UNIVERSITY OF CALIFORNIA, SAN DIEGO

Effective Health Monitoring Strategies for Complex Structures

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Structural Engineering

by

Colin Michael Haynes

Committee in charge:

Professor Michael D. Todd, Chair
Professor William S. Hodgkiss
Professor Hyonny Kim
Professor Francesco Lanza di Scalea
Professor Daniel M. Tartakovsky

2014

Copyright

Colin Michael Haynes, 2014

All rights reserved.

The Dissertation of Colin Michael Haynes is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2014

TABLE OF CONTENTS

Signature Page iii

Table of Contents iv

List of Abbreviations vi

List of Symbols vii

List of Figures viii

List of Tables xi

Acknowledgements xii

Vita xv

Abstract of the Dissertation xvii

Chapter 1 Introduction 1

1.1. Structural Health Monitoring 1

1.2. Review of Existing Structural Health Monitoring Technologies 3

1.3. Introduction to Ultrasonic Guided Waves 5

1.4. Motivation for Statistical Approach 7

1.5. Data Cleansing 8

1.5.1. Baseline Subtraction 8

1.5.2. Environmental Compensation 9

1.5.3. Sparse Array Signal Processing 11

1.6. Contributions of the Dissertation Work 11

Chapter 2 Statistical Framework for Structural Health Monitoring 13

2.1. Introduction to Detection Theory 13

2.2. Theoretical Detector Derivation 15

2.3. Time-of-Flight Localization 19

2.4. Proposed Localization Paradigm 21

2.5. Maximum-Likelihood Estimation Procedure 23

2.6. Alternative Statistical Formulations 27

2.7. Sensor Fusion 34

Chapter 3 Experimental Validation 41

3.1. Bolted Frame Structure 41

3.1.1. Geometry 42
3.1.2. Data Acquisition 42
3.1.3. Test Protocol 44
3.2. Fuselage Rib Structure 45
3.2.1. Geometry 46
3.2.2. Data Acquisition 47
3.2.3. Test Protocol 47
3.3. Comparative Studies 48
3.3.1. Detection 48
3.3.2. Localization 55
3.3.3. Sensor Fusion 70

Chapter 4 Empirical Estimation of Wave Scattering 77
4.1. Bolt Bearing Damage Scattering Characterization 78
4.1.1. Experimental Methods 79
4.1.2. Scattering Results 86
4.1.3. Angular Dependence of Scattering 92
4.1.4. Frequency Dependence of Scattering 96
4.1.5. Load Dependence of Scattering 97
4.2. Rivet Scattering Characterization 101
4.2.1. Specimen Description & Data Acquisition 103
4.2.2. Polar Wavenumber Processing 104
4.2.3. Scattering Matrix Results 106
4.3. Low-order Feature Modeling 111
4.3.1. Beam spread function 112

Chapter 5 Supporting Topics for System Implementation 115
5.1. Uncertainty Quantification 115
5.1.1. Motivation for Uncertainty Quantification 115
5.1.2. Application Example 116
5.2. Bayesian Experimental Design & Optimal Sensor Placement 122

Chapter 6 Conclusions and Future Work 130
6.1. Conclusions 130
6.2. Future Work 131

References 133

LIST OF ABBREVIATIONS

BED	Bayesian Experimental Design
CFRP	Carbon Fiber Reinforced Polymer
FDR	Fisher Discriminant Ratio
FFT	Fast Fourier Transform
i.i.d.	Independent and Identically Distributed
MFC	Macro Fiber Composite
MLE	Maximum Likelihood Estimate/Estimation
NDE	Non-destructive Evaluation
NP	Neyman-Pearson (Criterion)
NRE	Normalized Residual Energy
OBS	Optimal Baseline Subtraction
PDF	Probability Density Function
PZT	Lead-zirconate-titanate
ROC	Receiver Operating Characteristic
SHM	Structural Health Monitoring
SNR	Signal to Noise Ratio
UGW	Ultrasonic Guided Wave
UQ	Uncertainty Quantification

LIST OF SYMBOLS

t, n	Time / index reserved for time samples
i, j	Indices reserved for actuating/receiving transducers
z	Index reserved for different trials
τ, η	Time of flight / index of time of flight
N	Total number (usually, of time points)
M	Total number of transducers
x	Baseline-subtracted, enveloped time series measurement
y	Baseline time series measurement
$\mathbf{v}$	Feature vector
$\mathbf{V}$	Matrix of feature values
$\mathrm{p}(\ )$	Probability density / likelihood function
r	Geometric coordinate
γ	Decision threshold
$\mathbf{\Gamma}$	Matrix of decision thresholds
μ	Mean value
σ	Rayleigh parameter / standard deviation
P^D	Probability of detection
P^{FA}	Probability of false alarm
f_s	Sampling frequency
$\mathbf{I}$	Localization image
w	Window function
k	Index reserved for damage modes
θ	True damage state
d	Decision state with respect to damage
e	Experiment / SHM system design
B	Beam spread function
$\mathrm{C}^{(\mathrm{d})}(d, \theta)$	Decision cost function
$\mathrm{C}^{(\mathrm{e})}(e)$	System design cost function
S	Guided wave scatter matrix

LIST OF FIGURES

Figure 1 – Example histogram for binary hypothesis testing illustration....14

Figure 2 - Flowchart illustrating proposed damage localization paradigm....22

Figure 3 - Side-by-side comparison of enveloped residual and MLE filter result....26

Figure 4 - MLE Rayleigh-step model truncation study....29

Figure 5 - Typical baseline and residual signals, divided into distinct sections: 1 – EMI, 2 – before direct arrival, 3 – before first arriving damage reflection, 4 – increasing residual energy, 5 – decreasing residual energy....30

Figure 6 - Rayleigh ramp model....31

Figure 7 - Non-constant Rayleigh model (may be proportional to baseline signal)....31

Figure 8 - Rician model, scale parameter on left and noncentrality parameter on right....32

Figure 9 - Comparison of performance of two MLE filters applied to a simulated signal....33

Figure 10 - Fusion hierarchy as applied to damage detection....35

Figure 11 - Bolted frame geometry overview....42

Figure 12 - Sensor layout in bolted frame structure....43

Figure 13 - Actuation signal for the bolted frame structure....44

Figure 14 - Magnet damage applied to the bolted frame structure....45

Figure 15 - Fuselage rib overview....46

Figure 16 - Simplified geometry of fuselage rib structure, with sensors and damage locations noted....46

Figure 17 - Crack damage introduced in fuselage rib structure....48

Figure 18 - Detection metric ROC for frame bolt damage....50

Figure 19 - Detection metric ROC for frame magnet damage....51

Figure 20 - Detection metric ROC for fuselage rib crack damage....51

Figure 21 - Detection metric ROC for fuselage rib hole damage....52

Figure 22 - Detection results in bolted frame for magnet damage (left) and bolt damage (right) 54

Figure 23 - Detection results in fuselage rib for hole damage (left) and crack damage (right) 55

Figure 24 - Localization performance curves for bolted frame structure 62

Figure 25 - Localization performance curves for fuselage rib structure 63

Figure 26 - Likelihood map for damage location prediction, frame trial 133. All dimensions are given in mm. 65

Figure 27 - Likelihood map for damage location prediction, frame trial 187. All dimensions are given in mm. 66

Figure 28 - Average localization error by region on the bolted frame structure. Areas with insufficient data are shown cross-hatched. All units are given in mm. 67

Figure 29 - Likelihood map for damage location prediction, fuselage rib trial 10 68

Figure 30 - Likelihood map for damage location prediction, fuselage rib trial 22 69

Figure 31 - Likelihood map for damage location prediction, fuselage rib trial 34 69

Figure 32 - Likelihood map for damage location prediction, fuselage rib trial 46 69

Figure 33 - Two forms of sensor fusion applied to the frame magnet damage: L1 norm (left) and voting (right) 70

Figure 34 - Overview of tensile testing apparatus, plate 101 80

Figure 35 - Overview of MFC array, plate 101 81

Figure 36 - Load vs. displacement curves for plate 101 84

Figure 37 - Load vs. displacement curve for plate 102 85

Figure 38 - Electron microscope photograph of bolt bearing damage cross-section 86

Figure 39 - Time series results for sensor pair 7-6, plate 102 87

Figure 40 - Time series result for sensor pair 2-6, plate 102 88

Figure 41 - Amplitude scattering matrix, actuator 6, plate 101 89

Figure 42 - Amplitude scattering matrix, actuator 6, plate 102 89

Figure 43 - Baseline-subtraction scattering matrix, actuator 6, plate 101 90

Figure 44 - Baseline-subtraction scattering matrix, actuator 6, plate 102 91

Figure 45 - Interrogation angle comparison for plate 101, 300 kHz 93

Figure 46 - Interrogation angle comparison for plate 102, 300 kHz 94

Figure 47 - Plate 102 unloaded results, damage level 3 only 95

Figure 48 - Class separation vs. actuation frequency, plate 102 97

Figure 49 – NRE metric computed for plate 102 trials 101

Figure 50 - Riveted specimen overview (laser spot is visible and appears larger because it is in motion) 103

Figure 51 - Detail of scan area, with rivet numbers and sensor location noted 104

Figure 52 - Flowchart illustrating wavenumber processing procedure 105

Figure 53 - Overview of wave propagation in riveted specimen 106

Figure 54 - Scattering matrix estimation for four example rivets 108

Figure 55 - Circularity of scattering profiles for each rivet 109

Figure 56 - Average correlation coefficient for each rivet with all other rivets 110

Figure 57 - Average time series for all MFC sensors separated by 135° 119

Figure 58 - R^2 values for signal amplitude, time of arrival, and NRE metric 121

Figure 59 - Optimal number of sensors, varying missed detection to false alarm cost ratios 125

Figure 60 - Four-sensor layouts for different missed detection costs. Red dots indicate sensor locations. 126

Figure 61 - ROC of optimal four-sensor configurations for magnet damage alone. The square markers indicate the point on those ROC curves that provided the optimal performance for a particular cost. 127

Figure 62 - ROCs for magnet and bolt damage, considered jointly 128

LIST OF TABLES

Table 1 - Description of features used in ROC comparison....49

Table 2 - Mean prediction error for the bolted frame structure....60

Table 3 - Mean prediction error for the fuselage rib structure....61

Table 4 - Normalized AUC, bolted frame localization performance curves....62

Table 5 - Normalized AUC, fuselage rib localization performance curves....63

Table 6 - Comparison of fusion techniques -- probability of detection for false alarm rate of zero....72

Table 7- Frame structure localization results for various fusion and modeling choices....75

Table 8 - Fuselage rib structure localization results for various fusion and modeling choices....75

Table 9 - Geometrical properties estimated from high-resolution scan....118

ACKNOWLEDGEMENTS

I'd like to begin by acknowledging my advisor, Professor Michael Todd. I could not have asked for a better mentor, and I owe him much for all of his guidance over the years. His leadership consistently challenged me to be a better engineer and provided the right environment for me to achieve my goals.

It has been my great privilege to be a member of the UCSD SHM research group for my graduate studies. The friendships, technical discussions, and overall comradery of this group have made the PhD process both more enjoyable and more productive, and for that I am very grateful. In particular, Eric Flynn, Scott Ouellette, Richard Do, Zhu Mao, Mike Yeager, Sydney Sroka, Joao Cheong all contributed directly to this dissertation, in ways small and great.

I'd also like to thank all of my family and friends who supported me throughout the PhD process. My father and mother, John and Donna Haynes, are the most important in this category. Their unwavering love, support, and encouragement at every step have been incredible blessings. Without them, I would not be where I am now.

I have also been fortunate to have a number of great international collaborations. From my time at the University of Bristol in the United Kingdom, I must thank Anthony Croxford and Paul Wilcox for their superb hospitality and mentorship. In the summer of 2012, I had the distinct pleasure of working with Professor Nobuo Takeda and his research group at the University of Tokyo. For that experience, my thanks go particularly to Takeaki Nadabe, all of the other members of the Smart Composites System Laboratory, and the Production Technology Research Group under Professor Yoji Okabe. In 2013, I spent the summer at Chonbuk National University in Jeonju, Korea. My hosts there were exceedingly gracious, and I thank Professor JR Lee, Andy Chong, Nitam Sunuwar, Jae Kyeong Jang, and the rest of the members of the Engineering Institute – Korea.

The research presented in this dissertation was primarily funded by the National Science Foundation, through a Graduate Research Fellowship and an East Asia & Pacific Summer Institutes Fellowship. The research was also supported by the Japanese Society for the Promotion of Science and the Agency for Defense Development in Korea via the Leading Foreign Research Institute Recruitment Program of the National Research Foundation of Korea.

A portion of Chapter 2 and a portion of Chapter 3 have been published in Structural Health Monitoring: An International Journal, Colin Haynes, Michael Todd, Eric Flynn, and Anthony Croxford, 2013. The title of this paper is "Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation". The dissertation author was the primary investigator and author of this paper.

Separate portions of Chapter 2 and Chapter 3 have been submitted for publication in Mechanical Systems and Signal Processing, Colin Haynes and Michael Todd. The title of this paper is "Enhanced Damage Localization for Complex Structures through Statistical Modeling and Sensor Fusion". The dissertation author was the primary investigator and author of this paper.

A portion of Chapter 4 has been published in Structural Health Monitoring: An International Journal, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2014. The title of this paper is "Monitoring of Bearing Failure in Composite Bolted Connections Using Ultrasonic Guided Waves: A Parametric Study". The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Key Engineering Materials, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2013. The title of this paper is "Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite

Bolted Connections." The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Proc. of the International Workshop on Structural Health Monitoring, Colin Haynes and Michael Todd, 2013. The title of this paper is "Effect of Applied Load on Guided Wave Monitoring of a Composite Bolted Joint". The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Proc. of the Asia-Pacific Workshop on Structural Health Monitoring, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2012. The title of this paper is "Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite Bolted Connections". The dissertation author was the primary investigator and author of this paper.

A portion of Chapter 5 has been published in Proc. of SPIE, Colin Haynes and Michael Todd, 2013. The title of this paper is "Uncertainty quantification of a guided wave structural health monitoring system for composite bolted joints". The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 5 has been published in Proc. of the European Workshop on Structural Health Monitoring, Colin Haynes and Michael Todd, 2012. The title of this paper is "Bayesian Experimental Design for Damage Detection in a Bolted Frame". The dissertation author was the primary investigator and author of this paper.

VITA

2009	B.S. in Structural Engineering	University of California, San Diego
2010	M.S. in Structural Engineering	University of California, San Diego
2014	Ph.D. in Structural Engineering	University of California, San Diego

PUBLICATIONS

*Work that has been incorporated into this dissertation

Peer Reviewed Journal Articles

*Haynes C, Todd M, Flynn E, Croxford A. "Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation." Structural Health Monitoring. 2013 Mar 1; 12(2):141–152.

*Haynes C, Nadabe T, Takeda N, Todd M. "Monitoring of Bearing Failure in Composite Bolted Connections Using Ultrasonic Guided Waves: A Parametric Study." Structural Health Monitoring. 2014 Jan 1; 13(1):94–105.

*Haynes C, Nadabe T, Takeda N, Todd M. "Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite Bolted Connections." Key Engineering Materials. 2013; 558:314–322.

Journal Articles in Progress

*Haynes C and Todd M. "Enhanced Damage Localization through Statistical Modeling and Sensor Fusion for Ultrasonic Inspection of Complex Structures." Mechanical Systems and Signal Processing.

Conference Proceedings

Haynes C, Yeager M, Todd M, Lee JR. "Monitoring Bolt Torque Levels through Signal Processing of Full-Field Ultrasonic Data." SPIE San Diego, CA: 2014.

*Haynes C, Todd M. "Effect of Applied Load on Guided Wave Monitoring of a Composite Bolted Joint." International Workshop on Structural Health Monitoring: 2013.

Do R, Haynes C, Todd M, Gregory W, Key C. “Efficient Detection Methods on a Composite Plate with Interior Embedded Fiber Optic Sensors via Impact Test.” International Workshop on Structural Health Monitoring, Stanford University: 2013.

*Haynes C, Todd M. “Uncertainty quantification of a guided wave structural health monitoring system for composite bolted joints.” SPIE San Diego, CA: 2013

*Haynes C, Nadabe T, Takeda N, Todd M. “Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite Bolted Connections.” Asia-Pacific Workshop on Structural Health Monitoring, 2012.

*Haynes C, Todd M. Bayesian Experimental Design for Damage Detection in a Bolted Frame. European Workshop on Structural Health Monitoring, Dresden, Germany: 2012

Haynes C, Todd M. Bayesian probabilistic modeling for damage assessment in a bolted frame. SPIE San Diego, CA: 2012

Haynes C, Todd M, Napolitano K. Uncertainty Quantification of Weighted Residual Method in Loads Estimation. Jacksonville, Florida: 2012.

Todd MD, Haynes CM, Flynn EB. (2011). "Bayesian Experimental Design Approach to Optimization in Structural Health Monitoring." Proc. World Congress on Advances in Structural Engineering and Mechanics. Seoul, Korea: 2011.

Haynes C, Todd M, Napolitano K, Matt H. Optimal Sensor Placement and Weighted Residual Method in Loads Estimation Supporting Structural Health Monitoring. International Workshop on Structural Health Monitoring: 2011.

Haynes C, Konchuba N, Park G, Farinholt K. Modeling, Estimation, and Monitoring of Force Transmission in Wind Turbines. Jackson: 2010.

FIELDS OF STUDY

Major Field: Structural Engineering

Studies in Structural Health Monitoring
Professors Michael Todd and Charles Farrar

Studies in Structural Dynamics
Professor Michael Todd

Studies in Signal Processing
Professor William Hodgkiss

Studies in Computational Mechanics
Professors David Benson

ABSTRACT OF THE DISSERTATION

Effective Health Monitoring Strategies for Complex Structures

by

Colin Michael Haynes

Doctor of Philosophy in Structural Engineering

University of California, San Diego, 2014

Professor Michael Todd, Chair

The aim of this dissertation is to develop structural health monitoring (SHM) strategies specifically suited for application to structures of significant geometric or material complexity. To accomplish this aim, a statistical framework is developed that characterizes and idealizes such environments and then finds optimal approaches from a detection theory standpoint. An optimal detector for discriminating damaged and undamaged cases in complex structures is derived. A maximum-likelihood approach to estimating the first arrival point is presented as part of a larger framework for damage localization. Novel sensor fusion approaches are also developed to account for the varying uncertainties associated with each sensor pair in the array.

Extensive empirical validation of the statistical framework is provided through experiments on multiple testbeds with diverse geometries, materials, and damage modes. The proposed approaches to damage detection and localization are proved to be effective through quantitative comparison with existing approaches from the literature.

Characterizing the wave scattering properties (also called scattering matrices) from geometric features is an important topic for guided wave SHM, whether the features are damage modes or not. Two different sensing methodologies are presented for constructing empirical estimates of the scattering matrices. Additionally, other topics of critical importance to SHM system design are addressed—namely, uncertainty quantification, Bayesian experimental design, and optimal sensor placement. The result is a significant step forward in advancing SHM toward structures of realistic complexity.

Chapter 1
Introduction

1.1. Structural Health Monitoring

Structural health monitoring (SHM) refers to the process of in-service embedded monitoring for effectively managing structural systems with respect to the potential occurrence of damage. The goal of a SHM system is to decrease the cost of sustaining safe operations by facilitating condition-based maintenance [1]. In other words, the principal objective is to direct decision-making based on the current health of the structure. While related, SHM systems are distinguished in concept from non-destructive evaluation (NDE) approaches in two major respects. First, SHM systems utilize embedded sensors to provide monitoring without the need to take the system offline for inspection. Second, SHM systems are focused on more autonomous operation, reducing or even eliminating the need for expert interpretation of results.

It is now common practice to present SHM in terms of a multi-step, statistical pattern recognition process. Viewed through such a lens, there exists a natural hierarchy of information about a structure's state that may be exploited in order to provide a meaningful diagnosis and

thereby make well-informed decisions about performance, maintenance, or operational fitness [2]. This hierarchy is summarized as follows.

1. **Damage detection** – The most basic SHM systems must provide some indication about whether the structure has become damaged. Because all materials have inherent defects, such an assessment must be based upon pre-established definitions of the target damage resulting from an operational evaluation of the system [3]. Damage detection is addressed theoretically in Section 2.2 and experimental validation is provided in Section 3.3.1.

2. **Damage localization** – Following the question of existence is the question of damage location. A novel approach to damage localization is presented in Section 2.4 and a comparative study of various localization techniques is given in Section 3.3.2

3. **Damage extent and damage type** – Questions regarding characterization of the damage are next addressed. To answer these questions requires the implementation of a model of the damage with respect to the inspection technique used. Chapter 4 addresses methods of damage characterization that are useful for accomplishing this step.

4. **Damage prognosis** – The ultimate goal for practical implementation of SHM systems is to use information about the damage to drive decision-making, perhaps through estimating failure probabilities or the structure's remaining useful life. While a tremendous incentive exists to develop such technology, solutions have proven elusive for most applications [4].

In order to answer the previous questions and satisfy the objective of the SHM system, a four-step process has been described by Farrar and Worden [1].

1. Operational evaluation – The economic and life-safety case for implementing SHM must be understood in as much detail as possible, so as to establish the parameters that will drive the system design. The operating environment, expected damage modes, and any limitations on data acquisition must be considered in this step.

2. Data acquisition – The sensing modality must be chosen along with all of the particulars of the data acquisition process. This step must also consider any necessary data cleansing and normalization that will help in isolating the damage-sensitive aspects of the acquired signals.

3. Feature selection – "Features" in the SHM literature are the metrics used to distinguish damaged and undamaged data. This dissertation will adopt a detection-theory-based approach to identifying appropriate features, and will refer to the process alternately as detector design.

4. Statistical model development – Once the features have been selected, statistical models of the damaged and undamaged distributions are necessary to provide estimates of the structural health with quantifiable confidence.

Although all steps must be considered, this dissertation is primarily focused on the last two steps in this outline – that is, with developing appropriate metrics and statistical models for enhancing damage detection, localization, and characterization.

1.2. Review of Existing Structural Health Monitoring Technologies

Due to the potential cost savings offered by its realization, there has been much interest in the field of SHM recently across many different application domains. Ensuring the continued serviceability of civil infrastructure (such as bridges and buildings) is a significant challenge

both in the United States and globally that has attracted attention to SHM [5, 6]. Pipeline inspection has been an active research field, driven partly by very significant interest in developing the technology for the oil and gas industry [7–9]. In the aerospace industry, several application areas have garnered significant interest. Rapid inspection of satellite structures for pre-launch verification has made use of SHM/NDE technologies [10, 11]. Chia *et al*. presented work toward developing smart hangar technology, where noncontact measurements would be used to inspect aircraft structures [12]. Advanced composite materials have generated substantial interest in the research community due primarily to their increased usage in both military and commercial aircraft – only a few examples from this extensive research field are cited for reference [13–15]. Next-generation marine vessels are also adopting composite materials and for similar reasons there exists a need for capable monitoring systems [16].

Many different sensing methodologies have been deployed for these various applications in order to generate the data necessary for SHM. Vibration-based SHM is a vast field based on monitoring the dynamic response of structures to a variety of inputs – a comprehensive review was carried out by Doebling *et al*. [17]. There are many methods based upon strain measurements, which can be sensed by conventional gages or other techniques [18]. Optical fibers (especially fiber Bragg grating strain sensors) have been of particular interest for strain-based SHM due to a number of inherent advantages, including survivability, immunity to electromagnetic interference, and ease of embedding into various structures [19].

There are also hosts of supporting technologies that are necessary for the final implementation of SHM. For larger structures, wireless sensor networks are very attractive and often necessary to avoid the complications encountered when using traditional cable connections [20]. Often, battery power is impractical for deployment over the entire design life of the structure, giving rise to the field of energy harvesting. By using various ambient energy

sources—of which the most common is mechanical vibration, but which may also include solar, thermal, or chemical sources—the power demands of the sensor nodes can often be met indefinitely [21, 22]. Damage prognosis, which has been recognized as a key step in moving SHM to practice, stands at the intersection of failure modeling, hazard analysis, and probabilistic risk assessment [4].

In this dissertation, the research has been driven less by a specific application than by a perceived need in all application areas to address objectives of SHM in a way that recognizes the inherent uncertainties in the process. That said, most of the structures and specimens used for experimental testing are either directly taken from aerospace structures or motivated by challenges in the aerospace industry. To solve the challenges that arise in these applications, this dissertation will focus on active sensing systems based on ultrasonic guided waves, which are described in the next section.

1.3. Introduction to Ultrasonic Guided Waves

Ultrasonic testing represents one of the most prevalent inspection techniques for NDE and SHM [23]. The term "ultrasonic" simply refers to the frequency range of the waves, which is higher than the range of human hearing. The work in this dissertation makes use of frequencies in the range of 50-500 kHz, although for other applications either higher or lower frequencies may be warranted.

The ultrasonic inspection process begins with some kind of transducer converting a voltage pulse into a mechanical stress wave. The transducer is usually a piezoelectric element, which transforms applied voltages to mechanical strain and vice versa. When bonded to a structure, the sensors are capable of both exciting ultrasonic waves into the host material and sensing incident waves. When the ultrasonic waves are guided by the boundaries of a one-

dimensional (1D) or two-dimensional (2D) propagation medium, they belong to a unique class known as ultrasonic guided waves (UGWs). Guided waves are formed by the resonance of waves in the material as they reflect off the boundaries. The particular class of guided wave that is of interest for many UGW SHM applications is that of Lamb waves, which propagate in plate-like structures [24, 25]. There are two classes of Lamb wave modes—symmetric modes and antisymmetric modes. Although the number of modes in each class is infinite, higher modes can only exist at higher excitation frequencies. Therefore, below a certain cut-off frequency-thickness product, only the fundamental antisymmetric (A_0) and fundamental symmetric (S_0) modes exist. All of the experiments considered in this dissertation are below the cut-off frequency-thicknesses product of their respective media, and therefore these are the only two modes generally discussed. Which mode is dominant for a particular application primarily relates to the transducer design, although the so-called "mode ratio" may be influenced by other factors as well.

Guided waves are of particular interest in SHM for a number of reasons. Among nondestructive testing methods, UGW inspection offers relatively large coverage areas per sensor with appropriate system design, particularly considering the corresponding sensitivity to small damage [26]. UGW inspection is often still feasible when access to the part is limited. The transducers used for generating UGW are easily embeddable and relatively inexpensive [27].

The propagation of UGW is governed by several parameters, most notably material properties (especially elastic moduli and density), the product of excitation frequency and plate thickness, temperature, material stresses, and boundary conditions. In general, the velocity of UGW is a frequency-dependent parameter, giving rise to the phenomenon known as dispersion. Because of the dispersive character of these waves, any wave with non-zero frequency

bandwidth (i.e., all real waves) will have a distinct group and a phase velocity. The group velocity is the speed at which the envelope of the wave packet propagates, while the phase velocity is the speed of the particles within the wave packet. Because of their dispersive character, guided wave packets tend to spread in time as they propagate in space, making signal interpretation much more challenging [28].

For monitoring of damage, the wavelength of interrogation is perhaps the most critical parameter. In order to maintain sensitivity to a particular form of damage, the wavelengths must be of the same order as the damage scale [3, 29]. The efficacy of UGW as an inspection technique is ultimately dependent on the fact that the wave scattering changes when damage is present. One of the most significant challenges associated with the UGW SHM approach is that all geometric irregularities (stiffener elements, material changes, boundaries, etc.) cause changes in the wave propagation. Most of these geometric features represent impedance mismatches that cause waves to scatter in the same way as the target damage, which further complicates the accurate inspection of geometrically-complex structures. Note that geometric or structural "features" should not be confused with signal "features" introduced previously, which in this dissertation will usually be referred to as "metrics" or "detectors" to avoid ambiguity.

1.4. Motivation for Statistical Approach

There are two fundamental approaches that may be taken to analyze the structural response to UGW inspection. The first is to construct a high-fidelity model of the wave propagation, beginning from the fundamental wave equation and taking into account all of the boundary conditions and properties of the medium necessary to accurately predict the expected response. Analytical solutions are generally difficult to obtain, but several authors have

obtained such solutions for very specific types of scatterers [30–32]. Where exact solutions prove intractable, finite element modeling and other related numerical methods have been employed by many authors to characterize wave propagation, especially the interaction of those waves with defects [29, 33–35].

The second approach is to adopt a statistical viewpoint, treating the signals as random processes and attempting to characterize them sufficiently to facilitate the SHM pattern recognition procedure. For particularly complex structures, analytical solutions can quickly become intractable and numerical simulations infeasible, whereas statistical methods tend to flourish where uncertainty is high. As a result, the statistical approach has been taken in this dissertation. In some areas, this strategy implies developing simplified approximate models to characterize complex phenomena [36]. Other authors have proposed a variety of methods that minimize the *a priori* information required for UGW SHM, in a similar spirit [37–40].

1.5. Data Cleansing

1.5.1. Baseline Subtraction

Because of the inherently comparative nature of the questions posed in the SHM process, it is necessary to ask how the signals produced by UGW sparse sensor arrays are changed by the presence of damage. It has already been mentioned that to produce damage-sensitive signals, UGW inspection relies on wave scattering from the damage (in much the same way as a sonar or radar system does). Because any change in mechanical impedance causes UGW to scatter, there seem to be few damage modes to which UGW are not sensitive in some way.

On the other hand, to distinguish damaged from undamaged signals, some comparison with a reference is essential. Baseline subtraction proposes to subtract from each newly-recorded signal a previous signal recorded in a known reference state. The result, known as the "residual" signal, may then be operated upon. In theory, if the structure and the environmental conditions have not changed, the new signal will be very similar to the baseline, resulting in a residual that is very close to zero. If there are significant reflections in the residual signal, they are interpreted as reflections from damage that was not present in the baseline state. Naturally, any real case will be subject to some noise processes that cause imprecise baseline subtraction and non-zero residuals which, without appropriate adjustment, can also be interpreted as damage reflections.

1.5.2. Environmental Compensation

In order to create a successful UGW SHM system, it is necessary to discriminate benign changes in the signals from those caused by damage. UGW are impacted by many factors, including changes in the structure, sensor properties, or bond properties. The structure may see changes in loading profile, boundary conditions, or the vibration environment, to name a few examples. Sensors are susceptible to load, temperature, and usually to electromagnetic fields. The bond condition may change over time with any parameter that influences epoxy, most notably humidity and temperature. For a SHM system to provide robust damage information over the life of the structure, all of these factors (and more) must be taken into account and procedures must be put in place to mitigate their effect.

Of all of the potential environmental changes that affect the signal without indicating damage, temperature effects are perhaps the most significant and most studied. Temperature affects both the group and phase velocity of UGW propagating in the structure, as well as changing the bond properties and the transducer properties themselves. As reported by

Konstantinidis *et al*., the velocity changes with temperature arise from two sources [40]. Thermal expansion changes the dimensions of the structure, including the thickness, density, and propagation distance, and the wave velocities change according to the temperature dependence of the elastic modulus of the material. Temperature can also degrade the piezoelectric properties of the sensors, depending on the particular material used for actuation, leading to a decrease in signal amplitude at high temperatures [41]. Changes in adhesive properties can also decrease the received amplitudes in a similar fashion.

Environmental compensation is not the focus of this dissertation, and all of the experiments were carried out in a lab setting where environmental parameters could be controlled. However, it is not a topic that can entirely be ignored when utilizing baseline subtraction techniques, which are especially sensitive to changes in wave velocity. Therefore, optimal baseline subtraction (OBS), as described by Croxford *et al*. [42], is applied in the majority of the tests considered here. The OBS procedure involves first saving a database of measurements distributed across the range of environmental conditions anticipated for the operation of the structure. The structure must be in a reference (presumed undamaged) condition as this set of baselines is acquired. For each test thereafter, each of the baselines is subtracted from the current test result, yielding the residual signal. The implicit assumption behind all SHM methods is that changes due to damage must be somehow orthogonal to the benign changes occurring due to changes in the operating environment—otherwise it would be impossible to identify damage. Baseline subtraction techniques assume, therefore, that the damage will produce changes in the signal that are not contained in the database of baseline measurements. The OBS procedure chooses the baseline that minimizes the norm of the residual, reasoning that such a minimization will choose the baseline taken under most similar environmental condition to the current test.

1.5.3. Sparse Array Signal Processing

In designing a sensor system to acquire UGW data from a structure, another decision point is encountered regarding the necessary density of instrumentation. In some applications, sensors are needed only to cover a very limited area and may be clustered very closely around particular "hotspots" where damage initiation is anticipated. Some systems are designed to have very dense arrays of sensors covering a wider area over which damage may appear. Still others may cluster sensors into nodes that are suitable for coherent beamforming of incoming waves, by which the direction of the incoming wave may be determined as well as the range [43]. The sensor system with the greatest ratio of coverage area to sensor count is the sparse array system, whereby each instrumented location consists of only a single transducer, and each pair of transducers is separated by a distance much larger than the wavelength of interrogation. Sparse array techniques can provide excellent coverage area per sensor, thus offering cost and weight savings over methods utilizing denser configurations [44]. The decision as to which type of system is optimal is ultimately based on the specific application. This dissertation will focus on the last class – sparse arrays.

In a sparse sensor array, the relative phase between sensors is highly uncertain. As part of the data cleansing process, after the OBS procedure is implemented, the data are therefore enveloped to remove phase information. The enveloping procedure involves taking the magnitude of the analytic signal, implemented by the Hilbert transform. As Croxford *et al*. noted, it is important to envelope after baseline subtraction to achieve the best results [44].

1.6. Contributions of the Dissertation Work

The significant contributions of this dissertation are summarized as follows:

1. A statistical framework for SHM is developed and demonstrated.

2. An optimal damage detector for sparse arrays and minimum *a priori* knowledge is derived and validated.
3. A novel sensor fusion approach for enhanced localization is developed and validated.
4. Scattering matrix characterizations for complex geometries and damage modes are studied, leading to the development of geometrical models for improved localization.
5. Critical supporting topics that will enable SHM to bridge the gap to practice are addressed in uncertainty quantification and Bayesian experimental design.

Chapter 2
Statistical Framework for Structural Health Monitoring

2.1. Introduction to Detection Theory

SHM is naturally framed in terms of a statistical pattern recognition framework. Therefore, it is natural to begin with an introduction to detection theory. For a given detection problem, the decision to be made is whether damage is present or not present. This dichotomy immediately suggests the application of a binary hypothesis test, and in fact this dissertation will borrow substantially from the field of detection theory in order to find the most effective SHM strategies [45].

Given the binary hypothesis test, there are four decision states that may occur: undamaged trials may be labeled undamaged (true negatives), undamaged trials labeled damaged (false positives or false alarms), damaged trials labeled undamaged (false negatives), and damaged trials labeled as damaged (true positives or correct detections). The probability of false alarm (P^{FA}) is then simply defined as the ratio of the number of false alarms to the total number of undamaged trials. Similarly, the probability of detection (P^{D}) is the ratio of correct detections to damaged trials. There is an inherent trade-off between the probability of detection and the probability of false alarm – both cannot be minimized simultaneously. Figure 1 shows a graphical illustration of the tradeoff, where the red distribution is the damaged class and the

blue distribution is undamaged. Usually, the detection problem is formulated as maximizing the probability of detection with a fixed acceptable probability of false alarm. This approach to optimal decision-making is known as the Neyman-Pearson (NP) criterion for damage detection [45].

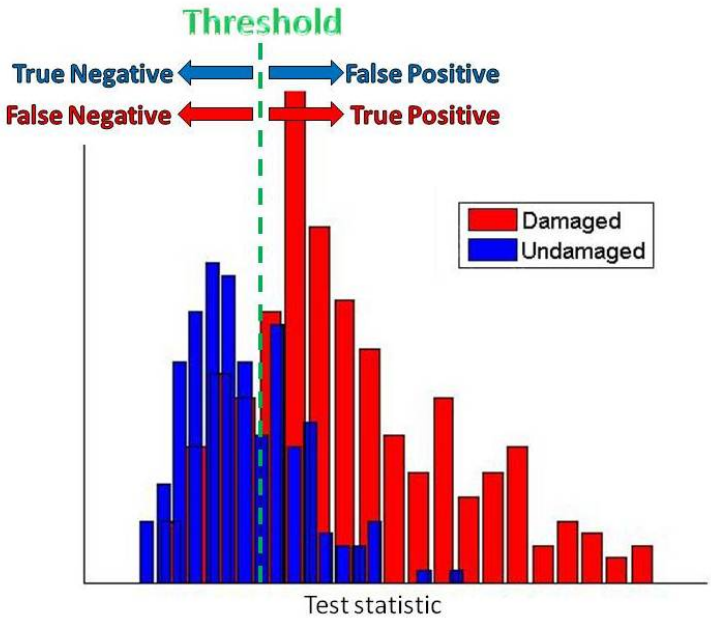

Figure 1 – Example histogram for binary hypothesis testing illustration

To provide quantifiable comparisons of different binary classifiers (or in the terminology of the SHM field, different features for damage detection), receiver operating characteristic (ROC) analysis is used. For a given binary detection problem the ROC compares the detection rate to the false alarm rate parametrically, independent of any specific underlying probabilistic structure [46]. To calculate the ROC curves, a threshold value is first chosen and each feature value is compared to that threshold. Then, the probability of detection and the probability of false alarm are calculated based on the resulting classifications. Varying the

threshold value from the minimum to the maximum feature values while repeating the preceding calculation at each point will generate the desired ROC curves.

Typically, ROC curves are plotted with probability of detection on the vertical axis and the probability of false alarm on the horizontal axis. A perfect detector would have a true positive rate of unity regardless of the false alarm rate; because this rarely occurs in practice, it is helpful to compare which detectors approach the performance of the perfect detector. Qualitatively, best performance corresponds to curves that are closer to the upper-left corner of the graph, with the 45-degree line representing the uninformed classifier (i.e. random guessing). For situations where a simple quantification of the detector performance (independent of threshold value) is desired, it is possible to reduce the ROC performance to a single scalar value. Perhaps the most common of these measures is the area under the curve (AUC), which as the name implies is simply the area under the ROC curve. Since ROC curves are plotted on the unit square, and the lower-bound, random-guessing AUC value is 0.5, the AUC should always be between 0.5 and 1.

2.2. Theoretical Detector Derivation

Beginning with the first portion of the SHM taxonomy, question of whether damage is present in the structure is addressed. In this dissertation, damage detection will be addressed in a global sense that is entirely separate from the subsequent question of damage location. That is, the entire sensor array for the system (however defined for a particular application) will render a decision of “damaged” or “undamaged” after each test. In order to derive an optimal detector using a classical detection theory approach, certain assumptions must be made about the signal and the underlying background noise process in which the signal appears. These assumptions are informed by experimental experience with sparse-array UGW SHM systems as

applied to the type of structures of interest for this dissertation. Rather than attempt to rigorously attempt to prove each of the assumptions at the outset, the derivation is completed first and the result is later validated empirically.

To begin, the data acquisition process is assumed to be linear, or near enough to linear that any nonlinearities do not have a substantial effect on the output signals. Next, the underlying noise process is assumed to produce Gaussian, independent, and identically distributed (i.i.d.) samples in time. After the first arrival (that is, the time when the first wave packet reflected from damage arrives) from any given sensor pair, the received signal is assumed to be composed of a complicated superposition of scattered waveforms, each of similar shape to the actuation pulse. According to the central limit theorem, the measured waveform may then also be modeled as a set of Gaussian i.i.d. samples in time. Note that the last assumption specializes the derivation to those structures where the arrivals of different wave packets are not clearly separable – that is, the structure and sensor system must be such that the signal is the superposition of many reflections beginning with the first arrival. Such structures are of great interest in moving SHM to practical application where the geometry is usually quite complex.

The first step after acquiring the raw data is to perform data cleansing by removing non-damage influences. The data are windowed in time to exclude data does not include damage information and to limit the quantity of data processed. Next, considering that the frequency content outside of the excitation bandwidth provides no information on the damage state (again, implicitly assuming linear wave propagation and scattering), the waveforms are match-filtered to the actuation signal. The OBS procedure, as described in Section 1.5.1, is applied to the filtered signals to yield residual signals. Since it is a simple subtraction, OBS is a linear process, so the output data are still Gaussian-distributed. Furthermore, according to the purpose

of the OBS procedure, the resulting residual should be noise only for undamaged records and the damage signal plus noise for damaged records, setting up the conditions for the binary hypothesis test.

The residual signals are next enveloped to remove phase by taking the magnitude of the Hilbert transform, as per the discussion in Section 1.5.3. The sum of the squares of independent, Gaussian-distributed components follows a Rayleigh distribution. The envelope values are, in fact, the sum of the analytic signal's real and imaginary parts squared, which are Gaussian i.i.d. variates. Therefore, the final envelope values are Rayleigh-distributed with Rayleigh parameter σ_0 for the H_0 (noise-only) hypothesis and σ_1 for the H_1 (signal + noise) hypothesis.

The optimal detector for choosing between the H_0 hypothesis and the H_1 hypothesis—that is, the detector that satisfies the NP criterion by maximizing the detection rate given a fixed false alarm rate—can be found through the likelihood ratio test [45]. The likelihood ratio is given by:

$$L(\mathbf{x}) = \frac{p(\mathbf{x}; H_1)}{p(\mathbf{x}; H_0)} > \gamma \tag{2.1}$$

where $\mathbf{x}$ represents the residual signal of length N data points, p is the probability density function (PDF) of $\mathbf{x}$, and γ is the decision threshold. The PDF of a Rayleigh distributed variable is given by:

$$p(x) = \frac{x}{\sigma^2} e^{-\frac{x^2}{2\sigma^2}}, x > 0 \tag{2.2}$$

Substituting, the likelihood ratio becomes:

$$L(\mathbf{x}) = \frac{\prod_{n=1}^{N} \frac{\mathbf{x}[n]}{\sigma_1^2} e^{-\frac{\mathbf{x}^2[n]}{2\sigma_1^2}}}{\prod_{n=1}^{N} \frac{\mathbf{x}[n]}{\sigma_0^2} e^{-\frac{\mathbf{x}^2[n]}{2\sigma_0^2}}} \tag{2.3}$$

$$L(\mathbf{x}) = \frac{\frac{\prod_{n=1}^{N} \mathbf{x}[n]}{\sigma_1^2} \exp\left(-\frac{1}{2\sigma_1^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right)}{\frac{\prod_{n=1}^{N} \mathbf{x}[n]}{\sigma_0^2} \exp\left(-\frac{1}{2\sigma_0^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right)} > \gamma \tag{2.4}$$

Taking the natural logarithm of both sides and rearranging:

$$L(\mathbf{x}) = \left(\ln\left(\frac{\prod_{n=1}^{N} \mathbf{x}[n]}{\sigma_1^2} \right) - \frac{1}{2\sigma_1^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right) - \left(\ln\left(\frac{\prod_{n=1}^{N} \mathbf{x}[n]}{\sigma_0^2} \right) - \frac{1}{2\sigma_0^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right) > \ln \gamma \tag{2.5}$$

$$L(\mathbf{x}) = \left(\sum_{n=1}^{N} \mathbf{x}[n] - \ln\left(\sigma_1^2\right) - \frac{1}{2\sigma_1^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right) - \left(\sum_{n=1}^{N} \mathbf{x}[n] - \ln\left(\sigma_0^2\right) - \frac{1}{2\sigma_0^2} \sum_{n=1}^{N} \mathbf{x}^2[n] \right) > \ln \gamma \tag{2.6}$$

Note that the first terms in each set of parentheses cancel out. Finally, incorporating all terms that do not depend on the data into a new threshold γ', the expression becomes:

$$\sum_{n=1}^{N} \mathbf{x}^2[n] > \gamma' \tag{2.7}$$

which is referred to as the energy detector [45]. Because imperfect baseline subtraction is usually the dominant source of noise in UGW SHM systems [44], the noise will have very

similar frequency content to the damage reflections the system must detect. Thus, it is reasonable that detecting an increase in energy in the residual signal is the optimal solution for the situation of sparse arrays in geometrically-complex structures. However, this proposition must still be evaluated experimentally to validate the assumptions made. Validation is provided in Section 3.3.1, where the data-driven method of ROC comparison will show that the energy detector outperforms a selection of other detectors from the literature.

Note that typically in application, it is advantageous to normalize the energy values for each sensor pair according to the energy contained in their baseline signals. While this does not affect the detection rate per se, it does make the detector output more consistent and easier to compare across sensor pairs and tests. Therefore, the final form of the detector that will be used for damage detection throughout this dissertation will be referred to as the normalized residual energy (NRE), and is given by the following expression:

$$\frac{\sum_{n} \mathbf{x}^2[n]}{\sum_{n} \mathbf{y}^2[n]} \tag{2.8}$$

where **y**[n] is the baseline signal.

2.3. Time-of-Flight Localization

When the presence of damage in a global sense has been thus established, it is natural to examine the question of damage location next. In considering UGW SHM systems, there are three basic approaches to localizing damage: tomographic methods, elliptical methods, and hyperbolic methods [47]. The first class of methods involves either scanning a transducer (whether contact or air-coupled) over the structure, or using laser excitation to scan through

each point of interest on the structure [48]. The last two classes are based on sparse array sensor systems. Methods that estimate the time of arrival of the wave packet scattered from the damage are known as elliptical methods, and these are perhaps the most common in the SHM literature [49, 50]. Hyperbolic methods are typically those that depend on cross-correlations of sensor responses, or time differences of arrival [51]. In this dissertation, methods from all three approaches are considered, with particular focus being placed on the ellipse method as the basis for the proposed localization paradigm for complex structures.

First-arrival localization bases the location estimate on the time-of-flight from the (known) actuation time until the time at which the first damage reflection is received. Knowing the locations of the sensors and the group velocity of the wave establishes the start point, end point, and total distance traveled. With that information, the location of the reflector can be narrowed to an ellipse with the two sensors at the foci and a semi-major axis corresponding to the distance the waves are predicted to travel in the estimated time to first arrival. Such procedures have become known in the SHM literature as elliptical methods.

It can easily be inferred that the effectiveness of the ellipse method is highly dependent on the estimation of the group velocity, the sensor locations, and the time of first arrival. The group velocity can sometimes be a difficult parameter to estimate, particularly due to its sensitivity to environmental conditions. However, these uncertainties have been very well-studied, as have the methods of estimating the group velocity. The ellipse method has even been extended to fully anisotropic plates [52]. While sensor locations are usually known with a relatively high degree of certainty, some authors have explicitly considered how to manage this uncertainty as well [37]. For the structures considered here, the most significant source of uncertainty in the ellipse method is the estimation of the first-arrival time. Because realistic structures tend to be more geometrically complex than simple, uniform plates, damage

reflections quickly become obscured by reflections from benign geometric features (so-called "secondary reflections") [53]. First-arrival estimation in such environments is addressed fully in Section 2.5.

However, the ellipse method is not used directly because with each sensor pair predicting only a line, the uncertainty of the prediction is not taken into account, and the problem of how to deal with ambiguous overlapping regions remains unaddressed. Some authors have proposed superimposing a second dimension on the ellipse to represent the uncertainty [54]. Instead, to create localization images in the present work, a delay-and-sum beamforming approach is used. In such an approach, an image of the structure is created where each pixel value represents each sensor pair's signal value (or a transformation of the signal values, as presented in Section 2.5) delayed according to the expected time-of-flight to that location and then summed over all sensor pairs. Mathematically, the expression is given by:

$$\mathbf{I}(r) = \sum_i \sum_j \mathbf{V}_{ij}\left(\tau_{ij}(r)\right) \tag{2.9}$$

where p represents a particular point on the structure, $\mathbf{I}$ represents the vector of pixel values, $\mathbf{V}$ represents the signals used for the delay-and-sum, and τ represents the time-of-flight to each point. The indices i and j denote the actuating and receiving sensors.

2.4. Proposed Localization Paradigm

Figure 2 shows a visual representation of the overall processing procedure used in the proposed localization paradigm. It begins with the data being collected on the left and proceeds until it concludes with a damage location estimate on the right. Each of the orange, double-lined nodes in the figure represents a step where one of various techniques may be chosen to

enhance the SHM system performance. The process centers on a weighted delay-and-sum beamforming step, taking as inputs the first-arrival estimation result and weights from the sensor fusion technique, beam spread function, feature models, and damage models. The beam spread function and feature modeling steps are based on the geometry of the system and are grouped accordingly – these are discussed in Section 4.3. The sensor fusion procedure, which can augment both damage detection and localization, is described in Section 2.7. Damage modeling—particularly empirical approaches to estimating the scattering properties of various defects—is discussed in Sections 4.1 and 4.2. Thus, this dissertation presents a comprehensive methodology for damage localization using UGW SHM systems on complex structures.

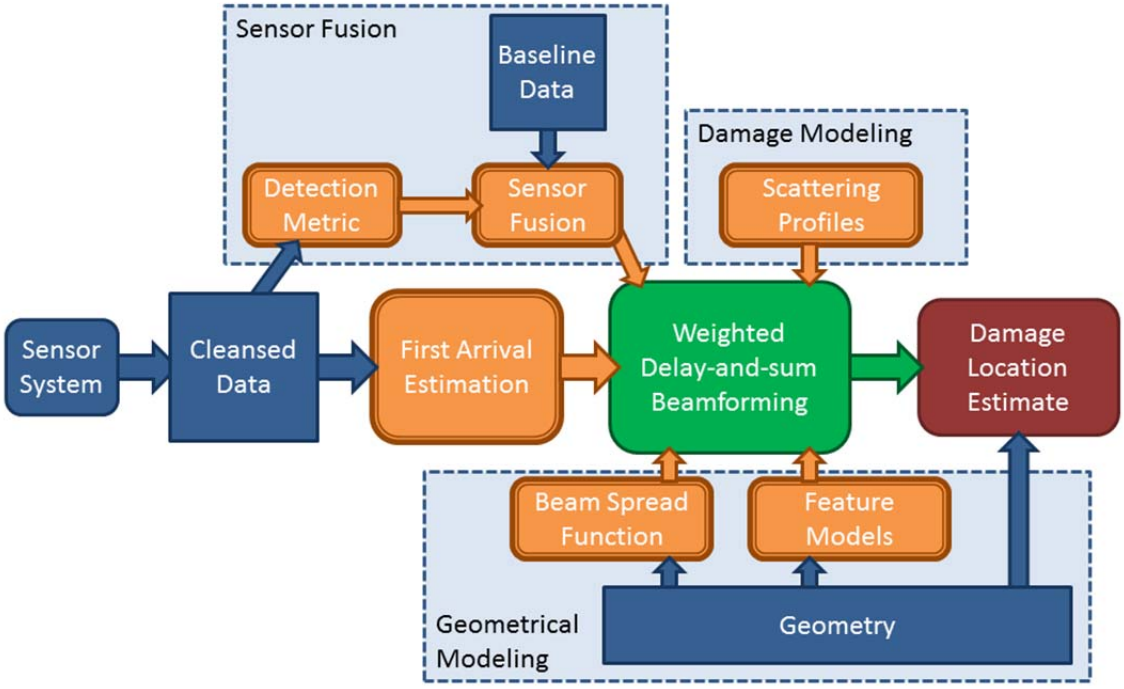

Figure 2 - Flowchart illustrating proposed damage localization paradigm

2.5. Maximum-Likelihood Estimation Procedure

Given that procedure that has been outlined for estimating damage location, the procedure for identifying the first-arrival point becomes crucially important. In simple applications, very straightforward methods may be applied by establishing time windows during which the first arrival must occur and amplitude thresholds based upon the expected noise floor. However, challenges arise when the secondary reflections become closer in time and amplitude to the first damage reflection, or when *a priori* knowledge is limited for established appropriate windows and thresholds. This situation is the case considered in this dissertation.

To implement first-arrival localization, this dissertation builds upon previous work by Flynn *et al*. [55]. Flynn proposed a maximum-likelihood estimation approach and proceeded to propose some assumptions about the signal properties of the residual signals encountered in UGW SHM systems – assumptions that correspond closely with the assumptions proposed for damage detection in Section 2.2. The key assumptions of this approach are as follows: the noise process of the baseline-subtracted residual signal may be described by a Gaussian-distributed random process. Before, the first-arriving damage reflection, only noise is present, and therefore the signal itself is modeled as Gaussian-distributed. After the first arrival, due to the complex, overlapping echoes of waves reflected from the damage, the central limit theorem is invoked to say the signal may also be described as a Gaussian-distributed random process whose samples are i.i.d. Enveloping the data requires summing the magnitude and phase of the signal in quadrature, resulting finally in two Raleigh-distributed random processes—one before the first arrival and one after the first arrival. The assumptions of the baseline-subtraction process imply that the Rayleigh parameter of the second section should be greater than the first section. The first-arrival point, therefore, can be estimated by finding the transition point between the two statistical regimes.

Computing the maximum likelihood estimate (MLE) of each of the two Rayleigh parameters, a likelihood ratio test can be implemented to describe the likelihood that each time point is the first arrival. The MLE of Rayleigh parameters is a well-known expression, given by taking the value of the Rayleigh parameter that maximizes the likelihood expression:

$$\hat{\sigma} = \arg\max_{\sigma}\left[L(\mathbf{x},\sigma)\right] = \arg\max_{\sigma}\left[\frac{\mathbf{x}}{\sigma^2}e^{-\frac{\mathbf{x}^2}{2\sigma^2}}\right] \tag{2.10}$$

Because the samples of the time series are independent, the total likelihood function may be expressed as the product of the individual likelihood functions. Taking the log of the likelihood (which is a standard procedure because the log is a monotonic function and therefore does not change the maximization result), the following expression is obtained:

$$\log\left[\prod_{n=1}^{N}\frac{\mathbf{x}[n]}{\sigma^2}e^{-\frac{\mathbf{x}^2[n]}{2\sigma^2}}\right] = \sum_{n=1}^{N}\left(\log\left[\mathbf{x}[n]\right] - 2\log\left[\sigma\right] - \frac{\mathbf{x}^2[n]}{2\sigma^2}\right) \tag{2.11}$$

Taking the derivative with respect to σ and setting equal to zero, the solution may be obtained as follows:

$$0 = -\frac{2N}{\hat{\sigma}} + \frac{\sum_{n=1}^{N}\mathbf{x}^2[n]}{\hat{\sigma}^3} \tag{2.12}$$

$$\hat{\sigma} = \sqrt{\frac{\sum_{n=1}^{N}\mathbf{x}^2[n]}{2N}} \tag{2.13}$$

Now, the likelihood ratio test can be formulated to provide the MLE of the arrival point in the signal. Assuming the arrival point to correspond to a time-of-flight of $n = \eta$, the total likelihood function of all the independent samples can be broken into two parts:

$$L(\mathbf{x}) = \prod_{n=1}^{\eta} \frac{\mathbf{x}[n]}{\sigma_1^2(\eta)} e^{-\frac{\mathbf{x}^2[n]}{2\sigma_1^2(\eta)}} \prod_{\eta+1}^{N} \frac{\mathbf{x}[n]}{\sigma_2^2(\eta)} e^{-\frac{\mathbf{x}^2[n]}{2\sigma_2^2(\eta)}} \tag{2.14}$$

Note that each Rayleigh parameter is a function of the time-of-flight index η. This likelihood ratio should not be confused with the one given by Equation (2.4), which was comparing undamaged and damaged hypotheses.

$$\begin{aligned} L(\mathbf{x}) = &\sum_{n=1}^{\eta} \left(\log\left[\mathbf{x}[n]\right] - 2\log\left[\sigma_1(\eta)\right] - \frac{\mathbf{x}^2[n]}{2\sigma_1^2(\eta)} \right) \\ &+ \sum_{n=\eta+1}^{N} \left(\log\left[\mathbf{x}[n]\right] - 2\log\left[\sigma_2(\eta)\right] - \frac{\mathbf{x}^2[n]}{2\sigma_2^2(\eta)} \right) \end{aligned} \tag{2.15}$$

Grouping the terms and substituting the MLE estimates of σ_1 and σ_2, the final expression is obtained as follows:

$$\begin{aligned} L(\mathbf{x}) = &2\sum_{n} \log\left[\mathbf{x}[n]\right] - 2(\eta - 1)\log\left[\hat{\sigma}_1[n]\right] \\ &-2(N - \eta)\log\left[\hat{\sigma}_2[n]\right] - N \end{aligned} \tag{2.16}$$

Note that all terms that do not depend on η may be dropped from the final form, because they do not affect the maximum or therefore the resulting location estimate. Because the maximum of the resulting vector is the predicted arrival point, the output of this filter may then be used in a delay-and sum procedure to produce a predicted damage location in the

structure. An example residual signal (taken from the bolted frame structure) and the corresponding MLE filter result are given in Figure 3. The green dotted line indicates the maximum of the MLE result, which is the first-arrival estimate according to this procedure. Despite the complex multitude of secondary reflections, the process correctly identifies the first reflected wave packet.

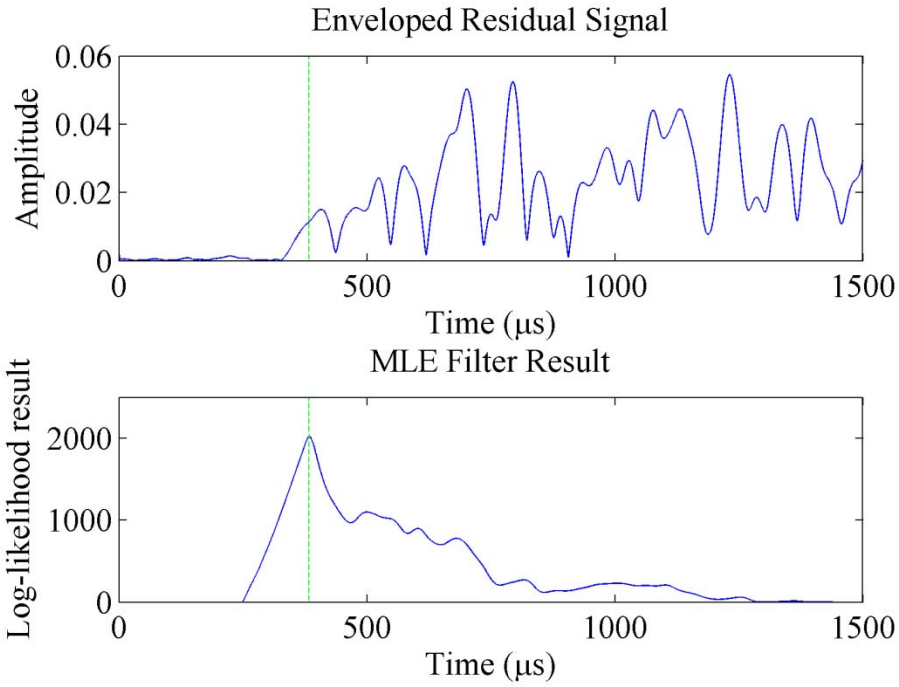

Figure 3 - Side-by-side comparison of enveloped residual and MLE filter result

There are some other key assumptions that go into the first-arrival estimation procedure that should be noted here. As is common to delay-and-sum formulations, this approach assumes single-mode propagation. That is, one mode is dominantly actuated and received by the sensor array, and this mode has a known group velocity. This dissertation is also concerned only with point-like scatterers assumed to occur at one location. While it is possible that this approach

could be extended to multiple damage locations or even possibly to larger areas of damage, such damage modes are outside the present scope.

Another fundamental premise of the estimation is that the signal is broken into two sections– before and after the first arrival of a significant damage reflection. Thus, the method inherently assumes that sensor pairs at arbitrary locations are actually able to observe damage reflections—that is, the amplitude of the first-arriving wave (and the subsequent reflections) should be greater than the noise floor in a statistically-significant sense. However, the impact of sensor pairs not observing the damage can be mitigated with proper application of sensor fusion principles. Also note that this derivation differs slightly from the original formulation in that it does not attempt to integrate the information from different sensor pairs. Flynn *et al*. asserted that the results from each sensor pair could also be considered independent likelihood functions and that the resulting log-likelihood estimates should be summed together accordingly (hence the delay-and-sum technique was incorporated from the start) [55]. Here, however, the MLE estimation procedure is considered to be a kind of nonlinear filter that may be applied to each signal. This viewpoint allows flexibility to scale the contributions of each sensor pair to the delay-and-sum procedure through sensor fusion without weakening the validity of the imaging scheme in any way. Sensor fusion turns out to be quite important in structures of significant complexity, as will be demonstrated in Section 3.3.2.

2.6. Alternative Statistical Formulations

While this dissertation will demonstrate the effectiveness of the MLE procedure for damage localization, it may be possible to improve performance even further by examining the underlying assumptions more closely. Because the MLE first-arrival estimation procedure plays a large part in the localization framework, it is important to understand the input variables

that have a significant influence over the filter result. Modeling the statistics of UGW residual signals as characterized by pre- and post-arrival Rayleigh distributions with a sharp transition (a "Rayleigh step" model) is one possible choice – to discover other plausible formulations, the statistical regimes of the signal must be characterized more closely. However, it should be noted at the outset that despite providing valuable insight into the foundational assumptions of the MLE formulation, a better-performing formulation was not obtained, and the first-arrival estimation presented in Section 2.5 is used in the remainder of this dissertation.

All of the comparison methods to be presented in Section 3.3.2 make use of time windowing of various kinds to limit the impact of the secondary reflections. Therefore, it is a fair question to ask how sensitive the MLE formulation is to changes in the time windowing. Figure 4 shows an example of the same ultrasonic signal truncated to a number of points between 2000 and 6000 (corresponding to 800 and 2400 μs). Circles highlight the maxima, which are chosen as the estimated first-arrival points. (Again, recall that vertical shifts do not affect the MLE formulation.) The figure shows that MLE is sensitive to the signal length, but that the filter shape varies as smoothly as can be expected. Moreover, the resulting estimation selects one of three points which are the most likely separation depending on the length of the signal that is retained. When used in the delay-and-sum architecture, the resulting estimation is less sensitive to the time window chosen than most of the competing methods.

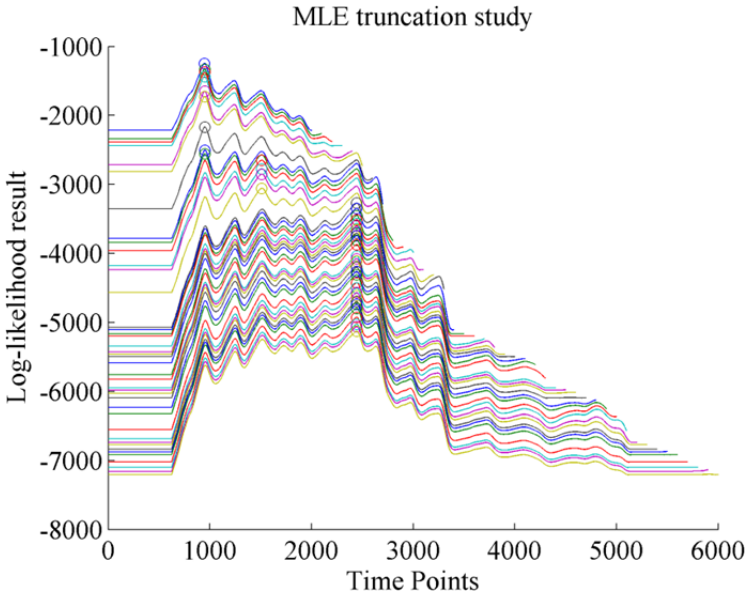

Figure 4 - MLE Rayleigh-step model truncation study

Figure 5 shows a representative residual signal from a damaged trial, broken down into five sections. The first regime is composed of the pre-actuation time (if any is recorded) and the actuation pulse that may cause electro-magnetic interference. Then follows an interval where the direct arrival has not reached the receiving sensor. Next, depending on the damage location, there may be an interval after the direct arrival but before the first arrival of the damage reflected wave. This third regime is still in comprehended in the "noise-only" definition for damage detection, but typically has a higher residual energy level than the time before the first direct arrival. Finally, there is the regime following the first damage reflection, after which the complex superposition of secondary damage reflections occurs. At first, the magnitude of these reflections tends to increase –the first-arriving wave packet is usually not the highest amplitude. The increase may be related to increasingly imprecise baseline subtraction as the reflections create a very different interference pattern. Referring to the example of the baseline signal, however, it is noted that the amplitude of the two subtracting signals begins to decrease rapidly

as the echoes begin to die out. Therefore, a fifth regime may be observed where the residual signal again tends to decrease toward the end of the record, if the signal is not truncated before this transition occurs.

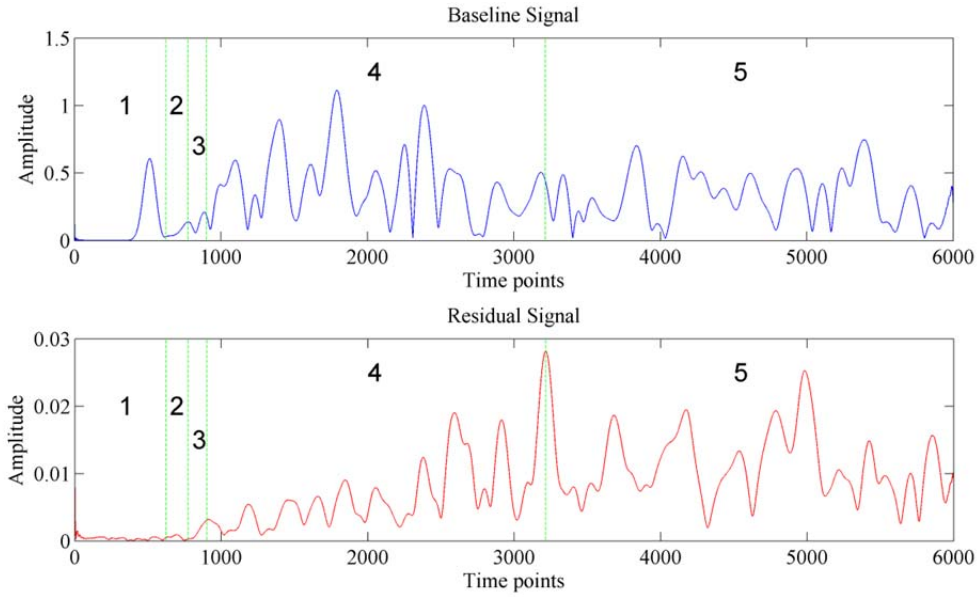

Figure 5 - Typical baseline and residual signals, divided into distinct sections: 1 – EMI, 2 – before direct arrival, 3 – before first arriving damage reflection, 4 – increasing residual energy, 5 – decreasing residual energy

In seeking to more accurately model the statistics of the residual signal for UGW SHM systems, three alternative formulations were developed:

1. Rayleigh ramp – Rather than simply having a step in the Rayleigh parameter at the first-arrival point, this model proposes that the Rayleigh parameter after the first arrival should grow in time. The concept is to better approximate signals where the true first arrival is much smaller in amplitude than the subsequent echoes. The growth in time, for simplicity, is simply taken to be linear. The model is described graphically in Figure 6.

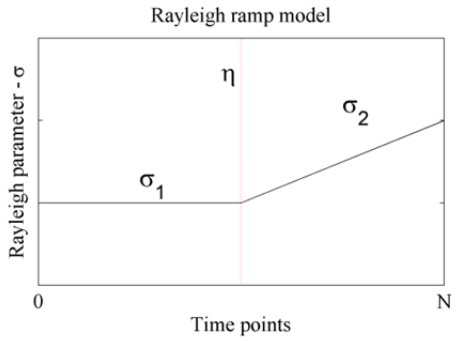

Figure 6 - Rayleigh ramp model

2. Baseline-proportional Rayleigh – This model assumes that the Rayleigh parameter is variable throughout the signal and is proportional to the baseline signal amplitude, which is intended to represent a noise process associated with imperfect baseline subtraction. The proportionality constant is modeled as having a step increase at the first-arrival point. The trend of Rayleigh parameter varies from signal to signal, but an illustration is given in Figure 7.

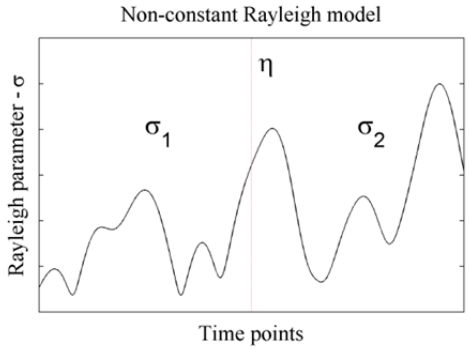

Figure 7 - Non-constant Rayleigh model (may be proportional to baseline signal)

3. Rician – Whereas the standard MLE approach is based on a Rayleigh distributed signal, a Rician model may also be used. In communications theory, a Rician fading channel is distinguished from a Rayleigh fading channel in that there is a dominant (or direct) arrival surrounded by lesser echoes. The Rician distribution therefore has an additional

parameter, known as the noncentrality parameter, that expresses the amplitude relationship between the direct and indirect arrivals. When the noncentrality parameter is zero, the distribution becomes a Raleigh distribution. It was thought that by using such a model, the additional parameter could be related to the expected arrival amplitude (generated from some simple models of the structure) to improve performance. The model is illustrated in Figure 8.

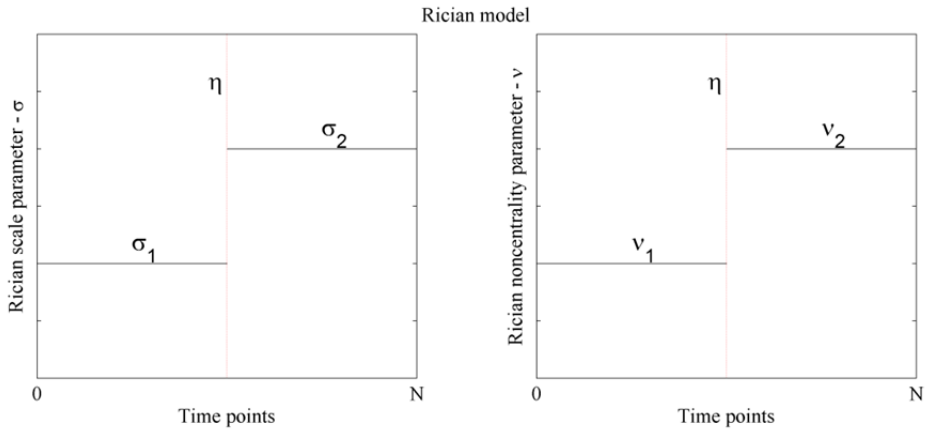

Figure 8 - Rician model; scale parameter on left and noncentrality parameter on right

The effectiveness of each of the statistical formulations has been evaluated, beginning with synthetic data generated based on the statistical assumptions, as well as using sample signals collected from UGW SHM systems. Testing with synthetic data involved using random number generation to build signals that matched exactly the assumptions (e.g., one Rayleigh parameter pre-arrival and a linearly-increasing Rayleigh parameter post-arrival). Because the arrival point is specified and known, the Monte Carlo simulations of the prediction accuracy can be generated for each of the filters. An example of such a test is given in Figure 9. In this figure, the known arrival point is designated by the red dotted line, and the maxima of both the Rayleigh ramp and Rayleigh step models are shown. The Rayleigh ramp exhibits a tendency to

grow again after peaking, as well as producing erratic values toward the beginning of the signal, both of which degraded its performance. The Rayleigh step is shown as a reference, and although in this particular case the predicted points are very close to one another, after Monte Carlo simulation, the Rayleigh step proved to be more effective. Filter performance was also compared qualitatively by application to test data where the damage location is similarly known in advance.

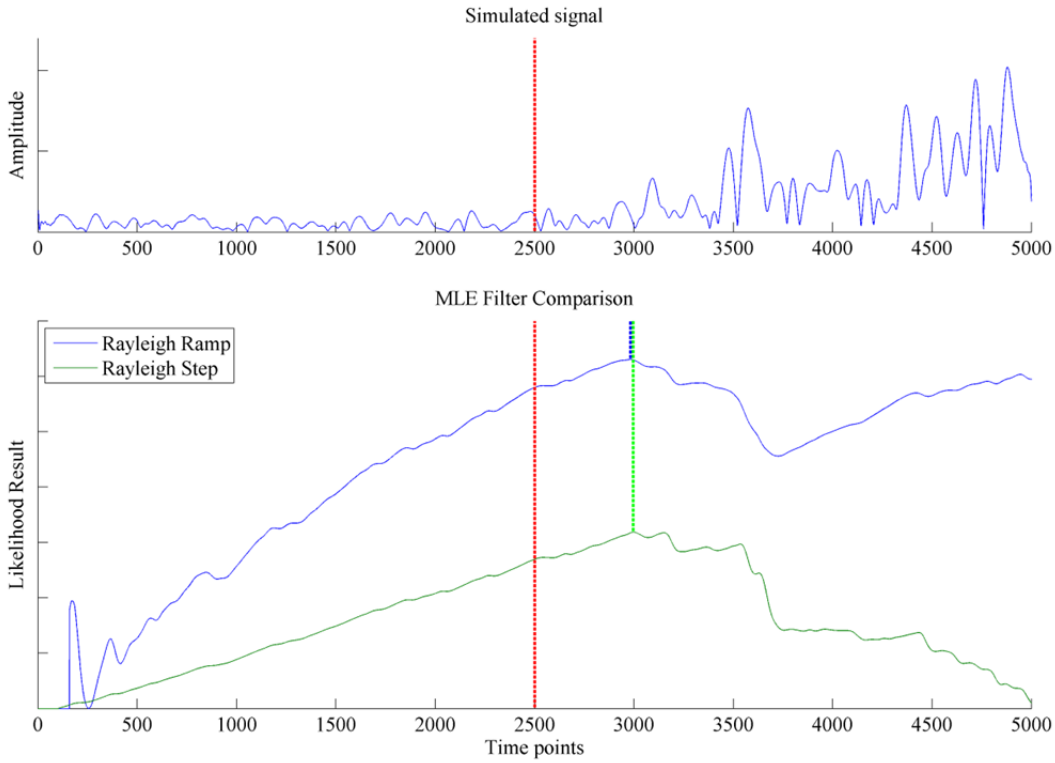

Figure 9 - Comparison of performance of two MLE filters applied to a simulated signal

The alternative formulations fail to outperform the previous MLE filtering procedure for signals that are fair approximations of those observed in UGW testing. Detailed analysis of the reasons for this failure has not been rigorously established, but it seems that the simplicity of the Rayleigh-step assumption lends robustness which produces a better result. Through

qualitative analysis, this section has examined some of the properties of baseline-subtracted residual signals and the MLE formulations that may be used to model them. It is left to future work to explore the formulations proposed here or new ones that may be able to further enhance first-arrival estimation accuracy.

2.7. Sensor Fusion

The (often) large volumes of data obtained by SHM systems must in every case be reduced, ultimately, to a decision about the health of the structure. Almost always, the SHM system must combine data from different sources, and these sources of data are rarely completely uniform. However, methods for how to deal with data sources that are not commensurate (i.e., synchronously sampled and the same data type) or which have different associated uncertainties have received relatively little attention in the SHM literature. These are the challenges that sensor fusion techniques seek to address.

Sensor fusion for this dissertation is defined as the process of integrating data from multiple sources in a way that enhances the decision-making performance of the SHM system. An overview of the general sensor fusion process along with many application examples may be found in Hall [56]. While sensor fusion is important for both damage detection and damage localization, how the procedure is applied is somewhat differently, and so the approach taken for each application will be discussed separately.

This dissertation will make use of the sensor fusion hierarchy outlined by Liu and Wang [57] and subsequently applied to a UGW SHM system by Lu and Michaels [58]. According to this hierarchy, there are three levels into which sensor fusion techniques may be classified – data-level, feature-level, and decision-level. The objective of sensor fusion process is always to take the full data matrix (which for UGW SHM systems has a size corresponding to the number

of time samples by the number of sensor pairs by the number of trials) and reduce it down to a set of decisions. Which level of fusion is chosen chiefly governs the order that the each of the dimensions is integrated. As such, it is a choice – only one of the levels of sensor fusion may be implemented at a time. Figure 10 illustrates graphically what sensor fusion as applied to damage detection looks like. Each of the yellow "Compute feature" boxes represents the step of computing the NRE metric as described in Section 2.2. Blue boxes indicate that some outside information may be added to the detection process. The red boxes are where each of the sensor fusion techniques discussed in the following sections is implemented. Each of the levels is next defined more fully, along with the example methods that were applied for the damage detection study.

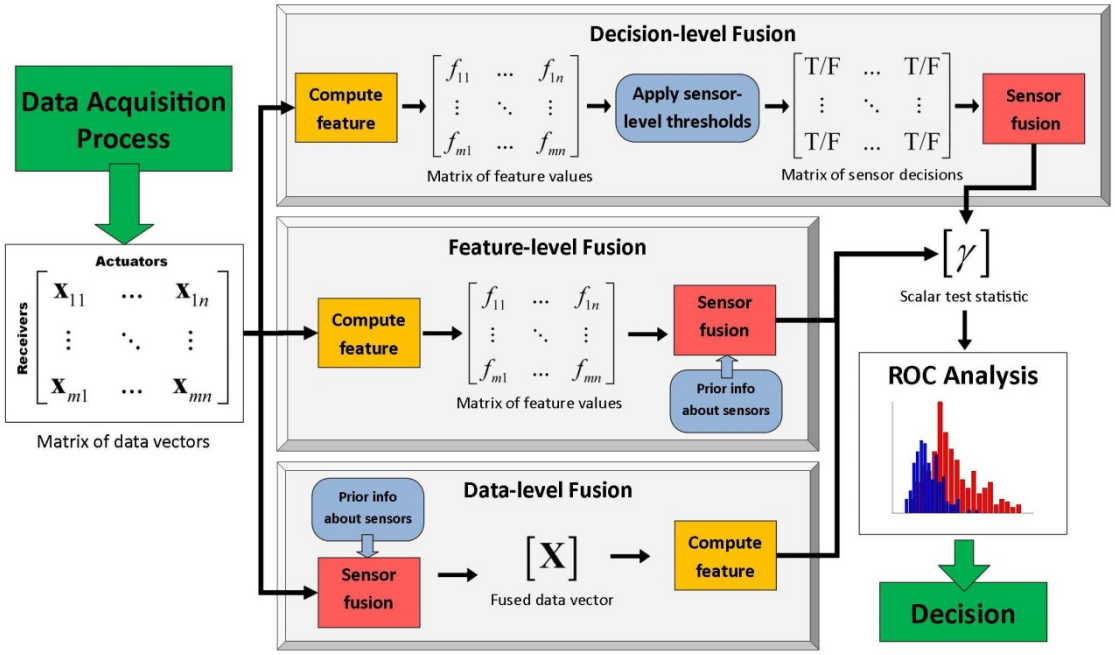

Figure 10 - Fusion hierarchy as applied to damage detection

Data-level fusion is the lowest level on which the fusion process may be performed. Data-level methods combine raw data from each of the sensors directly. As such, these

methods are often computationally expensive to execute. Furthermore, they may be difficult to implement for some SHM systems because of the requirement that all of the data be commensurate. One data-level fusion technique is implemented here, which involves concatenating the waveforms from all sensor pairs together before processing and calculating the normalized residual energy metric. Because the metric does not depend on time per se and the signals are commensurate in this case, concatenating the signals is not problematic. For the frame structure, this technique is further subdivided to consider pulse-echo and pitch-catch cases separately.

Feature-level fusion involves calculating feature values (i.e., applying the appropriate detector) for each sensor pair individually and then combining the features to make decisions. Feature-level fusion is retains a more of the underlying information than the decision-level fusion, but is substantially easier to apply than data-level fusion. Many of the common sensor fusion methods for SHM, such as simply averaging feature levels across all sensor pairs, belong to this category. For illustration purposes, three different matrix norms are applied to the data – the L^1 norm (or arithmetic mean), the L^2 Norm (or sum square), and the L^∞ Norm (or maximum value).

Decision-level fusion is the highest-level fusion in this hierarchy. In these methods, each sensor makes its own decision on the damage state independently, and the results from all of the decisions are then fused. In this way, there is no coordination between the data acquired from different sensors. Voting is one decision-level technique that is considered here. Any decision-level technique, including voting, involves determining the statistics of each sensor pair individually based on the baseline data. Then, detection thresholds may be established for each sensor pair allowing the decisions to be made independently. For the case of voting, after the decisions have been made, the number of waveforms whose feature value exceeds the

predetermined threshold is counted. The vote count then becomes the metric distinguishing the damaged and undamaged classes. A hybrid approach may also be used, whereby the features of all waveforms pertaining to either an actuator or a sensor are aggregated using some matrix norm (or other feature-level technique), and each sensor (rather than sensor pair) gets a vote. The hybrid approach is reminiscent of those proposed by Su *et al*. [59, 60].

For the localization paradigm presented previously, the role of sensor fusion is to appropriately balance the contributions of the MLE filtering result of each sensor pair in the delay-and-sum beamforming step. As such, the natural choice is techniques based on decision-level fusion. The sensor fusion step is particularly important for larger, complex structures instrumented with a sparse array because damage in arbitrary locations does not cause significant changes in the signal between some of the sensor pairs. Unaddressed, the lack of damage observability causes the maximum-likelihood estimate of the first arrival to break down and thus degrades localization performance.

Therefore, the NP criterion is applied to establish thresholds for the signals from each sensor pair separately. With such thresholds established from the baseline data, only the sensors that detect the presence of damage are selected for use in the localization procedure. A very similar approach has been proposed by Niri *et al*. [61] in the context of acoustic emission – here, the method will be referred to as "NP binary fusion". The binary hypothesis tests are carried out in the same way that was outlined in Section 2.2, utilizing the normalized residual energy (NRE) as the theoretically-optimal damage detection metric for systems with minimal *a priori* knowledge and high signal complexity [62].

Next, the matrix of all threshold values for each sensor pair (denoted as $\boldsymbol{\Gamma}$) is computed by taking the maximum across all of the undamaged trials as shown in Equation (1.16). The index (z) represents the different trials.

$$\boldsymbol{\Gamma}_{ij} = \max\left(\mathbf{V}_{ij}^{(z)},\ \forall\ \mathbf{V}^{(z)} \in \text{undamaged}\right) \tag{1.16}$$

Then the NRE values are compared to the thresholds to form **D**:

$$\mathbf{D}_{ij} \equiv 10^{\left(\mathbf{V}_{ij} - \boldsymbol{\Gamma}_{ij}\right)/20} \tag{1.16}$$

Finally, the scaling factor applied to each pair, **F**, is defined to be one for all pairs that exceed the threshold and zero otherwise for NP binary fusion:

$$\mathbf{F}_{ij}^{\text{NP binary}} = \begin{cases} 1,\ \mathbf{D}_{ij} > 1 \\ 0,\ \mathbf{D}_{ij} \leq 1 \end{cases} \tag{1.16}$$

Increasing the sophistication another step, the contribution of each sensor pair's data to the delay-and-sum procedure may be weighted in accordance with some estimate of confidence that the sensor pair can observe the damage. Niri *et al*. [61] at this juncture rely on a model of their structure, from which they predict the signal-to-noise ratio (SNR) of each sensor *a priori* by a Monte Carlo simulation. However, this solution is not directly applicable to baseline-subtraction techniques, where SNR is more difficult to define and would require models of the uncertainty in the subtraction process (more on this topic in Section 4.3.1). Here, therefore, a data-driven method is developed. The only data available from which to estimate the confidence of each sensor pair are the metric values computed on each waveform. By comparing the current metric value to the sensor-level threshold, a statistical distance from the baseline set may be computed, yielding a set of "confidence factors". This method, as described mathematically in Equation (1.16), will be referred to as "NP direct scaling" fusion.

$$\mathbf{F}_{ij}^{\text{NP direct scaling}} = \begin{cases} \mathbf{D}_{ij}, \ \mathbf{D}_{ij} > 1 \\ 0, \ \mathbf{D}_{ij} \leq 1 \end{cases} \tag{1.16}$$

These confidence factors are analogous to the SNR for passive sensing systems, since they represent by what margin a particular measurement has emerged from the "noise floor" of metric values that belong to the undamaged set. No guarantee of optimality is given because such a guarantee would have to consider the relationship between uncertainty in the input and output of the MLE first-arrival estimation. This relationship is a complicated question outside the scope of this dissertation, but recommended for future study. In this case, the statistical distance used is simply the ratio already defined as the NRE metric, but other scaling functions are possible. For example, one alternate choice of scaling factors might be the number of standard deviations of the baseline NRE values from the threshold. Again, further study to understand the propagation of uncertainty through the MLE first-arrival estimation process is required to determine what the optimal scaling technique might be. However, considering that the present approach utilizes the only helpful information available—that is, the comparison with threshold levels for each pair—it is reasonable to expect improved performance. The data, presented in Section 3.3.3, support this supposition.

In order to apply the confidence factors correctly, prior to applying them to the data, the minimum of each MLE filter result is subtracted from itself to make the minimum value zero. This normalization procedure does not impact the delay-and-sum result because the MLE formulation is invariant to additive shifts. Further note that the range of values produced by the MLE filter is difficult to relate to any particular property of the underlying physical signal. Therefore, scaling by confidence factors does not weaken the validity of the imaging scheme in any way.

A portion of this chapter has been published in Structural Health Monitoring: An International Journal, Colin Haynes, Michael Todd, Eric Flynn, and Anthony Croxford, 2013. The title of this paper is “Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation”. The dissertation author was the primary investigator and author of this paper.

A separate portion this chapter has been has been submitted for publication in Mechanical Systems and Signal Processing, Colin Haynes and Michael Todd. The title of this paper is “Enhanced Damage Localization for Complex Structures through Statistical Modeling and Sensor Fusion”. The dissertation author was the primary investigator and author of this paper.

Chapter 3
Experimental Validation

3.1. Bolted Frame Structure

The first structure upon which the statistical SHM framework will be tested is a three-story bolted frame. Several elements of the design are of interest for the present application. While bolted joints do transmit UGW energy between elements, it is extremely difficult to precisely predict the nature of that interaction, due to the high uncertainty of parameters like contact pressure, surface roughness, etc. The geometry is also somewhat unusual for UGW SHM applications, where the majority of work is done in plate-like structures. The frame would best be classified as a 1D manifold, which is to say that each element is 1D but they are connected to one another in 3D space. This type of geometry introduces multipath uncertainty because there are typically several paths that can be taken from one point to another. Finally, the two types of "damage" introduced to the frame were non-destructive in nature. Because the damage modes were reversible, it was possible to collect much more data than would have been feasible for destructive tests, allowing sufficiently large data sets for meaningful statistical analysis.

3.1.1. Geometry

The bolted frame is composed of 19 elements made of 2-inch-wide by 1/8-inch-thick (5.1 cm by 3.2 mm) plates with 12 and 24 inch (30.5 cm and 61.0 cm) lengths. The elements connect through angle brackets on each end with two ¼-inch-diameter (6.35 mm) bolts per connection. Both the elements and the angle brackets are made of steel. A picture of the geometry is given in Figure 11.

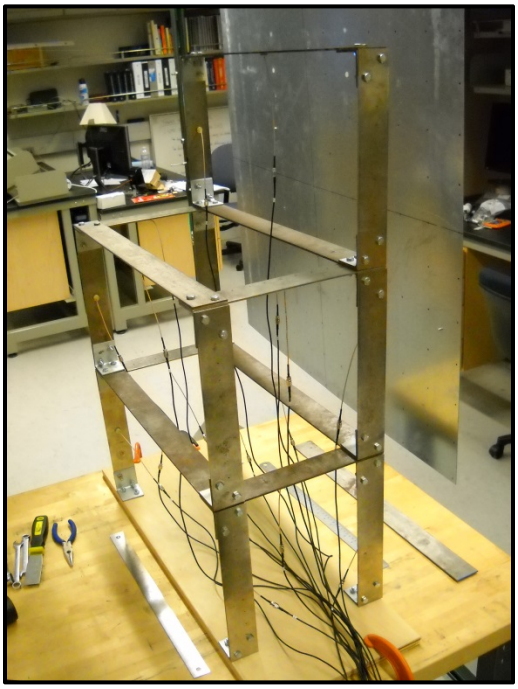

Figure 11 - Bolted frame geometry overview

3.1.2. Data Acquisition

For this test, 16 piezoceramic sensor-actuators were positioned on the structure as shown in Figure 12. While the structure itself is possesses substantial symmetry, the sensors were deliberately applied asymmetrically. Such asymmetry was introduced to help minimize multipath ambiguity which may complicate source localization in the medium.

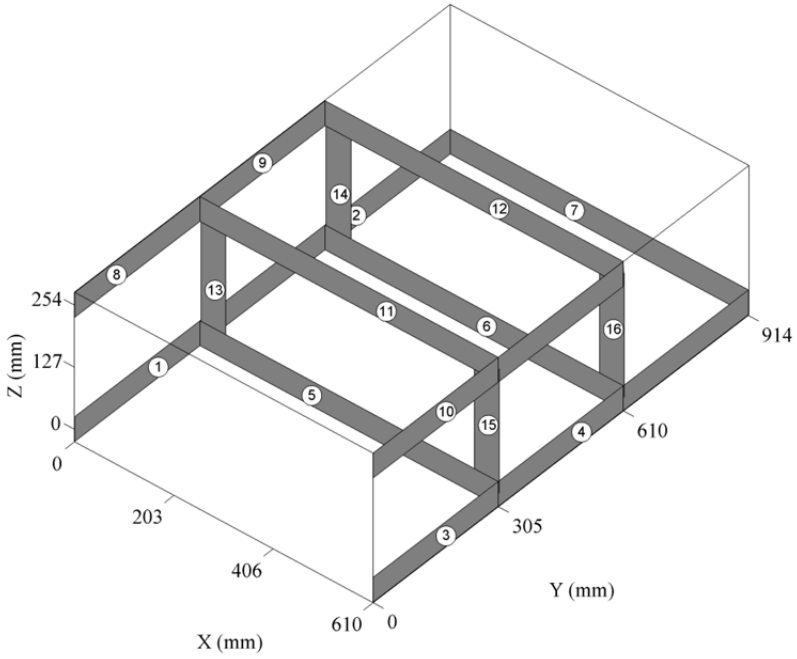

Figure 12 - Sensor layout in bolted frame structure

A National Instruments data acquisition system recorded both pitch-catch and pulse-echo for the entire 16x16 array with 14 bits of precision at a 2.5 MHz sampling frequency. The chosen actuation signal, shown in Figure 13, was a Gaussian-modulated sinusoid with a center frequency of 135 kHz. Each of these lead-zirconate-titanate (PZT) discs had a diameter of 12 mm and each was protected by sandwiching it between a magnet of the same diameter and an epoxy backing that also protected the electrical leads. The magnet was included as an aid for quickly attaching the sensors to ferromagnetic structures while retaining the ability to move them later. For these tests, however, each of the sensors was bonded to the steel element using Elmer's glue to improve the bond condition and therefore the power of the transmitted and received signals. Finally, it was determined that these sensors excite predominantly antisymmetric (A_0) mode waves in the 3.18-mm-thick plates.

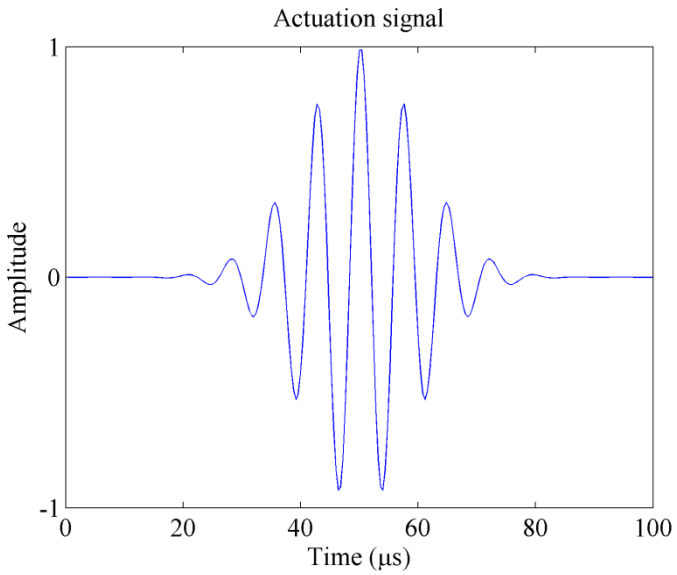

Figure 13 - Actuation signal for the bolted frame structure

3.1.3. Test Protocol

Two modes of damage were introduced in the frame structure. First, fastener preload loss was simulated by loosening the bolts in two stages (designated "partially loosened" and "no preload"). The bolt torque levels were not rigorously quantified for the test, but the tightening procedure was kept as consistent as possible across all trials. The process of loosening a bolt produces a change in contact pressure, which causes a very dramatic change in the UGW propagation, as was subsequently verified by full-field measurements of the wavefield [63]. In all, 20 bolts were loosened, accounting for at least one bolt at each node of the structure. Before each trial, 2 baselines and 8 undamaged records were recorded, followed by 5 records at the partially-loosened stage and 5 more at the no-preload stage.

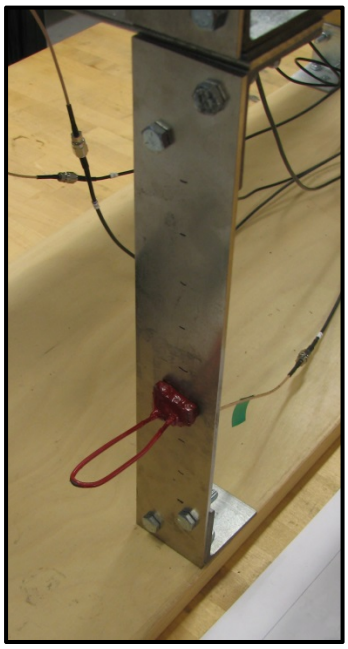

Figure 14 - Magnet damage applied to the bolted frame structure

The second damage type consisted in placing a magnet on the structure, as shown in Figure 14, which causes a local change in the mechanical impedance. Because any mechanical impedance change serves to scatter UGWs, the magnet has a similar impact to what corrosion or other forms of real damage might, while retaining the benefit of being completely reversible [64]. In all, the magnet was placed in 261 locations, essentially every location on each of the elements at one-inch (2.5 cm) intervals. Interspersed with the damaged cases were collected 261 undamaged records, a quarter of which were designated as the baseline set.

3.2. Fuselage Rib Structure

The second structure selected for testing was a section of fuselage rib from a British Aerospace BAe 146 aircraft, pictured in Figure 15. As an actual component from a commercial aircraft, this structure is well-suited for this type of study. Therefore, both forms of damage

introduced (holes and cracks) were done in a destructive manner to more closely represent damage modes that may be of practical interest.

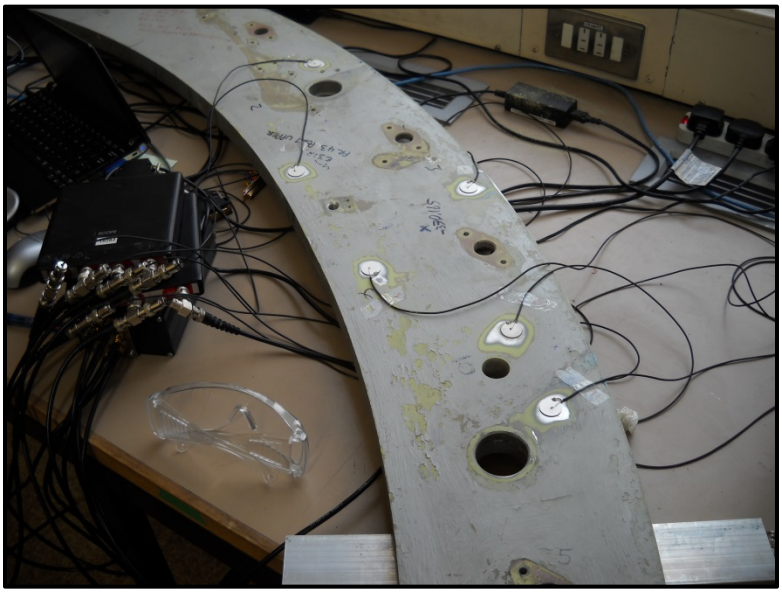

Figure 15 - Fuselage rib overview

3.2.1. Geometry

Details of the fuselage rib geometry are given in Figure 16. The longest dimension of the structure is about 1.7 m. The typical thickness of the top surface is 2.5 mm, and the typical stringer thickness is the same. However, many areas of non-uniform thickness, different stringer arrangements and sizes, and different ducts are all evident, defying a simple geometrical description.

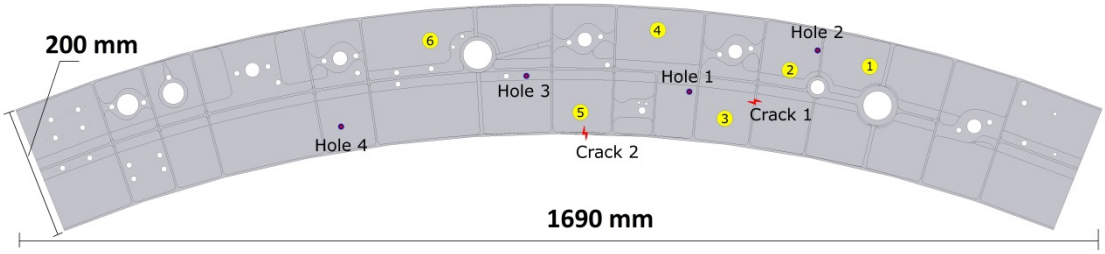

Figure 16 - Simplified geometry of fuselage rib structure, with sensors and damage locations noted

3.2.2. Data Acquisition

Six PZT discs were directly applied to the top surface of the structure with cyanoacrylate adhesive in the arrangement shown in Figure 16. Again, the sensors were deliberately placed with no symmetry, and are arranged in such a way that different pairs have different line-of-sight obstructions to test the effect of the obstacles on system performance. The excitation signal was a 5-cycle, Gaussian-modulated sinusoid centered at 190 kHz for this test. Due to data acquisition limitations, the sensors were used exclusively in pitch-catch mode. The PZT discs had a diameter of 20 mm and were designed to excite the first symmetric Lamb wave mode (S_0) with at least 10 times the magnitude of the first antisymmetric (A_0) mode.

3.2.3. Test Protocol

Damage was introduced in two modes: through-holes and simulated cracks. The holes were drilled in 4 stages, progressively increasing the diameter to 3.4, 4.8, 6.0, and 8.0 mm. The damage locations were chosen to present increasing levels of difficulty for the sensor network, with the final hole being located a considerable distance outside the sensor array. Before each hole location was drilled, 10 baseline records were saved. The damage, therefore, does not appear as cumulative, but each hole is considered individually. After each increase in diameter, 3 records were acquired.

Cracks were simulated by using a hacksaw to notch the stringers (through the whole thickness) on the reverse side of the fuselage rib. Notches were introduced in two stages, the first being about half the depth of the stiffener, and the second being the full depth. Again, before each crack was begun, 10 baseline records were saved, and after each increment of damage, 5 records were acquired. An example of the simulated crack damage is shown in Figure 17.

Figure 17 - Crack damage introduced in fuselage rib structure

3.3. Comparative Studies

3.3.1. Detection

In order to validate the optimality of the detector presented previously, ROC curves are computed for the data collected from the two testbeds. In this way, the energy detector may be compared empirically to other candidate detectors, and to this end six detectors were selected from the SHM literature. Taking $\mathbf{x}[n]$ to be the residual signal and $\mathbf{y}[n]$ to be the original baseline, the formulas for these six features are given in Table 1, where f_0 is the frequency bin corresponding to the waveform's center frequency, w is a window function (in this case, a Hamming window), and μ and σ represent the means and standard deviations, respectively, of the subscripted waveforms. Finally, τ_0 is the time of arrival of the first wave packet, determined by applying the MLE filtering approach. Note that the correlation coefficient uses original waveforms before baseline subtraction, denoted by $\mathbf{x}_y[n]$. Because damage detection is an inherently comparative process, each formula is essentially a variation on the ratio of the

residual to the original baseline either in the time domain or frequency domain. The result of each formula (except the correlation coefficient) may therefore be converted to a dB value for convenience. Dividing by the baseline feature value normalizes the values according to the amount of energy that can propagate between two given sensors.

Table 1 - Description of features used in ROC comparison

Name	Equation	Reference
Normalized Maximum Residual	$\frac{\max\left(\left\|\mathbf{x}[n]\right\|\right)}{\left\|\mathbf{y}[\tau_0]\right\|}$	Croxford *et al*. [42]
Maximum amplitude	$\frac{\max\left(\left\|\mathbf{x}[n]\right\|\right)}{\max\left(\left\|\mathbf{y}[n]\right\|\right)}$	Michaels *et al*. [65]
Normalized Residual Energy	$\frac{\sum_n \mathbf{x}[n]^2}{\sum_n \mathbf{y}[n]^2}$	Kay [45]
FFT at center frequency	$\frac{\mathrm{FFT}\left(\mathbf{x}[n]\right)\Big\|_{f=f0}}{\mathrm{FFT}\left(\mathbf{y}[n]\right)\Big\|_{f=f0}}$	Kay [45]
Correlation coefficient	$1-\frac{\sum\left(\mathbf{x}_{\mathrm{y}}[n]-\mu_{\mathrm{x}}\right)\left(\mathbf{y}[n]-\right.}{\sigma_{\mathrm{x}}\sigma_{\mathrm{y}}}$	Lu and Michaels [58]
Short Time Fourier Transform	$\frac{FFT(\mathbf{x}[\tau]w(\tau-t))\Big\|_{f=f0}}{FFT(\mathbf{y}[\tau]w(\tau-t))\Big\|_{f=f0}}$	Ihn and Chang [66]

In order to calculate these ROC curves, the feature values are first calculated for each sensor pair and each record separately. That is, no sensor fusion is implemented, but the binary

hypothesis test is applied to each actuator-sensor combination in turn and the results are aggregated. Instead of having a 16-sensor system for the frame structure, therefore, this methodology essentially sums together the detection performance of 16x15 = 240 pitch-catch pairs and 16 pulse-echo sensors. The detection problem therefore much more difficult because it measures the ability of each sensor pair to independently detect damage at any location. For the purposes of comparing metrics, however, it is instructive to look at the data before fusion.

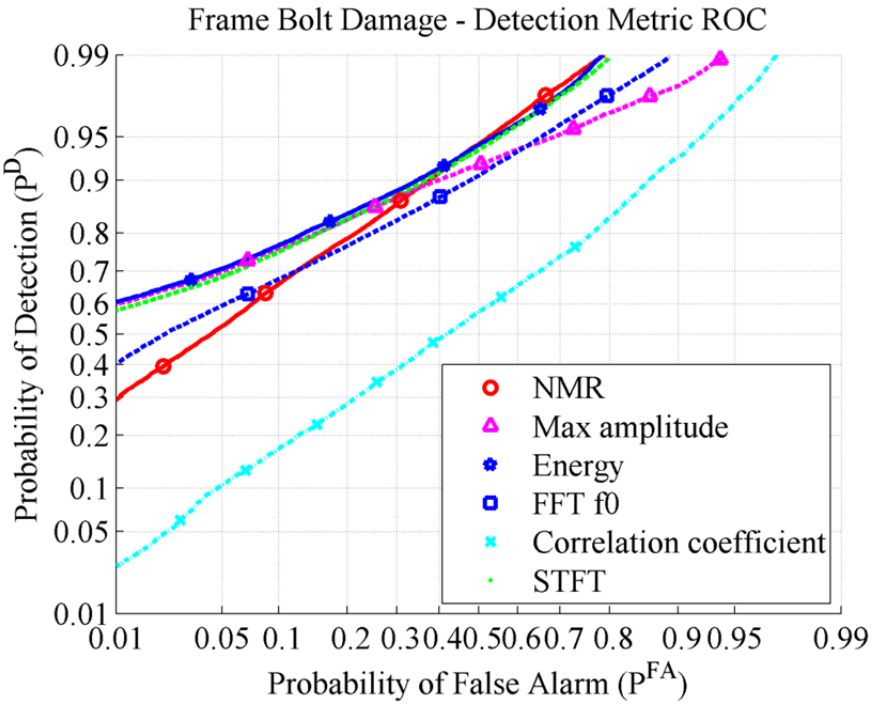

Figure 18 - Detection metric ROC for frame bolt damage

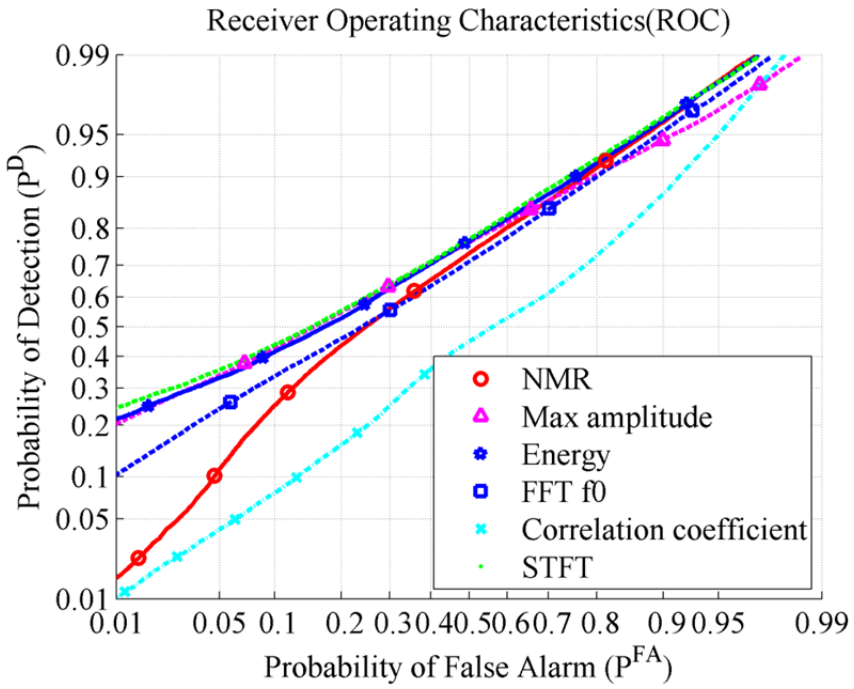

Figure 19 - Detection metric ROC for frame magnet damage

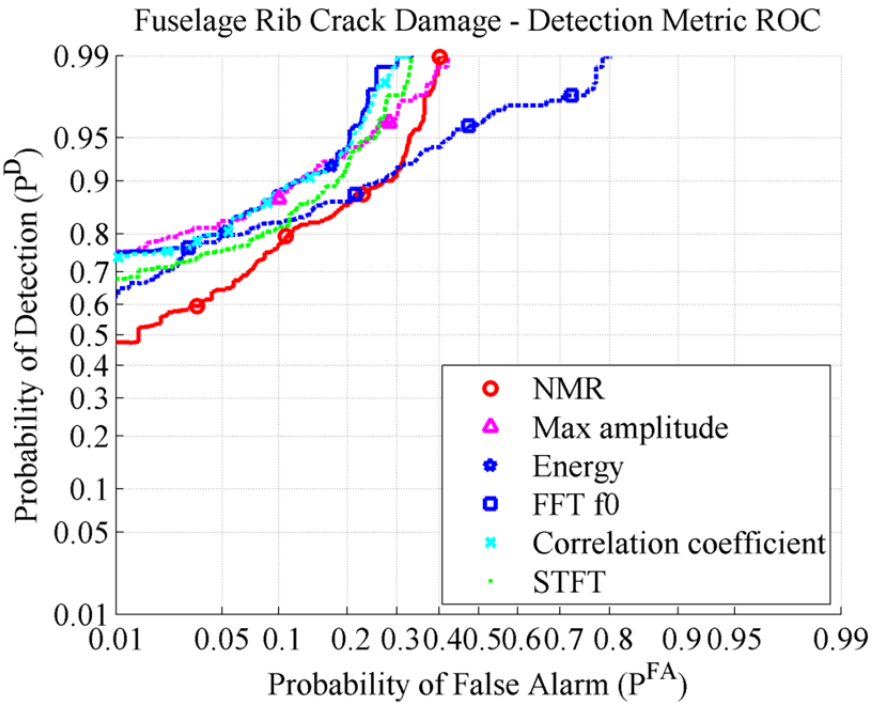

Figure 20 - Detection metric ROC for fuselage rib crack damage

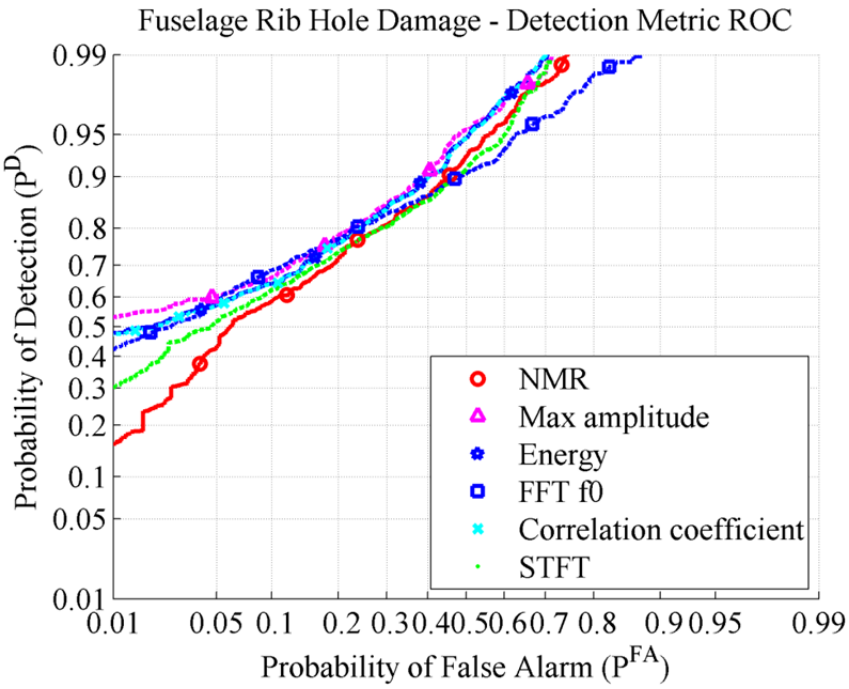

Figure 21 - Detection metric ROC for fuselage rib hole damage

The ROC comparisons are presented in Figure 18 for the frame structure bolt damage, Figure 19 for frame magnet damage, Figure 20 for the fuselage rib crack damage, and Figure 21 for the fuselage rib hole damage. As predicted theoretically, the energy metric does generally perform best for both structures, especially for low false alarm rates (which is the region of interest for most SHM applications). Interestingly, the maximum amplitude metric performs similarly well for both structures. One possible explanation for this result is the particular character of the reflections – if the damage happens to result in very sharp peaks of energy, the maximum amplitude metric would be expected to perform well. Of course, the fact that the energy metric is theoretically optimal for the given problem definition does not preclude other metrics from performing just as well, provided they have the same statistics. Also note that because the correlation coefficient metric does not use baseline subtraction, there is no

opportunity for environmental compensation, which may explain its poor performance in the frame structure tests.

When the data are fused together, the detection results only improve. For the bolted frame structure, 223 of the 261 magnet damage trials are detected for a threshold value that yields no false alarms (over 85% P^D). Most of the damage trials that are not detected are on one of the three elements that were not instrumented with a sensor. That is, the detection (and the localization) performance of the SHM system varies based on sensor coverage and damage position. Figure 22 shows the detection plot used ultimately to decide which magnet cases were correctly detected and passed to the localization algorithm. In the plot, the NRE metric is on the y-axis and the x-axis is labeled "Trial number", which is simply all of the tests aggregated into the appropriate damaged or undamaged class (with no implication of the order in which the trials were recorded). Clearly, the bolt damage can easily be detected, but note also that each cluster on the damaged side represents a different bolt location, demonstrating again the variability of the metric with damage position. The undamaged cases in the two examples have very different means, and the cause is the different test procedures. For bolt damage, retightening the bolts is not sufficient to bring the structure back to exactly the same undamaged state—the contact condition of the joint inevitably changes—so the undamaged trials have higher residual levels after baseline subtraction.

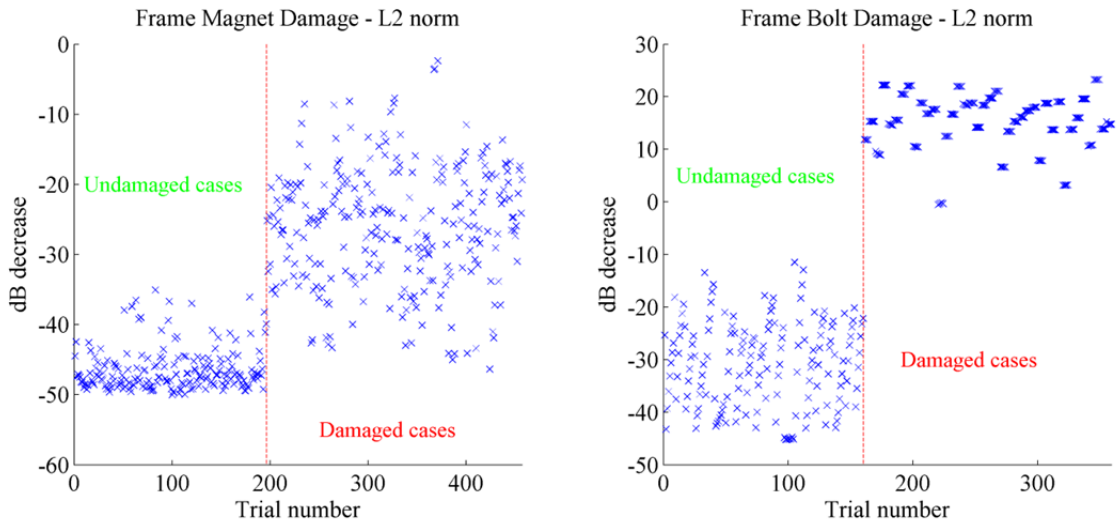

Figure 22 - Detection results in bolted frame for magnet damage (left) and bolt damage (right)

Figure 23 shows the damage detection results for the fuselage rib and reveals that both damage modes were detectable. As expected, the damage closer to the center of the sensor network produced a relatively greater shift in the fused detector value. All but the most remote hole (hole 4) yielded perfect detectability for the fused data using a simple L^1 norm (more on optimal sensor fusion choices in Section 3.3.3). Furthermore, each step of increasing damage size is manifested as an upward jump in the metric value. Such trends give reason to believe a measure of damage extent could be derived from the data in the future. Finally, despite not directly affecting the primary waveguide in the top surface of the fuselage rib, the crack damage on the stringers was also perfectly detectable, producing a change in residual level on the same order as the through-holes in the top surface.

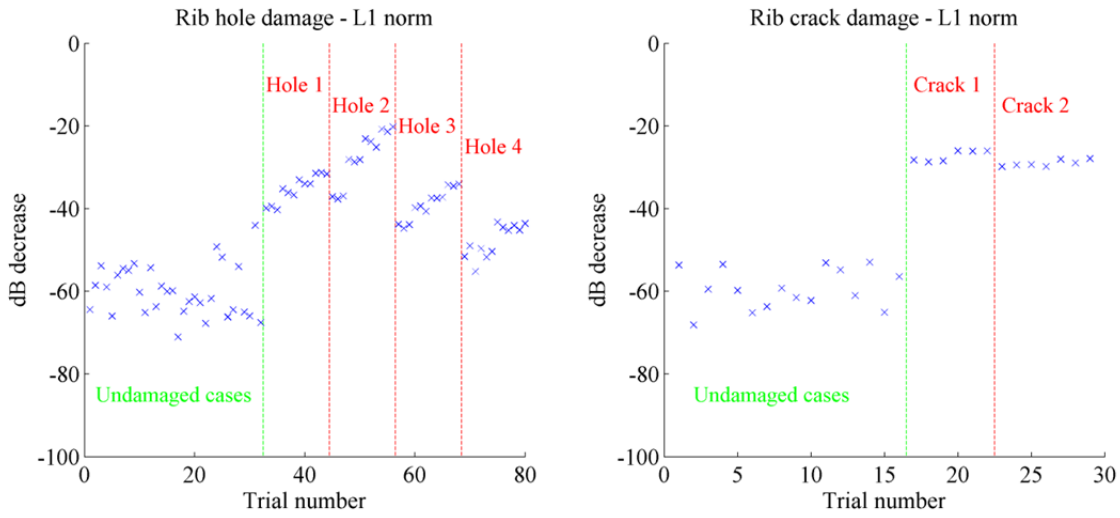

Figure 23 - Detection results in fuselage rib for hole damage (left) and crack damage (right)

Given the preceding results, the optimality of the energy metric for UGW inspection in geometrically-complex environments has been established both theoretically and empirically.

3.3.2. Localization

In order to establish the effectiveness of the previously described localization paradigm, a comparison with other localization techniques from the literature is presented. A similar comparative study was performed by Flynn *et al*. previously [55], and the current list of methods is the same except for the addition of the MVDR method. Each of the methods is described and references are provided – for more details, the reader is referred to Flynn *et al*. or the references given.

All of the presented techniques use data filtered about the input frequency band and then baseline-subtracted. Each signal is enveloped by taking the magnitude of the Hilbert transform, unless otherwise mentioned. The WTOA, TDOA, and WTP methods also apply an exponential window on the data to suppress later reflections. Because the boundary reflections and inaccurate baseline subtraction become more dominant with time, the purpose of the

exponential time window is to mitigate those effects by giving the information earlier in the signal a relatively higher weighting. The time constant for the exponential decay was chosen as 80 μs for the fuselage rib and 200 μs for the frame structure (according to the approximate ratio of group velocities). Note that exponential windowing, and many of the other methods, are highly dependent on tuning parameters like exponential constants or time windows. Each of the methods is implemented as follows:

TOA – the time of arrival method or the ellipse method is the most basic approach, simply computing the delay-and-sum result for the enveloped waveform [67]. The TOA formulation is described mathematically by Equation (3.1):

$$\mathbf{I}^{TOA}(r) = \sum_i \sum_j \mathbf{x}_{ij}\left(\tau_{ij}(r)\right) \tag{3.1}$$

WTOA – the windowed time of arrival method adds an exponential time window to the TOA method in order to suppress secondary reflections. The WTOA formulation is described mathematically by Equation (3.2):

$$\mathbf{I}^{WTOA}(r) = \sum_i \sum_j \mathbf{x}_{ij}\left(\tau_{ij}(r)\right) w\left(\tau_{ij}(r)\right) \tag{3.2}$$

TDOA – the time difference of arrival method is based on correlations between each two sensor pairs that share exactly one sensor in common [51]. Because the algorithm does not require knowledge of the time origin of the source signal, TDOA is commonly used by passive systems for source localization. Because the method operates on two sensor pairs simultaneously, the fusion approaches are not implemented to avoid a potentially false comparison. The TDOA formulation is described mathematically by Equation (3.3):

$$\mathbf{I}^{TDOA}(r) = \sum_{i=1}^{M} \sum_{\substack{m=1 \\ m \neq i}}^{M} \sum_{\substack{j=i+1 \\ j \neq m}}^{M} \rho_{im,mj} \left(\tau_{im}(r) - \tau_{mj}(r) \right) \tag{3.3}$$

RAPID – the reconstruction algorithm for the probabilistic inspection of damage uses cross-correlations between baseline and test signals (without enveloping) as the basis for localizing damage [68]. Each transducer pair contributes to elliptical sub-regions, the size of which is controlled by an additional factor, β. This parameter was set to a value of 1.45 based on calibrating it in post-processing to maximize the localization accuracy (a procedure which generally is not feasible). The RAPID formulation is described mathematically by Equation (3.4):

$$\begin{aligned} \mathbf{I}^{RAPID}(r) &= \sum_{i} \sum_{j} \left(1 - \rho_{ij}^{R}(r)\right) G\left(\frac{\beta - R_{ij}}{\beta - 1} \right) \\ G(x) &\equiv \begin{cases} x, & x > 0 \\ 0, & \text{otherwise} \end{cases} \\ R_{ij} &= \tau_{ij}(r) / \tau(i,j) \end{aligned} \tag{3.4}$$

where R is the ratio of the time-of-flight to the pixel of interest to the direct-arrival time between the sensors and $\rho_{ij}^{R}(r)$ is the zero-lag cross-correlation of the baseline and test signals for sensor pair ij. RAPID also applies a time window to the cross-correlation the arrival time of the point of interest (again, in an attempt to minimize the impact of any reflected signals). Here, the window width was chosen to be about the length of one actuation pulse. For the frame structure, this window was 50 μs, and for the fuselage rib, it was 30 μs.

EA – the energy arrival method is based on a ratio of energies, both in a time window beginning at the estimated arrival time and cumulatively up to that point [69]. In effect, the energy in the window centered at the expected arrival time is inversely weighted by the energy

preceding that time window. The EA method is similar to the MLE approach in that it searches for the time point where the previous signal energy is minimized compared to the energy of the reflection, but it does not consider the reflections occurring after the narrow time window. The EA formulation is described mathematically by Equation (3.5):

$$\mathbf{I}^{EA}(r) = \sum_i \sum_j \left(\frac{\sum_{n=\eta_{ij}-n_w}^{\eta_{ij}+n_w} x_{ij}^2 - \sum_{n=1}^{\eta_{ij}+n_w} x_{ij}^2}{\sum_{n=\eta_{ij}-n_w}^{\eta_{ij}+n_w} x_{ij}^2 + \sum_{n=1}^{\eta_{ij}+n_w} x_{ij}^2} \right) \tag{3.5}$$

TP – the total product method takes the short-time Fourier transform of each signal, computed at the actuation frequency and centered on the expected arrival time [66]. The results are then multiplied across all sensor pairs. As was previously noted by Flynn *et al*., this is approximately equivalent to the product of the filtered and enveloped waveform. Therefore, the simpler implementation has again been adopted. Because this technique is multiplicative, applying weighting factors to the sensor pairs (as proposed for the sensor fusion methods) does not change the result. There is therefore no difference in this method between NP binary and NP direct scaling fusion techniques. The EA formulation is described mathematically by Equation (3.6):

$$\mathbf{I}^{TP}(r) = \prod_i \prod_j S_{ij}\left(\omega_0, \eta_{ij}\right) \cong \prod_i \prod_j x_{ij}^2\left(\tau_{ij}(r)\right) \tag{3.6}$$

WTP – the weighted total product method adds an exponential time window to the TP method, as described by Equation (3.7):

$$\mathbf{I}^{WTP}(r) = \prod_i \prod_j x_{ij}^2\left(\tau_{ij}(r)\right) w\left(\tau_{ij}(r)\right) \tag{3.7}$$

MVDR – the minimum variance distortionless response method computes a set of optimal weights such that the energy that is not present in the "look direction" of any given pixel is minimized. The objective, therefore, is to minimize imaging artifacts outside a particular look direction. Where appropriate, phase information, dispersion compensation, and scattering fields may also be integrated directly into the approach to enhance the imaging results. However, to perform a fair comparison with the other techniques presented, these options were not implemented. Beam spreading, as defined by Equation (4.4), is typically included in the computation of the initial "steering vectors." While it was not implemented for the frame structure to better represent the physics, it was included for the fuselage rib, as the method was found to perform marginally better including effect. Furthermore, the MVDR method's optimization over the look vector is analogous to the sensor weighting provided by the sensor fusion techniques, so these methods were not applied. The mathematics of MVDR imaging is more complicated than the other methods, but a full description is given by Hall and Michaels [70].

Table 2 - Mean prediction error for the bolted frame structure

Frame Structure - Mean Prediction Error (mm)									
MLE	TOA	WTOA	TDOA	RAPID	EA	TP	WTP	MVDR	Fusion Technique
250	1019	730	624	636	926	1010	743	423	All PC channels
132	1004	645	-	222	943	1077	391	612	NP binary
120	970	614	-	141	747	-	-	-	NP direct scaling

Table 2 compares the results obtained by applying each of the preceding localization algorithms to the frame structure with no feature modeling applied. To construct a quantitative comparison a convenient and descriptive measure of the localization accuracy may be calculated by simply taking the distance between the predicted location and the true damage location. The results for both NP binary fusion and NP direct scaling fusion methods are presented, along with the results obtained using all pitch-catch sensor pairs. A more detailed analysis of the effect of fusion techniques is given in Section 3.3.3. For those methods where the fusion method chosen either does not influence the results or the application of the method is undefined, the entries are not repeated. For a structure of this complexity, the TOA, WTOA, TDOA, EA, and TP methods are all unable to outperform the completely uninformed localizer, which for the frame structure produces a mean error of 644 mm, as presented in Table 7 in the next section. (The uninformed localizer is a reference technique that takes no information from the sensing system. In such a situation, the best damage prediction possible is to always predict the “geometric centroid” of the structure, which is to say the point that is closest, on average, to every other possible damage location.) The MLE formulation achieves the best localization accuracy for these tests. In fact, only the RAPID and WTP methods are able to approach the level of performance achieved by MLE. Also note that for every localization algorithm, NP direct scaling fusion improved localization performance over that obtained by NP binary fusion.

Table 3 compares the results from each localization algorithm applied to the fuselage rib structure, again with no feature modeling. The performance of the MLE formulation is again superior, this time by a factor of more than two (in terms of mean prediction error) over the second-best method. For each structure, the second-best algorithm is RAPID, while none of the other methods exhibited similar performance on both of the structures tested.

Table 3 - Mean prediction error for the fuselage rib structure

Fuselage Rib Structure - Mean Prediction Error (mm)									
MLE	TOA	WTOA	TDOA	RAPID	EA	TP	WTP	MVDR	Fusion Technique
194	573	339	296	336	259	473	303	414	All PC channels
74	601	309	-	149	188	630	337	382	NP binary
63	443	324	-	141	163	-	-	-	NP direct scaling

To visualize their performance, localization performance curves are computed for each algorithm. These plots express the fraction of damage cases where the technique is able to localize the error to within a specific error bound. They may be interpreted similar to receiver operating characteristic curves, in that the best performance implies the curves should approach the top left corner, but note that the x-axis is not bounded except by the size of the structure. The localization performance curves for the bolted frame structure are given in Figure 24and the corresponding curves for the fuselage rib are given in Figure 25. Similar to ROC analysis, the performance represented by these curves can be represented by an AUC metric, except the area must be integrate from zero to the maximum error of the x-axis. Normalizing such that perfect localization would produce a value of one, the comparison results are again presented using this metric in Table 4 and Table 5.

Table 4 - Normalized AUC, bolted frame localization performance curves

Frame Structure - Normalized AUC									
MLE	TOA	WTOA	TDOA	RAPID	EA	TP	WTP	MVDR	Fusion Technique
0.857	0.412	0.580	0.641	0.633	0.466	0.417	0.572	0.757	All PC channels
0.925	0.420	0.629	-	0.874	0.456	0.378	0.776	0.647	NP binary
0.931	0.440	0.647	-	0.919	0.570	-	-	-	NP direct scaling

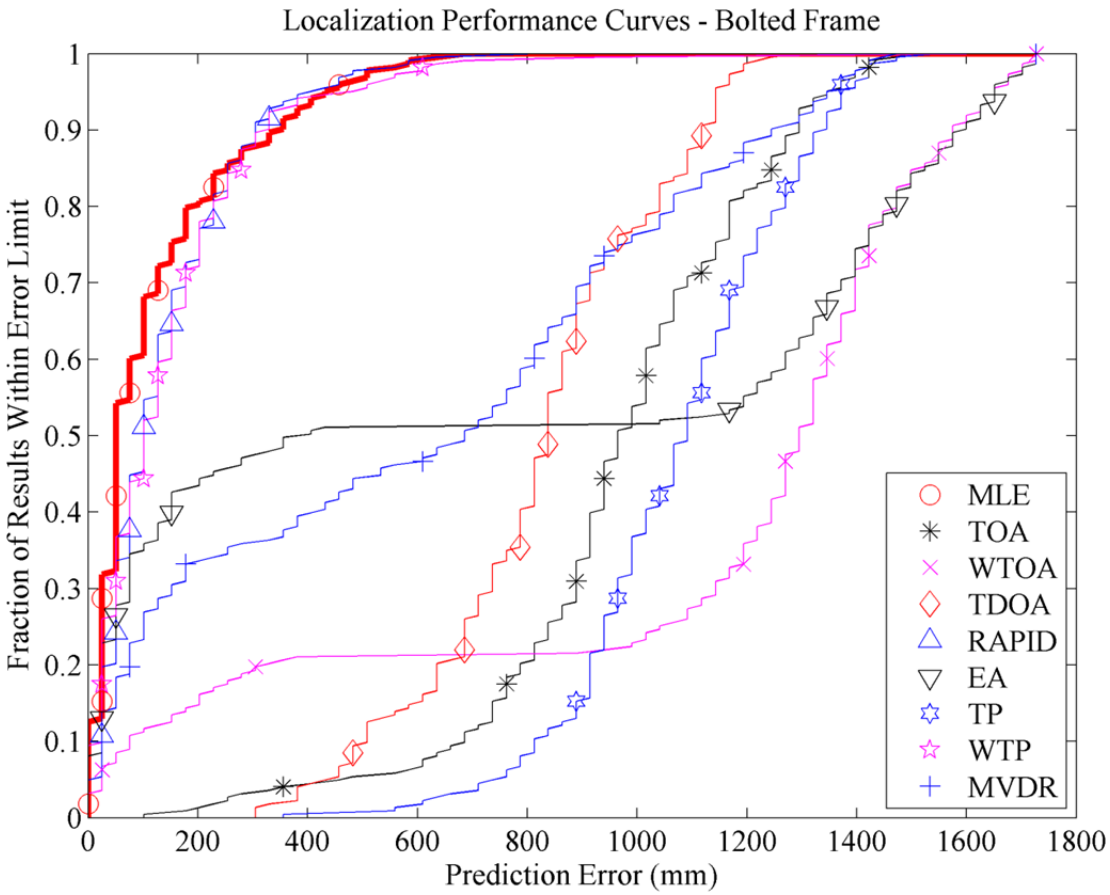

Figure 24 - Localization performance curves for bolted frame structure

Table 5 - Normalized AUC, fuselage rib localization performance curves

Fuselage Rib Structure - Normalized AUC									
MLE	TOA	WTOA	TDOA	RAPID	EA	TP	WTP	MVDR	Fusion Technique
0.785	0.346	0.615	0.667	0.618	0.709	0.462	0.657	0.529	All PC channels
0.921	0.313	0.648	-	0.832	0.788	0.280	0.617	0.567	NP binary
0.929	0.495	0.630	-	0.838	0.818	-	-	-	NP direct scaling

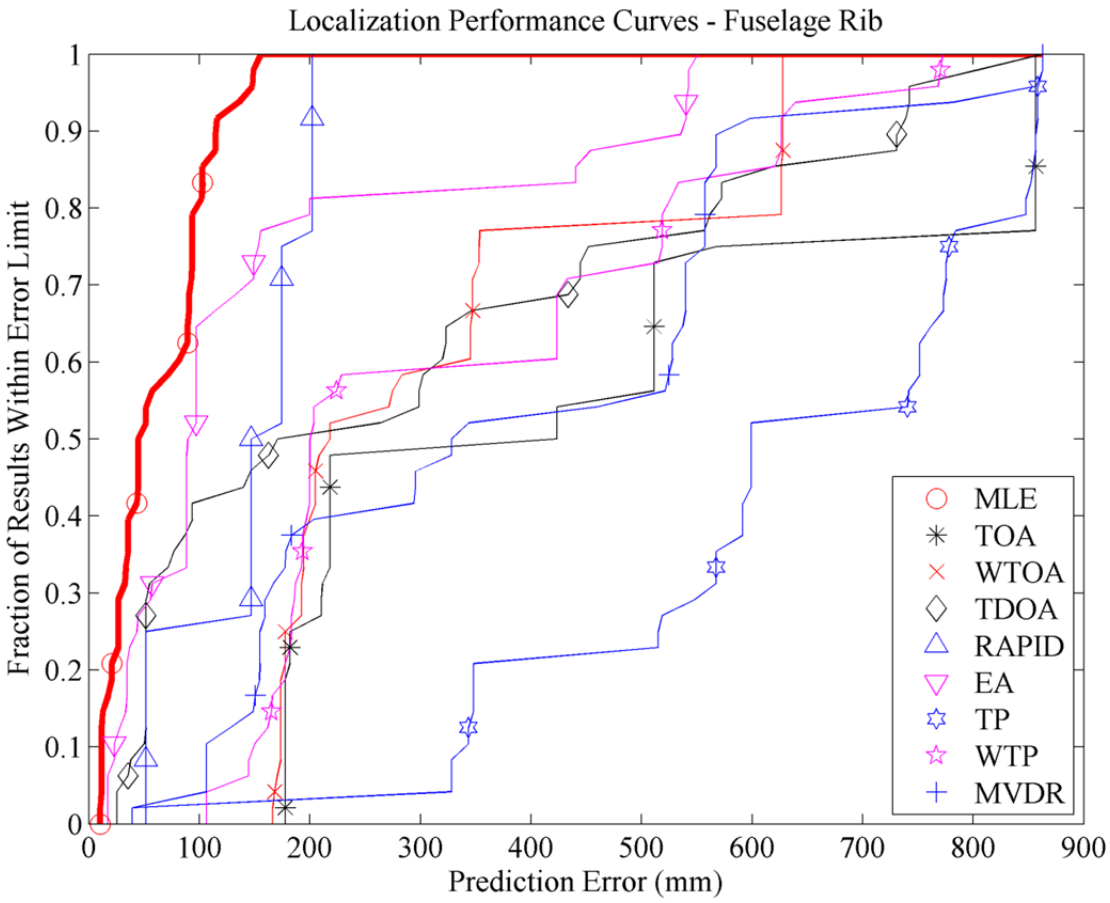

Figure 25 - Localization performance curves for fuselage rib structure

Each method of comparison suggests the same conclusion – the MLE procedure produces the most accurate results of any algorithm implemented here. The RAPID method

performs second-best, and no other method is able to produce comparable results for both the bolted frame and the fuselage rib. These conclusions hold independent of the particular sensor fusion method selected, a question examined more closely in the following section.

In the UGW SHM literature, it is common to think of delay-and-sum as generating images of the damage, where the highest pixel value or values represent the damage location. To visualize the MLE filtering results in a consistent manner, a final graphical view of the results is presented. The following damage localization figures are maps of the log-likelihood that the damage is at a given point on the structure. Because the damage detection question has been answered affirmatively in each of these cases, the location of the maximum likelihood is taken to be the predicted damage location. In the plots, sensors are noted as white circles with their corresponding channel numbers. To generate a convenient color-scale, the log-likelihood output from the algorithm has been normalized to a linear scale between 0 and 1. However, it is important to remember that these values do not represent a probability in any sense, only a normalized likelihood that is useful for visualizing how a particular maximum point came to be selected. These maps are also not meant to represent damage maps directly, and in their current form they cannot be used to estimate the size or extent of damage.

Two damage likelihood maps (of the 223 total) were chosen as representative examples for the bolted frame. Figure 26 shows an example where the damage location estimate directly coincides with the prediction. The pixel values in the image are significantly higher on the element containing the damage than on the other elements. The values also decrease farther away from the damage, as would be expected for a good localization procedure.

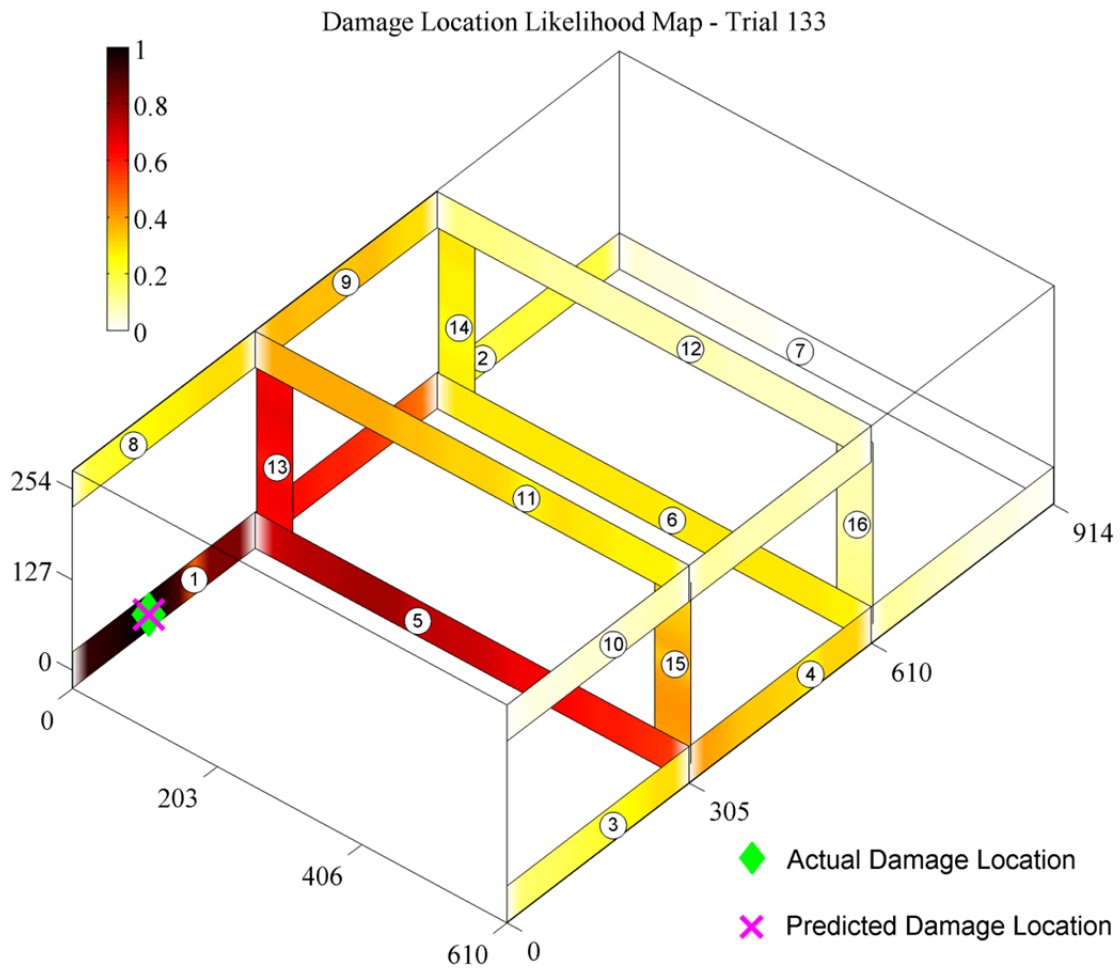

Figure 26 - Likelihood map for damage location prediction, frame trial 133. All dimensions are given in mm.

Figure 27 is an example case where the estimation is less accurate. In that scenario, the damage was located on an element on which no sensor was located, a situation that tends to produce higher estimation errors (again, as expected). The UGW SHM system overage is not uniform – that is, damage toward the center of the sensor array will be more easily observed than those on the outer edges.

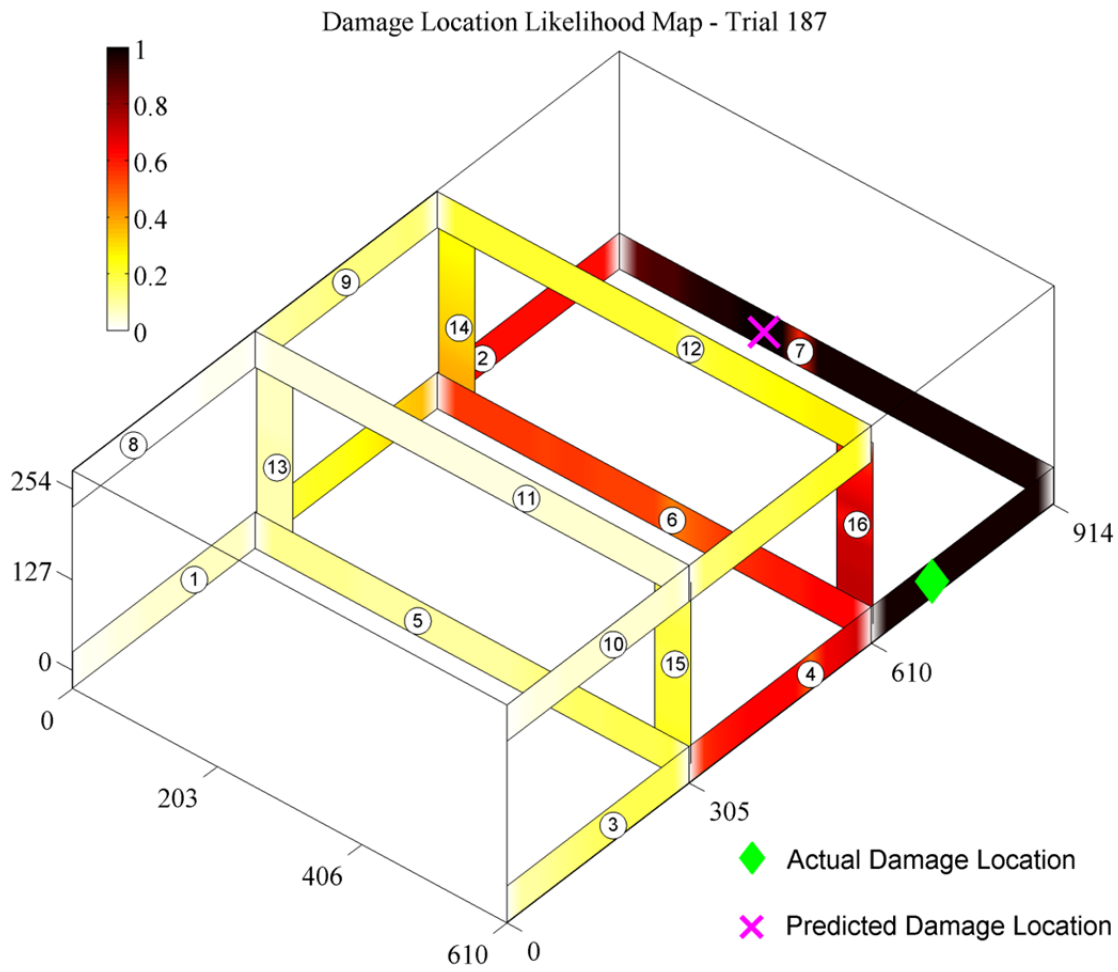

Figure 27 - Likelihood map for damage location prediction, frame trial 187. All dimensions are given in mm.

Thus, the accuracy of the localization is dependent on the location of the damage on the structure, as demonstrated by Figure 28. In this figure, each element of the bolted frame is broken into three and the average localization error on each segment computed.

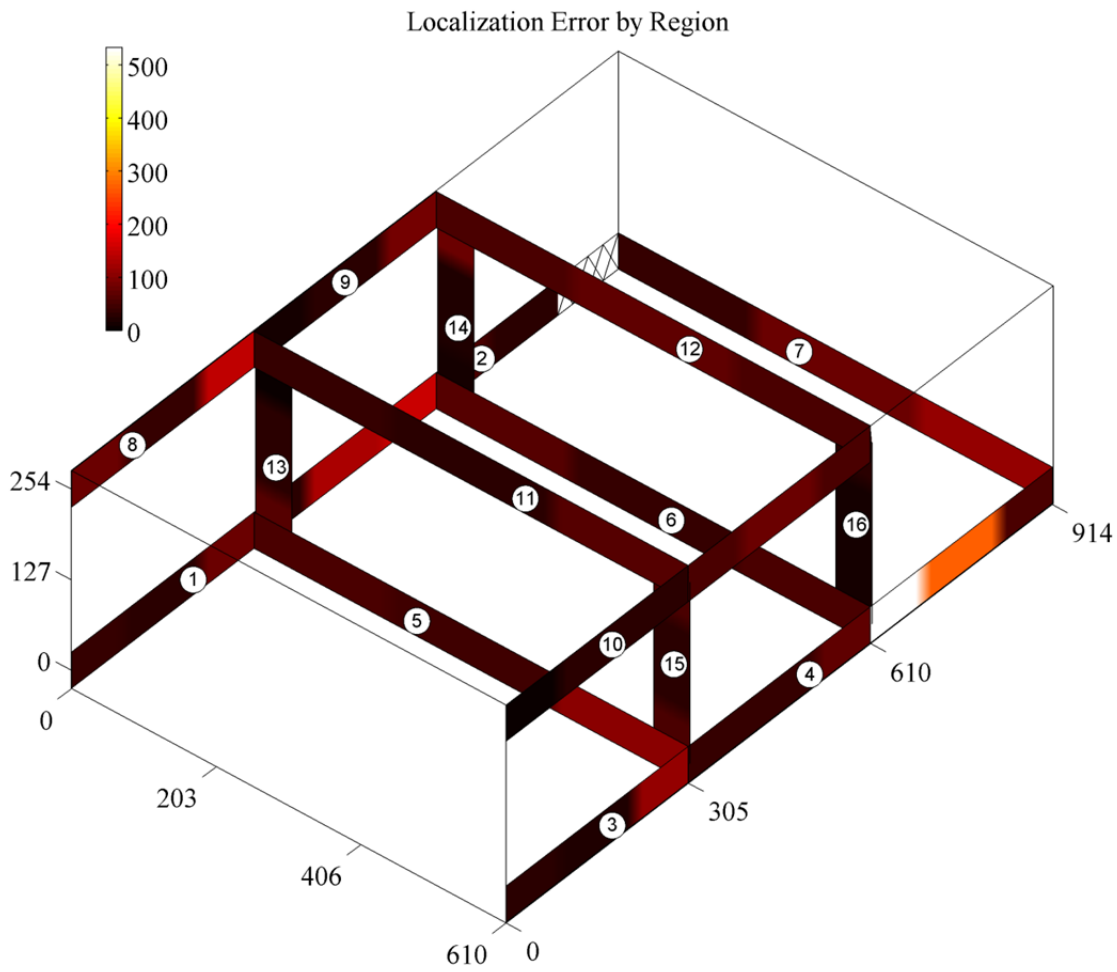

Figure 28 - Average localization error by region on the bolted frame structure. Areas with insufficient data are shown cross-hatched. All units are given in mm.

Similar images have been generated for the fuselage rib structure as well. In these figures, black circles represent major through-holes in the part and the solid black lines represent the larger stringers behind the instrumented surface. The part's actual geometry is more complex—smaller stringers and thickness changes were not considered for simplicity—but these considerations incorporate the some of the most significant geometrical obstacles to UGW inspection. As in the previous case, the likelihood values are normalized to fall between 0 and 1, where 1 represents the predicted damage location.

One example from each hole location on the fuselage rib, each at the maximum hole diameter, is presented for illustrative purposes. The first two hole locations, shown in Figure 29 and Figure 30, are located very precisely due to their central location with respect to the sensor array. Figure 31 shows a more challenging hole location, in that there is a large through-hole nearby which obscures the direct line of sight of the closest sensors. It is clear that the likelihood map is distorted in that region because the localization algorithm incorporates a basic model of the obstacle. In fact, similar distortions may be observed around the geometric features in each of the likelihood maps. Finally, Figure 32 depicts the final hole location, which is positioned well outside of the sensor array. Because the array viewed from that angle is nearly linear, the range to the damage can be predicted with much more accuracy than the angular component of the location. As a result, there is a roughly elliptical band of high likelihoods with nearly the correct range, but because the predicted position is uncertain and fluctuates within that band, a higher prediction error is observed for the last hole location.

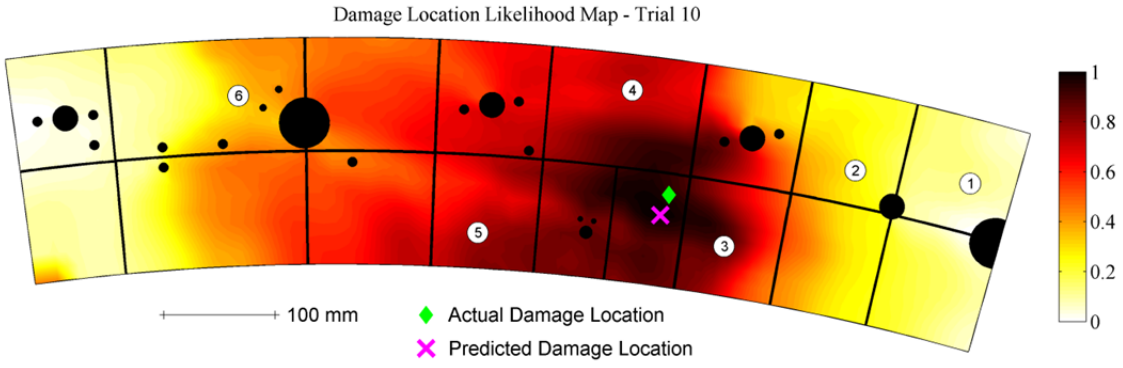

Figure 29 - Likelihood map for damage location prediction, fuselage rib trial 10

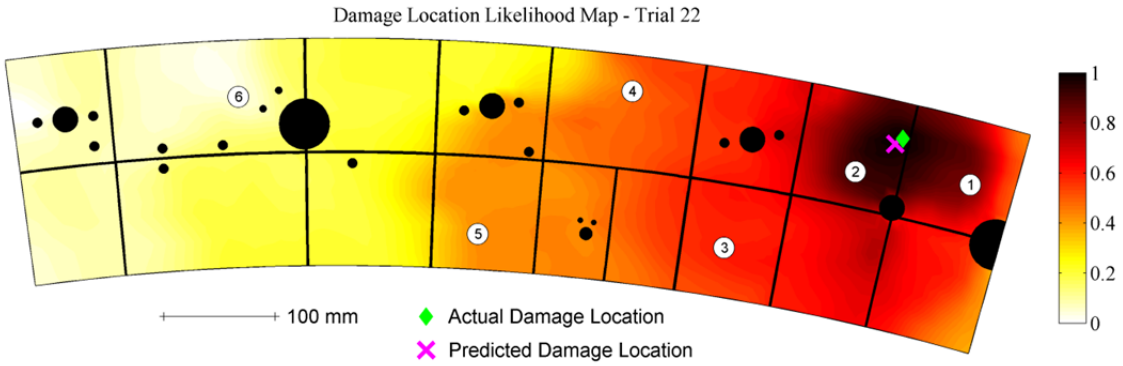

Figure 30 - Likelihood map for damage location prediction, fuselage rib trial 22

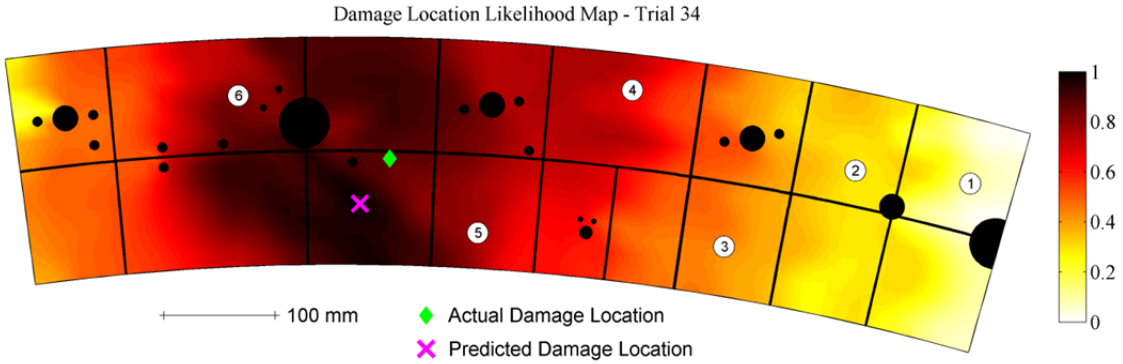

Figure 31 - Likelihood map for damage location prediction, fuselage rib trial 34

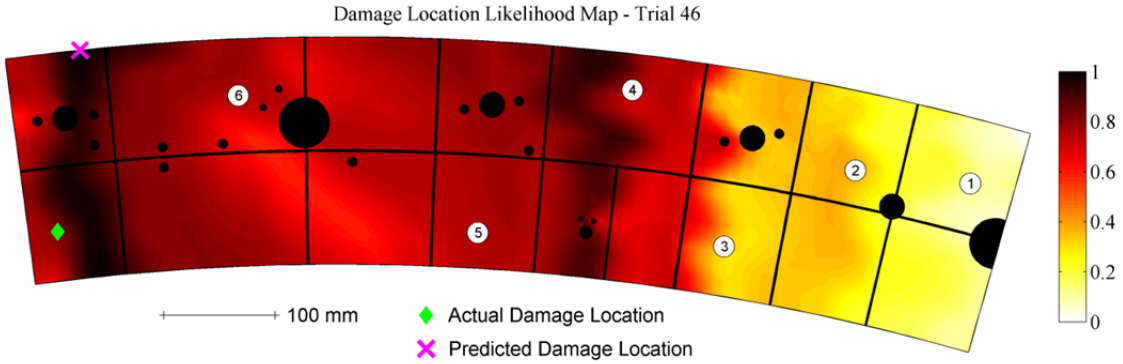

Figure 32 - Likelihood map for damage location prediction, fuselage rib trial 46

3.3.3. Sensor Fusion

As was introduced in Section 2.7, sensor fusion is defined as the process of integrating data from multiple sources in a way that enhances the decision-making performance of the SHM system. Which techniques will produce the most effective sensor fusion is an application-dependent question. For damage detection, the objective is to increase the class separability between the undamaged and damaged data. This principle is illustrated graphically in Figure 33, which shows two forms of sensor fusion applied to the frame magnet damage. Because in practice the class labels are unknown, a single threshold value must be chosen, drawing a horizontal line across the given plots. Therefore, the best detector is one that minimizes the overlap of the two clusters when projected onto the y-axis. Clearly, the voting fusion technique on the right side does a much better job of separating the two classes in metric value.

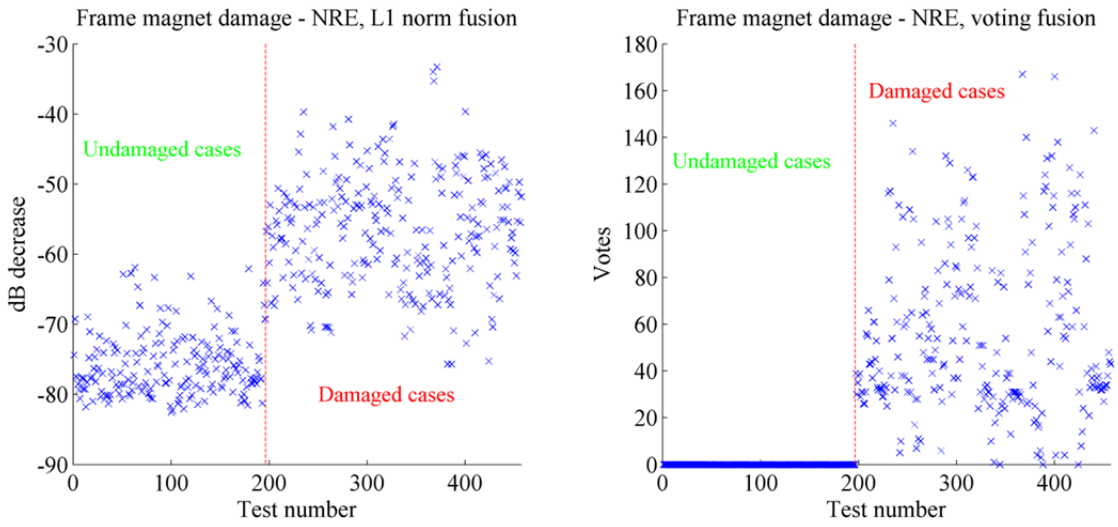

Figure 33 - Two forms of sensor fusion applied to the frame magnet damage: L1 norm (left) and voting (right)

In order to quantify the degree of separation (or the ease of correctly classifying a particular measurement as damaged or undamaged), a metric known as the Fisher Discriminant Ratio (FDR) is introduced [71]. The FDR is defined as:

$$FDR = \frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2} \tag{3.8}$$

where the symbols μ and σ represent the means and standard deviations of the respective classes. The data after fusion is reduced to a vector of scalar values, one for each trial. As such, calculating the mean and standard deviations of the undamaged and damaged classes is straightforward. For the damage types that included multiple levels of damage, all damaged cases were lumped together into the damaged class for the purposes of this comparison. Note that the FDR increases as the means get farther apart and as the variances decrease, both of which make the classes more easily distinguishable from one another. Therefore, the most desirable fusion technique is the one with the highest FDR value.

Table 6 - Comparison of fusion techniques -- probability of detection for false alarm rate of zero

	Frame Magnet FDR	**Frame Bolt FDR**	**Fuselage Rib FDR**
Data-level Fusion			
Energy sum over sensors	3.271	9.538	N/A
Energy sum (pitch-catch pairs only)	4.128	9.441	2.786
Energy sum (pulse-echo pairs only)	0.485	5.000	N/A
Feature-level Fusion			
L^1 Norm (linear mean, correct)	4.576	15.923	5.969
L^2 Norm (sum square)	5.127	22.889	5.133
L^∞ Norm (maximum)	3.674	29.273	4.539
Decision-level Fusion			
Simple voting	2.138	3.956	8.114
Voting - L^1 Norm on Receivers	0.838	2.649	6.363
Voting - L^2 Norm on Receivers	3.346	0.487	9.016
Voting - L^∞ Norm on Receivers	3.736	1.141	7.336
Voting - L^1 Norm on Actuators	0.810	2.617	6.306
Voting - L^2 Norm on Actuators	3.119	0.481	10.140
Voting - L^∞ Norm on Actuators	3.431	1.137	6.793

A selection of sensor fusion techniques from each level were applied to the test results from the bolted frame and the fuselage rib. The resulting comparison of FDR results is presented in Table 6. All of the voting fusion techniques, it should be noted, require threshold values for each sensor pair. Using a separate threshold for each of the N^2 sensors naturally requires more baseline data be retained than simply one threshold value on the final fused feature value. Finally, none of these methods seek to add information to the detection problem

– they simply map the information in the data to a decision, hopefully in the most effective way possible.

From the fusion results, it appears that the best technique is different for each type of damage. The best performing fusion technique for the frame magnet damage is the L^2 norm, whereas for bolt damage in the same structure the best technique is the L^∞ norm. In the fuselage rib, the voting using L^2 norm on actuators is the most effective in this experiment. The differences in optimal fusion techniques may be due to a number of factors: the specific differences in scattering behavior of the different damage mechanisms and the sensor density with respect to the size of the structure are the most likely candidates. A more detailed characterization of the scattering properties may be an important step in formulating the optimal fusion techniques *a priori*.

The different matrix norm techniques seem to exhibit similar performance, despite indicating vastly different approaches to integrating data. The maximum, or L^∞ norm, takes only information from the sensor that shows the highest possibility of damage, which may be why it is more effective in the very sparse array of the frame structure. On the other extreme, the mean technique weights every sensor equally in calculating the final feature value. Voting schemes do seem to work very well for each structure and damage type. The hybrid methods voting/matrix norm methods are sometimes more and sometimes less effective, depending on the particular case considered. The challenge when using a voting scheme for a SHM system is determining what the threshold values for each sensor pair should be, since the final probability of detection is only as good as the selected threshold values. Data-level fusion did not seem to yield particularly good results for any of the damage cases. Because of the details of the data acquisition process, the pulse-echo waveforms are substantially noisier, and for the frame structure it can be seen that their inclusion can actually hurt the overall detection performance.

Of course, it is important not to overgeneralize these trends from limited sample sizes, but this comparison is a good empirical starting point towards developing optimal sensor fusion for damage detection.

For damage localization, the objective of sensor fusion is instead to optimally combine data through the delay-and-sum process in order to produce accurate location estimates. With the MLE first-arrival localization approach having been validated for the proposed application in Section 3.3.2, the next step is to test the performance of the proposed sensor fusion approach outlined in Section 2.7. Table 7 and Table 8 show results for the MLE formulation with different sensor fusion strategies implemented. Trials 100 and 200 give the reference results obtained using no localization information (i.e. simply picking the geometric centroid each time). Next, the cases using all pitch-catch (PC) channels is presented, which represents the default sensor fusion technique for most delay-and-sum localization. Followed those cases are the results obtained by using the location of the sensor receiving the most votes in the sensor-by-sensor detection threshold process – again, though not a localization technique based on delay-and-sum, it provides a useful reference against which to compare the subsequent techniques. An appropriate UGW SHM localization strategy would be expected to outperform all of the first three techniques. As expected, without appropriate sensor fusion techniques, the ellipse methods of localization (in cases 101 and 201) for such complex structures are very inaccurate, and methods not based on delay and sum (as for 102 and 202) can be more accurate. Finally, results are presented for NP binary fusion and NP direct scaling fusion, both with and without the geometric feature modeling, which will be described in Section 4.3. For simplicity, the parameters were calculated using NP direct scaling only and kept constant for both of the fusion methods presented. Two metrics of localization performance are presented for each case: the average prediction error and the AUC metric.

Table 7- Frame structure localization results for various fusion and modeling choices

Run #	Mean Error (mm)	AUC	Delay (μs)	Attenuation	Fusion Technique
100	649	0.562	0	1	Centroid (no info)
101	251	0.832	0	1	All PC channels
102	133	0.911	0	1	Most popular sensor
103	136	0.911	0	1	NP binary
104	121	0.919	0	1	NP direct scaling
105	69.4	0.955	5.8	0.595	NP binary
106	65.4	0.957	5.8	0.595	NP direct scaling

Table 8 - Fuselage rib structure localization results for various fusion and modeling choices

Run #	Mean Error (mm)	AUC	Delay (μs)	Attenuation	Fusion Technique
200	254	0.703	[0, 0]	[1,1]	Centroid (no info)
201	196	0.780	[0, 0]	[1,1]	All PC channels
202	140	0.838	[0, 0]	[1,1]	Most popular sensor
203	78.9	0.916	[0, 0]	[1,1]	NP binary
204	59.6	0.931	[0, 0]	[1,1]	NP direct scaling
205	72.5	0.922	[1.44, 0.010]	[0.995, 0.995]	NP binary
206	56.2	0.936	[1.44, 0.010]	[0.995, 0.995]	NP direct scaling

For both the frame and rib structures, NP direct scaling outperformed NP binary sensor fusion, just as both of those techniques outperformed the results using all sensor pairs. The implementation of some form of feature modeling likewise produced a drop in the localization

error, most significant for the frame structure. The fuselage rib was not as dramatically influenced by feature modeling, most likely because the quantity and diversity of the features made the parameter estimates less precise. The final average localization error in the frame structure was a mere 65 mm, compared to a maximum path distance in the structure of 1.8 meters. The UGW wavelength in the frame is about 25 mm—an average error about 2.5 times the wavelength is an excellent result. The results in the fuselage rib structure are similarly excellent, with the average error decreasing from 196 mm with no sensor fusion or feature modeling down to just 56 mm when NP direct scaling and feature modeling were implemented.

A portion of this chapter has been published in Structural Health Monitoring: An International Journal, Colin Haynes, Michael Todd, Eric Flynn, and Anthony Croxford, 2013. The title of this paper is “Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation”. The dissertation author was the primary investigator and author of this paper.

A separate portion this chapter has been has been submitted for publication in Mechanical Systems and Signal Processing, Colin Haynes and Michael Todd. The title of this paper is “Enhanced Damage Localization for Complex Structures through Statistical Modeling and Sensor Fusion”. The dissertation author was the primary investigator and author of this paper.

Chapter 4
Empirical Estimation of Wave Scattering

Any inhomogeneity in the mechanical impedance of a structure will cause elastic waves to scatter in a way that is governed by material properties, wave properties, and the specific geometry of the scatterer. This chapter presents an empirical approach to estimating scattering profiles, both of specific types of damage and of geometric features. The aim of developing such profiles is to facilitate the SHM process by allowing quantitative prediction of how ultrasonic signals will interact with scatterers. Understanding these interactions can aid in SHM system design (such as optimal sensor placement and tailoring actuation signals), as well as in more accurate damage identification. Damage localization has been shown to benefit from models of directional types of damage, as well as modeling of the geometric features themselves [36, 72]. For damage characterization, modeling is indispensable to producing a prediction of the structural state [73].

A scattering matrix is a way of quantifying the magnitude and directionality of such an interaction, whether due to defects or to structural features. Since the terminology of the "scattering matrix" was borrowed from electromagnetic theory [74, 75], the interaction of ultrasonic waves with scatterers has been studied by many authors. An overview of some of the different ways these phenomena may be described was given by Wilcox and Velichko [35].

Scattering matrix methods typically compare the magnitude or energy of the wavefront scattered in all directions as a function of incident direction [76]. Usually these representations consider the far-field scattering behavior of the target feature [35]. The most general description of the scattering of a feature would be a function of incident angle, reflected angle, wave mode, frequency, and even possibly the radius at which the response is measured. In focusing on empirical techniques for estimating these scattering matrices, this work may complement theoretical or numerical approaches where solutions can be found. For more complex situations where analytical models cannot be constructed, experimental techniques may be the only feasible option. Although there are many sensing modalities that may be used, two examples are presented here—bolt bearing damage interrogated by fixed piezoelectric sensors and rivet scattering inspected through scanning-laser-generated ultrasound. Finally, the chapter concludes with a section on how geometrical models may be used to improve the damage localization process.

4.1. Bolt Bearing Damage Scattering Characterization

For an example of estimating scattering matrices from experiment, the case of bolt bearing damage in composite materials is considered. In particular, interaction of Lamb waves with bolt bearing damage in carbon fiber reinforced polymer (CFRP) bolted joints has been studied parametrically. The use of CFRP materials for aerospace structures has become widespread in due to their superior specific strength and specific stiffness. Moreover, mechanically-fastened connections, and specifically bolted joints, are the predominant connection type for primary structures made of advanced composites [77]. However, such connections are not as well understood as their isotropic counterparts, and the characterization of the failure mechanisms of CFRP bolted connections remains an active research area. As

CFRP bolted joints become more common features in aircraft design, it becomes increasingly important to develop systems capable of monitoring their failure mechanisms. Bolt bearing failure is of particular interest because connections that fail in this mode tend to have the highest strength, and so this is a preferable failure mode from a design perspective [78].

If the goals of SHM are to be realized, inspection techniques must exist that are sensitive to the damage mechanisms of interest – hence, the first objective of the current work, which is to establish the sensitivity of UGW inspection to bearing damage. Once that objective has been satisfied, the second goal is to determine the scattering properties of the damage as a function of the various parameters that may affect the UGW interaction.

Several of the aspects of the problem have been studied previously, including: the interaction of guided waves with various scatterers, the effect of applied load (i.e. the stress state of the material) on guided wave propagation [79, 80], bolt bearing damage in CFRP materials [77, 81], and the inspection of bolted joints [82, 83]. However, this study was the first, to the author's knowledge, to propose inspecting bolt bearing damage in CFRP through the use of a UGW system.

4.1.1. Experimental Methods

To begin experimental investigation, two CFRP test specimens were fabricated, which will be referred to as plates 101 and 102. The laminates consisted of eight T700S/2592 prepreg layers in a $[-45\ 0\ 45\ 90]_s$ stacking sequence. The final laminate measured 200 mm x 120 mm x 1.15 mm with 6-mm-diameter holes drilled to bolt to the tensile testing fixture in such a way that will introduce damage at a predictable location. The entire test fixture is shown in Figure 34. Several 40-mm-square CFRP washers on either side of the plate allow the other connecting elements to grip the specimens without contacting the surface of the plates, which is necessary to avoid putting pressure directly onto the sensors that are used to instrument them. The

double-shear testing configuration puts the highest stress on one side of a single target hole, ultimately introducing bolt bearing failure.

Figure 34 - Overview of tensile testing apparatus, plate 101

Each test specimen was instrumented with a uniformly-spaced circular array of eight macro fiber composite (MFC) sensors. MFCs are favored in composite applications because of their conformability to curved surfaces, their directional actuation characteristics, and their durability. The usefulness of MFCs in guided wave SHM for such applications has been demonstrated by many studies [84, 85]. The MFCs used for this experiment were all of the M2814-P2 type cut to a final dimension of 14 mm wide by 6 mm long, bonded to the surface of the laminate with a standard two-part epoxy adhesive. The sensors were positioned such that the front edge of the MFC was 30 mm away from the center of the target hole, oriented with the

principal axis of the MFC passing through the center of the hole. An overview of the sensor layout is given in Figure 35. Positioning the MFCs in a uniform circular array around the target hole where the defect was to be introduced facilitates the investigation of the directionality of the interaction. For reflections from a perfect cylindrical hole, the scattering matrix result can be calculated theoretically (at least for isotropic materials) or determined numerically using techniques such as finite element modeling or semi-analytical finite element modeling [32]. However, such results are not readily available in the literature for anisotropic materials, especially for complex defect structures like CFRP bolt bearing damage. As a result, the optimal angles of interrogation under different conditions are determined here by taking an experimental approach to estimating the scattering matrices.

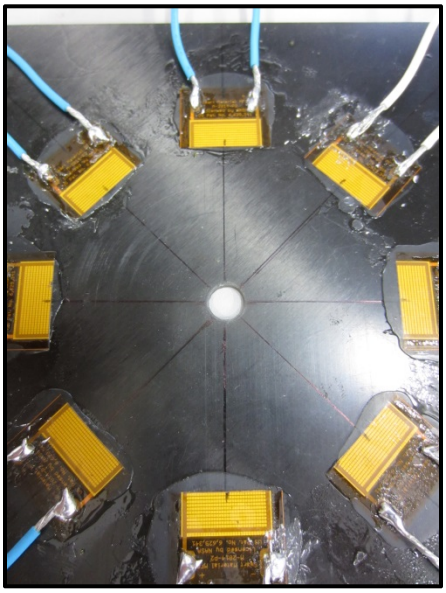

Figure 35 - Overview of MFC array, plate 101

Ultrasonic data were acquired by using each MFC to actuate sequentially, receiving with all other sensors upon each actuation. The actuation signal was chosen as a three-cycle Gaussian-modulated sinusoid. Generally, the optimal frequency for actuation of a UGW SHM

system may be selected by considering the sensor properties, the material properties of the medium, and the ratio of the target defect size to the wavelength. Determining the amplitudes and directions of the guided wave modes generated by an MFC sensor is a difficult question, though Collet *et al*. have presented some useful information obtained from finite element analysis [86]. From the results of that study, it may be inferred that the forward-propagation direction (the direction that will interact first with the bolt hole) is dominated by the S_0 mode of propagation. While the present work is not focused on detailed analysis of the wave propagation, this information is helpful for predicting the approximate group velocity, and thus the time window during which the scattered waves will be received. Indeed, rather than attempt to calculate the dispersion curves for this particular quasi-isotropic CFRP layup, the velocity of the dominant mode was determined from a simple experiment in a feature-less plate of the same material. Bonding two identical MFC sensors 60 mm apart on the surrogate plate, the same actuation parameters were used to launch and receive a pulse. From a basic time-of-flight calculation, the velocity was found to be about 4.0 km/s.

For plate 101, the data acquisition system did not support pulse-echo measurements, but all pitch-catch channels were recorded at a sampling frequency of 10 MHz. A center frequency of 300 kHz was chosen primarily because it provided a relatively high-amplitude response for the forward-propagating wave. Plate 102, however, was tested with a different data acquisition system, enabling the acquisition of both the pulse-echo responses and a parametric study of different actuation frequencies to be undertaken. The second data acquisition system operated at a sampling frequency of 2.5 MHz.

Each waveform was then denoised by applying a matched filter—that is, convolving the received signal with the input signal. This filtering procedure, as mentioned before, has the effect of eliminating frequency content other than that contained in the actuation signal, which

implicitly assumes linear data acquisition. Each signal is then enveloped by taking the magnitude of the Hilbert transform of the signal. Given that the sensor locations and wave speed have been determined, it is a straightforward calculation to determine the time window inside which the direct-arriving wave should fall. The other arrivals are either insensitive to damage (because the path does not interact with the hole) or more sensitive to non-damage-related parameters than the direct arrival (because the path includes edge reflections). Therefore, appropriate time windows are applied to select the proper wave packet, with the same window being used for all sensor pairs (since the nominal distances are the same and directional velocity differences may be neglected). Because no model for the effect of bolt bearing damage on guided waves in composites is available, it is difficult to determine an optimal detector *a priori.* Therefore, as a detector that has been shown to be effective in situations with minimal prior knowledge, the NRE metric presented in Section 2.2 is applied to the present situation as well.

Both of the specimens were subjected to a displacement-controlled tensile test until three different damage levels were introduced. The first level corresponded to the first significant load drop observed in the load-displacement curve. The following two damage levels were based on loading to a prescribed displacement beyond the initial load drop. While every effort was taken to ensure the consistency of the damage modes compared, the two plates used a slightly different test procedure. Tests for plate 101 were conducted at the University of Tokyo. After each load increment, plate 101 was unloaded and removed from the bolt fixture for the ultrasonic interrogation. Because all of the ultrasonic tests were conducted on the plate once removed from the test fixture, the boundary conditions were kept consistent. The damage in plate 102, by contrast, was introduced at the University of California, San Diego via a single tensile test, during which the loading was merely paused to take ultrasonic readings. Therefore,

the lateral restraint condition remained consistent for the plate 102 tests, although it was necessarily different from the plate 101 condition. Lateral restraint has been shown to particularly influence damage structure and the load level when damage initiates [87]. However, introducing the damage in a single tensile test allowed the effect of load state on the UGW inspection to be considered.

Figure 36 and Figure 37 show the load-displacement curves for each of the tests conducted. During each pause at the discrete loading levels for plate 102, the test apparatus relaxed slightly, causing the load drops for constant displacements that can be seen at regular displacement intervals (particularly at higher load values). However, the ultrasonic interrogation was not performed until the load level was stable enough (relative to the load increment of the tensile test) for all sensors to take their readings at the same nominal load state.

Figure 36 - Load vs. displacement curves for plate 101

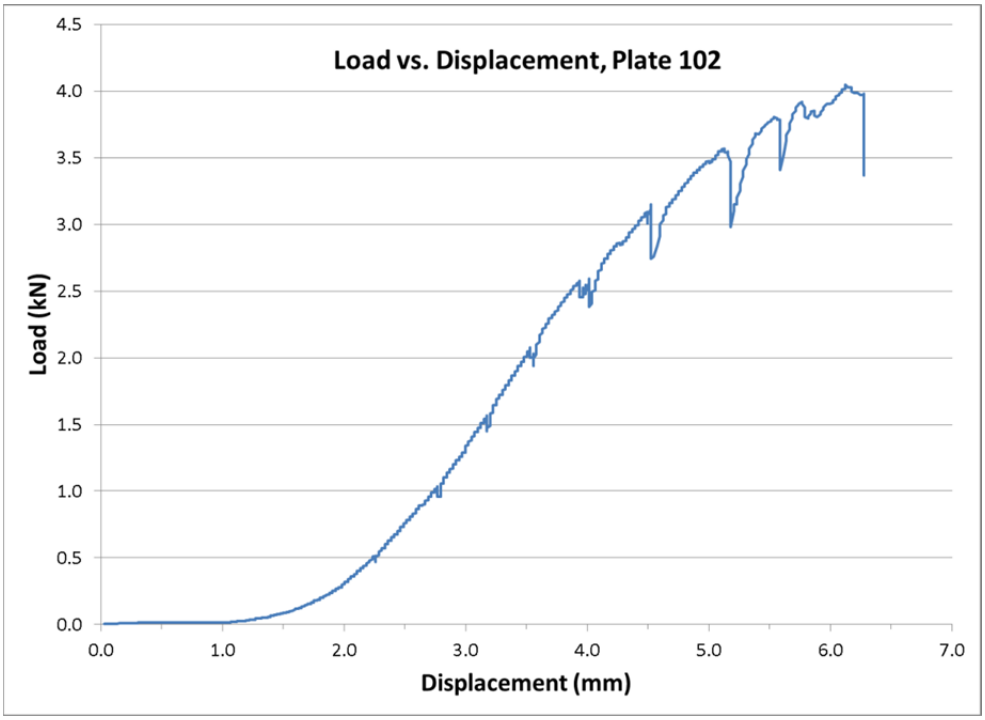

Figure 37 - Load vs. displacement curve for plate 102

Following the plate 101 tensile test, the specimen was cut along the line of damage progression, polished, and photographed with an electron microscope in order to capture the cross-section of the damage. The resulting photograph is presented in Figure 38. Note the periodic buckling of the zero-degree fibers (second layers from the outside), as well as the "X" patterns in the inner layers denoting shear failure [88]. The damage microstructure is presented to provide a more specific understanding of the type of damage that the UGW inspection attempts to identify in this experiment. The particular damage structure is not of primary importance, however – each specimen will exhibit slight variations to which a SHM system must be robust to be of practical interest.

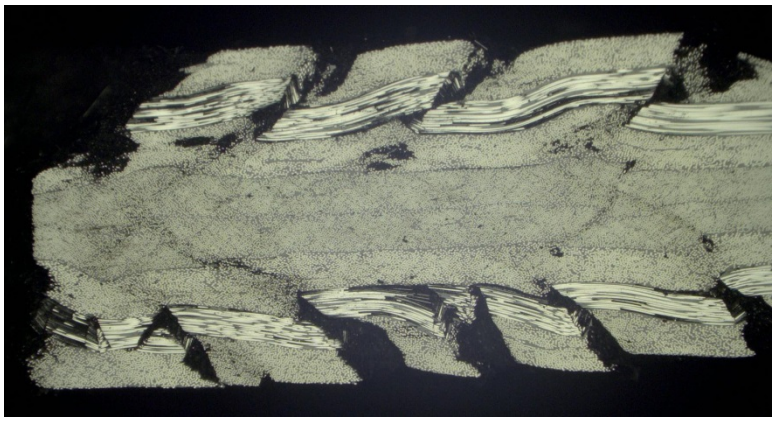

Figure 38 - Electron microscope photograph of bolt bearing damage cross-section

4.1.2. Scattering Results

For plate 101, three baseline trials were acquired before any loading began. Three more trials were conducted after each of the five load levels, of which two remained in the elastic region and three introduced progressive levels of damage. For plate 102, 18 trials were run before any loading. Among these baselines were various configurations, including trials with the specimen being placed in and removed from the bolt fixture. After this set of data was recorded, all further trials were conducted with the specimen in the test fixture under different levels of load. A baseline tensile test was then conducted to 2 kN, well within the elastic region of the specimen, with two baseline trials at every 0.5 kN interval. Finally, the main tensile test was conducted to failure (as shown in Figure 37), with trials at each 0.5 kN interval of loading also. These trials compose the undamaged data and three levels of damaged data, with eight additional unloaded trials after the damage had been introduced.

A representative group of waveforms taken from plate 102 is included in Figure 39. The figure shows the filtered and enveloped waveforms recorded by the 7-6 sensor pair for all the trials during the final tensile test. The waveforms are labeled according to the load state under which they were taken, which include all cases from unloaded to the introduction of

damage at the end of the test. The boundaries of the time window applied to the signal are denoted by green lines – the peak that falls between the lines is the reflection of the dominant mode with the target hole. The wave packet arriving around 10 μs is the result of the proximity of the MFCs to one another; despite their directional nature, some energy still travels directly between adjacent sensors, thus arriving earlier than the intended first arrival. Arrivals after 30 μs represent reflections from the boundaries of the small test specimen. Closer inspection of the peak of the wave packet of interest, shown in the inset, reveals that the amplitudes vary both with load and with damage level. An analogous plot is given in Figure 40 for sensor pair 2-6. This pair, rather than being adjacent, is on opposite sides of the target hole. As such, the first peak observed in the target hole reflection, because that path is also the direct path. From the inset, it may be observed that the sensitivity of the waveform to load effects is different from the other direction.

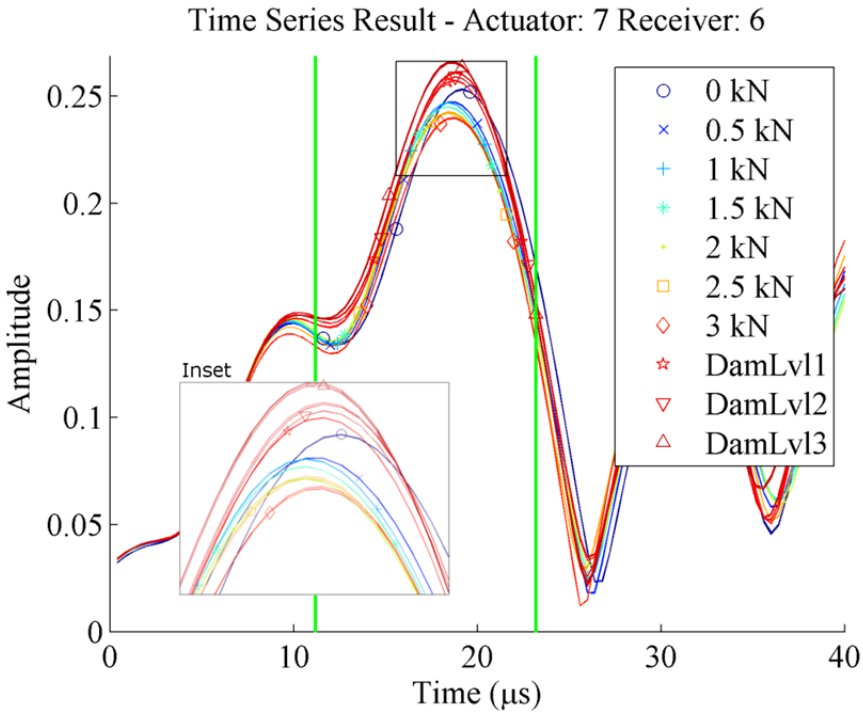

Figure 39 - Time series results for sensor pair 7-6, plate 102

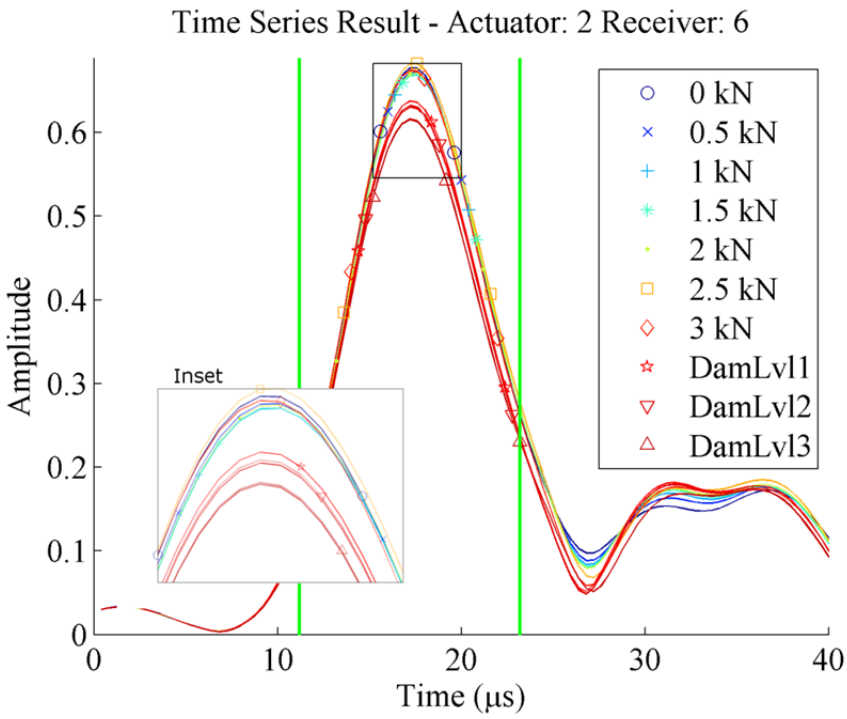

Figure 40 - Time series result for sensor pair 2-6, plate 102

Now that the underlying data has been examined, scatter plots are generated to attempt a quantification of the scattering directionality under various conditions. One representative scatter plot is given for each specimen based on the first-arrival-amplitude of the target hole reflection. Note that the coordinate system is defined such that the damage progresses in the 180° direction. Each plot can only show the results for one actuator, which is labeled in the figure with a red square. All examples here are using sensor 6, and each sensor's response plotted in the direction of that sensor. The plots for plate 102 also include pulse-echo channels, although these channels have a higher noise floor because they are significantly influenced by electro-magnetic interference effects from the actuation. While the amplitudes of these pulse-echo channels are therefore more variable than the respective pitch-catch channels, some cases still showed statistical separation between damaged and undamaged cases. As a result, the pulse-echo channels have been included where available.

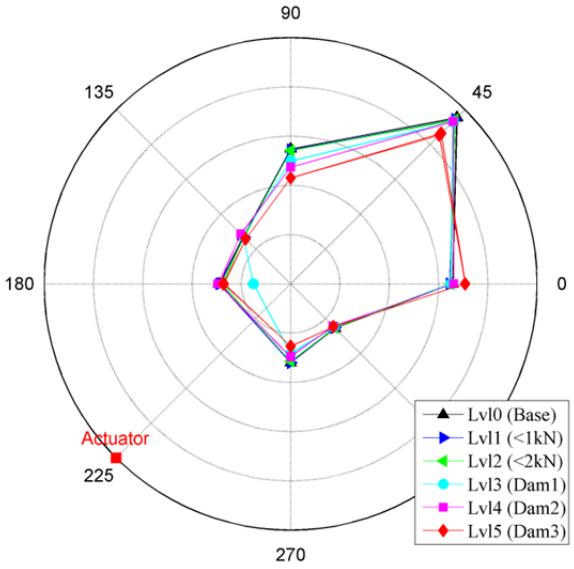

Figure 41 - Amplitude scattering matrix, actuator 6, plate 101

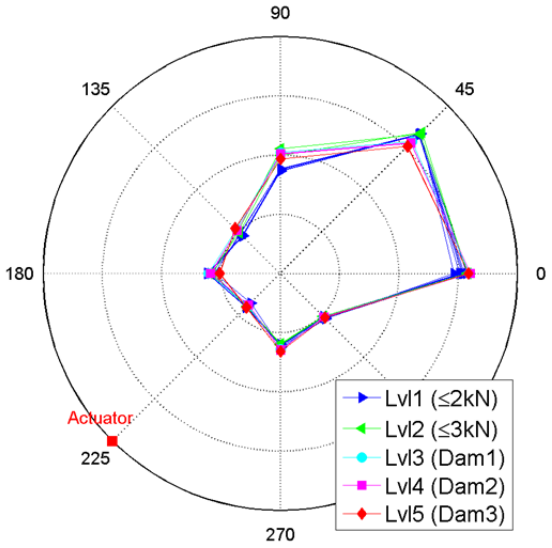

Figure 42 - Amplitude scattering matrix, actuator 6, plate 102

It is evident that the experimental scattering matrix results, even for the undamaged cases, are not quite symmetric. Only a slight directional dependence would be expected from the quasi-isotropic laminate (as demonstrated by Salas and Cesnik [13]), and the results are not entirely consistent with the theoretical predictions for this case. The presence of the bolt fixture

is unlikely to be the cause, since the plate 101 tests (shown in Figure 41) were conducted after removing the specimen from the bolt fixture (thus there was no contact around the bolt hole). Plate 102 tests (shown in Figure 42) were all conducted under a consistent lateral restraint condition since the plate was not removed from the fixture, so it is reasonable to conclude that lateral restraint is not the most likely cause of the asymmetry cause. Instead, discrepancies among the MFC sensors such as slight variations in location, orientation, and bond condition are probably the main factors contributing to the asymmetry. An uncertainty quantification study was also conducted to study the main effects that impacted the UGW system [89] – the results are given in Section 0. In both cases, there are angles for which some separation between damage states is noted, but it is difficult to determine a definite pattern from the first-arrival-amplitude scattering matrix comparisons directly.

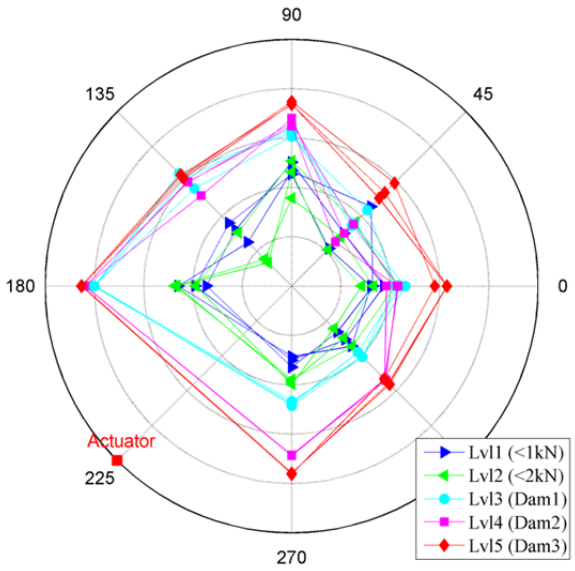

Figure 43 - Baseline-subtraction scattering matrix, actuator 6, plate 101

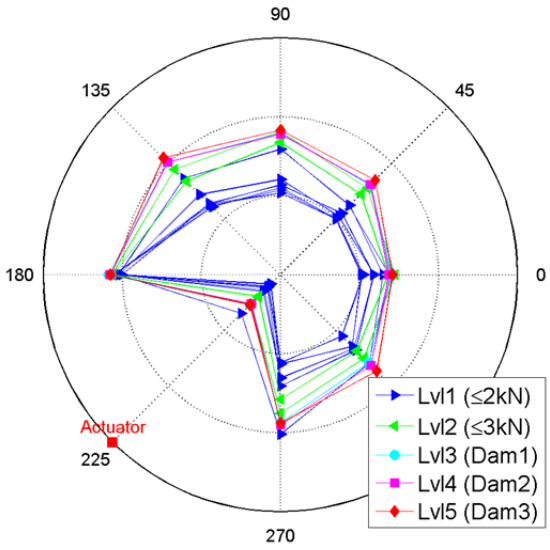

Figure 44 - Baseline-subtraction scattering matrix, actuator 6, plate 102

To visualize the damage progression another way, baseline subtraction was applied to the data and the NRE metric was used to construct scattering matrices. The resulting ratios are typically expressed in decibels, so the minimum value (across all sensors and trials) was subtracted to shift the negative dB values (which are inconvenient for polar plots) to positive values. The resulting plot indicates the smallest deviation from the baseline by the smallest magnitude of the radial component, which is still expressed in dB units. The plots shown are qualitative in nature, meant to highlight the patterns of residual scatter for each angle of interrogation.

First, it is clear from the baseline-subtracted results that the residual energy becomes greater as the damage level increases. Certain angles exhibit a more pronounced change with damage level, whereas the distinction is not as clear for others. Comparing the results for the two specimens against one another, however, the observed patterns are quite dissimilar. As demonstrated by Figure 43 for plate 101, the reflections that bounce back from the damaged side of the hole produce a relatively large separation among damage cases. However, the result

from plate 102 shown in Figure 44 suggests the opposite – waves propagating past the hole and received on the opposite side from the actuator produce the best separation. In the next section, the angular dependence is investigated more quantitatively and a possible cause of the discrepancy is discussed.

4.1.3. Angular Dependence of Scattering

To provide a quantification of performance at each input and output angle, the class separation between the data from each damage level and the undamaged data is computed for each sensor pair. Because of the small sample sizes when considering each sensor pair and frequency independently (there are relatively few trials for each level, particularly for the damaged cases), it is difficult to estimate the intra-class variability accurately. Additionally, the NRE metric is non-negative and should always increase as damage is introduced. Therefore, a modified version of the Student's t-test statistic is introduced as an appropriate class separation metric:

$$\text{Modified Student's t-test statistic } = \sum_z \frac{\delta}{\sigma^{(u)}}$$
$$\delta_z = \begin{cases} \mathbf{v}_z - \mu^{(u)}, & \mathbf{v}_z - \mu^{(u)} > 0 \\ 0, & \mathbf{v}_z - \mu^{(u)} < 0 \end{cases}, \tag{4.1}$$

where $\sigma^{(u)}$ is the standard deviation of the undamaged class, $\mu^{(u)}$ is the mean of the undamaged class, and $\mathbf{v}_z$ represents each metric value for the damaged class. Note that if the sample value is less than the mean of the undamaged class, the separation is set to zero so that it does not contribute to the summation. The class separation metric is not intended to be used for detection—it is instead used to compare various acquisition parameters given knowledge of the class labels for each test.

Figure 45 (which is taken from plate 101 data) compares the class separation for each sensor pair. Separate plots are given for each of the three damage levels, with all of the data acquired at a 300 kHz actuation frequency. The values of the distance measure may seem quite high because there is excellent separation among classes for specific sensor pairs. However, the scale of the vertical axis is not as important as the comparison between interrogation angles. From this point of view, it is evident that the best separation for all damage levels is obtained by using the sensors on the damaged side of the hole (sensors 4, 5 and 6) both to actuate and to receive. That is, the results are best when waves are generated from the damaged side, reflecting off of the damage and being received by a nearby sensor. These sensor pairs will henceforth be referred to as "backscatter angles".

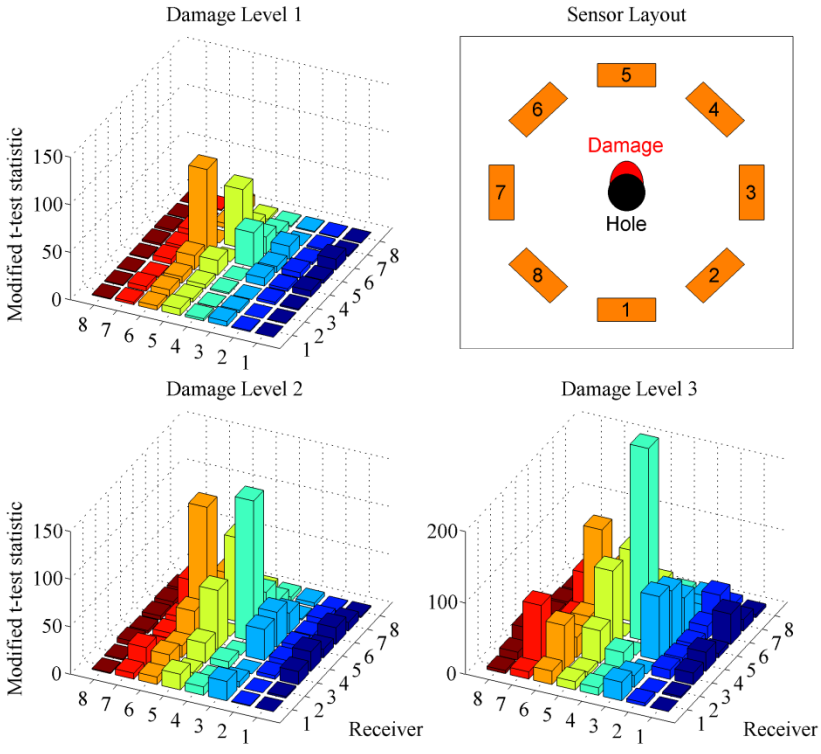

Figure 45 - Interrogation angle comparison for plate 101, 300 kHz

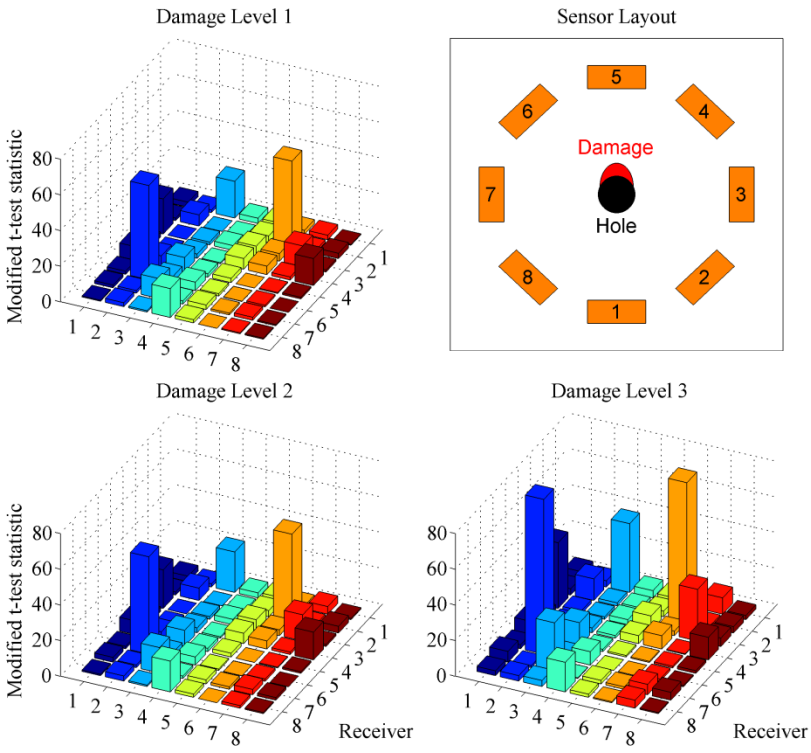

Figure 46 - Interrogation angle comparison for plate 102, 300 kHz

Figure 46 shows the analogous data for plate 102, using only the data acquired with the 300 kHz actuation frequency to match the previous plot. However, the observed pattern is not the same for the plate 102 tests. Rather than the backscatter angles, the data suggest that the through-transmission paths (the sensors separated by 180 degrees) are the best choice for separating damaged trials. Because the plates were fabricated and instrumented according to identical procedures, the most likely cause for the dramatic difference plate 102 was under load when the measurements were taken, while the plate 101 tests were always conducted with the load removed. The loading results in a different contact condition between the bolt and the bearing surface of the plate, causing the waves to interact with the hole differently. This hypothesis was confirmed by comparing results from plate 102 taken in an unloaded state (removed from the bolt fixture) before and after the introduction of damage. The resulting

pattern, given in Figure 47, is much more similar to the plate 101 results than to the plate 102 results obtained under load. Such influences may be difficult to predict in advance with modeling approaches for predicting the scattering from defects, such as that presented by Fromme and Sayir [32]. Results of this type highlight the motivation for taking the empirical approach implemented here. Of course, any given practical application will likely be different as well; for example, most connections involve more bolts than just one, resulting in more complex contact conditions. These considerations, while worthy of further investigation, are beyond the scope of this dissertation.

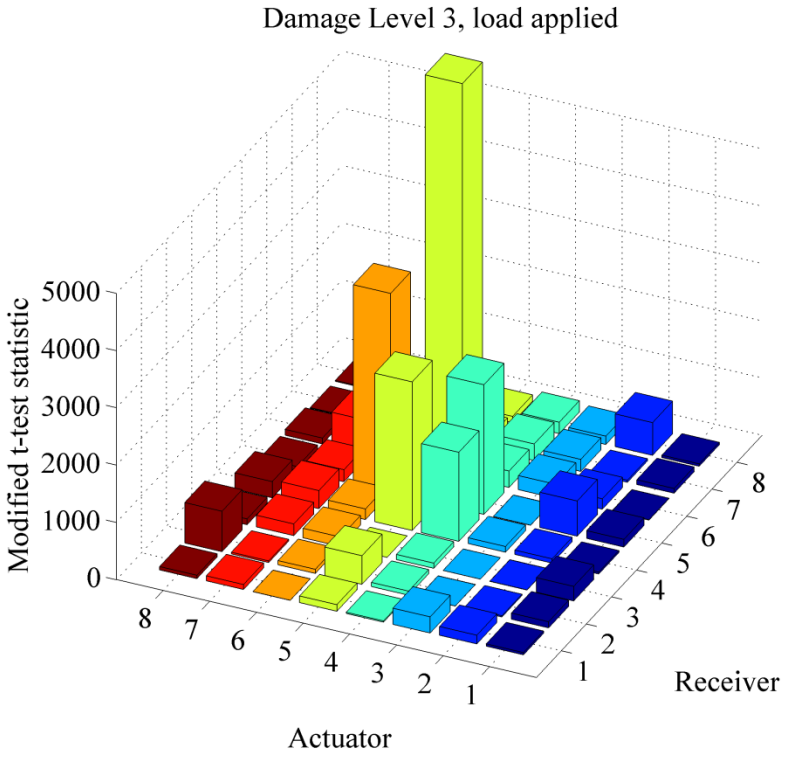

Figure 47 - Plate 102 unloaded results, damage level 3 only

4.1.4. Frequency Dependence of Scattering

With the optimal angles of interrogation thus established, the effect of actuation frequency is now examined for the plate 102 data. The same modified t-test metric as before was computed over the average NRE level of all sensor pairs. The results are given in Figure 48, where the class separation values for each damage level (1-3) are stacked on one another. Again, the higher the t-test metric, the better separation is achieved – with that in mind, the results imply that there is a peak in system performance around 325 kHz. The trend in Figure 48 is quite similar to the trend in the amplitude of the received signals (which is dependent primarily on the sensor properties). That higher signal amplitude would improve detection performance is not surprising. On the other hand, the wavelength of the interrogation signal (which is often a parameter of prime importance in achieving sensitivity to damage) varies between approximately 2.5-6 mm or about 0.4 to 1 hole-diameters. Therefore, the effect of signal amplitude seems to be more important in this case than any sensitivity gained by matching the wavelength more closely with the damage scale. Finally, this comparison provides some justification for the 300 kHz actuation frequency selected for previous comparisons.

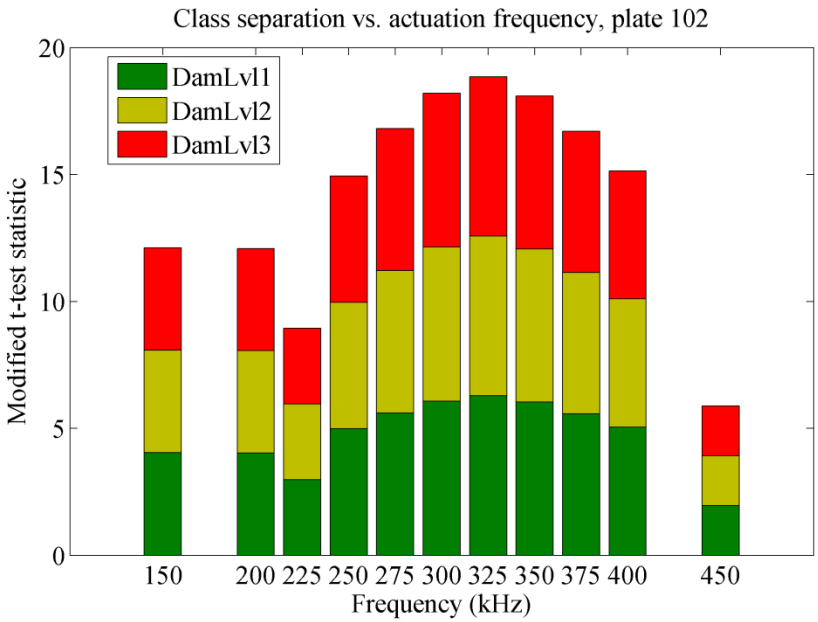

Figure 48 - Class separation vs. actuation frequency, plate 102

4.1.5. Load Dependence of Scattering

For the plate 102 tests, the recorded signals include the effect of applied load, and the majority of the tests were conducted without removing the specimen from the tensile testing machine. Since results have already been presented that suggest the scattering properties are a function of the load state, the relationship is explored further. Applied stress on a material is known to have a nonlinear and anisotropic effect on wave propagation, referred to as the acoustoelastic effect. This effect causes slight shift in both the group and phase velocity of UGW; measured in all directions for a unidirectional CFRP laminate, one study found the maximum velocity change to be around 1.9×10^{-4} /MPa [90]. Using this value as an estimated upper bound on the acoustoelastic effect in the present test article, it can be shown that such a phase difference would be nearly impossible to observe given the sampling rate of the present system and the small size of the specimen – the details of that calculation follow.

It has already been noted that the estimated group velocity in the material is 4.0 km/s. Therefore, the acoustoelastic velocity change would be around 0.74 m/s/MPa. Bearing damage initiated in this specimen at a load of 3.5 kN. Considering this force to be applied directly to the half-cylindrical surface of the bolt hole in direct contact with the bolt, the maximum stress in the material at failure can be found. However, this stress level is not representative average stress state of the material the UGW propagates through. In fact, much of the material towards the plate edge is unstressed, whereas material between the bolt holes (but far from either one) sees a stress equal to the load over the cross-sectional area. Finite element analysis providing a more detailed picture of the stress state has been carried out by Takeda *et al*. [18]. To construct a conservative estimate of the true average stress on any given path, the average of the maximum stress and the uniform tensile stress is used. Using the foregoing values, equation (4.2) gives the mathematical expression for calculating the change in velocity and thence the resulting shift in the arriving wave (expressed in the number of samples):

$$\begin{aligned} &\Delta v = c\hat{\sigma},\ \hat{\sigma} = \frac{1}{2}\left(\frac{2P}{\pi bd} + \frac{P}{bw}\right) \\ &\Rightarrow \Delta n = \left(\frac{s}{v} - \frac{s}{v+\Delta v}\right) f_s \end{aligned} \tag{4.2}$$

where v represents the nominal wave speed, c the acoustoelastic coefficient, P the tensile load, σ the stress, d the hole diameter, b the plate thickness, w the plate width, s the total propagation distance of the wave, f_s the sampling frequency, and finally n the number of samples until the first arrival occurs. Using the conservative values, the resulting velocity change would be approximately 35 m/s. At a sampling rate of 2.5 MHz, this change in group velocity would result in a time shift of just one-third of a sample in the arriving signals. Given that the true

nonlinear acoustoelastic effect should be lower than this upper bound, in the presence of any noise, the current system cannot detect a difference in response due to this effect directly. Lee *et al*. do propose a way of finding small phase shifts through a zero-crossing method [91], but such an estimation would still not be practical in this case due to the complicated stress state around the bolt hole and higher noise floors due to limitations on signal averaging.

Although the velocity change may not be observable, it has been noted in the plate 102 tests that a significant change is observed in the amplitude of the waves with load. This primary effect is may be due to the interaction of the bolt itself with the guided waves, which can be reasonably assumed to change with load level, or even with the effect of load on the transducer or bonding properties.

Because there is a noticeable change in the received signals for different levels of load, it is necessary to identify such changes and compensate for them in order to detect the formation of bolt bearing damage. The results confirm that it is possible, using the OBS procedure to find the most similar baseline, to classify lower load levels from the load level of the selected baseline. In effect, this is equivalent to using a k-nearest neighbor classifier with k=1. In this way, load states within the allowable operating range can be tracked with this system—all undamaged records in the load range where baselines existed were assigned the correct load state by this procedure. Higher loads for which baselines were not obtained do show an increase in the residual level, which, given proper threshold selection, would provide warning that the load exceeds the maximum allowable. Of course, the practice of using OBS would become more difficult if the load state was more complex than the current uniaxial condition since correspondingly more baselines would be required – however, for the present case, this compensation method proves effective.

Finally, bolt bearing damage causes a further increase in residual level, which can be detected by appropriate threshold selection according to the usual procedure. To demonstrate the damage detection potential of the system, Figure 49 shows the NRE metric values for the plate 102 trials. Only the result for 300 kHz is given, but all of the other frequencies showed similar results (i.e. the residual levels for all damaged tests were greater than the undamaged tests). For the example in Figure 49, it is clear that if a threshold level of around 0.12 had been chosen the SHM system would yield 100% correct detection with no false alarms for this limited sample set. It should be noted that this comparison does not use any intelligent sensor fusion scheme, but simply uses the L^1 norm of the sensor pairs on the feature level. The results are therefore likely suboptimal for an eight-sensor system, but on the other hand, that is not likely the arrangement that would be selected for an actual SHM application. Of course, such a scenario would not be as simple and perfect detection is not be possible in practice, but this result validates the capability of the UGW approach to correctly identify bolt bearing failure.

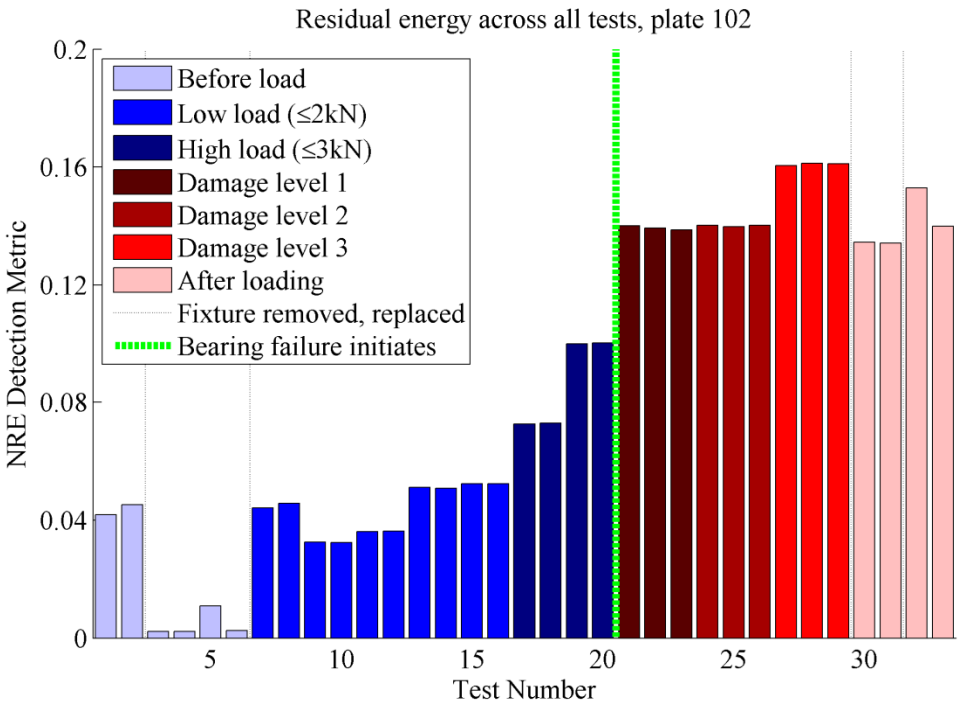

Figure 49 – NRE metric computed for plate 102 trials

A detailed investigation of the scattering properties of bolt bearing damage has thus been carried out. Numerous parameters that may influence the UGW-damage interaction were considered, including the input and output wave angles, actuation frequency (or wavelength), and load state. The results indicate that not only are UGW systems potentially sensitive to this unique form of damage, but that the resulting changes in scattering properties can be quantified. Furthermore, this work has shown that damage detection is not only possible, but that there is the potential to use the OBS approach to monitor the load state of the bolted connection.

4.2. Rivet Scattering Characterization

In building an empirical scattering model, the spatial density of measured data is a critical parameter. Where movable transducers are infeasible, noncontact methods become particularly advantageous. The next example presents work that was done to estimate the

scattering matrix of rivets on an aerospace component, both with and without artificial crack damage.

The sensing modality chosen for the rivet scattering characterization was a scanning-laser-generated ultrasound system [9]. This technique represents a very different approach from the sparse sensor arrays used in the majority of this dissertation. Instead on having a few scattered sensing sites, the wavefield can be visualized with high spatial resolution over the entire area of interest. For quantitatively characterizing the scattering properties, such full-field data is very useful, as will become apparent in the subsequent sections.

The scanning laser system used in this experiment generates UGWs in the structure by directing a Q-switched pulsed laser so that the beam impinges on the surface. These UGWs propagate throughout the structure and are received by a single, surface-mounted PZT sensor. A two-dimensional scanning mirror is then used to perform a raster scan over the area of the structure that is to be inspected. By invoking elastodynamic reciprocity (for the case of linear wave propagation), the resulting data can be interpreted as a full-field measurement of the structural response from UGW actuated at the fixed PZT sensor location.

Scanning-laser-generated ultrasound systems have several advantages which make them attractive for such a study [92, 93]. Such systems can rapidly produce a very high spatial density of information. The systems are noncontact and may be used on structures tens of meters distant, depending on the ability to focus the laser beam. Producing videos of the wave propagation is also a very intuitive way of interpreting guided wave data. However, the technology remains more of a non-destructive evaluation method, since the inspections require that the somewhat bulky and expensive laser system be deployed when inspection is called for. Furthermore, the large quantities of data only amplify the need for adequate signal processing techniques for automating the damage identification procedure.

4.2.1. Specimen Description & Data Acquisition

For this test, a riveted tail section from a military aircraft is considered. The section is made up of a carbon fiber composite material with unknown properties, the skin of which is fixed to the substructure via several rivet lines. Each rivet has a diameter of 5 mm, and they are all spaced at 22 mm apart, on center. An overview of the structure is given in Figure 50. Two rivet lines, comprising 20 total rivets, were chosen for analysis by the laser-generated UGW. Figure 51 shows a detail of the scan area, with the rivet numbering scheme and the surface-mounted PZT sensor location noted.

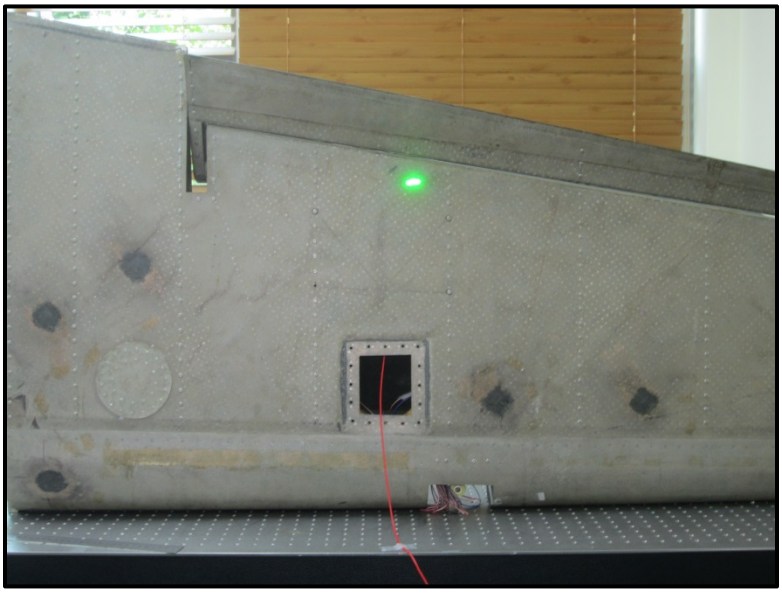

Figure 50 - Riveted specimen overview (laser spot is visible and appears larger because it is in motion)

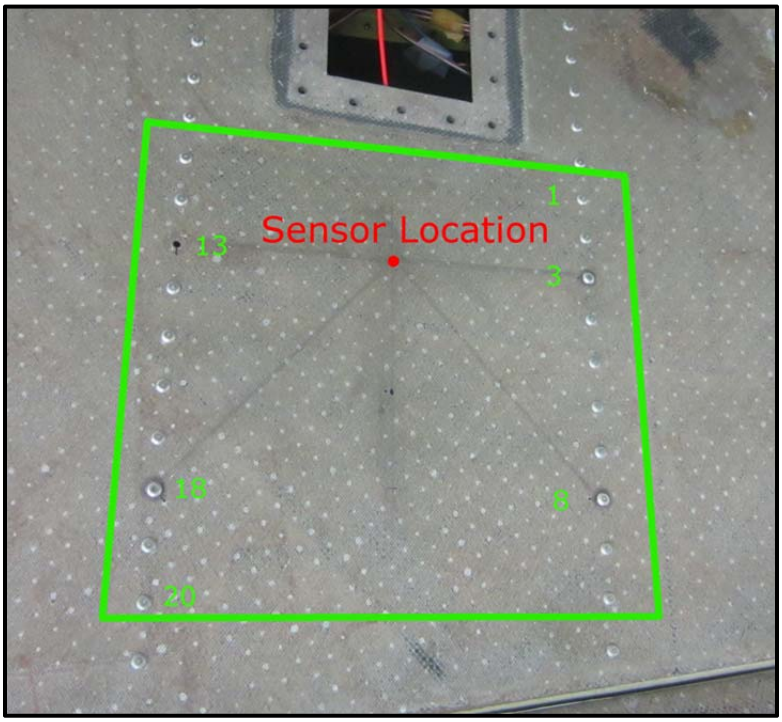

Figure 51 - Detail of scan area, with rivet numbers and sensor location noted

Four rivets, designated as rivets 3, 8, 13, and 18, were artificially damaged before any data acquisition was initiated. At each of the four locations, a simulated crack was machined, each with a different length and direction. Rivet 13 was completely removed. One surface crack was also machined independent from any rivet-related damage, but that damage mode is not considered here.

4.2.2. Polar Wavenumber Processing

Because circular PZT sensors are among the most common transducers for SHM, and because these PZT sensors generate radially-propagating wavefields, it is natural to consider a polar coordinate system in processing full-field data with one sensor as the source. Many small defects or geometrical features (such as rivets) are also characterized by radial scattering patterns. In order to produce a quantitative estimate of the scattering profile of both undamaged

and damaged rivets, a polar-wavenumber processing method is proposed. A graphical representation of the procedure is also given in Figure 52. Rather than (or in addition to) removing the incident wavefield through either time- or wavenumber-domain processing [94], the data is re-interpolated to a polar coordinate system, centered on each rivet in turn. A 3D Fourier transform yields the frequency-wavenumber data, after which a 3D window is applied to filter out all waves that do not have positive (outgoing) radial wavenumbers, effectively isolating waves emanating from the rivet. If desired, another filter may be applied to select only the frequencies of interest, but in this case the frequency filters applied were analog filters on the data acquisition hardware (which band-pass filtered between 40 and 150 kHz). Finally, an inverse 3D Fourier transform returns the data to the space-time domain (or the feature extraction process can be implemented in the wavenumber domain directly).

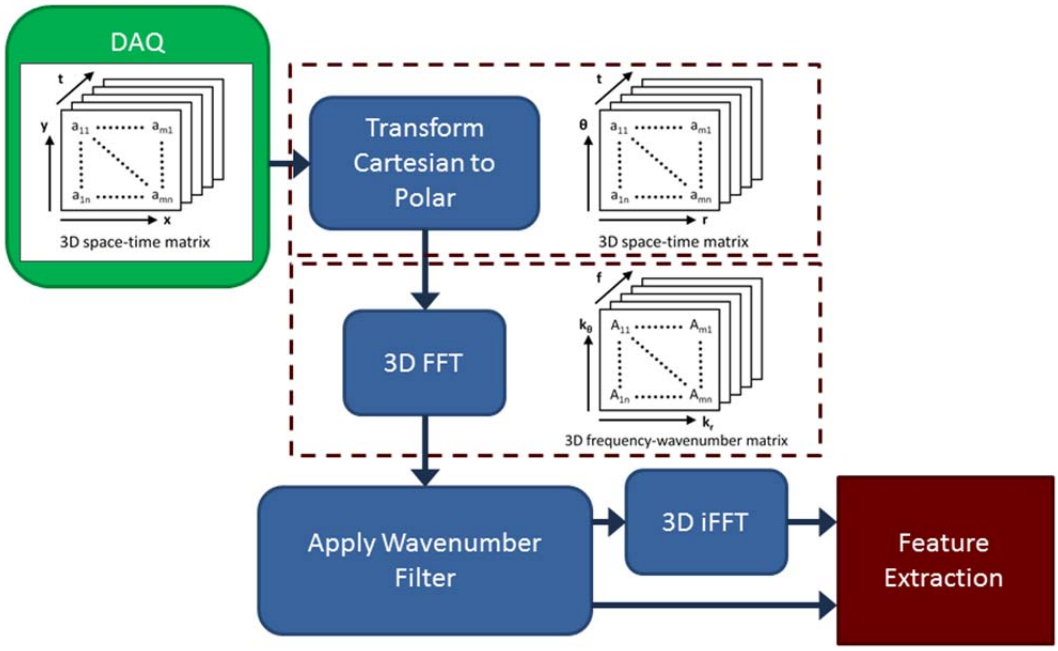

Figure 52 - Flowchart illustrating wavenumber processing procedure

4.2.3. Scattering Matrix Results

Figure 53 presents an overview of the global wave propagation in the riveted specimen. At this specific time (55 μs after the actuation began), the faster-propagating S0 mode has interacted with all of the rivets and begun to scatter. Roughly circular scatter patterns may be observed around each of the rivets, although these and the wave propagation are distorted by the slight anisotropy of the composite material. The slower but higher-amplitude A0 mode is also evident as it emanates from the sensor location.

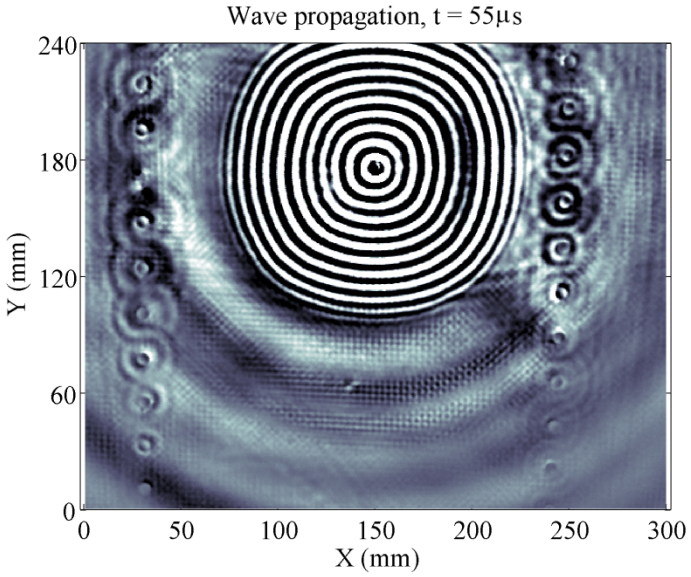

Figure 53 - Overview of wave propagation in riveted specimen

The scattering pattern of each rivet is estimated by establishing a spatial window of 40 mm x 40 mm centered on the rivet (which is the maximum size of the window such that the windows do not include the adjacent rivets). After implementing the polar wavenumber transformation on the data in that window, the scattering pattern at a fixed radius of 4 mm is extracted. The resulting representations are therefore rather near-field, at about twice the rivet radius, but this parameter could be adjusted as needed for a particular application. The data is

re-interpolated on that circumference and the energy of the wavefield is calculated over a specific time window, chosen to correspond with the expected arrival of the first scattered waves from the rivet. Note that the "scattering matrices" presented here are really a slice of a higher order tensor, because they are sampled at only one incident wave direction (because that transducer is fixed). A more complete characterization would use several transducers to highlight several incident wave angles, because UGW sensitivity to crack damage in particular is known to be highly directional-dependent. However, due to time constraints, this study focuses on only one incident angle, as well as the S0 mode (because it is the first to arrive, whereas the A0 arrival is superimposed with the S0 waves as well). However, it is feasible to perform mode selection in the frequency-wavenumber domain to use both modes, as has been demonstrated in previous work [95].

The scattering pattern of a few of the rivets is presented, where each has been rotated such that the incident wave approaches the rivet from the 0-degree direction. In Figure 54, the results from two undamaged and two damaged rivets, one from each rivet line, are depicted. The damage presence and direction is indicated by the solid red line. The notable feature of the undamaged plots is that most of the energy is scattered directly back to the source. While it is difficult to definitively say based on its orientation what the damage does to the scattering matrix, neither of the damaged cases follows the same tendency of scattering energy directly backward. Note that the absolute amplitudes of the plots are not expected to be equal because each rivet is a different distance and direction away from the source, so the incident waves have different energies.

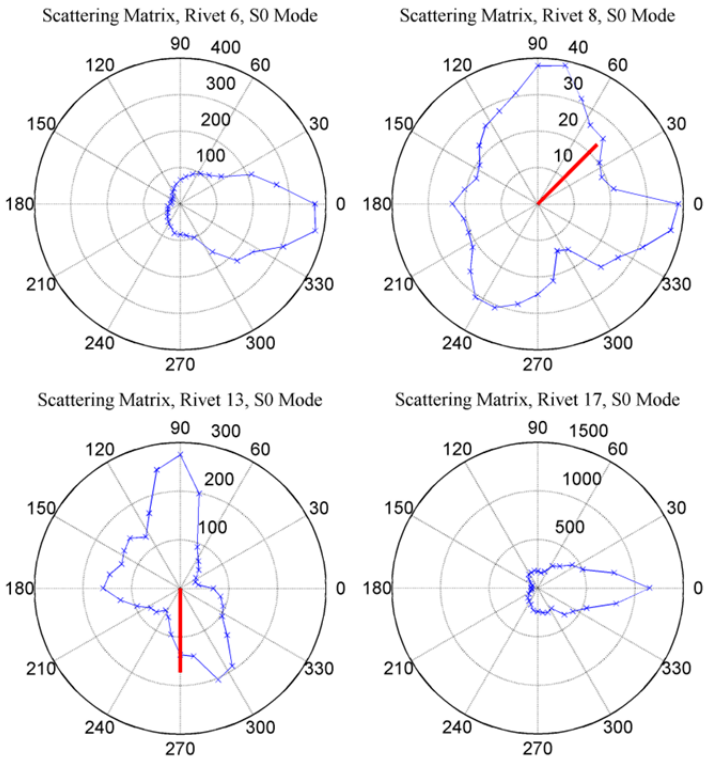

Figure 54 - Scattering matrix estimation for four example rivets

In order to interpret the changes in rivet scattering matrices from a damage detection perspective, two different metrics were implemented. The first metric, known as circularity, measures how similar the shape of the scattering profile is to a circle. While it is not expected that the scattering pattern should be circular *per se*, such a metric gives an effective measure of whether sudden deviations are present (which are observed at the angles corresponding to crack damage). Circularity, which is a common metric in the digital image processing field, is defined according to:

$$v^C = \frac{p^2}{A} \tag{4.3}$$

where p is the perimeter and A is the area of the shape [96]. For a perfect circle, the value of the circularity would be a lower-bound of $v^C = 4\pi$. Higher values indicate less circular objects. The result of computing this metric for each rivet's scattering profile is given in Figure 55. The magenta triangles represent the damaged rivets, the green stars represent healthy rivets, and the dotted black line denotes the partition between the two rivet lines. Rivets 13 and 18 are easily distinguished from the undamaged cases by this metric, but rivets 3 and 8 are less well-separated.

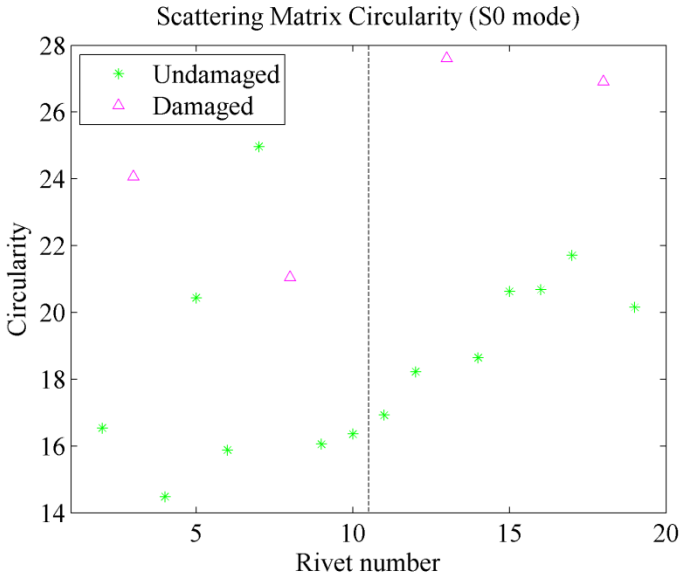

Figure 55 - Circularity of scattering profiles for each rivet

The second metric is better for a self-referencing damage detection paradigm. For this metric, the average correlation coefficient between a particular rivet and all of the other rivets is computed. Assuming that most of the rivets will be intact for situations where SHM is useful, damaged rivets should have scattering patterns that are less similar to the majority of healthy rivets. Figure 56 presents the comparison of average correlation coefficients.

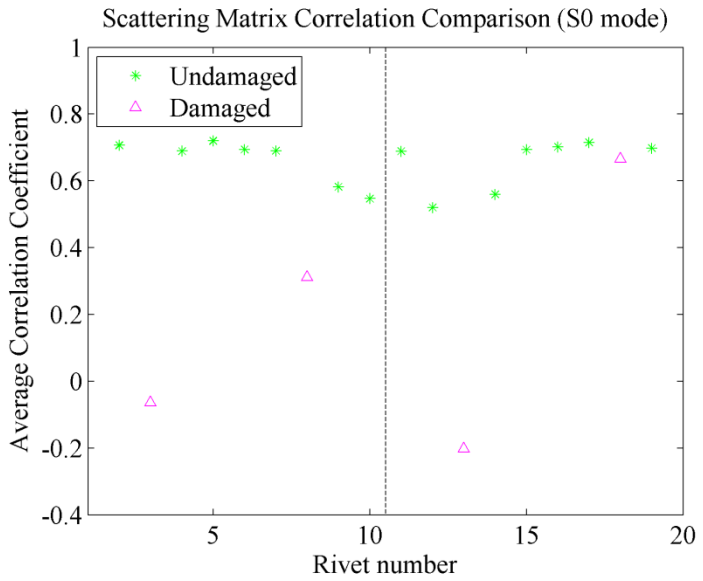

Figure 56 - Average correlation coefficient for each rivet with all other rivets

The results of the correlation coefficient comparison indicate that the damage is creating a significant change in the scattering pattern. The healthy rivets have scattering matrix correlation coefficients clustered in the range of 0.5-0.75. By contrast, rivets 3, 8, and 13 have values all less than 0.4, giving a clear indication that something is wrong with the rivet. Interestingly, however, rivet 18 does not exhibit the expected decrease, despite being clearly separable for the circularity metric.

While more study would be required to completely characterize the rivet scattering pattern and the changes corresponding to damage, the results have shown that full-field UGW systems have the potential to extract empirical scattering matrix estimates. Preliminary approaches to damage detection using these empirical estimates were also shown to produce reasonable results.

4.3. Low-order Feature Modeling

Knowledge of the geometry of a specimen is always critical to a successful UGW SHM system, though it is not always possible to generate detailed models of the wave propagation for every situation. Estimation of scattering properties is therefore not only helpful for modeling damage, but also for modeling geometric features. Where there are opportunities to add physical knowledge of the medium into post-processing, the system's performance tends to be enhanced. For example, Flynn *et al*. demonstrated that implementing simple models of such geometric features can provide significant improvement in detection performance, and validated that conclusion for a 2D structure with a bonded stringer [36]. Flynn *et al*. also showed dramatically decreasing returns from further increases in model complexity. These modeling approaches can be viewed as an attempt to estimate the amplitude of response that should be expected for each possible damage location and then utilize that estimate to enhance the predictions. Alternatively, the models could be interpreted as a weighting of the localization procedure according to the uncertainty in the localization estimate. Generally, the more features are in the wave path, the more uncertain that result becomes when received.

Applying a similar approach to the bolted frame and fuselage rib structures, simple models were developed to represent the major geometric features in each of the testbeds. Direct measurement is the most straightforward way of obtaining the scattering properties of different structural features, but is not always practical in every situation. Where it is not feasible to measure, it is possible to estimate those properties through supervised learning. In this case, the model parameters were estimated by a global minimization process where the average localization error was used as the fitness function for the search. Lacking direct measurement of feature parameters for the fuselage rib and frame structures, the set of damaged trials for each structure was split into "training" and "test" cases. For the frame, one in four trials were used in

training; for the fuselage rib, one trial each from the first, third, and fourth hole diameters was selected, also dedicating one-fourth of the cases for training. Only these training cases were used in generated the model parameters, and to calculate the results presented in Section 3.3.2, only test cases are used. Such a procedure is standard for machine-learning problems. Because the bolted frame has only two parameters, a direct search method was computationally feasible, whereas for the fuselage rib's four parameters, a genetic algorithm routine was utilized. Calibrating the values in a *post hoc* fashion is necessary because some kind of training data is always required in SHM applications where detailed models are not available or are highly uncertain. However, similar to the previous study, the localization results are only weakly sensitive to changes in the particular modeling parameters. In other words, merely including some recognition of the presence of geometric features is enough to realize most of the performance benefit. The impact on the localization performance of applying appropriate model parameters is presented.

4.3.1. Beam spread function

Another modeling consideration that ought not be ignored in this discussion involves how a wavefront loses energy as it travels, either from material attenuation or from geometric spreading of the wave energy [26]. Typically, the material attenuation in 2D metallic structures (usually modeled as a logarithmic decrease in wave amplitude with distance) is negligible compared to the geometric spreading [97, 98]. Here, the combination of factors that cause the amplitude of the wave to decrease with propagation distance will be referred to as the beam spread function. While different beam spread function models have been proposed by various authors, in a 2D medium the amplitude should vary according to:

$$B = \sqrt{\frac{1}{d_i(r)d_j(r)}} \tag{4.4}$$

where $d_i(r)$ is the distance from the actuating transducer to an arbitrary scattering point, and $d_j(r)$ is the distance from the scattering point to the receiving transducer [26, 36]. For a structure composed of 1D elements, there is no beam spread effect. Such models have been shown to be beneficial in situations where estimating the expected amplitude of the scattered waveform is important.

However, when the model given in Equation (4.4) was included in the localization procedure, the resulting prediction accuracy was not improved. To understand this result, one must recognize first that the MLE first-arrival localization result is calculated from the baseline-subtracted residual signal, which (contrary to ordinary signals) has a tendency to grow in time due to increasingly imprecise subtraction. The MLE filter result is also not directly an amplitude, but is instead a representation of the goodness-of-fit of a particular signal model. Therefore, in order to characterize the expected amplitude of the MLE filter result with distance, the true beam spread function should actually be composed of two terms – one decreasing due to the energy dissipation (both the smaller material attenuation and the geometric spreading) and one increasing due to increasingly imprecise baseline subtraction. Because the baseline-subtraction error is proportional to the underlying signal amplitude, however, this term would not be increasing for all time, but would eventually start to decay again. Baseline subtraction accuracy is also geometry dependent, as some boundary conditions are more changeable than others, and therefore a greater source of subtraction error. It is clear that the resulting composite function would be difficult to estimate *a priori*. A beam spread function of unity (which is in fact the default if no such modeling is considered) produced the best results and is used for all

localization in this dissertation. However, further work on how to most appropriately estimate the function of baseline-subtracted amplitude with distance is recommended.

A portion of this chapter has been published in Structural Health Monitoring: An International Journal, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2014. The title of this paper is "Monitoring of Bearing Failure in Composite Bolted Connections Using Ultrasonic Guided Waves: A Parametric Study". The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Key Engineering Materials, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2013. The title of this paper is "Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite Bolted Connections." The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Proc. of the International Workshop on Structural Health Monitoring, Colin Haynes and Michael Todd, 2013. The title of this paper is "Effect of Applied Load on Guided Wave Monitoring of a Composite Bolted Joint". The dissertation author was the primary investigator and author of this paper.

Another portion of Chapter 4 has been published in Proc. of the Asia-Pacific Workshop on Structural Health Monitoring, Colin Haynes, Takeaki Nadabe, Nobuo Takeda, and Michael Todd, 2012. The title of this paper is "Scattering Matrix Approach to Informing Damage Monitoring and Prognosis in Composite Bolted Connections". The dissertation author was the primary investigator and author of this paper.

Chapter 5
Supporting Topics for System Implementation

While the preceding sections address many central topics in the paradigm of UGW SHM, there are still many other topics that are important for transitioning SHM from research to practice. Two of those topics are discussed in this chapter: uncertainty quantification and Bayesian experimental design. One example is given for each field to demonstrate its application to the SHM methodology developed in this dissertation.

5.1. Uncertainty Quantification

5.1.1. Motivation for Uncertainty Quantification

Uncertainty in measurements has guided every aspect of the SHM strategy presented so far, but in some situations it is critical to have a well-quantified understanding. The goal of uncertainty quantification (UQ) is to characterize the effect of both parametric variability and non-parametric uncertainty on the output of an experiment or simulation [99]. Particularly for

higher-cost structures, efficiency and performance can be improved while reducing costs if estimates of uncertainty can be reliably obtained.

Many different approaches have been developed for UQ in the area of structural dynamics and SHM [100]. Perhaps the most basic and commonly used technique is the Monte Carlo simulation. Monte Carlo simulations obtain outputs (which may be from experiment, but are more commonly from simulation) based on randomly sampling each of the input parameters of interest. Monte Carlo simulations (and the related techniques that more efficiently sample the design space) therefore overcome a lack of knowledge of either the system or its inputs through repeated sampling. On the other hand, if the distributions of the inputs to an estimation process are known, complete information of the uncertainty in the outputs may be obtained by analytically propagating PDFs through to find output distributions. Often, however, the system is not well-defined enough for such analyses, or the analytical solution proves intractable. Sensitivity analyses, both global and local in scope, are often quite useful to quantify how results vary with changes in input. Analysis of variance studies are another major class of methods that seek to determine the main factors in driving the uncertainty of the output. While other methods have been utilized in connection with the work of this dissertation, for illustrative purposes one example of analysis of variance has been chosen for detailed description in the following section.

5.1.2. Application Example

In the course of empirically estimating scattering profiles for the bolt bearing damage in Section 4.1, it was observed that the resulting scattering matrices differed substantially from the expected behavior for a circular through-hole scatterer (for which analytical solutions are available). An undamaged circular through-hole would be expected to produce results symmetric about the direction of the incident waves, for example, but this was often not true of

the experimental result. Therefore, UQ techniques are applied to help identify the causes for some of those discrepancies and understand how the sources of uncertainty may be affecting the damage detection problem.

UQ can be furthered by model-driven approaches, which can generate surrogate data helpful for exploring the relationship between input and output variance. To be helpful in this case, however, a model would require very detailed modeling of the MFC sensors, their bond condition, and the bolt bearing damage. Such a model is not currently available—creating one would be time consuming and difficult to validate. Instead, measured responses are compared with measured parameters to evaluate their influence on the results.

The first step in successfully understanding the uncertainty is to assess the factors that may be influencing the result. Sensor bond condition, the sensor length, and the orientation and position of the sensors are all factors expected to affect the amplitude of the received signals. The state of the bolt fixture (particularly the applied load) was also observed to affect the amplitude. For the arrival time, which is directly related to the group velocity, the temperature and the position of the sensors (both the radial and rotational coordinates) are expected to produce changes. However, not all of the parameters that affect the results can be included in the analysis because some cannot be measured – most notably, no information is available on the bond condition.

For this study, particular emphasis was placed on determining the impact of the location, orientation, and dimensions of the MFC sensors on the UGW results. In order to measure these parameters as accurately as possible, a high-resolution scan of the specimen was made. Each of the geometric parameters was measured from that image, along with an estimated uncertainty for each. For the coordinate locations, the center of the leading edge of each MFC was established and used for the measurements (consistent with how the distances

were measured for the original installation). The resulting estimates are presented in Table 9. Column 2 gives the nominal angular positions of each sensor, whereas the values in column 3 represent the deviation of the measured values from those in column 2. Orientation error represents the angular deviation of the sensor direction, as estimated from the directions of the MFC fibers, from the radius it should be aligned with. Finally, recognizing that the sensors are not perfectly rectangular, an average length and width of the active area of the sensors are given. The bottom row of the table presents the mean and standard deviation of the values for all eight sensors. Where appropriate, the mean of the absolute values are used so that the mean of the angle and orientation errors represents how much those measures tend to deviate from the presumed values.

Table 9 - Geometrical properties estimated from high-resolution scan

Sensor	Nominal Angle (deg)	Angle Error (deg)	Radial Distance (mm)	Orientation Error (deg)	Length (mm)	Width (mm)
1	0	-1.4	29.57	0.88	6.24	13.65
2	45	-2.1	29.99	1.45	6.35	13.71
3	90	-0.8	30.12	0.08	6.32	13.72
4	135	0.2	29.46	-0.83	6.19	13.62
5	180	0.0	30.04	-0.88	6.24	13.63
6	225	-0.4	30.27	-0.38	6.40	13.63
7	270	-0.2	30.14	-0.65	6.26	13.70
8	315	-0.7	30.30	-0.72	6.22	13.64
Avg	N/A	0.71±0.66	29.99±0.29	0.73±0.37	6.28±0.07	13.66±0.04

In order to analyze the variance in the result due to the factors just quantified, features must likewise be selected from the received time series for comparison. The three features considered are: the amplitude of the target hole reflection, the arrival time of the same, and the NRE metric computed on the relevant time window. These three metrics together should capture the relevant features of the response variance for the following analysis. To give a qualitative impression of the variability of the time series on which these metrics are computed, Figure 57 shows the average of all of the time series for MFC sensors separated by 135°. The red dotted lines show the limits of one standard deviation, and the mean and standard deviations of the arrival estimate are shown in green.

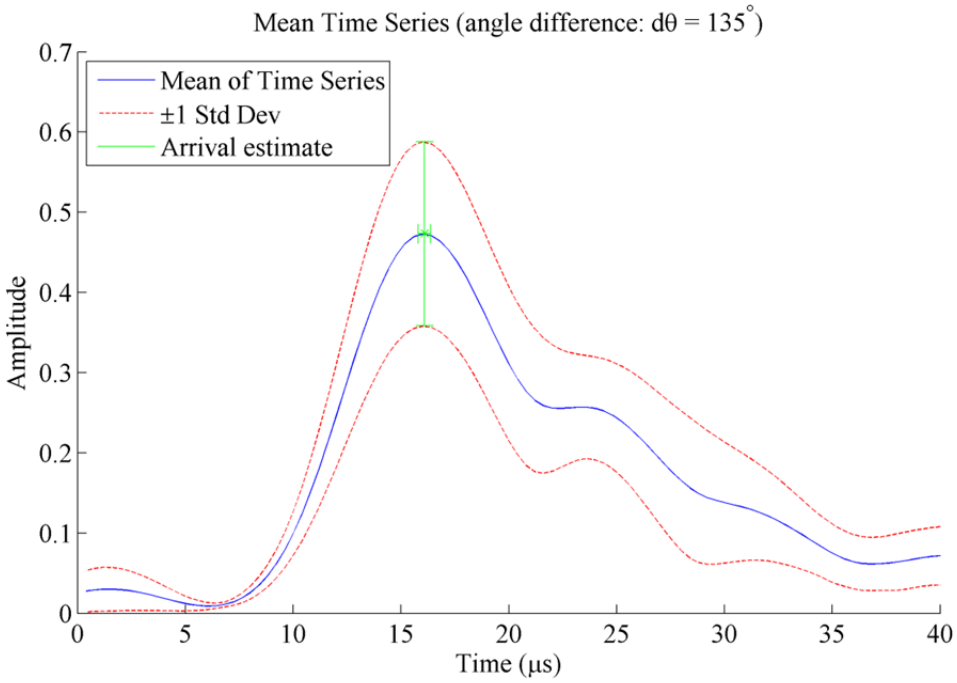

Figure 57 - Average time series for all MFC sensors separated by 135°

For the analysis of variance study, the R^2 statistic introduced to quantify the effect that each parameter has on the overall variability in the system. For a particular factor, the R^2

statistic represents the ratio of the sum of the conditional variance with the factor held fixed at each of its levels to the overall variance of the feature. Essentially, the result is related to the significance of the considered variable in determining the overall variance of the response, with higher values indicating a more significant parameter. The mathematical representation is given in Equation (5.1):

$$R^2 = 1 - \frac{\sum_{a=1}^{N_{levels}} \sum_{b=1}^{N_{trials}/N_{levels}} \left(v_b^{(a)} - \overline{v^{(a)}} \right)^2}{\sum_{b=1}^{N_{trials}} \left(v_b - \bar{v} \right)^2} \tag{5.1}$$

where v represents the feature values, the superscript denotes a particular level of the variable value, and an over-bar denotes the mean value. N_{levels} is the number of unique values the variable takes and N_{trials} is the number of undamaged tests.

The factors used to calculate the R^2 statistics are each derived from the measurements presented in Table 9. "R" is the radial distance of each sensor from the center of the hole. "Load" is the tensile load applied to the specimen during each test. Orientation error is given by "ϕ_{12}", which calculated as the sum of the absolute values of the orientation errors of the receiving and sensing MFC. Sensor length, denoted "L_{12}", is defined similarly. The nominal difference in angle between the pair is represented by "$\Delta\theta$", and "ε_θ" is the error between the measured to the nominal difference in angle. Finally, "Random" is a reference vector included to provide a verification check on the computation. For this category, the R^2 statistic is computed from a set of values generated from a normal distribution – because a random vector is not related to the output result, it is expected to yield a result near zero. Each set of values is broken down into five "levels" (by percentile) to facilitate the comparison by ensuring that each factor has the same number of levels. The results are presented in Figure 58.

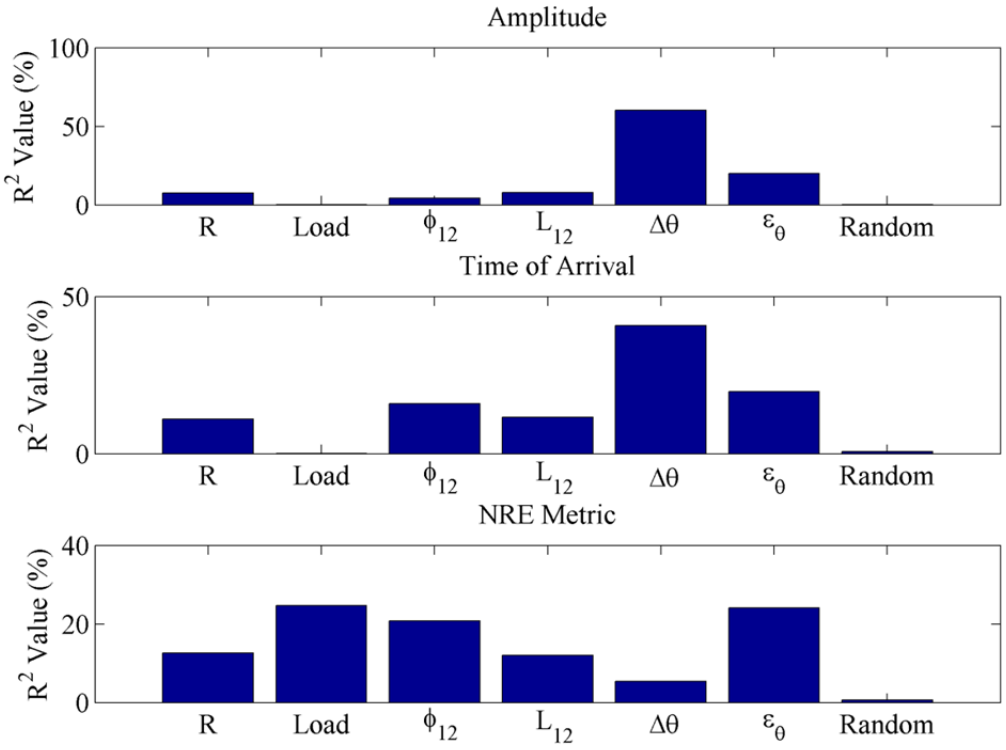

Figure 58 - R^2 values for signal amplitude, time of arrival, and NRE metric

First, it is clear that the random vector does not contribute significantly to the overall variance, so the rest of the factors may be examined. For the signal amplitude, by far the most significant factor is the nominal angle difference. Considering that the scattering pattern from the hole is primarily driven by this parameter, it is not unreasonable to think that it may drive the variance as well, so this is not a surprising result. Angle difference is also the dominant factor in the variance of the estimated time of arrival, most likely for the same reason. For both of amplitude and time of arrival, the error in the angle difference is the next most significant factor. Interestingly, the applied load has almost no influence on the variance in either of the first two cases. By contrast, load is the most significant effect when considering the NRE metric, which possibly explains why detection performance was found to be substantially affected by the presence of load. The influence of angle difference, on the other hand,

dramatically decreases for the NRE metric, suggesting that baseline subtraction mitigates the sensitivity to this parameter.

The results, unfortunately, are not conclusive in the sense of suggesting one particular parameter that will do the most to reduce the overall variance if it can be better controlled. All of the parameters are significant in affecting the results, so it must be concluded that care should be taken to place sensors as exactly as possible if the scattering matrices measured by fixed transducer locations are to be matched to their theoretical counterparts. As described in Chapter 4, there are other sensing methodologies that may be used to estimate scattering matrices, but for situations utilizing fixed sensor locations, this study gives insight into the expected uncertainty.

5.2. Bayesian Experimental Design & Optimal Sensor Placement

As expressed in Section 1.1, the goal of a SHM system is to decrease the cost of sustaining safe operations by facilitating condition-based maintenance. The primary drivers of implementing SHM, therefore, can be expressed in economic terms. Improved reliability and life-safety, reduced downtime due to unexpected maintenance demands, the ability to adjust operation based on the state of the structure, lower costs through maintenance and performance optimization – all of these may be viewed in terms of a cost-savings to the owner. Likewise, each decision made involves potential costs. Just in determining how to instrument the structure, many choices are involved: for instance, the price of the sensors and data acquisition equipment, added weight, power consumption, and false-alarm rates. Identifying and quantifying all of the costs associated with SHM allows for the system to be designed intelligently through a procedure known as Bayesian experimental design (BED) [101, 102].

The BED process will be illustrated in this section by applying it to another key supporting topic for SHM system deployment: optimal sensor placement.

The criterion for designing the best SHM system from the standpoint of Bayes cost may be expressed as:

$$\min_{e}\left(\sum_{a,b}\mathrm{C}^{\mathrm{d}}\left(d_a,\theta_b\right)\mathrm{P}\left(d_a\mid\theta_b,e\right)\mathrm{P}\left(\theta_b\right)+\mathrm{C}^{\mathrm{e}}\left(e\right)\right) \tag{5.2}$$

The symbol *e* refers to the SHM system design ("experiment" in BED terminology), involving whatever parameters are viewed as pertinent to the optimization. This may take into account a wide variety of design parameters, including the type, number, and location of sensors, the data acquisition used, and elements of the feature extraction process (such as the decision threshold). All possible decisions d_a taken based on the damage prognosis and all possible true states of the structure, θ_b are considered in the summation. The likelihood function $\mathrm{P}\left(d_a\mid\theta_b,e\right)$ is the term that is most often the focus of research in developing a SHM system. This function represents the probability that the system will make decision d_a given the actual state of the structure θ_b and the SHM system used to assess that true state (*e*). In other words, this function gives a measure of the prediction accuracy of the system, whether from probabilities of detection estimated from testing (data-driven) or from validated models (model-driven). The term $\mathrm{C}^{\mathrm{d}}\left(d_a,\theta_b\right)$, then, represents the cost of making a particular decision given the true state of the structure. Finally, the term $\mathrm{P}\left(\theta_b\right)$ is the prior probability of each state occurring (established through the system's operational evaluation), and $\mathrm{C}^{\mathrm{e}}\left(e\right)$ is the cost of implementing the SHM system.

The first step in applying the BED approach to SHM is to select the target states of the structure between which the system should discriminate [103]. Next, the possible actions that may be taken in response to information on the structure's state must be determined. Once these parameters have been established, the BED procedure can be used to optimize the SHM system in a number of different ways. To demonstrate the BED methodology, the bolted frame structure is used. As an example of the methodology, various costs are assumed and an optimal sensor layout will be established using a data-driven approach.

In this case, certain assumptions guide the BED framework. When the SHM system detects damage, the action will be to immediately inspect and repair, whereas no action will be taken for a negative reading. The likelihood function $\mathrm{P}\left(d_a \mid \theta_b, e\right)$ is estimated from ROC analysis of the test data described in Sections 0 and 3.2.3 – that is, the probabilities for a given system and damage state are found from the percentage of the test cases in that state and using that system arrangement that are classified in each category. Finally, the prior probabilities that each damage state will occur are assumed (both bolt and magnet damage are assigned equal prior probability). Because the bolted frame structure has no prior history of failure, loading, etc., the prior probabilities are taken to be equal for bolt and magnet damage modes.

With costs thus assigned, optimal system design can now be considered. The first example answers the question of how many sensors should be used in the SHM system. It is expected that as the cost of missed detection (false negative) increases, the BED framework would favor using more sensors to achieve a better coverage. Lower costs of missed detection would allow the cost of adding more sensors to drive the decision. This trend is indeed observed, as shown in Figure 59. As the missed detection to false alarm cost ratio increases, the optimal number of sensors (denoted by the vertical red line) continues to increase.

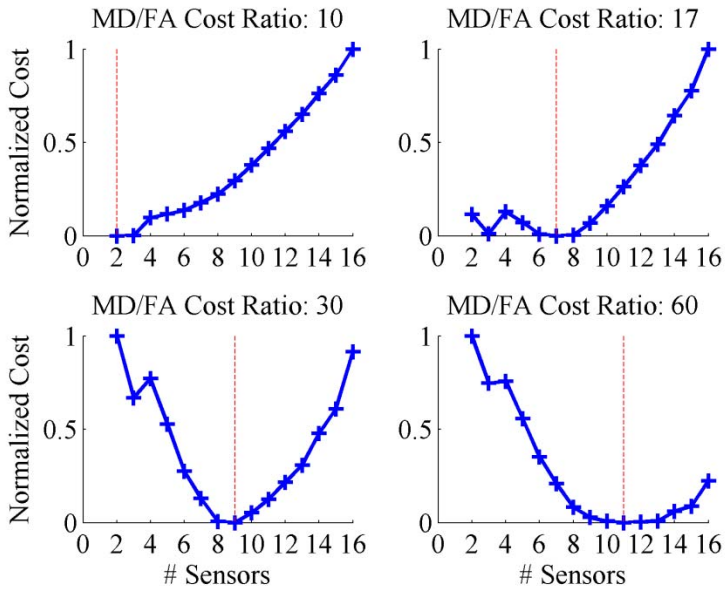

Figure 59 - Optimal number of sensors, varying missed detection to false alarm cost ratios

By contrast, if for weight, space, or other constraints, only a certain number of sensors may be used, the BED procedure can optimally select their locations based on the relative costs of false alarm and missed detection for each type of damage. Figure 60 shows sensor layouts where the bolted frame is constrained to have only four sensors and the missed detection cost is varied. Four sensors is an extremely sparse system for the size of the structure being monitored, which serves to highlight the trade-offs as the missed detection cost changes. When the missed detection cost is relatively low, as in the top-left example, the system chooses the most conservative arrangement possible. The sensors tend to cluster in smaller regions, achieving very accurate local prediction by sacrificing coverage area in order to minimize false alarms. By contrast, when the missed detection cost is relatively high, as in the bottom-right example, the BED framework chooses network that are more spread out, maximizing coverage and minimizing the areas where damage might escape detection altogether.

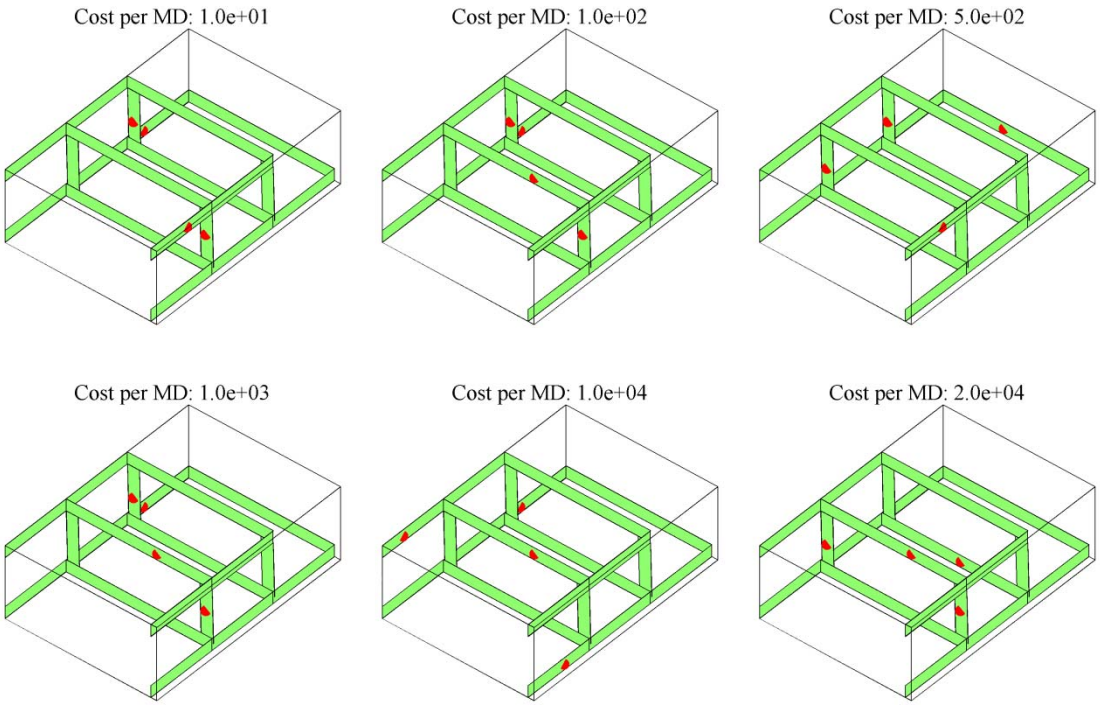

Figure 60 - Four-sensor layouts for different missed detection costs. Red dots indicate sensor locations.

Examining the ROCs directly can provide more insight into the trade-off process. Continuing to consider the case where the sensor system is constrained to exactly four sensors, Figure 61 presents results for magnet damage only. Each line on the plot represents the ROC of the sensor arrangement that was found to be optimal at one of the missed detection cost values (with the cost of false alarm held constant). As the missed detection cost increases, the system biases more towards solutions that have a better detection rate as well as the accompanying higher false alarm rate. Graphically, this trend appears as solutions being driven up and to the right. Furthermore, it is clear that the markers are always closest to the top left in their particular region—that is, the BED algorithm selects sensor arrangements that are *locally* optimal (in terms of the ROC) for the specified missed detection cost ratio. Similar results may be observed for the bolt damage case; however this case is omitted due to the near-perfect detectability of that damage type which makes comparison more difficult.

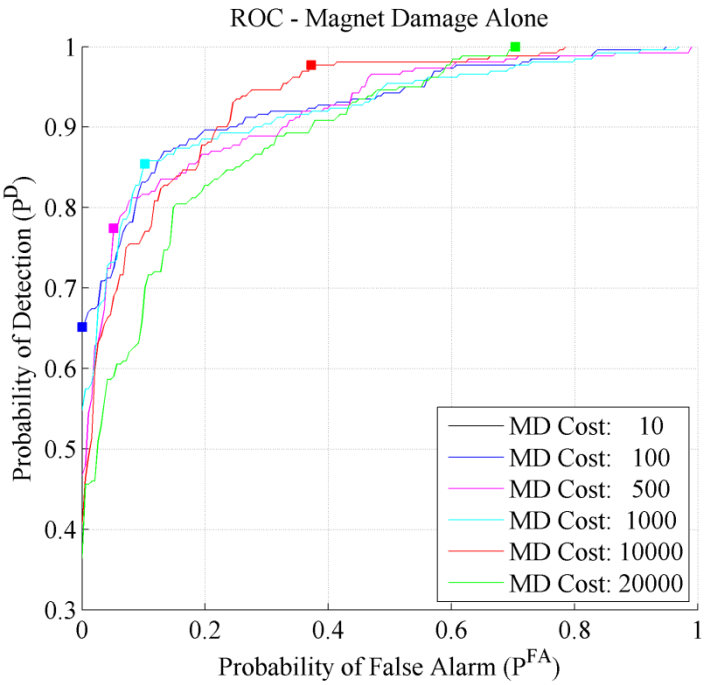

Figure 61 - ROC of optimal four-sensor configurations for magnet damage alone. The square markers indicate the point on those ROC curves that provided the optimal performance for a particular cost.

Looking at the two damage modes simultaneously presents a different challenge, however. The chosen points on each of the ROC plots independently are no longer necessarily the local optima. Because detecting the each form of damage requires a different strategy, the optimal solution now consists of a compromise. Figure 62 demonstrates this principle through the ROCs for each type of damage at several different missed detection cost ratios, where the layout was jointly optimized.

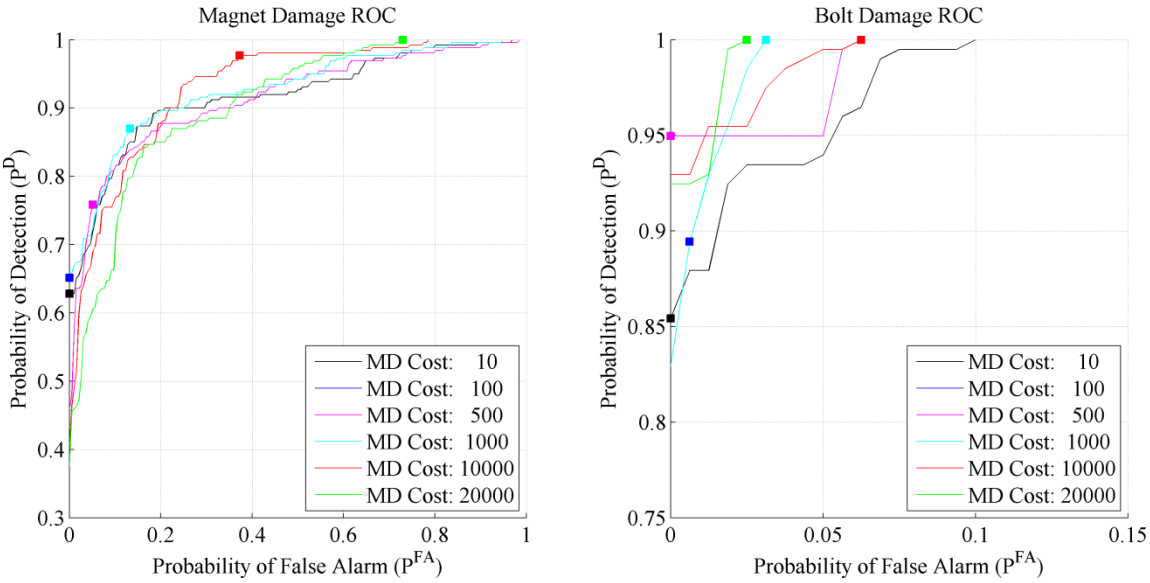

Figure 62 - ROCs for magnet and bolt damage, considered jointly

BED provides a useful framework designing SHM systems that balance the various trade-offs that arise. Using the example of the bolted frame structure, the principles of system selection through the BED process have been demonstrated. If the costs of various decisions and alternatives can be established, BED can be used to suggest optimal sensor layouts, actuation procedures, threshold values, and many other decisions. Here, the BED process was first used to determine the optimal number of sensors based on the costs of the sensors and the costs of incorrectly assessing the damage. As expected, more sensors are used when the decision costs increase relative to the sensor cost. Next, it was assumed that the sensor count is fixed at four sensors, and the optimal arrangements of those sensors were plotted. Furthermore, ROC analysis may be used to show that the system selects the optimal points on the ROC curve based on the cost ratios and the types of damage considered.

A portion of this chapter has been published in Proc. of SPIE, Colin Haynes and Michael Todd, 2013. The title of this paper is “Uncertainty quantification of a guided wave

structural health monitoring system for composite bolted joints". The dissertation author was the primary investigator and author of this paper.

Another portion of this chapter has been published in Proc. of the European Workshop on Structural Health Monitoring, Colin Haynes and Michael Todd, 2012. The title of this paper is "Bayesian Experimental Design for Damage Detection in a Bolted Frame". The dissertation author was the primary investigator and author of this paper.

Chapter 6
Conclusions and Future Work

6.1. Conclusions

Many UGW SHM studies to date have focused on testing various algorithms in simple, plate-like structures with few geometric scatterers. When applied to complex materials and geometries, the increased uncertainty often causes these algorithms to break down. Here, a statistical framework was proposed at the outset, accounting for the expected uncertainty by modeling the signals as random processes. Residual signals from guided wave processes can be broken into pre-arrival and post-arrival regimes. The principles of detection theory were then applied to derive an optimal detector for discriminating damaged and undamaged cases. Maximum likelihood estimation was utilized to estimate the first-arrival point, which was integrated into a comprehensive localization framework. Novel sensor fusion techniques were developed to improve localization accuracy by weighting sensor pars by their associated uncertainty.

In order to validate the presuppositions of the statistical framework, several experiments were carried out, including different materials, geometries, and damage types. A bolted frame structure made of steel was tested with reversible damage—both magnets and bolt loosening. A section of a metallic fuselage rib was tested destructively through machined through-holes and

cracks. For the combined data from these two tests, the statistical approaches to detection, localization, and sensor fusion were validated. Through extensive comparative studies, the approach was shown to outperform a selection of techniques from the literature. Data-driven comparisons like these lend strong support to the proposed SHM strategy. Another focus of the dissertation was to develop techniques for empirical estimation of scattering properties, as detailed in two different examples. A scattering matrix characterization of bolt bearing damage was constructed, investigating the effect of applied load, input and output wave angles, and actuation frequency on the results. Next, damaged and undamaged rivets were characterized using full-field UGW data.

Several supporting topics are also considered which are important for bridging the gap to practice. In order to translate SHM assessments into decisions, uncertainty quantification is an essential step. One example is presented for analysis of variance on the scattering matrix estimations. Finally, Bayesian experimental design provides a cohesive framework for decision making and balancing trade-offs in terms of quantified costs, and as such forms a good conclusion for tying together all of elements of a comprehensive SHM strategy.

In all, this dissertation represents a significant step forward in advancing SHM toward structures of realistic complexity. Through detection theory arguments and empirical studies, the statistical approach to managing the uncertainty in such environments has been shown to be effective.

6.2. Future Work

There are a number of topics suggested by this dissertation that may warrant further study. Currently, the localization formulations only consider point-like scatterers. Whether the MLE filtering and other localization processes may be extended to larger area damage is an

important question. The framework as presented currently are searching only for one damage site – for some applications, multi-site damage may be critical to investigate. Some kind of uncertainty analysis relating to the output of the MLE filtering process would also be important, especially in lending support to the proposed sensor fusion scheme. Relating to the scattering matrix analysis, digital image processing techniques may be of great help in processing full-field data, and should be examined for that purpose. Finally, more detailed investigation of the appropriate beam spread function for residual signals may lead to a further improvement in localization accuracy.

References

1. C.R. Farrar, K. Worden, An introduction to structural health monitoring, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences. 365, 1851; 2007, 303–315.

2. A. Rytter, Vibration based inspection of civil engineering structures, 1993.

3. K. Worden, C.R. Farrar, G. Manson, G. Park, The fundamental axioms of structural health monitoring, Proc. R. Soc. A. 463, 2082; 2007, 1639–1664.

4. C.R. Farrar, N.A.J. Lieven, Damage prognosis: the future of structural health monitoring, Phil. Trans. R. Soc. A. 365, 1851; 2007, 623–632.

5. P.C. Chang, A. Flatau, S.C. Liu, Review Paper: Health Monitoring of Civil Infrastructure, Structural Health Monitoring. 2, 3; 2003, 257–267.

6. J. Ou, H. Li, Structural Health Monitoring in mainland China: Review and Future Trends, Structural Health Monitoring. 9, 3; 2010, 219–231.

7. J. Davies, P. Cawley, The application of synthetic focusing for imaging crack-like defects in pipelines using guided waves, IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control. 56, 4; 2009, 759–771.

8. J.K. Van Velsor, H. Gao, J.L. Rose, Guided-wave tomographic imaging of defects in pipe using a probabilistic reconstruction algorithm, Insight: Non-Destructive Testing & Condition Monitoring. 49, 9; 2007, 532–537.

9. J.-R. Lee, S.Y. Chong, H. Jeong, C.-W. Kong, A time-of-flight mapping method for laser ultrasound guided in a pipe and its application to wall thinning visualization, NDT & E International. 44, 8; 2011, 680–691.

10. A. Zagrai, D. Doyle, B. Arritt, Embedded nonlinear ultrasonics for structural health monitoring of satellite joints, Proc. SPIE, 6935, 2008, San Diego, CA.

11. Z. Mao, M. Todd, A structural transmissibility measurements-based approach for system damage detection, Proc. SPIE, 7650, 2010, San Diego, CA.

12. C.C. Chia, J.-R. Lee, C.-Y. Park, H.-M. Jeong, Laser ultrasonic anomalous wave propagation imaging method with adjacent wave subtraction: Application to actual damages in composite wing, Optics & Laser Technology. 44, 2; 2012, 428–440.

13. K.I. Salas, C.E.S. Cesnik, Guided wave structural health monitoring using CLoVER transducers in composite materials, Smart Mater. Struct. 19, 1; 2010, 015014.

14. N. Takeda, Y. Okabe, J. Kuwahara, S. Kojima, T. Ogisu, Development of smart composite structures with small-diameter fiber Bragg grating sensors for damage detection: Quantitative evaluation of delamination length in CFRP laminates using Lamb wave sensing, Composites Science and Technology. 65, 15-16; 2005, 2575–2587.

15. S. Salamone, I. Bartoli, F.L.D. Scalea, S. Coccia, Guided-wave Health Monitoring of Aircraft Composite Panels under Changing Temperature, Journal of Intelligent Material Systems and Structures. 20, 9; 2009, 1079–1090.

16. R. Do, C. Haynes, M. Todd, W. Gregory, C. Key, Efficient Detection Methods on a Composite Plate with Interior Embedded Fiber Optic Sensors via Impact Test, Proc. International Workshop on Structural Health Monitoring, 2013, Stanford University.

17. S.W. Doebling, C.R. Farrar, M.B. Prime, D.W. Shevitz, Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review, 1996.

18. N. Takeda, S. Minakuchi, T. Nadabe, Distributed Strain Monitoring for Damage Evolution in CFRP Bolted Structures with Embedded Optical Fibers, Key Engineering Materials. 558, 2013, 252–259.

19. M. Majumder, T.K. Gangopadhyay, A.K. Chakraborty, K. Dasgupta, D.K. Bhattacharya, Fibre Bragg gratings in structural health monitoring—Present status and applications, Sensors and Actuators A: Physical. 147, 1; 2008, 150–164.

20. J.P. Lynch, K.J. Loh, A summary review of wireless sensors and sensor networks for structural health monitoring, Shock and Vibration Digest. 38, 2; 2006, 91–130.

21. N. Elvin, A. Erturk, Advances in energy harvesting methods, Springer, 2013.

22. A. Dewan, S.U. Ay, M.N. Karim, H. Beyenal, Alternative power sources for remote sensors: A review, Journal of Power Sources. 245, 2014, 129–143.

23. P.J. Shull, editor, Nondestructive Evaluation: Theory, Techniques, and Applications, CRC Press, Boca Raton, 2002.

24. H. Lamb, On waves in an elastic plate, Proceedings of the Royal Society of London. Series A. 93, 648; 1917, 114–128.

25. D.C. Worlton, Experimental Confirmation of Lamb Waves at Megacycle Frequencies, Journal of Applied Physics. 32, 6; 1961, 967–971.

26. P.D. Wilcox, M.J.S. Lowe, P. Cawley, Mode and Transducer Selection for Long Range Lamb Wave Inspection, Journal of Intelligent Material Systems and Structures. 12, 8; 2001, 553–565.

27. V. Giurgiutiu, J. Bao, W. Zhao, Piezoelectric wafer active sensor embedded ultrasonics in beams and plates, Experimental Mechanics. 43, 4; 2003, 428–449.

28. T. Kundu, editor, Ultrasonic Nondestructive Evaluation: Engineering and Biological Material Characterization, CRC Press, Boca Raton, 2004.

29. D.N. Alleyne, P. Cawley, The interaction of Lamb waves with defects, IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control. 39, 3; 1992, 381–397.

30. F.B. Cegla, A. Rohde, M. Veidt, Analytical prediction and experimental measurement for mode conversion and scattering of plate waves at non-symmetric circular blind holes in isotropic plates, Wave Motion. 45, 3; 2008, 162–177.

31. J.C.P. McKeon, M.K. Hinders, Lamb Wave Scattering from a Through Hole, Journal of Sound and Vibration. 224, 5; 1999, 843–862.

32. P. Fromme, M.B. Sayir, Measurement of the scattering of a Lamb wave by a through hole in a plate, The Journal of the Acoustical Society of America. 111, 3; 2002, 1165–1170.

33. O. Diligent, T. Grahn, A. Bostrom, P. Cawley, M.J.S. Lowe, The low-frequency reflection and scattering of the S[sub 0] Lamb mode from a circular through-thickness hole in a plate: Finite Element, analytical and experimental studies, The Journal of the Acoustical Society of America. 112, 6; 2002, 2589–2601.

34. G.R. Liu, A combined finite element/strip element method for analyzing elastic wave scattering by cracks and inclusions in laminates, Computational Mechanics. 28, 1; 2002, 76–82.

35. P.D. Wilcox, A. Velichko, Efficient frequency-domain finite element modeling of two-dimensional elastodynamic scattering, Journal of the Acoustical Society of America. 127, 1; 2010, 155–165.

36. E.B. Flynn, M.D. Todd, A.J. Croxford, B.W. Drinkwater, P.D. Wilcox, Enhanced detection through low-order stochastic modeling for guided-wave structural health monitoring, Structural Health Monitoring. 11, 2; 2012, 149–160.

37. J.S. Hall, J.E. Michaels, A model-based approach to dispersion and parameter estimation for ultrasonic guided waves, The Journal of the Acoustical Society of America. 127, 2; 2010, 920–930.

38. G.J. Jarmer, E.B. Flynn, M.D. Todd, Dispersion curve estimation via phased array beamforming methods, Journal of Intelligent Material Systems and Structures. 25, 9; 2014, 563-574.

39. M.D. Todd, E.B. Flynn, P.D. Wilcox, B.W. Drinkwater, A.J. Croxford, S. Kessler, Ultrasonic wave-based defect localization using probabilistic modeling, Proc. Review of Progress in Quantitative Nondestructive Evaluation, 31, 2011, 639–646, Burlington, VT.

40. G. Konstantinidis, P.D. Wilcox, B.W. Drinkwater, An Investigation Into the Temperature Stability of a Guided Wave Structural Health Monitoring System Using Permanently Attached Sensors, IEEE Sensors Journal. 7, 5; 2007, 905–912.

41. M.J. Schulz, M.J. Sundaresan, J. Mcmichael, D. Clayton, R. Sadler, B. Nagel, Piezoelectric Materials at Elevated Temperature, Journal of Intelligent Material Systems and Structures. 14, 11; 2003, 693–705.

42. A.J. Croxford, J. Moll, P.D. Wilcox, J.E. Michaels, Efficient temperature compensation strategies for guided wave structural health monitoring, Ultrasonics. 50, 4-5; 2010, 517–528.

43. H.L. Van Trees, Optimum Array Processing - Part IV, Detection, Estimation, and Modulation Theory, John Wiley & Sons, New York, 2002.

44. A.J. Croxford, P.D. Wilcox, B.W. Drinkwater, G. Konstantinidis, Strategies for guided-wave structural health monitoring, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science. 463, 2087; 2007, 2961–2981.

45. S. Kay, Fundamentals of Statistical Signal Processing, Vol. 2, Prentice Hall PTR, Upper Saddle River, New Jersey, 1998.

46. T. Fawcett, An introduction to ROC analysis, Pattern Recognition Letters. 27, 8; 2006, 861–874.

47. J.S. Hall, J.E. Michaels, Multi-path guided wave imaging for in situ monitoring of complex structures, AIP Conference Proceedings, 1511, 2013, 707–713.

48. T.R. Hay, R.L. Royer, H. Gao, X. Zhao, J.L. Rose, A comparison of embedded sensor Lamb wave ultrasonic tomography approaches for material loss detection, Smart Mater. Struct. 15, 4; 2006, 946.

49. C.T. Ng, M. Veidt, A Lamb-wave-based technique for damage detection in composite laminates, Smart Materials and Structures. 18, 7; 2009, 074006.

50. C.H. Wang, J.T. Rose, F.-K. Chang, A synthetic time-reversal imaging method for structural health monitoring, Smart Mater. Struct. 13, 2; 2004, 415.

51. J.E. Michaels, A.J. Croxford, P.D. Wilcox, Imaging algorithms for locating damage via in situ ultrasonic sensors, Proc. IEEE Sensors Applications Symposium, 2008, 63–67.

52. J. Moll, R.T. Schulte, B. Hartmann, C.-P. Fritzen, O. Nelles, Multi-site damage localization in anisotropic plate-like structures using an active guided wave structural health monitoring system, Smart Mater. Struct. 19, 4; 2010, 045022.

53. T. Clarke, P. Cawley, Enhancing the defect localization capability of a guided wave SHM system applied to a complex structure, Structural Health Monitoring. 10, 3; 2011, 247–259.

54. C. Zhou, Z. Su, L. Cheng, Probability-based diagnostic imaging using hybrid features extracted from ultrasonic Lamb wave signals, Smart Mater. Struct. 20, 12; 2011, 125005.

55. E.B. Flynn, M.D. Todd, P.D. Wilcox, B.W. Drinkwater, A.J. Croxford, Maximum-likelihood estimation of damage location in guided-wave structural health monitoring, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science. 467, 2133; 2011, 2575–2596.

56. David L. Hall, Mathematical Techniques in Multisensor Data Fusion, Artech House, Boston, 1992.

57. Q. Liu, H. Wang, A case study on multisensor data fusion for imbalance diagnosis of rotating machinery, Artif. Intell. Eng. Des. Anal. Manuf. 15, 3; 2001, 203–210.

58. Y. Lu, J.E. Michaels, Feature Extraction and Sensor Fusion for Ultrasonic Structural Health Monitoring Under Changing Environmental Conditions, IEEE Sensors Journal. 9, 11; 2009, 1462–1471.

59. Z. Su, X. Wang, L. Cheng, L. Yu, Z. Chen, On Selection of Data Fusion Schemes for Structural Damage Evaluation, Structural Health Monitoring. 8, 3; 2009, 223 –241.

60. Z. Su, X. Wang, Z. Chen, L. Ye, A hierarchical data fusion scheme for identifying multi-damage in composite structures with a built-in sensor network, Smart Materials and Structures. 16, 6; 2007, 2067–2079.

61. E.D. Niri, A. Farhidzadeh, S. Salamone, Adaptive multisensor data fusion for acoustic emission source localization in noisy environment, Structural Health Monitoring. 12, 1; 2013, 59–77.

62. C. Haynes, M.D. Todd, E. Flynn, A. Croxford, Statistically-based damage detection in geometrically-complex structures using ultrasonic interrogation, Structural Health Monitoring. 12, 2; 2013, 141–152.

63. C. Haynes, M. Yeager, M. Todd, J.-R. Lee, Monitoring bolt torque levels through signal processing of full-field ultrasonic data, Proc. SPIE, 2014, San Diego, CA.

64. E.B. Flynn, M.D. Todd, Optimal Placement of Piezoelectric Actuators and Sensors for Detecting Damage in Plate Structures, Journal of Intelligent Material Systems and Structures. 21, 3; 2010, 265–274.

65. J.E. Michaels, A.C. Cobb, T.E. Michaels, A comparison of feature-based classifiers for ultrasonic structural health monitoring, Proc. SPIE, 5394, 2004, 363–374, Bellingham, WA.

66. J.-B. Ihn, F.-K. Chang, Pitch-catch Active Sensing Methods in Structural Health Monitoring for Aircraft Structures, Structural Health Monitoring. 7, 1; 2008, 5 –19.

67. J.E. Michaels, T.E. Michaels, Guided wave signal processing and image fusion for in situ damage localization in plates, Wave Motion. 44, 6; 2007, 482–492.

68. H. Gao, Y. Shi, J.L. Rose, Guided Wave Tomography on an Aircraft Wing with Leave in Place Sensors, AIP Conference Proceedings, 760, 2005, 1788–1794.

69. J.E. Michaels, T.E. Michaels, Damage Localization in Inhomogeneous Plates Using a Sparse Array of Ultrasonic Transducers, AIP Conference Proceedings, 894, 2007, 846–853.

70. J.S. Hall, J.E. Michaels, Minimum variance ultrasonic imaging applied to an in situ sparse guided wave array, IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control. 57, 10; 2010, 2311 –2323.

71. R.J. Schalkoff, Pattern Recognition: Statistical, Structural, and Neural Approaches, John Wiley & Sons, Inc., New York, 1992.

72. C. Zhou, Z. Su, L. Cheng, Quantitative evaluation of orientation-specific damage using elastic waves and probability-based diagnostic imaging, Mechanical Systems and Signal Processing. 25, 6; 2011, 2135–2156.

73. N. Takeda, Characterization of microscopic damage in composite laminates and real-time monitoring by embedded optical fiber sensors, International Journal of Fatigue. 24, 2-4; 2002, 281–289.

74. B.A. Auld, General electromechanical reciprocity relations applied to the calculation of elastic wave scattering coefficients, Wave Motion. 1, 1; 1979, 3–10.

75. G.S. Kino, The application of reciprocity theory to scattering of acoustic waves by flaws, Journal of Applied Physics. 49, 6; 1978, 3190–3199.

76. J. Zhang, B. Drinkwater, P. Wilcox, Defect characterization using an ultrasonic array to measure the scattering coefficient matrix, Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on. 55, 10; 2008, 2254 –2265.

77. S.D. Thoppul, J. Finegan, R.F. Gibson, Mechanics of mechanically fastened joints in polymer–matrix composite structures – A review, Composites Science and Technology. 69, 3–4; 2009, 301–329.

78. Y. Xiao, T. Ishikawa, Bearing strength and failure behavior of bolted composite joints (part I: Experimental investigation), Composites Science and Technology. 65, 7–8; 2005, 1022–1031.

79. A. Zagrai, V. Gigineishvili, W.A. Kruse, A. Murray, D. Doyle, W. Reynolds, et al., Acousto-elastic measurements and baseline-free assessment of bolted joints using guided waves in space structures, Proc. SPIE, 765017, 2010, San Diego, CA.

80. X. Chen, J.E. Michaels, S.J. Lee, T.E. Michaels, Load-differential imaging for detection and localization of fatigue cracks using Lamb waves, NDT & E International. 51, 2012, 142–149.

81. P.P. Camanho, M. Lambert, A design methodology for mechanically fastened joints in laminated composite materials, Composites Science and Technology. 66, 15; 2006, 3004–3020.

82. S. Minakuchi, T. Nakamura, T. Nadabe, M. Nishikawa, N. Takeda, M. Kishi, et al., Damage Detection of CFRP Bolted Joints Using Embedded Optical Fibers with BOCDA System, Journal of the Japan Society for Aeronautical and Space Sciences. 59, 690; 2011, 176–182.

83. D. Doyle, A. Zagrai, B. Arritt, H. Çakan, Damage Detection in Bolted Space Structures, Journal of Intelligent Material Systems and Structures. 21, 3; 2010, 251–264.

84. F. Lanza di Scalea, H. Matt, I. Bartoli, S. Coccia, G. Park, C. Farrar, Health Monitoring of UAV Wing Skin-to-spar Joints using Guided Waves and Macro Fiber Composite Transducers, Journal of Intelligent Material Systems and Structures. 18, 4; 2007, 373–388.

85. Y. Okabe, K. Fujibayashi, M. Shimazaki, H. Soejima, Damage Detection in Aircraft Composite Materials Using a Built-in Broadband Ultrasonic Propagation System, Journal of System Design and Dynamics. 5, 5; 2011, 966–981.

86. M. Collet, M. Ruzzene, K.A. Cunefare, Generation of Lamb waves through surface mounted macro-fiber composite transducers, Smart Materials and Structures. 20, 2; 2011.

87. H.-T. Sun, F.-K. Chang, X. Qing, The Response of Composite Joints with Bolt-Clamping Loads, Part II: Model Verification, Journal of Composite Materials. 36, 1; 2002, 69–92.

88. T. Nadabe, M. Nishikawa, T. Nakamura, N. Takeda, Damage Evolution Mechanism in Bolted Joints of CFRP Laminates, Proc. 24th ASC Technical Conference, 2009, 174–187, Delaware.

89. C. Haynes, M. Todd, Uncertainty quantification of a guided wave structural health monitoring system for composite bolted joints, Proc. SPIE, 8695, 2013, San Diego, CA.

90. W.H. Prosser, Ultrasonic Characterization of the Nonlinear Elastic Properties of Unidirectional Graphite/Epoxy Composites. NASA Langley Research Center, 1987.

91. S.J. Lee, N. Gandhi, J.E. Michaels, T.E. Michaels, Comparison of the Effects of Applied Loads and Temperature Variations on Guided Wave Propagation, AIP Conference Proceedings. 1335, 1; 2011, 175–182.

92. J.-R. Lee, C. Ciang Chia, C.-Y. Park, H. Jeong, Laser ultrasonic anomalous wave propagation imaging method with adjacent wave subtraction: Algorithm, Optics & Laser Technology. 44, 5; 2012, 1507–1515.

93. M. Ruzzene, Frequency–wavenumber domain filtering for improved damage visualization, Smart Mater. Struct. 16, 6; 2007, 2116.

94. T.E. Michaels, M. Ruzzene, J.E. Michaels, Frequency-wavenumber domain methods for analysis of incident and scattered guided wave fields, Proc. SPIE, 7295, 2009, San Diego, CA.

95. E.B. Flynn, S.Y. Chong, G.J. Jarmer, J.-R. Lee, Full Field Measurement of High Order Ultrasonic Guided Wave Modes for Detecting and Visualizing Structure Defects, Proc. Asia Pacific Vibration Conference, 2013, Jeju, Korea.

96. B. Jahne, Digital Image Processing, Springer-Verlag, Berlin, 1991.

97. R.P. Dalton, P. Cawley, M.J.S. Lowe, The Potential of Guided Waves for Monitoring Large Areas of Metallic Aircraft Fuselage Structure, Journal of Nondestructive Evaluation. 20, 1; 2001, 29–46.

98. Lester W. Schmerr, Jr., Fundamentals of ultrasonic nondestructive evaluation: a modeling approach, Plenum Press, New York, 1998.

99. S.W. Doebling, F.M. Hemez, Overview of uncertainty assessment for structural health monitoring, Proc. 3rd international workshop on structural health monitoring, 2001, Stanford University.

100. B.R. Mace, K. Worden, G. Manson, Uncertainty in structural dynamics, Journal of Sound and Vibration. 288, 3; 2005, 423–429.

101. K. Chaloner, I. Verdinelli, Bayesian Experimental Design: A Review, Statistical Science. 10, 3; 1995, 273–304.

102. D.V. Lindley, Bayesian Statistics, a Review, Society for Industrial and Applied Mathematics, Philadelphia, 1972.

103. M.D. Todd, C.M. Haynes, E.B. Flynn, Bayesian Experimental Design Approach to Optimization in Structural Health Monitoring, Proc. World Congress on Advances in Structural Engineering and Mechanics, 2011, Seoul, Korea.