater chance that all singular values based on true poles were included. As the LRHA algorithm was iterated, noise was suppressed and some more singular values dropped out of the computation, but most or all of the true (signal-based, rather than noise-based) singular values remained.

With P offset by +0.3, more spurious estimates were introduced at SNR of 25 dB and below. The poles estimated at 25 dB SNR with P offset by +0.3 are shown in Figure 3.31. Given the number reported at 25 dB SNR, it may be inferred that some




Figure 3.31 Estimated ($\cdot$) and actual ($\diamond$) poles, 25 dB SNR, $P = 1.6$ (offset +0.3), using M-TLS-MPM with 4 iterations of LRHA, lines indicate ω zone limits

spurious poles were reported above that SNR level as well. Spurious poles tended to have much smaller damping rates $|\Omega|$ than the actual poles. As before, they did not occur in the same position from one measurement to another. The computed mean-square error increased at many SNR levels, compared to the no-offset case, because of the spurious poles in the same zone as each actual pole.

Figures 3.32 through 3.37 show the mean-squared error (MSE) of the pole estimates. Each actual pole was associated with only those estimated poles that fell within its zone, based on ω . The MSE was normalized by dividing it by the actual parameter value squared. The result was reported in decibels.

Figures 3.32 and 3.33 show the MSE found with P set to the prescribed value of $\text{SNR}/20$. The error on the damping rates Ω_m was much greater, relative to the parameter values, than the error on the oscillating rates ω_m at every level of SNR. The worst error was found on the real axis pole Ω_1 . To put the MSE numbers in perspective, a relative error of -26 dB equates to an error of 5% of the parameter




Figure 3.32 MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20$ (no offset)




Figure 3.33 MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20$ (no offset)




Figure 3.34 MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 + 0.2$




Figure 3.35 MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 + 0.2$




Figure 3.36 MSE (dB) of $\hat{\Omega}$, normalized by Ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 - 0.2$




Figure 3.37 MSE (dB) of $\hat{\omega}$, normalized by ω^2 , vs SNR, found via M-TLS-MPM with 4 iterations of LRHA and $P = \text{SNR}/20 - 0.2$

value. It can then be seen that the error on Ω_1 is likely to be greater than 5% when the SNR is less than 21 dB. The other damping parameters are likely to have greater than 5% relative error only when the SNR is under 18 dB. The curves in Figure 3.33 show that the ω estimates are likely to have less than 5% relative error whenever the SNR is greater than 1 dB, but the curves should not be trusted fully for SNR levels below 5 dB because the MSE is shown to decrease for some parameters below that point.

For an SNR of 9 dB and higher, the MSE plots in Figures 3.32 and 3.33 were straight lines. This pattern was seen previously in the case of KT data in Section 3.2.1, when the MSE was close to the Cramer-Rao bound (CRB). Since the straight-line MSE behavior was observed for both KT and synthetic data, it may be the case that the MSE for synthetic data in the straight-line region was also close to the CRB, although there is not sufficient evidence to draw that conclusion. The CRB was not computed for synthetic data, so the MSE performance relative to the CRB was not verified in this case. At low SNR, the MSE decreased in some cases (Ω_2 , Ω_3 , ω_3) because some actual poles did not have estimates associated with them. Those estimates that were present at higher SNR but missed below 9 dB would have been ones that had larger errors.

Figures 3.34 and 3.35 show the MSE found with P offset from the prescribed value of SNR/20 by +0.2. In this case, it was anticipated that additional, spurious pole estimates would be reported. Given the simple pole-association technique (by zones based on ω), the spurious poles were associated with the actual poles even if good estimates were also reported, resulting in larger MSE whenever spurious poles were present. The actual poles in the signal were still estimated with small error, and were consistent from measurement to measurement.

The elevated MSE values seen in Figures 3.34 and 3.35 were more indicative of a problem in data association than in pole estimation. As discussed above, pole estimates were associated with actual poles if their imaginary components $\hat{\omega}$ fell

within broad limits around each actual ω . A more elaborate data association algorithm might decide that some of the poles within the limits around a pole were not related to that pole. Such unassociated poles would not be included in the MSE computation. For the purposes of the P evaluation, however, it was sufficient to use the simple data-association algorithm of sorting poles by ω .

Figures 3.36 and 3.37 show the MSE found with P offset from the prescribed value of $\text{SNR}/20$ by -0.2 . In this case, it was anticipated that valid pole estimates would frequently fail to be reported at low SNR levels. At 11 dB and above, the MSE was the same as that found in the no-offset case. Below 11 dB, fewer poles were estimated. Below 9 dB SNR, the fourth pole pair ($\Omega_4 \pm j\omega_4$) was not estimated at all. There were no poles at SNR of 1 or 3 dB because P was set equal to zero, so at most one pole could have been estimated, if it corresponded to the largest singular value of the Hankel data matrix.

In this section it is shown that the significance criterion P , used to select the number of poles in M-TLS-MPM, has a substantial effect on the algorithm's output. The best performance is when P is set according to the SNR of the data, as in $P=\text{SNR}/20$. At low SNR levels, a slight deviation (offset) away from the prescribed value will result in either additional, spurious pole estimates (positive offset) or the loss of valid estimates (negative offset).

3.3.4 Signal Reconstruction and Error. The discussion of MPM variants in the previous sections only considered the ability to estimate poles in a damped-sine signal. Both poles and residues are unknown, so some attention should be paid to finding the coefficients (residues). The problems are separable, and once an estimate of the poles is available, it may be used to estimate the coefficients (residues). Every published variant of MPM estimates the poles first and then the residues [49, 51, 69, 119]. Each author finds the residues via a least squares solution using a Vandermonde matrix, shown in Equation (3.9).

To find the residue estimates $\hat{\hat{C}}$, the Vandermonde matrix [121] associated with the pole estimates is developed. This matrix reconstructs elements of the underlying signal based on the estimated poles. For each estimated pole $\hat{s}_n$, the term $\hat{z}_n = e^{\hat{s}_n \Delta t}$ is found. Note that $\hat{z}_n$ is an estimate of z_n which can be viewed in terms of the z -transform, a frequency domain analysis technique used with digital data [62]. The factor z_n represents the change in the n th term of the pole series (the signal) after one sample period. To find the effect after two, three, etc. sample periods, z_n (if constant, as it is here) is squared, cubed, etc. To find the effect of one term n of the SEM pole series over K sample periods, the following vector is formed:

$$[\hat{z}_n^0 \ \hat{z}_n^1 \ \hat{z}_n^2 \ \hat{z}_n^3 \ \cdots \ \hat{z}_n^K]^T \quad (3.8)$$

To account for the additive effect of each pole term in the series, several vectors are arranged into a Vandermonde matrix [45]. For N poles and K measurements, the matrix is

$$\mathbf{Z} = \begin{pmatrix} \hat{z}_1^0 & \hat{z}_2^0 & \hat{z}_3^0 & \hat{z}_4^0 & \cdots & \hat{z}_N^0 \\ \hat{z}_1^1 & \hat{z}_2^1 & \hat{z}_3^1 & \hat{z}_4^1 & \cdots & \hat{z}_N^1 \\ \hat{z}_1^2 & \hat{z}_2^2 & \hat{z}_3^2 & \hat{z}_4^2 & \cdots & \hat{z}_N^2 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \hat{z}_1^{K-1} & \hat{z}_2^{K-1} & \hat{z}_3^{K-1} & \hat{z}_4^{K-1} & \cdots & \hat{z}_N^{K-1} \end{pmatrix} \quad (3.9)$$

where

$$\hat{z}_n = \exp(\hat{s}_n \Delta t) \quad (3.10)$$

Of course, $\hat{z}_n^0 = 1$ for all $\hat{z}_n$. Estimates of the complex coefficients $\hat{\hat{C}}$ can be found using the least-squares solution [120]

$$\hat{\hat{C}} = (\mathbf{Z}^{T*} \mathbf{Z})^{-1} \mathbf{Z}^{T*} \bar{y} \quad (3.11)$$

For a real signal, all complex coefficients will appear in conjugate pairs. The solution in Equation (3.11) only applies when a matrix is thin, or overdetermined, i.e., with more rows than columns [121]. In our problem, more measurements than poles are needed to use Equation (3.11) to solve for the coefficients. If there were more poles than measurements, the coefficients could still be estimated. For this underdetermined matrix a least-norm solution would be used as follows [121]:

$$\hat{\hat{C}} = \mathbf{Z}(\mathbf{Z}\mathbf{Z}^{T*})^{-1}\bar{y} \quad (3.12)$$

As an alternative to the least-squares (LS) solution in Equation (3.11), a weighted least-squares (WLS) solution can be found using

$$\hat{\hat{C}} = (\mathbf{Z}^{T*}\mathbf{W}\mathbf{Z})^{-1}\mathbf{Z}^{T*}\mathbf{W}\bar{y} \quad (3.13)$$

In the WLS solution a weight matrix $\mathbf{W}$ is chosen, which must be nonsingular and symmetric [121]. If the weights are chosen so that $\mathbf{W} = \mathbf{I}$, the WLS solution is equivalent to an unweighted LS solution.

The WLS criterion was attempted in a few cases in which reconstruction error was poor. The weight matrices selected were diagonal matrices in which the elements on the main diagonal increased either linearly, as $1/I$, $2/I$, $\dots$, I/I for a signal with I samples, or as the square of the linearly increasing weight, $(1/I)^2$, $(2/I)^2$, $\dots$, $(I/I)^2$. The purpose for these choices of weights was to force more attention to the later parts of the time window, where signal strength faded. In some cases, reconstructed signals fit the original signal better early in the time window.

There was no thorough test of WLS effectiveness against unweighted LS effectiveness. In many cases, reconstruction error was determined to be acceptable, so WLS was unnecessary. In a few cases of poor reconstruction error in which WLS was attempted, the resulting error was the same as, or greater than, the reconstruction error found with unweighted LS. Further tests with a variety of weight matrices,

such as exponentially-tapered diagonal elements, may be a useful direction for future researchers.

Once poles and coefficients are estimated, it is simple to reconstruct the signal itself. The Vandermonde matrix is multiplied by the estimated coefficients:

$$\bar{y}_{REC} = \mathbf{Z}\hat{\mathbf{C}} \quad (3.14)$$

The elements of $\bar{y}_{REC}$ are a reconstructed version of the data samples $y(t_0)$, $y(t_1)$, $\dots y(t_{K-1})$. The more accurately the poles and coefficients are estimated (and the number of significant poles correctly determined), the more closely the reconstructed signal will match the original radar signal.

Reconstruction error is an important measure of signature quality, which helps select an appropriate model complexity, in terms of number of poles. Another measure of signature quality is how well the pole estimates match up to poles found via theory or other pole estimation techniques. However, comparison to other data sources can only be accomplished in the limited set of cases for which such data is available. Reconstruction error can be computed for late-time signals from any target, making it the most widely-applicable measure of signal quality.

In the target recognition algorithm, it is advantageous that the signatures should be as simple (comprising as few poles) as practically possible. Small size makes the computations in the filter simpler and quicker. However, if the number of poles used to reconstruct the signal is reduced, reconstruction may suffer.

The reconstruction error of the signal was the NTSE defined in Section 3.3.3, which was the temporal average of the squared error (in this case, the difference between the reconstructed and unmodified signals), divided by the mean-square signal. Normalized error was used because the signal strength varied greatly from one target to another, or even between two measurements of the same target at different azimuths. Normalized error allowed the reconstruction of two measurements to be

compared even if one signal was ten or more times the amplitude of another. The reconstruction error was computed as:

$$\begin{aligned}
e &= \frac{\sum_i [y(t_i) - y_R(t_i)]^2}{\sum_i y^2(t_i)} = \frac{\sum_i [y(t_i) - \sum_n C_n \exp(s_n t_i)]^2}{\sum_i y^2(t_i)} \\
&= \frac{\sum_i \{y(t_i) - 2 \sum_m \exp(\Omega_m t_i) [A_m \cos(\omega_m t_i) - B_m \sin(\omega_m t_i)]\}^2}{\sum_i y^2(t_i)}
\end{aligned} \tag{3.15}$$

Each radar signal's reconstruction was evaluated with two criteria, one subjective and the other quantitative. The first was a subjective judgment of a plot of the original and reconstructed signal together. If they overlaid each other with only small gaps, the reconstruction was good. It could then be assumed that one of two things was true. Either a sufficient set of the poles underlying the signal had been adequately estimated, or sufficient poles had been used that a forced fit had been achieved, regardless of the quality of individual pole estimates.

The possibility of a forced fit came up in early experiments with the Matrix Pencil Method. Given several hundred poles, the reconstructed signal could closely match every noise-induced twist and turn, even in a signal which did not necessarily fit the decaying exponential signal model. However, that possibility could be discounted because only 10 to 20 poles were estimated with the M-TLS-MPM code, rather than hundreds.

The second criterion used to assess signal reconstruction was the reconstruction error shown in Equation (3.15). An example signal is shown in Figure 3.38. This data was taken from a circular cross-section cylinder, three inches long and one inch in diameter, with the cylinder's axis parallel to the $\tilde{z}$ axis and 42 degrees from broadside incidence. The incident field was ϕ -polarized. The reconstruction error




Figure 3.38 Overlay of measured (solid) and reconstructed (dashed) signals, showing gaps: $P=1.3$, 14 poles used, reconstruction error = 4.8%.

calculated for this data set was 0.0477. The poles were found using the M-TLS-MPM algorithm discussed in Section 3.2 with $P = 1.3$ and four iterations of LRHA. The signal was reconstructed using WLS with linearly increasing weights.

Given the error seen in Figure 3.38, either the poles were not estimated correctly, enough of the correct ones were not included to capture the signal adequately, or noise in the signal prevented a match. While the signal does not appear to contain significant white noise, other forms of noise are present including Gibbs phenomenon and narrowband noise in the frequency-domain simulation data. The data was smoothed in the frequency domain to reduce step discontinuities that cause Gibbs phenomenon in the time-domain data, but some ringing was still present. Much better reconstruction was seen in the θ -polarized scattering from the same cylinder at the same angle. The scattered and reconstructed signals from the θ -polarized case are shown in Figure 3.39. The reconstruction error for this signal was 0.0020.




Figure 3.39 Overlay of measured (solid) and reconstructed (dashed) signals, showing gaps: $P=1.1$, 4 poles used, reconstruction error = 0.2%.

Increasing P slightly and allowing small number of additional poles to be estimated did not necessarily decrease the reconstruction error, and frequently increased it. If the increase in P allowed a single additional pole to be estimated, a real axis pole may have been lost and replaced by a pole pair. Even if the same number of poles was estimated, the pole estimates changed in some cases, resulting in a different fit between the reconstructed and noisy signal.

In this section it was shown that estimated poles can be used to reconstruct a signal matching the original signal to within a small amount of error. Subjective and quantitative criteria were developed. The quantitative criterion, reconstruction error, is a useful guide to how well the set of estimated poles captures the dynamics of the received signal. It should be supplemented with a subjective comparison of how well the reconstructed signal matches the received signal, particularly in cases where the reconstruction error is near a predetermined limit. Researchers using M-TLS-MPM should experiment with various P levels to achieve acceptable levels

of reconstruction error. When some estimate of the SNR of the late-time signal is available, P should be set in the range $\text{SNR}/20 \pm 0.2$. In this research, P was adjusted so that signal reconstruction error was on the order of 1%.

3.3.5 Application to Theoretical Sphere Data. The goal of all the experimentation with M-TLS-MPM was to provide a technique that will work with measured radar data. Such data consists of a theoretically infinite number of poles, an unknown number of which are significant.

An intermediate step before applying M-TLS-MPM to radar data was to apply it to theoretical sphere data. The sphere data was processed the same way as the radar data (as explained in Appendix D), but the late-time signal would be weaker than that of some other targets. Spheres are known to be weakly resonant relative to other objects such as wire dipoles [59]. The dominant scattering effect from a sphere in late time is the creeping wave [59]. Effects that are present in late time for other targets, such as travelling waves, multiple diffractions, etc., are not present in sphere scattering [59].

Scattering from a sphere is determined in the frequency domain by the Mie Series [129]. The Mie Series is described fully for the backscatter case in Section A.1 of Appendix A, and its convergence was examined. Mie Series data was transformed from the frequency domain into the time domain via the Inverse Fast Fourier Transform, as explained in Appendix D. The same procedure was performed on all radar data simulated or measured in the frequency domain. Since the pole locations have been found analytically for the sphere, the M-TLS-MPM results were compared to the analytic pole results. Unlike the Synthetic Data discussed previously, there is an infinite number of poles in the backscattered signal from a sphere. Only a few of these were significant enough to be found by M-TLS-MPM. The sphere's poles are discussed further in Section A.2 of Appendix A.

Mie-Series data was generated from 0.10 to 18.0 GHz every 0.01 GHz, for a sphere of three-inch radius. The frequency-domain data was tapered with a Hanning window prior to employment of an IFFT to find the impulse response. The Hanning window was used to suppress Gibbs phenomenon that would otherwise be caused by the step discontinuity at the end of the available data. The choice of the Hanning window over other windows and the window's effects on M-TLS-MPM results are both discussed in Section D.3 of Appendix D.

In Figure 3.40, it can be seen that M-TLS-MPM is able to estimate the sphere poles very well if a very high significance level P (explained in Section 2.3.2) is used. There was excellent agreement for the poles with the smallest damping rate $|\Omega|$, as seen by the $\times$ symbols that fell inside the circles. The theoretical poles with larger-magnitude damping rates were not estimated, most likely because their residues were small, meaning that their scattering modes were weak and their effect on overall scattering was insignificant. The set of $\times$ symbols running in a roughly horizontal line at around $\omega = 110$ Grad/s were due to Gibbs phenomenon, as explained in Section D.3 of Appendix D. The upper limit of the frequencies used to generate this data was 18 GHz, or 113.1 Grad/s. No poles were found at ω frequencies higher than 113.1 Grad/s. Since no noise was added to this data, there were potentially up to 15 digits of precision available, so P was set to 15.

Given the small number of spurious poles to the right of the theoretical poles in Figure 3.40, the significance level of $P = 15$ was slightly high for the data. For realistic data, however, it would be extremely high. Setting $P = 15$ implies a SNR of approximately 300 dB, far in excess of the SNR likely to be seen in an actual, late-time radar signal.

Figure 3.41 shows the same data, with $P = 6$. No noise has been added, but with the lower significance level, some valid singular values were suppressed prior to pole estimation. These are seen as circles with no estimate ($\times$) nearby. Some of the poles that were still significant had poorly-estimated damping rates. These are




Figure 3.40 Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 15$, and theoretical poles ($\circ$)

shown by the estimates ($\times$) to the right of the theoretical poles ($\circ$) in the center of the figure, with ω between 30 and 65 Grad/s. Many poles are missed or estimated badly. These poles were associated with weak scattering modes, so they contribute little to the late-time signature. A good signature for the sphere could be developed with a subset of the poles that are associated with larger residues. Using all the estimated poles, the late-time scattered signal was reconstructed. The reconstruction error was extremely low, 1.08×10^{-12} , indicating essentially perfect reconstruction. The signal was reconstructed again, this time with only the poles that were estimated well. Reconstruction error with only the well-estimated poles was 5.44×10^{-12} , still indicating very good reconstruction. The poles that were estimated poorly in this case contributed very little to the scattered signal.

With P set to a value more appropriate for the other data sources, very few poles were estimated. Figures 3.42 and 3.43 show the results for $P = 2$ and $P = 1$, corresponding to SNR values of 40 and 20 dB, respectively. The vertical scale of




Figure 3.41 Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 6$, and theoretical poles ($\circ$)

the plots has been changed to focus on the area where poles were estimated. Only theoretical poles with damping rates $\Omega > -20$ GN/s were shown to improve the figure's clarity. For the $P = 2$ case, the reconstruction error was 3.87%. For the $P = 1$ case, the reconstruction error was 27.35%. In both cases the reconstructed signal appeared to be sufficiently different from the theoretical signal that target recognition performance for a sphere and these pole estimates would likely be poor.

While various techniques have been used in the literature to find the poles of simple-shape targets, much less attention has been paid to determining which poles have a significant effect on late-time scattering. In the case of the sphere of three-inch radius, few poles are significant, given SNR levels from 20 to 40 dB, corresponding to $P = 1$ to 2.

It appeared that one reason the higher frequency poles were lost was the way the data was processed. The Hanning window, applied to the sphere data prior to the IFFT, reduced the amplitude of the higher-frequency data. The Hanning




Figure 3.42 Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 2$, and theoretical poles ($\circ$)




Figure 3.43 Estimated poles ($\times$) from 3-inch radius sphere, Mie Series data, using M-TLS-MPM, no noise added, $P = 1$, and theoretical poles ($\circ$)

window is explained in Section D.3 of Appendix D. With higher-frequency data suppressed, the power in late-time scattering modes at those frequencies would also have been suppressed to some extent. The singular values in the Hankel data matrix associated with higher frequencies would have been reduced, and then eliminated if too much lower than the largest singular value. The same window was applied to all radar data, suppressing the higher-frequency poles. A Hanning window is commonly applied to radar data prior to application of the IFFT.

To examine the effect of the Hanning window, the Mie Series data was transformed into the time domain again, this time with no window applied. The poles found for the unwindowed $P = 2$ case were similar to the $P = 2$ case with a Hanning window, except that two additional poles were estimated that had ω values near the upper limit of the data. The poles are shown in Figure 3.44 and the late-time signal is shown in Figure 3.45. The reconstructed signal is also shown in the figure, but the reconstruction error was so small, 0.0066%, that the difference between the signals is hard to discern. The high-frequency oscillations that look somewhat like white noise in the picture are an example of Gibbs phenomenon. The additional freedom in reconstruction offered by the extra high- ω poles were sufficient to allow an excellent match.

From this section it can be seen that, although a large number of poles may be identified for a target, only a few of these will have such a significant effect on the late-time scattering that they will be estimated accurately by the M-TLS-MPM. The problem should not be seen as a failure of the M-TLS-MPM in relation to other MPM variants because its performance was shown to be comparable to or better than MMP3 and TLS-MPM in Section 3.2.1. The technique was shown to be capable of estimating poles accurately from synthetic data (damped sinusoid series) in Section 3.3.3. For certain targets with weak late-time scattering, such as the sphere, this technique will produce poor pole estimation and poor signal reconstruction unless the available data is extremely clean. For strong late-time




Figure 3.44 Estimated poles ($\times$) from 3-inch radius sphere, unwindowed Mie Series data, using M-TLS-MPM, no noise added, $P = 2$, and theoretical poles ($\circ$)




Figure 3.45 Late time signal and reconstructed signal from 3-inch radius sphere, unwindowed Mie Series data, using M-TLS-MPM, no noise added

scattering, such as the θ -polarized scattering from a cylinder seen in Section 3.3.4, the answer provided by M-TLS-MPM is sufficient to reconstruct the signal well.

3.4 *Combining Data from Various Azimuths*

Previous sections in this chapter discussed estimation of poles from a single signal measured at one aspect angle (azimuth and elevation). Poles estimated from a single aspect of a target may be used to identify that target [102, 103], assuming that all of the relevant poles of the target were significant at the aspect angle considered. However, the overall goal of this research is an aspect-independent signature comprising every significant pole at many or all aspect angles. In this section an algorithm which develops a multiple-aspect signature is discussed.

An algorithm called Pole Association and Averaging (PA&A) was developed to combine poles from several measurements of a single target, taken at different aspect angles. PA&A averaged nearby poles, using all poles found from each measurement of a target.

The key to combining poles from different measurements or simulations at different aspect angles is data association, the process of determining which bits of data (pole estimates in this case) correspond to the same underlying reality (an actual pole). In Section 3.3.3, estimated poles were associated with actual poles of synthetic data by simply comparing the $\hat{\omega}$ value to some established zones, shown in Figures 3.24 through 3.27 as horizontal lines. The result was that spurious poles were included in the MSE computations, increasing error. To produce a single signature that represented a target from a range of aspect angles, a more careful algorithm was required, one that took into account differences in Ω as well as ω . Since actual poles were unknown for these targets, zones could not be established in advance. Even if theoretical poles were known, no information was available to show which ones would be most significant.

Given the results with synthetic data in Section 3.3.3, particularly Figures 3.24 through 3.26, it was anticipated that the pole estimates for a target measured or simulated from several angles would fall in a small region of the s -plane around the actual pole location. The pole estimates had $\hat{\omega}$ values both above and below the true parameter value, and appeared to be unbiased. A mean of several $\hat{\omega}$ estimates would therefore yield a result close to the actual ω value.

While the $\hat{\omega}$ estimates appeared to be unbiased, the $\hat{\Omega}$ estimates were clearly biased toward lower values. A mean of the $\hat{\Omega}$ values associated with a particular actual Ω would be strongly biased downward as well, since the lower (larger-magnitude) $\hat{\Omega}$ values would dominate the mean. The $\hat{\Omega}$ values are displayed on a logarithmic scale in this document to account for their wide variation. A mean of $\log \hat{\Omega}$ would be less affected by the downward bias.

For each pole, PA&A determined which, if any, poles from different measurements were near it. Nearness was determined by arbitrary tolerance limits (a ‘box’) in Ω and ω . Since the tolerance limits on Ω were different from the limits on ω , the box was rectangular rather than square. For each pole found within the box, a new box was drawn and all poles within it were combined into a single group. The number of poles in a group was used to determine whether the group likely represented a true pole or a collection of spurious poles that fell nearly together. The mean value of each group of poles was treated as a single, aspect-invariant pole.

Care was taken to avoid adding two poles from the same measurement into the same group, because nearby poles from the same measurement represented different scattering effects. Groups were split into two or more separate groups to avoid averaging poles from the same measurement.

The quality of the multiple-azimuth pole set was evaluated by using the new poles to reconstruct all the measured signals. The signals were reconstructed as explained in Section 3.3.4.

Polarization presented a dilemma. Each target was measured or simulated from two orthogonal polarizations, identified as θ or ϕ and explained in Appendices B and C. The scattered signal from a target could be very different in amplitude and shape depending on the polarization viewed. A clear example was the cylinder scattering shown in Figures 3.38 and 3.39. The θ -polarized scattering in Figure 3.39 appeared to be dominated by a single damped sine, while the ϕ -polarized scattering in Figure 3.38 was more complicated. While the same poles are theoretically present regardless of polarization, the coefficients are definitely polarization-dependent [6]. The result is that modes considered significant at one polarization may be completely insignificant at the orthogonal polarization, particularly in the case of cylinders. A pole-estimation algorithm such as M-TLS-MPM would be unable to estimate the poles associated with insignificant scattering. Therefore, poles were combined separately for each polarization. This resulted in two separate signatures for each target. For best results, an operational target recognition system could alternate the polarization used to illuminate a target on successive pulses. The antenna required for dual-polarization illumination is more complex than that for single-polarization but is certainly feasible [16, 81].

Examples using simulated radar data are shown in the next section. Poles are estimated from several aspect angles via M-TLS-MPM. Combined poles are estimated via PA&A to develop target signatures.

3.5 *Developing Target Signatures with M-TLS-MPM*

The M-TLS-MPM technique was used to estimate the poles of several targets. These included targets measured as described in Section B.1 of Appendix B or targets simulated with various techniques described in Section C.2 of Appendix C. Polarization is given in terms of the spherical coordinates θ , the angle away from the $\tilde{z}$ -axis, and ϕ , the angle around the $\tilde{z}$ -axis. Diagrams of target geometry, including coordinate axes, are given in Appendices B and C.

The significance level P was set so that the reconstruction error would be on the order of 1%. If much more error was present, the set of poles was considered insufficient to represent the late-time scattering adequately. If the error was much less than that, it was possible that scattering could be modelled with fewer poles, resulting in smaller matrices and less computational load in the target recognition system, as explained in the next chapter.

For each target, P was set differently, although a commonly-used setting was $P = 1.1$. Setting P higher than that often resulted in the reporting of spurious poles. The spurious poles were marked by both large variability in Ω and higher (less negative) Ω values, particularly higher than theoretical poles, if known. Setting P lower than that frequently resulted in increased reconstruction error, most likely due to failing to estimate significant poles. The increased error percentage and the visual agreement between the reconstructed signal and the original signal were such that target recognition could have been adversely affected.

Poles were estimated for data from each aspect angle, but the poles from a few aspect angles were left out of the signature computation for each target. Data from the excluded aspect angles was used later as test data. In this way, test data did not directly influence the signature computation and the test data could be independent from the model, resulting in more meaningful tests.

Simulated scattering from three right circular cylinders was examined. Simulation techniques are described in Section C.2 of Appendix C. The M-TLS-MPM results are compared to published numerical results for cylinder poles, which are discussed in Section C.1.1 of Appendix C. The published results were found via a completely different technique from MPM, which did not utilize late time scattering, but instead relied on analysis of coupling matrices [86]. The θ -polarized results (magnetic field perpendicular to the cylinder's axis) were dominated by a single damped sine term, which was consistent with published numerical results for cylinder poles. The ϕ -polarized signals (electric field perpendicular to the cylinder's axis) did not

exhibit the same behavior as the θ -polarized signals at the same aspect angles. That is because different scattering phenomena, creeping waves, dominate the ϕ -polarized case. While the same poles are theoretically present, the overall late-time signal was much weaker in the ϕ case, so late-time results were dominated by effects other than the SEM poles, such as noise and Gibbs oscillations. For the 3" long cylinder, the θ -polarized scattering was on the order of 50 times greater in amplitude than the ϕ -polarized scattering, or on the order of 2500 times greater (about 34 dB greater) in power.

An observer familiar with forced resonances (from steady-state illumination) on a cylinder may expect that the unforced resonances are similar in frequency. Forced resonances include resonant modes along the length of the cylinder such that the cylinder length equals integer multiples of a half wavelength. Such modes would be most strongly excited by θ -polarized incident waves. In addition, there will be resonant modes such that the circumference equals integer multiples of the wavelength. These modes would be most strongly excited by ϕ -polarized incident waves [4, 59]. These modes do not correspond to the poles found by Merchant, et al. [86] or Vechinski and Shumpert [137]. Natural (unforced) resonant frequencies do vary with cylinder dimensions, but are not found at the same oscillating frequencies as the forced resonant modes. For a right circular cylinder of 1" diameter, setting wavelength equal to circumference πd results in a frequency of $\omega = (2\pi c)/(\pi d) = 23.62$ Grad/s (c is the speed of light). For a cylinder of 5" length, setting $\lambda/2$ equal to length L results in $\omega = (2\pi c)/(2L) = 7.42$ Grad/s.

The poles and combined poles found for the cylinder of length 5 inches and diameter 1 inch are shown as dots in Figures 3.46 and 3.47. Combined poles, those produced by the PA&A algorithm from the poles estimated for each azimuth, are shown as squares. The bottom of each image includes part of the lower half-plane. Some poles computed by Merchant, et al. [86], specifically those with the smallest-magnitude damping and oscillating rates, are shown as circles. Two pole pairs es-

estimated via M-TLS-MPM are close to Merchant's pole pairs. For this target and θ polarization, reconstruction error was typically on the order of 0.01% and always under 0.3%. With error that small, reducing P was an option, but with only three pole pairs in the signature, there was no need to reduce the number of poles. For ϕ polarization, 14 combined poles were found and reconstruction error ranged from 3% to 5%.

As was noted in Section 3.3.3, the estimation error on the Ω parameter of each pole was greater, relative to parameter size, than the error on the ω parameter. That led to the horizontal pattern of pole estimates (dots) in Figures 3.24, 3.25, and 3.26. In Figure 3.47, horizontal displacement of pole estimates can again be seen. The weakness of the ϕ -polarized signal leads to a smaller SNR and greater error on the Ω estimate than was seen in the θ -polarized case in Figure 3.46.

A three-inch cylinder, with a diameter of one inch, was also examined. Results are shown in Figures 3.48 and 3.49. Prior to the results shown in these figures, poles were found for $P = 2$, but after comparison with lower- P results, it was determined that many of the poles found for $P = 2$ were spurious. The spurious poles could be distinguished by their smaller-magnitude damping rates. For $P = 1.1$, only two pole pairs were found, one of which matched the lowest- ω published pole well. Reconstruction error was less than 1% even with only two combined pole pairs. Reconstruction error was under 0.4% for θ polarization and ranged from 4% to 9% for ϕ polarization, for which 14 combined poles were found. As was seen with the five-inch cylinder, larger relative errors on Ω lead to a horizontal pattern of dots seen in Figure 3.49.

Scattering from a six-inch-long cylinder, with a diameter of two inches, was also simulated. The poles and combined poles estimated for this target are shown in Figures 3.50 and 3.51. This cylinder had the same proportion of length to radius as the three-inch cylinder. In theory, the poles found for this target would have half the values of ω and Ω as the three-inch cylinder. The combined poles found for




Figure 3.46 Poles estimated ($\cdot$) from 5'' by 1'' diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.47 Poles estimated ($\cdot$) from 5'' by 1'' diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.48 Poles estimated ($\cdot$) from 3'' by 1'' diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.49 Poles estimated ($\cdot$) from 3'' by 1'' diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)

θ polarization in Figure 3.50 matched the published poles well. As with the other cylinders, the ϕ polarization combined poles shown in Figure 3.51 did not match published results, because the much-weaker ϕ -polarized scattering was more affected by noise than the θ -polarized scattering. The estimated real axis combined pole in the ϕ case was a fairly close match to the published real axis pole.

A 4.5" square plate with a thickness of 0.1" was illuminated in a simulation with ϕ -plane polarization. The poles and combined poles found for the plate are shown in Figure 3.52 and 3.53 for ϕ and θ polarizations, respectively. The bottom of Figure 3.52 includes part of the lower half-plane, so some lower-half-plane poles close to the real axis are shown near their upper-half-plane counterparts. Published poles determined by Sun, Chen, Nyquist and Rothwell [130] are shown as well. The poles estimated via M-TLS-MPM (shown as dots) occurred in tight clusters, indicating good agreement between the various aspect angles. One of the pairs found was a good match to the published pole pair at $-2.175 \pm j5.331$. Reconstruction error with the combined poles (shown as squares) was under 1.4%.

The plate target was examined again, this time with θ -plane polarization at the same aspect angles. Results are shown in Figure 3.53. One combined pole pair was a good match to both the ϕ -plane measurement and the lowest published pole at $-2.175 \pm j5.331$. The other combined pole had about the same $|\omega|$ value as a pole found in the ϕ case, but greater $|\Omega|$. Again, since a portion of the lower half-plane is shown, the lower right corner of Figure 3.53 includes both halves of some pole pairs. Reconstruction error for this target was under 0.4%.

The poles found via M-TLS-MPM were compared to the poles computed by Long [68] for cubes and rectangular blocks. Long listed normalized poles for a cube and two rectangular blocks with side lengths in the same proportion as the targets named Block 112 and Block 123. These targets are discussed in detail in Appendix C. The combined poles did not match the published poles at all. The damping rates observed for Long's poles were much greater in magnitude than the damping rates




Figure 3.50 Poles estimated ($\cdot$) from 6'' by 2'' diameter cylinder, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.51 Poles estimated ($\cdot$) from 6'' by 2'' diameter cylinder, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.52 Poles estimated ($\cdot$) from 4.5" square plate, ϕ -pol, using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.53 Poles estimated ($\cdot$) from 4.5" square plate, θ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)

estimated by M-TLS-MPM. However, the signals were easily reconstructed using the combined poles. The discrepancy between Long’s poles and the ones estimated here may be due to differences between the Method of Moments implementations used by Long and used in this research. The discrepancy is not due to the PA&A procedure, because the individual poles estimated by M-TLS-MPM at each azimuth did not match those found by Long. Long simulated the surface currents on a section of each block, with an eighth of the total surface area (one quarter of the area of three faces), including one corner. He determined the response of the complete object by symmetry, which is a common procedure in electromagnetic analysis [105, 144] and is analogous to decomposing a function $f(x)$ into odd and even components. Long’s choice of boundary conditions at the edges of the sections may have changed the current modes present in the answer, but it is not clear from the article [68] which boundary conditions were used.

For this research, the response of the entire surface was simulated for each block. Simulating the entire surface of the object eliminated any errors that could have been introduced by combining simulated portions of the object.

It should also be noted that Long’s procedure for finding poles is completely different in concept from M-TLS-MPM. Long found poles by finding zeros of the mutual impedance matrix developed in a scattering simulation. His method makes no use of late-time scattering, so it is unable to determine which poles are associated with significant scattering, as opposed to weak, insignificant scattering. Given what was seen in the case of the sphere and cylinders, it is likely that the higher- ω poles were associated with weaker scattering modes, and were too weak to show up. Long’s technique is also unable to make use of measured scattering data. In contrast, M-TLS-MPM relies on late-time scattering data from simulations or measurements, which may contain various forms of interference in addition to actual scattering. While M-TLS-MPM is somewhat resistant to white noise, it will react to narrowband interference such as Gibbs oscillations. Such interference may result in pole estimates

that represent the interference as much as the resonant scattering. The simulated data may be corrupted by Gibbs Phenomenon, although a smoothing procedure was applied to it to reduce the effect of step discontinuities found in the frequency-domain data. For some targets, the received signals are shown in the same plots as signals reconstructed using estimated poles. The plots demonstrate that the estimated poles accurately model the received signals.

The poles and combined poles found for Block 112 for ϕ polarization are shown in Figure 3.54 along with published poles. The estimated poles did not match the published poles. For this target, reconstruction error was around 1.5% at all aspect angles with six poles.

Figure 3.55 shows the poles and combined poles for θ -polarized scattering from Block 112. Different poles were found, although the pole pair in the neighborhood of $-0.4 \text{ GN/s} \pm j0.8 \text{ Grad/s}$ was close to the lowest published pole. Nine poles were found and reconstruction error ranged from 0.6% to 4.3%.

Examples of signal reconstruction for Block 112 are shown in Figures 3.56 and 3.57. The target aspect used in Figure 3.56 was $\theta = 36^\circ$ and $\phi = 0^\circ$. For Figure 3.57 the target aspect was $\theta = 39^\circ$ and $\phi = 0^\circ$. The coordinate system for each target is explained in Appendix C. The target aspects chosen for these figures had reconstruction errors that were typical of all measurements for this target at the indicated polarizations. The poles used to reconstruct the signals were adequate to capture the signals' dynamics with little error. The amplitude of the late-time scattered signal was on the order of 0.003 millivolts per meter for both polarizations, making it comparable to the ϕ -polarized scattering from cylinders.

The poles found for Block 123 are shown in Figure 3.58. The incident polarization was in the ϕ direction. Published poles [68] are shown as circles. One of the pole pairs estimated via M-TLS-MPM is close to the published pole pair at $-3.4614 \text{ GN/s} \pm j14.0453 \text{ Grad/s}$. For this target, reconstruction error was under 5% (under 3% at most aspect angles) with 14 poles.




Figure 3.54 Poles estimated ($\cdot$) from Block 112, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.55 Poles estimated ($\cdot$) from Block 112, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.56 Signal reconstruction example for Block 112, ϕ -pol, error 1.6%




Figure 3.57 Signal reconstruction example for Block 112, θ -pol, error 2.9%




Figure 3.58 Poles estimated ($\cdot$) from Block 123, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.3$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.59 Poles estimated ($\cdot$) from Block 123, θ -pol, using M-TLS-MPM, no noise added, $P = 1.6$, combined poles ($\square$), and published poles ($\circ$)

The θ -polarized scattering from Block 123 was analyzed, and the poles and combined poles are shown in Figure 3.59. The combined poles were different from the published poles, although in the θ -polarized case, one pole pair was close to the lowest- ω published pole at $-4.5784 \text{ GN/s} \pm j8.5858 \text{ Grad/s}$. Since a portion of the lower half-plane is shown, the lower edge of Figure 3.59 includes the complex conjugate of one pole pair at approximately $-0.45 \text{ GN/s} - j1.5 \text{ Grad/s}$. Reconstruction error for most aspect angles (22 out of 32) was under 4%, although a few data sets had much higher reconstruction error. The worst case was 45% error, indicating that the combined set of poles did not represent the target well at those aspect angles and θ polarization. A good set of poles has not yet been achieved for this target, when illuminated with θ polarization. Block 123 may require different tolerances in the PA&A algorithm, resulting in more poles being included in the signature. Another approach is to use more than one signature at each polarization (θ and ϕ) to describe this target, which is undesirable because a goal of this research was to represent a target with the most compact signature possible. Still, it may be required for this target. If multiple signatures are used at each polarization for Block 123, it will be necessary to determine how many signatures to use, and which aspect angles to associate with each signature. Further refinement of the Block 123 signature via a different P setting in M-TLS-MPM, modification of PA&A and/or development of separate signatures is left to future research.

Examples of signal reconstruction for Block 123 are shown in Figures 3.60 and 3.61. The poles used to reconstruct the signals were adequate to capture the signals' dynamics with little error. For the θ -polarized case, the signal reconstruction shown has higher error than desired, but this fit may still allow target recognition. The data sets that had higher reconstruction error than this one will probably result in incorrect target recognition. The amplitude of the scattered signal in Figure 3.60 was on the order of 0.001 mV/m and that shown in Figure 3.61 was on the order of

0.010 mV/m. The scattered signals were comparable in amplitude (and thus power) to the ϕ -polarized scattering from cylinders.

The poles and PA&A combined poles found for the 3" cube are shown in Figure 3.62. The incident polarization was in the ϕ direction. Published poles are shown as circles. One of the pole pairs estimated via M-TLS-MPM is close to the published pole pair at $-2.3689 \text{ GN/s} \pm j14.3571 \text{ Grad/s}$. For this target, six poles were found and reconstruction error ranged from 0.8% to 3.0%.

The θ -polarized scattering from the 3" cube was analyzed, and the poles and combined poles are shown in Figure 3.63. Six poles were found, which were different from the ones found in the ϕ polarization case. The estimated poles fell in the lower- ω region of the s -plane, with most ω components less than 5 Grad/s. Reconstruction error for most aspect angles was under 3%, although a few data sets had error between 4% and 16%.

Examples of signal reconstruction for the cube are shown in Figures 3.64 and 3.65. The error was sufficiently small at all aspect angles for ϕ -polarization and most angles for θ -polarization, that target recognition using this set of poles should be effective. The signal amplitude shown in Figures 3.64 and 3.65 was on the order of 0.010 mV/m, making it comparable to that found from ϕ -polarized scattering from cylinders.

A target called Combo A was made up of Block 112 and Block 123, as described in Section C.2 of Appendix C. The poles found for Combo A for ϕ and θ are shown in Figures 3.66 and 3.67, respectively. There are no published poles for this target. For the ϕ case, 10 poles (five pole pairs) were found and reconstruction error was under 1%. In the θ case, six poles (three pairs) were found and reconstruction error was under 1.6%, even though P was set to the lowest value used for any targets, 0.7. For this target, the estimated poles were compared to the estimated poles of the two objects which made up the target. The poles for Block 112 and Block 123 are shown as circles in Figures 3.66 and 3.67.




Figure 3.60 Signal reconstruction example for Block 123, ϕ -pol, error 2.9%




Figure 3.61 Signal reconstruction example for Block 123, θ -pol, error 4.8%




Figure 3.62 Poles estimated ($\cdot$) from cube, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.63 Poles estimated ($\cdot$) from cube, θ -pol, using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and published poles ($\circ$)




Figure 3.64 Signal reconstruction example for 3'' cube, ϕ -pol, error 2.2%




Figure 3.65 Signal reconstruction example for 3'' cube, θ -pol, error 4.1%




Figure 3.66 Poles estimated ($\cdot$) from Combo A, ϕ -pol using M-TLS-MPM, no noise added, $P = 1.1$, combined poles ($\square$), and combined poles of components ($\circ$)




Figure 3.67 Poles estimated ($\cdot$) from Combo A, θ -pol, using M-TLS-MPM, no noise added, $P = 0.7$, combined poles ($\square$), and combined poles of components ($\circ$)

For ϕ polarization, there is close correspondence between a pole pair at approximately $-0.35 \text{ GN/s} \pm j7.5 \text{ Grad/s}$ and a pole pair found for Block 112. Several pole estimates contributing to the combined pole pair at $-0.42 \text{ GN/s} \pm j3.5 \text{ Grad/s}$ are close to poles from both Block 112 and Block 123. The combined pole could be due to scattering from either of these component targets. The combined pole at $-0.4 \text{ GN/s} \pm j15 \text{ Grad/s}$ is close to a pole from Block 123 at $-0.3 \text{ GN/s} \pm j15 \text{ Grad/s}$. For θ polarization, the combined pole pair near $-0.55 \text{ GN/s} \pm j7.4 \text{ Grad/s}$ is close to the combined pole from Block 112 at $-0.4 \text{ GN/s} \pm j7.5 \text{ Grad/s}$.

The poles for Combo A appeared to consist of some poles estimated for Block 112, some poles estimated for Block 123, and some other poles. That is consistent with theory. A target composed of multiple objects may have late-time scattering dominated by the interaction between the targets, or by the target signatures. However, conclusions drawn for this target must be considered preliminary. As was discussed above, some adjustment to the PA&A algorithm may be necessary to better capture the signature of Block 123. Some of the combined poles in the Block 123 signature may move, so the comparison between the Combo A signature and the Block 123 signature will change. This analysis is left to future research.

A target called Combo B was made up of a 4.5" square plate and a Block 112. It is described further in Section C.2 of Appendix C. The poles found for Combo B for ϕ and θ are shown in Figures 3.68 and 3.69, respectively. There are no published poles for this target. As with Combo A, the estimated poles were compared to the estimated poles of the two objects which made up the target. The poles for the square plate and Block 112 are shown as circles in Figures 3.68 and 3.69.

For the ϕ case, nine poles were found and reconstruction error was under 1%. The combined pole at approximately $-1.5 \text{ GN/s} \pm j9.0 \text{ Grad/s}$ was an excellent match to a pole from the square plate. There are individual poles displaced in Ω value from a square plate combined pole at $-0.055 \text{ GN/s} \pm j1.5 \text{ Grad/s}$. If more poles were estimated in this area (such as via higher P setting in M-TLS-MPM), or if PA&A




Figure 3.68 Poles estimated ($\cdot$) from Combo B, ϕ -pol using M-TLS-MPM, no noise added, $P = 0.9$, combined poles ($\square$), and combined poles of components ($\circ$)




Figure 3.69 Poles estimated ($\cdot$) from Combo B, θ -pol, using M-TLS-MPM, no noise added, $P = 1.2$, combined poles ($\square$), and combined poles of components ($\circ$)

tolerances were different, a combined pole could have been produced here. In the θ case, five poles were found and reconstruction error was under 6%. The error was particularly bad in the θ case, so various significance settings up to $P = 2$ were considered to attempt to find an error minimum. There was no significant change in error from $P = 1.2$, with five poles, to $P = 2.0$, with nine poles. Given the dispersed nature of the estimated poles, it is likely that some spurious poles were estimated. Reducing P below 1.2 would focus on poles representing actual scattering dynamics, although error is unlikely to decrease. Varying the tolerance box and other parameters in the PA&A algorithm would result in the combined pole at $-0.8 \text{ GN/s} \pm j7.5 \text{ Grad/s}$ being interpreted as two or more nearby poles. One of these would be closer to the Block 112 pole at $-0.4 \text{ GN/s} \pm j7.5$. Further development of the signatures for the Combo B target is left to future research.

To summarize the results in this section, the M-TLS-MPM technique was applied to the scattered signals of several targets. The targets were primarily of simple geometric shapes, but some were combinations of simple shapes. The results were mixed. In some cases, results matched theory and other published results well. In other cases, the pole estimates did not match well, but signals reconstructed with the estimated poles matched the original signals well.

For sphere theoretical data, M-TLS-MPM can find many poles accurately, but only in the case of very low noise. Sphere resonance modes are weaker at higher frequencies, and are reduced by the window applied to the data in the frequency domain, so even small amounts of noise mask their effects. Similar results were found for simulated scattering data from cylinders. For each cylinder, two to four poles were found for the θ -polarized scattering that matched previously published poles. Other published poles were likely associated with weak scattering modes, thus were not significant contributors to late-time scattering. Signal reconstruction was good, even with those few poles, indicating that the estimated poles adequately modelled the important dynamics of the signal.

As with the sphere, only the published poles for the cylinders and square plate with the smallest $|\omega|$ and $|\Omega|$ were significant enough to be estimated. The scattering from a square plate matched the lowest- ω published pole pair well given either θ or ϕ polarization. The cylinders matched the lowest- ω published poles in the analysis of θ polarization scattering. The ϕ -polarized scattering from circular cylinders did not match the published poles because scattering at that polarization was weak and dominated by Gibbs oscillation and noise. The ϕ -polarized scattering was in the neighborhood of 34 dB lower in power than the θ -polarized scattering. However, the received signals were reconstructed well with the estimated poles, as was true in the θ polarization case. The poles estimated for ϕ polarization may represent noise and Gibbs more than the cylinder's actual resonance. Given the much weaker late-time signal from cylinders at ϕ polarization, it is important to estimate target signatures using two orthogonal polarizations. A target recognition system exploiting late-time scattering should also make use of two orthogonal polarizations, to ensure the θ polarization is captured.

The poles found via M-TLS-MPM did not match the poles computed by Long [68] for cubes and rectangular blocks. However, the signals were easily reconstructed using those poles, indicating that the estimated poles represented the dynamics actually present in the signals, although these were not the poles reported in the literature via Long's technique.

Long simulated portions of the target and deduced surface currents on the entire object via symmetry. The boundary conditions he used at the symmetry planes are not clearly apparent from the article [68]. It would be useful to revisit Long's work, simulating blocks via MoM or some other computational technique and finding poles via the impedance matrix. Such work is left to future research.

The frequency-domain data for the cube and rectangular blocks contained some step discontinuities which would give rise to a significant amount of Gibbs phenomenon, which could give rise to spurious pole estimates in M-TLS-MPM. For the

block and cube targets, the signal power in late time was comparable in amplitude and power to that of the ϕ -polarized scattering from cylinders. Like the cylinders (given ϕ -polarization), the poles found for the cube and blocks failed to match the published poles found via impedance matrix analysis. Since the signal was so weak, it is plausible that noise and Gibbs oscillations were significant contributors in late time, and that the poles estimated by M-TLS-MPM represented Gibbs oscillations more than the late-time scattering predicted by SEM theory. That should not be viewed as a failure or error in the simulation used in this research, but as a limitation of the applicability of its results. The same simulation code was used for the square plate, with good pole estimation and signal reconstruction results.

If late-time backscatter scattering for the cube and blocks was so weak that it could not be found among Gibbs oscillations, then the signals for those targets may be too weak for proper pole estimation via M-TLS-MPM or any similar technique. Scattered signals from directions other than the backscatter direction (opposite of the incident direction) may provide stronger signals in late time. If stronger (and less noise-dominated) signals are not available for analysis, weakly resonant targets such as isolated cubes and blocks may be unsuited for identification by the target recognition technique described in Chapter IV. The technique may be limited to more strongly resonant targets.

Although Gibbs oscillation appeared to be a significant source of error in this research, it will be less apparent in an actual application of this technique. Gibbs phenomenon is found in signals that have been translated from the frequency domain into the time domain. Step discontinuities in the frequency domain produce the oscillations [62]. In actual operation of a system using this algorithm, the signals of interest would be generated by an impulse radar, and would not contain Gibbs oscillations. The signals would contain other forms of noise, some of which can be characterized as additive, white, Gaussian noise, to which M-TLS-MPM exhibits some resistance.

The poles found by M-TLS-MPM tended to be stable over a wide variety of θ aspect angles. They were not precisely aspect-independent, but tended to have little variation in ω . The variation in Ω was somewhat larger relative to the parameter value, but still consistent with the results seen for synthetic data at approximately 9 dB SNR. Signatures were developed by combining the poles found at various aspect angles for a given target at a single polarization.

When poles appeared near the same location when estimated from data from many azimuths, the PA&A algorithm was able to average them and produce a combined pole that represented an average of the estimates of all poles in that location. The parameter values used in PA&A were selected to account for the pole distributions seen in most targets, i.e., separate clusters of poles. For some targets, such as Block 123, the poles did not appear in neat, distinguishable clusters. In that case, parameter variations within PA&A could have a significant effect on the combined poles used in the target signature. More research is needed to determine the proper settings within PA&A to find the combined poles for those targets that lead to the best signal reconstruction. Data from M-TLS-MPM with several settings of the P parameter should be input to PA&A to look for interactions in parameter variations in both programs.

3.6 Chapter Summary

M-TLS-MPM is an improved variant of MPM which combines features of TLS-MPM and MMP3. It makes use of the iterative LRHA to suppress noise in the signal prior to eigenvalue estimation. The LRHA converges quickly, and little change is seen in parameter estimates after 4 iterations.

In tests with complex-valued signal data, the technique compared favorably with both MMP3 and TLS-MPM. For noise power less than -13 dB (SNR greater than +9.4 dB), the benefit was slight because TLS-MPM and MMP3 were already achieving parameter estimation errors close to the Cramer-Rao bound computed by

Kumaresan and Tufts for this data. At -11 dB noise power (+7.4 dB SNR), the estimates of ω were superior to those found either by TLS-MPM or MMP3 by 10 to 15 dB.

In tests with real-valued signals, M-TLS-MPM was able to estimate the parameters of a signal with three pole pairs and a real axis pole correctly. The results depended on an accurate assessment of the SNR of the data, which was used to set a significance threshold that determined the number of poles estimated. The ω estimates appeared to be unbiased, although some bias was apparent in the estimates of Ω . Variance was greater in the Ω case as well. The worst bias and variance were seen in the estimate of the real axis pole.

This research is the first to characterize the noise-suppression ability of the LRHA. It was found that a single iteration of the LRHA can reduce noise (defined as MNTSE) by 11.7 dB. An additional iteration results in further noise reduction of 0.5 dB, but after that there is no improvement.

For theoretical sphere poles and published poles of cylinders and plates, the higher- ω poles appeared to be associated with weaker scattering modes, so M-TLS-MPM was unable to form estimates of them. Lower- ω poles were estimated accurately for the sphere and plate, as well as the θ -polarized case for the cylinders. For the plate and θ -polarized cylinders, signal reconstruction was good, even with only one or two pole pairs, indicating that the estimated poles adequately modelled the important dynamics of the signal.

For the sphere, signal reconstruction was poor. That is due to the large number of scattering modes that are marginally significant. The poor signal reconstruction may prevent or limit the use of this algorithm with spherical or nearly-spherical targets. Fortunately, existing targets of interest, such as land mines or exterior components of vehicles, tend to be of other shapes, such as cylinders, rectangular plates, and thin wires.

Scattering from several simple-shape targets was simulated in the frequency domain using two implementations of the method of moments. The late-time portions of the scattered signals were analyzed using the M-TLS-MPM algorithm. Signatures were developed by combining the poles found at various aspect angles for a given target at a single polarization.

The ϕ -polarized scattering (perpendicular to the axis) from circular cylinders did not match the published poles because scattering at that polarization was weak and dominated by Gibbs oscillation and noise. The weakness of ϕ -polarized, late-time scattering from cylinders is a problem that can be solved by using pulses of different polarizations in the radar transmitter used for target recognition. If pulses in a short sequence have differing polarization (e.g., alternating vertically-polarized pulses with horizontally-polarized pulses) then the response from a cylindrical object will be dominated by the stronger, θ -polarized scattering.

The poles found via M-TLS-MPM did not match the poles computed by Long [68] for cubes and rectangular blocks. However, the signals were easily reconstructed using those poles, indicating that the actual poles of the signals were found, although these were not the poles reported in the literature via Long's technique. The discrepancy between Long's poles and the ones estimated here may have been due to the weakness of the late-time scattered signal from those targets. The frequency-domain data for the cube and rectangular blocks used in this research contained some step discontinuities which would give rise to some Gibbs oscillation, which may have been a significant part of the simulation result in late time, given the true signal's weakness. A signal that is generated in the time domain, such as ultra-wideband, short-pulse radar scattering, will not have Gibbs phenomenon. Such scattering will still be subject to other forms of noise. Noise approximating the characteristics of additive, white, Gaussian noise can be expected. This technique is designed to be somewhat resistant to such noise, and in fact is more resistant than the previous MMP3 and TLS-MPM methods, as was shown in Section 3.2.1.

The poles found by M-TLS-MPM were stable over a wide variety of θ angles for most targets. They were not precisely aspect-independent, but tended to have little variation in ω . The variation in Ω was somewhat larger relative to the parameter value, but still consistent with the results seen for synthetic data at approximately 9 dB SNR.

The pole estimation technique presented in this chapter is a critical component of an overall target recognition method. The quality of the signature library built with the M-TLS-MPM algorithm directly affects the performance of the target recognition algorithm developed in the next chapter.

IV. Reading from the Library: Stochastic Estimation Research

4.1 Overview

The target recognition algorithm involves comparing a signal to a known set of signal models. The set of signal models is termed a library. The sources of the data from which the library is built were presented in Chapter III, and the signatures used to fill the library were developed. In this chapter signals are compared to signatures in the library. The targets for which signatures were included in the library were simple shapes and combinations of two simple shapes. Targets are discussed more completely in Appendix B (for measured targets) and Appendix C (for simulated targets). For each test of the algorithm, from four to ten targets were included in the library. Typically four were used so that display and analysis of results would be accomplished more easily.

To recognize the target, several Kalman filters were implemented in a parallel structure known as MMAE-MAP [76]. Each filter contained a linear system model for a target type, based on the natural frequencies (Ω_m, ω_m) estimated as presented in Chapter III. Each filter's model was a set of simultaneous linear differential equations which were solved in an equivalent discrete-time form, i.e. via difference equations. A diagram of the target recognition algorithm was shown in Figure 1.1, and is repeated for review in Figure 4.1:

The signal model, with data sampled at evenly-spaced points t_i , is represented by:

$$y(t_i) = 2 \sum_{m=1}^M \exp(\Omega_m t_i) [A_m \cos(\omega_m t_i) - B_m \sin(\omega_m t_i)] u(t_i - t_0) \quad (4.1)$$

The unit step function $u(t_i)$ zeros out data prior to the start of late time, t_0 . The scattered radar data from times prior to t_0 did not fit the signal model in Equation (4.1), so it was not used by this algorithm. Each target can have a different




Figure 4.1 Target Recognition System Diagram

number of significant poles in its late-time scattered signal. For signal models used to represent measured data, typically up to ten damped sinusoid terms and up to one non-oscillating ($\omega_m = 0$) term were used.

In addition to the frequencies, each term in the signal is characterized by two coefficients (A_m, B_m). The values of these coefficients are unknown, but they can be assumed constant during the duration of the scattered data. They depend on the aspect angle, which is the angle from which the target is viewed. They are also dependent upon the polarization of the scattered field. The radar receiver is assumed to have a linear polarization, and the target's orientation relative to the antenna polarization is assumed not to change during the duration of a signal. The coefficients determine the amplitude of the signal. If the late-time scattered signal is viewed as the solution to a differential equation, then the coefficients are the initial conditions assigned to that equation.

A concern early in the research was that lack of knowledge of the parameters would impede the target recognition process. A more elaborate version of the target recognition algorithm was considered, which attempted to estimate the coefficients along with the signal. The coefficients were specific to each target, so coefficient estimation had to be performed within each Kalman filter to be effective. The elaborate version implemented an adaptive parameter estimation (APE) technique that found the maximum likelihood estimates of states and uncertain parameters via Rao Scoring [76]. This technique was first applied to Kalman filters by Abramson [99] for parameters in the covariance matrices of the measurement noise ($\mathbf{R}$) and the dynamics noise ($\mathbf{Q}_d$). It was applied by Maybeck [75] to the state transition (Φ), control input ($\mathbf{B}_d$) and measurement output ($\mathbf{H}$) matrices of the system model. Each filter in the MMAE-MAP filter bank included its own APE subroutine. This led to a great increase in both code complexity and required computation. It was later determined that an appropriate filter tuning strategy could allow a Kalman filter with good frequency estimates to reconstruct the signal, even with incorrect coefficient values and no APE. A much simpler recognition algorithm without APE was possible, so the algorithm implementing APE was abandoned prior to a complete evaluation.

The synthetic data described in Section 3.3.1 was used to test the recognition algorithm in addition to measured and computational data. Filter models were developed based on the poles estimated in Chapter III. For the performance evaluation, the following metrics were compared: fraction of attempts resulting in correct identification, fraction of attempts resulting in incorrect identification, and fraction of attempts in which the algorithm had insufficient information to declare an identification. Results for low-SNR cases are also displayed in terms of ‘confusion matrices,’ as explained in Section 4.3.2. Varying amounts of noise were added to the signal to gauge the effects of noise on the performance of the MMAE-MAP and E-Pulse [112, 113] algorithms.

The technique developed in this research applies only to late-time scattering. That is because late-time scattering can be represented in a form that easily converts into a simple, analytic transfer function with a limited number of significant terms that can be identified. If the early-time scattering could be turned into an analytic transfer function, a state-space model could be developed from the transfer function and a very similar algorithm could be applied to all scattering. Instead, while the frequency-domain scattering from an object can be viewed as a transfer function [87], it is not an analytic one, but a set of complex-valued points measured or simulated at various frequencies.

Linear system models for target late-time scattering are presented in Section 4.2. The MMAE-MAP target recognition algorithm is presented along with test results in Section 4.3. For comparison, a target recognition algorithm using the E-Pulse technique is presented in Section 4.5.

4.2 *Applying the Kalman Filter*

To develop the linear system model, a transfer function is found for each frequency term (pole-pair transform) of Equation (2.29). This result can be obtained via the sum of two conjugate terms of Equation (2.16) or via the Laplace transform of one term of Equation (4.1):

$$\begin{aligned} T_m(s) &= \frac{2A_m(s - \Omega_m) - 2B_m\omega_m}{(s - \Omega_m)^2 + \omega_m^2} \\ &= \frac{2A_ms - 2(A_m\Omega_m + B_m\omega_m)}{s^2 - 2\Omega_ms + (\Omega_m^2 + \omega_m^2)} \end{aligned} \tag{4.2}$$

Equation (4.2) describes a second-order dynamic system. Examples of second-order systems include spring-mass systems and parallel RLC circuits [30], both of which are described below.

When discussing second-order systems, the term ‘damping’ is often used. A system may be described as underdamped, overdamped, or critically damped, based

on the damping ratio ζ . The term ‘damping’ used in this way is distinct from the damping frequency or rate Ω , the real part of the natural frequency s , although Ω is related to ζ . An underdamped system has a response which oscillates and has a decaying magnitude, in other words, a damped sinusoidal response. Overdamped and critically damped systems have a non-oscillating response. The quantities associated with natural resonance and system response concepts are compared and contrasted in Table 4.2. The concept of damping ratio is not used in this research; it is presented here only to distinguish it from the damping rate Ω .

	Natural Resonance Concepts	Second-order System Concepts
Natural Frequency	$ s $	$\omega_{natural}$
Damping Ratio	$ \Omega / s $	ζ
Damping Frequency (Rate)	$ \Omega $	$\zeta\omega_{natural}$
Oscillating Frequency	ω	$\omega_{natural}\sqrt{1-\zeta^2}$

Table 4.1 Comparison of Resonance and Response Concepts

Equation (4.2) is a linear combination of two rational transfer functions. A rational function is the ratio of two polynomials. A 2nd-order transfer function applicable to a variety of physical systems has the form [30]:

$$Y(s) = \frac{s + C}{s^2 + 2\zeta\omega_{natural}s + \omega_{natural}^2} \quad (4.3)$$

where C is a constant. The roots of the denominator of Equation (4.3) are:

$$s, s^* = -\zeta\omega_{natural} \pm \omega_{natural}\sqrt{1-\zeta^2} = \Omega \pm j\omega \quad (4.4)$$

Two examples where Equations (4.3) and (4.4) apply are spring-mass systems and parallel RLC circuits, which consist of resistive, inductive, and capacitive elements. For a spring-mass system, $\omega_{natural} = \sqrt{K/M}$, and $\zeta = f/(2\sqrt{KM})$, where K=spring constant, M = mass, and f is friction [30]. For a parallel RLC circuit, $\omega_{natural} = \sqrt{LC}^{-1}$, $\zeta = \sqrt{LC}/(2RC)$, $\Omega = (-2RC)^{-1}$, and $\omega = \sqrt{(LC)^{-1} - (4R^2C^2)^{-1}}$,

where R is resistance, L is inductance, and C is capacitance [95]. Either way, the differential equation is of the same form, so the Laplace-domain transfer function is of the same form.

Returning to the transfer function for late-time scattering, it is noted that some signatures include poles on the real axis, which correspond to non-oscillating exponential decay in the time domain. If ω_m is zero, the transfer function simplifies somewhat. The conjugate pole pair then becomes a single pole on the $-\Omega$ axis, with a real residue (coefficient) and the transfer function:

$$T_m(s) = \frac{2A_m(s - \Omega_m)}{(s - \Omega_m)^2} = \frac{2A_m}{s - \Omega_m} \quad (4.5)$$

Each conjugate pole pair, i.e., each pair of poles s_m, s_m^* in which $s_m^* = \Omega_m - j\omega_m$, was represented by two states based on the Ω_m and ω_m frequencies from Equation (4.1). Each real axis pole will be represented by a single state. The coefficients A_m and B_m from Equation (4.1) were placed in the control input vector. The control vector $\bar{u}$ represented the waveform of the incident radar pulse, i.e., the pulse transmitted from the radar to the target.

The measurement $\bar{y}(t_i)$, control input $\bar{u}(t_i)$, and residual $\bar{r}(t_i)$ for the experiments conducted, had only one element each, but are shown herein as vectors instead of scalars for consistency with the Kalman filter equations shown in Section 2.4.1. The residual covariance $\mathbf{A}$ and the measurement noise covariance $\mathbf{R}$ also had one element each, so the term ‘variance’ applies to each of them, rather than ‘covariance’, but they are shown as matrices to be consistent with Section 2.4.1.

The target signatures used in this research may consist of ten or more conjugate pairs of poles, plus possibly one or two real axis poles. For ease of presentation, an example with only five poles is shown. A five-pole model with two pole pairs and one real pole in observer canonical form [106] looks like this:

$$\mathbf{F} = \begin{pmatrix} 2\Omega_1 & 1 & 0 & 0 & 0 \\ -\Omega_1^2 - \omega_1^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2\Omega_2 & 1 & 0 \\ 0 & 0 & -\Omega_2^2 - \omega_2^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \Omega_3 \end{pmatrix} \quad (4.6)$$

$$\mathbf{B} = \begin{pmatrix} 2A_1 \\ -2(A_1\Omega_1 + B_1\omega_1) \\ 2A_2 \\ -2(A_2\Omega_2 + B_2\omega_2) \\ 2A_3 \end{pmatrix} \quad (4.7)$$

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 \end{pmatrix} \quad (4.8)$$

$$\mathbf{D} = 0 \quad (4.9)$$

It is also possible to formulate the model in control canonical form [106] as was done by Pascoe, Wood, Maybeck, and Wood [102, 103].

The signal model is a continuous-time process with sampled data. The system dynamics model equations were developed in continuous-time form and converted to sampled-data form. To develop a linear system model representing this signal, frequencies were expressed in units of gigradians per second (Grad/s) and giganepers per second (GN/s), or tens of Grad/s and tens of GN/s, or hundreds of Grad/s and hundreds of GN/s. To be consistent, the sample period used to generate the filter equations were in units of nanoseconds (ns), tenths of ns, or hundredths of ns. The order of magnitude chosen for the units depended on the size of the highest frequency of the target. Since frequencies (Ω_m, ω_m) in the signal model are always multiplied

by times ($t_i = i\Delta t$), the products ($\Omega_m t_i, \omega_m t_i$) are unaffected by the change in units, as are the z -Transform frequencies $z_m, z_m^* = \exp(\Omega_m t_i \pm j\omega_m t_i)$.

The system dynamics matrix $\mathbf{F}$ contains elements defined as $2\Omega_m$ and $-(\Omega_m^2 + \omega_m^2)$. If these quantities were instead expressed in units of nepers per second and radians per second, the magnitudes of these terms would be on the order of 10^{10} and 10^{20} , in the same matrix as constants equal to 1. Such a matrix would be ill-conditioned, meaning that computations involving the matrix would be very sensitive to perturbations [38]. An appropriate scalar associated with the condition of a matrix for numerical precision in mathematical operations (multiplication, addition, etc.) may be expressed as a condition number, which is the ratio of the largest singular value to the smallest [73, 77]. Using frequency units of gigaradians and giganepers, and time units of nanoseconds in the signal model, the matrices of the linear system model can be well-conditioned. The choice of units affected the size of the off-diagonal terms of the state transition matrix Φ , which changed the condition number of that matrix as well, but to a lesser extent than the $\mathbf{F}$ matrix.

The condition of the $\mathbf{F}$ matrix was not directly a concern, since $\mathbf{F}$ did not appear explicitly in the filter equations, but the Φ matrix derived from $\mathbf{F}$ was definitely a concern. The Φ matrix affected the state covariance $\mathbf{P}$ during the propagation phase of the filter, as shown in Equation (2.65).

A dynamics noise model appropriate for the system was generated. This noise represented uncertainty inherent in the linear system model that represents the actual system. For the late-time transient signal shown in Equation (4.1), the most significant form of uncertainty is the possibility of errors in the frequencies Ω_m and ω_m of each term in the damped sinusoid series. Another source of uncertainty is the presence of unmodeled terms. While there are an infinite number of damped sinusoid terms associated with a particular target, only a finite number of terms are significant in a given stream of data due to noise and the limits of precision. Other terms exist, but are insignificant due to small size of the coefficients, A_m and B_m .

Uncertainty due to insignificant terms is not modeled; such uncertainty is considered part of the noise. To generate a noise model, consider a two-state system representing a single pole pair. A five-state system is shown subsequently. Development of the noise model began with the dynamics equation in continuous-time form:

$$\frac{d\bar{x}(t)}{dt} = \mathbf{F}\bar{x}(t) + \mathbf{B}\bar{u}(t) + \mathbf{G}\bar{w}(t) \quad (4.10)$$

Substituting in two-state partitions of the in Equations (4.6) and (4.7), the dynamics equation becomes:

$$\begin{aligned} \frac{d}{dt} \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} &= \begin{pmatrix} 2\Omega_1 & 1 \\ -\Omega_1^2 - \omega_1^2 & 0 \end{pmatrix} \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} \\ &+ \begin{pmatrix} 2A_1 \\ -2(A_1\Omega_1 + B_1\omega_1) \end{pmatrix} u(t) + \begin{pmatrix} p_1 \\ -p_1\Omega_1 + q_1\omega_1 \end{pmatrix} w_1(t) \end{aligned} \quad (4.11)$$

and $w_1(t)$ is of strength Q_1 . The noise strength is defined by

$$E\{w_1(t)w_1(t + \tau)\} = Q_1\delta(\tau) \quad (4.12)$$

where $E\{f\}$ is the expected value of a random variable f and δ is the Dirac delta function.

A five-state version of the dynamics noise model is:

$$\mathbf{G}\bar{w}(t) = \begin{pmatrix} p_1 & 0 & 0 \\ -p_1\Omega_1 + q_1\omega_1 & 0 & 0 \\ 0 & p_2 & 0 \\ 0 & -p_2\Omega_2 + q_2\omega_2 & 0 \\ 0 & 0 & p_3 \end{pmatrix} \begin{pmatrix} w_1(t) \\ w_2(t) \\ w_3(t) \end{pmatrix} \quad (4.13)$$

The dynamics noise $\bar{w}(t_i)$ has an assumed strength of

$$\mathbf{Q} = \begin{pmatrix} Q_1 & 0 & 0 \\ 0 & Q_2 & 0 \\ 0 & 0 & Q_3 \end{pmatrix} \quad (4.14)$$

The structure of the $\mathbf{G}$ and $\mathbf{Q}$ matrices reflects uncertainty in the system model due to errors in the frequencies Ω_m and ω_m , $m = 1, 2, 3$. The terms p_m and q_m were based on the uncertainty in the parameters as estimated by the M-TLS-MPM algorithm as discussed in Chapter III. The Q_m weight associated with each mode is used for coarse tuning to reflect the quality of the estimates on any pair of states (or single state for a non-oscillating term).

The strength of the dynamics noise assumed for each term of the signal model in Equation (4.1) is independent of the noise assumed for the other terms. The values Q_m ($m = 1, 2, 3$), in the $\mathbf{Q}$ matrix are tuned to allow the state or pair of states representing each term of the signal model to match that term. The matrix $\mathbf{G}\mathbf{Q}\mathbf{G}^T$ is a block-diagonal matrix with a two-by-two block for each pair of states representing an oscillating term and a one-by-one block representing a non-oscillating term, if any:

$$\mathbf{G}\mathbf{Q}\mathbf{G}^T = \begin{pmatrix} p_1^2 Q_1 & \Gamma_1 Q_1 & 0 & 0 & 0 \\ \Gamma_1 Q_1 & \Upsilon_1 Q_1 & 0 & 0 & 0 \\ 0 & 0 & p_2^2 Q_2 & \Gamma_2 Q_2 & 0 \\ 0 & 0 & \Gamma_2 Q_2 & \Upsilon_2 Q_2 & 0 \\ 0 & 0 & 0 & 0 & p_3^2 Q_3 \end{pmatrix} \quad (4.15)$$

where

$$\Gamma_m = (-p_m^2 \Omega_m + q_m^2 \omega_m), \quad m = 1, 2 \quad (4.16)$$

and

$$\Upsilon_m = (p_m^2 \Omega_m^2 - 2p_m \Omega_m q_m \omega_m + q_m^2 \omega_m^2), \quad m = 1, 2 \quad (4.17)$$

Note that since Ω is always negative, the product $-p\Omega$ is always positive and greater uncertainty on Ω increases uncertainty overall.

A state space representation of a transfer function is not unique: an infinite number of state-space representations are possible. It is therefore incumbent upon the modeler to choose a representation which suits his purpose. The Kalman filters were initially implemented in the observer canonical form [106]. This form is attractive because all states are observable in this form, which can be a very convenient property. Observability refers to the ability to reconstruct or estimate states of a given system, given knowledge of measurements and control inputs. Observability is the ability to determine states of a system uniquely from a finite number of observations of the system's output [26].

The stability of the system must also be considered. The Routh-Hurwitz criterion is applied to assess stability [106]. Since the system consists of the sum of independent damped sinusoids, each of which is the solution to a differential equation of the same form, stability can be shown for one term and thereby demonstrated for all terms. With each term shown to be stable, it can be asserted that the sum of these terms is also stable. To check for stability, the Routh array is assembled for the characteristic polynomial (the denominator) of the transfer function, and shown below.

s^2	1	$\Omega^2 + \omega^2$
s^1	-2Ω	0
s^0	$\Omega^2 + \omega^2$	0

Table 4.2 Routh Array

The Routh-Hurwitz criterion is that the number of polynomial roots in the right half-plane is equal to the number of sign changes in the first column of the array. Since Ω is constrained to be negative, the value of the term -2Ω is positive, so there are no right-half plane roots. It is known from the unforced nature of the response that Ω must be negative for physically realizable late-time scattering.

In addition to system model stability, the stability of the filter derived from that model must also be considered. The filter's stability may be established by first demonstrating the observability and controllability of the system model [76]. The observer canonical form is intended to provide complete observability for all the states. Complete observability is the property of a system model (not of an actual system) that the value of all states over a time period can be determined exactly by the measurement and control information known during that time period.

A more stringent criterion is that of stochastic observability, which also implies complete observability [77]. For a system to be stochastically observable, there must exist real parameters α , β , and integer N such that $0 < \alpha < \infty$, $0 < \beta < \infty$, $0 < N < \infty$, and

$$\alpha \mathbf{I} \leq \sum_{j=i-N+1}^i \Phi^T(t_j, t_i) \mathbf{H}^T(t_j) \mathbf{R}^{-1}(t_j) \mathbf{H}(t_j) \Phi(t_j, t_i) \leq \beta \mathbf{I} \quad (4.18)$$

where the inequality $\mathbf{X} \leq \mathbf{Y}$ indicates that $\mathbf{Y} - \mathbf{X} \geq \mathbf{0}$, that is, $[\mathbf{Y} - \mathbf{X}]$ is positive semidefinite. Positive semidefiniteness can be verified by checking that all eigenvalues of the matrix are nonnegative.

Here it is useful to review some properties of the state transition matrix $\Phi(t_i, t_j)$. The transition after several time steps $\Phi(t_i, t_j)$ is the product of the transitions for the individual time steps:

$$\Phi(t_{i+N}, t_i) = \Phi(t_{i+N}, t_{i+N-1}) \Phi(t_{i+N-1}, t_{i+N-2}) \cdots \Phi(t_{i+1}, t_i) \quad (4.19)$$

In this research, $\Phi(t_i, t_j)$ depends only on the time difference $(t_i - t_j)$, so the state transition for each sample period, $\Phi(t_{i+1}, t_i)$, is the same matrix. The transition after N sample periods, from t_i to t_{i+N} , is thus $[\Phi(t_{i+1}, t_i)]^N$. Reversing the order of the times, as in $\Phi(t_i, t_{i+1})$, indicates a backward transition. The backward transition $\Phi(t_i, t_{i+1})$ is simply the inverse of the forward state transition matrix $\Phi(t_{i+1}, t_i)$. A transition over zero time intervals is given by $\Phi(t_i, t_i) = \mathbf{I}$, the identity matrix. Given

a time-invariant state transition matrix $\Phi(t_{i+1}, t_i)$ (i.e., derived from a constant $\mathbf{F}$ matrix), along with constant $\mathbf{H}$ and $\mathbf{R}$, the test of stochastic observability simplifies somewhat, to the requirement that all the eigenvalues of:

$$\sum_{j=i-N+1}^i \{[\Phi^{-1}(t_{i+1}, t_i)]^{i-j}\}^T \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} [\Phi^{-1}(t_{i+1}, t_i)]^{i-j} \quad (4.20)$$

are real, positive, and finite.

Controllability is the property of a system model (not of an actual system) that the value of any state at an initial time can be driven to any desired final value in finite time using a control input vector [26]. If the value of all states can be changed as desired by the control inputs commanded during that time period, the model is termed completely controllable. A more stringent condition is stochastic controllability, which implies complete controllability [77]. For a system to be stochastically controllable, there must exist real parameters α , β , and integer N such that $0 < \alpha < \infty$, $0 < \beta < \infty$, $0 < N < \infty$, and

$$\alpha \mathbf{I} < \sum_{j=i-N+1}^i \Phi(t_i, t_j) \mathbf{G}_d(t_{j-1}) \mathbf{Q}_d(t_{j-1}) \mathbf{G}_d^T(t_{j-1}) \Phi^T(t_i, t_j) < \beta \mathbf{I} \quad (4.21)$$

Given the time-invariant nature of the one-sample state transition matrix $\Phi(t_{i+1}, t_i)$ as well as $\mathbf{G}_d$ and $\mathbf{Q}_d$, the test of stochastic controllability simplifies a bit, to the requirement that all the eigenvalues of:

$$\alpha \mathbf{I} < \sum_{j=i-N+1}^i \Phi(t_{i+1}, t_i)^{i-j} \mathbf{G}_d \mathbf{Q}_d \mathbf{G}_d^T [\Phi(t_{i+1}, t_i)^{i-j}]^T < \beta \mathbf{I} \quad (4.22)$$

are real, positive, and finite.

The stochastic observability and stochastic controllability of each model was verified prior to use. The lesser standards of complete observability and complete controllability were fulfilled by implication. Because the models were stochastically

observable and controllable, the Kalman filters based on those models were therefore uniformly asymptotically globally stable [77].

The state transition matrix $\Phi(t_{i+1}, t_i)$ must now be computed. In this problem, $\mathbf{F}$ is constant, so Φ is the inverse Laplace transform of the resolvent, $\Phi^L(s) = [s\mathbf{I} - \mathbf{F}]^{-1}$ [77]. In the five-pole example in observer canonical form, $\Phi(t_i, t_{i-1})$ is thus the inverse Laplace transform of

$$\begin{aligned} \Phi^L(s) &= \begin{pmatrix} s - 2\Omega_1 & -1 & 0 & 0 & 0 \\ \Omega_1^2 + \omega_1^2 & s & 0 & 0 & 0 \\ 0 & 0 & s - 2\Omega_2 & -1 & 0 \\ 0 & 0 & \Omega_2^2 + \omega_2^2 & s & 0 \\ 0 & 0 & 0 & 0 & s - \Omega_3 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} \Phi_1^L(s) & \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{2 \times 2} & \Phi_2^L(s) & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{1 \times 2} & \mathbf{0}_{1 \times 2} & \Phi_3^L(s) \end{pmatrix} \end{aligned} \quad (4.23)$$

The inverted matrix is too large to be displayed all together on the page, so it is displayed in blocks. The bold zeros are zero matrices of the indicated dimensions. The blocks $\Phi_1^L(s)$ and $\Phi_2^L(s)$ are defined as

$$\Phi_m^L(s) = \begin{pmatrix} \frac{s}{s^2 - 2s\Omega_m + \Omega_m^2 + \omega_m^2} & \frac{1}{s^2 - 2s\Omega_m + \Omega_m^2 + \omega_m^2} \\ \frac{-\Omega_m^2 - \omega_m^2}{s^2 - 2s\Omega_m + \Omega_m^2 + \omega_m^2} & \frac{s - 2\Omega_m}{s^2 - 2s\Omega_m + \Omega_m^2 + \omega_m^2} \end{pmatrix} \quad (4.24)$$

for $m = 1, 2$, and the one-element block $\Phi_3^L(s)$ is defined as

$$\Phi_3^L(s) = \frac{1}{s - \Omega_3} \quad (4.25)$$

The matrix is then converted into the time domain:

$$\Phi(t - t_0) = \begin{pmatrix} \Phi_1(t') & \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{2 \times 2} & \Phi_2(t') & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{1 \times 2} & \mathbf{0}_{1 \times 2} & \Phi_3(t') \end{pmatrix} \quad (4.26)$$

in which $t' = t - t_0$ and, for $m = 1, 2$,

$$\Phi_m(t') = \begin{pmatrix} \cos(\omega_m t') - \frac{\Omega_m}{\omega_m} \sin(\omega_m t') & \frac{1}{\omega_m} \sin(\omega_m t') \\ -\frac{\Omega_m^2 + \omega_m^2}{\omega_m} \sin(\omega_m t') & \cos(\omega_m t') + \frac{\Omega_m}{\omega_m} \sin(\omega_m t') \end{pmatrix} \exp(\Omega_m t') \quad (4.27)$$

and $\Phi_3 = \exp(\Omega_3 t')$.

Matrices Φ and $\mathbf{B}_d$ are computed from the continuous-time matrices $\mathbf{F}$ and $\mathbf{B}$, as explained in Section 2.4.1. Equations are shown in a form specific to this research problem in that some matrices are shown below as constants. Time dependence is shown explicitly below, so matrices with no time dependence shown are constant. Because $\mathbf{F}$ is a constant, the state transition matrix Φ for propagating over one sample period is the same for all sample periods, and so it will be denoted as $\Phi(\Delta t)$ rather than $\Phi(t_i - t_{i-1})$. The state transition matrix for time periods other than one sample period will be shown with explicit time arguments. The discrete-time control input matrix $\mathbf{B}_d$ is found via:

$$\mathbf{B}_d = \int_{t_{i-1}}^{t_i} \Phi(t_i - \tau) \mathbf{B} d\tau \quad (4.28)$$

Since $\mathbf{Q}_d$ has as many rows and columns as the state vector, $\mathbf{G}_d$ can be chosen equal to an identity matrix and left out of the equation. The Kalman filter propagation equations are [77]:

$$\hat{x}(t_i^-) = \Phi(\Delta t) \hat{x}(t_{i-1}^+) + \mathbf{B}_d \bar{u}(t_{i-1}) \quad (4.29)$$

$$\mathbf{P}(t_i^-) = \Phi(\Delta t) \mathbf{P}(t_{i-1}^+) \Phi^T(\Delta t) + \mathbf{Q}_d \quad (4.30)$$

$$\mathbf{Q}_d = \int_{t_{i-1}}^{t_i} \Phi(t_i - \tau) \mathbf{G} \mathbf{Q} \mathbf{G}^T \Phi^T(t_i - \tau) d\tau \quad (4.31)$$

The Kalman filter update equations for the system, in which $\mathbf{H}$ and $\mathbf{R}$ are constant, are [77]:

$$\mathbf{A}(t_i) = \mathbf{H} \mathbf{P}(t_i^-) \mathbf{H}^T + \mathbf{R} \quad (4.32)$$

$$\mathbf{K}(t_i) = \mathbf{P}(t_i^-) \mathbf{H}^T [\mathbf{A}(t_i)]^{-1} \quad (4.33)$$

$$\hat{\mathbf{x}}(t_i^+) = \hat{\mathbf{x}}(t_i^-) + \mathbf{K}(t_i) [\bar{\mathbf{y}}(t_i) - \mathbf{H} \hat{\mathbf{x}}(t_i^-)] \quad (4.34)$$

$$\mathbf{P}(t_i^+) = [\mathbf{I} - \mathbf{K}(t_i) \mathbf{H}]^T \mathbf{P}(t_i^-) [\mathbf{I} - \mathbf{K}(t_i) \mathbf{H}] + \mathbf{K}(t_i) \mathbf{R} \mathbf{K}^T(t_i) \quad (4.35)$$

Collectively, these equations describe a Kalman filter which can read a noise-corrupted signal of the form shown in Equation (4.1), with the assumed frequencies, and reconstruct it optimally. While five states were shown here, a Kalman filter for any number of poles can be constructed by augmenting the matrices in this section. To account for additional poles, the state vector $\hat{\mathbf{x}}$ adds two states for each pole pair and one for each real axis pole. The $\mathbf{H}$ matrix grows by one column, and $\mathbf{B}$ grows by one row, for each state added to $\hat{\mathbf{x}}$. The various square matrices add both a row and a column for each state appended to $\hat{\mathbf{x}}$, but retain their sparse character. Every element is zero except for an element or a two-by-two block on the main diagonal. The filters developed in this research had from ten to twenty states. Most of these were associated with conjugate pairs of poles, but some models had one or two states associated with a real axis pole or non-oscillating decay term.

A consideration for a user of this algorithm is that matrices are built with knowledge of the sample period. If a different sample period is used, the matrices must be recomputed. A consequence of this is that filters with different assumed

sample periods would require additional manipulation before being included in the same filter bank. This is not a serious limitation. Once all filter matrices are compatible with the sample period of the measured data, there should be no impact on the online performance of the algorithm. Another approach is to resample a data set, via interpolation, to make its sample period compatible with the filter. To simplify the research, filters developed for targets measured or simulated with different incident bandwidths were not tested simultaneously in the target recognition algorithm. Note that it is the bandwidth *after* zero padding (explained in Appendix D) that determines the sample period via the IFFT. Slight changes in the measured or simulated bandwidth, that don't change the zero-padded bandwidth, will not change the sample period.

4.2.1 Filter Tuning. A Kalman filter must be tuned to work properly. The tuning process involves adjusting values within the dynamics noise covariance $\mathbf{Q}_d$ as well as the initial value of the state covariance matrix $\mathbf{P}$. The measurement noise covariance matrix, $\mathbf{R}$, is usually set equal to the covariance of the noise actually encountered on the measurements. An additional tuning consideration in this research was parameter units. As discussed in Section 4.2, the units of Ω , ω , and Δt were chosen to ensure a well-conditioned Φ matrix. For some filters the units of these quantities were changed to improve performance.

For the target recognition algorithm, the ultimate goal of tuning was not the best filter performance considered in isolation, but the filter's performance in the context of multiple model adaptive estimation. Good state estimation was desirable, but good target recognition was essential. Thus, when the filter-assumed parameter values were correct, excellent state estimation and appropriate residual properties should ensue. However, when those parameter values were incorrectly assumed, it would be desired for that to be reflected strongly in the residuals, rather than good state estimation performance being preserved *despite* such wrong assumptions. By so tuning the elemental filters, the various elemental filter models would be readily

distinguished from one another and the probability flows within the MMAE would yield desired levels of adaptability and target recognition. The filters were tuned to work properly, but some then required additional tuning to work successfully in the target recognition algorithm.

The signal model is a sum of independent mode terms. Taking that into account, the noise model assumes that the uncertainty on one mode doesn't affect the uncertainty on other modes. That leads to a block-diagonal form of the $\mathbf{Q}_d$ matrix. Within each two-state mode, the uncertainty of one state is related to the uncertainty of the other, so the non-zero blocks of $\mathbf{Q}_d$ do have terms off the main diagonal. In this research, tuning involved selecting the order of magnitude of the $\mathbf{Q}$ matrix, then manipulating the elements on the diagonal of $\mathbf{Q}$ to change the noise on individual modes. To review, the continuous-time dynamics noise strength $\mathbf{Q}$ is related to the discrete-time dynamics noise covariance $\mathbf{Q}_d$ via:

$$\mathbf{Q}_d = \int_{t_{i-1}}^{t_i} \Phi(t_i - \tau) \mathbf{G} \mathbf{Q} \mathbf{G}^T \Phi^T(t_i - \tau) d\tau \quad (4.36)$$

The noise model includes three parameters, p_m , q_m , and Q_m for each mode m . The mode is the state or pair of states representing a single term in the damped sinusoid signal model. There are noise parameters associated with damping (p_m), oscillating (q_m), and the mode overall (Q_m). It was desired to find an appropriate ratio of p_m and q_m that would work for most modes of most filters, requiring adjustment only in a few cases. Then tuning would reduce to adjusting the overall noise on each mode.

For each mode, p_m was set equal to the MSE of Ω_m (as estimated by M-TLS-MPM), normalized by $|\Omega_m|$. Similarly, q_m was set to the MSE of the M-TLS-MPM estimate of ω_m , divided by ω_m . These settings were chosen so that uncertainty in the estimate would drive uncertainty in the model in a simple, straightforward way.

White, Gaussian, pseudo-random noise was added to the signal $\bar{y}(t_i)$ input to the filter. The variance of this noise was set so that a desired SNR was achieved. It

was assumed that some noise was present in the simulated or measured signal prior to the addition of white noise, but the added noise had a larger variance than any previously-existing noise. Previously-existing noise would be the result of any minor errors in the scattering simulation, propagated into the time domain via the IFFT. Such noise would be band-limited, rather than white, and would not necessarily be well-modeled as Gaussian. The signal-to-noise ratio was computed with regard to the added white, Gaussian noise only. The single element of the constant 1-by-1 $\mathbf{R}$ matrix was set equal to the variance of the added noise.

In the initial plan for this research, it was assumed that a parameter estimation routine would be necessary within each filter to estimate the unknown coefficients A_m and B_m associated with each damped sinusoid term in the signal model. The planned parameter estimation technique was the Rao Scoring approximation to the Newton-Raphson solution of the maximum likelihood equations [76]. The tuning strategy used in the algorithm allowed the states to adapt to the signal, so parameter estimation was found to be unnecessary, and a significant amount of additional computation was avoided.

To explain the tuning strategy, consider first a simplified explanation of the measurement update process of a Kalman filter. The state update equation, Equation (2.72), is cast in a different form to show clearly that it is a weighted average of the previous states and the measurements:

$$\hat{\hat{x}}(t_i^+) = [\mathbf{I} - \mathbf{K}(t_i)\mathbf{H}(t_i)]\hat{\hat{x}}(t_i^-) + \mathbf{K}(t_i)\bar{y}(t_i) \quad (4.37)$$

At measurement update, the filter updates the estimated state vector with a weighted average of the state immediately prior to update, and the current measurement. The weight in the equation is the Kalman gain $\mathbf{K}(t_i)$. If the values in the gain matrix $\mathbf{K}(t_i)$ are small, the measurement has little effect on the updated value. If the Kalman gain is larger, the residual is weighed more heavily and the state is changed significantly after update.

The goal of the tuning strategy was to force the filter to weigh the measurement heavily at first, so for a brief initial time the measurement would drive the value of $\hat{x}(t_i^+)$ regardless of the system's dynamics. The filter would 'lock on' to the incoming measurement quickly. Because of this tuning choice, the initial residuals were tiny, preventing the MMAE equation from distinguishing between filters. After the initial period, the gain decreased significantly, allowing the system dynamics to influence the filter more heavily, so that filters with incorrect system dynamics models would be unable to adapt to the incoming measurements, generating larger residuals.

To achieve the appropriate $\mathbf{K}(t_i)$ values over the time window, two techniques were considered. One of them was the use of time-varying covariance matrices $\mathbf{Q}_d$ and $\mathbf{R}$. This technique was tested, but not used in the final version of the algorithm because a less complicated approach was available.

A simpler way to achieve the same effect was to use a large initial value for the state covariance $\mathbf{P}(t_0)$, along with smaller values of $\mathbf{Q}_d$ than called for by conservative filter tuning. Conservative tuning refers here to setting the value of the dynamics noise strength $\mathbf{Q}$ or covariance $\mathbf{Q}_d$ high enough to force a filter to weigh incoming measurements heavily, thus making it responsive to changes in system behavior. For effective target recognition, it was desired that filters with models that did not match the incoming signal be prevented from adapting to that signal. Over time the state covariance decreased toward a steady state value. The measurement and system noise covariances $\mathbf{R}$ and $\mathbf{Q}_d$ were both set to constant values. Each filter adjusted the values in the state covariance matrix downward from their initial values toward lower, steady-state values. The Kalman gain exhibited the desired behavior with the least required computation of the tuning methodologies considered.

When tuning a filter, one criterion that was applied was that the individual output states should appear to be noise-free in the later part of the time window, after the initial transient. The output states were those states corresponding to nonzero elements of the $\mathbf{H}$ matrix. Once the filter settled down, it needed to be able

to reject signals that didn't match its dynamics model. If an output state appeared to be affected by noise after the filter's initial transient period, it was too responsive to the measurement (i.e., the gain was too high) to reject incorrect dynamics in a signal adequately. The value of Q_m (a parameter on the main diagonal of the $\mathbf{Q}$ matrix) was reduced for the indicated mode until the output state appeared clean.

It was expected that, in a properly-tuned filter for which the model matched the signal, the size of the quantity $\bar{r}^T(t_i)\mathbf{A}^{-1}(t_i)\bar{r}(t_i)$ (the square of the scaled residuals) would be on the order of the number of elements in the measurement vector. In this case, with scalar measurements, the squared scaled residuals will be on the order of one. That is because $\mathbf{A}$ is the residual covariance (variance in this scalar-measurement case), so residuals will be on the order of the standard deviation, $\sqrt{\mathbf{A}}$. When the inverse of the standard deviation is used to scale the residuals, the result should be on the order of one.

An additional parameter that was adjusted as part of filter tuning was the set of units of the frequencies and the sample period. Some filters, typically those with fewer than five states, ran better with frequency units of Grad/s and GN/s, and the sample period in ns, but most operated well with units of tens of Grad/s, tens of GN/s, and tenths of ns. For the synthetic data, units of hundreds of Grad/s, hundreds of GN/s, and hundredths of ns were used. Changing the units affected the magnitude of elements of the system dynamics matrix $\mathbf{F}$, which in turn affected the state transition matrix $\Phi(\Delta t)$ computed from $\mathbf{F}$. Changes in $\Phi(\Delta t)$ affected the rate of change in the state covariance $\mathbf{P}(t_i)$ via Equation (4.30). Since $\mathbf{P}(t_i)$ was used to compute the Kalman gain $\mathbf{K}(t_i)$ as seen in Equations (4.32) and (4.33), changes in units affected the way each filter performed over time.

4.2.2 Filter Tests with Synthetic Data. Some details of the filter's operation are explored in this section. In each case, data is shown from a single simulation run as well as overall results from a set of 100 runs. Including both types of plots allows

a clear demonstration of the filter's operation in the presence of noise as well as an indication that the results obtained were typical, rather than exceptional. Three experiments were conducted to demonstrate filter effectiveness. The first involved the filter's response to a signal that fits the filter's system model. The second experiment involves a perturbed signal which still fits the filter's system model, but has a different amplitude on one signal mode. The third experiment shows how the filter responds when the signal does not match the filter's system model.

Figure 4.2 shows a noisy signal and the filter's reconstruction of it. The filter's Ω_m and ω_m parameters were estimated using M-TLS-MPM and Synthetic Data Version 1 with 13 dB SNR. The signal input to the filter had an SNR of 3 dB. The filter model's parameters adequately reflected the parameters of the underlying signal, so the filter performed well despite the presence of significant noise. Figure 4.3 shows the mean, plus or minus one standard deviation, of 100 runs of the signal plus additive, white, Gaussian noise. The figure also shows the mean, plus or minus one standard deviation, of the filter's response. Except for an initial transient period, it can be seen in Figure 4.3 that the mean of the filter's reconstruction corresponded approximately to the mean of the 100 noisy signals. The mean of the reconstruction also corresponded approximately to the uncorrupted signal. Means are not shown in the picture for clarity, but the means are sufficiently discernable from the standard deviation bounds.

As desired, the scaled residuals for this filter appeared to be zero-mean white noise, with an amplitude on the order of one. The scaled residuals are shown in Figure 4.4. The mean, plus or minus one standard deviation, of the scaled residuals is shown in Figure 4.5. Here it can be seen clearly that the residuals were correctly scaled after an initial transient period. In other words, the filter's calculation of the residual variance $\mathbf{A}(t_i)$ was accurate.

The filter has only a scalar measurement, but must use that limited information to reconstruct the various modes that make up the data. As seen in Figures 4.6




Figure 4.2 Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\hat{x}}$




Figure 4.3 Mean and standard deviation of Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\hat{x}}$




Figure 4.4 Scaled residuals for Synthetic Data Version 1




Figure 4.5 Mean and standard deviation of scaled residuals for Synthetic Data Version 1

through 4.13, the filter's output states matched the modes that made up the noise-free signal. The output states are those elements of the state estimate vector $\hat{\mathbf{x}}(t_i)$ corresponding to a non-zero element in the measurement matrix $\mathbf{H}$. These are states $\hat{x}_1(t_i)$, $\hat{x}_2(t_i)$, $\hat{x}_4(t_i)$, and $\hat{x}_6(t_i)$, corresponding to the non-oscillating mode (pole on the real axis of the s -plane) and three damped sinusoid modes. With poor tuning, the filter had some difficulty reconstructing the non-oscillating mode, and reconstructed the three oscillating modes with some phase error in each output state.

Figures 4.6 and 4.7 show the filter's response to the non-oscillating mode in Synthetic Data version 1. After the initial transient, the filter's mode settles at a point higher than the actual non-oscillating mode, and declines exponentially from there. Note that the comparison in Figure 4.7 is to a portion of the Version 1 data which has not been corrupted by noise; the data itself is presented, rather than the mean and deviation of many runs.

Figures 4.8 and 4.9 show the filter's response to the lowest-frequency oscillating mode in the received signal. The filter reconstructs this mode well, although there is a slightly smaller amplitude and the phase of the reconstructed mode leads the actual mode.

Figures 4.10 and 4.11 show the filter's response to the middle-frequency oscillating mode in the received signal. The filter reconstruction follows the signal mode well, with only a slightly smaller amplitude and a slight phase delay.

Figures 4.12 and 4.13 show the filter's response to the highest-frequency oscillating mode in the received signal. The filter did a good job reconstructing this mode, although there was a smaller amplitude and a slight phase delay. The slight reduction in amplitude of the three oscillating modes, combined with the slightly higher amplitude of the non-oscillating mode, allowed the filter to reconstruct the total signal properly.




Figure 4.6 Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode




Figure 4.7 Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode




Figure 4.8 Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode




Figure 4.9 Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode




Figure 4.10 Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode




Figure 4.11 Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode




Figure 4.12 Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode




Figure 4.13 Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode

In the second experiment, the filter's response to a modified signal was considered. The coefficients associated with each mode change along with the aspect from which the target is viewed. It is therefore useful to consider the filter's reaction to a weak mode. The same Synthetic Data Version 1 used previously was modified; the coefficient of the third mode (estimated by states $\hat{x}_4$ and $\hat{x}_5$, output state $\hat{x}_4$) multiplied by 0.01 to suppress it without actually eliminating the mode. The results are shown in Figures 4.14 through 4.25. In this experiment, the filter reacted properly to the near-absence of the weak mode. The filter was able to reconstruct the three non-suppressed modes adequately.

Figures 4.14 and 4.15 show the filter's reconstruction of the signal with the third mode (middle oscillating mode) suppressed. The signal input to the filter had an SNR of 3 dB, as in the previous experiment. The frequency and damping parameters of the signal and filter are the same as in the previous experiment.

As was seen in the previous experiment, the scaled residuals for this filter appeared to be zero-mean white noise, with an amplitude on the order of one. The scaled residuals are shown in Figure 4.16. The mean, plus or minus one standard deviation, of the scaled residuals is shown in Figure 4.17. The figure shows that the scaled residuals had a standard deviation of approximately one, as desired and as seen in the first experiment.

Figures 4.18 and 4.19 show the filter's response to the non-oscillating mode in Synthetic Data version 1. After the initial transient, the filter's mode approaches the signal's non-oscillating mode closely, and declines exponentially from there at the appropriate rate, following the signal's mode. This mode responds better in this experiment than it did in the previous one.

Figures 4.20 and 4.21 show the filter's response to the lowest-frequency oscillating mode in the received signal. The filter reconstructs this mode well, although there is a slightly smaller amplitude and the phase of the reconstructed mode leads the actual mode.




Figure 4.14 Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, third mode suppressed




Figure 4.15 Mean and standard deviation of Synthetic Data Version 1, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, third mode suppressed




Figure 4.16 Scaled residuals for Synthetic Data Version 1, 3 dB SNR, third mode suppressed




Figure 4.17 Mean and standard deviation of Scaled residuals for Synthetic Data Version 1, 3 dB SNR, third mode suppressed




Figure 4.18 Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed




Figure 4.19 Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed




Figure 4.20 Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed




Figure 4.21 Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed

Figures 4.22 and 4.23 show the filter's response to the suppressed middle-frequency oscillating mode. After the very large initial transient, this filter mode assumes a small amplitude to match the signal mode. Because the signal mode's amplitude is so small, it is well below a single standard deviation of the filter's mode.

Figures 4.24 and 4.25 show the filter's response to the highest-frequency oscillating mode in the received signal. The filter's mode has the same amplitude as the signal's mode, and there is only a slight phase delay.

In Figures 4.2 through 4.25, it can be seen that a properly tuned filter with the right signal parameters can estimate a signal very well. Another key consideration for target recognition is how well a filter reacts to the wrong signal, i.e., a signal not matching the filter's assumed parameters. The desire is that the filter will react poorly. In the third experiment, we consider a filter's response to the wrong signal.

Figures 4.28 through 4.37 show results from a filter reacting to a signal it was not designed to match. The filter was based on estimated poles from Synthetic Data Version 1 at 13 dB SNR. The signal received was Version 2. As in the first two experiments, the SNR of the signal was 3 dB. The signal and the filter were each based on one non-oscillating and three damped sinusoid modes, but the frequencies were different. The frequencies for Versions 1 and 2 are given in Table 3.1.

Figures 4.26 and 4.27 show the filter's reconstruction of the signal. In Figure 4.26 some mismatch between the dashed line (the filter) and the noisy solid line (the signal) can be seen in the left side of the graph, between 0.2 and 0.6 nanoseconds. Figure 4.27 shows somewhat more clearly that the mean of the filter's reconstruction does not follow the mean of the signal.

It is expected that the scaled residual plot will stay mostly between the amplitude limits of 1 and -1, if the underlying model's assumptions were correct. The scaled residuals from the Version 1 filter reacting to the Version 2 signal are shown




Figure 4.22 Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode, third mode (this mode) suppressed




Figure 4.23 Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 1, and noise-free mode, third mode (this mode) suppressed




Figure 4.24 Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed




Figure 4.25 Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 1, and noise-free mode, third mode suppressed




Figure 4.26 Synthetic Data Version 2, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, filter based on Version 1 frequencies




Figure 4.27 Mean and standard deviation of Synthetic Data Version 2, 3 dB SNR, and Kalman Filter reconstruction $\mathbf{H}\hat{\mathbf{x}}$, filter based on Version 1 frequencies

in Figure 4.28. The residuals in this case are a bit larger than unity, and some dynamics are apparent in the left side of the graph, between 0.1 and 0.8 nanoseconds. In a well-matched filter, the scaled residuals should take on the appearance of white Gaussian noise because most or all traces of the signal have been subtracted out.

The mean of the residuals, plus or minus one standard deviation, is shown in Figure 4.29. Unlike the previous two experiments, the scaled residuals for this filter contained significant dynamics, indicating that the filter did not account for the signal's dynamics. After one nanosecond, the residuals sustain a standard deviation of approximately one, but that is only due to the fading signal strength. After one nanosecond, there is not much left in the noise-corrupted signal except noise.

Here it is seen that a Kalman filter that is matched to a signal different from the received signal will generate scaled residuals that are large and non-white (non-constant power spectral density), where it was seen earlier that a filter matched to the received signal will produce residuals that are white. The ability to distinguish filters that match the signal poorly from those that match the signal well is the basis of the target recognition process.

Figures 4.30 through 4.37 show the performance of individual filter modes compared to the modes that made up the noise-free signal. In each figure, the filter's modes had significant error relative to the modes in the underlying signal. Both phase and amplitude errors are pronounced in each oscillating case.

Figures 4.30 and 4.31 show the filter's response to the non-oscillating mode in Synthetic Data version 1. This mode experiences a longer initial transient than was present in earlier experiments. After the initial transient, the filter's mode is too high in amplitude and declines at a different rate from the signal's mode.

Figures 4.32 and 4.33 show the filter's response to the lowest-frequency oscillating mode in the received signal. The filter's mode never matches the signal's mode. The filter has the wrong amplitude, the wrong frequency, and the wrong phase.




Figure 4.28 Scaled residuals for Synthetic Data Version 2, 3 dB SNR, filter based on Version 1 frequencies




Figure 4.29 Mean and standard deviation of Scaled residuals for Synthetic Data Version 2, 3 dB SNR, filter based on Version 1 frequencies




Figure 4.30 Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.31 Mean and standard deviation of Kalman Filter state $\hat{x}_1$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.32 Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.33 Mean and standard deviation of Kalman Filter state $\hat{x}_2$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.34 Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.35 Mean and standard deviation of Kalman Filter state $\hat{x}_4$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.36 Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies




Figure 4.37 Mean and standard deviation of Kalman Filter state $\hat{x}_6$ for Synthetic Data Version 2, and noise-free mode, filter based on Version 1 frequencies

Figures 4.34 and 4.35 show the filter's response to the middle-frequency oscillating mode. The filter's mode has an amplitude similar to the signal's mode, but the frequency and phase are off, so the match is poor.

Figures 4.36 and 4.37 show the filter's response to the highest-frequency oscillating mode in the received signal. The filter's mode has a much smaller amplitude than the signal's mode, and the frequency and phase are off. The filter's mode is unable to match the signal's mode.

Kalman filters built with frequencies estimated from synthetic data are capable of estimating the matched signal in the presence of noise. They are also capable of rejecting an unmatched signal, resulting in large scaled residuals. Thus, properly tuned Kalman filters are useful building blocks in a target recognition algorithm based on an MMAE architecture.

4.3 *Applying Multiple Model Adaptive Estimation*

The Multiple Model Adaptive Estimation (MMAE) algorithm is the key to the target recognition algorithm. MMAE is frequently used to estimate the value of one or more unknown parameters in a linear system model. Here MMAE is not used to estimate individual parameters, but instead to select which set of fixed, known parameters best fits the target, given a library of sets of parameters. The maximum *a posteriori* (MAP) version of MMAE [76] is implemented, which selects the system model within a single filter as the one most representative of the system being modeled. That is in contrast to the standard or Bayesian blended version of MMAE, which generates a single state estimate based on a probability-weighted average of the state estimates generated by each elemental filter.

The hypothesis conditional probability (HCP) computation routine within the MMAE process determines which model does the best job of estimating the values of the true state variables. The filter with the highest HCP is the one considered closest to the true system in the sense of the prediction distance defined by Baram [5]. In

fact, each filter's state estimate $\hat{x}(t_i)$ is only used within the filter that generates it. The target recognition decision uses the HCP $p(t_i)$ of each filter, but does not directly use the associated state estimate.

Filter covariances were tuned for best target recognition performance, rather than for the best state estimation. This involved modifying dynamics driving noise covariance in the filters to improve parameter estimation [89]. It also involved changing the size of initial values in the covariance matrix.

As seen in Equations (2.75) through (2.78), the HCP at each sample point was weighted by the HCP of the previous sample point. If the HCP of any one filter ever was reduced to zero, it would be weighted by zero on each succeeding update, and the filter could never be declared correct. It was possible for the probability of any single filter, even one with the correct dynamics model, to go to zero early in the algorithm's operation.

To force the MMAE system to be more responsive to new inputs, a minimum probability of $L = 0.01$ was imposed on each filter [76, 79]. Mathematically, since the filters represent mutually-exclusive possibilities (one filter is assumed correct and all others must logically be incorrect), the probabilities associated with all of the filters must sum to one. After the lower bound was imposed, the HCP of each filter was rescaled by dividing by the sum of HCP values, to ensure that all the HCP values would sum to one. Increasing L forced the algorithm to be more sensitive over the entire time window. Setting $L = 0$ would result in an unmodified algorithm. Minimum probabilities are usually used in MMAE systems to keep the system alert to changing parameter values. In this problem, parameter values did not change, but it was desirable to prevent the algorithm from excluding information from any one filter prior to each filter reaching steady state operation. A lower bound of 0.001 has been used in some failure detection and identification systems [79, 84]. These were flight control systems which would operate over several minutes of flight time, with excitation provided by control inputs during the period of operation. In this

research, no control inputs were used, and the late-time signal decayed rapidly. Thus, a higher value ($L = 0.01$) was used for the lower bound. It was desired to allow each filter to recover more rapidly from a very low probability that might be assigned prior to filter convergence to a good parameter estimate.

To evaluate the noise resistance of the target recognition algorithm, target signatures were developed to represent the Synthetic Data Versions 1 through 4. Several test data sets were created by adding Gaussian noise to the clean Synthetic Data. The noisy data was analyzed using the M-TLS-MPM described in Section 3.2. Four iterations of the LRHA were used because four iterations produced a reduction in the MSE of parameter estimates compared to the three-iteration case in Section 3.2.1, while the additional reduction in MSE for a fifth iteration was insignificant. The power of the added noise was set so that the signal-to-noise ratio (SNR) varied from -15 to 25 dB. The significance level P was set to its prescribed value, $P = \text{SNR}/20$, rounded up to the next tenth. There were 100 noisy samples used at each SNR level.

Target filters were built for all four versions of Synthetic Data described in Section 3.3.1. The parameters Ω_m and ω_m used to build the filters were estimated from noisy synthetic data at various SNR settings, using the M-TLS-MPM algorithm with four iterations of the LRHA. The filters were built using the parameter estimates and MSE results found using signals with 25 dB SNR and the same signals with 13 dB SNR. The filters based on 25 dB SNR data were tuned differently from the filters for which the pole estimates were based on 13 dB data. The results, in terms of percent correct identification, are shown in Figures 4.38 and 4.39.

Figure 4.38 shows the results found with filters developed from 25 dB SNR data. The effectiveness of the algorithm varied as different input signals were used. For Version 2, the correct answer was found in 90% of cases when the SNR was -7 dB or higher, while for Version 3, the correct answer was found in 90% of attempts when the SNR was -15 dB or above.




Figure 4.38 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 25 dB data




Figure 4.39 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data

Figure 4.39 shows the results found with filters developed from 13 dB SNR data. The results were similar to the 25 dB case in that the percent of correct identifications was high when the SNR was above approximately -7 dB. For Version 1, the correct answer was found in 90% of attempts at all SNR levels examined.

For synthetic data, the algorithm is capable of associating the correct filter with the signal, even when the SNR of the signal is -7 dB. In Section 4.4, the technique is evaluated with simulated radar data.

4.3.1 Selecting a Target. To distinguish between filter models, the algorithm needs a single ‘goodness of fit’ indication from each one. All the data available from the filter must be distilled down to a single measure of effectiveness which supports a decision. The HCP of the correct filter is such a criterion. However, as the signal fades, the algorithm may be less able to distinguish between filters. As the signal strength fades, the term $\mathbf{H}_k \bar{x}_k(t_i)$ in the filter residual equation becomes smaller in magnitude, and the residual in each filter becomes nearly equal to the measurement. Each filter will then have the same residual. Since the term $\bar{r}_k^T \mathbf{A}_k^{-1} \bar{r}_k$ for filter k is the same as for all other filters, the HCP of each filter will be determined primarily by the coefficient β_k , first shown in Equation (2.76) and defined in Equation (2.77). This problem is known as ‘beta dominance’ [84, 136]. While the signal is still strong, the HCP will be determined by the square of the scaled residuals, $\bar{r}_k^T \mathbf{A}_k^{-1} \bar{r}_k$.

To solve that problem, at least three approaches are possible: changing the equation for the HCP, forming a weighted average of the HCP over time, and terminating the MMAE-MAP algorithm when each filter’s residual covariance has reached its steady-state value. The equation for the conditional probability density may be modified in an *ad hoc* way that reduces its dependence on the residual covariance $\mathbf{A}(t_i)$. The leading coefficient β_k shown in Equation (2.76) may be set equal to one [79, 84] so that the computation is less sensitive to variations in residual covari-

ance from filter to filter. The modified conditional density function will take this form:

$$f_{\bar{y}(t_i)|\bar{a},\bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_k, \bar{Y}_{i-1}) = \exp\{-\frac{1}{2}\bar{r}_k^T(t_i)\mathbf{A}_k^{-1}(t_i)\bar{r}_k(t_i)\} \quad (4.38)$$

The conditional density function, when modified in this way, no longer meets the mathematical definition of a probability density in that its integral is no longer equal to one. However, the sum of all filters' HCP values is one. This approach was found to improve the algorithm's performance in some cases discussed below. There was no additional computational burden.

Results of the modified algorithm in which β_k is set equal to one for each filter k are shown in Figure 4.40. This figure corresponds to the previous Figure 4.39, which shows the results of the unmodified algorithm. The modified algorithm's performance is slightly worse than that of the unmodified algorithm. In the unmodified case, each target is correctly identified in at least 90% of attempts when the SNR is greater than -5 dB. The modified algorithm requires -2 dB SNR to achieve the same results for all four signals.

Another approach that was considered was to form a weighted time-average of the HCP of each filter. For an appropriately-selected weight function, the period of time during which the algorithm had the best opportunity to perform properly could be emphasized. This technique provided accurate target recognition for a small increase in required computations [102, 103].

A third approach, which was chosen for the final version of the algorithm, was to terminate the target recognition algorithm once it was most likely to have reached a successful conclusion, yet before the signal strength faded. To determine the appropriate stopping criterion, the filter covariances were considered. Given the tuning employed, each filter's residual covariance started at a high value and quickly fell toward a steady state. At the time each filter reached steady state, the dynamics within each incorrect filter were forcing increases in the residuals, while the




Figure 4.40 MMAE-MAP with $\beta_k = 1$, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data

correct filter's residuals remained small. That allowed the MMAE-MAP algorithm to distinguish between filters with the best chance of a successful decision.

At each update, the target recognition algorithm compared the residual covariance of each filter to its value at the previous update. A ratio of the current to the previous covariance was formed. If the ratio for each filter was nearly one within a tolerance, and remained that way for a number of turns, the filter was considered to have reached steady state. The tolerance was 0.001, so residual covariance ratios between 0.999 and 1.001 indicated filter convergence if they persisted in this range for ten sample periods. When all filters were determined to be at steady state, the algorithm was terminated and a decision was declared.

This technique added some computation to each measurement update of the MMAE-MAP algorithm. It was intended to reduce computations overall by terminating the algorithm before the end of the data set was reached. In fact, it frequently did not save any computations. Because filters achieved steady state operation at

different times, the best possible result was to terminate after the longest-settling filter had settled. If that was within ten sample periods of the end of the data window, no savings were realized.

The reduction in computation due to the early-stopping technique depends on the length (number of sample periods) of the signal to be tested, as well as the tolerance used to declare filter convergence. It also depends on the number of sample periods the algorithm is continued after filter convergence is achieved. Relaxing the tolerance or reducing the number of sample periods after convergence results in a faster decision and computational savings, but may also result in declaring a decision before the algorithm has been driven to the right answer by differences in the residuals or the elemental filters. Future research will identify optimal settings for the tolerance and number of samples after steady-state to allow the best target identification performance, with a secondary objective of reducing computations slightly.

In this research, data sets of fixed length were input to the algorithm. In field operation, the length of the input signal will be determined by operational parameters including the distance separating the target of interest from other targets, and the SNR. The early-stopping technique will reduce computation in those cases in which long signals are available.

4.3.2 Confusion Matrices. When talking about sensors, engineers frequently use terms such as probability of detection and probability of false alarm. Since target detection is assumed in this research, these probabilities are not used. 100% detections and 0% false alarms were assumed before the target recognition algorithm was brought into use. A different measure of effectiveness is needed, one appropriate to a target recognition system rather than a target detection system. An example is a confusion matrix, sometimes termed a discrimination event matrix [2].

A confusion matrix is a way to display the results of several runs of a target recognition system. The matrix shows the fraction of attempts in which a recognition system (or ‘classifier’ [2]) declared a correct or incorrect target designation, based on the actual target identity. Ideally, the values in the main diagonal of the matrix are all equal to one, and values off the main diagonal are zero. Thus, a perfect confusion matrix (showing no confusion) is an identity matrix. A confusion matrix is a valuable tool to see where a classifier succeeds and fails. It takes into account much more information than the percentage of correct identifications, and can show biases toward or away from any one target type [2].

To fill in the elements of the matrix, several experiments were performed using late-time signals from the targets indicated in the left column. Each experiment resulted in one target type being declared. To produce this table, the final HCP of each filter was not retained. The only data used was which filter received the highest HCP at the time of target declaration. Confusion matrices ordinarily take into account either numbers or averages of discrete events such as target declarations [2]. A possible variant of a confusion matrix could show the time-average of the HCP for each target over an interval following target declaration. Such a matrix would require different analysis than a standard confusion matrix, since each target type would have a non-zero HCP, indicating confusion even when the correct target was declared. To simplify analysis, standard confusion matrices were used. Variants of the confusion matrix may be useful tools for future research into target recognition.

An example matrix for Synthetic Data is shown in Table 4.3. It is based on the same data used to generate Figure 4.39. The results were generated with the SNR of the measured signal set to -9 dB. Above -9 dB, results were so good that a confusion matrix would have little diagnostic value. From +11 to +15 dB SNR, the confusion matrix was an identity matrix.

In this case, when erroneous declarations were made, the filter for the Version 1 signal tended to be declared correct. It is possible that additional tuning could result

Correct Target	Declared Version 1	Declared Version 2	Declared Version 3	Declared Version 4
Version 1	0.96	0.03	0.01	0.00
Version 2	0.23	0.69	0.08	0.00
Version 3	0.14	0.01	0.84	0.01
Version 4	0.17	0.02	0.03	0.78

Table 4.3 Confusion matrix, synthetic data, SNR = -9 dB

in less bias toward the Version 1 filter at lower SNR, but different tuning also has the potential to reduce the effectiveness of the algorithm at higher SNR, as will be discussed in Section 4.4. If this kind of confusion occurred at a higher SNR level, it would be viewed as a problem and a correction would be attempted through tuning. In this case, the algorithm achieved good results for -7 dB SNR, much better than the E-Pulse technique using the same data, as discussed in Section 4.5.

A possible source of the bias was beta dominance [79], discussed previously. For this data, the same test was run with the HCP equation modified as shown in Equation (4.38). The results were identical to those found with the algorithm that computed HCP via Equation (2.76). This indicated that the source of the bias was not beta dominance.

A confusion matrix is very useful for diagnosing failures when they are significant. Examples of confusion matrices are shown in Section 4.4 for tests involving simulated radar data.

4.3.3 Verifying Target Identification via Chi-Squared Test. Since the goal of this research was accurate target recognition, it was useful to consider measures for how well the recognition performed. In operational use, the ‘ground truth’ or actual nature of the target is unknown. In that case, some measure of confidence is desired, which will allow the user of this algorithm to know how much trust to place in the result.

The χ^2 -test involves comparing a test statistic to a threshold. If the test statistic is less than the specified threshold, the hypothesis is shown to be supported by the data, and accepted as true [101]. The test statistic q_k for filter k is developed using the current residual and the previous $N - 1$ residuals. It is given by:

$$q_k = \sum_{j=i-N+1}^i \bar{r}_k^T(t_j) \mathbf{A}_k^{-1}(t_j) \bar{r}_k(t_j) \quad (4.39)$$

in which N is the length of the data sequence used to test for confidence. This variable has a χ^2 distribution with N degrees of freedom [44]. The probability distribution function of a χ^2 variable is given by [104]:

$$\chi^2(N) = f(x) = \frac{x^{(N/2)-1}}{2^{N/2}\Gamma(N/2)} \exp(-x/2) \quad (4.40)$$

Here the test was used on the residuals from the best matching filter, to verify that the residuals were actually well-described by the probability density function (PDF) they were assumed to fit. Individual filter scaled residual measurements were assumed to be normally distributed. Each filter's PDF, used in the HCP computation, was chosen to be normal. To review, the HCP was computed by:

$$p_k(t_i) = \frac{f_{\bar{y}(t_i)|\bar{a}, \bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_k, \bar{Y}_{i-1}) p_k(t_{i-1})}{\sum_{j=1}^K f_{\bar{y}(t_i)|\bar{a}, \bar{Y}(t_{i-1})}(\bar{y}_i | \bar{a}_j, \bar{Y}_{i-1}) p_j(t_{i-1})} \quad (4.41)$$

The significance of the test was set to 0.05 and the residuals from the last 15 sample periods were used. Given these settings, the probability was 0.95 that the sum of squared scaled residuals computed via Equation (4.39) would be equal to or less than 24.9958. In practice, for a significance of 0.05, the χ^2 test should indicate a failure in approximately 5% of trials when the correct filter model is declared. Thus, a single failure of the χ^2 test should not be interpreted as a definitive indication that the best-matching filter was a poor match to the target. Instead, a failure should

indicate that the target recognition attempt was inconclusive and the test should be applied again.

When the target model accurately depicts the signal's behavior, the residuals are white, Gaussian, zero-mean variables with covariance equal to the filter-computed covariance $\mathbf{A}(t_i)$ [77]. When the signal strength is low, such as during the later portion of the time window when SNR is low, the residuals of a filter may be dominated by white, Gaussian noise whether or not the filter is based on the correct model of the signal. In either case, the sum q_k described in Equation (4.39) may have been well-described by a χ^2 distribution. Thus, χ^2 failures occurred approximately 5% of the time for low-SNR trials, regardless of which filter was declared. That result indicates that the value of the χ^2 test diminishes as SNR is reduced.

Results of the χ^2 test were recorded. For the synthetic data tests discussed previously, the rate of failures ranged from 0% to 7% for all tested combinations of SNR and correct target, whether or not the correct target was declared.

Figure 4.41 shows the results of the χ^2 test in terms of the number of failures when the indicated target was the correct answer. The declared filter's residuals failed the test in approximately 5% of attempts. There was no obvious trend toward more or fewer failures as SNR varied. Results involving simulated radar data are discussed in Section 4.4.

An alternative approach to the χ^2 test involved distance in parameter space. To implement it would require a measure of the distance in parameter space from one model to another. The parameter space discussed here is an N -dimensional vector space in which each pole is a pair of parameters occupying its own dimension. Unfortunately, each model k may have a different number N_k of poles, leading to incomparably different parameter spaces. In addition, if two models have the same number of poles, there is no rationale to say that the first pole pair in the signature of Target A occupies the same dimension of parameter space as the first pole pair of Target B. Thus, a parameter-space approach was not implemented.




Figure 4.41 Percent of failures of χ^2 test, filters based on synthetic data with 13dB SNR

4.4 Target Recognition Using Radar Data

In target detection and recognition studies, it is useful to separate training data from test data. The target recognition system was built using a set of parameters estimated from various measurements via M-TLS-MPM. The signals used to test the target recognition algorithm were taken from the same data sets used to estimate the target poles. However, they were not used in the pole combination algorithm described in Section 3.4 that determined the signatures to be used in the filter. The test signals thus had no effect on the determination of the frequencies used in each filter.

The algorithm was tested with four component filters, each based on the poles estimated from simulated radar scattering data. Figures 4.42 to 4.45 show the output of the component filters in the target recognition algorithm. The solid line is the signal received by the filter, while the dashed line is the reconstructed signal $\mathbf{H}\hat{\hat{x}}(t_i)$.




Figure 4.42 Output of Filter 1 with signal model for simulated 1" $\times$ 1" $\times$ 2" block target; measurement (solid), reconstruction (dashed) and 1- σ bounds (dash-dot)




Figure 4.43 Output of Filter 2 with signal model for simulated 1" $\times$ 2" $\times$ 3" block target; measurement (solid), reconstruction (dashed) and 1- σ bounds (dash-dot)




Figure 4.44 Output of Filter 3 with signal model for simulated $1.1'' \times 1.9'' \times 2.8''$ block target; measurement (solid), reconstruction (dashed) and $1-\sigma$ bounds (dash-dot)




Figure 4.45 Output of Filter 4 with signal model for simulated $4.5''$ square plate target; measurement (solid), reconstruction (dashed) and $1-\sigma$ bounds (dash-dot)

The dotted lines above and below the plot are the one-standard-deviation bounds around $\mathbf{H}\hat{\mathbf{x}}(t_i)$. The bounds are computed as

$$\mathbf{H}\hat{\mathbf{x}}(t_i) \pm \sqrt[3]{\mathbf{A}(t_i)} \quad (4.42)$$

where $\mathbf{A}(t_i)$ is the residual covariance computed by each filter.

Filter 1 was the correct answer, meaning that the system model from which the filter was developed was based on the poles estimated from the same target that produced the received signal. The signal was scattered data from a rectangular block of dimensions 1" by 1" by 2", known as Block 112. It was simulated using horizontal polarization with the long dimension of the block oriented 42 degrees away from the radar. The target is described further in Section B.1. The time window shown in the plot includes only late-time scattering.

The other filters were based on the poles estimated from simulations of other block or plate targets. Each one was oriented so that the scattered field's direction of propagation was in the $\tilde{z} - \tilde{x}$ plane, oriented 42 degrees away from the $\tilde{z}$ axis, with ϕ polarization. The targets and their orientations are described further in Section C.2. The target for Filter 2, known as Block 123, had dimensions 1" by 2" by 3". The target for Filter 3, known as Block 123B, had dimensions only slightly different from that of Filter 2, 1.1" by 1.9" by 2.8". The target for Filter 4 was a square plate with 4.5" sides and 0.1" thickness.

The HCP for the correct filter was initially low. Figure 4.46 shows a plot of the HCP of each filter over time. The values for the four filters are indicated by different line styles: solid, dashed, dotted, and dash-dot. Because the initial values of the filter covariance matrices $\mathbf{P}$ were set to a high value, each filter weighed the measurement more heavily than its own system model. Residuals were low and the residual covariance matrix $\mathbf{A}$ was initially large, so the scaled residuals $\sqrt[3]{\mathbf{A}^{-1}(t_i)}]^T \bar{\mathbf{r}}(t_i)$ were small. With small residuals, the MMAE algorithm did not identify the correct filter in the early part of the time window, but that was not a




Figure 4.46 HCP of Filter 1 (correct answer, solid line), Filter 2 (dashed), Filter 3 (dotted), Filter 4 (dash-dot). Correct filter has highest HCP at filter equilibrium.

problem. What mattered was the later part of the time window, as the filters neared steady-state operation, and the residuals became distinct. In Figure 4.46, after four nanoseconds, the correctly matched filter clearly and consistently had the highest HCP.

Noise was added to the signals to test the algorithm's ability to operate effectively in the presence of noise. The noise was generated by a Gaussian pseudo-random number generator, with the result multiplied by the desired noise standard deviation. Since a pseudo-random signal is in fact deterministic, it is repeatable. The same noise signal could be added to different signals, or the same noise repeated in subsequent tests. One hundred samples of noise were used in this test. One sample was added to the radar signal for each test run. The variance of the noise was adjusted to achieve the desired SNR. The radar signals associated with each of the filters were tested.

Signals from the four simulated targets (Block 112, Block 123, Block 123B, and the 4.5" square plate) were used to test the effectiveness of the algorithm in the presence of varying amounts of noise. Noise was added so that the SNR ranged from 15 dB to -13 dB. Signals corresponding to each of the four filters were used at each SNR level. Each signal was repeated 100 times with different noise samples. The number of times each filter correctly recognized its associated signal at each SNR level is shown in Figure 4.47. The filters for which results are shown in Figure 4.47 were generated based on ϕ -polarized data. The signals used to test the algorithm were also ϕ -polarized.




Figure 4.47 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, ϕ -pol

As seen in Figure 4.47, the target recognition algorithm was able to distinguish between targets, even in the presence of significant noise. Using a signal scattered from Block 123 (the dashed line), the algorithm achieved a 90% success rate when the SNR was approximately 4 dB or higher. For the other targets, a 90% success rate was achieved with much lower SNR levels: approximately 0 dB for Block 123B and -4

dB for Block 112. For the square plate, a high level of success was achieved regardless of the SNR. As noted in Section 3.5, reconstructing scattered signals using the Block 123 signature resulted in high reconstruction error, indicating that improvements to the signature were desirable. With a better signature, it is likely that the residuals from the Block 123 filter will be smaller and the MMAE algorithm's recognition performance for that target will improve.

The algorithm was tested again with filters representing the same four targets (Block 112, Block 123, Block 123B, and the 4.5" square plate), but built with frequencies estimated from θ -polarization. While the poles found for the cube and block targets in Section 3.5 did not match those published in the literature, some poles were found and signal reconstruction was good for both the cube and Block 112, so the poles were used to test the target recognition technique. The number of times each filter correctly recognized its associated signal at each SNR level is shown in Figure 4.48. The test signals were θ -polarized. When these target signatures were developed, higher reconstruction errors were found for the θ -polarized case, so it was expected that the target recognition algorithm would have greater difficulty matching the signals correctly. The SNR again ranged from 15 dB to -13 dB.

In this test the algorithm failed to declare the Block 123 target at high SNR. As noted previously, the signature for Block 123 led to poor reconstruction error, and may need more development. Short of redefining the target's signature, the MMAE result for Block 123 may be addressed with change in filter tuning. The problem is examined further with the use of confusion matrices as introduced in Section 4.3.2. As indicated by Miller [89], tuning a filter for proper operation in isolation from other filters does not guarantee its successful operation in the context of an MMAE system. It is likely that tuning of individual filters was too conservative, and a change of values in the $\mathbf{Q}_d$ matrix of the poorly performing filters was required. An additional experiment with the same filters but different tuning is presented later in this section.




Figure 4.48 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, θ -pol

Table 4.4 shows an example of a confusion matrix for the targets shown in Figure 4.48. In this case, the filters were developed from θ -polarized scattering data. The SNR was -3 dB. The test signals used were also θ -polarized. Significant confusion did occur at -3 dB SNR between Block 123 and Block 123B.

Another confusion matrix is shown in Table 4.5. The data was taken with the SNR set to 3 dB, which is the point at which the algorithm performed the best.

Correct Target	Declared Bk 112	Declared Bk 123	Declared Bk 123B	Declared Sq Plt
Block 112	0.48	0.04	0.01	0.47
Block 123	0.04	0.31	0.01	0.64
Block 123B	0.32	0.32	0.09	0.27
Square Plate	0.02	0.02	0.01	0.95

Table 4.4 Confusion matrix, simulated radar data, SNR = -3 dB, θ -Polarized Signals

Correct Target	Declared Bk 112	Declared Bk 123	Declared Bk 123B	Declared Sq Plt
Block 112	0.95	0.01	0.02	0.02
Block 123	0.15	0.81	0.01	0.03
Block 123B	0.23	0.21	0.47	0.09
Square Plate	0.01	0.01	0.00	0.98

Table 4.5 Confusion matrix, simulated radar data, SNR = +3 dB, θ -Polarized Signals

Correct Target	Declared Bk 112	Declared Bk 123	Declared Bk 123B	Declared Sq Plt
Block 112	0.99	0.00	0.01	0.00
Block 123	0.47	0.24	0.26	0.03
Block 123B	0.03	0.00	0.97	0.00
Square Plate	0.00	0.00	0.00	1.00

Table 4.6 Confusion matrix, simulated radar data, SNR = 11 dB, θ -Polarized Signals

Although less confusion is evident at this SNR, Block 123B was still erroneously declared more often than it was correctly declared.

An interesting problem occurred with the θ -polarized data for the blocks and plate. The Block 123 target (dimensions 1" by 2" by 3") was erroneously declared to be Block 112 (dimensions 1" by 1" by 2") at high SNR levels, although it was declared correctly more often at lower levels, as shown in Figure 4.48. The confusion matrix for the target recognition results at 11 dB SNR was examined to help understand where problems occurred. The matrix is shown in Table 4.6. At this SNR level, the algorithm correctly identified the square plate as well as Blocks 112 and 123B (dimensions 1.1" by 1.9" by 2.8"). The signal from Block 123 was improperly associated with Block 112 more often than Block 123B. This problem can likely be solved through either retuning the filter for Block 123 or revisiting the target signature developed via M-TLS-MPM and the pole combination algorithm.

While Block 123B was confused with Block 123, both Blocks 123 and 123B were confused with Block 112. The confusion was not due to beta dominance. An

HCP for each elemental filter was determined using both Equation (2.76) and Equation (4.38). The first equation allows beta dominance to occur, while the second equation eliminates it. The same target was declared regardless of the equation used to compute each HCP.

The test was repeated with different filter tuning, and the results are shown in Figure 4.49. The new tuning consisted of adjustments to the values on the main diagonal of the $\mathbf{Q}$ matrix, which were then used to compute the $\mathbf{Q}_d$ matrix. There was significant confusion between Block 123B and Block 112 as well as between Block 123B and Block 123, as can be seen in Table 4.6. This experiment's results were better overall than the previous one involving these targets, detailed in Figure 4.48 and Table 4.4, but the confusion involving Block 123 indicates that further refinement of both the signature and the filter tuning is necessary. Due to time constraints, such refinement is left to future research.

Another test involved several body-of-revolution targets, with a θ -polarized incident field. The targets used were a right circular cylinder with length 5" and diameter 1", a right circular cylinder with length 6" and diameter 2", a combination of cylinders called a tophat (described in Section C.2 of Appendix C), and a cone of length 2" and base diameter 2". The percentage of correct identifications at various SNR levels is shown in Figure 4.50. The results were good, with better than 90% correct identifications when the SNR was greater than 1 dB for three target signals and better than 9 dB for all target signals.

A test was conducted with a different set of body-of-revolution targets using ϕ polarization. The targets were the Cone, Cylinder A, Cylinder B, Cylinder C. Results are shown in Figure 4.51. Except for Cylinder B, the targets required higher SNR levels to achieve a 90% successful rate of identification than tests with bodies of revolution using θ polarization. As was discussed in Section 3.5, the ϕ -polarized scattered signals from bodies of revolution are weaker than the θ -polarized signals. Since ϕ -polarized scattering is weaker, it may be difficult to obtain the higher-SNR




Figure 4.49 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated blocks and plate, θ -pol




Figure 4.50 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated cylinders and cone, θ -pol




Figure 4.51 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated Cone, Cylinder A, Cylinder B, Cylinder C, ϕ -pol

signals in practice. Further tuning may improve performance, although it may also be necessary to revisit the signature development for some of these targets.

In another test, block targets were mixed with bodies of revolution. The targets were the Cube, Combo A, Cone, Cylinder C. Results are shown in Figure 4.52. The signals were θ -polarized. This experiment produced good recognition results, with 90% or better correct identification when the SNR was above 3 dB. It showed that the technique is not limited to distinguishing among sets of similar objects.

The χ^2 test was used with the Cube, Combo A, Cone, Cylinder C targets, to see if it was any more useful with simulated radar data than it was with the synthetic data discussed in Section 4.3.3. Results are shown in Figure 4.53. As before, the χ^2 test was failed in approximately 4% of cases, regardless of SNR and regardless of whether the algorithm identified the correct target or not. As with the synthetic data, a χ^2 test provided no useful way to determine the confidence that the algorithm provided the correct answer.




Figure 4.52 MMAE-MAP, Percent Correct ID vs SNR (dB), Signatures: Simulated Cube, Combo A, Cone, Cylinder C, θ -pol




Figure 4.53 Percent of failures of χ^2 test vs SNR (dB), Signatures: Simulated Cube, Combo A, Cone, Cylinder C, θ -pol

The target recognition algorithm may be modified to employ a technique called Inter-Residual Distance Feedback (IRDF) [70]. This technique adjusts either filter tuning or gain online so that the distinguishability of residuals is increased. IRDF was developed by Lund for continuous-time filters and adapted for discrete-time use by Miller [89] and Vasquez [136]. IRDF was not employed in this research due to time constraints, but is a useful area for exploration by future researchers.

The signatures used in this research represented each target only from a single polarization. In actual operation of this system, it may be necessary to illuminate each target twice, from two orthogonal polarizations, to excite a sufficient number of scattering modes so that good target recognition may be achieved.

This research demonstrates that confusion matrices can indicate where filter retuning may be necessary. They are also a guide for users of the algorithm to situations in which the results may not be used with confidence.

4.5 Comparison to the E-Pulse

A resonance annihilation (RA) technique was implemented using the Extinction Pulse (E-Pulse) [112, 113]. The E-Pulse technique was implemented in two stages, pulse generation and pulse convolution. The pulse was generated with estimated poles s_n , $n = 1, 2, \dots, N$ of an object, found via the M-TLS-MPM. A rectangular pulse basis function was used because it combined ease of implementation with noise-suppressing ability superior to that of a delta basis function. The basis function was defined by

$$\beta(t) = \begin{cases} g, & (n-1)T \leq t \leq nT \\ 0, & \text{otherwise} \end{cases} \quad (4.43)$$

where g was a weight function used to normalize the pulse.

The Laplace Transform of the basis function was

$$\begin{aligned}
\beta^L(s) &= \int_{(n-1)T}^{nT} g \exp(-st) dt \\
&= -(g/s)[\exp(-snT) - \exp(-s(n-1)T)] \\
&= g \exp(-snT)[\exp(-sT) - 1]/s
\end{aligned} \tag{4.44}$$

A coefficient α_m was applied to each basis function. The coefficient vector $\bar{\alpha}$ was determined by

$$\begin{pmatrix} \beta_1^L(s_1) & \beta_2^L(s_1) & \cdots & \beta_N^L(s_1) \\ \beta_1^L(s_2) & \beta_2^L(s_2) & \cdots & \beta_N^L(s_2) \\ \vdots & \vdots & \ddots & \vdots \\ \beta_1^L(s_N) & \beta_2^L(s_N) & \cdots & \beta_N^L(s_N) \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_N \end{pmatrix} = \begin{pmatrix} -e^{fL}(s_1) \\ -e^{fL}(s_2) \\ \vdots \\ -e^{fL}(s_N) \end{pmatrix} \tag{4.45}$$

in which e^{fL} was the Laplace-domain forcing function. The forcing function was chosen to have the same form as the basis function, with an amplitude of 1.

The length of the basis function T was chosen to meet two sets of constraints. The first constraint was $T = p\Delta t$, $p = 1, 2, \dots$. This constraint was necessary to accommodate the E-Pulse technique to data sampled every Δt seconds. The second set of constraints was that $T \neq k\pi/\omega_n$, $k = 1, 2, \dots$, $n = 1, 2, \dots, N$. This was necessary to ensure the β matrix was not singular. The pulse length T was not constrained to be the same for each target's E-Pulse.

Figures 4.54 through 4.57 show the E-Pulses generated for the four versions of Synthetic Data. The poles used were those estimated via M-TLS-MPM from data with 25 dB SNR. Although the Version 1 E-Pulse shown in Figure 4.54 bears some superficial resemblance to that shown in Figure 4.55, the convolution results for those two pulses were very different. The closest similarity in convolution results, as determined by the Energy Discrimination Number, was between Version 1 and Version 4, although the E-Pulse shown in Figure 4.57 bears no obvious resemblance




Figure 4.54 E-Pulse for Synthetic Data Version 1, estimated from 25 dB data




Figure 4.55 E-Pulse for Synthetic Data Version 2, estimated from 25 dB data




Figure 4.56 E-Pulse for Synthetic Data Version 3, estimated from 25 dB data




Figure 4.57 E-Pulse for Synthetic Data Version 4, estimated from 25 dB data

to that in Figure 4.54. Given the similarity in EDN, confusion between these two targets can be expected.

Figures 4.58 through 4.61 show the results of convolving the E-Pulse from each version of Synthetic Data with a noisy signal representing Version 1. The eight Version 1 signals each had different mode coefficients, as shown in Table 3.3. These are shown to indicate the nature of the results from the E-Pulse technique. The vertical lines shown in Figures 4.58 through 4.61 are the limits after which the convolved signal should go to zero if the pulse eliminates all late-time resonance. Figure 4.58 shows several convolutions of the Version 1 signal with the E-Pulse based on Version 1 estimated poles. This is the correct answer, so the portion of the convolved signal to the right of the vertical line should be greatly reduced. In a noise-free case it should be equal to zero. In this case it is weak, but not equal to zero due to the added noise.

Figures 4.59 through 4.61 are incorrect matches between the E-Pulse and the received signal. Because the E-Pulses used to generate these figures were developed for a different set of frequencies than were present in the signal, the convolved signal to the right of the vertical line was not expected to be reduced as much as seen in Figure 4.58. The signal on the right side of each figure oscillated in a manner consistent with a damped sinusoid, indicating that the E-Pulse failed to eliminate all late-time resonance. This is an indication that the technique was working properly.

The oscillations on the right sides of Figures 4.60 and 4.61 were less pronounced than that of Figure 4.59, but were still obviously due to damped sinusoidal oscillation rather than noise. The closest match to Version 1 in E-Pulse experiments was Version 4, indicating that confusion between these two targets is possible.

Figures 4.62 and 4.63 show the results of target recognition tests using the E-Pulse in the presence of various amounts of noise. The pulses were generated for Synthetic Data based on parameters estimated via M-TLS-MPM at 25 dB SNR and at 13 dB SNR. Noise was added so that the SNR ranged from 23 dB to 1 dB. Signals




Figure 4.58 Eight convolutions of E-Pulse for Synthetic Data Version 1 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start




Figure 4.59 Eight convolutions of E-Pulse for Synthetic Data Version 2 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start




Figure 4.60 Eight convolutions of E-Pulse for Synthetic Data Version 3 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start




Figure 4.61 Eight convolutions of E-Pulse for Synthetic Data Version 4 with Noisy Version 1 Signal, 19 dB SNR, vertical line at late time start




Figure 4.62 E-Pulse, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 25 dB data




Figure 4.63 E-Pulse, Percent Correct ID vs SNR (dB), Signatures: Synthetic Data Versions 1,2,3, & 4 estimated from 13 dB data

corresponding to each of the four filters were used at each SNR level. Each signal was repeated 100 times with different noise samples. The number of times each filter correctly recognized its associated signal at each SNR level is shown in the figures.

The performance of the E-Pulse was significantly worse than that of the algorithm based on MMAE-MAP. Figure 4.62 can be compared to Figure 4.38 to see the SNR required for successful target recognition by the E-Pulse and MMAE-MAP techniques, respectively. The same signatures (sets of oscillation and damping rates) were used to represent the targets in each case. A similar comparison can be made between Figures 4.63 and 4.39, in which signatures were estimated from 13 dB SNR data.

Using synthetic data and estimated poles, the E-Pulse technique needed approximately 12 dB SNR or better to achieve accurate identification in over 90% of cases. This compares poorly to the MMAE-MAP technique, which requires an SNR of -7 dB or higher to achieve the same result. It is superior to the E-Pulse results found by Mooney, *et. al* [92], in which 20 to 42 dB SNR was required to achieve a 90% rate of successful identification, although differences in the target data and frequencies prevent a direct comparison. Mooney, *et. al* simulated scattering from wire targets via a combined SEM and Method of Moments (MoM) technique [92]. The frequencies used to generate the E-Pulse for each target were found via the same SEM/MoM technique.

4.6 Chapter Summary

A target recognition algorithm is presented which exploits the late-time scattering predicted by the Singularity Expansion Method. The target recognition system is based on MMAE and uses a Kalman filter to represent each target signature in a library. It can successfully distinguish between simple-shape, conducting targets in free space in the presence of a significant amount of white, Gaussian noise. The technique is able to operate in much greater levels of noise than a target recognition

technique based on the Extinction Pulse which used the same pole estimates. Further research may broaden the scope of available targets to include those of more complex shape and that are near or inside media other than free space.

This algorithm declares an unknown target to match one of the target signatures within its library. For a correct identification to be made, the library must contain the correct signature. That requires users of this technique to determine which target types are likely to be encountered, or otherwise of sufficient interest, and ensure that signatures for those target types are developed and the corresponding filters included in the library. Use of this algorithm with an insufficient library will generate false reports.

The success of the target recognition system depends on accurate estimation of the poles present in the late-time signal. Since the pole estimation technique M-TLS-MPM had difficulty estimating poles for the weak scattering from curved surfaces, the target recognition system has difficulty with targets such as spheres. In addition, M-TLS-MPM had difficulty estimating higher-frequency poles in many cases. Improvements in M-TLS-MPM or the pole combining algorithm PA&A will allow identification of more valid poles and the use of higher-order models in the target recognition algorithm. Further development of some signatures, such as that of the Block 123 or Combo A targets, may improve performance of the target recognition algorithm when the elemental filter representing that target is used.

A target's late-time scattered signal consists, in theory, of an infinite number of scattering modes. These modes can notionally be divided into significant, marginally significant, and insignificant modes. The significance of a given scattering mode is due to its amplitude relative to other scattering modes and noise. M-TLS-MPM and similar pole estimation programs that are based on analysis of the late-time signal will not estimate poles for insignificant modes. Such algorithms may or may not estimate poles for modes that are marginally significant, depending on parameter values in the algorithm and the amount of noise in a signal.

If poles for marginal modes are not estimated, the target filters in the MMAE algorithm are in effect based on reduced order models. In some cases in this research, unmodeled terms appeared to lead to incorrect identifications. Adjusting parameters to allow more poles to be estimated allows the possibility that spurious pole estimates, based on noise, will be included in the signature. These additional poles could also lead to incorrect identifications. Results of the suppressed mode experiment in Section 4.2.2 indicate that states representing weakly excited modes will react properly to the weak excitation. Those results may tempt a user of this algorithm to set pole estimation parameters in a 'loose' manner and deliberately allow some noise-induced, spurious poles, knowing that they are likely to be weakly excited when the target recognition system is used. That is a false notion, since what matters is the performance of the MMAE system as a whole, not that of a single filter. Spurious pole estimates lead to additional states that may react to a component of a signal which the filter was not intended to match. Users of this algorithm must balance a desire for a complete description of the target's signature against a desire to eliminate spurious poles.

Successful use of the algorithm also depends on proper tuning of the noise model. This tuning consists primarily of adjustments to the dynamics noise covariance in each Kalman filter. An additional tuning option involves varying the units of the parameters in the system dynamics matrix to change its numerical conditioning, which affects the filter performance. Tuning is a process involving trial and error. Filters should be tuned in two stages. First, they should be tuned individually, to ensure they follow the desired signals and fail to follow undesired ones. They must also be tuned to work well as elements of an MMAE system.

V. Conclusion and Recommendations

The target recognition system presented here, based on Multiple Model Adaptive Estimation (MMAE), can associate late-time radar data with a signature developed in advance. The process is simple and computationally efficient compared to techniques involving synthetic aperture radar. The recognition algorithm has better noise resistance than the Extinction Pulse (E-Pulse) [18, 52, 92, 112, 113, 117] and related techniques [14, 36, 58, 117] that have been the subject of most research into target recognition using late-time scattering from transient radars.

The target recognition technique has been demonstrated using simulated scattering data from simple geometric shapes, such as cylinders and blocks of conducting material. These shapes may be found as substructures on various ground vehicles, or as field equipment or weapons such as land mines.

The target recognition technique is resistant to white, Gaussian noise, so it would be resistant to intentional interference (jamming) using white noise or wide-band noise. It is not designed to be resistant to narrowband interference, so intentional interference using single tones near the oscillating rate of targets in the library may be more successful in causing false target declarations than white noise.

To recognize targets, it is necessary first to have an adequate target model for each, based on their natural resonant frequencies. The more accurately a target's significant frequencies are measured, the better the target recognition results will be. To provide for better estimates, an improved estimation algorithm, the Modified Total Least-Squares Matrix Pencil Method (M-TLS-MPM), has been demonstrated. M-TLS-MPM has improved performance in noise compared to other versions of the Matrix Pencil Method (MPM), such as TLS-MPM [47, 120] and MMP3 [69]. M-TLS-MPM is more noise sensitive than the target recognition algorithm, however, so it should be used under controlled conditions to minimize the effect of noise.

While poles estimated with M-TLS-MPM did not in all cases match previously published results, the poles estimated for each target illuminated at various aspect angles tended to occur in approximately the same locations in the complex plane. That is consistent with the prediction in SEM theory that poles are aspect-independent. A simple algorithm, Pole Association and Averaging (PA&A), has been developed to combine frequencies estimated via M-TLS-MPM from several aspect angles of the target into a single set of frequencies representing the target at a broad range of aspect angles.

Signals reconstructed using the multiple-aspect (PA&A) poles followed the scattered signals with little error. There were cases in which some poles estimated for a target did not correspond to poles at other azimuths. The PA&A algorithm rejected those non-corresponding poles and generated a set of average poles for the target that still led to good signal reconstruction in most cases. Since the rejected pole estimates were not significant contributors to reconstruction, they may have been the result of noise. The frequencies found by M-TLS-MPM and PA&A are used to develop the component elemental filters of the MMAE target recognition algorithm.

5.1 *Contributions*

This research has contributed to existing knowledge in the areas of signal processing and stochastic estimation. The overall contribution of this research is the application of stochastic estimation theory, specifically the MMAE-MAP technique, to the problem of target recognition by late-time resonance. An MMAE-MAP algorithm was developed, implemented, and tested as the basis for a successful target recognition algorithm. The algorithm can successfully distinguish the targets used in this research from one another when the signal-to-noise ratio (SNR) is approximately 3 dB, with few or no errors.

The MMAE-MAP algorithm has much better resistance to additive, white, Gaussian noise than the E-Pulse algorithm [113]. In tests using damped sine series data, the MMAE-MAP target recognition technique was able to achieve a success rate of 90% when the minimum SNR was between -15 and -7 dB, depending on the signal used. The E-Pulse technique required a minimum SNR between +7 and +11 dB, depending on the signal, to achieve a 90% rate of successful identification for synthetic data.

In the signal processing area, a further contribution is the Modified Total Least-Squares Matrix Pencil Method (M-TLS-MPM), a novel variant of the Matrix Pencil Method (MPM) [49]. M-TLS-MPM combines elements of two MPM variants, the Total Least-Squares Matrix Pencil Method (TLS-MPM) [120] and the Modified Matrix Pencil Method 3 (MMP3) [69], into a single algorithm with improved parameter estimation performance in noise. In tests with complex damped exponential data, the M-TLS-MPM algorithm achieved lower mean-square error (MSE) in parameter estimates than either TLS-MPM or MMP3 when the SNR was 11.4 dB or less. It approached the Cramer-Rao bound at SNR levels above 11.4 dB. The M-TLS-MPM and the MMP3 use an iterative technique known as the Low-Rank Hankel Approximation (LRHA) [66, 69]. The effect of a varied number of iterations on parameter estimation was explored. The reduction in MSE of the parameter estimate was not significant after the fourth iteration for the cases studied. Most of the reduction in parameter MSE is attributed to the first iteration.

This algorithm could contribute to Air Force efforts to detect and identify vehicles and air defense weapons hidden in foliage or under camouflage, using reconnaissance aircraft or unmanned air vehicles. Other potential applications of military interest involve scanning packages and personnel at a distance for the presence of weapons. It may also enable rapid detection and identification of land mines and unexploded ordnance from aircraft.

5.2 *Recommendations for Future Research*

Research was conducted in the areas of signal processing and stochastic estimation. The recommendations for further research fall naturally into these two areas, in addition to the field of electromagnetics. The recommendations are organized by area in the following sections.

5.2.1 Recommendations for Future Electromagnetics Research. The targets measured and simulated in this research have been treated in a free-space environment. In other words, all measurements and predictions have separated the target from other targets and media, save for the styrofoam mounts used in measurement ranges. In a real-world application, the target of interest may be a ground vehicle, a water vehicle, an underground structure, or a land mine. Hanson and Baum [42] have shown how to approximate the poles of an object within or adjacent to a lossy dielectric medium such as earth or sea water. Their result is a perturbation formula based on an exact integral equation.

It will be useful to measure targets with short-pulse illumination and compare to swept-frequency results transformed into the time domain. Most likely the targets to be measured with short-pulses will be larger than those discussed herein, so it will be necessary to scale target size and frequencies. That is easily accomplished with swept-frequency results. If the object is longer, wider, and deeper by a factor of x , the wavelength must be longer by a factor of x , so the frequency must be lower by a factor of x . In this way, a three-inch-long cylinder with a one-inch diameter, measured from 2 to 18 GHz is equivalent to a three-foot-long cylinder with a one-foot diameter, measured from 166 MHz to 1.5 GHz.

In addition to short-pulse measurements, it will be useful to perform time domain simulation of CAD models of the targets. Limited results from TEMPUS [122] were considered, but the results were reported in the frequency domain and reconverted to time domain. Direct reporting of time-domain results removes two

steps, FFT and IFFT, which contribute to uncertainty in the result. Direct time-domain results may then be compared to frequency domain results transformed into the time domain.

If this target recognition algorithm is implemented in a ground-to-ground situation, the target may be partially obscured by terrain. If only a portion of a target is visible to the radar, some substructures will not be directly illuminated, and their ringing returns will be reduced significantly. That will skew the perceived signature toward the more strongly illuminated poles. The target recognition algorithm may then be unable to identify the obscured target properly. Additional research can explore the ability of the algorithm to discriminate between several targets when the targets are partially obscured.

Lehman [63] developed statistical models for fields inside cavities. This approach applies various simplifying assumptions to the natural modes present in a fully-enclosed cavity (no apertures) of general shape. Lehman did not attempt to determine poles for any specific target, but his techniques hold promise to find analytic solutions, based on statistical treatment, for a wider variety of objects than is possible using only deterministic approaches. If the target recognition technique developed herein is applied to ground vehicles, the scattered signal from some vehicles (tanks, artillery, etc.) may include a significant cavity response from the main gun at some aspect angles. If the technique is applied to identification of concealed firearms carried by an individual person, the cavity response from the barrel of the weapon may be significant for sufficiently high frequencies. Further statistical or analytic development of late-time scattering from cavities may assist in development of signatures for targets with such cavities.

It may be possible to improve the effective bandwidth of radar scattering data using a technique called bandwidth extrapolation (BWE, a form of super resolution) [19, 29]. BWE is a technique ordinarily associated with synthetic aperture radar (SAR) which can lead to finer details in a SAR or inverse SAR image. It

may be useful in the analysis of the measured data. From the author's experience using BWE, the effective bandwidth of an inverse SAR image can be improved by a factor of two with only a minor addition of spurious features (correlated noise) in the image. The applicability of BWE to the late-time resonance due to swept frequency measurements may be considered. Other super-resolution techniques, including various Maximum Entropy Method techniques, have been applied to the ARX, IV and ARMAX algorithms with little success [25].

Here are some tasks appropriate for future research:

1. Study theoretical effect of lossy dielectric half space on target signatures so the algorithm can be applied to targets on or under ground or sea. Measure targets embedded in moist soil, inside a styrofoam container, to trace the effect of the soil.
2. Measure and simulate more complicated targets to move algorithm to real world applicability.
3. Measure and simulate targets with resistive or dielectric coatings. This will allow extension of the target recognition technique to targets other than highly-conductive metallic bodies.
4. Measure targets with short-pulse illumination (collecting time-domain data) and compare the resulting late time return to the results of swept-frequency illumination (frequency-domain data) which has been transformed to the time domain via IFFT.
5. Simulate targets in a time-domain simulation such as TEMPUS [122] and compare the resulting late time return to the results of a frequency domain simulation such as CARLOS [138] or JRMBOR [111], involving the same target geometry, after transformation into the time domain via IFFT.
6. Develop SEM theoretical poles for the interior of a circular cross-section cavity which is open at one end. A scalar Green's function for the cavity is available

[22, 126] which can be converted into a dyadic Green's function [6]. This will allow application of the target recognition technique to objects such as tanks and artillery pieces, for which waveguide response may be significant at some frequencies and aspect angles.

7. Apply the statistical electromagnetics techniques of Lehman [63] to find poles analytically for various objects including cavities.
8. Apply bandwidth extrapolation [19, 29] to the IFFT algorithm to see if it is possible to improve the amount of useful data without distorting the late-time signal.
9. For efficient analysis of targets over wider bandwidths, apply nonlinear (logarithmic) sampling to the IFFT algorithm to maintain adequate sampling of lower frequencies while reducing the number of high-frequency samples.
10. Estimate the poles of a cube and blocks with edge lengths in the proportions 1:1:2 and 1:2:3 by analysis of the impedance matrix in a method of moments simulation. Simulate the entire object in each case, and compare the results to those of Long [68].
11. Simulate and analyze the signals scattered from a cube and blocks with edge lengths in the proportions 1:1:2 and 1:2:3 at angles other than backscatter. Find aspect angles at which scattering is strongest, and estimate poles from the scattered signals at those angles via M-TLS-MPM.

5.2.2 Recommendations for Future Signal Processing Research. The Matrix Pencil Method and its variants are designed to estimate the parameters of a summation of complex exponential terms with constant coefficients. These correspond to first-order poles in the Laplace domain. If the target response contains second or higher-order poles, the time-domain signal will contain complex exponential terms with time-varying coefficients. Conducting objects which are coated in a resistive material will have second-order poles in their late-time response [132], and

the scattering modes corresponding to these poles may be significant. To estimate poles from such a target, it may be necessary to modify the Matrix Pencil Method (and variants thereof, such as M-TLS-MPM) significantly.

The late-time response from objects near to or embedded in a lossy dielectric half-space includes branch points [37]. If the scattering mode associated with these branch points is significant, further modifications of the Matrix Pencil Method (and its variants) may be necessary.

Here are some tasks appropriate for future research:

1. Study the effect of narrowband or colored noise on estimates of target signatures.
2. Measure and simulate more complex targets to move algorithm to real world applicability.
3. Study the effect of a lossy dielectric half space on target signatures (changed or added poles) so the algorithm can be applied to targets on or under the ground or sea, or small targets adjacent to a human body.
4. Study the effect of resistive or dielectric coatings on estimated signatures. Determine whether Matrix Pencil Method variants can be modified to account for second order poles or other phenomena.
5. Apply Bandwidth Extrapolation [19, 29] to the radar data and note the effect on the Matrix Pencil Method.
6. Develop guidelines for the number of iterations to use in the Low-Rank Hankel Approximation (LRHA).
7. Compare pole estimation performance in noise of MPM variants against that of resonance annihilation techniques.
8. Compute the Cramer-Rao bound for synthetic data and recompute it for Kumaresan-Tufts data.

9. Develop the signature of the Block 123 and Combo A targets again using different parameter values in the M-TLS-MPM and PA&A codes. Also develop multiple signatures representing these targets from limited sets of aspect angles.

5.2.3 Recommendations for Future Stochastic Estimation Research. The Kalman filter algorithm discussed in Section 2.4 is known as the Joseph form of the filter. Other variations exist which have better numerical stability, such as the U-D covariance factorization form [77]. In the U-D factor form, the covariance matrix $\mathbf{P}(t_i^-)$ is factored into component matrices $\mathbf{U}(t_i^-)$ and $\mathbf{D}(t_i^-)$ such that $\mathbf{P}(t_i^-) = \mathbf{U}(t_i^-)\mathbf{D}(t_i^-)\mathbf{U}^T(t_i^-)$. The algorithm computes and uses $\hat{\hat{x}}$, $\mathbf{U}$ and $\mathbf{D}$ rather than $\hat{x}$ and $\mathbf{P}$. The U-D factor form provides improved numerical stability and guarantees that the covariance of the state estimate is positive semidefinite. In addition, the numerical precision required to represent computations with the $\mathbf{U}$ and $\mathbf{D}$ matrices accurately is less than that needed for computations with the $\mathbf{P}$ matrix. The disadvantage of the U-D technique is an increase in required computations compared to a conventional Kalman filter. However, the number of computations required is actually less than the number required for the Joseph form Kalman filters implemented herein.

Neither the conventional nor the Joseph form of the Kalman filter is numerically stable, but the U-D factor form can be shown to be numerically stable [77]. Numerical stability as discussed here means that the computed result of the algorithm applied to a given problem corresponds to the exact solution of a problem that is slightly perturbed from the given problem. This is the standard definition of Wilkerson [141].

The radar signal is assumed to be corrupted by white noise, but the author has conducted no experiments to verify that the white assumption is a good one. There may be colored noise in measured radar data either in experimentation or in use of this algorithm in realistic conditions. The noise in computational codes is not

strictly white. It will consist of some white noise in addition to some band-limited, or colored noise. The target recognition algorithm could be modified to account for colored noise.

If a few targets are difficult to distinguish due to nearly identical frequencies, a hierarchical system can be adopted [139]. In this case, the algorithm would identify the target by comparison to a preliminary model which could fit a few actual targets. Then it could initiate a new set of filters, each of which used a more detailed linear system model of the target's response. A more detailed model would include frequencies which may have been left out of the preliminary model. A variant of the hierarchical approach known as MMAE with Filter Spawning [33] has been applied to failure detection and identification problems. Once a failure is detected, the MMAE with Filter Spawning algorithm adds filters to the bank. These 'spawned' filters represent partial, rather than complete, failures. Spawned filters in the context of target recognition could be based on more elaborate, higher-fidelity, models of the scattered signal from various targets. If the declared target is known to have a high probability of being confused with another target, higher-fidelity target models may be employed when subsequent pulses are used to illuminate the target. An alternative to a hierarchical approach may be a moving-bank MMAE algorithm, which could incorporate a very large number of linear system models, but would only initiate those which are considered likely to be correct, based on scaled residuals [76,136]. The hierarchical, the filter spawning, or the moving-bank MMAE system may be a good way to deal with non-nominal configurations of targets, i.e., small changes in target configuration that make its signature deviate from the standard. In the case of an armored vehicle target, a non-nominal configuration could be an open hatch or a turned turret.

Lund [70] has proposed an inter-residual distance feedback technique to increase distinguishability between filters in MMAE. This technique adjusts the noise covariance matrices of individual Kalman filters to keep a measure of inter-residual

distance above a specified minimum value. The technique is intended to increase distinguishability of the elemental filters within an MMAE algorithm.

Here are some tasks appropriate for future research.

1. Replace existing Kalman filters with U-D factored form algorithm which provides improved numerical precision and stability, and exploits sparse nature of matrices.
2. Apply filter spawning and moving-bank concepts to allow efficient use of larger target signature libraries by the MMAE-MAP algorithm.
3. Perform additional filter tuning experiments to get best possible performance from the system.
4. Port code from MATLAB to C++ and compile it to improve speed of execution. Benchmark execution speed against other target recognition techniques using similar computing power.
5. Study effect of narrowband or colored noise on target distinguishability
6. Replace existing Kalman filters in the MMAE-MAP algorithm with Kalman-based optimal smoothers which provide improved noise resistance. Apply a U-D formulation of the Kalman smoother for improved numerical precision and stability.
7. Redevelop the filter equations to account for a mix of first and second order poles.
8. Find the time domain contribution due to branch point singularities in the s plane. Redevelop the filter equations to account for branch point contributions.
9. Apply Inter-Residual Distance Feedback [70,89] to improve distinguishability between filters.

Appendix A. Additional Discussion on Electromagnetic Theory

In this appendix certain aspects of electromagnetic theory are discussed which were explored to a limited extent during research, but were not central to the overall topic of target recognition via late-time transients. The sections below cover a mathematical consideration not addressed in the Singularity Expansion Method (SEM) sphere analytic development, a critique of the application of SEM theory to wire loops, and some practical considerations regarding the start of late time.

A.1 Mie Series Convergence

The complex field scattered from an ideal PEC sphere at any frequency is found using the Mie series [129]. The Mie series specified to backscattering is

$$E^r(\omega) = \frac{1}{2k} \sum_{n=1}^{\infty} (-1)^{n+1} \frac{2n+1}{h_n^{(2)}(ka) \frac{\partial h_n^{(2)}(ka)}{\partial(ka)}} \quad (\text{A.1})$$

in which $k = \omega/c$, the sphere radius is a , and $h_n^{(2)}(x) = j_n(x) - jy_n(x)$ is the Spherical Hankel Function of the second type [4]:

$$h_n^{(2)}(x) = \left(\frac{\pi}{2x}\right)^{1/2} H_{n+1/2}^{(2)}(x) = \left(\frac{\pi}{2x}\right)^{1/2} [J_{n+1/2}(x) - jY_{n+1/2}(x)] \quad (\text{A.2})$$

In the Mie-series code used to compute the sphere scattering, the field is multiplied by a factor of $2\sqrt{\pi}$. That allows the user to find RCS by multiplying the field by its conjugate. Essentially, the factor 4π found in calculations of RCS is introduced here and split between the field and its conjugate. This convention is widely used in radar processing software and computational electromagnetics codes. If power is calculated using these fields, the magnitude must be reduced by 10.99 dB to get the correct value.

The Mie series converges quickly at lower frequencies but requires more terms at higher frequencies. To use it efficiently, a guideline is needed for the correct number

of terms to employ to get the desired convergence result. The sphere radius was set to 3" and compared the results for 10, 20, 30, 40, 50, 55, and 60-term truncations of the series. It was found that for 18 GHz, the 40 and 50 term series agree to 6 significant digits, while the 55 and 60 term series agree to 17 digits at 18 GHz. A most efficient computation of the Mie series would increase the number of terms used as frequency rises. In fact, at lower frequencies the series may run into overflow or underflow numeric errors if too many terms are employed. For frequencies below .05 GHz and a sphere of 3" radius, fewer than 20 terms must be used, at least in the MATLAB[®] formulation.

A simple rule of thumb can be developed to approximate the number of terms required for series convergence. The number of terms required increases with either increasing radius or decreasing wavelength. In the Mie series, the radius a and wavelength λ terms always appear in the ratio $ka = 2\pi a/\lambda$ in the arguments of functions. That means a useful rule of thumb for convergence should be based on the ratio of radius to wavelength. The number of terms required can be related directly to this ratio to find a relation between a desired level of convergence and the number of terms necessary to achieve it.

It is interesting to note that the Mie series relies on spherical Hankel functions. The large argument asymptotic form of these functions is based on a square root of the argument.

Table A.1 shows the results of another convergence test. Convergence was verified using a Mie series code and adjusting the number of terms until there was complete agreement between answers to fifteen digits. The point of that was to try to fit a curve to the convergence limit. Various functions of a/λ were compared to see if the number of terms required can be modelled simply. Here f is frequency in GHz, a is the sphere radius, c is the speed of light in free space and λ is the wavelength. A reasonable fit within the indicated frequency range is found via either the square or cube root of a/λ . The function *ceil* is the smallest integer which is greater than

frequency (GHz)	a/λ ($a = 3''$)	$ceil(500a/\lambda)$	$ceil(25\sqrt{a/\lambda})$	$ceil(32\sqrt[3]{a/\lambda})$	Conver- gence
0.01	0.00254	2	2	5	4
0.03	0.00762	4	3	7	5
0.10	0.0254	13	4	10	6
0.30	0.0762	39	7	14	7
1.00	0.254	128	13	21	12
3.00	0.762	382	22	30	20
10.00	2.54	1271	40	44	37
30.00	7.62	3813	70	63	76

Table A.1 Number of terms of the Mie series required for 15-digit convergence, compared to functions of a/λ

or equal to the argument; it rounds up. A very good fit can be achieved over a wide bandwidth if the number of terms used is equal to:

$$2 + ceil(24\sqrt{a/\lambda}) \quad (\text{A.3})$$

A.2 Sphere Theoretical Poles

In this section the theoretical development of sphere late-time scattering via SEM is discussed. Theoretical sphere late-time scattering is used to test the pole estimation algorithms in Chapter III. Some additions to the SEM theory are made and theoretical results are compared to measured results from a sphere. A sphere is the one target for which a complete analytic description exists of both narrowband scattering (the Mie series) and late-time, ultra-wideband scattering. Comparisons to theory allow assessment of the quality of the measured and computed sphere data. Theory is also compared to computational results for a sphere.

Compared to other simple shapes, the sphere is weakly resonant. The poles of the sphere have larger damping rates $|\Omega|$ than other targets of similar size such as plates, wires, and cylinders. Therefore, the exponential terms associated with each pole die out more quickly than other targets' exponential terms. That makes the

theoretical sphere a difficult test case for the pole estimation techniques introduced in Chapter III.

The original development of SEM included a complete theoretical development for a sphere, as well as a short list of poles and a graph of a few pole locations [6]. Tesche published a more extensive graph of pole locations in 1973 [131]. Here are the theoretical SEM poles for a sphere assumed to be a perfect electrical conductor. The poles are the frequencies $s_{q,n,n'}$ which satisfy one of two equations. Those with subscript $q = 1$ affect surface current but not surface charge density; they satisfy:

$$k_n(s_{1,n,n'}a/c) = 0, \quad n = 1, 2, \dots \quad (\text{A.4})$$

$$-floor(n/2) \leq n' \leq floor(n/2), \quad n' \neq 0 \text{ if } n \text{ even}$$

The *floor* function in Equation (A.4) returns the largest integer less than or equal to the argument; it ‘rounds down’. Those with subscript $q = 2$ affect both surface current and charge density. The $q = 2$ poles satisfy:

$$\frac{d}{ds}[(a/c)s_{2,n,n'}k_n(s_{2,n,n'}a/c)] = 0, \quad n = 1, 2, \dots, \quad (\text{A.5})$$

$$-floor((n+1)/2) \leq n' \leq floor((n+1)/2), \quad n' \neq 0 \text{ if } n \text{ odd}$$

The poles in the two cases give rise to the solenoidal ($q = 1$) and irrotational ($q = 2$) terms in the surface current density. Both $q = 1$ and $q = 2$ poles should be present in the field scattered from a sphere. In each of these equations, k_n is defined by

$$\begin{aligned} k_n(\zeta) &= \frac{\exp(-\zeta)}{\zeta} \sum_{p=0}^n \frac{(n+p)!}{p!(n+p)!} (2\zeta)^{-p} \\ &= -(j^{-n})h_n^{(2)}(-j\zeta) \end{aligned} \quad (\text{A.6})$$

in which $h_n^{(2)}(\zeta)$ is the spherical Hankel function and

$$\zeta = \frac{a}{c} s_{q,n,n'} \quad (\text{A.7})$$

To take the derivative of a complex function with respect to a complex variable, differentiability must first be proved. This is accomplished in Section A.3 for the derivative shown in Equation (A.5):

$$\begin{aligned} \frac{d}{d\zeta} [\zeta k_n(\zeta)] &= \frac{d}{d\zeta} \left[\exp(-\zeta) \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} (2\zeta)^{-p} \right] \\ &= \frac{d}{d\zeta} [\exp(-\zeta)] \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} (2\zeta)^{-p} \\ &\quad + \exp(-\zeta) \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} 2^{-p} \frac{d\zeta^{-p}}{d\zeta} \\ &= -\exp(-\zeta) \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} (2\zeta)^{-p} \\ &\quad + \exp(-\zeta) \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} 2^{-p} (-p) \zeta^{-p-1} \\ &= -\exp(-\zeta) \sum_{p=0}^n \frac{(n+p)!}{p! (n+p)!} \zeta^{-p} \left(1 + \frac{p}{\zeta}\right) \end{aligned} \quad (\text{A.8})$$

The poles can be found numerically from Equations (A.6) and (A.8). The series in each equation is simply a polynomial in ζ ; the factor $-\exp(-\zeta)$ has no effect on the zeros. The equation may be multiplied by ζ^n to put it in a more convenient form. It is then straightforward to calculate the roots of each polynomial numerically.

The first several sphere poles are listed below for the cases $q = 1$ and $q = 2$. The $q = 1$ poles do not appear in the surface charge equations in Baum's development of the sphere poles, while the $q = 2$ poles do [6].

The fourth column of Table A.2 represents the normalized frequencies in ζ , while the fifth lists the frequencies for a six-inch diameter sphere. The units in the fifth column are in giganepers per second and gigaradians per second.

q	n	n'	ζ	$s_{q,n,n'} \text{ (Gn/s} \pm j\text{Grad/s)}$
1	1	0	$-1.0000 \pm j0$	$-3.9370 \pm j0$
1	2	1	$-1.5000 \pm j0.8660$	$-5.9055 \pm j3.4095$
1	3	0	$-2.3222 \pm j0$	$-9.1425 \pm j0$
1	3	1	$-1.8389 \pm j1.7544$	$-7.2398 \pm j6.9070$
1	4	1	$-2.8962 \pm j0.8672$	$-11.4024 \pm j3.4143$
1	4	2	$-2.1038 \pm j2.6574$	$-8.2826 \pm j10.4623$
2	1	1	$-0.5000 \pm j0.8660$	$-1.9685 \pm j3.4095$
2	2	0	$-1.5961 \pm j0$	$-6.2837 \pm j0$
2	2	1	$-0.7020 \pm j1.8073$	$-2.7636 \pm j7.1155$
2	3	1	$-2.1571 \pm j0.8706$	$-8.4927 \pm j3.4274$
2	3	2	$-0.8429 \pm j2.7579$	$-3.3184 \pm j10.8577$
2	4	0	$-2.9487 \pm j0$	$-11.6092 \pm j0$
2	4	1	$-0.9542 \pm j3.7148$	$-3.7568 \pm j14.6251$

Table A.2 Sphere Theoretical Poles

There are in theory an infinite number of poles. Those found for $q = 1$ have a real component Ω with larger magnitude than comparable poles from $q = 2$. This observation is based on the poles for $n = 1$ through $n = 27$, which follow a definite pattern. Poles associated with a higher index n depart from this pattern as noted below. Baum noted the poles for a given n fall approximately in an arc from $-jnc/a$ to $+jnc/a$ through $-0.66nc/a$. In Figure A.1 it can be seen that this approximation holds through $n = 27$.

Poles for $n > 27$ did not fit the pattern predicted by Baum due to limits of numerical precision in the computational software package. The poles for $n = 1 : 35$ are shown in Figure A.2 and those for $n = 1 : 85$ are shown in Figure A.3. MATLAB[®] is unable to calculate the poles beyond $n = 85$, because of numerical underflow or overflow somewhere in the calculation. With higher-precision computations, more of the computed poles would fit the pattern seen for the $n = 1 : 27$ poles. MATLAB[®] computes roots of a polynomial by computing eigenvalues of the polynomial's companion matrix [45, 73]. Interestingly, the $q = 1$ and $q = 2$ poles reached convergence limits at the same value of n although they had different functional forms.




Figure A.1 Theoretical poles for 3" radius sphere for n=1 through n=27, q=1 (·) and q=2 (+)




Figure A.2 Theoretical poles for 3" radius sphere for n=1 through n=35, q=1 (·) and q=2 (+)




Figure A.3 Theoretical poles for 3" radius sphere for $n=1$ through $n=85$, $q=1$ (·) and $q=2$ (+)

The original pole derivation would lead one to assume that, as incident frequency increased, the resonant $q = 2$ poles forming a curve closest to the $\Omega = 0$ axis would continue in a more or less straight line toward $\omega = \infty$ and $\omega = -\infty$. The poles probably in fact fit this pattern, but the numerical precision required to find them is beyond the 15 digits available to MATLAB®.

In any event, the poles closest to the $j\omega$ axis are the ones associated with the slowest-decaying, or longest-duration signals. These $q = 2$ poles are the ones which have the most significant effect on the late-time signal because their effect extends the furthest. The size of the coefficients associated with these poles should also be considered to confirm that they are in fact the most significant.

Figure A.4 shows the upper half-plane sphere poles with an oscillating frequency $\omega < 36\pi$ gigaradians per second and a damping rate $\Omega > -8.4$ giganepers per second. These are the poles expected to be excited by frequencies up to 18 gigahertz. The magnitudes for the coefficients estimated for those poles are shown




Figure A.4 Theoretical poles with $\Omega > -8.4$ for 3" radius sphere

in Figure A.5. Only a few poles account for a large portion of the signal magnitude. The sphere late-time scattering was reconstructed using only the poles with the largest coefficients, numbered 2, 31, 32, 33, and 34 in Figures A.4 and A.5. The reconstruction error using only these poles was 0.64 percent. The most significant contributors to the signal were not the $q = 2$ poles with the smallest $|\Omega|$ and a range of $|\omega|$ values, but a few poles, primarily with $q = 1$, with the smallest $|\Omega|$ and $|\omega|$.

The poles numbered 31 through 34 in Figure A.4 correspond to four of the $q = 1$ poles shown in Table A.2. The poles numbered 1 through 5 in Figure A.4 correspond to five of the $q = 2$ poles shown in Table A.2. Poles numbered 6 through 30 in Figure A.4 are $q = 2$ poles with $n \geq 4$ that are not shown in Table A.2 .

Figure A.6 shows a portion of the late-time scattering for an ideal sphere with a diameter of six inches. The large negative return at approximately 5" downrange from the center of the sphere is the creeping wave return due to current induced on the illuminated side of the sphere continuing around the unilluminated side and back




Figure A.5 Estimated coefficient magnitudes for theoretical poles with lowest $|\Omega|$ for 3" radius sphere




Figure A.6 Theoretical late-time return from a 3" radius sphere, incident illumination 0.1 to 18 GHz

to the illuminated side, where it radiates energy back to the source of illumination [59].

A.3 Proof of Differentiability for Sphere SEM Poles

Here a mathematical consideration is introduced which has not been explicitly included in the literature to date [6, 131]. Taking the derivative of $k_n(\zeta)$ requires that its differentiability with respect to a complex variable ζ first be established using the Cauchy-Riemann equations. If a complex function of a complex variable satisfies these equations, it is differentiable [64, 125]. This is more easily shown using the summation form of the definition in Equation (A.6).

First, it must be possible to identify real and imaginary components of $\zeta k_n(\zeta)$ as well as modulus (complex magnitude) and phase components. Let $\zeta = x + jy$, so

$$\zeta k_n(\zeta) = (x + jy)k_n(x + jy) = \sum_{p=0}^n \frac{(n+p)! 2^{-p}}{p! (n+p)!} (x + jy)^{-p} \exp(-x - jy) \quad (\text{A.9})$$

The differentiability of each term may now be evaluated separately because differentiation is a linear operator in the complex plane as well as the real line [64, 125]. Differentiability is then determined for

$$(x + jy)^{-p} \exp(-x - jy) = |x + jy|^{-p} \exp(-x) \exp(-jy) \exp(-jp \arg(x + jy)) \quad (\text{A.10})$$

so each term of $\zeta k_n(\zeta)$ has modulus m and phase ϕ defined as

$$m = |x + jy|^{-p} \exp(-x), \quad \phi = -y - p \arg(x + jy) \quad (\text{A.11})$$

and rectangular components $u(x, y), v(x, y)$ defined as

$$u = m \cos(\phi), \quad v = m \sin(\phi) \quad (\text{A.12})$$

The Cauchy-Riemann equations [64, 125] are now applied

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}, \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x} \quad (\text{A.13})$$

which in their polar form are

$$\frac{\partial u}{\partial m} = \frac{1}{m} \frac{\partial v}{\partial \phi}, \quad \frac{\partial v}{\partial m} = -\frac{1}{m} \frac{\partial u}{\partial \phi} \quad (\text{A.14})$$

The polar form equations are easily verified

$$\frac{\partial u}{\partial m} = \cos(\phi) = \frac{1}{m} \frac{\partial m \sin(\phi)}{\partial \phi} = \frac{1}{m} \frac{\partial v}{\partial \phi} \quad (\text{A.15})$$

$$\frac{\partial v}{\partial m} = \sin(\phi) = \frac{1}{m} \frac{\partial m \cos(\phi)}{\partial \phi} = \frac{1}{m} \frac{\partial u}{\partial \phi} \quad (\text{A.16})$$

Each term of Equation (A.9) is differentiable by ζ , so because of the linearity discussed above, the entire equation is as well.

Appendix B. Measurement Techniques

In this appendix, measurement and data processing techniques are discussed in detail. Topics covered include data collection, data calibration, and transformation of frequency-domain data into the time domain.

A few words of caution are necessary about the conclusions drawn from measured data. Physical targets deviate from theory. Theoretical targets are perfect electrical conductors, with exactly zero resistance to the movement of current on their surfaces or interiors. Measured targets were composed of metal (or covered with metallic paint) with high, rather than perfect, conductivity. Also, target shapes may diverge from the ideal assumption, as can be seen for the sphere and the thin wire. The illumination also deviates from theory. Instead of a time-domain impulse, wide-band frequency-domain illumination is transformed into a non-causal, time-domain short pulse by the IFFT. This process involves arbitrary choices of zero padding and windowing that have some effect on the answer. Getting good results out of radar measurement ranges is difficult and involves both art and science, even for the much larger returns associated with early-time scattering. The late-time return is much weaker, making the acquisition of good data that much more difficult.

A variety of simple shapes were measured in various RCS measurement facilities. The targets were measured using a swept frequency transmitter and a coherent receiver. The result was the complex, frequency-domain scattering at each frequency in the bandwidth. Each complex data point had an in-phase component, formed by mixing the received signal with the incident wave, and a quadrature component, formed by mixing the incident wave with a signal 90 degrees out of phase with the incident wave, at the same frequency. In this way, a real measured signal can be used to generate complex “I and Q” data. Frequency-domain measurements were used to develop equivalent late-time signature data from each target at several azimuth angles. The targets used are discussed in Section B.1.

In each case, the target was measured in an anechoic chamber, which is an enclosed room in which the walls are lined with radar absorbing material. The target was mounted on top of a slanted pedestal of ogival cross-section, or in one case (at Kirtland AFB) hung from the ceiling via wires. The antenna was linearly polarized and was rotated to provide θ or ϕ polarizations. Figure B.1 above shows the basic components and data flow in a stepped-frequency radar system [85].




Figure B.1 Stepped-Frequency Radar System

Measurements were taken in the Air Force Research Laboratory (AFRL) RCS range at Wright-Patterson AFB. The incident frequencies ranged from 2.00 to 18.00 gigahertz, every 0.01 GHz. Each pulse was repeated 16384 times and coherently integrated to boost the signal to noise ratio. This level of integration could improve the signal to noise ratio of each measurement by 42 dB, assuming a perfect integrator. In a theoretically perfect integrator, signal strength is multiplied by the number of integrations, while the independent noise samples on each pulse sum incoherently and cancel each other out. Therefore, noise power does not increase, while the noise power is multiplied by the number of pulses integrated. In reality, coherent integration will provide a lesser, though still significant, advantage [59].

For each target, there was a target file (`.tar`) for each measured azimuth. There was also a target background file (`.bkg`) with only the target's styrofoam mount. Each measurement day there were two calibration (`.cal` file) measurements,

one at each polarization, involving a calibration target with known scattering. The calibration target was a squat cylinder, 4.5 inches in diameter and 2.1 inches in height. Background measurements (.cbk file) involving the calibration target's styrofoam mount were also collected.

The data was calibrated according to this formula

$$\bar{E}_{reduced}^{Fr}(\omega) = \frac{\bar{E}_{tar}^{Fr}(\omega) - \bar{E}_{bkg}^{Fr}(\omega)}{\bar{E}_{cal}^{Fr}(\omega) - \bar{E}_{cbk}^{Fr}(\omega)} \bar{E}_{exact}^{Fr}(\omega) \quad (\text{B.1})$$

The superscript F indicates the fields are expressed in the Fourier domain. The superscript r indicates that these are scattered fields. The subscripts on each electric field vector refer to the data file from which the field data was taken. Each field is a complex variable rendered in terms of real and imaginary, or “I and Q” components. These components refer to the field strength which is in phase (I) with a local oscillator or $\frac{\pi}{2}$ radians out of phase (in quadrature, Q) with the oscillator. Each electric field is considered a function of frequency rather than time because each is measured at a single frequency, without reference to a time variable.

If both the incident field data and the scattered field data are transformed into the time domain, the equivalent time-dependent incident $[\bar{E}_{reduced}^i(t)]$ and scattered $[\bar{E}_{reduced}^r(t)]$ fields are produced. In theory, the equivalent scattered fields will be the same ones found had the target been illuminated with the equivalent incident field and the time-domain response been measured. The Inverse Discrete Fourier Transform used to convert data into the time domain is discussed further in Appendix D.

B.1 Target Descriptions

Several targets were measured at the RCS measurement facilities of the Air Force Research Laboratory Sensors Directorate at Wright-Patterson Air Force Base, Ohio. One target was measured at the Air Force Research Laboratory Directed Energy Directorate at Kirtland Air Force Base, New Mexico.

Very small targets were used so that the wavelength of the available radar measurement range would match the lower-frequency modes of the targets. The bandwidth of the range was 2 to 18 GHz, so the wavelengths of the transmitted signals varied from 15.00 down to 1.67 centimeters. This technique is intended to be used against larger targets measured with longer wavelength (lower-frequency) illumination.

Table B.1 shows the dimensions of the targets measured at the AFRL/SNS range. The block and cube targets were measured with one face down on the pedestal. The azimuth reference was the point where one side was normal to the angle of incidence. The Block 123 target was measured twice, once with the 0.1999" by 0.6008" face toward the radar, the other with the 0.4003" by 0.6008" face toward the radar. In each case, The target was rotated from 0 to 90 degrees every 5 degrees. Measurements were taken with θ and ϕ polarization.

Target	diameter	length	width	height
Sphere	0.9995	-	-	-
Sphere	0.5000	-	-	-
Cylinder	0.2503	1.001	-	-
Cylinder	0.0629	1.270	-	-
Cylinder	0.0935	0.808	-	-
Cylinder	0.0937	1.196	-	-
Block 112	-	0.2506	0.2506	0.5016
Block 123	-	0.1999	0.4003	0.6008
Cube	-	0.4999	0.5003	0.5003

Table B.1 Targets measured with bandwidth 2-18 GHz, dimensions in inches

The measurement geometry is shown in Figure B.2. A cylinder target is shown to illustrate. Polarization is discussed in terms of the spherical coordinates θ , the angle away from the $\tilde{z}$ -axis, and ϕ , the angle around the $\tilde{z}$ -axis. Targets were set on a turntable in the measurement chamber. The $\tilde{z}$ -axis of the target was parallel to the turntable, so rotating the turntable varied the angle θ . Figure B.2 shows ϕ polarization, in which the electric field is perpendicular to the plane formed by the $\tilde{z}$ -axis and the incident field's direction of propagation, $\tilde{r}^i$ (the $\tilde{z}$ - $\tilde{r}^i$ plane, which was




Figure B.2 Target measurement geometry illustrating ϕ polarization

the scan plane for measured targets). In the case of θ polarization, the magnetic field is perpendicular to the scan plane.

Appendix C. Computational Techniques

C.1 Computational Results in the Literature

The Singularity Expansion Method (SEM) [18], explained in Section 2.2.1, has been applied to a limited number of targets so far. There are theoretical results for a sphere [6] and wire loops [20, 134], plus results based on the Method of Moments or other integral equations for a variety of other targets: polygonal wire loops [124], elliptical wire loops [114], wire dipoles [110, 110, 131], cylinders, rectangular plates, rectangular boxes, etc. Most of these targets were assumed to be isolated in free space. Some solutions exist for wires in the presence of ground planes [133], objects in or near lossy, dielectric media [12, 37, 42], coupled system consisting of two wires or two loops [43], etc.

There is no purely analytic description of the poles for objects other than the sphere or wire loop. The lack of further purely-analytic development is due to the difficulty in finding the exact Green's function needed to develop the impedance dyadic used in Equation (2.12). That difficulty is not unique to SEM. There are very few purely analytic results of narrowband scattering, or of early-time, ultra-wideband scattering. However, poles have been identified for various simple objects based on scattering simulations.

C.1.1 Cylinders and Cones. Current solutions on the surface of a body of revolution must have a dependence on $\exp(jm\phi)$, where m is an integer and ϕ is the angle of any position on the surface, in the plane of rotation, with respect to an arbitrary angle reference. This dependence is necessary because every $2m\pi$ rotation of an observation point around the surface of a rotationally symmetric body returns the observer to its starting location. Surface current and scattering are therefore composed of azimuthal modes [4, 105].

Merchant, Moser, Nagl and Uberall [86] computed the poles of right circular cylinders and spheroids. First they computed scattering using the Waterman T-matrix method, a technique applicable only to bodies of revolution. To compute poles, they computed the zeros of the determinant of the coupling matrix of the expansion coefficients. The poles from the $m = 0$ through $m = 4$ modes were reported.

Vechinski and Shumpert [137] also covered various bodies of revolution. They computed scattering via electric, magnetic, and combined field integral equation (EFIE, MFIE, CFIE) formulations of a body of revolution code. The code they employed found separate solutions for various azimuthal modes. The poles from the $m = 0$ and $m = 1$ modes were reported. The poles were found as the zeros of the mutual impedance matrix. Vechinski and Shumpert found discrepancies with some of the $m = 1$ mode poles found by Merchant, *et al* [86]. However, the poles used for comparison in this work were from the $m = 0$ mode, for which the two sets of authors agreed.

Brooks, Maier and Vechinski [25] provided pole estimates for a cone. Poles were found via a body of revolution, method of moments code that found separate solutions for the various rotational modes. The poles for the $m = 0$, $m = 1$, and $m = 2$ modes were displayed in a graph rather than a table, so some error may have been introduced during the process of reading the poles from the graph.

C.1.2 Thin Wire Dipoles. Various results exist for thin wire poles, using various techniques. Different results are compared for later comparison with measured results.

Tesche [131] found late-time scattering from thin wire dipoles based on the Method of Moments (MOM), but determined the poles using different techniques. He conducted frequency-domain analysis of the wire to find surface currents. He then found the time domain surface current and scattering via an Inverse Fast Fourier

Transform. The poles he estimated are shown in the first column of Table C.1.2. Long, Peng, Wen and Xie [67] also found surface currents via a MOM technique. They employed a Kuhn algorithm to the mutual coupling matrix developed for MOM [67]. The poles they find are similar to Tesche's. The oscillating frequency is related to the length of the wire, with harmonics present.

Sun, Chen, Nyquist, and Rothwell [130] analyzed rectangular plates of varying widths, using thin wires as a limiting case of very narrow plates. The poles found for a thin wire compare favorably to those found by Tesche [131] and Long [67], as shown in Table C.1.2. The target response was simulated using an electric field integral equation and exploiting target symmetry. Poles were found via a theorem relating the variation of the argument of a complex function integrated along a contour to the number of zeros and poles inside the contour. Only four poles are shown, all in the $\omega > 0$ region, although the paper indicates more exist. The existence of the conjugate poles of the four poles listed in Table C.1.2 should be inferred.

Tesche	Long, <i>et al</i>	Sun, <i>et al</i>
$-0.2576 \pm j2.9091$	$-0.2702 \pm j2.9405$	$-0.2571 \pm j2.7344$
$-0.3770 \pm j5.9596$	$-0.4178 \pm j6.0758$	$-0.4000 \pm j5.7813$
$-0.4618 \pm j8.9441$	$-0.5529 \pm j9.2551$	$-0.4857 \pm j8.8889$
$-0.5309 \pm j12.1077$	$-0.6440 \pm j12.4281$	$-0.5714 \pm j12.0635$

Table C.1 First Four Pole Pairs of Wire Dipole, sd/c ,

C.1.3 Rectangular Plates. Computed poles exist for rectangular plates of various widths, including square plates. The square plate results were compared to measured results for six-inch and twelve-inch square plates. Computational results were also compared for a six-inch square plate.

Sun, Chen, Nyquist, and Rothwell [130] analyzed rectangular plates. The target response was simulated using an electric field integral equation and exploiting target symmetry. Poles were found via a theorem relating the variation of the argument of a complex function integrated along a contour to the number of zeros and

poles inside the contour. Only four poles are shown, all in the $\omega > 0$ region, although the paper indicates more exist. The existence of the conjugate poles of the four poles shown here may be inferred. Frequencies are normalized as $sa/c = (\Omega + j\omega)a/c$. The data shows Ω increasing with increasing target width from the thin wire poles shown in Table C.1.2 to the square plate poles shown in Table C.2. The increasing Ω means the current dissipates more quickly. The authors did not propose an explanation for this phenomenon. The oscillating component ω changes little as the plate width widens.

The frequencies were not listed in a table, but only displayed on a graph. That means some error may have been added in the process of reading the graphs. Their normalized poles are shown in Table C.2.

Square Plate, sides= a sa/c
$-.8286 \pm j2.0313$
$-1.8857 \pm j5.0000$
$-1.6857 \pm j8.7302$
$-2.0000 \pm j11.9048$

Table C.2 First Four Computed Pole Pairs For Square Plate

C.2 Target Simulation

It was initially desired to simulate targets using a time domain code. As an alternative, frequency domain codes were considered. Time domain techniques include finite volume time domain (FVTD) [35] and finite difference time domain (FDTD) [105]. Techniques applicable to either the frequency or time domain include the method of moments (MoM) and Finite Element Method (FEM) [105]. To use a frequency domain code in a way comparable to the frequency-domain measurements, several hundred runs must be made per azimuth per target. Candidate codes were the TEMPUS [122] package (FVTD) by HyperComp, FISC (MoM) by SAIC-Demaco, the TSAR [82] package originally developed by Lawrence Livermore

National Laboratory, and method of moments codes called JRMBOR [111] and CARLOS [138]. Target simulation provided late-time signature data from each target at several azimuth angles.

Developing or modifying a computational code was outside the scope of this research. Only commercially available codes which were already hosted on computers available to AFIT students were considered. These include a PC network and various SGI and Sun workstations at AFIT. Supercomputers were also available at the Major Shared Resource Center (MSRC) computer facility run by the Aeronautical Systems Center (ASC) at Wright-Patterson AFB, Ohio. By connecting to the ASC MSRC, it was possible to gain access to the MSRC run by the Army Research Laboratory (ARL) at Aberdeen Proving Ground, Maryland.

The geometry for target illumination is shown in Figure C.1. Polarization direction is given in terms of the spherical coordinate variables θ or ϕ for scattering simulations. Here θ is the angle away from the $\tilde{z}$ axis, while ϕ is the angle around the $\tilde{z}$ axis, starting ($\phi = 0$) at the $\tilde{x}$ axis and increasing toward the $\tilde{y}$ axis. Polarization is the direction of the incident $\bar{E}$ field, which is perpendicular to both the direction of propagation and the $\bar{H}$ field in the case of plane waves. All computational targets had a defined set of axes, and the directions of illumination and polarization were defined with respect to the target's axes. Illumination directions were in arcs of increasing θ from 30 to 75 degrees, at either $\phi = 0$ (the $\tilde{x}$ - $\tilde{z}$ plane) or $\phi = 90$ (the $\tilde{y}$ - $\tilde{z}$ plane). Polarization was parallel to either the θ or ϕ vector at that direction of illumination.

It was originally desired to simulate targets using a finite-volume, time-domain code known as TEMPUS. Since the phenomenon of interest in this research was associated with time-domain scattering, a time-domain simulation appeared to be most appropriate. It would allow a direct approach to simulating the phenomenon, rather than combining multiple, independent, steady-state simulations at single frequencies and finding the time domain via IFFT. It was also thought that a time domain code




Figure C.1 Coordinates and geometry of illumination

would allow efficient use of computational resources, allowing a single run per aspect angle, per target. Some risk of inaccurate results was possible, since Miller states that time-domain models accumulate errors in late time [88].

Difficulties arose with the volume grid required around each target. Frequency-domain simulations characterize the surface of a conducting target as a connected set of facets, called a surface grid. For time-domain simulations, the space around the target must also be characterized, as a set of connected faceted shapes,, such as tetrahedrons or cubes, called a volume grid. Gridding targets for use with TEMPUS posed a series of technical difficulties, including the development and translation of target geometry files, and operation of various volume grid tools. The technical problems were not solved successfully. Continued development of the TEMPUS software package will make this an attractive option for future researchers.

Instead, two frequency domain codes, JRMBOR and CARLOS, were used. Both are described below. JRMBOR, CARLOS, and TEMPUS all compute the surface currents on the object. JRMBOR and CARLOS do that in the frequency domain, while TEMPUS computes currents in the time domain. TEMPUS then

converts the surface currents into the frequency domain. TEMPUS computes the far scattered field via a far field transform, which is a solution of the radiation integral in terms of a Fourier Transform. This transform operates between scatterer surface variables and far-field angle variables, rather than between time and frequency as in the more common Fourier transform application.

Cylinder, sphere and cone targets were simulated in the frequency domain using a code called JRMBOR. This code is limited to bodies of revolution such as circular cylinders, cones, ellipsoids, and coaxial combinations thereof. A sphere, three cylinders of varying length and radius, a cone, and one target which combined cylinders, were simulated. The tophat and cone targets are illustrated in Figures C.2 and C.3, respectively.

Target	radius	length
cylinder A	0.5	3.0
cylinder B	0.5	1.0
cylinder C	1.0	2.0
cone	1.0	1.9
tophat	2.25	2.1

Table C.3 Dimensions (inches) of targets simulated with CARLOS code

The output of JRMBOR was in a frequency domain format in magnitude (dB) and phase (degrees) with five digits of precision. The data was converted into real and imaginary (I and Q) components, then fed into the IFFT algorithm described in Section D.2.

Mie series sphere data was compared to data produced by a Method of Moments code called JRMBOR [111]. The reason for this comparison was to increase confidence that JRMBOR provides useful data. Mie series data was expected to agree much better with JRMBOR data than with measured data. In fact, the power magnitude found from the two data sources matched very well. However, there appeared to be a 90 degree shift on the phase plots, as shown in Figure C.4. This may have been due to an improperly set phase reference in a JRMBOR input file.




Figure C.2 Geometry of 'Tophat' Target




Figure C.3 Geometry of Cone Target




Figure C.4 Phase difference, JRMBOR (no phase correction) minus Mie, degrees

The phase error was corrected by simply multiplying the frequency-domain data by $\exp(-j\pi/2)$, which shifts the phase of the data by -90 degrees.

The time domain field was different in shape due to the phase shift in the original JRMBOR data. The frequency range of the data used to generate Figure C.5 was 0.1 to 18 GHz.

A 90 degree phase shift was applied to JRMBOR data on the single sphere. The result was an excellent match to the Mie series data. JRMBOR data for all targets was then phase shifted to ensure consistency. The sphere late-time signal developed from phase-corrected JRMBOR data matched late-time data based on the Mie series very well. Reconstruction error was 0.17 percent.

The CARLOS software package [138] was used to simulate targets that could not be simulated using JRMBOR. These included various rectangular objects such as cubes, blocks, and plates, and combinations thereof. CARLOS is a frequency-domain, method of moments code suitable for general shapes. The surface of each




Figure C.5 Sphere backscattered fields, time domain, Mie (dotted) and JRMBOR (solid) with no phase correction, zero range is center of sphere

target is divided into triangular facets in a grid. A grid generation program suitable for rectangular boxes was implemented. The grid generation program was used to produce a new grid for every frequency at which scattering was simulated. For the rectangular shapes, several simulation runs were conducted using various criteria for the minimum grid density in points per wavelength (ppw) along each edge. The minimum grid density was set to 12, 15, and 20 ppw. CARLOS solves the MoM computation via Galerkin's method [138], in which the same function is used for basis and testing [105].

Scattering from several targets was simulated with CARLOS. The targets are shown in Table C.4 with dimensions, except for the targets called Combo A and Combo B, which are described below. All of them were illuminated from 0.010 to 5.000 GHz every 0.002 GHz. They were illuminated in both the $\tilde{z} - \tilde{x}$ plane and the $\tilde{z} - \tilde{y}$ plane at θ angles from 30 to 75 degrees every 3 degrees.

Target	X	Y	Z
Cube	3.0	3.0	3.0
Block 112	1.0	1.0	2.0
Block 123	1.0	2.0	3.0
Block 123B	1.1	1.9	2.8
Plate	4.5	4.5	0.1
Combo A	-	-	-
Combo B	-	-	-

Table C.4 Dimensions (inches) of targets simulated with CARLOS code

Target Block 123B was designed to be similar to Block 123. The purpose of Block 123B was to study distinguishability of two targets with the same overall shape but slightly different measurements.

The targets Combo A and Combo B combined two objects separated by 0.5 inches. Combo A consisted of a Block 123 and a Block 112 oriented as shown in Figure C.6. Combo B consisted of a Plate and a Block 112 oriented as shown in Figure C.7.

The various grid densities resulted in slight differences in the signal in the frequency domain. These differences were centered at only a few wavelengths. In fact, at several frequencies, there was no difference between the 12, 15, and 20 ppw cases. Differences tended to be larger at lower frequencies, as shown in Figure C.8.

In the time domain, changes were apparent in the scattered field. Particularly, the late-time field is noticeably different for the 12, 15 and 20 ppw cases. Figure C.9 shows the late-time scattering developed from 12 ppw and 20 ppw data. In each case, the late-time signal behaved in accordance with SEM theory in that it could be represented well by a finite number of damped sinusoid terms and possibly some non-oscillating decay terms. In addition, the same or nearly the same pole locations could be estimated from various aspect-angle samples of the same target. But the poles for grid densities of 12 and 15 ppw were different from the 20 ppw case, which




Figure C.6 Geometry of 'Combo A' Target




Figure C.7 Geometry of 'Combo B' Target




Figure C.8 Real and imaginary components of frequency domain data from 3" cube, computed with grid densities of 12 ppw (solid line) and 20 ppw (dashed), focused on lower frequencies




Figure C.9 Time domain data from 3" cube, computed with grid densities of 12 ppw (solid line) and 20 ppw (dashed), focusing on late time

was taken to be the best answer because it was the most finely sampled simulation. The target signatures used in this research were developed from the 20 ppw targets.

Appendix D. Data Processing

D.1 The Inverse Discrete Fourier Transform

Scattering data was collected in the frequency domain. In other words, each data point represented the complex scattered field at a single frequency. However, the scattering phenomenon of interest is seen in time-domain scattering. It was therefore necessary to transform data from the time domain to the frequency domain using an Inverse Discrete Fourier Transform (IDFT), or a special case of the IDFT known as an Inverse Fast Fourier Transform (IFFT) [62, 100].

The IFFT of the reflected energy is known loosely as the impulse response. Here the term is used loosely because a true impulse has infinite bandwidth and zero duration, while the effective incident wave in this case merely has wide bandwidth and short practical time duration. While Inverse Laplace Transforms are discussed elsewhere in this document, the appropriate transform for this application is a variant of the Inverse Fourier Transform. Each measured data point is associated with real oscillating frequencies ω only, not complex frequencies $s = \Omega + j\omega$. The Inverse Fourier Transform and its discrete variant the IDFT assume frequencies are purely oscillating, so they are well-suited to the data. In addition, the IDFT is more amenable to implementation in software than an Inverse Laplace Transform. There are well known algorithms for the IDFT and IFFT, but there is no widely-used Inverse Discrete Laplace Transform.

An algorithm was implemented to perform an IFFT on the data and isolate the late-time response. The script generated negative frequency data, weighted it with a window function, added zero padding and shifted the negative frequency data to higher positive frequencies. It then transformed the data, computed the time and range scales for the time domain data, and isolated a portion of the late-time response.

Given the (padded) bandwidth $W = N\Delta t$, the total time and sample time of the signal resulting from the IFFT can be determined. The total time is T , found via

$$T = N\Delta t = \frac{N}{N\Delta f} = \frac{1}{\Delta f} \quad (\text{D.1})$$

and the sample time is

$$\Delta t = T/N = \frac{1}{N\Delta f} = 1/W \quad (\text{D.2})$$

After the IFFT, the data is in a new form representing the downrange image of the target as it would be seen by a radar transmitting an equivalent short pulse. The downrange image shows the field strength in volts versus the number of points. The number of points can be scaled and shifted appropriately to represent either time or distance downrange. The downrange image is complex, but only the real portion of the signal is used. That is because the equivalent signal transmitted by a short pulse radar (loosely, impulse radar) is a real signal. In a properly implemented IFFT, as discussed in Section D.2, the imaginary component of the data is much smaller than the real component, on the order of 300 dB down, so it can be ignored safely.

A script was developed to perform an IFFT on the data and isolate the late-time response. The script generated negative frequency data, weighted it with a window function, added zero padding and shifted the negative frequency data to higher positive frequencies as discussed in Section D.2. It then transformed the data with an existing IFFT function, computed the time and range scales for the time domain data, and isolated a portion of the late-time response.

$$n = [f_{high} - f_{low}]/\Delta t + 1 \quad (\text{D.3})$$

data points, prior to including negative frequencies and padding with zeros as explained in Section D.2. Instead of n points, there are now $N > n$ points. Since N

data points are present in the frequency domain signal, the time domain signal will include N data points as well, spaced Δt seconds apart.

Given the (padded) bandwidth $W = N\Delta t$, the total time and sample time of the signal resulting from the IFFT can be determined. The total time is T , found via

$$T = N\Delta t = \frac{N}{N\Delta f} = \frac{1}{\Delta f} \quad (\text{D.4})$$

and the sample time is

$$\Delta t = T/N = \frac{1}{N\Delta f} = 1/W \quad (\text{D.5})$$

D.2 Negative Frequency Data and the IDFT

For the IDFT (or its IFFT variant) to produce real data in the time domain, it must have data in complex-conjugate pairs in the frequency domain. The pairs consist of data points collected at the positive and negative of each frequency. The positive frequency data is that actually measured or simulated, while each negative frequency data point is the complex conjugate of the data at each corresponding positive frequency data point.

The necessary role of the negative frequencies can be seen in the development of the Discrete Fourier Transform (DFT), which is defined as

$$F_m = \sum_{n=0}^{N-1} f_n \exp(-j2\pi mn/N) \quad (\text{D.6})$$

The Inverse Discrete Fourier Transform (IDFT) is defined as

$$f_n = \frac{1}{N} \sum_{m=0}^{N-1} F_m \exp(j2\pi mn/N) \quad (\text{D.7})$$

Here note that $f_n = \frac{T}{N}f(n\Delta t)$ and $F_m = F(m\Delta\omega)$. The total time interval is T and the number of samples is N . For a sampled signal which is real in the

time-domain, f_n is a sequence of real numbers, and $F_{-m} = F_m^*$. Again, the negative frequency components are the conjugates of the positive frequency components.

If an IDFT is applied to complex data, both the negative and positive frequency components are needed to produce a real signal in the time domain. To reproduce the equivalent time-domain signal, it is only necessary to conjugate the data already measured. The output of the IDFT is real in theory, but numerically there is a very small imaginary component, about 300 dB down from the real signal, which can be ignored.

D.3 Frequency Data Windowing

A problem specific to measured data is that there is unlikely to be frequency data near the origin. For measurements taken from 2 to 18 GHz, there is conjugate data from -18 to -2, but there is a gap from just above -2 GHz to just under 2 GHz. This gap must be filled with zeros at data points spaced every Δf apart to ensure the frequency components are seen at the right frequencies by the IDFT algorithm.

If there is an abrupt transition from data to zeros, the time-domain result will be corrupted by Gibbs phenomenon [62]. To suppress Gibbs oscillations, various functions are used to weight the frequency-domain data prior to implementing the IFFT. These window functions taper the magnitude of the data from its peak in the center, down to zero or a small number at the edges. Window functions allow for a smooth transition from measured data to the zero-valued points used to pad the data. Various windows include the Hann, Blackman, Hamming, Kaiser, Bartlett, etc. There are some windows that are clearly inappropriate for use with an IFFT. A boxcar window, essentially no window at all, provides equal weight to each measured point and no taper, ensuring a harsh transition from data to zeros. The Hamming window tapers the data magnitude, but not all the way to zero. The end points of the window are on the order of 0.08, compared to a maximum value of 1. That makes the Hamming window less attractive than the Hann (often called ‘Hanning’)

window. The Kaiser window takes on various shapes based on the β parameter. For $\beta < 7$, the Kaiser window has high end points, like the Hamming window or higher, so it is also less attractive than the Hann window.

One preferred window must be selected. The window should eliminate abrupt transitions between data and zeros to suppress Gibbs phenomenon. The Gaussian window has an advantage over the Hann (Hanning) window because it transforms (via FFT or IFFT) to a Gaussian window in the opposite domain. That assumes a single window centered at the origin, which works very well when the measurement bandwidth runs close to DC ($\omega = 0$). Early in the research, it was assumed that the lower extent of the measured bandwidth had to be tapered as well, because measurements conducted in the Air Force Institute of Technology range ran from 6.2 to 18.2 GHz, leaving an enormous gap in the middle. The newer measurements leave a much smaller gap (-2 to 2 GHz), but it is still considerable.

The effect of the window in the time domain must also be considered. Each set of measurements was processed with one of three different frequency-domain windows. These were a boxcar (effectively no window), a Hann (Hanning) window, and a Gaussian window.

Sphere scattering data was generated in the frequency domain using the Mie series. The window configurations used were a single Gaussian centered at zero frequency, a pair of Gaussian windows centered at 10 GHz to cover the measured data and its conjugate at negative frequencies, and a pair of Hann windows. Twin frequency windows are important because there was a significant gap in data, from 0 to 2 GHz. Considering the negative frequencies, the gap was 4 GHz wide, from -2 to 2 GHz. The transition from data to zero padding would be a sharp one if a taper was not applied.

Use of twin Hann windows resulted in good looking time-domain sphere data. The single Gaussian window resulted in a very different looking sphere with a single pulse up front instead of a triplet, then a very low level in the middle, then a

singlet creeping wave return. Oddly, the poles estimated for the two time domain data sets were about the same. The key difference in the appearance of the time domain signals was the use of two windows instead of one centered at the origin. The lower frequency contribution will be reduced if twin windows are used, whether the windows are Gaussian, Hann, or Blackman.

The best result seen for the Mie series data from 2-18 GHz was with the twin Gaussian windows, shown in Figure D.1. Close to that was the result for twin Hann windows. The single Gaussian window, shown in Figure D.2, only got a few of the poles equal to the theoretical result. Most likely, the Gibbs phenomenon caused by the jump at ± 2 GHz is to blame. Some distortion was seen in the time domain with the 2-18 GHz Mie series and a single window.

Another set of sphere data was generated via the Mie series from 0.010 to 5.000 GHz, sampled every 0.002 GHz. With the much narrower frequency gap from -0.1 to 0.1 GHz, it was more appropriate to use a single window centered at the origin to taper the data. The poles estimated for this data when tapered with various windows are shown in Figure D.3.

The Blackman window formula is [21, 74]:

$$0.42 + 0.50 \cos(\pi k/K) + 0.08 \cos(2\pi k/K), \quad |k| < K, \quad 0 \text{ otherwise} \quad (\text{D.8})$$

The Hann window [62], named for Julius Von Hann but frequently called the Hanning window [21], is found via:

$$0.50 + 0.50 \cos(\pi k/K), \quad |k| < K, \quad 0 \text{ otherwise} \quad (\text{D.9})$$

The exponential window used in this comparison was given by:

$$\exp[-(5\pi/2) (k^2/K^2)], \quad |k| < K, \quad 0 \text{ otherwise} \quad (\text{D.10})$$




Figure D.1 Sphere time domain response calculated with twin Hann windows




Figure D.2 Sphere time domain response calculated with single Gaussian window




Figure D.3 Poles estimated from sphere late-time, Mie-series data from 0.010 to 5.000 GHz every 0.002 GHz. Various windows applied prior to IFFT: Hanning ($\cdot$), Gaussian ($\times$), Blackman ($+$). Theoretical sphere poles shown as circles.

Of the three windows studied, there was no clear advantage between the Blackman and Hanning windows. Either of those appeared to be a better choice than the Gaussian window used in the analysis. There are many possible implementations of the Gaussian window, but the problems associated with this implementation would also be present in other formulations. A single Hanning window was used for all simulated targets.

D.4 Data Smoothing

A technique was implemented to suppress noise in frequency domain measured data. The effect of frequency-domain data smoothing on the time-domain signal was examined as well as its effect on pole estimation.

Smoothing was employed on simulated and measured data. The simulated data frequently had small jump discontinuities. These occurred at some frequencies

where the grid density on the object's surface changed significantly. The quality of measured data from the AFRL/SNS range is ordinarily vary good, but since the targets measured were extremely small, the scattered signals were weak.

The script smoothed data by point averaging. Each component of each point was replaced by an average of its N nearest neighbors on either side, a total of $2N+1$ points. The size of N depended on the bandwidth and sampling interval, as well as the extent of the noise or step discontinuities apparent in the data. Measured data had a bandwidth of 2.000 to 18.000 GHz and a sampling interval of 0.010 GHz. For measured data, $N = 5$ so a band of ± 0.050 GHz around each frequency was used. Simulated data had a bandwidth of 0.010 to 5.000 GHz and a sampling interval of 0.002 GHz. For simulated data, $N = 15$ so a band of ± 0.030 GHz around each frequency was used.

At the high and low ends of the data, the algorithm used as many points as were available, retaining symmetry around the data point being smoothed. For example, the fourth data point could only be averaged with the three points below it, so only the three points above it were used. The average of those seven points (the current point was included) was written to a new file. An IFFT was performed on the smoothed data and the M-TLS-MPM algorithm was applied to the resulting late-time data. The poles from smoothed data had same ω but larger (more negative) Ω for most poles.

Figures D.4 and D.5 show the results of smoothing frequency domain data. The target in question was a cube with sides of 0.5 inches, illuminated with θ -polarized fields. The cube was oriented 25 degrees off of broadside. Two curves are shown in each figure. They are the real and imaginary components of the scattered field. Figure D.4 shows the entire bandwidth from 2 to 18 GHz to put the scale into perspective, while Figure D.5 shows a close-up on one part of the real component of the field to show more clearly the effectiveness of smoothing. The example chosen was typical for the small targets that were measured.




Figure D.4 Smoothed (dashed) and raw (solid) measured data, Frequency Domain, 2-18 GHz, real and imaginary components shown




Figure D.5 Smoothed (dashed) and raw (solid) measured data, Frequency Domain, 15-18 GHz, real component shown

A potential problem introduced by the point-smoothing approach was levelling of peaks. At the top of a peak in the frequency-domain signal, most or all neighboring points are below the central point being smoothed. The average of that point and its neighbors is then biased toward zero. To solve that, a curve-fitting algorithm was attempted. It was intended to fit quadratic curves to sections of the data, and replace the original data with the curves. A difficulty inherent in this method was forcing the end points of the curves to match. Curves were then allowed to overlap and the overlapping portions were averaged to produce a single smooth signal. The increased complexity of this approach, as compared to point smoothing, indicated that a significant amount of development work would have been required. Further work on this approach was abandoned due to time constraints.

The effect of smoothing on the time domain signal was also examined. Figure D.6 shows the time domain signals resulting from the smoothed and raw frequency domain data from the half-inch cube. The difference between the two signals was very small.

The poles present in the smoothed and unsmoothed data were estimated using the M-TLS-MPM algorithm discussed in Section 3.2. The results, shown in Figure D.7, indicated that many of the poles moved toward larger $|\Omega|$, indicating more rapid damping of the sinusoids. No poles were found above 60 gigradians per second, although the bandwidth extended up to 113 Grad/s. That is consistent with the results shown in Section 3.3.5, where the M-TLS-MPM algorithm was unable to estimate the higher- $|\omega|$ poles of the sphere when the significance criterion P was reduced. Higher- $|\omega|$ scattering modes tend to have less energy, thus are harder to estimate, than lower- $|\omega|$ modes in the sphere data.




Figure D.6 Smoothed (solid) and raw (dashed) measured data, Time Domain, half-inch cube




Figure D.7 Poles estimated from smoothed (+) and raw (x) measured data, half-inch cube, $P = 0.8$

Bibliography

1. Agee, F. J., Baum, C. E., Prather, W. D., Lehr, J. M., O'Loughlin, J. P., Burger, J. W., Shoenberg, J. S. H., Scholfield, D. W., Torres, R. J., Hull, J. P., and Gaudet, J. "Ultra-Wideband Transmitter Research," *IEEE Transactions on Plasma Science*, PS-26(3):860–873 (June 1998).
2. Alsing, S. G. *The Evaluation of Competing Classifiers*. PhD dissertation, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, March 2000.
3. Auton, J. R. and Blaricum, M. L. V. *Investigation of Procedures for Automatic Resonance Extraction From Noisy Transient Electromagnetic Data: Translation of Prony's Original Paper and Bibliography of Prony's Method*. Mathematics Notes, Note 79, Volume III, General Research Corporation, September 1981.
4. Balanis, C. A. *Advanced Engineering Electromagnetics*. New York: John Wiley and Sons, 1989.
5. Baram, Y. "Distance Measures for Stochastic Models," *International Journal of Control*, 33(1):149–157 (January 1981).
6. Baum, C. E. *On the Singularity Expansion Method for the Solution of Electromagnetic Interaction Problems*. Interaction Notes, Note 88, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, 11 December 1971.
7. Baum, C. E. *On the Singularity Expansion Method for the Case of First Order Poles*. Interaction Notes, Note 129, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, 14 October 1972.
8. Baum, C. E. "Emerging Technology for Transient and Broad-Band Analysis and Synthesis of Antennas and Scatterers," *Proceedings of the IEEE*, 64(11):1598–1616 (November 1976).
9. Baum, C. E. *Transient Electromagnetic Fields*, chapter 3. The Singularity Expansion Method, 129–177. Springer-Verlag, 1976.
10. Baum, C. E. *Electromagnetic Scattering*, chapter 15. Toward an Engineering Theory of Electromagnetic Scattering: The Singularity and Eigenmode Expansion Methods, 571–651. Academic Press, 1978.
11. Baum, C. E. *Representation of Surface Current Density and Far Scattering in EEM and SEM With Entire Functions*. Interaction Notes, Note 486, Kirtland Air Force Base, NM: Phillips Laboratory, 24 February 1992.
12. Baum, C. E. *The SEM Representation of Scattering From Perfectly Conducting Targets in Simple Lossy Media*. Interaction Notes, Note 492, Kirtland Air Force Base, NM: Phillips Laboratory, 21 April 1993.

13. Baum, C. E., editor. *Detection and Identification of Visually Obscured Targets*. Taylor and Francis, 1999.
14. Baum, C. E. "Direct Construction of a ξ -Pulse from Natural Frequencies and Evaluation of the Late-Time Residuals." *Ultra-Wideband, Short-Pulse Electromagnetics 4: Proceedings of the Fourth International Conference on Ultra-Wideband, Short-Pulse Electromagnetics, 1998*. 349–360. New York: Kluwer Academic / Plenum Publishers, 1999.
15. Baum, C. E. *An Observation Concerning the Entire Function in SEM Scattering*. Interaction Notes, Note 567, Kirtland Air Force Base, NM: Air Force Research Laboratory, 8 April 2001.
16. Baum, C. E., Farr, E. G., and Giri, D. V. *Review of Radio Science 1996-1999*, chapter Review of Impulse-Radiating Antennas, 403–439. Oxford, UK: Oxford University Press, 1999.
17. Baum, C. E. and Pearson, L. W. "On the Convergence and Numerical Sensitivity of the SEM Pole Series in Early-Time Scattering Response," *Electromagnetics*, 1(3):209–228 (1981).
18. Baum, C. E., Rothwell, E. J., Chen, K.-M., and Nyquist, D. P. "The Singularity Expansion Method and its Application to Target Identification," *Proceedings of the IEEE*, 79(10):1481–1492 (October 1991).
19. Benitz, G. R. "High Definition Vector Imaging," *Lincoln Laboratory Journal*, 10(2):1008–1015 (June 1997).
20. Blackburn, R. F. and Wilton, D. R. "Analysis and Synthesis of an Impedance-Loaded Loop Antenna Using the Singularity Expansion Method," *IEEE Transactions on Antennas and Propagation*, AP-26(1):136–140 (January 1978).
21. Blackman, R. B. and Tukey, J. W. *The Measurement of Power Spectra from the Point of View of Communications Engineering*. New York: Dover Publications, Inc., 1959.
22. Bombardt, J. N. J. and Spohn, D. J. *The Interior Electrostatic Field of a Semi-Infinite Cylindrical Cavity in a Grounded, Conducting Half-Space*. Interaction Notes, Note 122, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, May 1972. Previously Published as HDL-TM-72-12 by US Army Harry Diamond Laboratories.
23. Bowman, J. J., Senior, T. B. A., and Uslenghi, P. L. E., editors. *Electromagnetic and Acoustic Scattering by Simple Shapes* (Revised Edition). Hemisphere Publishing Corp, 1987.
24. Brooks, J. W. and Maier, M. W. "Object Classification by System Identification and Feature Extraction Methods Applied to Estimation of SEM Parameters." *Proceedings of the IEEE National Radar Conference*. 200–205. 1994.

25. Brooks, J. W., Maier, M. W., and Vechinski, S. R. "Applying System Identification and Neural Networks to the Efficient Discrimination of Unexploded Ordnance." *Proceedings of the IEEE Aerospace Conference 1997* 2. 449–467. 1997.
26. Bryson, A. E. and Ho, Y.-C. *Applied Optimal Control*. Waltham, MA: Blaisdell Publishing Company, 1969.
27. Clemmow, P. *The Plane Wave Spectrum Representation of Electromagnetic Fields*. Pergamon Press, 1966.
28. Coxeter, H. S. M. *Projective Geometry* (First Edition). New York: Blaisdell Publishing Company, 1964.
29. Cuomo, K. M., Piou, J. E., and Mayhan, J. T. "Ultrawide-band Coherent Processing," *IEEE Transactions on Antennas and Propagation*, AP-47(6):1094–1107 (June 1999).
30. Dorf, R. C. *Modern Control Systems* (Fourth Edition). Reading, MA: Addison-Wesley Publishing Company, 1986.
31. Emre, E. and Seo, J. "A Unifying Approach to Multitarget Tracking," *IEEE Transactions on Aerospace and Electronic Systems*, AES-25(4):520–527 (July 1989).
32. *FCC Rules and Regulations*. Title 47, Code of Federal Regulations, Part 15, Federal Communications Commission, May 30 2002.
33. Fisher, K. A. and Maybeck, P. S. "Multiple Model Adaptive Estimation with Filter Spawning," *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):755–768 (July 2002).
34. Fuller, D. F. and Andrew J. Terzuoli, J. "1-D Feature Extraction Using a Dispersive Scattering Center Parametric Model." *Digest of the 1998 IEEE/APS International Symposium*. 1296–1299. June 21–26 1998.
35. Gaitonde, D. and Shang, J. S. "Optimized Compact-Difference-Based Finite-Volume Schemes for Linear Wave Phenomena," *Journal of Computational Physics*, 138(2):617–643 (December 1997).
36. Gallego, A., Ruiz, D. P., and Carrion, M. C. "E-pulse scheme based on higher-order statistics for radar target discrimination in the presence of colored noise," *Electronics Letters*, 32(4):396 (15 February 1996).
37. Giri, D. V. and Tesche, F. M. "On the Use of Singularity Expansion Method for Analysis of Antennas in Conducting Media," *Electromagnetics*, 1(4):455–471 (1981).
38. Golub, G. H. and Loan, C. V. *Matrix Computations* (Second Edition). Baltimore, MD: The Johns Hopkins University Press, 1989.

39. Gustafson, J. A. and Maybeck, P. S. "Flexible Spacestructure Control via Moving-Bank Multiple Model Algorithms," *IEEE Transactions on Aerospace and Electronic Systems*, AES-30(4):750–757 (July 1994).
40. Hambley, A. R. *An Introduction to Communication Systems*. New York: Computer Science Press, 1990.
41. Hanlon, P. D. *Practical Implementation of Multiple Model Adaptive Estimation Using Neyman-Pearson Based Hypothesis Testing and Spectral Estimation Tools*. PhD dissertation, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, September 1996.
42. Hanson, G. W. and Baum, C. E. *Perturbation Formula for the Natural Frequencies of an Object in the Presence of a Layered Medium*. Interaction Notes, Note 532, Kirtland Air Force Base, NM: Phillips Laboratory, 8 October 1997.
43. Hanson, G. W. and Baum, C. E. "Asymptotic Analysis of the Natural System Modes of Coupled Bodies in the Large-Separation Low-Frequency Regime," *IEEE Transactions on Antennas and Propagation*, AP-47(1):101–111 (January 1999).
44. Hogg, R. V. and Craig, A. T. *Introduction to Mathematical Statistics* (Fifth Edition). Upper Saddle River, NJ: Prentice Hall, 1995.
45. Hohn, F. E. *Elementary Matrix Algebra*. New York: The MacMillan Company, 1958.
46. Horn, R. A. and Johnson, C. R. *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1985.
47. Hua, Y. and Sarkar, T. "On the Total Least-Squares Linear Prediction Method for Frequency Estimation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-38(12):2186–2189 (December 1990).
48. Hua, Y. and Sarkar, T. K. "Generalized Pencil-of-Function Method for Extracting Poles of an EM System from Its Transient Response," *IEEE Transactions on Antennas and Propagation*, AP-37(2):229–234 (February 1989).
49. Hua, Y. and Sarkar, T. K. "Matrix Pencil Method for Estimating Parameters of Exponentially Damped/Undamped Sinusoids in Noise," *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-38(5):814–824 (May 1990).
50. Hua, Y. and Sarkar, T. K. "A Perturbation Property of the TLS-LP Method," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-38(11):2004–2005 (November 1990).
51. Hua, Y. and Sarkar, T. K. "On SVD for Estimating Generalized Eigenvalues of Singular Matrix Pencil in Noise," *IEEE Transactions on Signal Processing*, SP-39(4):892–900 (April 1991).

52. Ilavarasan, P., Ross, J. E., Rothwell, E. J., Chen, K.-M., and Nyquist, D. P. "Performance of an Automated Radar Target Discrimination Scheme Using E Pulses and S Pulses," *IEEE Transactions on Antennas and Propagation*, AP-41(5):582–588 (May 1993).
53. Jain, V. K. "Filter Analysis by use of Pencil-of-Functions: Part I," *IEEE Transactions on Circuits and Systems*, CAS-21(3):574–579 (September 1974).
54. Jain, V. K., Sarkar, T. K., and Weiner, D. D. "Rational Modeling by Pencil-of-Function Method," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-31(3):564–573 (June 1983).
55. *Joint Publication 1-02, Department of Defense Dictionary of Military and Associated Terms*. The Joint Staff, 12 April 2001.
56. Kailath, T. "An Innovations Approach to Least Squares Estimation - Part I: Linear Filtering in Additive White Noise," *IEEE Transactions on Automatic Control*, AC-13(6):646–655 (December 1968).
57. Kailath, T. "The Innovations Approach to Detection and Estimation Theory," *Proceedings of the IEEE*, 58(5):680–695 (May 1970).
58. Kennaugh, E. M. "The K-Pulse Concept," *IEEE Transactions on Antennas and Propagation*, AP-35(3):327–331 (March 1981).
59. Knott, E. F., Shaeffer, J. F., and Tuley, M. T. *Radar Cross Section* (Second Edition). Boston: Artech House, 1993.
60. Kraus, J. D. *Electromagnetics* (Third Edition). New York: McGraw-Hill Book Company, 1984.
61. Kumaresan, R. and Tufts, D. W. "Estimating the Poles of Exponentially-Damped Sinusoids and Pole-Zero Modeling in Noise," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-30(6):833–840 (December 1982).
62. Lathi, B. P. *Signals and Systems*. Carmichael, CA: Berkeley-Cambridge Press, 1987.
63. Lehman, T. H. *A Statistical Theory of Electromagnetic Fields in Complex Cavities*. Interaction Notes, Note 494, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, May 1993.
64. LePage, W. R. *Complex Variables and the Laplace Transform for Engineers*. McGraw-Hill, 1961.
65. Lewis, F. L. and Syrmos, V. L. *Optimal Control* (Second Edition). New York: Wiley-Interscience, 1995.

66. Li, Y., Liu, K. J. R., and Razavilar, J. "A Parameter Estimation Scheme for Damped Sinusoidal Signals Based on Low-Rank Hankel Approximation," *IEEE Transactions on Signal Processing*, SP-45(2):481–486 (February 1997).
67. Long, Y., Peng, Z., Wen, X., and Xie, C. "Pivoting Method for Efficient Determination of Poles in Singularity Expansion Method," *Electronics Letters*, 28(13):1279–1280 (18th June 1992).
68. Long, Y. "Determination of the Natural Frequencies for Conducting Rectangular Boxes," *IEEE Transactions on Antennas and Propagation*, AP-42(7):1016–1021 (July 1994).
69. Lu, B., Wei, D., Evans, B. L., and Bovik, A. C. "Improved Matrix Pencil Methods." *Conference Record of the Thirty-Second Asilomar Conference on Signals, Systems and Computers, 1997*. 1433–1437. Piscataway, NJ: IEEE Press, Nov 1-4 1998.
70. Lund, E. J., Balchen, J. G., and Foss, B. A. "Multiple Model estimation with inter-residual distance feedback," *Modeling, Identification and Control*, 13(3):127–140 (July 1992).
71. Magill, D. T. "Optimal adaptive estimation of sampled stochastic processes," *IEEE Transactions on Automatic Control*, AC-10(4):434–439 (October 1965).
72. Mathews, J. and Walker, R. L. *Mathematical Methods of Physics*. Reading, MA: Addison-Wesley, 1971.
73. *MATLAB Function Reference* (Version 6 Edition), 1, 2, & 3. Natick, MA: The Mathworks, June 2001.
74. *Signal Processing Toolbox User's Guide For Use with MATLAB* (Version 6 Edition). Natick, MA: The Mathworks, December 2002.
75. Maybeck, P. S. *Combined State and Parameter Estimation for On-Line Applications*. PhD dissertation, Massachusetts Institute of Technology, Cambridge, MA, February 1972.
76. Maybeck, P. S. *Stochastic Models, Estimation, and Control, Volume 2*. Academic Press, 1979.
77. Maybeck, P. S. *Stochastic Models, Estimation, and Control, Volume 1*. Academic Press, 1979.
78. Maybeck, P. S., Floyd, R. M., and Moseley, A. "Synthesis and Performance Evaluation Tools for CGT/PI Advanced Digital Flight Control Systems." *Proceedings of the IEEE National Aerospace and Electronics Conference*. 1259–1266. Dayton, OH: IEEE Press, May 1983.
79. Maybeck, P. S. and Hanlon, P. D. "Performance Enhancement of a Multiple Model Adaptive Estimator," *IEEE Transactions on Aerospace and Electronic Systems*, AES-31(4):1240–1254 (October 1995).

80. Maybeck, P. S., Miller, W. G., and Howey, J. M. "Robustness Enhancement for LQG Digital Flight Controller Design." *Proceedings of the IEEE National Aerospace and Electronics Conference*. 518–525. Dayton, OH: IEEE Press, May 1984.
81. McGrath, D. T. and Baum, C. E. "Scanning and Impedance Properties of TEM Horn Arrays for Transient Radiation," *IEEE Transactions on Antennas and Propagation*, AP-47(3):469–473 (March 1999).
82. McLeod, R. R. *Temporal Scattering and Response Software Users' Manual, Version 2.1(R)*. Technical Report UCRL-MA-104861, Lawrence Livermore National Laboratory, September 1990.
83. McNamara, D. A., Pistorius, C. W., and Malherbe, J. *Introduction to The Uniform Geometric Theory of Diffraction*. Norwood, MA: Artech House, 1990.
84. Menke, T. E. and Maybeck, P. S. "Sensor/Actuator Failure Detection in the Vista F-16 by Multiple Model Adaptive Estimation," *IEEE Transactions on Aerospace and Electronic Systems*, AES-31(4):1218–1229 (October 1995).
85. Mensa, D. L. *High Resolution Radar Cross-Section Imaging*. Norwood, MA: Artech House, 1991.
86. Merchant, B. L., Moser, P. J., Nagl, A., and Uberall, H. "Singularity Expansion Representations of Fields and Currents in Transient Scattering," *IEEE Transactions on Antennas and Propagation*, AP-32(5):466–473 (May 1984).
87. Miller, E. K. "A Selective Survey of Computational Electromagnetics," *IEEE Transactions on Antennas and Propagation*, AP-36(9):1281–1305 (September 1988).
88. Miller, E. K. "Solving Bigger Problems - By Decreasing the Operations Count and Increasing the Computation Bandwidth," *Proceedings of the IEEE*, 79(10):1493–1504 (October 1991).
89. Miller, M. M. *Modified Multiple Model Adaptive Estimation (M^3AE) for simultaneous Parameter and State Estimation*. PhD dissertation, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, March 1998.
90. Mooney, J. E., Ding, Z., and Riggs, L. "Robust Target Identification Using a Generalized Likelihood Ratio Test." *Ultra-Wideband, Short-Pulse Electromagnetics 3: Proceedings of the Third International Conference on Ultra-Wideband, Short-Pulse Electromagnetics, 1996*. 343–350. New York: Plenum Press, 1997.
91. Mooney, J. E., Ding, Z., and Riggs, L. S. "Robust Target Identification in White Gaussian Noise for Ultra Wide-Band Radar Systems," *IEEE Transactions on Antennas and Propagation*, AP-46(12):1817–1823 (December 1998).

92. Mooney, J. E., Ding, Z., and Riggs, L. S. "Performance Analysis of an Automated E Pulse Target Discrimination Scheme," *IEEE Transactions on Antennas and Propagation*, AP-48(3) (March 2000).
93. Nag, S., Fluhler, H., and Barnes, M. "Preliminary Interferometric Images of Moving Targets obtained using a Time-Modulated Ultra-Wide Band Through-Wall Penetration Radar." *Proceedings of the 2001 IEEE Radar Conference*. 64–69. Piscataway, NJ: IEEE Press, 2001.
94. *Practical Aspects of Kalman Filtering Implementation*. AGARD Lecture Series 82, North Atlantic Treaty Organization Advisory Group for Aerospace Research and Development, May 1976.
95. Nilsson, J. W. *Electric Circuits* (Second Edition). Reading, MA: Addison-Wesley Publishing Company, 1986.
96. Nishiyama, K. "A Nonlinear Multiple Complex Sinusoid Estimator," *IEICE Transactions on Fundamentals*, E83-A(3):497–506 (March 2000).
97. Noble, B. and Daniel, J. W. *Applied Linear Algebra* (Third Edition). Englewood Cliffs, NJ: Prentice-Hall, 1988.
98. Osborne, M. R. and Smith, G. K. "A Modified Prony Algorithm for Exponential Function Fitting," *SIAM Journal on Scientific Computing*, 16(1):119–138 (1995).
99. P. D. Abramson, J. *Simultaneous Estimation of the State and Noise Statistics in Linear and Dynamic Systems*. PhD dissertation, Massachusetts Institute of Technology, Cambridge, MA, May 1968.
100. Papoulis, A. *Signal Analysis*. McGraw-Hill, 1977.
101. Papoulis, A. *Probability, Random Variables, and Stochastic Processes* (Third Edition). New York: McGraw-Hill, Inc, 1991.
102. Pascoe, K. J., Wood, W. D., Maybeck, P. S., and Wood, A. W. "Target Recognition Using Late Time Scattering." *Progress In Electromagnetics Research Symposium, 2002*. Cambridge, Massachusetts: The Electromagnetics Academy, July 2002.
103. Pascoe, K. J., Wood, W. D., Maybeck, P. S., and Wood, A. W. "Target Recognition Via Late Time Scattering." *Ultra-Wideband, Short-Pulse Electromagnetics 6: Proceedings of the Ultra-Wideband, Short-Pulse Electromagnetics 6, held at the American Electromagnetics 2002 Conference*, edited by Eric L. Mokole, et al. 463–470. New York: Kluwer Academic Press / Plenum Publishers, 2003.
104. Peebles, P. V. *Probability, Random Variables, and Random Signal Processes* (Third Edition). New York: McGraw-Hill, Inc, 1993.

105. Peterson, A. F., Ray, S. L., and Mittra, R. *Computational Methods for Electromagnetics*. Piscataway, NJ: IEEE Press, 1998.
106. Phillips, C. L. and Harbor, R. D. *Feedback Control Systems* (Second edition Edition). Prentice-Hall, 1991.
107. Potter, L. C., Chiang, D.-M., Carrière, R., and Gerry, M. J. "A GTD-Based Parametric Model for Radar Scattering," *IEEE Transactions on Antennas and Propagation*, AP-43(10):1058–1066 (October 1995).
108. Prather, W. D., Baum, C. E., Lehr, J. M., O'Loughlin, J. P., Tyo, S., Shoenberg, J. S. H., Torres, R. J., Tran, T. C., Scholfield, D. W., Gaudet, J., and Burger, J. W. "Ultra-Wideband Source and Antenna Research," *IEEE Transactions on Plasma Science*, PS-28(5):1624–1629 (October 2000).
109. Razavilar, J., Li, Y., and Liu, K. J. R. "Spectral Estimation Based on Structured Low Rank Matrix Pencil." *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing* 5. 2503 –2506. 1996.
110. Richards, M. A., Shumpert, T. H., and Riggs, L. S. "A Modal Radar Cross Section of Thin-Wire Targets via the Singularity Expansion Method," *IEEE Transactions on Antennas and Propagation*, AP-40(10):1256–1260 (October 1992).
111. Rogers, J. *JRMBOR: Computation of Scattering from a General Body of Revolution, Vol. 1: Formulation and MoM Algorithm, Vol. 2: Test Cases, Vol. 3: Code Listing*. Technical Report JR-2, Greenbelt MD: Atlantic Aerospace Corporation, 1990.
112. Rothwell, E. J. *Radar Target Discrimination Using the Extinction-Pulse Technique*. PhD dissertation, Michigan State University, East Lansing, MI, 1985.
113. Rothwell, E. J., Chen, K.-M., Nyquist, D. P., and Sun, W. "Frequency Domain E-Pulse Synthesis and Target Discrimination," *IEEE Transactions on Antennas and Propagation*, AP-35(4):426–434 (April 1987).
114. Rothwell, E. J. and Gharsallah, N. "Determination of the Natural Frequencies of a Thin Wire Elliptical Loop," *IEEE Transactions on Antennas and Propagation*, AP-23(6):777–781 (November 1975).
115. Roy, R. and Kailath, T. "ESPRIT - Estimation of Signal Parameters Via Rotational Invariance Techniques," *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-37(7):984–995 (July 1989).
116. Ruck, G. T., Barrick, D. E., Stewart, W. D., and Kirchbaum, C. K. *Radar Cross Section Handbook, 1 and 2*. Plenum Press, 1970.
117. Ruiz, D. P., Gallego, A., and Carrion, M. C. "Extinction Pulse and Resonance Annihilation Filter: Two Methods for Radar Target Discrimination," *Radio Science*, 34(1):93–102 (January-February 1999).

118. Ruiz, D. P., Gallego, A., and Carrion, M. C. "Scattered Signal Enhancement Algorithm Applied to Radar Target Discrimination Schemes," *IEICE Transactions on Communications*, E82-B(11):1858–1866 (November 1999).
119. Sarkar, T. K., Park, S., Koh, J., and Rao, S. M. "Application of the Matrix Pencil Method for Estimating the SEM (Singularity Expansion Method) Poles of Source-Free Transient Responses from Multiple Look Directions," *IEEE Transactions on Antennas and Propagation*, AP-48(4):612–618 (April 2000).
120. Sarkar, T. K. and Pereira, O. "Using the Matrix Pencil Method to Estimate the Parameters of a Sum of Complex Exponentials," *IEEE Antennas and Propagation Magazine*, 37(1):48–55 (February 1995).
121. Scharf, L. L. *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Reading, MA: Addison-Wesley Publishing Company, 1991.
122. Shankar, V., Rowell, C. M., Kabakian, A., and Hall, W. F. *Software User Manual for the TEMPUS CEM Solver Code Suite, Version 1.0*. Technical Report, HyPerComp, Inc., September 2000.
123. Shilov, G. E. *Elementary Real and Complex Analysis, 1*. The MIT Press, 1973.
124. Shumpert, T. H., Riggs, L. S., and Lindsey, J. M. "Singularity Expansion Method Analysis of Regular Polygonal Loops," *IEEE Transactions on Antennas and Propagation*, AP-38(8):1302–1306 (August 1990).
125. Silverman, R. A. *Introductory Complex Analysis*. New York: Dover Publications, Inc., 1972.
126. Spohn, D. J. *An integral Equation for the Interior Electromagnetic Field of a Semi-Infinite Cylindrical Waveguide in a Conducting Half-Space Illuminated by a Normally Incident Plane Wave*. Interaction Notes, Note 242, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, December 1974. Previously Published as HDL-TR-1678 by US Army Harry Diamond Laboratories.
127. Stakgold, I. *Green's Functions and Boundary Value problems* (Second Edition). New York: Wiley Interscience, 1998.
128. Stimson, G. W. *Introduction to Airborne Radar* (Second Edition). Mendham, NJ: SciTech Publishing, 1998.
129. Stratton, J. A. *Electromagnetic Theory*. McGraw Hill Book Company, 1941.
130. Sun, W., Chen, K.-M., Nyquist, D. P., and Rothwell, E. J. "Determination of the Natural Modes for a Rectangular Plate," *IEEE Transactions on Antennas and Propagation*, AP-38(5):643–652 (May 1990).
131. Tesche, F. M. "On the Analysis of Scattering and Antenna Problems Using the Singularity Expansion Method," *IEEE Transactions on Antennas and Propagation*, AP-21(1):53–62 (January 1973).

132. Tesche, F. M. "The Far-Field Response of a Step-Excited Linear Antenna Using SEM," *IEEE Transactions on Antennas and Propagation*, AP-23(11):834–848 (November 1975).
133. Umashankar, K. R., Shumpert, T. H., and Wilton, D. R. "Scattering by a Thin Wire Parallel to a Ground Plane Using the Singularity Expansion Method," *IEEE Transactions on Antennas and Propagation*, AP-23(2):178–184 (March 1975).
134. Umashankar, K. R. and Wilton, D. R. *Transient Characterization of Circular Loops Using Singularity Expansion Method*. Interaction Notes, Note 259, Kirtland Air Force Base, NM: Air Force Weapons Laboratory, August 1974.
135. Vasquez, J. R. *Detection of spoofing, jamming, or failure of a Global Positioning System*. MS thesis, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, December 1992.
136. Vasquez, J. R. *New Algorithms for Moving-Bank Multiple Model Adaptive Estimation*. PhD dissertation, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, May 1998.
137. Vechinski, S. R. and Shumpert, T. H. "Natural Resonances of Conducting Bodies of Revolution," *IEEE Transactions on Antennas and Propagation*, AP-38(7):1133–1136 (July 1990).
138. Volakis, J. L. "Carlos-3D; A General-Purpose three-dimensional method-of-moments scattering code," *IEEE Antennas and Propagation Magazine*, 35(2):69–71 (April 1993).
139. Watanabe, K. and Tzafestas, S. G. "A Hierarchical Multiple Model Adaptive Control of Discrete-Time Stochastic Systems for Sensor and Actuator Uncertainties," *Automatica*, 26(5):875–886 (1990).
140. White, N. A., Maybeck, P. S., and DeVilbiss, S. L. "MMAE Detection of Interference/Jamming and Spoofing in a DGPS-Aided Inertial System," *IEEE Transactions on Aerospace and Electronic Systems*, AES-34(4):1208–1217 (October 1998).
141. Wilkinson, J. H. *Rounding Errors in Algebraic Processes*. Englewood Cliffs, NJ: Prentice-Hall, 1963.
142. Willsky, A. S. "A Survey of Design Methods for Failure Detection in Dynamic Systems," *Automatica*, 12:601–611 (December 1976).
143. Willsky, A. S. and Jones, H. L. "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems," *IEEE Transactions on Automatic Control*, AC-21(1):108–112 (February 1976).

144. Wilton, D. R. “Review of Current Status and Trends in the Use of Integral Equations in Computational Electromagnetics,” *Electromagnetics*, 12(3-4):287–341 (1992).

Vita

Major Kenneth J. Pascoe grew up in Saint Clair Shores, Michigan. He graduated from South Lake High School in 1985. Following that he studied Electrical Engineering at Michigan Technological University, graduating Cum Laude in 1989. While at Michigan Tech, he was a member of Detachment 400 of the Air Force Reserve Officer Training Corps. Upon graduation, he was commissioned as an Air Force officer.

He served as an air weapons controller in the 726th Air Control Squadron at Homestead Air Force Base (AFB), Florida and Shaw AFB, South Carolina. He moved to Wright-Patterson AFB, Ohio, in 1993. While there he served as an emergency actions officer in the Command Center of Air Force Materiel Command and as a radar signature engineer in the National Air Intelligence Center. In 1996 he moved to Detachment 1 of the 31st Test and Evaluation Squadron, Kirtland AFB, New Mexico, where he planned and conducted operational tests, computer simulations, and evaluations of various Air Force weapon systems. In 1999 he returned to Wright-Patterson AFB and entered doctoral studies at the Air Force Institute of Technology. Upon completion of full-time studies, he returned to Kirtland AFB to conduct research in the High Power Microwaves Division of the Air Force Research Laboratory.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 074-0188	
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 the 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 a 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) 15-06-2004		2. REPORT TYPE Ph.D. Dissertation		3. DATES COVERED (From – To) Oct 2002 – May 2004	
4. TITLE AND SUBTITLE TARGET RECOGNITION USING LATE TIME RETURNS FROM ULTRA-WIDEBAND, SHORT-PULSE RADAR				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Kenneth J. Pascoe, Major, USAF				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 641 WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/DS/ENG/04-04	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/DEHP Mr. William Prather 3550 Aberdeen Ave SE Kirtland AFB NM 87117-5776 DSN: 246-0416				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) N/A	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The goal of this research is to develop algorithms that recognize targets by exploiting properties in the late-time resonance induced by ultra-wide band radar signals. A new variant of the Matrix Pencil Method algorithm is developed that identifies complex resonant frequencies present in the scattered signal. Kalman filters are developed to represent the dynamics of the signals scattered from several target types. The Multiple Model Adaptive Estimation algorithm uses the Kalman filters to recognize targets. The target recognition algorithm is shown to be successful in the presence of noise. The performance of the new algorithms is compared to that of previously published algorithms.					
15. SUBJECT TERMS Ultra Wide Band, Matrix Pencil, Multiple Model Adaptive Estimation, Kalman Filtering, Singularity Expansion Method					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
U	U	U	UU	316	Michael J. Havrilla, Ph.D. (ENG) (937) 255-3636, ext. 4582; e-mail: michael.havrilla@afit.edu