ULTRASONIC GUIDED WAVE IMAGING VIA SPARSE RECONSTRUCTION

A Thesis
Presented to
The Academic Faculty

by

Ross M. Levine

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology
May 2014

ULTRASONIC GUIDED WAVE IMAGING VIA SPARSE RECONSTRUCTION

Approved by:

Professor Jennifer E. Michaels,
Advisor
Committee Chair
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Justin Romberg
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Waymond R. Scott Jr.
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Thomas E. Michaels
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Professor Massimo Ruzzene
School of Aerospace Engineering
Georgia Institute of Technology

Date Approved: March 5, 2014

ACKNOWLEDGEMENTS

I would first like to express my thanks to my advisor, Dr. Jennifer E. Michaels, for her support over the past several years. It is only with your guidance that I have been able to successfully complete my studies as a graduate student. I have learned so much from you, including new ways to look at problems and how to better communicate ideas; my mind is more able to focus and I have become more disciplined. These gifts have made me into a better person, and I will always remember your kindness just as much as I will your brilliance.

Though it has been difficult, I was never alone in my journey. It was an honor to spend my time in QUEST lab with some of the brightest minds I have ever known. I am especially glad to have met Xin Chen, who started graduate school the same semester as I. It has been a true pleasure to work alongside you, and I am in your debt for your help with all of the experiments in this thesis. I must also mention my dear friend, Navneet Gandhi, who was the social glue of our lab before his graduation; I really miss having you around. And thank also goes Dr. Sang Jun Lee, our lab's postdoc for several years; working with you was fantastic fun!

I would also like to thank Dr. Thomas E. Michaels, for his support and insight, as well as all of the other QUEST lab students with whom I have worked, including James Hall, Ler Gullayanon, Leo Lu, Shiv Chawla, Xander Dawson, Jason Zutty, Mitch Costley, Honglei Li, and Fan Shi. It has been great to know you all and I will miss you terribly.

Thank you to the other members committee, Dr. Justin Romberg, Dr. Waymond R. Scott Jr., and Dr. Massimo Ruzzene. I received great insight during my proposal exam that helped me to carry on my research and successfully complete this

document. I appreciate your time and helpfulness.

I offer my sincere gratitude to my friends and family. Thank you, mom, for always pushing me to succeed at bigger and better things; I may not always have appreciated your forcefulness at the time, but in retrospect I see your parental wisdom and would not be here without you. I sincerely thank my stepfather, Paul, for seeing my true potential and encouraging me to periodically look back at how far I had come. Thank you, Dad, for always being there for me when I needed help – just knowing you were always available at a moment's notice made this task so much easier. And thank you Susan, my stepmother, for your understanding and for making me always feel welcome in your home. Thank you Nanny and Poppy, for your constant love and for great vacations! Coming to visit you in Florida was always perfect to get rid of the stress of grad student life. And thanks go out to all of my friends, who gave nothing but encouragement over the course of my studies. I especially want to thank my brother, Alex, whose lifelong friendship is priceless; thanks for making me laugh every time I would go back home to Kentucky. Good luck with the rest of your studies.

Finally, I must thank my fiancèe, Sara. I love you and your support over these past years has been immeasurable. It is not an exaggeration to say that this would not have been possible without you. You encouraged me on my best days and comforted me on my worst; you relaxed me on my busiest days and chided me on my laziest; you grounded me on my most easy-going days and calmed me on my angriest. At times, I would despair at the thought of the mountain of work that was still to come; there was part of me that believed that this day would never come. Your belief in me gave me the confidence to continue the long march towards eventual graduation. I am so grateful to have you in my life. You are wonderful.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

Symbol	Space/Type	Description
$\mathbf{f}_S$ and $\mathbf{f}_A$	set	cutoff frequencies for Lamb wave symmetric and antisymmetric modes
c_l	$\mathbb{R}_{>0}$	longitudinal wave velocity
c_t	$\mathbb{R}_{>0}$	transverse (shear) wave velocity
h	$\mathbb{R}_{>0}$	plate thickness
$c_p(f)$	$\mathbb{R} \to \mathbb{R}$	Lamb wave dispersion curve (velocity as function of frequency)
d	$\mathbb{R}_{\geq 0}$	propagation distance
d_{ref}	$\mathbb{R}_{>0}$	arbitrary reference distance
$v_0[t]$ or $\mathbf{v}_0$	$\mathbb{R}^L$	source excitation function
$\mathcal{P}_{c_p \atop \mathbf{a} \to \mathbf{b}}$	$\mathbb{R}^L \to \mathbb{R}^L$	propagation operator from $\mathbf{a}$ to $\mathbf{b}$
$\mathcal{P}^{-1}_{c_p \atop \mathbf{a} \to \mathbf{b}}$	$\mathbb{R}^L \to \mathbb{R}^L$	reverse propagation operator to $\mathbf{a}$ from $\mathbf{b}$
N	$\mathbb{Z}_{>0}$	total number of pixels (indexed by n)
$\mathbf{p}_n$	$\mathbb{R}^2$	X-Y coordinate of pixel n
N_T	$\mathbb{Z}_{>0}$	total number of transducers
P	$\mathbb{Z}_{>0}$	total number of transducer pairs (indexed by i)
$\mathbf{s}_i$	$\mathbb{R}^2$	X-Y coordinate of source (transmitter) for pair i
$\mathbf{r}_i$	$\mathbb{R}^2$	X-Y coordinate of receiver for pair i
K	$\mathbb{Z}_{\geq 0}$	total number of scatterers (indexed by k)
$\mathbf{q}_k$	$\mathbb{R}^2$	X-Y coordinate of scatterer k
$v_i^{\text{direct}}[t; \mathbf{s}, \mathbf{r}]$	$\mathbb{R}^L$	direct arrival for pair i
$H[f; \theta^{\text{in}}, \theta^{\text{out}}]$	$\mathbb{R}^L$	scattering pattern (frequency response)

Symbol	Space/Type	Description
$h\left[t;\theta^{\text{in}},\theta^{\text{out}}\right]$	$\mathbb{R}^L$	scattering pattern (impulse response)
$v^{\text{scattered}}\left[t;\cdots\right]$	$\mathbb{R}^L$	scattered signal due to scatterer(s)
$\theta_{\mathbf{a},\mathbf{b}}$	$[0,2\pi)$	angle on X-Y plane of $\mathbf{b}-\mathbf{a}$
$\mathbf{y}_i^{\text{BL}}/\mathbf{y}_i^{\text{FU}}/\mathbf{y}_i$	$\mathbb{R}^L$	measured baseline/follow-up/differential signal for pair i
$\mathbf{e}_i^{\text{BL}}/\mathbf{e}_i^{\text{FU}}/\mathbf{e}_i$	$\mathbb{R}^L$	noise component of signal for pair i
$y_{i,\mathbf{p}}^{\text{Back}}\left[t\right]$	$\mathbb{R}^L$	back-propagated signal for pair i assuming damage at $\mathbf{p}$
$w_{i,\mathbf{p}}$	$\mathbb{R}$	DAS weight for pair i for pixel $\mathbf{p}$
$y_{\mathbf{p}}^{\text{DAS}}\left[t\right]$	$\mathbb{R}^L$	DAS summed signals assuming damage at $\mathbf{p}$
$z_{\mathbf{p}}^{\text{DAS}}$	$\mathbb{R}$	DAS pixel value at $\mathbf{p}$
$\mathbf{y}$	$\mathbb{R}^M=\mathbb{R}^{LP}$	measured signal (concatenation of $\mathbf{y}_1,\ldots,\mathbf{y}_P$)
$\mathbf{x}$	$\mathbb{R}^N$ or $\mathbb{C}^N$	unknown vector with scattering coefficient (or block of coefficients) at each pixel
$\tilde{\mathbf{x}}$	$\mathbb{R}^N$ or $\mathbb{C}^N$	estimate of $\mathbf{x}$
$\mathbf{A}$	$\mathbb{R}^{M\times N}$ $=\mathbb{R}^{LP\times N}$	dictionary matrix; also used as a "generic" matrix for explaining concepts
$\mathbf{A}_i$	$\mathbb{R}^{L\times N}$	dictionary submatrix corresponding to pair i
$\mathbf{a}_{i,n}$	$\mathbb{R}^L$	dictionary submatrix entry corresponding to pair i at pixel $\mathbf{p}_n$
$\mathbf{a}_n$	$\mathbb{R}^L$	dictionary column corresponding to pixel $\mathbf{p}_n$
$\mathbf{D}/\bar{\mathbf{D}}/\hat{\mathbf{D}}$	diagonal matrix	diagonal matrix of dictionary norms
$\bar{\mathbf{y}}/\bar{\mathbf{a}}/\bar{\mathbf{e}}$	nonnegative vector	envelope-detected vector

Symbol	Space/Type	Description
$\hat{\mathbf{y}}/\hat{\mathbf{a}}/\hat{\mathbf{e}}$	complex vector	analytic representation of vector
$\bar{\mathbf{A}}$	$\mathbb{R}^{LP \times N}$	envelope-detected dictionary matrix
$\hat{\mathbf{A}}$	$\mathbb{C}^{LP \times N\Lambda P}$	block dictionary matrix
σ	$\mathbb{R}_{\geq 0}$	BPDN relaxation parameter
σ'	$\mathbb{R}_{\geq 0}$	adjusted BPDN relaxation parameter after dictionary projection
S	$\mathbb{R}_{\geq 0}$	sparsity of a vector
$\delta_S\left(\mathbf{A}\right)$	$\mathbb{R}_{\geq 0}$	S-RIP constant for matrix $\mathbf{A}$
$\mu\left(\mathbf{A}\right)$	$[0,1]$	coherence of matrix $\mathbf{A}$
$\mathbf{P}_{\mathbf{A}}$	$\mathbb{R}^{M \times M}$ or $\mathbb{C}^{M \times M}$	projection matrix onto the column space of $\mathbf{A}$
$\mathcal{G}\left(n\right)$	group index	group index of column n of the block dictionary
G_i	group index	group index number i
$\delta_{S\vert\mathcal{G}}\left(\hat{\mathbf{A}}\right)$	$\mathbb{R}_{\geq 0}$	block-S-RIP constant for matrix $\hat{\mathbf{A}}$
$\hat{\mathbf{A}}_{G_i}$	$\mathbb{C}^{LP \times \#G_i}$	submatrix consisting only of the columns of $\hat{\mathbf{A}}$ with group index G_i; here $\#G_i$ is the number of such columns
$\mathbf{x}_{G_i}$	$\mathbb{C}^{\#G_i}$	subvector consisting only of the entries of $\mathbf{x}$ with group index G_i
$\mathbf{y}^{\Vert\mathbf{A}}$	$\mathbb{R}^L$ or $\mathbb{C}^L$	projection of $\mathbf{y}$ onto the column space of $\mathbf{A}$
$\mathbf{y}^{\perp\mathbf{A}}$	$\mathbb{R}^L$ or $\mathbb{C}^L$	projection of $\mathbf{y}$ onto the left null space of $\mathbf{A}$
$E_{\Vert}$	$[0,1]$	fraction of signal energy remaining after projection
θ_k	$[0,2\pi)$	kth principal angle between two subspaces
ρ_k	$[0,1]$	cosine of θ_k; kth canonical correlation

SUMMARY

The field of structural health monitoring (SHM) is concerned with the continu-
ous, long-term assessment of structural integrity. So-called "smart structures" have
embedded sensing elements that interrogate the structural framework and supply
time-history data to algorithms designed to recognize the appearance and progres-
sion of damage. One commonly investigated SHM technique uses guided ultrasonic
waves, which travel through the structure and interact with damage. Measured sig-
nals are then analyzed in software for detection, estimation, and characterization of
damage. A particularly attractive configuration for such a system uses a spatially-
distributed array of fixed piezoelectric transducers; such a setup is inexpensive and
can cover large areas. Typically, one or more sets of prerecorded baseline signals are
recorded when the structure is in a known state, with imaging methods operating on
differences between follow-up measurements and these baselines.

Conventionally, images are created using the well-known delay-and-sum imaging
algorithm, which back-propagates signals and then adds them together. This al-
gorithm often performs poorly when multiple sites of damage are present or when
interference is present in the signals due to multipath effects or poor baseline sub-
traction. Presented in this dissertation is a different class of algorithms that rely on
sparse reconstruction, which attempts to solve inverse problems for which the solu-
tion is known to have structure. Using an intelligently-selected redundant dictionary,
signals can be decomposed into a small number of atoms that have some physical
meaning. For this problem, that meaning is location-based, allowing a large recon-
struction coefficient to directly correspond to damage at some location. A new class of

Lamb wave SHM imaging methods is developed that uses this construction of sparsity to produce imaging results that are superior to conventional delay-and-sum methods.

Two types of sparse imaging techniques are demonstrated in this work. The first, which relies on sparse reconstruction of raw signals, signal analytic representations, or signal envelopes, uses an *a priori* assumption of scattering behavior to generate a redundant dictionary matrix where each column corresponds to a pixel in the two-dimensional image. The measured signals are modeled as a linear combination of a small number of dictionary columns, with damage at a particular pixel indicated by a non-zero coefficient for its corresponding column. The second method extends this concept by using multidimensional models for each pixel, with each possible location on the discretized region of interest corresponding to a "block" in the dictionary matrix instead of a single column. This block-sparse method does not require any advance knowledge or assumptions of scattering behavior.

The contributions of this work include:

- Formulation of damage detection and imaging as a sparse (and block-sparse) reconstruction problem;

- Analysis of the effects of envelope detection on noisy signals that are used for sparse reconstruction;

- Experimental verification of the methods, including using nondispersive dictionaries to demonstrate robustness; and

- A method to extract scattering patterns from block-sparse imaging results.

The analysis and experimental results presented demonstrate the validity of the assumption of damage sparsity. Additionally, experiments show that images generated with sparse methods are superior to those created with delay-and-sum methods; sparsity-based techniques are experimentally shown to be tolerant of propagation

model mismatch. The block-sparse method described here also allows the extraction of scattering patterns from its reconstruction coefficients, which can be used for damage characterization; this capability is verified with experimental data.

CHAPTER I

INTRODUCTION

1.1 *Overview*

Ultrasonic testing has become a reliable and accepted method of damage interrogation for over a half-century. It is possible to excite ultrasonic waves in a structure such as an aircraft wing and examine the resultant wavefield for reflections from internal flaws, which may be physically inaccessible or too small to detect by visual inspection. Traditional ultrasonic inspection, however, is often intrusive and costly; for example, aircraft are partially disassembled for scheduled inspections. The more recent field of structural health monitoring envisions a different paradigm, where sensors are integrated into the structure and regularly interrogate it for damage. These data can be monitored over the life of the structure to detect any new damage formation and to monitor the severity of existing damage. If damage is detected, a more thorough offline inspection can then be performed.

One particularly attractive method of performing ultrasonic SHM is through the use of guided waves. In a plate-like structure, these waves are named *Lamb waves*, after mathematician Horace Lamb, who published the first theoretical description and analysis of their behavior [1] almost a century ago. Though the potential benefits of Lamb wave ultrasonic testing were recognized in the mid-20th century [2], it took many decades before such techniques became practical. This is for a variety of reasons, including the need for more powerful computers and a more complete understanding of Lamb wave excitation, measurement, and behavior. Research into nondestructive testing and evaluation with guided ultrasonic waves, including Lamb waves, began in the late 1980s and matured throughout the 1990s; a good summary

of the techniques developed is presented in [3]. Early investigations into guided wave SHM focused on Lamb wave behavior resulting from various manifestations of damage in different material types and shapes [4, 5, 6], including analysis of the advantages and disadvantages of different Lamb wave modes with respect to damage type. Some of the first damage-detection methods included tomography [7, 8] and, later, the use of phased arrays [9, 10, 11], sparse arrays [12, 13, 14, 15], and synthetic aperture techniques [16, 17].

Of particular note is the work by Wang *et al.* in 2004 [12] that introduced the use of signal baselines in the context of sparse array imaging. This research also adapted delay-and-sum beamforming, a well-known radar technique, for use with guided waves for the problem of damage detection and localization with a sparse array. Delay-and-sum techniques are still considered the gold standard in SHM due to their robustness and conceptual simplicity, but the performance of such methods is limited, especially in cases with multiple sites of damage or high levels of interference or clutter. While various improvements have been proposed, e.g. with statistical models or other *a priori* assumptions, to date there has been no effort to incorporate the reasonable assumption of damage *sparsity*. This research utilizes the sparsity assumption – that is, the assumption that damage is limited to a small number of discrete locations – to improve detection and localization of damage in plate-like structures.

Recently, interest in the field of *sparse reconstruction* has received considerable attention, in part due to the rapid development of the related field of *compressed sensing*. The fundamental idea behind sparse reconstruction is that most "interesting" signals can be represented with a very small number of coefficients, provided their representation is chosen in a smart way. Instead of using least-squares methods, linear inverse problems are solved using methods that produce sparse solutions (i.e., vectors whose entries are mostly zero-valued). This *a priori* assumption of a sparse signal can often yield remarkable results when compared to least-squares methods;

under certain circumstances, it is even possible to guarantee exact recovery [18]. A sparse solution usually makes the most sense in the context of an *overcomplete dictionary* representation, which contains a very large number of columns, or *atoms*. The measurement vector is assumed to be linear combination of a small subset of the atoms; the "hard" portion of the problem is properly selecting which elements belong in this subset for a given measurement.

1.2 The Problem of Interest

1.2.1 Introduction to the Problem

The work here considers a linear, elastic plate in a low-noise laboratory environment and uses a sparse array of transducers. For analytical purposes, the plate is considered to be of infinite size, and edge reflections from real data are treated as interference (i.e., coherent noise). Some closed region of the plate is designated the region of interest (ROI), inside which it is desired to detect any damage. A total of N_T transducers are affixed to the plate; to simplify notation, the transfer function of each transducer is ignored (i.e., assumed to be 1), though incorporating known transducer transfer functions into the algorithms presented is straightforward. This results in a total of $P = N_T (N_T - 1)/2$ unique transducer pairs, with some arbitrary ordering, $1, 2, \ldots, P$. Within each pair, one transducer is designated the source and the other the receiver; let $\mathbf{s}_i$ represent the 2-dimensional location of the source for pair i, and let $\mathbf{r}_i$ represent the location of the receiver.

It is assumed that the material properties of the plate and transducers are known; in particular, there is a frequency that allows single-mode propagation as described in [19], and it is assumed that the dispersion curve of that mode, $c_p[f]$, is known. (If the dispersion curve is not known, the group velocity can be used instead; the effects of such a substitution are analyzed in Chapter 6.) Details about Lamb wave propagation and scattering are presented in Section 2.2.

3

1.2.2 Baseline and Follow-Up Measurements

For the following sections, it is be useful to use vector notation; appropriate choices of sampling frequency and signal duration are assumed. A baseline measurement set is taken while the plate is in a known condition, with one transducer transmitting at a time while the others record received waves. Reciprocal signals (i.e., from a pair's receiver to its source) may be recorded, but are not used for reconstruction, as the principle of reciprocity dictates that the two reciprocal signals should be the same. (However, these signal pairs can be used for other purposes, such as detecting transducer failure [20].) Denote the excitation signal $\mathbf{v}_0 \in \mathbb{R}^L$, and the measured waveform for pair i as $\mathbf{y}_i^{\mathrm{BL}} \in \mathbb{R}^L$.

At some future time, a follow-up measurement is performed, under the same conditions and in the same fashion as the baseline measurement. Denote the follow-up wavefield measurement for pair i as $\mathbf{y}_i^{\mathrm{FU}} \in \mathbb{R}^L$. If no damage was introduced, these signals should be equal to the baseline signals (except for noise); otherwise, differences are assumed to be caused by newly-introduced damage.

1.2.3 Residual Measurement

At this point, baseline subtraction is performed; this is further described in Section 2.3.2; the residual signal vectors are denoted $\mathbf{y}_i$. The simplest type of baseline subtraction is simple subtraction, where the residual signals are computed as $\mathbf{y}_i = \mathbf{y}_i^{\mathrm{FU}} - \mathbf{y}_i^{\mathrm{BL}}$. All signal processing is performed directly on the residual measurements, which represent the changes to the plate since the baseline was recorded.

1.3 Notation

Throughout this text, bold, capital letters denote matrices (e.g., the matrix $\mathbf{A}$). Bold, lower-case letters denote vectors; this notation is used interchangeably with signal notation whenever it is convenient (so, $\mathbf{v}_0$ and $v_0[t]$ both refer to the same vector,

but in different contexts). A column of a matrix is referenced by a lower-case letter and subscript (e.g., if context provides, $\mathbf{a}_i$ is the ith column of the matrix $\mathbf{A}$). Scalars are never bold and can use the same notation to refer to an entry of a vector (e.g., x_i is the ith entry of $\mathbf{x}$). All figures use a custom colorset that appears grayscale when printed in black and white.

1.4 Organization

This research utilizes sparse reconstruction techniques for the problems of damage detection and localization. First, a history of both the problem and of sparse reconstruction is presented in detail in Chapters 2 and 3. Next, it is shown in Chapter 4 that a redundant dictionary of location-based signal components provides a representation in which scattered signals are sufficiently sparse, and that furthermore solving a linear inverse problem with sparse reconstruction is equivalent to finding sites of damage. This technique is then extended in Chapter 5 to incorporate a multidimensional linear model that uses block-sparse reconstruction to allow for robust damage detection as well as characterization. Extensive experimental results are shown in Chapter 6 and compared to existing delay-and-sum techniques, and concluding remarks are made in Chapter 7.

CHAPTER II

BACKGROUND ON ULTRASONICS

2.1   *Introduction*

Nondestructive testing and evaluation (NDT&E) is a collective term for many techniques that inspect a structure without damaging it. Methods include the use of X-rays, eddy currents, and, of particular importance to the proposed research, ultrasonic waves. Ultrasonic testing is used extensively in many industries; for example, it is common in the aircraft industry, where it is used in a wide variety of materials [21, pp. 8-21]. A particular advantage is the ability to detect subsurface flaws in many materials without harming the structure or posing health risks to the operator [22]. The most common form of ultrasonic NDT&E is through the use of bulk waves that propagate through the material, either in a specific direction or in a spherical pattern, depending on the method of excitation. For example, a piezoelectric transducer may be used to excite an ultrasonic wave which propagates into the material. This wave is then reflected at the opposite surface and the amplitude response is recorded at the surface, either by the same transducer (a configuration known as *pulse-echo*) or a second transducer (called *through-transmission* when the transducers are on opposite sides of the material, or *pitch-catch* when they are on the same side). If a flaw is present, either additional reflections will be present in the recorded signal, or an expected echo will be missing or modified; both are consequences of interaction of the ultrasonic wave with the flaw.

While well-established, these conventional methods can be quite expensive and time-consuming. For example, in a pulse-echo configuration, the entire structure needs to be "scanned," since only flaws in close proximity to the transducer will be

detected. Structures examined with this method are typically inspected by hand or with the help of robotic scanning systems. In the U.S. airline industry, the FAA mandates inspections either annually or as part of a continuous airworthness maintenance program [23]. Most airlines choose the latter option, which typically requires frequent but less thorough routine inspections and occasional exhaustive inspections [21, pp. 8-15], which require the aircraft to be disassembled for comprehensive testing. Other structures such as bridges can be quite difficult to examine, since there are many areas which inspectors cannot access. Even in cases where manual inspection is readily available, operators must be well-trained to discriminate between scattered signals from damage and other types of signal scattering such as from rivets or material boundaries.

The field of structural health monitoring (SHM) focuses on continuous monitoring of structures to reduce or eliminate the need for these costly and time-consuming manual tests. Time-history data is often emphasized in many SHM systems to attempt to determine and predict the integrity of a structure and determine when maintenance or replacement is necessary. Ultrasonic methods are one of several techniques used in SHM, and are especially suited for aircraft. In contrast, other sensing methods are often more appropriate for general monitoring of very large structures such as bridges; for example, after the I-35W Mississippi River Bridge tragically collapsed in 2007, it was replaced with a "smart bridge" that incorporated strain gauges, accelerometers, and fiber optic cable [24]. Nevertheless, ultrasonic techniques work well for portions of these civil structures, such as welds [25] and portions of the structure with complex geometry [26]. Ultrasonic guided waves are also commonly used to inspect pipes [27], where the one-dimensional propagation enables very long travel distances [28].

2.2   Guided Ultrasonic Waves

2.2.1   Introduction

In linear, isotropic solids, there are two different types of ultrasonic bulk waves: longitudinal (typically called *P-waves*, short for primary or pressure waves) and transverse (known as *S-waves*, short for secondary or shear waves). These two types of waves have different propagation velocities, c_l and c_t, respectively; with the exception of certain metamaterials, $c_l > c_t$ in solids [29, p. 124]. Mathematician Horace Lamb predicted the existence of ultrasonic guided waves in plates in 1917 that could form as a result of the interaction of the two types of bulk waves with the plate boundaries [1]. These waves use their solid medium as a waveguide; i.e, they propagate cylindrically outward as a two-dimensional wave within the structure when excited at a point.

D. C. Worlton proposed an ultrasonic inspection method in the late 1950s [2] using ultrasonic guided waves. At the time, they were predicted only in theory; in 1961 he published the first experimental verification of their existence [30]. Lamb waves are particularly attractive because of their ability to propagate over relatively large distances, typically a meter or more, and can be sensitive to damage throughout the plate thickness [4]. Lamb wave behavior, however, is complicated: the waves are dispersive (i.e., different frequency components propagate at different velocities), and like many types of guided waves, Lamb waves exist as a countably infinite number of propagation modes. Each of these modes arises when the longitudinal and shear bulk waves interact with the two plate surfaces in such a way that they coalesce into a propagating mode.

2.2.2   Lamb Wave Propagation Modes

Strictly speaking, the term "Lamb wave" has historically referred only to waves with out-of-plane motion in a flat, linear, isotropic plate; however, the term is often applied

more loosely to include the horizontally-polarized SH waves and/or any ultrasonic guided wave in a plate-like structure, including layered plates and pipes with low curvature. Here, the term is used only to refer to the more strict definition, with other waves referred to simply as guided waves or quasi-Lamb waves.

A flat, linear, isotropic plate supports an infinite number of Lamb wave propagation modes [31, p. 70]; the behavior of these propagation modes depends on the plate thickness h; the longitudinal and transverse (shear) bulk wave velocities, c_l and c_t, respectively; and the frequency f. A Lamb wave has a *class* and an integral *order*. There are two classes of propagation modes, symmetric modes and antisymmetric modes, whose names refer to the displacement profile of a propagating wave. The propagating Lamb wave modes have whole-number orders; the symmetric and antisymmetric modes of order n are denoted S_n and A_n, respectively. With the exception of the zero-order S_0 and A_0 modes, which may exist at all frequencies, a Lamb wave mode may not exist below its corresponding nascent frequency. The nascent frequencies of symmetric and antisymmetric modes are of the form [31, p. 71]

$$
\begin{aligned}
\mathbf{f}_\text{S} &= \left\{ 0.5\frac{c_l}{h},\ \frac{c_t}{h},\ 1.5\frac{c_l}{h},\ 2\frac{c_t}{h},\ 2.5\frac{c_l}{h}, \ldots \right\} \\
\mathbf{f}_\text{A} &= \left\{ 0.5\frac{c_t}{h},\ \frac{c_l}{h},\ 1.5\frac{c_t}{h},\ 2\frac{c_l}{h},\ 2.5\frac{c_t}{h}, \ldots \right\},
\end{aligned}
\tag{1}
$$

where h is the plate thickness. Note that since c_l and c_t are unequal, the cutoff frequencies for each mode class are not necessarily in the order presented in Eq. 1; the symmetric and antisymmetric mode orders are always numbered by increasing nascent frequency. In addition to the propagating Lamb wave modes, there are an infinite number of non-propagating (evanescent) modes which exist only in the near-field of their point of excitation or at plate boundaries. These waves are not typically used for damage detection and their existence may often be neglected.

Signals that contain multiple modes are significantly more difficult to analyze than those that contain a single mode. At low enough frequencies, only the S_0 and A_0 Lamb wave modes may exist. Since $c_l > c_t$ in normal solids, any excitation

at frequencies below $f_{A_1} = 0.5\frac{c_t}{h}$ will eliminate any higher-order Lamb wave modes, limiting any signal at such frequencies to two modes. Furthermore, it is often possible to select a frequency for which a transducer is "tuned" to one of the two fundamental modes, suppressing the other almost completely to create an almost purely single-mode Lamb wave excitation [19]. This allows Lamb wave SHM systems to use waves that are approximately single-mode to simplify analysis, though the specific frequency required depends on the structure in question as well as the transducer geometry.

In addition to symmetric and antisymmetric Lamb wave modes, there are also SH waves which have horizontal polarization, perpendicular to the direction of propagation [32, p. 190]. These guided waves are also used for SHM [33], but their use is not described here. Circular transducers with vertical or radial polarization, such as those used for the experiments here, do not generate or effectively measure SH waves, allowing this type of guided wave to be safely ignored.

2.2.3 Dispersion

Each Lamb wave mode has its own dispersion profile, which describes the phase and group velocities of the wave mode as a function of frequency. These profiles, also called dispersion curves, are most often calculated with computer software such as DISPERSE [34] or Vallen Dispersion [35]. The dispersion curve of a Lamb wave mode expresses the angular frequency ω as a function of wavenumber k and determines how the shape of a wave pulse changes as it propagates. The phase velocity c_p can be expressed through the relation $c_p := \omega/k = 2\pi f/k$; the group velocity is defined as $c_g := \partial\omega/\partial k$. For a specific Lamb wave mode with a known phase velocity dispersion curve $c_p(f)$, it is possible to describe the far-field shape of a cylindrically-propagating wave [32, p. 220] after it has propagated a distance d from its source:

$$v(t) = \mathcal{F}^{-1}\left\{ \left(\frac{d}{d_{\text{ref}}}\right)^{-\frac{1}{2}} \left(\mathcal{F}\big\{v_0(t)\big\}(f)\right) \exp\left(\frac{-\mathrm{i}2\pi f d}{c_p(f)}\right)\right\}, \qquad (2)$$

where $v_0(t)$ is the time-domain waveform excited by the source, d_{ref} is some reference distance, $\mathcal{F}\{\cdot\}$ denotes the discrete Fourier transform, and $\text{i} = \sqrt{-1}$. The first term in Eq. 2 describes amplitude decay via geometric spreading, while the last term describes propagation behavior of individual frequencies. It is convenient to express Eq. 2 more concisely and in discrete time. Let $v_0[t]$ be the source excitation, sampled at an appropriate frequency (i.e., above the Nyquist rate). Define the propagation operator $\mathcal{P}_{c_p}$ in terms of v_0 to represent the shape of a single-mode guided wave packet which is excited at location $\mathbf{p}_1$ and measured at location $\mathbf{p}_2$ with a linear, isotropic plate as a medium:

$$\left(\underset{\mathbf{p_1} \to \mathbf{p_2}}{\mathcal{P}_{c_p}} v_0 \right)[t] := \mathcal{F}^{-1} \left\{ \left(\frac{\|\mathbf{p}_2 - \mathbf{p}_1\|_2}{d_{\text{ref}}} \right)^{-\frac{1}{2}} \mathcal{F}\{v_0\}[f] \, \exp\left(\frac{-\text{i}2\pi f \|\mathbf{p}_2 - \mathbf{p}_1\|_2}{c_p[f]} \right) \right\}. \tag{3}$$

The inverse operation $\mathcal{P}_{c_p}^{-1}$ indicates reverse propagation and is defined

$$\left(\underset{\mathbf{p_1} \to \mathbf{p_2}}{\mathcal{P}_{c_p}^{-1}} v_0 \right)[t] := \mathcal{F}^{-1} \left\{ \left(\frac{\|\mathbf{p}_2 - \mathbf{p}_1\|_2}{d_{\text{ref}}} \right)^{+\frac{1}{2}} \mathcal{F}\{v_0\}[f] \, \exp\left(\frac{+\text{i}2\pi f \|\mathbf{p}_2 - \mathbf{p}_1\|_2}{c_p[f]} \right) \right\}. \tag{4}$$

Due to the superposition principle, these operations are linear in the argument v_0.

2.2.4 Guided Wave Interaction with Damage

Consider an undamaged, infinite plate with two ideal point-transducers; one is the source (transmitter) and is located at $\mathbf{s} := [s_x, s_y]^\mathsf{T}$ and the other is the receiver and is located at $\mathbf{r} := [r_x, r_y]^\mathsf{T}$. Suppose a single-mode Lamb wave is excited by the source transducer, with a time-domain waveform $v_0[t]$. The waveform measured by the receiver, ignoring noise and material attenuation, is then measured as

$$v^{\text{direct}}[t] = \underset{\mathbf{s} \to \mathbf{r}}{\mathcal{P}_{c_p}} v_0[t]. \tag{5}$$

A guided wave is scattered when it interacts with a defect within the medium of propagation. The behavior of this secondary wave is a function of the geometry of the defect, and may depend on numerous factors, including incident angle, Lamb wave

mode, and frequency. The scattered wave may be highly directional and is typically phase-shifted from the original wave. It is also possible for waves of other modes to appear, a phenomenon known as *mode conversion.*

Small scatterers are often modeled as point scatterers that act using a linear scaling called a *scattering pattern,* which is the frequency response as a function of the incoming and outgoing angles. The scattering pattern is sometimes represented as a two-dimensional lookup table known as a *scattering matrix* [36]. The point scattering assumption is valid in the far field of the defect (outside the range of the Lamb wave evanescent modes, which is typically on the order of one wavelength) and when the order of magnitude of the size of the scatterer is the same as, or smaller than, that of the Lamb wave mode's wavelength. More formally, consider once again the two transducers located at $\mathbf{s}$ and $\mathbf{r}$, and consider some scatterer at location $\mathbf{q}$. The model of the scattering pattern $H\left[f; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right]$ assumes that the scattered portion of the wavefield measurement by the receiving transducer is (once again, ignoring noise)

$$
\begin{aligned}
v^{\mathrm{scattered}}\left[t; \mathbf{s}, \mathbf{r}, \mathbf{q}, H\right] &= \mathcal{P}_{c_p}_{\mathbf{q} \to \mathbf{r}} \left(h\left[t; \theta_{\mathbf{s},\mathbf{q}}, \theta_{\mathbf{q},\mathbf{r}}\right] * \mathcal{P}_{c_p}_{\mathbf{s} \to \mathbf{q}} v_0\left[t\right] \right) \\
&= h\left[t; \theta_{\mathbf{s},\mathbf{q}}, \theta_{\mathbf{q},\mathbf{r}}\right] * \mathcal{P}_{c_p}_{\mathbf{q} \to \mathbf{r}} \mathcal{P}_{c_p}_{\mathbf{s} \to \mathbf{q}} v_0\left[t\right],
\end{aligned}
\tag{6}
$$

where $\theta_{\mathbf{p}_1,\mathbf{p}_2} = \angle\left(\mathbf{p}_2 - \mathbf{p}_1\right)$ and $h = \mathcal{F}^{-1}\{H\}$. (Note that in practice, the convolution in Eq. 6 is performed by multiplying by $H\left[f; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right]$ in the frequency domain, but this is tedious to repeatedly express with notation used here. Instead, the scattering impulse response is used for conceptual simplicity.) H may depend on additional parameters that are not explicitly indicated here. If a defect induces mode conversion, there will be multiple scattering patterns; for example $H_{\mathrm{S_0} \to \mathrm{A_0}}\left[f; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right]$ would contain the amplitude and phase shift information of an $\mathrm{S_0}$-to-$\mathrm{A_0}$ mode conversion. Since the excitation frequency is assumed to be chosen to tune a transducer to a single mode on both generation and reception, any mode conversion can be neglected. In addition to analytical results of simple damage shapes, estimates of scattering patterns

for many types of damage such as through-holes and cracks are commonly obtained using finite element simulations or from wavefield experiments using scanning laser vibrometry (e.g., [37, 38, 39, 40, 41]); however, in many cases it is impossible to predict in advance what type of damage will occur in a structure. A small library of precomputed scattering patterns is of limited use in these cases.

A common simplifying assumption in damage detection is the *single-scattering* assumption, where multiple defects are assumed to only interact with the direct wave; second-order scattering is neglected. For K scatterers, where scatterer k is at location $\mathbf{q}_k$, the combined scattered waveform is approximated as

$$v^{\text{scattered}}\left[t; \mathbf{s}, \mathbf{r}, \{\mathbf{q}_k\}, \{H_k\}\right] \approx \sum_{k=1}^{K} \left\{ h_k\left[t; \theta_{\mathbf{s}, \mathbf{q}_k}, \theta_{\mathbf{q}_k, \mathbf{r}}\right] \; \ast \; \underset{\mathbf{q}_k \to \mathbf{r}}{\mathcal{P}_{c_p}} \; \underset{\mathbf{s} \to \mathbf{q}_k}{\mathcal{P}_{c_p}} \; v_0\left[t\right] \right\}. \qquad (7)$$

This sort of simplification is related to the *Born approximation*, which is used in a diverse set of fields related to scattering theory, from synthetic-aperture radar to quantum mechanics, where a scattered field is assumed to be a function only of the incident field. In addition to a reduction in analytical and computational complexity, this simplification also linearizes $v^{\text{scattered}}$ in terms of each piece of damage, which is necessary to use sparse reconstruction methods, and is implicit in many other imaging algorithms, such as DAS. This type of assumption is only valid when the magnitude of the scattered field is much smaller than that of the incident field. For the case of Lamb waves, the scattered fields of interest (i.e., when damage is small) are typically one to two orders of magnitude below that of the incident field.

2.3 Guided Wave Structural Health Monitoring

Recently, the concept of structural health monitoring (SHM) has been considered, either as a replacement for or supplement to conventional NDT&E. The fundamental concept of SHM is the implementation of automated material interrogation processes that provide continuous monitoring of a structure, such as a pipe, bridge, or aircraft. An SHM system thus provides a set of historical data for the structure, allowing

for detection of changes over time, and even the ability to, for example, track the growth of a crack and determine when it poses a risk of material failure. The eventual goal of SHM research would lead to fewer unnecessary periodic inspections (i.e., testing *smarter* instead of *harder*), or even, eventually, a near-complete replacement of existing NDT&E methods for some industries.

Guided wave SHM is well-established in pipes, for which wave propagation is often essentially one-dimensional. The lack of geometric spreading enables guided waves to travel distances of tens of meters or more from a single point of excitation, and simple time-of-flight calculations can be performed to determine the locations of scatterers that correspond to signal echoes. In contrast, damage detection and localization for two-dimensional structures such as plates is more difficult and is less mature.

2.3.1 Common Lamb Wave SHM Configurations

2.3.1.1 Sparse Arrays

One potential approach to Lamb wave SHM is the use of a sparse, or spatially-distributed, array of N_T fixed piezoelectric transducers, in either a repeating pattern [42] or randomized arrangement [43, 44], with each transducer capable of generating and receiving Lamb waves [45]. The shape and location of the array can be selected according to different criteria; for example, to minimize the effects of geometric spreading [42], to avoid the creation of areas with different levels of sensitivity to damage [46], or to maximize the probability of detecting certain types of scatterers [47]. The spatially-distributed array is a commonly proposed configuration because of its low cost and ease of implementation. In typical operation, measurements are taken in a round-robin procedure: one transducer generates a Lamb wave while the other $N_T - 1$ transducers record signals; this process is repeated N_T times, so that each transducer gets a turn as the transmitting element. This results in a total of $P = N_T (N_T - 1)/2$ unique transducer pairs. The acoustic principle of reciprocity states that both signals from a single pair (i.e., both A→B and B→A) should be

14

identical; therefore a total of P signals are recorded, excluding the pulse-echo signals which are not used here. Such a set of signals is referred to as a *measurement set*. This research is performed with a sparse array configuration, but there is nothing that restricts the concepts developed here to such a setup.

Sparse arrays are inexpensive because they use a small number of elements. They also have the advantage of interrogating damage at a wide variety of angles, receiving forward- and back-scattered signals. One downside to this method is that it can be difficult to distinguish between direct arrivals and forward-scattered signals from flaws in the material. The most common way to differentiate these two components is by subtracting a prerecorded baseline signal set that contains only the direct arrival [12].

2.3.1.2 Phased Arrays

The concept of a phased array was first created by physicist Karl Ferdinand Braun in 1905 [48], who shared the Nobel Prize in Physics in 1909 with Guglielmo Marconi "in recognition of their contributions to the development of wireless telegraphy." Phased arrays are widely used for electromagnetic waves, from communications to radar, and also for acoustic waves; e.g., sonar [49, p. 84]. Phased arrays of various geometries are used in ultrasonics for both bulk waves [50, 51, 52, 53] and guided waves [9, 10, 11]. A phased array uses many transmitting elements to "steer" a beam in a particular direction; the array is designed so that constructive interference will maximize the signal energy in that direction. This steering can be accomplished in various ways. Some phased arrays have fixed elements which are steered solely by the element spacing and delay lines; these arrays are usually manually steered, e.g. by gimbal mounting for radar arrays. Other phased arrays are dynamic and can have a programmable phase shift or time delay assigned to individual elements, which can be used to electronically steer the beam. Most ultrasonic phased arrays are dynamic and

use time-delay steering, including many medical ultrasound devices as well as some arrays used for inspection, and use a compact 1- or 2-dimensional array of transducer elements; however, a number of devices feature a common pulsing mode that excites all transducers simultaneously to generate images more quickly than sweeping a beam over a wide area.

In addition to the capability of steered excitation, a dynamic phased array can be used in a mode similar to that of a sparse array, with each element pulsing one-at-a-time to collect a full set of pairwise signals. This acquisition method is often called full matrix capture (FMC) and typically includes reciprocal pairs and pulse-echo signals. Beamforming can then be performed using the total focusing method (TFM) [52], which is essentially equivalent to performing DAS imaging on the FMC signals.

Phased arrays have the advantage of being compact and steerable when used for Lamb wave inspection, but only receive back-scattered signals from damage. Depending on the distance from the array, excitation can be complicated when interrogating areas close to the array because of the cylindrical (or, for bulk waves, spherical) shape of the propagation [54]. Phased arrays are also expensive compared to the relative simplicity of the sparse array setup.

If a movable excitation method is available (e.g., an air-coupled transducer on a CNC system), synthetic aperture methods can be used to emulate a phased array [55, 56]; alternatively, a phased array can be used in a synthetic aperture mode by collecting pulse-echo signals one-at-a-time instead of performing full matrix capture. Synthetic aperture ultrasonic techniques are borrowed from radar, where the concept has existed for over 50 years [57]. The use of delay-and-sum imaging on ultrasonic synthetic aperture data is often referred to as the synthetic aperture focusing technique (SAFT).

2.3.2 Baseline Subtraction

For many SHM problems, a minimum of two measurement sets are recorded: a baseline measurement set, which is taken when the structure is known to be damage-free (or some other known state), and subsequent measurement sets to detect if any damage has been introduced. If geometrical reflections are ignored, the baseline measurements are simply the direct arrival signals as described in Eq. 5. Each baseline signal $i = 1, 2, \ldots, P$ can be represented as

$$y_i^{\mathrm{BL}}[t] = \mathcal{P}_{c_p} \atop \mathbf{s}_i \to \mathbf{r}_i} v_0[t] + e_i^{\mathrm{BL}}[t], \tag{8}$$

where $\mathbf{s}_i$ and $\mathbf{r}_i$ are the locations of the transmitter and receiver for pair i, $v_0[t]$ is the excitation function (which is assumed to be the same for every pair), and $e_i^{\mathrm{BL}}[t]$ is a term that encompasses noise and interference (clutter). If K scatterers are introduced at locations $\mathbf{q}_k$, the follow-up measurement $y_i^{\mathrm{FU}}[t]$ is a sum of the direct arrival and the scattered signals:

$$y_i^{\mathrm{FU}}[t] = \mathcal{P}_{c_p} \atop \mathbf{s}_i \to \mathbf{r}_i} v_0[t] + \sum_{k=1}^{K} \left\{ h_k[t; \theta_{\mathbf{s}, \mathbf{q}_k}, \theta_{\mathbf{q}_k, \mathbf{r}}] * \mathcal{P}_{c_p} \atop \mathbf{q}_k \to \mathbf{r}} \mathcal{P}_{c_p} \atop \mathbf{s} \to \mathbf{q}_k} v_0[t] \right\} + e_i^{\mathrm{FU}}[t]. \tag{9}$$

Here any second-order scattering effects are considered to be negligible and are therefore incorporated into the noise term. This equation also assumes no mode conversion (or, that it is also represented in the noise term).

Since SHM techniques attempt to detect changes over time, it is often useful to consider the difference between two measurement sets. If no environmental changes are present, a differential measurement set can be obtained by simply subtracting each signal in the second measurement from the corresponding baseline [13]. Using this method and the same assumptions as in Eqs. 8-9, the signals $y_i[t]$ in this differential

set can be expressed as

$$
\begin{aligned}
y_i[t] &:= y_i^{\mathrm{FU}}[t] - y_i^{\mathrm{BL}}[t] \\
&= \sum_{k=1}^{K} \left\{ h_k\left[t; \theta_{\mathbf{s},\mathbf{q}_k}, \theta_{\mathbf{q}_k,\mathbf{r}}\right] \ * \ \underset{\mathbf{q}_k \to \mathbf{r}}{\mathcal{P}_{c_p}} \ \underset{\mathbf{s} \to \mathbf{q}_k}{\mathcal{P}_{c_p}} \ v_0[t] \right\} + \left(e_i^{\mathrm{FU}}[t] - e_i^{\mathrm{BL}}[t]\right) \qquad (10) \\
&= v^{\mathrm{scattered}}\left[t; \mathbf{s}_i, \mathbf{r}_i, \{\mathbf{q}_k\}, \{H_k\}\right] + e_i[t],
\end{aligned}
$$

which, under ideal conditions, completely eliminates the direct arrival and leaves only the scattered signals of Eq. 7 (plus noise).

Baseline subtraction under changing environmental conditions is more complicated. For example, if the surrounding temperature changes between the baseline and follow-up measurements, the signals will be mismatched; this is because the material's thickness and bulk wave velocities are functions of temperature, and these properties affect the Lamb wave mode's dispersion curves. When such a mismatch is present, simple subtraction will produce large artifacts and can mask scattered signals [42], because a change in temperature affects the arrival times of signals [58]. Adaptive baseline subtraction is an area of active research; e.g., [59, 60]. Applied loads also affect Lamb wave propagation [61, 62], and the differences between signals at different loads can be used in place of residuals from baselines [63]. Regardless of the actual method used in practice, the analyses presented in later chapters assume ideal baseline subtraction.

2.3.2.1 *Optimal Baseline Selection*

One common adaptive strategy to correctly subtract baseline signals is *optimal baseline selection* (OBS) [60, 64]. Instead of a single measurement set of baseline signals, multiple baseline sets are recorded at various environmental conditions. When follow-up measurements are recorded, they are matched to the closest baseline set, using some distance metric. Such sets of baselines are easy to obtain in a laboratory, but it may be infeasible for larger structures that are already in-use. Another problem with OBS is that there may be a large number of environmental parameters that all

18

have different effects on signals (e.g., temperature, structural loading, humidity), with each additional parameter geometrically increasing the number of required baselines. Even for a small number of such parameters, selecting an appropriate range of values and level of discretization can be difficult. In some cases, interpolation can be used to synthesize baselines whose signal parameters do not match those of the prerecorded sets [65, 66].

2.3.2.2 Baseline Signal Stretch

Another technique is *baseline signal stretch* (BSS). Since the primary effect of a temperature change is to stretch or compress the signal in time [60], BSS attempts to correct this by finding an optimal time-stretch parameter β that matches a baseline to the corresponding measurement. This transformation is often performed in the frequency domain. BSS can be combined with OBS in a two-step process, where OBS is first used to select the closest baseline set, and then BSS is used to fine-tune the signals [60]. Various improvements to this two-step process also exist; e.g., [65].

2.3.2.3 Load Differential Imaging

In lieu of a baseline, signals can be acquired when the structure is subject to different loads. For example, a tensile force on a structure will open cracks, increasing the magnitude of their reflectivity [67]; if measurement sets are recorded at two different loads in the presence of such a crack, the differential measurement will contain a scattered signal due to the change in the crack's scattering behavior. Load differential imaging can be performed at multiple loads, and can be used to estimate crack severity [63].

2.3.3 Lamb Wave Detection Methods

The simplest method of damage detection is baseline comparison, where follow-up measurements are compared to baselines from a known state. This comparison can

examine the energy in the residual, or the maximum of the cross-correlation, or measured modal parameters, and return a value that quantifies the difference. This value, often called a *damage index* [68], is then compared to a preselected threshold.

Various baseline-free detection methods also exist. Some techniques usually attempt to detect echoes that could be explained by damage, or sources of mode conversion that correspond to potential scatterers [69]. Other algorithms analyze modal properties [70], special transducer configurations [71], or detection of symmetry breaking in symmetrical structures. One method of baseline-free damage detection is the use of time reversal [72], which is a two step process: the transmitting element in a transducer pair first generates a Lamb wave with a known excitation; the transducers then switch roles, and the original receiving element transmits a time-reversed copy of the wavefield measurement that was previously recorded. If the structure behaves linearly, the signal received by the original transmitter will be identical to the original waveform due to the acoustic principle of reciprocity; differences are due to nonlinearities that could be indicative of damage. Since baseline-free techniques do not use prerecorded signals from a known state, they are usually applicable to structures that have simple shapes, such as beams and pipes, or have some type of symmetry.

2.3.4   Lamb Wave Localization

There are a variety of Lamb wave SHM damage localization methods, which are capable of determining the locations of any sites of damage (and can also be used for detection). The pioneering work by Wang, *et al.* [12] used a delay-and-sum (DAS) technique, which is sometimes called *synthetic time-reversal, point-focusing*, or *time-domain beamforming*, on differential signals; delay-and-sum methods are well-established in the radar community, and are commonly used in Lamb wave SHM due to their conceptual simplicity and computational efficiency. Other methods include compact phased array beamforming [10] as well as adaptive methods, including

minimum-variance imaging [73] (known as *Capon beamforming* in the radar commu-
nity), stochastic modeling [74], and multi-path deconvolution [75].

Lamb wave tomography is also commonly investigated, in which a region of interest
is considered to be a scalar field of some wave property such as slowness, attenuation,
or diffraction intensity [76]. An inverse problem is then solved to recover these values;
these often rely on parallel projection, diffraction methods, fan-beam geometry, or
iterative algorithms [77, 78, 79]. One popular algorithm, RAPID (reconstruction
algorithm for probabilistic inspection of damage), is not a true tomographic method,
but instead is an ad-hoc method that distributes signal differences over elliptical
regions [80].

A detailed description of DAS is provided, since it is foundation for the new work
shown in this thesis. Delay-and-sum imaging is performed over a discretized grid of
pixels and works by reverse-propagating signals along their hypothetical path from
source to scatterer to receiver and adding all such signals together. If a scatterer is
actually located at the pixel of interest, the back-propagated signals should add con-
structively; if no scatterer is present, the signals will not overlap or will destructively
interfere.

Let $v_0[t]$ be a windowed toneburst (e.g., with a Hann window), and let t_0 be the
time corresponding to the maximum of the window function. Calculate the differential
measurements $y_i[t]$ from the baseline and follow-up measurement sets. Assume that
there is a single scatterer at location $\mathbf{q}$ with an unknown scattering pattern.

Consider the pixel at X-Y location $\mathbf{p}$. If a scatterer is (hypothetically) present
at $\mathbf{p}$, the received scattered signal for pair i will have traveled a total distance of
$\|\mathbf{r}_i - \mathbf{p}\|_2 + \|\mathbf{p} - \mathbf{s}_i\|_2$. Back-propagating the signals through their assumed paths
results in the signals

$$y_{i,\mathbf{p}}^{\text{Back}}[t] := \mathcal{P}_{c_p}^{-1} \underset{\mathbf{s}_i \to \mathbf{p}}{} \mathcal{P}_{c_p}^{-1} \underset{\mathbf{p} \to \mathbf{r}_i}{} y_i[t]. \tag{11}$$

At this point, the signals are summed together. Since the scattering pattern is

unknown in this case, the signal envelopes are used instead:

$$y_{\mathbf{p}}^{\text{DAS}}[t] := \sum_{i=1}^{P} w_{i,\mathbf{p}} \left| y_{i,\mathbf{p}}^{\text{Back}}[t] + \mathrm{i}\mathcal{H}\left\{y_{i,\mathbf{p}}^{\text{Back}}\right\}[t] \right|, \tag{12}$$

where $\mathcal{H}\{\cdot\}$ represents the Hilbert transform and $w_{i,\mathbf{p}}$ is a weighting variable. Finally, the actual pixel value $z_{\mathbf{p}}^{\text{DAS}}$ is determined by using the time index where the maximum value is expected:

$$z_{\mathbf{p}}^{\text{DAS}} := y_{\mathbf{p}}^{\text{DAS}}[t_0]. \tag{13}$$

If $\mathbf{p}$ is the actual location of the scatterer (i.e., $\mathbf{p} = \mathbf{q}$), the inverse propagation from Eq. 11 will cancel the propagation in Eqs. 7-10:

$$y_{i,\mathbf{q}}^{\text{Back}}[t] = h_k\left[t; \theta_{\mathbf{s}_i,\mathbf{q}}, \theta_{\mathbf{q},\mathbf{r}_i}\right] \ast v_0[t]. \tag{14}$$

The result of Eq. 12 when $\mathbf{p} = \mathbf{q}$ is to sum P weighted copies of the original signal envelope, each subject to the scatterer's impulse response at the appropriate angle. If the scatterer behaves as expected (i.e., like a point scatterer), the pixel value $z_{\mathbf{p}}^{\text{DAS}}$ will be very high. In contrast, locations away from $\mathbf{q}$ will have smaller pixel values, because the back-propagated signals will not align.

As noted, this is only an example of one simple delay-and-sum method. DAS algorithms all share this general structure, but the specifics of the implementation can vary. For example, some versions use only the group velocity to reverse-propagate in the time domain, others integrate over some time window instead of simply taking the value at t_0, and various methods of weighting the signals exist.

CHAPTER III

BACKGROUND ON SPARSE RECONSTRUCTION

3.1 Introduction

A reasonable, but hitherto unconsidered, assumption for many operating structures is damage *sparsity*. For example, it is extremely unlikely that an aircraft has damage nearly everywhere while it is in flight; cracks or other flaws almost always appear individually and grow slowly over time, until one flaw undergoes a sudden rapid expansion that results in structural failure. While the structure is pre-failure, there may be several damage sites at material stress points or locations that received some sort of prior trauma, but most areas will necessarily be damage-free to permit operation. As posed here, the sparsity assumption is integral to a detection algorithm that takes a differential measurement as input and attempts to find a small set of locations that could contain damage consistent with that measurement. This can present several advantages over current methods; for example, results using DAS methods often have a spot size on the order of 75-100 mm (e.g., results in [43]) and easily get "overwhelmed" in the presence of multiple defects or geometrical boundaries.

3.2 Sparse Reconstruction and Compressed Sensing

The name *sparse reconstruction* refers to a class of techniques for solving linear inverse problems when it is known that the solution is sparse [81]. Consider the inverse problem $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{x}$ is the unknown vector. If $\mathbf{A}$ is square and full-rank, there is exactly one solution; if underdetermined, there are infinitely many solutions; and if overdetermined, there is no solution. The conventional least-squares answer to this problem is $\tilde{\mathbf{x}} = \mathbf{A}^{+}\mathbf{y}$, where $\mathbf{A}^{+}$ is the Moore-Penrose pseudoinverse. If $\mathbf{A}$ is

underdetermined ("short and fat"), this operation will select the solution $\tilde{\mathbf{x}}$ with the least energy; if overdetermined ("tall and skinny"), it will select the "closest" $\tilde{\mathbf{x}}$ using Euclidean distance as a metric.

In many cases, it is known that $\mathbf{x}$ has a sparse structure; i.e., that most components will be zero (or in the case of noisy measurements, very small). Especially if $\mathbf{A}$ is underdetermined, the least-squares method is wholly inappropriate and will return meaningless results [82]. In these situations, a sparse solver is preferred; in essence, such a method must be able to (1) select a limited number of components of $\tilde{\mathbf{x}}$ that should be assigned nonzero values, and (2) determine what those values are. A sparse reconstruction algorithm will therefore impose an additional constraint on the optimization problem; for example, it might be required that $\mathbf{x}$ contain no more than S nonzero entries.

Many "interesting" signals have some sort of structure and can be represented in a cleverly-selected basis for which they are sparse. One early example is the representation of seismic layers; in 1988 Santosa and Symes successfully determined underground impedance profiles by assuming a piecewise-constant representation (i.e., the derivative is sparse) [83]. The JPEG 2000 image standard [84] compresses images by representing them in the 2D wavelet domain, which is known to be mostly sparse for images such as photographs.

Another particularly common case where sparse reconstruction is appropriate is when signals can be represented with a redundant dictionary. In these situations, it is known that $\mathbf{y}$ is a linear combination of a limited number of vectors, or *atoms*, in this dictionary, but it is not known in advance which vectors will be present, nor their coefficients. Many problems can be represented this way, such as speech recognition (where the dictionary could be a collection of phonemes or Gabor functions [85]), array-based radar detection (where it might be a collection of steering vectors over the azimuth-elevation plane), and ground-penetrating radar (where dictionary

entries correspond to received signals [86]). The methods presented also use redundant dictionaries; here each dictionary entry will correspond to a potential scatterer at a different location over some area of interest. Sparse reconstruction has previously been applied to ultrasonic measurements in this context for denoising as well as mode separation of signals [87, 88].

The field of sparse reconstruction has received renewed attention due to the emerging field of compressed sensing (also called compressive sampling), which concerns signal acquisition using fewer samples than traditional Nyquist sampling. Under this paradigm, sampling is generalized to multiplication by an acquisition matrix $\mathbf{A}$; while Nyquist sampling can be represented using the identity matrix, it is possible to take fewer measurements (i.e., reduce the number of rows in $\mathbf{A}$) by using diffuse sampling methods. The seminal research by Candès, Romberg, Tao, and Donoho derived conditions under which exact recovery of a sparse signal can be guaranteed, including the structure of $\mathbf{A}$, the number of rows (samples) it must contain, and how to use sparse reconstruction to perform the recovery [89, 90, 81]. The fundamental property that $\mathbf{A}$ must satisfy is called the *restricted isometry property* (RIP) [18]. The restricted isometry constant $\delta_S(\mathbf{A})$ for S-sparse vectors is the minimum value that satisfies the inequality

$$(1 - \delta_S(\mathbf{A})) \|\mathbf{x}\|^2 \leq \|\mathbf{A}\mathbf{x}\|^2 \leq (1 + \delta_S(\mathbf{A})) \|\mathbf{x}\|^2 \quad \text{for all } S\text{-sparse } \mathbf{x}. \quad (15)$$

$\delta_S(\mathbf{A})$ is, roughly speaking, a measure of energy preservation when $\mathbf{A}$ is applied to sparse vectors. A vector is considered S-sparse if it has S or fewer nonzero components.

One issue with the RIP is that the constant $\delta_S(\mathbf{A})$ is itself NP-hard to compute. A fundamental insight in compressed sensing is that random matrices (for example, matrices where each element is a zero-mean Gaussian random variable) satisfy the RIP with overwhelming probability, even when they are greatly underdetermined. Specifically, if $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a random matrix, where $M \in \mathcal{O}(S \log(N/s))$, then with

25

overwhelming probability, every S-sparse $\mathbf{x} \in \mathbb{R}^N$ can be recovered exactly by solving an ℓ_1-norm optimization problem. When $\mathbf{A}$ is deterministic, it is more practical to instead calculate the *coherence* $\mu\left(\mathbf{A}\right) = \max\limits_{i \neq j} \left| \frac{\mathbf{a}_i^{\mathsf{H}} \mathbf{a}_j}{\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2} \right|$. Coherence is not as rigorous as the RIP, but can still be used to estimate reconstruction performance. Reconstruction is more difficult with a highly-coherent dictionary because it contains columns that are similar.

3.3 Algorithms for Sparse Reconstruction and Compressed Sensing

Regardless of its computability, if the RIP constant $\delta_{2S}\left(\mathbf{A}\right)$ is less than 1, such as in Gaussian and Bernoulli random matrices, then the optimization problem

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_0 \ \text{ subject to } \mathbf{y} = \mathbf{Ax}, \tag{16}$$

will exactly recover $\mathbf{x}$ if it is S-sparse [18]; here the ℓ_0 pseudo-norm is equal to the number of nonzero entries in $\mathbf{x}$. Equation 16 is NP-hard and can only be solved by "brute force" methods, and is therefore infeasible for all but the most trivial cases. As a result, several algorithms exist that approximate Eq. 16. Note that the RIP must be satisfied for vectors that are have twice as many nonzero entries as $\mathbf{x}$, since the difference between two S-sparse vectors is $2S$-sparse.

It should be additionally noted that exact recovery is not always necessary. In many engineering problems, there is some tolerance that is allowed or expected. For example, in the damage detection problem, it might be acceptable for a scatterer's reported position to be "off" by 10-20 mm or for the solver to incorrectly detect its amplitude. In these problems, even a matrix that does not satisfy the RIP can perform extremely well.

3.3.1 Basis Pursuit

As previously stated, if the RIP-$2S$ constant is less than one, the sparse vector $\mathbf{x}$ is the unique solution to Eq. 16, which can only be solved by exhaustive search. However,

under only slightly more strict RIP conditions [18], the solution to this problem is exactly equal that of

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \ \text{subject to} \ \mathbf{y} = \mathbf{A}\mathbf{x}, \tag{17}$$

where $\|\mathbf{x}\|_1 = |x_1| + |x_2| + \cdots + |x_N|$. This well-known sparse recovery problem, known as basis pursuit (BP) [91], can be solved with a linear program in polynomial time.

In most cases, there is some amount of noise in $\mathbf{y}$; this is typically modeled as additive noise of the form $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$. There are several convex optimization problems related to Eq. 17 that are designed to handle noise:

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \ \text{subject to} \ \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \sigma \tag{18}$$

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 + \lambda \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \tag{19}$$

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \ \text{subject to} \ \|\mathbf{x}\|_1 \leq \tau \tag{20}$$

The three optimization problems above are equivalent under certain conditions. In particular, for a given noise level σ, there exist λ and τ for which Eqs. 18, 19, and 20 will yield the same result, but the relation between the three parameters is unknown in general, unless the problem is already solved. The names for these three problems are often used interchangeably; Eq. 18 is often called basis pursuit denoising (BPDN), though in the original description by Chen, *et al.*, BPDN refers to Eq. 19 [91]. The term LASSO (least absolute shrinkage and selection operator) was originally used by Tibshirani to describe Eq. 20 [92], but now often refers to Eq. 19 in most usage instead. In this document, BPDN refers to Eq. 18. The other two optimization problems are not considered.

If the σ parameter is chosen such that $\|\mathbf{e}\|_2 \leq \sigma$, then the error $\|\tilde{\mathbf{x}} - \mathbf{x}\|_2$ is bounded; this bound is a function of σ, $\delta_{4S}(\mathbf{A})$, and the sparsity of $\tilde{\mathbf{x}}$ [93].

3.3.2 Matching Pursuit

While BP and BPDN have favorable performance guarantees, their execution time can be slow on some problems. A second class of sparse reconstruction algorithms called iterative greedy methods is much more conceptually simple, but the algorithms lack the performance guarantees of BP-based methods. Iterative greedy algorithms select appropriate columns of $\mathbf{A}$ one-at-a-time until some stopping criterion is met, so that a small subset of atoms has been selected. The most well-known such algorithm is matching pursuit (MP) [94], which repeatedly selects the column of $\mathbf{A}$ most correlated with $\mathbf{y}$ and subtracts a scaled version, removing that component of the signal. An improved version, orthogonal matching pursuit (OMP), uses orthogonal projections to improve the convergence rate of MP [95]. Many other modifications exist; for example, the well-known CoSaMP (compressive sampling matching pursuit) algorithm [96] is tailored specifically for recovering sparse vectors in compressed sensing problems.

Algorithm 1 Matching Pursuit

Input:	dictionary matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$, measurement $\mathbf{y} \in \mathbb{C}^m$, stopping criterion
Output:	sparse estimate $\tilde{\mathbf{x}} \in \mathbb{C}^n$

Algorithm:

$\quad \mathbf{r}_0 \leftarrow \mathbf{y}$
$\quad \tilde{\mathbf{x}} \leftarrow \mathbf{0}$
$\quad n \leftarrow 0$
$\quad$ **while** stopping criterion is not satisfied **do**
$\quad\quad n \leftarrow n + 1$
$\quad\quad i_{(n)} \leftarrow \underset{i \notin \{i_{(1)}, \ldots, i_{(n-1)}\}}{\arg\max} \left| \mathbf{a}_i^{\mathsf{H}} \mathbf{r}_{(n-1)} \right|$
$\quad\quad \tilde{x}_{i_{(n)}} \leftarrow \mathbf{a}_{i_{(n)}}^{\mathsf{H}} \mathbf{r}_{(n-1)}$
$\quad\quad \mathbf{r}_{(n)} \leftarrow \mathbf{r}_{(n-1)} - \tilde{x}_{i_{(n)}} \mathbf{a}_{i_{(n)}}$
$\quad$ **end while**
$\quad$ **return** $\tilde{\mathbf{x}}$

Matching pursuit is shown in Algorithm 1 and is extremely easy to understand. It solves a sparse problem the way many people might try to solve such a problem

by hand; it finds and removes the biggest signal component, then repeats for the next biggest signal component, and continues until the stopping condition is met. Common stopping criteria include an energy bound on the residual vector $\mathbf{r}_n$ or a fixed number of iterations.

Algorithm 2 Orthogonal Matching Pursuit

Input: dictionary matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$, measurement $\mathbf{y} \in \mathbb{C}^m$, stopping criterion
Output: sparse estimate $\tilde{\mathbf{x}} \in \mathbb{C}^n$

Algorithm:

$\quad \mathbf{r}_0 \leftarrow \mathbf{y}$
$\quad \tilde{\mathbf{x}} \leftarrow \mathbf{0}$
$\quad n \leftarrow 0$
$\quad$ **while** stopping criterion is not satisfied **do**
$\quad\quad n \leftarrow n + 1$
$\quad\quad i_{(n)} \leftarrow \arg\max_i \left| \mathbf{a}_i^{\mathsf{H}} \mathbf{r}_{(n-1)} \right|$
$\quad\quad \mathbf{S}_{(n)} \leftarrow \begin{bmatrix} \mathbf{a}_{i_{(1)}} & \mathbf{a}_{i_{(2)}} & \cdots & \mathbf{a}_{i_{(n)}} \end{bmatrix}$
$\quad\quad \mathbf{r}_{(n)} \leftarrow \left(\mathbf{I} - \mathbf{P}_{\mathbf{S}_{(n)}} \right) \mathbf{y}$
$\quad$ **end while**
$\quad \tilde{\mathbf{x}}^s \leftarrow \mathbf{S}_{(n)}^{+} \mathbf{y}$
$\quad$ **for** k from 1 to n **do**
$\quad\quad \tilde{x}_{i_{(k)}} \leftarrow \tilde{x}_{(k)}^s$
$\quad$ **end for**
$\quad$ **return** $\tilde{\mathbf{x}}$

Orthogonal matching pursuit, shown in Algorithm 2, uses a similar principle to MP, but it keeps a running list of all selected columns and updates the residual vector every iteration via projection. Here, one column from the dictionary is picked every iteration and appended to the "selected atoms" matrix $\mathbf{S}_{(n)}$. The measurement vector is projected onto the left null space of this matrix to get the new residual, which is orthogonal to $\mathbf{S}_{(n)}$'s column space. (The matrix $\mathbf{P}_{\mathbf{A}}$ denotes the projection matrix onto $\mathbf{A}$'s column space; i.e., $\mathbf{P}_{\mathbf{A}} := \mathbf{A}\mathbf{A}^{+}$.) When the stopping criterion is met, the corresponding entries of $\tilde{\mathbf{x}}$ are computed using the Moore-Penrose pseudoinverse.

Since matching pursuit algorithms work one column at a time, they are particularly bad at handling highly-coherent dictionaries. In these cases it is not uncommon for the algorithm to pick the "wrong" dictionary column on the first iteration because the available atoms are so similar. Depending on the problem and the degree of coherence, this may or may not be acceptable (for example, highly-correlated columns represent similar damage locations for this research). Basis pursuit methods are not immune to this problem, but they at least find a global solution to the optimization problem and come with some performance guarantees; they also tend to be much more robust to these situations.

3.3.3 Other Methods

There are also various less well-known or more specialized sparse solvers. One recent example are approximate message passing solvers [97]. These new methods apply a modified form of message passing that uses iterative soft thresholding to converge on a sparse solution. These techniques show great promise, but are not considered here, as BPDN and OMP worked sufficiently well. Other solvers use variational methods [98], and even more esoteric solvers use statistical or matrix completion methods. An extensive collection of solvers and other references is available online at the *sparse-and low-rank approximation wiki* [99].

3.4 Sparse Recovery of a Signal in a Union of Subspaces

One way to view sparse reconstruction is that it tries to find the smallest possible subspace to explain the measured signal; for methods that allow noise, the algorithm attempts to find the smallest possible subspace that explains some portion of the signal. In this sense, the dictionary matrix $\mathbf{A}$ is an overcomplete basis from which the smaller subspace is selected, and the nonzero entries of $\tilde{\mathbf{x}}$ correspond to the (non-orthogonal) basis vectors of the selected subspace. Sometimes, however, the signal does not solely lie in a single subspace, but a union of subspaces, with each

subspace corresponding to some sort of phenomenon. In other words, while "standard" sparse reconstruction selects vectors independently, it might make sense to partition the vectors in $\mathbf{A}$ into blocks and select the blocks independently instead of their constituents. For the remainder of this document, any use of unqualified term "sparse reconstruction" refers exclusively to the standard case without any block structure.

Examined in the more conventional sense, a union-of-subspaces problem is the result of a block structured problem. In this sort of problem, the columns of $\mathbf{A}$ and corresponding entries of $\mathbf{x}$ fall into natural "blocks." The atoms within each block do not exist in isolation; the entire block is collectively either "on" or "off," though each atom still has its own, independently-determined coefficient if its block is selected. If a block is not selected, all of its coefficients are set to zero. These problems are said to exhibit a property known as *block sparsity* or *group sparsity* [100]. A vector is said to be S-block-sparse if the total number of blocks with at least one nonzero entry is S or less. For these problems, such a dictionary matrix will be denoted $\hat{\mathbf{A}}$.

Suppose the dictionary matrix $\hat{\mathbf{A}}$ is divided into groups (blocks) by assigning each column n a *group index*, denoted $\mathcal{G}(n)$. Two dictionary columns n_1 and n_2 are in the same group if $\mathcal{G}(n_1) = \mathcal{G}(n_2)$. The group indices themselves are denoted $G_1, G_2, \ldots, G_\Gamma$. Since $\mathbf{x}$ shares the same group structure as $\hat{\mathbf{A}}$, the grouping $\mathcal{G}$ applies to it as well. As a slight abuse of notation, $\mathcal{G}$ will also be used as the set of all possible groups in statements (e.g., "for each $G \in \mathcal{G}$..." is equivalent to "for each $G \in \{G_1, G_2, \ldots, G_\Gamma\}$...").

It is often convenient to describe only the submatrix of $\hat{\mathbf{A}}$ or the subvector of $\mathbf{x}$ that corresponds to a group. Let $\mathbf{x}_{G_i}$, $G_i \in \mathcal{G}$ denote a subvector of $\mathbf{x}$ containing only the entries in group G_i, and let the matrix $\hat{\mathbf{A}}_{G_i}$ contain the columns of $\hat{\mathbf{A}}$ that correspond to that group. Table 1 shows an example of an arbitrary block structure.

The block-sparse problem is a generalization of its non-block counterpart and has many similar properties. There exists a block-RIP constant [100] which measures

Table 1: An example of group subvectors.

$\mathbf{x}$	group index	$\mathbf{x}_{G_1}$	$\mathbf{x}_{G_2}$
$\begin{bmatrix} 2 \\ 4 \\ 6 \\ 8 \\ 10 \end{bmatrix}$	1 1 2 1 2	$\begin{bmatrix} 2 \\ 4 \\ 8 \end{bmatrix}$	$\begin{bmatrix} 6 \\ 10 \end{bmatrix}$

energy preservation when restricted to block-sparse vectors; if the block-RIP constant for $2S$-block-sparse vectors is less than one, the original vector $\mathbf{x}$ can be uniquely recovered by an exhaustive search. Given $\hat{\mathbf{A}}$, $\mathbf{y}$, and $\mathcal{G}$, if $\mathbf{x}$ is $2S$-block-sparse and the block-RIP-2S constant $\delta_{2S|\mathcal{G}}\left(\hat{\mathbf{A}}\right) < 1$, then the optimization problem

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \sum_{G \in \mathcal{G}} \mathbb{I}\left(\|\mathbf{x}_G\|_2 > 0\right) \text{ subject to } \mathbf{y} = \hat{\mathbf{A}}\mathbf{x} \tag{21}$$

will exactly recover $\mathbf{x} = \tilde{\mathbf{x}}$; $\mathbb{I}(\cdot)$ represents the indicator function. Equation 21 uses a mixed ℓ_2/ℓ_0 norm, though it may not be immediately apparent due to the use of the indicator function; the equation can also be written as

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \left\| \left[\|\mathbf{x}_{G_1}\|_2 \quad \|\mathbf{x}_{G_2}\|_2 \quad \cdots \quad \|\mathbf{x}_{G_\Gamma}\|_2\right]^{\mathsf{T}} \right\|_0 \text{ subject to } \mathbf{y} = \hat{\mathbf{A}}\mathbf{x}. \tag{22}$$

In this form, it is more clear that block sparsity changes the problem to a norm-of-norms optimization. Regardless of how it is written, all Eqs. 21 and 22 do is minimize the number of active blocks. Once again, this is computationally infeasible for all but the smallest problems.

3.4.1 Block-Sparse Basis Pursuit

As with the non-block case, an argument can be made to use the ℓ_1 norm in place of the ℓ_0 norm. If $\mathbf{x}$ is S-block sparse and $\delta_{2S|\mathcal{G}}\left(\hat{\mathbf{A}}\right) < \sqrt{2} - 1$, then the solution to the ℓ_2/ℓ_1 optimization problem

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \sum_{G \in \mathcal{G}} \|\mathbf{x}_G\|_2 \text{ subject to } \mathbf{y} = \hat{\mathbf{A}}\mathbf{x} \tag{23}$$

is equal to $\mathbf{x}$ [100, 101]. In addition to the block-RIP constant, the matrix $\hat{\mathbf{A}}$ has measures of *block coherence*, which is a generalization of coherence that applies to subspaces, and *sub-coherence*, which is the maximum coherence within any one block [101] which can be used to determine if Eq. 23 can achieve perfect reconstruction. This condition, as well as the block-RIP condition, is sufficient (but not necessary) to guarantee exact recovery.

There is also a block-sparse BPDN problem

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_{G \in \mathcal{G}} \|\mathbf{x}_G\|_2 \ \text{ subject to } \ \left\| \mathbf{y} - \hat{\mathbf{A}}\mathbf{x} \right\|_2 \leq \sigma \tag{24}$$

that allows for noise.

3.4.2 Block-Sparse Matching Pursuit

In addition to block-sparse basis pursuit, there are various modifications of matching pursuit methods that are adapted to handle problems with block structure. Algorithm 3 shows an implementation of block OMP, which functions nearly identically to OMP, except it picks one block at a time instead of one vector at a time.

3.4.3 Block-Sparse Example: Touch-Tone Phone

3.4.3.1 Touch-Tone Phone Operation

Here is a simple example to demonstrate a block-sparse problem. Touch-tone phones use a standardized dual-tone multi-frequency signal circuit [102] to uniquely identify which key is pressed. Each button on the phone produces two different tones according to Table 2. The specifications also state that the actual frequencies transmitted must be within 1.8% of the nominal frequencies and that the ratio of the tones to any distortion must be at least 20 dB.

3.4.3.2 Detecting Touch Tones with Block-Sparse Reconstruction

Touch tones are commonly decoded with filterbanks and combinatorial logic, however in this example a different approach will be used. Let $y(t)$ be a received touch-tone

---

Algorithm 3 Block Orthogonal Matching Pursuit

Input: dictionary matrix $\hat{\mathbf{A}} \in \mathbb{C}^{m \times n}$, measurement $\mathbf{y} \in \mathbb{C}^m$, group assignment $\mathcal{G}$, stopping criterion

Output: sparse estimate $\tilde{\mathbf{x}} \in \mathbb{C}^n$

Algorithm:

$\quad \mathbf{r}_0 \leftarrow \mathbf{y}$
$\quad \tilde{\mathbf{x}} \leftarrow \mathbf{0}$
$\quad n \leftarrow 0$
$\quad$ **while** stopping criterion is not satisfied **do**
$\quad\quad n \leftarrow n + 1$
$\quad\quad G_{(n)} \leftarrow \underset{G \in \mathcal{G}}{\arg\max} \left\| \mathbf{P}_{\hat{\mathbf{A}}_G} \mathbf{r}_{(n-1)} \right\|_2$
$\quad\quad \mathbf{S}_{(n)} \leftarrow \begin{bmatrix} \hat{\mathbf{A}}_{G_{(1)}} & \hat{\mathbf{A}}_{G_{(2)}} & \cdots & \hat{\mathbf{A}}_{G_{(n)}} \end{bmatrix}$ (preserving block structure)
$\quad\quad \mathbf{r}_{(n)} \leftarrow \left(\mathbf{I} - \mathbf{P}_{\mathbf{S}_{(n)}} \right) \mathbf{y}$
$\quad$ **end while**
$\quad \tilde{\mathbf{x}}^s \leftarrow \mathbf{S}_{(n)}^+ \mathbf{y}$ (block structure of $\tilde{\mathbf{x}}^s$ should match $\mathbf{S}_{(n)}$)
$\quad$ **for** k from 1 to n **do**
$\quad\quad \tilde{\mathbf{x}}_{G_{(k)}} \leftarrow \tilde{\mathbf{x}}^s_{G_{(k)}}$
$\quad$ **end for**
$\quad$ **return** $\tilde{\mathbf{x}}$

---

signal from a single button, and denote its analytic representation $\hat{y}(t) = y(t) + i\mathcal{H}\{y\}(t)$. Ignoring noise, $\hat{y}(t)$ should lie in a two-dimensional subspace, with its upper and lower tones as its basis functions. For example, if the "1" button is pressed, then under ideal circumstances $\hat{y}(t) = A \exp(i2\pi(697)t) + B \exp(i2\pi(1209)t)$, where A and B are complex. In fact, $\hat{y}$ could lie in one of twelve two-dimensional subspaces – one for each button.

Table 2: International standards for touch-tone frequencies

		Upper Tone (Hz)		
		1209	1336	1447
Lower Tone (Hz)	697	1	2	3
	770	4	5	6
	852	7	8	9
	941	$*$	0	$\#$

This problem could be solved using (non-block) sparse reconstruction by constructing a dictionary matrix for the discretized vector $\hat{\mathbf{y}} \in \mathbb{C}^n$ (presumed to be sampled at sufficiently high frequency). Let $\mathbf{F}_{f_0} \in \mathbb{C}^n$ denote a complex vector with positive frequency component f_0 that is normalized to unit ℓ_2 norm. Then the dictionary matrix is:

$$\mathbf{A} = \begin{bmatrix} \mathbf{F}_{697} & \mathbf{F}_{770} & \mathbf{F}_{852} & \mathbf{F}_{941} & \mathbf{F}_{1209} & \mathbf{F}_{1336} & \mathbf{F}_{1447} \end{bmatrix}$$

Solving Eq. 17, or Eq. 19 in the presence of noise, will recover the frequency components in $\hat{y}(t)$. Of course, this isn't revolutionary – as n approaches infinity, the coherence of the dictionary approaches zero; for any reasonable sample length, $\tilde{\mathbf{x}} = \mathbf{A}^{\mathsf{H}}\hat{\mathbf{y}}$ recovers the components as well. After computing $\tilde{\mathbf{x}}$ with either method, its components are quantized to "0" or "1" and used to calculate the button pressed.

Another way to solve this problem is to directly find which of the twelve possible subspaces contains $\hat{\mathbf{y}}$, using a block dictionary:

$$\hat{\mathbf{A}} = \begin{bmatrix} \overbrace{\mathbf{F}_{697} \ \mathbf{F}_{1209}}^{1} & \overbrace{\mathbf{F}_{697} \ \mathbf{F}_{1336}}^{2} & \overbrace{\mathbf{F}_{697} \ \mathbf{F}_{1447}}^{3} & \overbrace{\mathbf{F}_{770} \ \mathbf{F}_{1209}}^{4} & \overbrace{\mathbf{F}_{770} \ \mathbf{F}_{1336}}^{5} & \overbrace{\mathbf{F}_{770} \ \mathbf{F}_{1447}}^{6} & \overbrace{\mathbf{F}_{852} \ \mathbf{F}_{1209}}^{7} & \overbrace{\mathbf{F}_{852} \ \mathbf{F}_{1336}}^{8} & \overbrace{\mathbf{F}_{852} \ \mathbf{F}_{1447}}^{9} & \overbrace{\mathbf{F}_{941} \ \mathbf{F}_{1209}}^{*} & \overbrace{\mathbf{F}_{941} \ \mathbf{F}_{1336}}^{0} & \overbrace{\mathbf{F}_{941} \ \mathbf{F}_{1447}}^{\#} \end{bmatrix}$$

This matrix has a mutual coherence of 1, because it contains duplicated columns. Since only one phone button is pressed at a time, the solution vector $\mathbf{x}$ to the equation $\hat{\mathbf{y}} = \hat{\mathbf{A}}\mathbf{x}$ should be 1-block-sparse. It can be easily demonstrated that this matrix does not satisfy the block-RIP by showing that $\delta_{2|\mathcal{G}}\left(\hat{\mathbf{A}}\right) \geq 1$; for example, the vector $\mathbf{c} = \begin{bmatrix} 1 \ 0 \ -1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \end{bmatrix}^{\mathsf{T}}$ is 2-block-sparse but lies in $\hat{\mathbf{A}}$'s null space. ($\hat{\mathbf{A}}$'s block coherence and sub-coherence are not low enough to guarantee exact recovery, either.) Regardless, block-BP (or, in the presence of noise, block-BPDN) can be used to determine the button pressed.

To show this, a touch-tone signal was generated in MATLAB for the "7" button with the parameters shown in Table 3. The noise component of the signal is additive white Gaussian noise. Figure 1 shows the generated signal.

Table 3: Touch-tone signal parameters for block-sparse example problem

Parameters	Value
Sampling Frequency	50 kHz
Signal Duration	10 ms
Upper Tone Frequency	1209 Hz
Upper Tone Amplitude	1
Upper Tone Phase	$-32°$
Lower Tone Frequency	852 Hz
Lower Tone Amplitude	1.3
Lower Tone Phase	$95°$
Signal-to-Noise Ratio	10 dB

After computing the analytic signal $\hat{\mathbf{y}}$, it is possible to recover the components of each button by solving Eq. 23 with $\mathbf{y} \leftarrow \hat{\mathbf{y}}$ and with each pair of frequencies belonging to the group that represents its button. Version 1.8 of the freely-available SPGL1 package for MATLAB [103] was used to perform the optimization, with $\sigma = 0.1 \|\hat{\mathbf{y}}\|$. Results are shown in Figure 2; even though the $\hat{\mathbf{A}}$ matrix is coherent, the coefficients are correctly assigned to block 7 due to the block structure of the problem.

3.4.3.3   Coefficient Denormalization

The coefficients in Figure 2 have been "denormalized" to account for the the normalization of $\hat{\mathbf{A}}$; i.e., the columns can be expressed as $\mathbf{F}_{f_0}[t] = \frac{1}{\alpha} \exp{(i2\pi f_0 t)}$, with the coefficient α used to normalize the vector. A matching signal component $\mathbf{y}_{f_0}[t] = A \exp{(i2\pi f_0 t)}$ will then be assigned a coefficient of αA, so the coefficient can be appropriately scaled by dividing by α.

3.4.3.4   Discussion

One advantage to the block-sparse method is that the results directly reflect the structure of the problem. The block with the highest norm, $\|\tilde{\mathbf{x}}_g\|_2, g \in \mathcal{G}$, is the button that is pressed. For a simple problem like this, such structure is not of large

benefit; however, for much larger problems with more complicated structures, block-sparse reconstruction allows the use of multidimensional models that directly correlate to the unknowns of the problem.

(An astute reader might realize that a simpler way to solve this problem would perhaps be to project the signal onto each of the twelve subspaces and measure which contained the largest signal energy. Such an operation is equivalent to running block-OMP for a single iteration.)
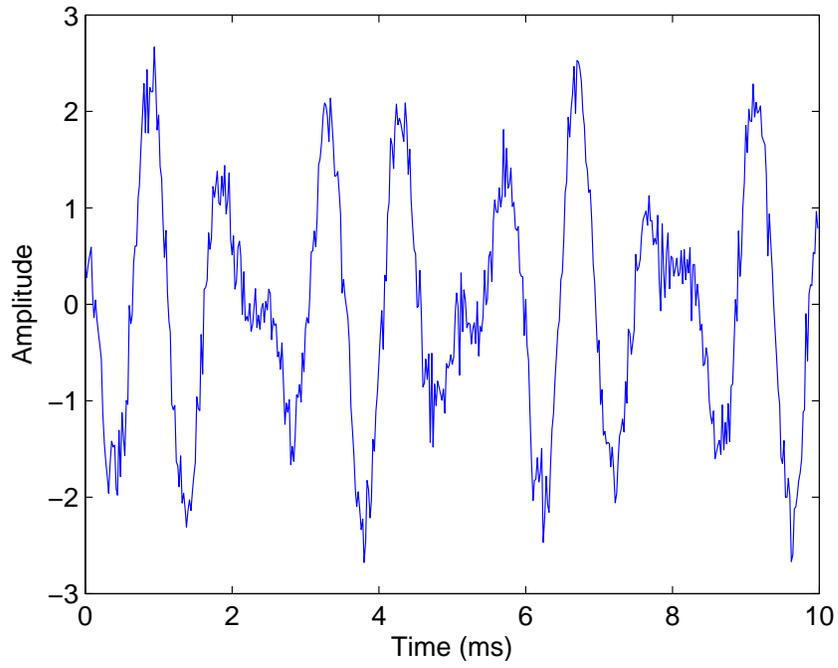
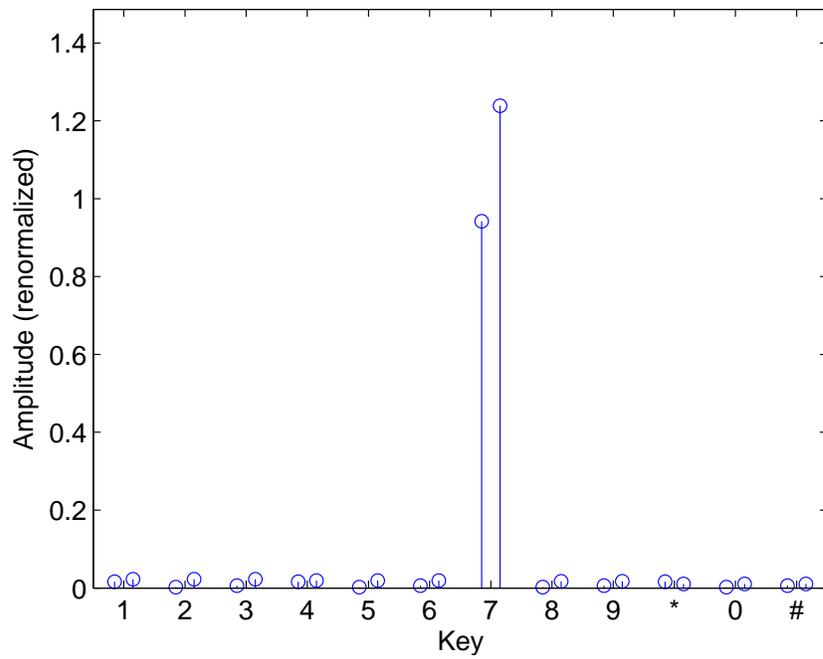Figure 1: Touch-tone phone signal generated in MATLAB.

Figure 2: Block-sparse reconstruction of touch-tone signal components. The signal components are grouped by block, with the lower frequency on the left and the upper frequency on the right.

CHAPTER IV

A SPARSE RECONSTRUCTION ALGORITHM FOR DAMAGE LOCALIZATION

4.1 Introduction

This chapter describes the implementation of a (non-block) sparse reconstruction method to detect damage, given a set of residual signals $\mathbf{y}_i$. Initial development of the method was presented in [104], with further analysis in [105].

To use sparse reconstruction methods, the damage detection problem must be formulated so that it conforms to the linear equation $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}$. The approach taken in this research is to use a redundant dictionary with columns that correspond to locations in the region of interest (ROI).

The use of sparse reconstruction to solve the problem of interest is presented as follows. First, a method is presented that relies on usage of a known scattering pattern. Next, strategies are described and analyzed for the more realistic scenario where scattering patterns are unknown in advance. An in-depth examination of concepts and considerations regarding the sparse solver follows; this section is tailored to BPDN (the recommended algorithm), but the analysis may apply to other methods as well. Finally, a discussion section contains various notes on the algorithm, including dictionary coherence and computational efficiency.

4.2 Sparse Reconstruction for Known Scattering Patterns

Initially, let us assume that only point scatterers occur in the region of interest, all sharing the same scattering pattern. This scattering pattern is known *a priori* and is denoted $H\left[f; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right]$, as defined in Section 2.2.4; its inverse Fourier transform is

denoted $h\left[t; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right]$. Moreover, environmental effects such as temperature changes are ignored; in particular, this means that the dispersion curves are assumed to match the nominal curves computed based on the material properties. Finally, the assumption of single-mode propagation remains in effect.

First, the ROI is discretized into N pixels, which are typically arranged as a 2D rectangular grid. The grid spacing has an effect on the reconstruction algorithm; a coarse spacing results in faster execution time and a less coherent dictionary, but may lead to pixel straddle effects if scatterers fall between pixels; a fine spacing increases execution time and increases the dictionary coherence. Let $\{\mathbf{p}_n\}_{n=1}^{N}$ be the set of such pixels, where $\mathbf{p}_n \in \mathbb{R}^2$. It is assumed that a potential scatterer is present at each pixel, with an unknown intensity coefficient. Let the vector $\mathbf{x} \in \mathbb{R}^N$ represent the intensity of the corresponding scatterer at $\mathbf{p}_n$; if $x_n = 0$, no damage is present, and if $x_n \neq 0$, damage is present with a reflectivity proportional to $|x_n|$.

4.2.1 Model for a Single Transducer Pair

The dictionary matrix $\mathbf{A}$ is constructed from several submatrices; one submatrix corresponds to each transducer pair. The vector $\mathbf{y}$ is constructed in the same way. Consider a single pair i and the corresponding differential measurement $\mathbf{y}_i \in \mathbb{R}^L$, where the source transducer is at position $\mathbf{s}_i$ and the receiver is at $\mathbf{r}_i$. Assume the presence of K scatterers that each lie exactly on one of the pixels; denote the pixel index of the kth scatterer n_k; In other words, the kth scatterer is located at $\mathbf{p}_{n_k}$. From Eq. 10, the expression for the differential signal $\mathbf{y}_i$ is:

$$
\begin{aligned}
\mathbf{y}_i &= \sum_{n=1}^{N} x_n \left\{ \mathbf{h}_n^i \ast \mathop{\mathcal{P}_{c_p}}_{\mathbf{p}_n \to \mathbf{r}_i} \mathop{\mathcal{P}_{c_p}}_{\mathbf{s}_i \to \mathbf{p}_n} \mathbf{v}_0 \right\} + \mathbf{e}_i \\
&= \sum_{k=1}^{K} x_{n_k} \left\{ \mathbf{h}_{n_k}^i \ast \mathop{\mathcal{P}_{c_p}}_{\mathbf{p}_{n_k} \to \mathbf{r}_i} \mathop{\mathcal{P}_{c_p}}_{\mathbf{s}_i \to \mathbf{p}_{n_k}} \mathbf{v}_0 \right\} + \mathbf{e}_i,
\end{aligned}
\tag{25}
$$

where $\mathbf{h}_n^i = h\left[t; \theta_{\mathbf{s}_i, \mathbf{p}_n}, \theta_{\mathbf{p}_n, \mathbf{r}_i}\right]$. It is possible to write this equation in matrix form as follows:

$$\mathbf{y}_i = \mathbf{A}_i \mathbf{x} + \mathbf{e}_i, \tag{26}$$

where $\mathbf{A}_i \in \mathbb{R}^{L \times N}$, with individual columns $\mathbf{a}_{i,n} = \mathbf{h}_n^i \ * \ \underset{\mathbf{p}_n \to \mathbf{r}_i}{\mathcal{P}_{c_p}} \ \underset{\mathbf{s}_i \to \mathbf{p}_n}{\mathcal{P}_{c_p}} \ \mathbf{v}_0$. (The matrix $\mathbf{A}_i$ is intentionally unnormalized.) Note that $\mathbf{a}_{i,n}$ can also be interpreted as the vector representation of $v^{\text{scattered}}\left[t; \mathbf{s}_i, \mathbf{r}_i, \mathbf{p}_n, H\right]$ from Eq. 6. In other words, each column of the dictionary is generated by simulating damage at the corresponding pixel. Here a nonzero x_n corresponds to the intensity of the scattered (residual) signal at location $\mathbf{p}_n$. Using this formulation, the dictionary submatrix $\mathbf{A}_i$ is highly coherent, since $\mathbf{a}_{i,n}$ is determined solely by the propagation distance $\|\mathbf{s}_i - \mathbf{p}_n\|_2 \ + \ \|\mathbf{p}_n - \mathbf{r}_i\|_2$, which takes equal values on elliptical contours with foci at $\mathbf{s}_i$ and $\mathbf{r}_i$; however Eq. 26 represents only a single transducer pair.

4.2.2 Model for all Transducer Pairs

Since there are P different pairs, all of which are linear in $\mathbf{x}$, it is possible to concatenate the system of equations to obtain

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} \ \text{ where } \ \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_P \end{bmatrix} \ \text{ and } \ \mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_P \end{bmatrix} \mathbf{D}^{-1}, \tag{27}$$

where $\mathbf{D}$ is a diagonal matrix used to normalize the columns of the dictionary matrix with diagonal elements $D_{n,n} = \left(\sum_{i=1}^{P} \|\mathbf{a}_{i,n}\|_2^2 \right)^{1/2}$. The matrix $\mathbf{A}$ has lower coherence than each of its submatrices, although its coherence can still be high, depending on the pixel grid spacing and the number of transducers as well as their locations. However, the correlation between individual columns of $\mathbf{A}$ strongly depends on the proximity of their corresponding pixels. Consider a scatterer located at $\mathbf{p}$ that results in the (concatenated) differential signals $\mathbf{y}_{\mathbf{p}}$. If the location of this scatterer is perturbed to

$\mathbf{p} + \Delta\mathbf{p}$, each constituent subvector in $\mathbf{y}_{\mathbf{p}+\Delta\mathbf{p}}$ will be slightly shifted in time either earlier (if the resulting source $\rightarrow$ scatterer $\rightarrow$ receiver path has been shortened) or later (if the path has been lengthened). Depending on the excitation function, propagation velocity, and pixel grid spacing, this time shift can be quite small, making columns of $\mathbf{A}$ that correspond to adjacent pixels highly correlated. This high coherence is normally a barrier to perfect reconstruction: since columns are highly correlated, an algorithm will have a hard time "knowing" which columns to select if many of them are similar; however, for this application, so-called "perfect" reconstruction is not necessary. Even though some columns of $\mathbf{A}$ are highly correlated, these columns almost always represent pixels that are very close to each other. In most cases, the results are acceptable even if the damage sites reported by the algorithm are off by a small distance. A dictionary with many highly-correlated columns that do not represent nearby pixels is often an indication of poor array design; for example, it could indicate that there are too few transducers.

As previously stated, the pixel density chosen for discretization has several effects on the reconstruction problem. In addition to increasing the coherence of the dictionary, a more-finely sampled grid will result in more columns in $\mathbf{A}$, which increases execution time and memory usage. Greedy algorithms like OMP have a linear time complexity with respect to the number of columns in the dictionary [106], but BP solvers have a higher (but still polynomial) time complexity – although, in practice their execution time can be comparable (e.g., [107]). On the other hand, having too coarse a pixel density will increase straddle effects, which results in a multiple pixels with decreased amplitudes if a scatterer lies in a "gap" between pixels.

4.2.3   Summary

For the case with a known scattering pattern, each submatrix $\mathbf{A}_i$ is a collection of differential signals, where damage is simulated at different pixels on the area of

interest. The full dictionary is the (normalized) concatenation of these submatrices, and similarly $\mathbf{y}$ is the concatenation of the individual differential signals. The vector $\mathbf{e}$ is the concatenation of the noise vectors, but this is not shown in Eq. 27. The matrix $\mathbf{A}$ can be precomputed. After obtaining $\mathbf{y}$ from baseline subtraction, it is possible to solve for $\mathbf{x}$ using a sparse solver such as BPDN or OMP.

4.3   Sparse Reconstruction for Unknown Scattering Patterns

For many (if not most) applications, *a priori* knowledge of the exact scattering pattern H is not possible. In these cases, there can be poor imaging performance if the assumed scattering pattern does not match an actual scatterer. Small differences between the nominal and actual dispersion curves and transducer locations can lead to phase shift errors as well. Since the scattering pattern contains both amplitude and phase information about the scatterer, the reconstruction can contain a high level of artifacting if, for example, the scatterer phase shift varies by angle but this is not modeled in H. For a single pair, this would not be an issue, because the reconstruction algorithm could be run with the analytic representation of signals and return a complex amplitude for the scatterer. With multiple pairs, however, there is an issue: suppose $\mathbf{A}$ is generated with the assumption of a uniform scatterer, but in reality there is a scatterer with an angle-dependent phase shift. This will cause phase issues in the dictionary: for some transducer pairs, $\mathbf{y}_i$ will be in phase with the appropriate column $\mathbf{a}_{i,n}$, but for others they will be out-of-phase. However, the structure of Eq. 27 allows only a single coefficient for the concatenated $\mathbf{a}_{i,n}$.

For example, Figure 3 shows the results of the sparse reconstruction algorithm on two sets of simulated data for an aluminum-6061 plate with eight attached transducers using a 5-cycle, 100 kHz Hann-windowed toneburst and the nominal dispersion curve for the A_0 mode. The scatterer in Figure 3(a) has a scattering pattern $H\left[f; \theta^{\mathrm{in}}, \theta^{\mathrm{out}}\right] = 1$, which matches the dictionary atoms; unsurprisingly, this scatterer

is detected perfectly, with only a single pixel visible on the scale shown. Figure 3(b) uses the same dictionary, but the scatterer instead has an amplitude of one but a phase shift of $\pi \cos\left(\theta^{\text{out}} - \theta^{\text{in}}\right)$. The results with this phase mismatch are of significantly poorer quality due to this mismatch. (These images were generated using the BPDN formulation in Eq. 37 with $\sigma' = 0.5\sigma'_{\text{max}}$ and without image denormalization. See Section 4.4 for details.)

4.3.1 Application of the Hilbert Transform

One way to reduce, but not eliminate, errors due to phase mismatch is to use a complex dictionary by using the Hilbert transform to obtain the analytic representation

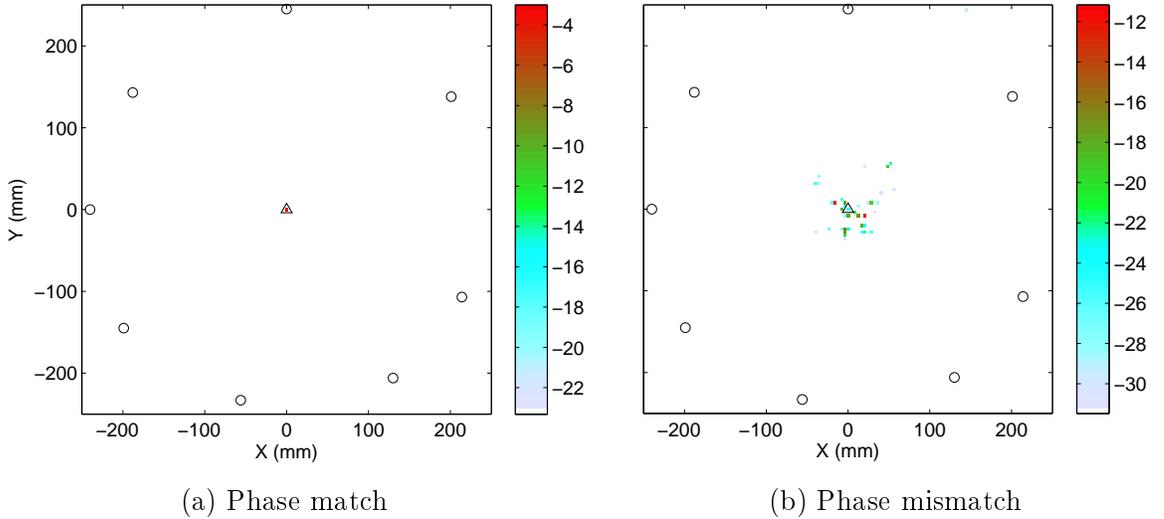(a) Phase match   (b) Phase mismatch

Figure 3: Simulated results using raw (unrectified) signals. (a) The scatterer acts omnidirectionally and with no phase shift, which matches the model used for the dictionary matrix. (b) The scatterer acts omnidirectionally, but with an angle-dependent phase shift. For both images, circles denote transducer locations and the triangle denotes the location of the scatterer. The images are shown on a 20 dB scale, normalized to the largest pixel value; note the difference in the scales of the two images. In both cases, $\sigma' = 0.5\sigma'_{\text{max}}$.

of signals. That is, instead of the form in Eq. 27, the optimization problem

$$
\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e} \ \text{ where } \ \mathbf{y} = \begin{bmatrix} \mathbf{y}_1 + i\mathcal{H}\left\{\mathbf{y}_1\right\} \\ \mathbf{y}_2 + i\mathcal{H}\left\{\mathbf{y}_2\right\} \\ \vdots \\ \mathbf{y}_P + i\mathcal{H}\left\{\mathbf{y}_P\right\} \end{bmatrix} \ \text{ and } \ \mathbf{A} = \begin{bmatrix} \mathbf{A}_1 + i\mathcal{H}\left\{\mathbf{A}_1\right\} \\ \mathbf{A}_2 + i\mathcal{H}\left\{\mathbf{A}_2\right\} \\ \vdots \\ \mathbf{A}_P + i\mathcal{H}\left\{\mathbf{A}_P\right\} \end{bmatrix} \mathbf{D}^{-1} \quad (28)
$$

should be solved, where $\mathcal{H}\left\{\cdot\right\}$ denotes the column-wise Hilbert transform. Note that in this formulation, the matrix $\mathbf{D}$ is scaled by a factor of $\sqrt{2}$. Using analytic representations allows a constant phase shift in $\mathbf{y}$ and also allows different pixels to constructively and destructively interfere. However, the underlying issue is not solved: if only some transducer pairs are out of phase of the dictionary, the solver will produce artifacts. Nevertheless, the additional degree of freedom in each pixel value produces better results overall. Figure 4 shows the differences between the two dictionary types, using the same configuration as Figure 3(b). Figure 4(a) shows the result with the real-valued dictionary. Figure 4(b) shows the result when using analytic representations. Artifacts are still present in the latter, but they are noticeably closer to the actual scatterer location. In addition, the image created with analytic representations is more sparse in the ℓ_1 sense.

4.3.2 Sparse Reconstruction of Signal Envelopes

4.3.2.1 Modified Problem for Signal Envelopes

A different approach to address the issue of phase mismatch is to adapt the method described in Section 4.2 to the use of envelope-detected signals, which are obtained from the absolute value of signal analytic representations. Let $\mathbf{z}$ be a time-domain signal with analytic representation $\hat{\mathbf{z}} = \mathbf{z} + i\mathcal{H}\left\{\mathbf{z}\right\}$. Taking the complex absolute value of $\hat{\mathbf{z}}$ yields the envelope of the signal, which is denoted $\bar{\mathbf{z}} = |\hat{\mathbf{z}}| = |\mathbf{z} + i\mathcal{H}\left\{\mathbf{z}\right\}|$. This envelope detection operation is nonlinear and discards all phase information in

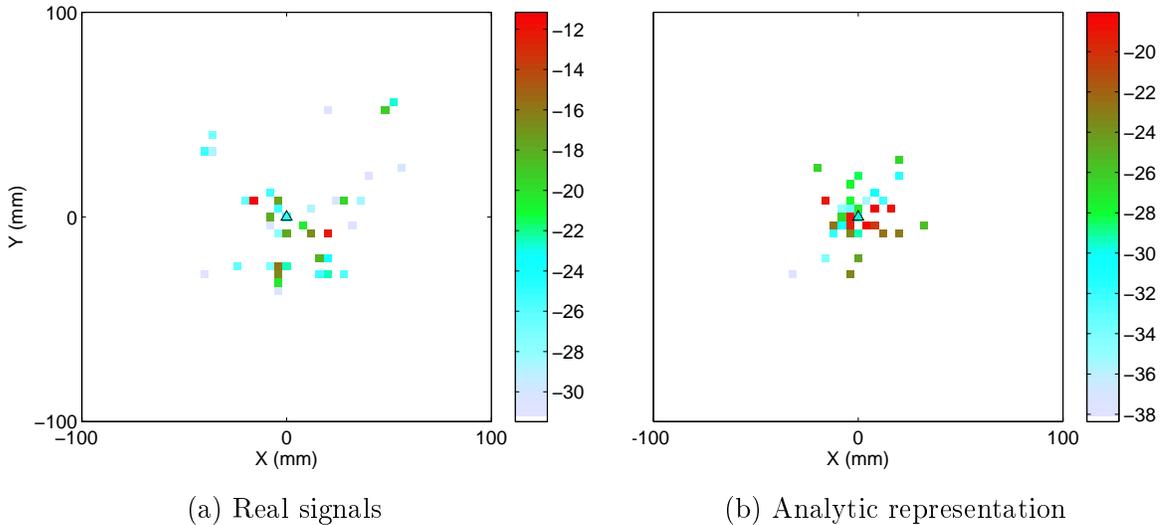(a) Real signals   (b) Analytic representation

Figure 4: Simulated results using raw (unrectified) signals. The scatterer acts omnidirectionally, but with an angle-dependent phase shift. Reconstruction performs poorly because the dictionary does not model this behavior. (a) Imaging result using real-valued signals and dictionary atoms; this is a zoomed version of Figure 3(b). (b) Imaging result using analytic representations. Figures are zoomed and shown on a 20 dB scale, normalized to the largest pixel value. In both cases, $\sigma' = 0.5\sigma'_{\max}$.

$\mathbf{z}$. The approach taken here is to solve the modified problem using signal envelopes,

$$\bar{\mathbf{y}} = \bar{\mathbf{A}}\mathbf{x} + \bar{\mathbf{e}} \ \text{ where } \ \bar{\mathbf{y}} = \begin{bmatrix} \bar{\mathbf{y}}_1 \\ \bar{\mathbf{y}}_2 \\ \vdots \\ \bar{\mathbf{y}}_P \end{bmatrix} \ \text{ and } \ \bar{\mathbf{A}} = \begin{bmatrix} \bar{\mathbf{A}}_1 \\ \bar{\mathbf{A}}_2 \\ \vdots \\ \bar{\mathbf{A}}_P \end{bmatrix} \bar{\mathbf{D}}^{-1}, \tag{29}$$

instead of Eq. 27; here $\hat{\mathbf{y}}$ contains the envelopes of the residual signals, the matrix $\bar{\mathbf{A}}_i = \begin{bmatrix} \bar{\mathbf{a}}_{i,1} & \bar{\mathbf{a}}_{i,2} & \cdots & \bar{\mathbf{a}}_{i,N} \end{bmatrix}$ is the envelope dictionary, and $\bar{D}_{n,n} = \left(\sum_{i=1}^{P} \|\bar{\mathbf{a}}_{i,n}\|_2^2 \right)^{1/2}$. Since the reformulated problem discards phase information, Eq. 29 is more robust to unknown scattering patterns. For scatterers that are at least somewhat omnidirectional, using the envelope-detected form with the assumption that $\left| H\left[\cdots\right] \right| = 1$ can significantly improve results. There is, of course, a price that comes with this operation. Since envelope detection is nonlinear, the signal $\bar{\mathbf{y}}$ is not simply the sum of the envelope-detected contributions due to each scatterer; this reduces the fidelity

46

of the problem reconstruction. In practice, performance is far better with signal envelopes than in the case where scatterer phase information is mismatched. Using signal envelopes for the data in Figure 3 completely eliminates the problems from phase mismatch and results in an image nearly identical to Figure 3(a) (not shown). Signal envelopes do not address scatterers with directionally-dependant amplitude.

4.3.2.2 Envelopes of Noisy Signals

A secondary issue with envelope detection is the transformation of the noise component of $\mathbf{y}$. Suppose that $\mathbf{e}$ is zero-mean, i.i.d Gaussian random noise with variance σ^2. The individual entries of $\mathbf{y}$ then have the distribution

$$y_l \sim \mathcal{N}\left([\mathbf{A}\mathbf{x}]_l, \sigma^2\right). \tag{30}$$

If the analytic representation of $\mathbf{y}$ is taken, the distribution is complex normal:

$$[\hat{\mathbf{y}}]_l \sim \mathcal{CN}\left(\left[\hat{\mathbf{A}}\mathbf{x}\right]_l, 2\sigma^2\right), \tag{31}$$

where $\hat{\mathbf{A}}$ holds the analytic representations of the columns of $\mathbf{A}$. Note that entries of the analytic signal will only be statistically independent if they are downsampled by a factor of two (to "cancel out" the extra redundancy introduced by the Hilbert transform). Taking the complex absolute value of the distribution in Eq. 31 results in a Rice distribution [108]:

$$\bar{y}_l \sim \text{Rice}\left(\left|\left[\hat{\mathbf{A}}\mathbf{x}\right]_l\right|, \sigma\right). \tag{32}$$

The Rice distribution has two interesting limiting cases. As the first argument (which represents the complex absolute value of the complex normal random variable and is often denoted ν) approaches infinity, the Rice distribution approaches a normal distribution:

$$\text{Rice}\left(\nu, \sigma\right) \xrightarrow{\text{approx.}} \mathcal{N}\left(\nu, \sigma^2\right) \text{ if } \nu >> \sigma. \tag{33}$$

47

On the other hand, if $\nu = 0$, the Rice distribution reduces to the Rayleigh distribution:

$$\text{Rice}\,(0, \sigma) = \text{Rayleigh}\,(\sigma).\qquad(34)$$

When the envelope-detection operation is performed on $\mathbf{y}$, it therefore distorts the noise component. In the parts of $\mathbf{y}$ that are mostly signal, the noise retains an approximately Gaussian distribution; however, sections that are dominated by noise become approximately Rayleigh-distributed instead. Since the Rayleigh distribution arises from taking an absolute value, it has a nonzero mean; if $R \sim \text{Rayleigh}\,(\sigma)$, then $E\,[R] = \sigma\sqrt{\pi/2}$. This nonzero noise mean can interfere with reconstruction results and must be compensated in some way. The approach initially taken is simple subtraction of this mean; that is, the problem

$$\left(\bar{\mathbf{y}} - \sigma \mathbb{1}\sqrt{\pi/2}\right) = \bar{\mathbf{A}}\mathbf{x} + \bar{\mathbf{e}}\qquad(35)$$

is solved instead. If the noise power is unknown, it is possible to estimate σ; for example,

$$\hat{\sigma} = \text{median}\,(\bar{\mathbf{y}})\sqrt{2/\pi}\qquad(36)$$

is a reasonable estimator when most of the time samples consist only of noise; the result is to subtract the median value from the noisy signal. Subtracting this quantity from the signal effectively converts the noise to zero-mean noise while reducing the amplitude of signal components in a manner similar to soft thresholding.

### 4.3.2.3 Effects of Uncompensated Rician Noise

Figure 5(a) shows sparse imaging results using signal envelopes on noisy data. Additive white Gaussian noise was added to the signal used in Figs. 3 and 4 to create a signal with 0 dB SNR within the bandwidth of the toneburst. The sparse reconstruction method is able to recover the location of the scatterer, but the spot size is approximately 75 mm in diameter and significant artifacts are present in the corners of the plate due to the nonzero mean. Figure 5(b) shows the first 600 μs of the noisy

signal, its projection onto the range of $\bar{\mathbf{A}}$, and the signal obtained by reconstruction, $\bar{\mathbf{A}}\bar{\mathbf{x}}$, for the first transducer pair. Since the noisy signal and its projection have nonzero mean, the solver "wants" to create a reconstructed signal with wide support. The pixels in the corners of the plate are the most time-delayed and therefore are selected to widen the support as much as possible; this phenomenon is also present when signals contain interference due to boundary reflections or poor baseline subtraction. Figures 5(c) and 5(d) show the results when the median is subtracted from $\bar{\mathbf{y}}$ before performing sparse reconstruction. Only two pixels are detected at the 20 dB level; the image is unquestionably superior to that of the uncompensated case. (All reconstructions are performed with $\sigma' = 0.5\sigma'_{\max}$; the energy fraction $E_{\parallel}$ is 0.24 for the uncompensated image and 0.10 for the compensated image. The energies are particularly low due to the high amount of noise in the signals; see Section 4.4 for details on these measures.)

Noise compensation is even more important for the case of multiple scatterers. Figure 6 compares imaging results ($\sigma' = 0.5\sigma'_{\max}$) for uncompensated and compensated noisy signal envelopes of simulated signals due to four scatterers with an in-bandwidth SNR of 0 dB. Imaging on the uncompensated signals ($E_{\parallel} = 0.23$) fails catastrophically, with three scatterers completely ignored; even at lower σ' values, reconstruction is unsuccessful. The image for the compensated signals ($E_{\parallel} = 0.08$) contains several lower-amplitude artifacts, but all of the scatterers are detected.

4.3.3   Summary

For the case of an unknown scatterer, a modest increase in imaging quality can be obtained by using signal analytic representations. However, the reconstruction is improved greatly by using signal envelopes instead of either raw (RF) signals or their analytic representations. Both the measurements and the dictionary columns are replaced with their envelopes for this method. This converts zero-mean noise to
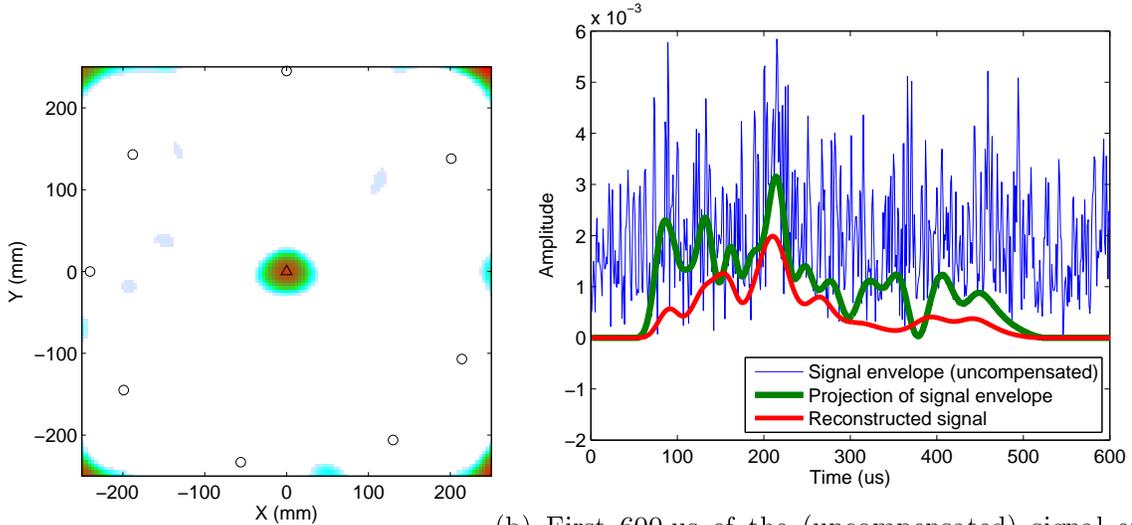
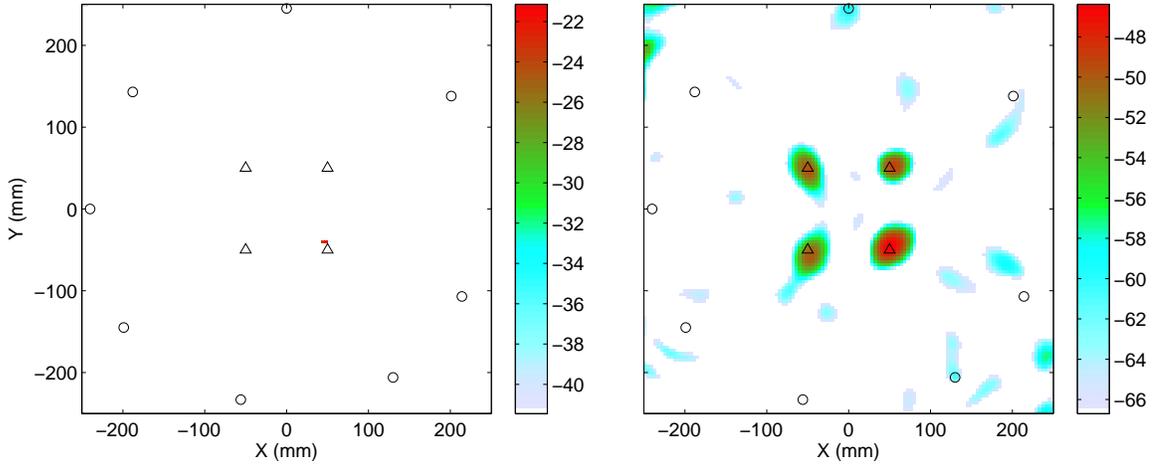(a) Reconstruction with uncompensated signal envelopes

(b) First 600 µs of the (uncompensated) signal envelope, projection, and reconstruction for transducer pair #1
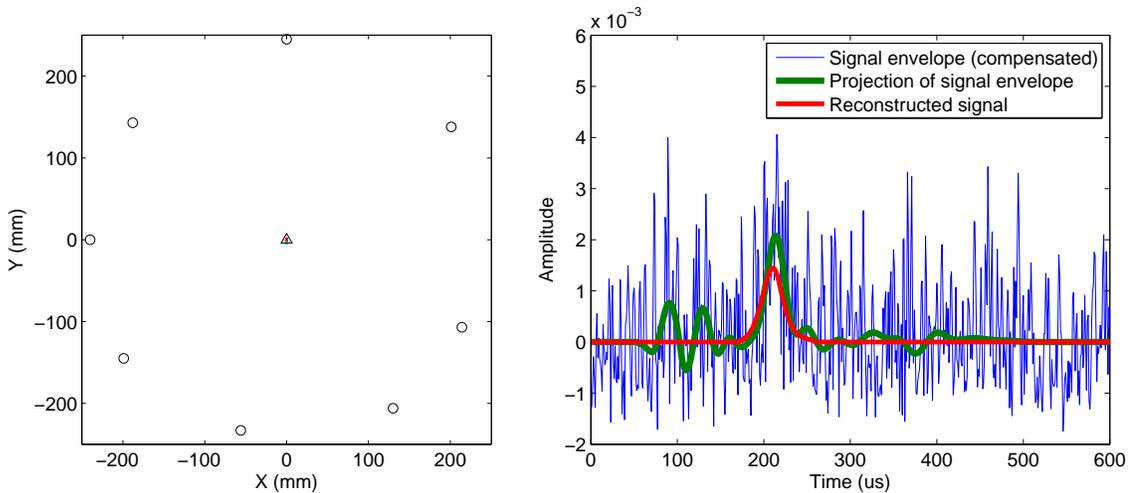
(c) Reconstruction with compensated signal envelopes

(d) First 600 µs of the (compensated) signal envelope, projection, and reconstruction for transducer pair #1

Figure 5: Simulated results using envelopes of noisy signals (0 dB SNR within the toneburst bandwidth). (a) Imaging result for signal envelopes with uncompensated noise. (b) The signal envelope (blue) has a non-zero mean. The solver picks the optimal signal (red) to match the projection of the signal envelope onto the column space of the dictionary (green). (c) Imaging result using envelope compensation. (d) Subtracting the median from the signal envelope results in zero-mean noise. Both reconstruction images are shown on a 20 dB scale, normalized to the largest pixel value. Circles denote the positions of transducer elements; the triangle is the location of the simulated scatterer. The colorbar is not shown to conserve page width.

(a) Reconstruction with uncompensated signal envelopes

(b) Reconstruction with compensated signal envelopes

Figure 6: Simulated results using envelopes of noisy signals (0 dB SNR within the toneburst bandwidth) for four scatterers. (a) Signal reconstruction with uncompensated envelopes. (b) Signal reconstruction with envelopes that have been compensated for noise. Both images are shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations; triangles denote scatterers.

nonzero-mean noise, for which compensation must be performed if a significant noise level is present. Taking a signal envelope is a nonlinear operation which degrades the performance of the reconstruction algorithm, but not as much as an incorrectly-modeled scatterer.

4.4 *Performing Sparse Reconstruction with BPDN*

In this section, the symbols $\mathbf{y}$, $\mathbf{A}$, and $\mathbf{D}$ refer to the measured signal, dictionary matrix, and normalization matrix, respectively, but all analysis applies even if using analytic representations or signal envelopes as described in Section 4.3.

4.4.1 Choice of σ and Column Space Projection

Once the problem is in the form of Eq. 27 (or Eqs. 28 or 29), a sparse reconstruction algorithm can recover the scattering coefficients at each pixel. For this problem, basis pursuit denoising works very well and is recommended. As shown in Eq. 18, BPDN

requires a parameter σ that offers a trade-off between sparsity and reconstruction fidelity; the solver finds the sparsest $\tilde{\mathbf{x}}$ that is within a distance of σ to the solution space of $\mathbf{y} = \mathbf{Ax}$. The value of σ must fall in the range $0 \leq \sigma < \|\mathbf{y}\|_2$; the trivial solution $\tilde{\mathbf{x}} = \mathbf{0}$ is the sparsest solution if $\sigma \geq \|\mathbf{y}\|_2$, while a choice of σ that is too low will likely ensure that no solution exists. ($\sigma = 0$ reduces the problem to basis pursuit.)

To show that there exists a low-valued σ for which no solution exists, split the vector $\mathbf{y}$ into two parts: $\mathbf{y}^{\|\mathbf{A}} = \mathbf{P_A y}$, the portion of $\mathbf{y}$ that lies in the column space of $\mathbf{A}$, and $\mathbf{y}^{\perp \mathbf{A}} = \mathbf{y} - \mathbf{y}^{\perp \mathbf{A}} = (\mathbf{I} - \mathbf{P_A})\,\mathbf{y}$, the remaining portion that lies in the left null space of $\mathbf{A}$, $\mathrm{null}\,(\mathbf{A}^\mathsf{T})$. The latter signal component, $\mathbf{y}^{\perp \mathbf{A}}$, is outside the range of $\mathbf{A}$ and therefore $\|\mathbf{y} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \geq \|\mathbf{y}^{\perp \mathbf{A}}\|_2$ for all possible $\tilde{\mathbf{x}}$. Therefore, no solution exists to Eq. 18 if $\sigma < \|\mathbf{y}^{\perp \mathbf{A}}\|_2$; the behavior of a particular solver varies in this case, with some returning a suboptimal solution and others returning an error. This problem does not occur in compressed sensing, where the $\mathbf{A}$ matrix typically has full row rank and thus $\mathrm{null}\,(\mathbf{A}^\mathsf{T}) = \{\mathbf{0}\}$.
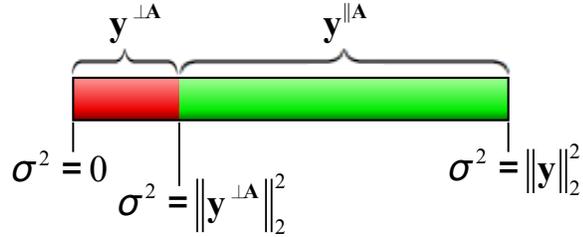

Figure 7: Graphical decomposition of the energy in $\mathbf{y}$ into orthogonal components and the role of σ in BPDN. The energy of $\mathbf{y}$ that lies in the column space of $\mathbf{A}$ is shown in green and the energy in the left null space of $\mathbf{A}$ is shown in red.

A graphical decomposition of the energy in $\mathbf{y}$ is shown in Figure 7, with the energy lying in the column space of $\mathbf{A}$ shown in green, and the remaining energy (i.e., energy in the left null space of $\mathbf{A}$) shown in red. A smaller value of σ corresponds to a smaller residual; a larger value results in a higher level of sparsity. A meaningful solution only

exists for a choice of σ^2 that lies in the green section. Portions of $\mathbf{y}$ that lie outside of $\mathbf{A}$'s column space can be from noise, model mismatch, or other effects such as pixel straddling. In some sense these parts of the signal can be ignored; in fact, solving Eq. 18 is equivalent to solving the optimization with only $\mathbf{y}^{\|\mathbf{A}}$ and an adjusted σ. In other words, the optimization problem

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ subject to } \left\|\mathbf{y}^{\|\mathbf{A}} - \mathbf{A}\mathbf{x}\right\|_2 \leq \sigma', \qquad (37)$$

where $\sigma' := \sqrt{\sigma^2 - \|\mathbf{y}^{\perp\mathbf{A}}\|_2^2}$, results in an identical solution to Eq. 18 if $\left\|\mathbf{y}^{\perp\mathbf{A}}\right\|_2 \leq \sigma \leq \|\mathbf{y}\|_2$.

It is convenient to express the optimization in this form because it simplifies the selection of the parameter to a valid range of $0 \leq \sigma' \leq \sigma'_{\max} := \left\|\mathbf{y}^{\|\mathbf{A}}\right\|_2$. Additionally, the energy fraction $E_\| := \dfrac{\|\mathbf{y}^{\|\mathbf{A}}\|_2^2}{\|\mathbf{y}\|_2^2}$ can be used to quantify the fit of the model. If $E_\|$ is very small, either the model is badly mismatched or no damage is present.

Finding an optimal choice of σ' is not straightforward. The measurements in $\mathbf{y}^{\|\mathbf{A}}$ can be affected by different types of model mismatch; for example, temperature changes can change the Lamb wave mode dispersion curve. If σ' is close to zero, model mismatch can lead to very large spot sizes. In contrast, if multiple scatterers are present, a value very close to $\sigma'_{\max}$ will cause weaker scatterers to be undetected in favor of a more parsimonious representation. For both simulations and experiments, the (arbitrarily chosen) value $\sigma' = 0.5\sigma'_{\max}$ seems to perform well for a small number of scatterers; in rough terms, this choice of σ' requires that the recovered signal $\tilde{\mathbf{y}} := \mathbf{A}\tilde{\mathbf{x}}$ should explain three-fourths of the energy in $\mathbf{y}$ that fits the dictionary model; in general, the fraction of $\mathbf{y}$ required to be explained by the solution is equal to $1 - (\sigma')^2$. A lower value of σ' should be used instead if several scatterers are suspected; this increases the reconstruction fidelity.

After choosing a value of σ' and computing $\mathbf{y}^{\|\mathbf{A}}$, the sparse solver can be used to recover $\tilde{\mathbf{x}}$. The SPGL1 solver for Matlab (available online [103] and described in

[109]) is simple to use and relatively fast; as of version 1.8, the function `spg_bpdn()` is used to solve the BPDN problem, with the syntax

```
x_tilde = spg_bpdn(A, y_A, sigma_prime[, opts]);
```

where `A` is the dictionary matrix $\mathbf{A}$, `y_A` is the vector $\mathbf{y}^{\|\mathbf{A}}$, and `sigma_prime` is σ'; the optional `opts` structure holds configuration settings such as the optimization tolerance. After solving for $\tilde{\mathbf{x}} =$ `x_tilde`, an image is generated by assigning the intensity $|\tilde{x}_n|$ to the pixel $\mathbf{p}_n$, typically on a logarithmic scale.

4.4.2 Image Denormalization

Recall that the matrix $\mathbf{A}$ was normalized with a diagonal matrix $\mathbf{D}$ which had the effect of dividing each column of the dictionary by its norm. The effect of this operation is to scale the coefficient $\tilde{x}_n$ by that norm; see Section 3.4.3.3 for an example. The denormalized pixel values are defined as

$$\tilde{\mathbf{x}}^{\mathrm{d}} := \tilde{\mathbf{x}}\mathbf{D}^{-1}. \tag{38}$$

The denormalized solution vector reflects the "true" scattering behavior in some sense. Without denormalization, pixels that correspond to dictionary atoms with larger norms (before normalization) are scaled more than other pixels; in other words, denormalization compensates for the effects of geometric spreading, which are not included in the normalized dictionary. For the case of Lamb waves, a shorter propagation distance results in a larger dictionary atom, so as a result of denormalization, pixels that are outside the transducer aperture have their pixel values scaled down, and pixels that are extremely close to a transducer have their pixel values scaled up. The former is typically beneficial, as artifacts are much more likely towards plate edges; the latter can be detrimental and create artifacts near transducer locations. Within the convex hull of the transducer locations, the effects of image denormalization are minimal. Most results shown here use the normalized pixel results (i.e., denormalization is not

performed); in practice, images generated with and without denormalization do not differ significantly.

4.5 Discussion

4.5.1 Dictionary Location-Dependence of Coherence

As previously stated, the dictionary matrices $\mathbf{A}$ and $\bar{\mathbf{A}}$ have a high coherence. This is especially true for the latter matrix, which has all non-negative columns and no phase information. The high correlation between $\bar{\mathbf{A}}$'s columns in particular prevents the use of several useful theorems regarding exact recovery. However, the coherence in the dictionary matrix is location-dependent; i.e., columns of the dictionary matrix are highly correlated if they represent pixels that are close together. Figure 8 shows an example of this location-dependence for the envelope dictionary described in Section 4.3.2 for the dictionary entry representing the point $(100, 100)$. The pixel value $z_{(x,y)}$ for a pixel in this image is the correlation of its corresponding dictionary atom with that of $(100, 100)$; that is, $z_{(x,y)} = \left| \mathbf{a}_{(x,y)}^{\mathsf{H}} \mathbf{a}_{(100,100)} \right|$, where $\mathbf{a}_{(x,y)}$ is the dictionary column corresponding to the point (x, y). The correlation is above 0.75 for pixels within roughly 50 mm, above 0.5 for pixels at a distance of less than 100 mm, and ranges from 0.1 to 0.5 elsewhere. The result of this structure is that errors in exact reconstruction are likely to be relatively close to the true location of the scatterer.

4.5.2 Computational Considerations

4.5.2.1 Sparse Matrices

Most excitations used in practice have compact support; in these cases, most entries of the dictionary matrix will be zero-valued. Many programming languages support sparse matrix data structures, which can use considerably less memory than storing the full matrix; MATLAB has native support for sparse matrices, and external libraries are available for many languages, including C++ [110] and Python [111]. To use these representations, the dictionary matrix must first be thresholded at a level above the
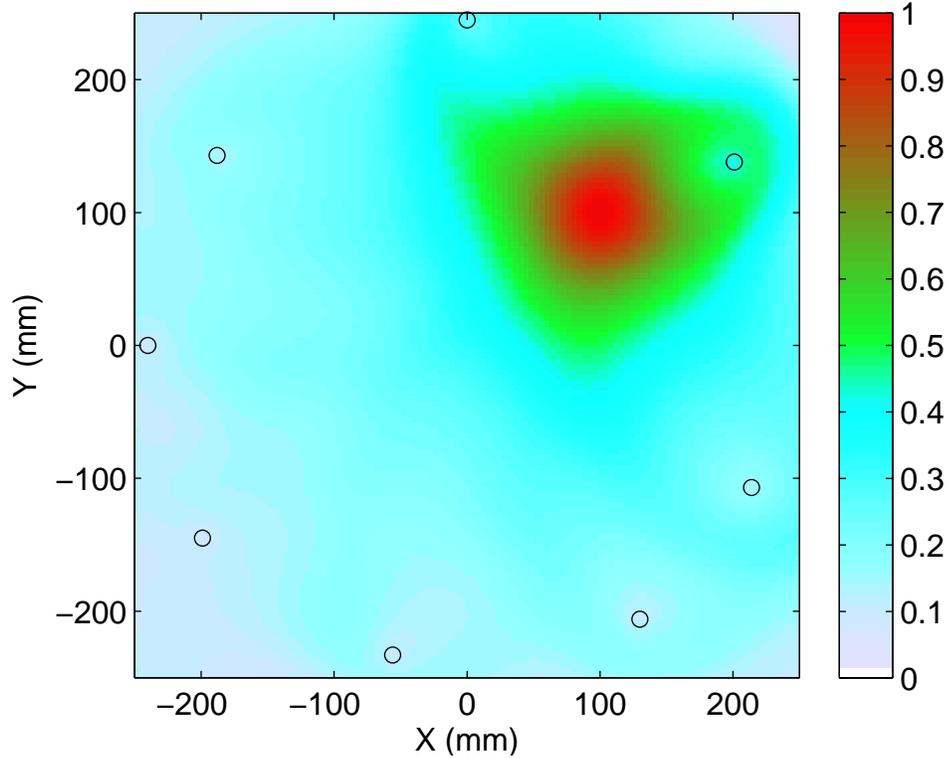
Figure 8: Correlation map of envelope dictionary from $(100, 100)$. The value of each pixel is equal to the correlation between that pixel's corresponding dictionary atom and that of $(100, 100)$.

rounding error caused by the Fourier transform and its inverse. All results shown in this work used MATLAB's sparse matrix implementation for their dictionaries, with a threshold value selected to preserve at least 99.9% of the energy in the matrix, measured using the square of the Frobenius norm.

4.5.2.2 Projection Matrix Approximation

As discussed in Section 4.4.1, it is useful to project the measured signal $\mathbf{y}$ onto the column space of $\mathbf{A}$. The naive approach is to use the definition of a projection matrix, which is $\mathbf{P_A} := \mathbf{A}\mathbf{A}^+$. However, this computation can be quite expensive, as it relies on a matrix inversion, and uses a large amount of memory. A better solution for large matrices is to use the singular value decomposition to approximate $\mathbf{P_A}$. Let $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathsf{H}}$, with the singular values ordered from largest to smallest, and let $\mathbf{U}_r$

denote the matrix that holds the first r columns of $\mathbf{U}$. Then, the matrix $\mathbf{U}_r \mathbf{U}_r^{\mathsf{H}}$ is the best r-rank approximation to $\mathbf{P_A}$, and the fraction of energy preserved in the approximation can be computed as $\frac{\left\| \mathbf{U}_r \mathbf{U}_r^{\mathsf{H}} \mathbf{A} \right\|_F^2}{\|\mathbf{A}\|_F^2}$, where $\|\cdot\|_F$ is the Frobenius norm. This approach saves computation time by avoiding a full matrix inverse and saves memory, since only $\mathbf{U}_r$ needs to be stored instead of the entire projection matrix. All results shown here use this approximation, with r selected such that $\frac{\left\| \mathbf{U}_r \mathbf{U}_r^{\mathsf{H}} \mathbf{A} \right\|_F^2}{\|\mathbf{A}\|_F^2} \geq$ 0.999.

4.5.2.3   Matrix Compression

The dictionary matrix $\mathbf{A}$ (or $\bar{\mathbf{A}}$) can be compressed using various orthogonal transforms and then performing thresholding. Preliminary testing shows that, for dictionaries of Hann-windowed tonebursts, the discrete wavelet transform (using, for example, Daubechies wavelets [112]) can allow a compression factor of 2-5 for unrectified dictionaries and 5-10 for envelope dictionaries. Other transforms such as the lapped orthogonal transform may be viable as well. Additionally, if using column space projection, the matrix of left-singular vectors $\mathbf{U}$ (or its approximation $\mathbf{U}_r$) can be used as an orthogonal basis. Other than testing for feasibility, these transforms were not used for the results presented.

CHAPTER V

BLOCK-SPARSE RECONSTRUCTION FOR ROBUST DAMAGE LOCALIZATION AND CHARACTERIZATION

5.1 Introduction

The problem with the sparse reconstruction method in Chapter 4 is that it strongly depends on the model of the scatterer. If the scatterer's behavior is unknown, as is commonly the case, an envelope operation is applied to remove problematic phase information. However, this also destroys the assumption of linearity, throws out valuable portions of the data, and does not address amplitude model mismatch. One solution to this problem is to use a multidimensional model for scattering. As discussed in Section 3.4, block-sparse methods are particularly suited for a union-of-subspaces problem. In this chapter, scattering is generalized to a linear model where each pixel has its own corresponding subspace in which the residual signals reside. By modeling the problem in this manner, scatterers with unknown behavior can be detected and possibly even characterized. Once again, an infinite plate is considered that supports a single Lamb wave mode with a known dispersion curve and an unspecified number of point scatterers. These assumptions are made for clarity and conciseness; more complicated models, e.g., those that incorporate multimode propagation or boundary reflections, are supported with this methodology. Unlike in the previous chapter, the scattering behavior of potential scatterers is not considered to be known *a priori*. Substantial portions of this chapter also appear in [113] and [114].

5.2 Fixed Damage Location, Single-Pair Scattering Model

5.2.1 One-Dimension-Per-Pair Model of Analytic Signal

First, consider only transducer pair i, whose transducers are affixed to the plate at locations $\mathbf{s}_i$ and $\mathbf{r}_i$. After a baseline measurement is taken, damage is introduced at a known pixel location $\mathbf{p}_n$. The simplest feasible scatterer model is point scattering with frequency-independent, but directionally-dependent, amplitude scaling and phase shifting. The most common Lamb wave excitation is a toneburst, because a narrow bandwidth reduces the effects of dispersion and allows mode tuning. In such cases when the excitation $\mathbf{v}_0$ is narrowband, the scattering model can be relaxed to a point scatterer whose directionally-dependent amplitude and phase responses are constant over the bandwidth of the pulse.

Consider two fixed transducers that interrogate a scatterer under this assumption. The simplest way to model the effects of scatterer on this particular transducer pair is with a constant coefficient. Because the scatterer has directional behavior, this coefficient varies with transducer placement; every transducer pair will measure a different scattering coefficient because it is interrogating at different incoming and outgoing angles. Since this coefficient must account for phase shift, it is convenient to consider the analytic representation of the differential signals, $\hat{\mathbf{y}}_i \in \mathbb{C}^L$, for $i = 1, 2, \ldots, P$. Then, the scattering pattern $H\left[f; \theta^{\text{in}}, \theta^{\text{out}}\right]$ is simply a complex constant, which will be denoted here $x_{i,n} \in \mathbb{C}$, and Eq. 6 can be rewritten:

$$\hat{v}^{\text{scattered}}\left[t; \mathbf{s}_i, \mathbf{r}_i, \mathbf{p}_n, H\right] = x_{i,n} \underset{\mathbf{p_n} \rightarrow \mathbf{r_i}}{\mathcal{P}_{c_p}} \underset{\mathbf{s_i} \rightarrow \mathbf{p_n}}{\mathcal{P}_{c_p}} \hat{v}_0\left[t\right]. \tag{39}$$

Here, $|x_{i,n}|$ is the amplitude of the scatterer and $\angle x_{i,n}$ is the phase shift. As shown in Eq. 10, under perfect baseline subtraction, the vectors $\mathbf{y}_i$ are equal to $v^{\text{scattered}}\left[t; \mathbf{s}, \mathbf{r}, \mathbf{q}, H\right]$ plus noise; thus $\hat{\mathbf{y}}_i$ is the analytic representation of that scattered signal, plus complex noise. Since the value of $x_{i,n}$ is unknown, it must be included in the scattering model. Because the location of damage is assumed to be known, the scattering coefficient is

the only variable and may be found using the equation

$$x_{i,n} = \frac{\hat{\mathbf{a}}_{i,n}^{\mathsf{H}} \, \hat{\mathbf{y}}_i}{\|\hat{\mathbf{a}}_{i,n}\|_2^2},$$ (40)

where $\hat{\mathbf{a}}_{i,n} = \mathcal{P}_{c_p} \atop \mathbf{p}_n \to \mathbf{r}_i \quad \mathcal{P}_{c_p} \atop \mathbf{s}_i \to \mathbf{p}_n} \hat{\mathbf{v}}_0$.

5.2.2   Generalized Model

If a higher-dimensional model for the scattered signal $\hat{\mathbf{y}}$ is available, Eq. 40 can be generalized. Let Λ be the dimension of the model, and let $\hat{\mathbf{A}}_{i,n} \in \mathbb{C}^{L \times \Lambda}$ contain the vectors that model $\hat{\mathbf{y}}_i$ as its columns, and denote the coefficients $\mathbf{x}_{i,n} \in \mathbb{C}^{\Lambda}$. The coefficients can be recovered by using least squares:

$$\mathbf{x}_{i,n} = \hat{\mathbf{A}}_{i,n}^{+} \, \hat{\mathbf{y}}_i.$$ (41)

As an example, this model could be used for multimode signals by using one column for each mode. It can also be used for block-sparse solvers unable to work with complex numbers; in this case, $\hat{\mathbf{A}}_{i,n}^{+}$ should have columns for the real and imaginary parts of $\hat{\mathbf{a}}_{i,n}$.

5.3   *Fixed Damage Location, Multiple-Pair Model*

In this section, all transducer pairs are considered, but the scatterer is still assumed to be at the known location $\mathbf{p}_n$. Regardless of its dimension, the model for the scattered signal $\hat{\mathbf{y}}_i$ is treated here as a matrix, $\hat{\mathbf{A}}_{i,n} \in \mathbb{C}^{L \times \Lambda}$; similarly, even if the one-dimension-per-pair model is used, the scattering coefficient is treated as a vector, $\mathbf{x}_{i,n} \in \mathbb{C}^{\Lambda}$. For the one-dimension-per-pair case, $\Lambda{=}1$.

Since each transducer pair has different incoming and outgoing angles, it cannot be assumed that there is only a single coefficient (or, for the generalized Λ-dimensional model, Λ coefficients). Instead, each pair is modeled as having an independent scattering coefficient (or set of Λ coefficients) and all of these coefficients are evaluated

simultaneously. This can be done by forming the extended linear equation

$$\underbrace{\begin{bmatrix} \hat{\mathbf{y}}_1 \\ \hat{\mathbf{y}}_2 \\ \vdots \\ \hat{\mathbf{y}}_P \end{bmatrix}}_{\hat{\mathbf{y}}} = \underbrace{\begin{bmatrix} \hat{\mathbf{A}}_{1,n} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \hat{\mathbf{A}}_{2,n} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \hat{\mathbf{A}}_{P,n} \end{bmatrix}}_{\hat{\mathbf{A}}_n} \underbrace{\begin{bmatrix} \mathbf{x}_{1,n} \\ \mathbf{x}_{2,n} \\ \vdots \\ \mathbf{x}_{P,n} \end{bmatrix}}_{\mathbf{x}_n} + \hat{\mathbf{e}} \quad (42)$$

where $\hat{\mathbf{y}} \in \mathbb{C}^{LP}$, $\hat{\mathbf{A}}_n \in \mathbb{C}^{LP \times \Lambda P}$, and $\mathbf{x}_n \in \mathbb{C}^{\Lambda P}$. As in Section 5.2.2, the scattering coefficients can be determined by using least squares:

$$\mathbf{x}_n = \hat{\mathbf{A}}_n^+ \hat{\mathbf{y}}. \quad (43)$$

The ΛP-dimensional model in Eq. 42 allows the scatterer to have different amplitude and phase shift between each pair due to the structure of the extended dictionary matrix $\hat{\mathbf{A}}_n$. This is important because many other algorithms require *a priori* characterization of potential scatterers. Since different types of damage may behave differently, it is usually not possible to a specific of scattering behavior. For example, large cracks have highly directional scattering, while through-holes scatter approximately omnidirectionally. In contrast, the $\hat{\mathbf{A}}_n$ matrix is completely parameterized by the distance from each transducer to the damage location, and requires no advance knowledge of the scattering pattern. Instead, the equation solution *is* the scattering pattern, sampled at the various angles dictated by the geometry of the transducers and the location of $\mathbf{p}_n$.

5.4 Unknown Damage Location, Multiple-Pair Model

Finally, consider the case where K scatterers are present at unknown locations on the otherwise damage-free plate; K need not be known in advance. As in Chapter 4, the ROI of the plate is discretized into N pixels. Each of these pixels has a different model matrix $\hat{\mathbf{A}}_n$, since each has a different set of transducer distances and therefore

different propagated signals. This matrix is described in Eq. 42. Using these models, an overcomplete block dictionary

$$\hat{\mathbf{A}}^{\mathrm{d}} := \begin{bmatrix} \hat{\mathbf{A}}_1 & \hat{\mathbf{A}}_2 & \cdots & \hat{\mathbf{A}}_N \end{bmatrix} \in \mathbb{C}^{LP \times N\Lambda P} \tag{44}$$

can be constructed, with columns from matrix $\hat{\mathbf{A}}_n$ assigned to group G_n. As in the sparse reconstruction case, the dictionary matrix must be normalized; the normalized block dictionary matrix is defined as

$$\hat{\mathbf{A}} := \hat{\mathbf{A}}^{\mathrm{d}} \hat{\mathbf{D}}^{-1} \tag{45}$$

where $\hat{\mathbf{D}} \in \mathbb{R}^{N\Lambda P \times N\Lambda P}$ is a diagonal matrix with $D_{j,j} = \left\| \left(\hat{\mathbf{a}}^{\mathrm{d}} \right)_j \right\|_2$, where $\left(\hat{\mathbf{a}}^{\mathrm{d}} \right)_j$ is the jth column of $\hat{\mathbf{A}}^{\mathrm{d}}$. The complete model of the K-scatterer differential signals is then

$$\hat{\mathbf{y}} = \underbrace{\begin{bmatrix} \hat{\mathbf{A}}_1 & \hat{\mathbf{A}}_2 & \cdots & \hat{\mathbf{A}}_N \end{bmatrix} \hat{\mathbf{D}}^{-1}}_{\hat{\mathbf{A}}} \underbrace{\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_N \end{bmatrix}}_{\mathbf{x}} + \hat{\mathbf{e}}, \tag{46}$$

where $\mathbf{x} \in \mathbb{C}^{N\Lambda P}$. If a scatterer is located at pixel n, the corresponding submatrix $\mathbf{x}_n$ will contain its scattering coefficients; if no scatterer is present, the coefficients should be all zero. Because it involves a redundant dictionary, Eq. 46 cannot be solved using least squares to obtain a meaningful result. Instead, the solver must (a) select which blocks should be active and (b) determine the coefficients within each active block. Provided that most of the structure is damage-free, a block-sparse solver will satisfy these criteria by selecting a sparse set of blocks that approximately describes the residual signal measurements.

If the model of one dimension per transducer pair is used (i.e., if $\Lambda = 1$) another way to express the block dictionary is by interleaving the pixel blocks. This results

in the block-diagonal dictionary

$$\hat{\mathbf{A}}^{\text{alt}} := \begin{bmatrix} \mathbf{A}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{A}_P \end{bmatrix} \in \mathbb{C}^{LP \times N\Lambda P}, \tag{47}$$

where each submatrix $\mathbf{A}_i \in \mathbb{C}^{L \times N}$ is a per-pair submatrix. Under this formulation, the block corresponding to pixel i consists of columns $i, i+N, i+2N, \ldots, i+(N-1)N$. The relationship between the non-block dictionary in Eq. 27 and the block dictionary is more readily apparent for this definition.

5.5   *Solving the Block-Sparse Model*

Analysis from Section 4.4 applies to the block-sparse case as well. In particular, the use of the vector $\hat{\mathbf{y}}^{\|\hat{\mathbf{A}}} = \mathbf{P}_{\hat{\mathbf{A}}}\hat{\mathbf{y}}$ is recommended, which allows a range of σ' from 0 to $\left\|\hat{\mathbf{y}}^{\|\hat{\mathbf{A}}}\right\|_2$. This results in the optimization problem

$$\tilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \sum_{G \in \mathcal{G}} \|\mathbf{x}_G\|_2 \quad \text{subject to} \quad \left\|\hat{\mathbf{y}}^{\|\hat{\mathbf{A}}} - \hat{\mathbf{A}}\mathbf{x}\right\|_2 \leq \sigma'. \tag{48}$$

The `spg_group()` function of SPGL1 1.8 can then be used to solve the block-sparse problem, with the syntax

```
x_tilde = spg_group(A_hat, y_hat_A, G, sigma_prime[, opts]);
```

where `A_hat` is the dictionary matrix $\hat{\mathbf{A}}$, `y_hat_A` is $\hat{\mathbf{y}}^{\|\hat{\mathbf{A}}}$, `sigma_prime` is σ', and the vector `G` holds the group assignments for `y_hat_A`. The optional `opts` structure holds configuration settings such as the optimization tolerance.

Using the block method, each pixel n has multiple coefficients in its subvector $\tilde{\mathbf{x}}_n$. To generate the image, assign the pixel n the intensity value $\|\tilde{\mathbf{x}}_n\|_2$. As before, the "denormalized" coefficients $\tilde{\mathbf{x}}^{\text{d}} = \hat{\mathbf{D}}^{-1}\tilde{\mathbf{x}}$ may be used in place of $\tilde{\mathbf{x}}$ for this image.

In addition to its ability to work with unknown scattering behavior, this block-sparse method can be used for damage characterization. The block coefficients in

the denormalized solution vector describe how the damage behaves between individual transducer pairs, allowing for the possibility of scattering pattern estimation. Examples are shown in the next chapter.

5.6 Discussion

5.6.1 Advantages over Sparse Reconstruction

For most applications, "standard" sparse reconstruction must be performed with signal envelopes as described in Section 4.3, both because the scattering pattern is unknown *a priori* and other contributions to phase may not be well-known. This worsens the noise conditions as described in Section 4.3.2.2, but more importantly, it introduces a nonlinearity into a problem that assumes linear scattering because the dictionary cannot incorporate destructive interference due to phase differences. For excitations with compact support, the effects on reconstruction will be minimal if scatterers are far apart; however, at some separation distance, the nonlinearity of the problem can mask scatterers. The block-sparse method is not subject to the same issue since the algorithm "chooses" appropriate signal phases as part of the solution.

Figure 9 compares results using sparse reconstruction with signal envelopes to block-sparse reconstruction using $\sigma' = 0.5\sigma'_{\max}$ with simulated signals of a 3.175 mm, Al-6061 plate with four equal-strength, omnidirectional scatterers forming a square with 50 mm side length. Amplitude differences are solely a function of transducer and scatterer placement. Because of the nonlinearity involved, the envelope method (Figure 9(a)) fails to detect the top-right scatterer, and the locations of the bottom two scatterers are biased towards each other. The block-sparse method (Figure 9(b)) is not subject to these problems, though the amplitude of the top-right scatterer is lower. At lower values of σ', the envelope method detects the otherwise masked scatterer, though only weakly (not shown). The excitation used was a 5-cycle, 100 kHz Hann-windowed toneburst using the A_0 mode, with a pixel resolution of 4 mm.

(a) Sparse reconstruction with envelopes   (b) Block-sparse reconstruction
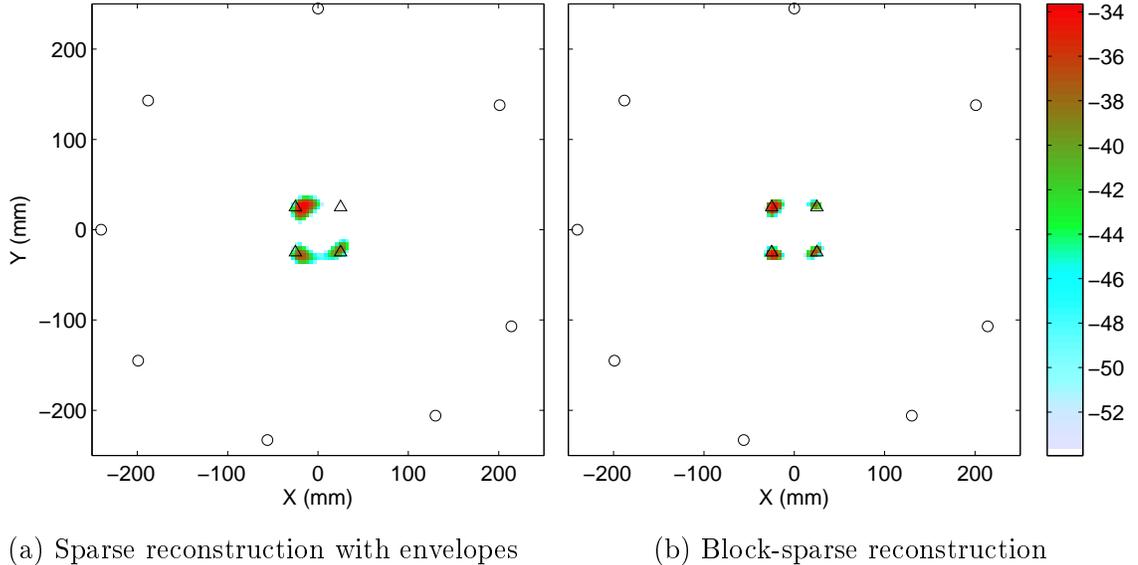
Figure 9: Simulated imaging results for four scatterers in close proximity. (a) Image generated with sparse reconstruction on signal envelopes. (b) Image generated with block-sparse reconstruction. For both images, circles denote transducer locations and triangles denote scatterer locations. All scatterers are of equal strength, though transducer placement can affect the pixel values. The images are shown on a 20 dB scale, normalized to the largest pixel value, and used $\sigma' = 0.5\sigma'_{\mathrm{max}}$.

Image denormalization is not performed for Figure 9(b), but the effects are negligible.

5.6.2   Dictionary Location-Dependence of Coherence

The coherence of $\hat{\mathbf{A}}$ is nearly 1 for the same reason that the submatrices in Section 4.2.1 are coherent. It is more meaningful, however, to examine a measure of similarity between dictionary blocks (subspaces) to determine how well coherence is localized. One such measure is the set of *principal angles* between two subspaces [115]. These represent the angles between individual dimensions of the two subspaces. The first (smallest) principal angle between two subspaces $\mathcal{U}$ and $\mathcal{V}$ is defined as

$$\cos\theta_1 := \min_{\substack{\mathbf{u}\in\mathcal{U} \\ \mathbf{v}\in\mathcal{V}}} \frac{\mathbf{u}^{\mathsf{H}}\mathbf{v}}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2}. \tag{49}$$

The minimizing vectors are denoted $\mathbf{u}_1$ and $\mathbf{v}_1$. The other principal angles are defined recursively:

$$\cos \theta_k := \min_{\substack{\mathbf{u} \in \mathcal{U} \\ \mathbf{v} \in \mathcal{V}}} \frac{\mathbf{u}^{\mathsf{H}} \mathbf{v}}{\|\mathbf{u}\|_2 \|\mathbf{v}\|_2} \text{ subject to } \mathbf{u} \perp \mathbf{u}_i, \mathbf{v} \perp \mathbf{v}_i, \text{ for } i = 1, 2, \ldots, k-1. \quad (50)$$

The resulting $D = \min \{\dim(\mathcal{U}), \dim(\mathcal{V})\}$ principal angles, $\{\theta_1, \theta_2, \ldots, \theta_D\}$, are in ascending order. The first principal angle θ_1 is the angle between the closest dimensions $\mathcal{U}$ and $\mathcal{V}$; θ_2 is the angle between the next-closest dimensions, and so on, up to the angle between the two least-correlated dimensions, θ_D. If $\dim(\mathcal{U}) = \dim(\mathcal{V}) = 1$, then θ_1 is just the angle between any two non-zero vectors in the two subspaces. The principal angles are also related to the projection matrices of the two subspaces. Let $\mathbf{P}_{\mathcal{U}}$ and $\mathbf{P}_{\mathcal{V}}$ be projection matrices onto $\mathcal{U}$ and $\mathcal{V}$, respectively. Then the D largest non-negative eigenvalues of $\mathbf{P}_{\mathcal{U}} - \mathbf{P}_{\mathcal{V}}$ are $\{\sin \theta_1, \sin \theta_2, \ldots, \sin \theta_D\}$.

The cosines of these principal angles, $\rho_k := \cos \theta_k$, represent correlation coefficients between the individual dimensions. In fact, the principal angles have an interesting relation to the statistical field of canonical correlation analysis [116]. Let $\mathcal{U} = \mathrm{range}(\mathbf{U})$ and $\mathcal{V} = \mathrm{range}(\mathbf{V})$ for some matrices $\mathbf{U}$ and $\mathbf{V}$. If the columns of $\mathbf{U}$ and $\mathbf{V}$ are zero-mean, then the values $\{\rho_1, \rho_2, \ldots, \rho_D\}$ are the canonical correlations of $\mathbf{U}^{\mathsf{T}}$ and $\mathbf{V}^{\mathsf{T}}$. In this case, the rows of $\mathbf{U}^{\mathsf{T}}$ and $\mathbf{V}^{\mathsf{T}}$ are considered to be random variables, and columns are considered to be observations.

As in Section 4.5.1, the correlations $\{\rho_k\}$ between two blocks depend on the distance between their corresponding pixels. Since the blocks in this dictionary have more degrees of freedom than the single vectors for the non-block case, the "worst" correlation between two blocks, ρ_1, will be higher than the corresponding correlation coefficient for the non-block dictionary. Figure 10 shows a correlation map similar to Figure 8, but for a block dictionary that uses one (complex) vector per transducer pair, creating a 28-dimensional subspace for each pixel. Other parameters are identical to those specified in Section 4.5.1.
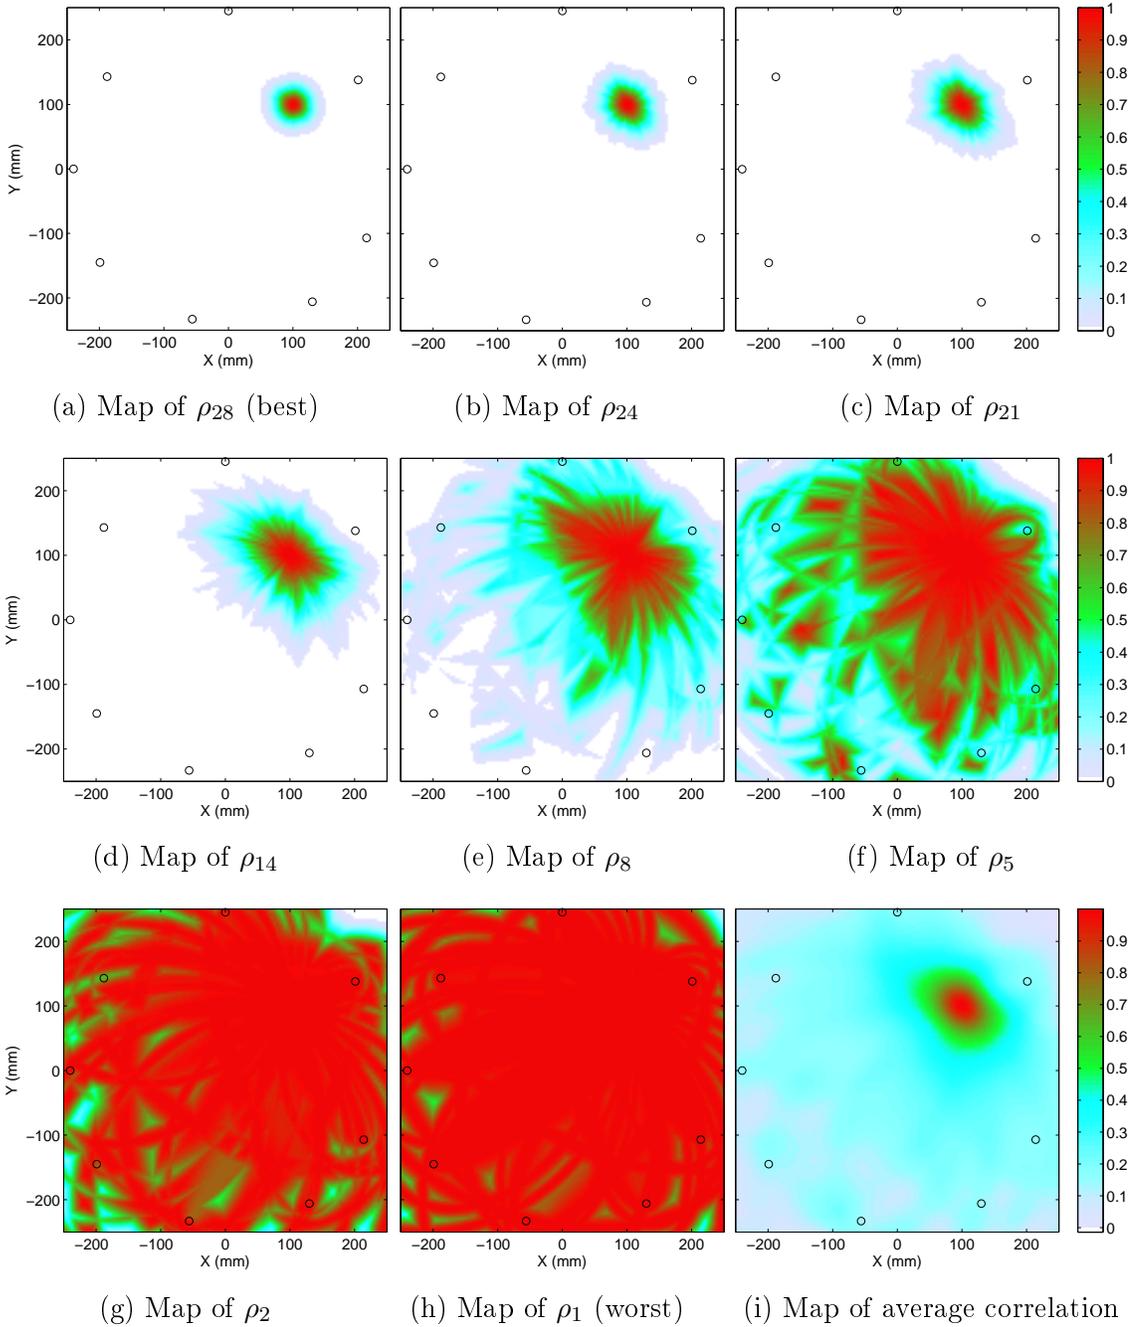
(a) Map of ρ_{28} (best)   (b) Map of ρ_{24} (c) Map of ρ_{21}

(d) Map of ρ_{14} (e) Map of ρ_8 (f) Map of ρ_5

(g) Map of ρ_2 (h) Map of ρ_1 (worst)   (i) Map of average correlation

Figure 10: Correlation maps of block dictionary from $(100, 100)$. The value of each pixel is equal to $\left| \rho_k^{(x,y),(100,100)} \right|$, the cosine of the kth principal angle between that pixel's corresponding dictionary block and that of $(100, 100)$. (a-h) Each map represents the correlation of a single pair of dimensions between the 28-dimensional subspaces. Pixels that are farther away have more uncorrelated dimensions. (i) Map of the average correlation $\frac{1}{28} \sum_{k=1}^{28} \left| \rho_k^{(x,y),(100,100)} \right|$.

To generate the figure, the principal angles between each subspace (block) in $\hat{\mathbf{A}}$ and the subspace corresponding to $(100, 100)$ were determined. The result was a set of 28 principal angles, $\theta_k^{(x,y),(100,100)}$, and their corresponding correlation coefficients, $\rho_k^{(x,y),(100,100)}$. The result is 28 different correlation maps (one for each principal angle dimension), where the map for the ρ_1 values shows the worst-case correlation and the map for the ρ_{28} values displays the best-case correlations. Eight of these maps are shown. Each displays a degree of locality: the least-correlated subspaces have a correlation of less than 0.5 at a distance of roughly 25 mm; the map corresponding to ρ_{14} (the middle correlation map) has a spot size of 100-150 mm; the highest-correlated maps have a correlation of nearly 1 everywhere. The near-total correlation of the subspaces in at least one dimension is a result of the elliptical nature of the data. Two pixels lying on the same ellipse with foci at any two transducer locations will share at least one dimension because the time-of-flight from one transducer to the pixel location to the other transducer is the same for both pixels. The figure clearly shows that most dimensions for close pixels are highly correlated. In contrast, the subspaces for two distant pixels have a number of dimensions that are essentially uncorrelated. The result of this sort of structure is that if there are errors in locating scatterers due to noise increases and modeling errors, the solver is more likely to select nearby pixels than distant ones. Figure 10(i) shows the average of all 28 maps and is comparable to the case with signal envelopes shown in Figure 8.

5.6.3   Computational Considerations

5.6.3.1   Sparse Matrices and Matrix Compression

The discussion in Sections 4.5.2.1 and 4.5.2.3 also applies to the block-sparse dictionary matrix $\hat{\mathbf{A}}$. This matrix has many more columns and is much more sparse than the dictionary matrix of the previous chapter, so the use of a sparse matrix structure is all but required for large problems.

5.6.3.2 Projection Matrix Approximation

As in Section 4.5.2.2, the projection matrix $\mathbf{P}_{\hat{\mathbf{A}}}$ can be approximated with the singular value decomposition. Due to the block structure of this dictionary, the projection can also be split into blocks to further reduce computation time. Let $\hat{\mathbf{A}}^i \in \mathbb{C}^{L \times N\Lambda P}$ denote the submatrix of $\hat{\mathbf{A}}$ that contains only those rows corresponding to pair i; the dictionary can be broken into P such blocks. The projection matrix for block i can be approximated as $\mathbf{P}_{\hat{\mathbf{A}}^i} \approx \mathbf{U}_{r_i}^i \left(\mathbf{U}_{r_i}^i \right)^{\mathsf{H}}$, where $\mathbf{U}_{r_i}^i$ is the first r_i columns of the left-singular matrix $\mathbf{U}^i$ of $\hat{\mathbf{A}}^i$. The projection matrix $\mathbf{P}_{\hat{\mathbf{A}}}$ can then be approximated with the block-diagonal matrix

$$\mathbf{P}_{\hat{\mathbf{A}}} \approx \mathbf{P}_{\hat{\mathbf{A}}}^{\{r_i\}} := \begin{bmatrix} \mathbf{P}_{\hat{\mathbf{A}}^1} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{P}_{\hat{\mathbf{A}}^2} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{P}_{\hat{\mathbf{A}}^P} \end{bmatrix}. \tag{51}$$

All block-sparse reconstruction dictionary projection results presented here use this procedure, with $\{r_i\}$ selected such that $\dfrac{\left\| \mathbf{P}_{\hat{\mathbf{A}}}^{\{r_i\}} \hat{\mathbf{A}} \right\|_F^2}{\left\| \hat{\mathbf{A}} \right\|_F^2} > 0.999$.

CHAPTER VI

RESULTS AND ANALYSIS

6.1 Experiment One

The first experiment was performed on an Aluminum-6061 plate with dimensions 1220 mm × 1220 mm × 3.175 mm. Eight 0.5 mm thick, 7 mm diameter transducers were affixed to locations in the center of the plate in a roughly circular pattern, with the coordinates listed in Table 4. All are listed in two dimensions, with the origin at the center of the plate. For this experiment, it was desired to minimize the effects of reflections from the plate edges; in addition to the large plate area, the sides of the plate were damped with duct-sealing compound. The acquisition setup is shown in Figure 11 and utilized an arbitrary waveform generator, signal multiplexer, amplifier, and digitizer. Additional details about the experiment are provided in [114].

To simulate damage, two 77.8 mm long steel rods were affixed to the plate with glue. The larger 9 mm diameter rod was attached at (30 mm, 40 mm), while the smaller, 6 mm rod was placed at (−20 mm, −80 mm). One set of baselines was recorded on the pristine plate; two follow-up measurements were taken: one after

Table 4: Transducer placement for experiments one and two

Transducer	X (mm)	Y (mm)
#1	-240	0
#2	-199	-145
#3	-56	-233
#4	130	-206
#5	214	-107
#6	201	138
#7	0	245
#8	-188	143
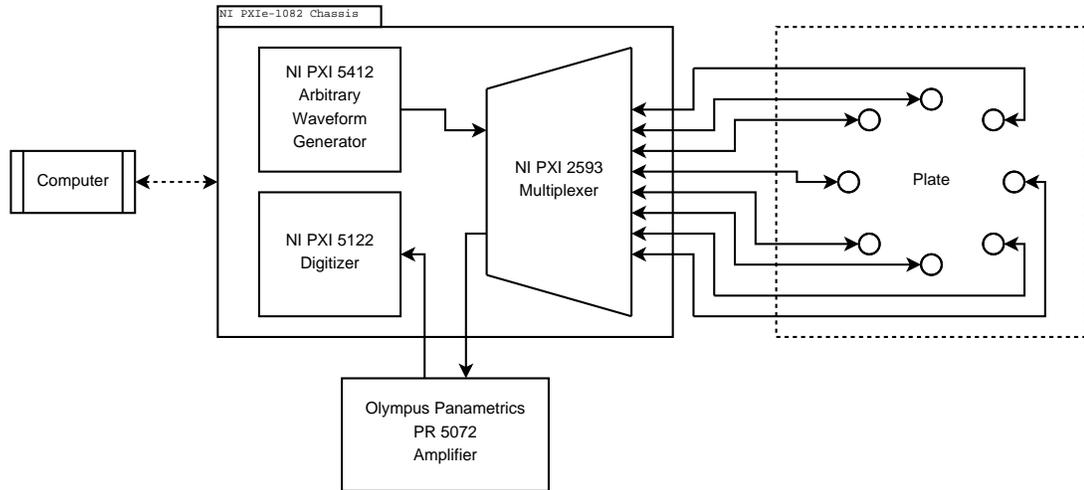
Figure 11: Block diagram of experimental system.

placing the first (larger) rod, and one after both rods were attached.

Excitation was performed with a 200 μs chirp over the 50 kHz-500 kHz range and sampled at 20 MHz. These signals were postprocessed to obtain the equivalent response [117] to a 5-cycle, Hann-windowed toneburst with center frequency 100 kHz; at this frequency, the A_0 mode is dominant. Finally, the signals were downsampled to a 1 MHz sampling rate and time windowed to 1 ms (1000 samples).

6.1.1 Delay-and-Sum

For comparison, an image was first generated using existing custom software capable of performing DAS imaging [43]. The software uses signal envelopes and does not incorporate dispersion, instead estimating the group velocity from experimental data. For these images, no weighting of any sort was applied – signals were simply time-shifted by the appropriate amount and added together. Figure 12 shows the DAS result for the experimental data, on a 20 dB scale. Both rods are detected, but their spot sizes are on the order of 75 mm in diameter and bleed into each other.

71
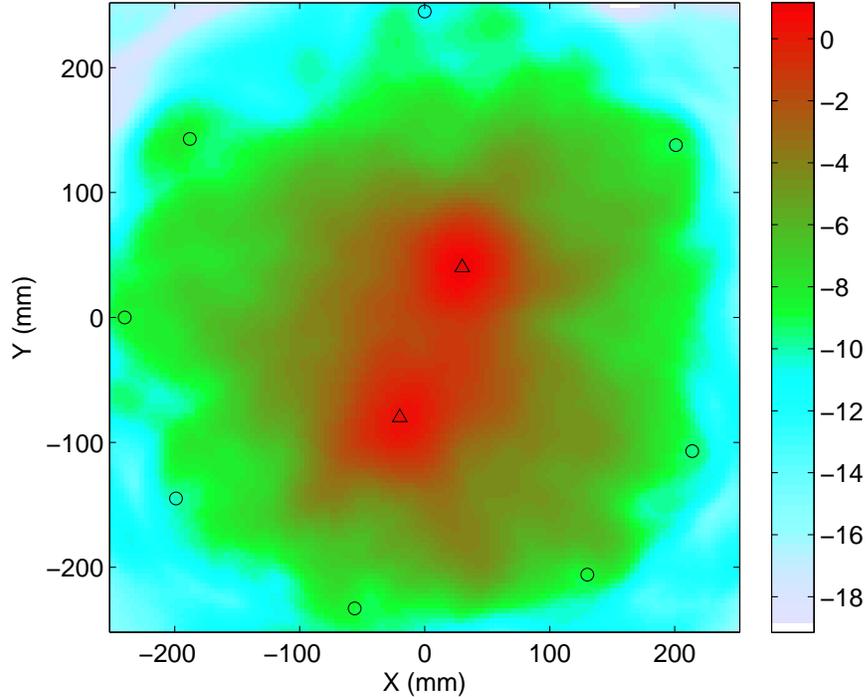
Figure 12: Result for experiment one (glued-on steel rods) using delay-and-sum imaging on signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.
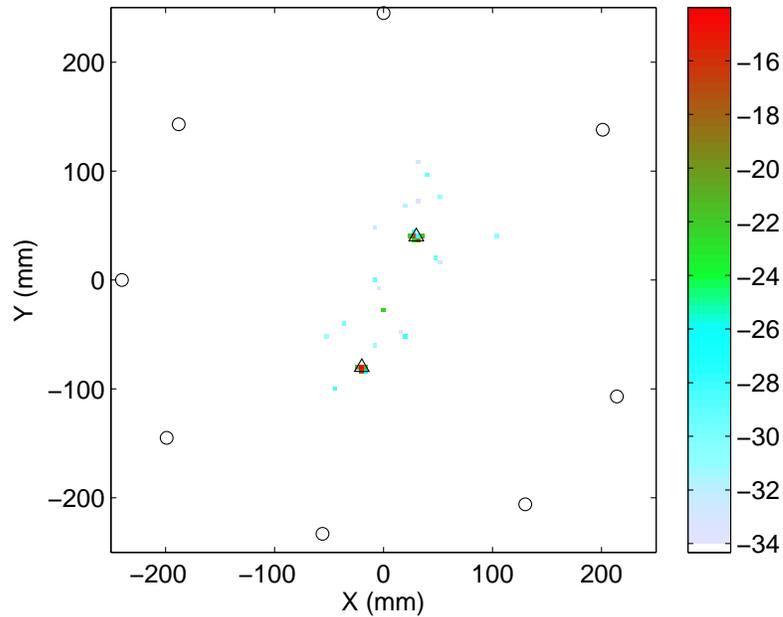
6.1.2 Sparse Reconstruction with Raw Signals

The sparse reconstruction dictionary $\mathbf{A}$ was formed using the nominally-computed A_0 dispersion curves and the same toneburst waveform, at the same 1 MHz sampling frequency and number of samples. The square -248 mm $\leq x, y \leq 248$ mm was designated as the area of interest and was discretized into 4 mm pixels. Scattering was assumed to be uniform and with no phase shift. The energy fraction after projection onto the real-valued dictionary is $E_{\parallel} = 0.64$, indicating that 64% of the energy in the concatenated residuals lies in the column space of the dictionary; this is a typical energy fraction for signals with good baseline subtraction and low levels of noise and interference. Reconstruction was performed using $\sigma' = 0.5\sigma'_{\max}$, meaning that the reconstructed signal was required to explain 75 percent of the signal energy that remained after projection. Imaging results are shown in Figure 13, with

no image denormalization. Reconstruction fails catastrophically, primarily due phase mismatch.

A second reconstruction was performed using the analytic representations of the differential signals and the dictionary matrix as described in Section 4.3.1. The energy fraction after projection is $E_\| = 0.64$; imaging results using $\sigma' = 0.5\sigma'_{max}$ and no denormalization are shown in Figure 14 and show considerable improvement, though artifacts are still present, including one with an intensity that is less than 4 dB below the largest pixel of the top scatterer.

6.1.3 Sparse Reconstruction with Signal Envelopes

Next, reconstruction was repeated using signal envelopes. The dictionary matrix $\bar{\mathbf{A}}$ was constructed as described in Section 4.3.2; it is the complex absolute value of the dictionary of analytic representations. Noise compensation was performed as described in Section 4.3.2.2. The projection energy fraction $E_\| = 0.65$; the result with $\sigma' = 0.5\sigma'_{max}$ and no denormalization is shown in Figure 15. The two scatterers are well localized and no artifacts are visible at the 20 dB level. When noise compensation is not performed, the spot sizes are slightly larger (not shown).

6.1.4 Sparse Reconstruction with Signal Envelopes with Unknown Dispersion Curves

To test robustness to model mismatch, the previous reconstruction was repeated with a dictionary that used only the theoretically-calculated A_0 group velocity; i.e., dispersion was not modeled. No time offset was used and the group velocity was confirmed to agree with experimental data. The energy fraction, $E_\| = 0.65$, is unchanged to two decimal places. Imaging results using the same parameters are shown in Figure 16 and are visually similar to the results with the dispersive dictionary, despite the simpler propagation model.
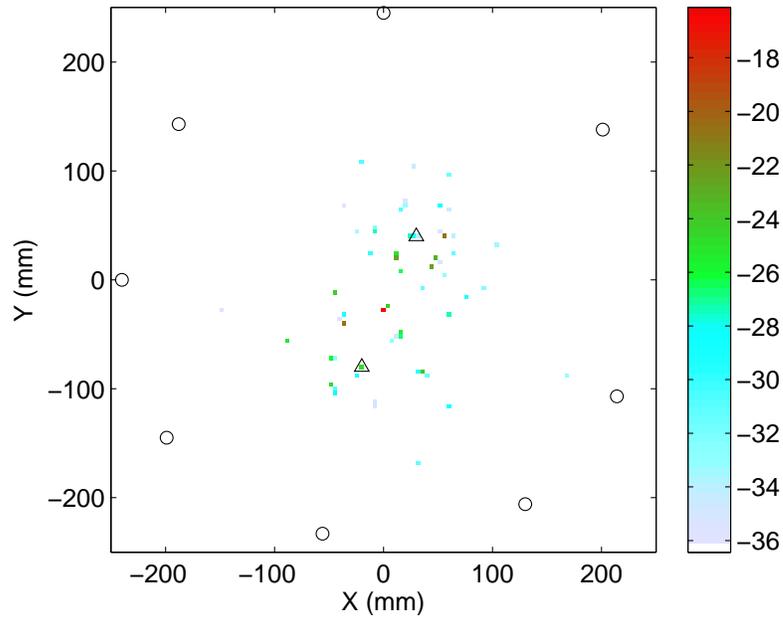
Figure 13: Result for experiment one (glued-on steel rods) using sparse reconstruction with RF signals, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.
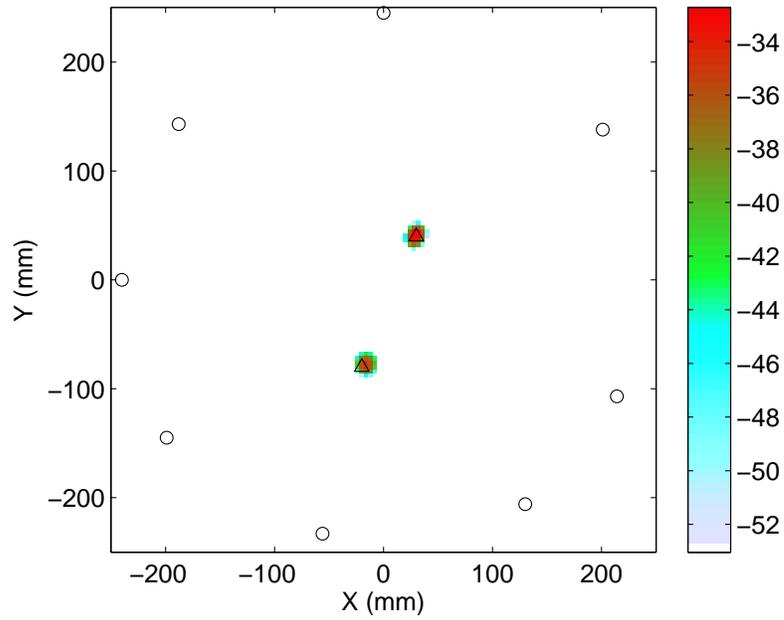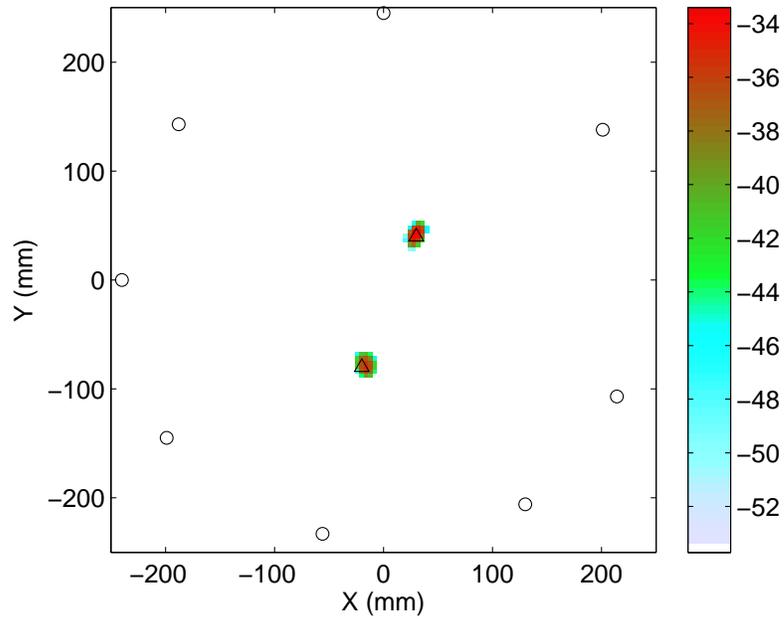

Figure 14: Result for experiment one (glued-on steel rods) using sparse reconstruction with analytic representations of RF signals, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.
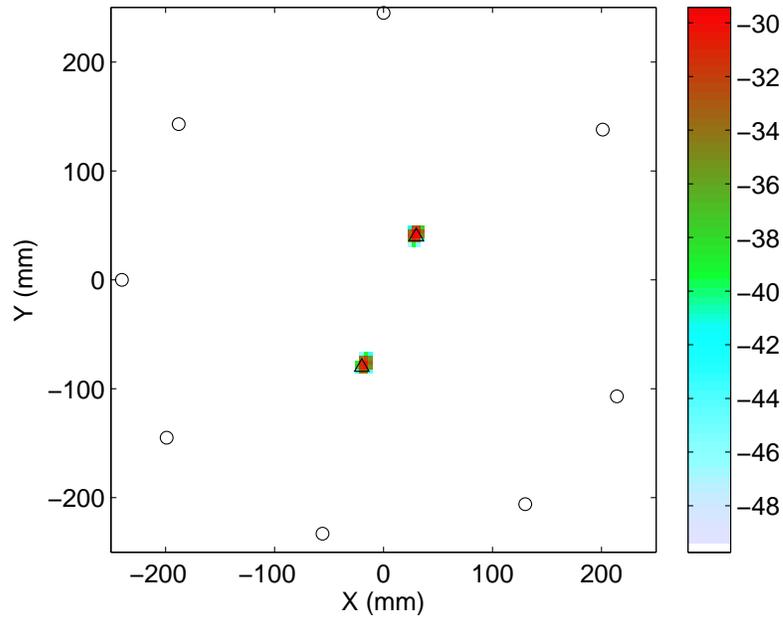
Figure 15: Result for experiment one (glued-on steel rods) using sparse reconstruction with signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.

Figure 16: Result for experiment one (glued-on steel rods) using sparse reconstruction with signal envelopes. Dispersion is not modeled in the dictionary. The image is shown on a 20 dB scale and is normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.

6.1.5 Block-Sparse Reconstruction

Reconstruction was performed using block-sparse reconstruction as described in Chapter 5. The complex-valued, one-dimension-per-pair model, presented in Section 5.2.1, was used to create the dictionary. The energy fraction, $E_\parallel = 0.65$, is similar to the previous non-block cases. Imaging results with $\sigma' = 0.5\sigma'_{\max}$ and without denormalization are shown in Figure 17. The spot sizes of the scatterers are slightly reduced compared to the non-block case with signal envelopes, but the results are otherwise similar.

6.1.6 Block-Sparse Reconstruction with Unknown Dispersion Curves

To compare robustness with the standard sparse reconstruction method, block-sparse reconstruction was repeated with a dictionary that did not model dispersion, in the same manner as Section 6.1.4. Again, $E_\parallel = 0.65$, and results using $\sigma' = 0.5\sigma'_{\max}$ and no denormalization are shown in Figure 18. The increase in spot sizes between the block-based images shown in Figures 17 and 18 is greater than that of the envelope images shown in Figures 15 and 16.

6.1.7 Discussion

A summary of the results from experiment one is shown in Table 5. Sparse reconstruction with real-valued signals failed to locate the scatterers and is not shown. The results for experiment one are very similar for block-sparse reconstruction and sparse reconstruction with signal envelopes. As discussed in Section 5.6.1, signal envelopes perform well when scatterers are far apart because the corresponding scattered signals are separated in time for many of the received signals. Additionally, the steel rods have an approximately omnidirectional scattering behavior, so the dictionary is well-matched to the data.

Figure 17: Result for experiment one (glued-on steel rods) using block-sparse reconstruction, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.
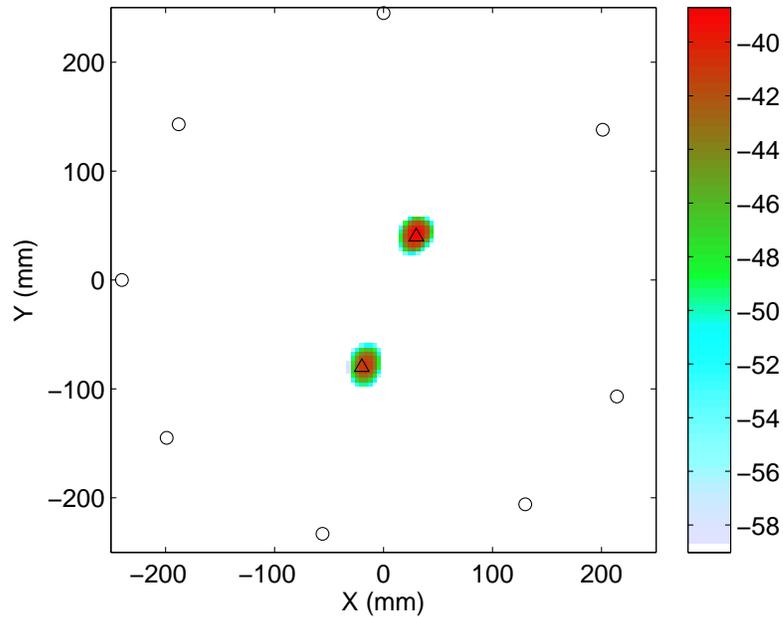

Figure 18: Result for experiment one (glued-on steel rods) using block-sparse reconstruction. Dispersion is not modeled in the dictionary. The image is shown on a 20 dB scale and is normalized to the largest pixel value. Circles denote transducer locations, and triangles denote the locations of the rods.

Table 5: Summarized results of experiment one (glued-on steel rods)

	Dispersion Curves			Group Velocity		
	RF Analytic	Envelope[a]	Block	Envelope-GV[a]	Block-GV	DAS
Projection energy fraction	0.64	0.65	0.65	0.65	0.65	N/A
Spot size[b] of top scatterer	32 mm^2	112 mm^2	80 mm^2	144 mm^2	272 mm^2	N/A[c]
Spot size[b] of bottom scatterer	16 mm^2	144 mm^2	80 mm^2	176 mm^2	304 mm^2	N/A[c]
Amplitude of largest artifact	-8.10 dB	N/A[d]	N/A[d]	N/A[d]	N/A[d]	N/A[e]

[a] With noise compensation.
[b] Total area of connected pixels within 3 dB of the scatterer's maximum pixel value.
[c] The 3 dB spots of both scatterers are merged together.
[d] No artifacts present at the 20 dB level.
[e] The large spot size of the scatterers obscures any potential artifacts.

6.2   Experiment Two

The second experiment was performed on the same plate as experiment one. The glued-on rods were removed and a 9.9 mm notch was hand-cut at a 58° orientation relative to the X-axis, using a drilled starter hole. The baseline signals were acquired after drilling the starter hole but before cutting the notch. The plate dimensions, transducer coordinates, and acquisition procedure were identical to the description in Section 6.1. Of important note is that all dictionaries used for reconstruction in experiment two are identical to those in the first experiment; the algorithms were simply run with different data files to determine their robustness to different types and locations of damage.

6.2.1   Delay-and-Sum

As in experiment one, a preliminary delay-and-sum image was generated as a baseline to which sparsity-based methods could be compared. Figure 19 shows the results using the same software as experiment one. Again, the spot size of the scatterer is quite large (on the order of 100 mm along its major axis) and a sizable portion of the ROI is within 10 dB of the maximum.

6.2.2   Sparse Reconstruction with Raw Signals

Results with real-valued signals were shown to be extremely poor in Section 6.1.2 and were not repeated for this experiment. Using analytic signal representations, the projection energy fraction is $E_\parallel = 0.67$. Imaging results using $\sigma' = 0.5\sigma'_{\max}$ and no denormalization are shown in Figure 20. The scatterer is well localized; however, several artifacts are present due to phase mismatches, in a similar fashion as experiment one.
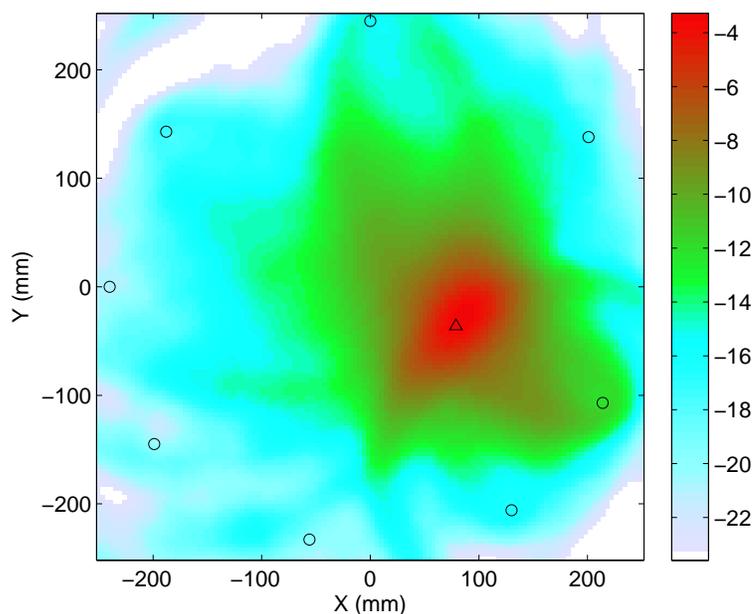
Figure 19: Result for experiment two (hand-cut notch) using delay-and-sum imaging on signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and the triangle denotes the location of the notch.
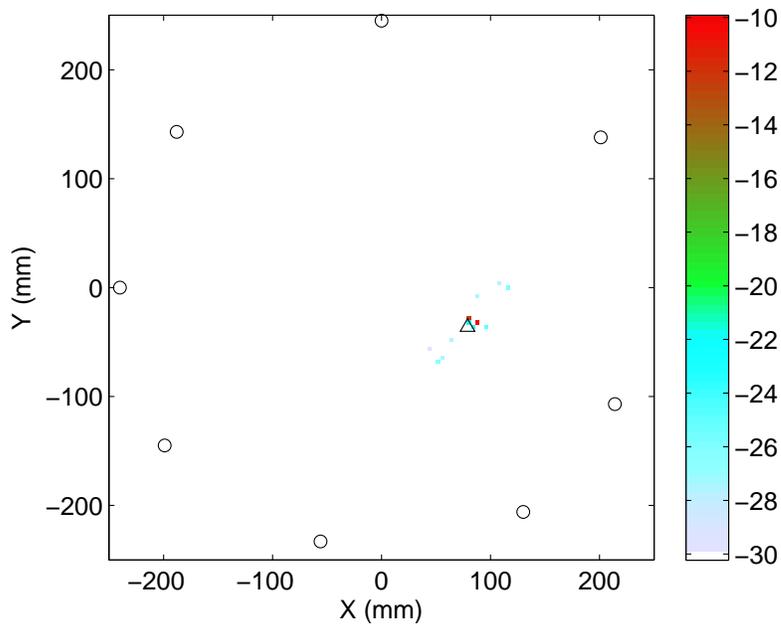

Figure 20: Result for experiment two (hand-cut notch) using sparse reconstruction with analytic representations of RF signals, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and the triangle denotes the location of the notch.

6.2.3 Sparse Reconstruction with Signal Envelopes

Reconstruction was then performed using signal envelopes, using the same dictionary as Section 6.1.3. Noise compensation was performed on the residual signals, which somewhat decreased the energy fraction to $E_{\parallel} = 0.59$; without noise compensation, the energy fraction is comparable to that of the dictionary of raw signals. Imaging results with $\sigma' = 0.5\sigma'_{\max}$ and with no denormalization are shown in Figure 21. The scatterer is clearly detected, but the spot appears elongated. Interestingly, the angle of the spot in the image is similar to that of the notch. It is unknown if this is a coincidence or is related to the 58° orientation. The small artifacts over transducers 5 and 7 are due to amplitude mismatch in the dictionary; the scatterer is highly directional, but the dictionary is modeled with an omnidirectional scatterer. Despite the artifacts, the use of signal envelopes successfully mitigates phase mismatch.

6.2.4 Sparse Reconstruction with Signal Envelopes with Unknown Dispersion Curves

As in experiment one, it was desired to determine robustness by using a dictionary that lacked modeling of dispersion by using only the nominal group velocity. The energy fraction is similar, at $E_{\parallel} = 0.60$ for the noise-compensated signals. Imaging results are shown, again using the noise-compensated signals with $\sigma' = 0.5\sigma'_{\max}$ and no denormalization, in Figure 22; the spot size has lengthened significantly. Again, artifacts are present at transducer locations due to model mismatch.

6.2.5 Block-Sparse Reconstruction

Next, reconstruction was performed using block-sparse reconstruction, using the same dictionary as Section 6.1.5. For this dictionary, $E_{\parallel} = 0.68$; imaging results with $\sigma' = 0.5\sigma'_{\max}$ and without denormalization are shown in Figure 23. The spot size of the notch is significantly smaller than for the envelope-based method; the flexibility of the multidimensional model eliminates the artifacts of that method, as well.
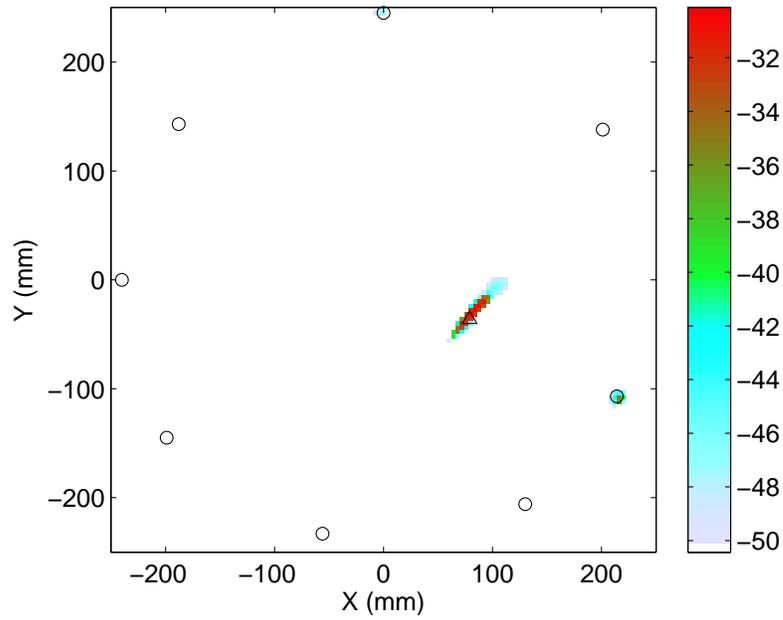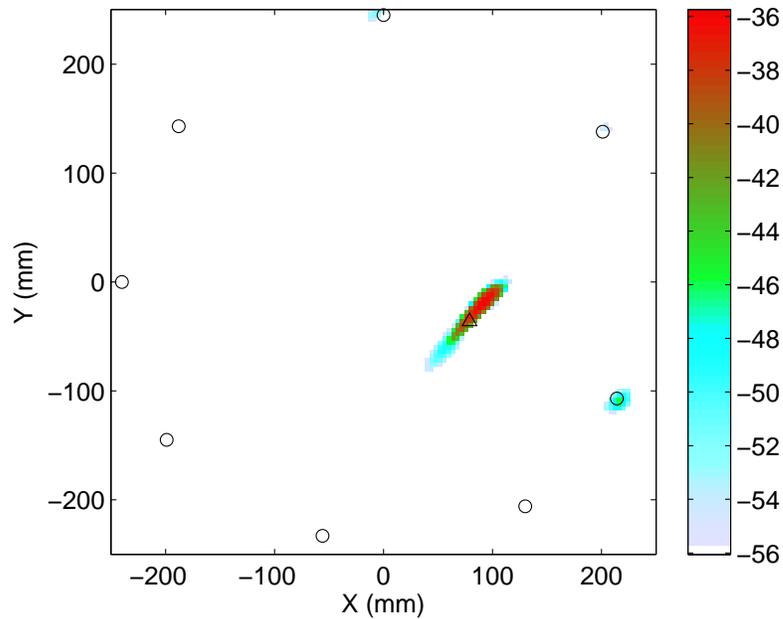
Figure 21: Result for experiment two (hand-cut notch) using sparse reconstruction with signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and the triangle denotes the location of the notch.
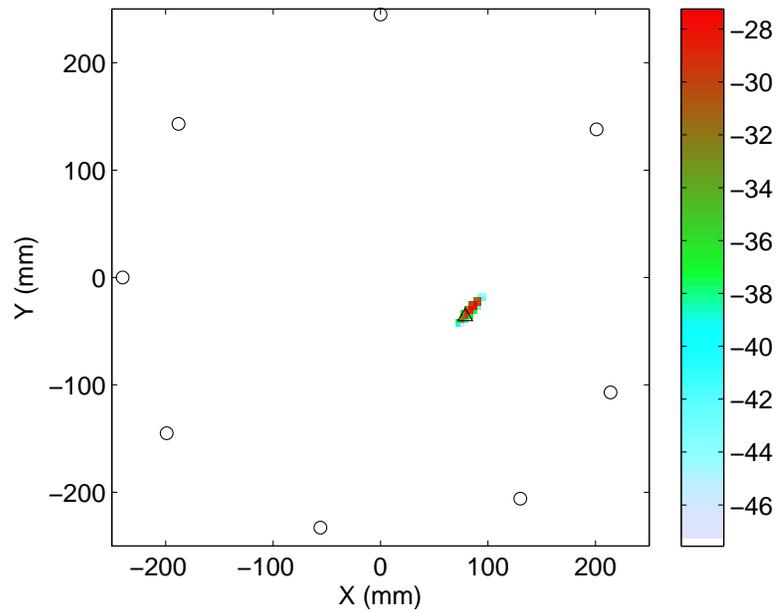

Figure 22: Result for experiment two (hand-cut notch) using sparse reconstruction with signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Dispersion is not modeled in the dictionary. Circles denote transducer locations, and the triangle denotes the location of the notch.

6.2.6 Block-Sparse Reconstruction with Unknown Dispersion Curves

Block-sparse reconstruction was repeated with a dictionary that did not model dispersion, to compare robustness with the envelope method; using this dictionary, $E_\parallel = 0.68$. The imaging result, which was generated using $\sigma' = 0.5\sigma'_{\max}$ and with no denormalization, is shown in Figure 24. In experiment one, the use of group velocities instead of dispersion curves affected the block-sparse method more than the non-block method with signal envelopes; in this experiment, the quality of the block-sparse image is only slightly reduced and the method performs better with this simplification than its non-block, envelope-based counterpart.

6.2.7 Discussion

A summary of the results of experiment two is shown in Table 6. Unlike experiment one, the results for experiment two are improved to some degree by the use of block-sparse reconstruction over the envelope method. This is primarily due to the flexibility of the former to handle directional scatterers without requiring a specific model. Nevertheless, the envelope-detected dictionary method is still capable of locating the hand-cut notch.

Figure 23: Result for experiment two (hand-cut notch) using block-sparse reconstruction, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and the triangle denotes the location of the notch.
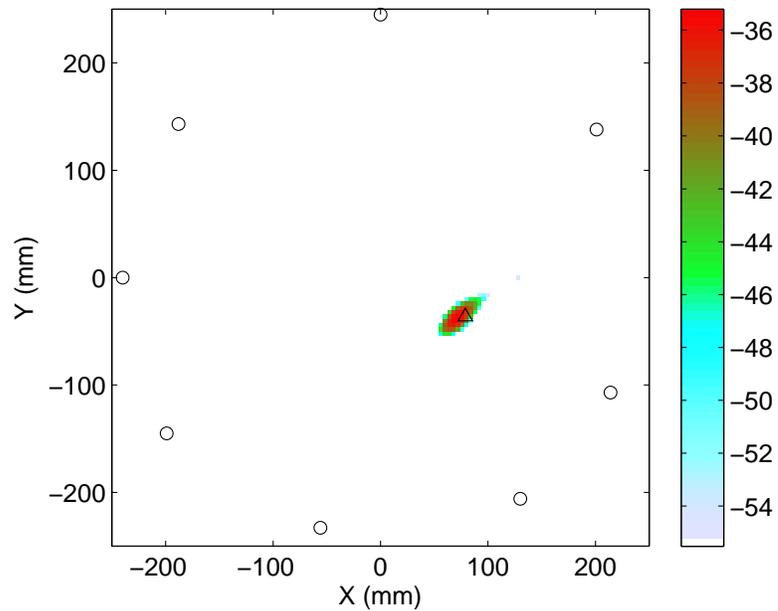

Figure 24: Result for experiment two (hand-cut notch) using block-sparse reconstruction, shown on a 20 dB scale, normalized to the largest pixel value. Dispersion is not modeled in the dictionary. Circles denote transducer locations, and the triangle denotes the location of the notch.

Table 6: Summarized results of experiment two (hand-cut notch)

	Dispersion Curves			Group Velocity		
	RF Analytic	Envelope[a]	Block	Envelope-GV[a]	Block-GV	DAS
Projection energy fraction	0.67	0.59	0.68	0.60	0.68	N/A
Scatterer spot size[b]	32^c mm^2	208 mm^2	96 mm^2	368 mm^2	256 mm^2	4992 mm^2
Amplitude of largest artifact	-16.07 dB	-5.44 dB	N/A[d]	-9.99 dB	-18.84 dB	N/A[e]

[a] With noise compensation.
[b] Total area of connected pixels within 3 dB of the scatterer's maximum pixel value.
[c] Includes a nearby, but non-adjacent, pixel.
[d] No artifacts present at the 20 dB level.
[e] The large spot size of the scatterers obscures any potential artifacts.

6.3　Experiment Three

The third experiment was conducted on a smaller plate to evaluate the effects of edge reflections. Six 0.5 mm thick, 7 mm diameter transducers were affixed to an Aluminum-6061 plate with dimensions 292 mm × 600 mm × 3.175 mm. The locations of the transducers are shown in Table 7. A 3.6 mm notch was hand-cut at a $-50°$ orientation from a 5.0 mm drilled starter hole at location $(-50.5\text{ mm}, -51.0\text{ mm})$ after baseline measurements were obtained. The acquisition setup was identical to the previous experiments (Figure 11), with the exception of the number of attached transducers. The square region $-146\text{ mm} \leq x, y \leq 146\text{ mm}$ was designated as the region of interest; unlike the previous two experiments, the ROI touches the left and right sides of the plate (but not the top and bottom).

For all imaging methods, the high levels of clutter in the signals (unmodeled boundary reflections) result in large artifacts at the edges of the ROI. Unless stated otherwise, all images in this section are normalized to the largest pixel that is near the scatterer's true location. This is to emphasize the relative magnitudes of artifacts relative to the actual defect. Artifacts that are above the maximum of the 20 dB scale are shown in dark brown.

Table 7: Transducer placement for experiment three

Transducer	X (mm)	Y (mm)
#1	-94.5	-41.0
#2	0.5	-108.0
#3	94.5	-63.0
#4	83.0	66.0
#5	0.5	110.0
#6	-89.0	47.0

6.3.1 Delay-and-Sum

As with previous experiments, the first image to be generated was a DAS image for comparison. Figure 25 shows the results. The scatterer is detected, but large artifacts fill the plate due to the numerous boundary reflections that are either imperfectly subtracted from their baselines or are secondary bounces due to the scatterer. The overall dynamic range of the image is 8.51 dB.
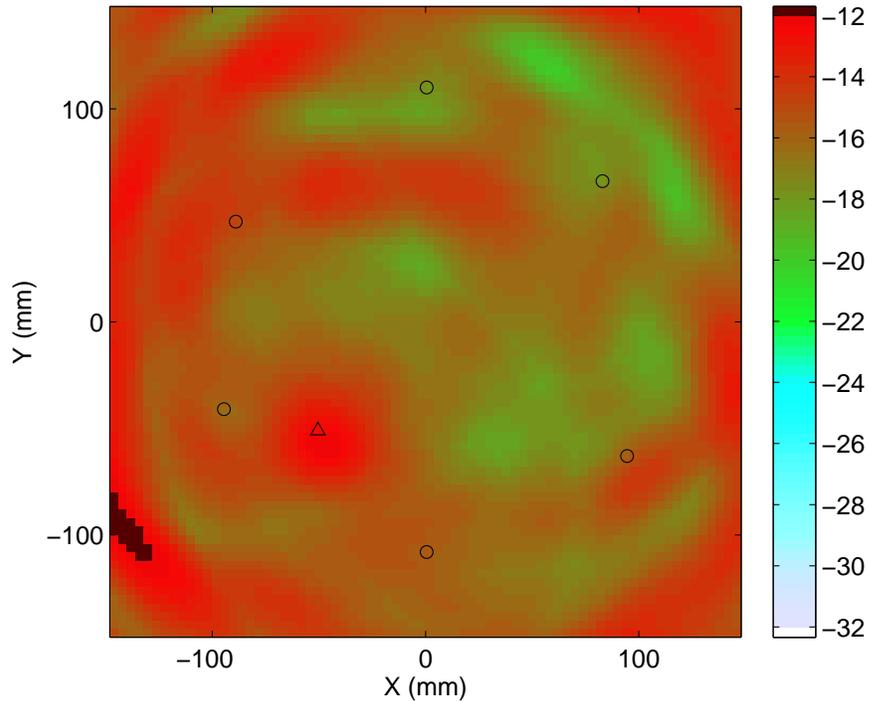

Figure 25: Result for experiment three (small plate) using delay-and-sum imaging on signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.

6.3.2 Sparse Reconstruction with Raw Signals

Due to the smaller size of the plate and imperfect baseline subtraction, a high degree of interference is present in the residual signals. As a result, the projection energy fraction is quite small: using the non-block, complex analytic form resulted in $E_\parallel =$
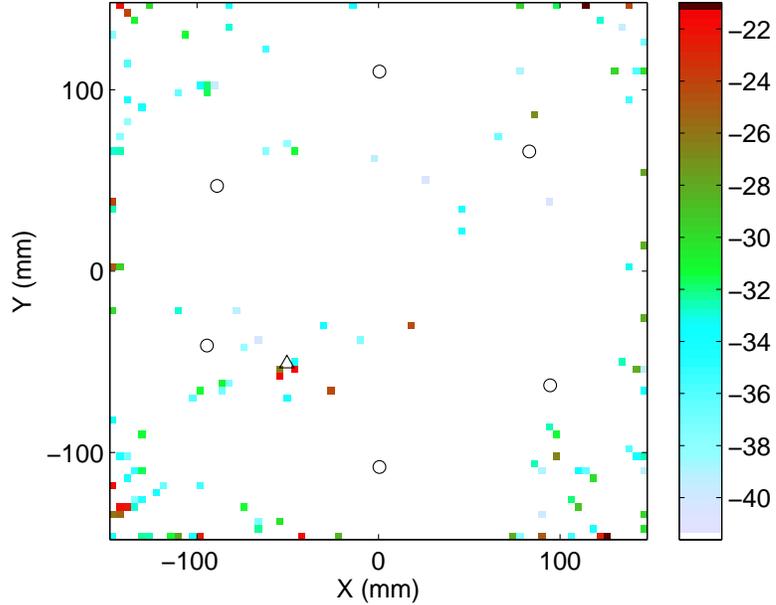
Figure 26: Result for experiment three (small plate) using sparse reconstruction with analytic representations of RF signals, shown on a 20 dB scale, normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.

0.073. Figure 26 shows results using this dictionary with $\sigma' = 0.5\sigma'_{\max}$. A few pixels near the scatterer are high-valued, but even the largest of these pixels is less than 3 dB above some of the artifacts in the image. Due to the signal interference, many pixels at the ROI corners and edges are higher-valued than the scatterer by 2 dB or more.

6.3.3 Sparse Reconstruction with Signal Envelopes

Next, reconstruction was performed using signal envelopes. Because of the high level of clutter in the signals, estimates of the noise level were unreliable, and noise compensation was not used. The energy fraction resulting from the envelope dictionary, $E_{\parallel} = 0.061$, is somewhat smaller than the corresponding result using raw signals. Results using $\sigma' = 0.5\sigma'_{\max}$ are shown in Figure 27; the image is an overall improvement from the image generated with raw signals. Several pixels are lit in the immediate vicinity of the scatterer, though many of these artifacts are present in locations that

correspond to the DAS image shown in Figure 25. The corner pixels have much higher values than with raw signals; the pixel in the bottom-left corner is 16 dB above the level of the largest pixel near the scatterer. The artifact in the upper-center portion of the plate (within the convex hull of the transducers) is more than 6 dB below the pixel value of the scatterer. Another artifact is present near transducer #1; this artifact also appears in the DAS image and is likely exacerbated by the unmodeled directionality of the scatterer.

6.3.4 Sparse Reconstruction with Signal Envelopes with Unknown Dispersion Curves

Results were then generated using an envelope-detected dictionary using only the group velocity instead of dispersion curves. The energy fraction in the dictionary's column space is similar, at $E_\parallel = 0.057$. Results are shown, with $\sigma' = 0.5\sigma'_{\max}$, in Figure 28. The spot sizes of the scatterer, central artifacts, and corner artifacts are significantly increased. The top-center artifact is only approximately 5 dB below the pixel intensity of the scatterer, and the amplitudes of some corner artifacts are more than 15 dB above that of the scatterer.

6.3.5 Block-Sparse Reconstruction

Next, reconstruction was performed using block-sparse reconstruction. The energy fraction due to the block dictionary is similar to that of the non-block dictionary with raw signals; $E_\parallel = 0.077$. Figure 29 shows the results when $\sigma' = 0.5\sigma'_{\max}$. While the spot size of the scatterer and artifacts is much larger than those of the corresponding image using envelope-detected signals, the artifact level is greatly reduced due to the flexibility of the block model: the largest artifact within the convex hull of the transducer locations is nearly 9 dB below the level of the scatterer, and the largest corner artifact is only 7.7 dB above the scatterer.

6.3.6 Block-Sparse Reconstruction with Unknown Dispersion Curves

Block-sparse reconstruction was again repeated with a dictionary that did not model dispersion, to compare to the the envelope method. The energy fraction is essentially unchanged, at $E_\parallel = 0.077$. Figure 30 shows the results using $\sigma' = 0.5\sigma'_{max}$. Image quality is similar to that of the dispersive dictionary; the largest artifact within the transducer aperture is 11 dB below the scatterer intensity, but the largest corner artifact is nearly the same amount above the scatterer (10.9 dB).

6.3.7 Discussion

Table 8 summarizes the performance of the various imaging methods. The comparison between signal envelopes and block-sparsity is an interesting one for this experiment. For the case of the dispersive dictionary, the envelope method exhibits superior results to the block method in terms of spot size, though the artifacts are somewhat higher in magnitude. In some sense, this result shows the strengths and weaknesses of the block-sparse algorithm: it has a more robust scatterer model, which decreases artifact magnitudes due to mismatch; however, the flexibility of the model leads to nearby pixels sharing many of the dimensions of their respective subspaces, which increases spot sizes.

When using group velocity instead of dispersion curves, the block-sparse method does better overall: with the exception of the artifact at $(-80\text{ mm}, 120\text{ mm})$, which is well outside the transducer aperture, all artifacts are decreased in magnitude compared to the envelope method.

The projection energy fractions for this experiment are much smaller than previous experiments due to the presence of unmodeled edge reflections. The differential signal envelopes for the first transducer pair in experiments two and three are compared in Figure 31. The signals from experiment two are relatively clean, with a large peak from the notch, as well as small peaks due to imperfect baseline subtraction. In
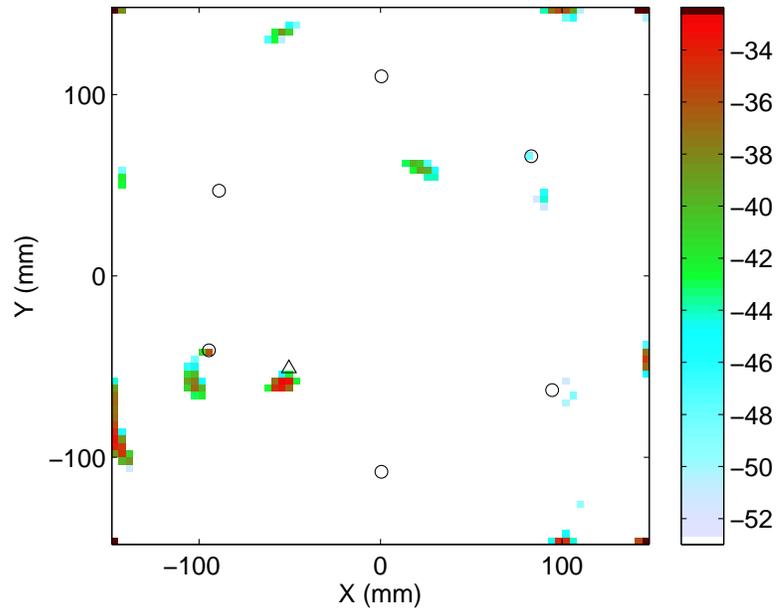
Figure 27: Result for experiment three (small plate) using sparse reconstruction with signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.
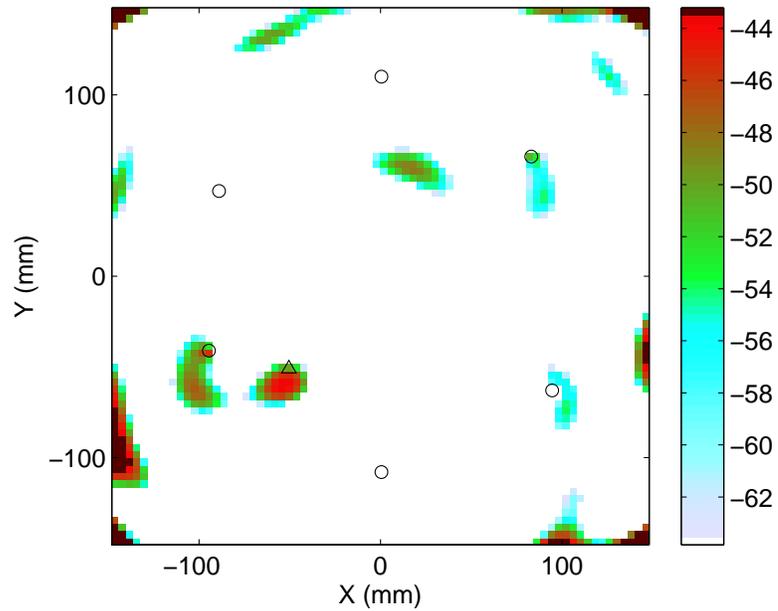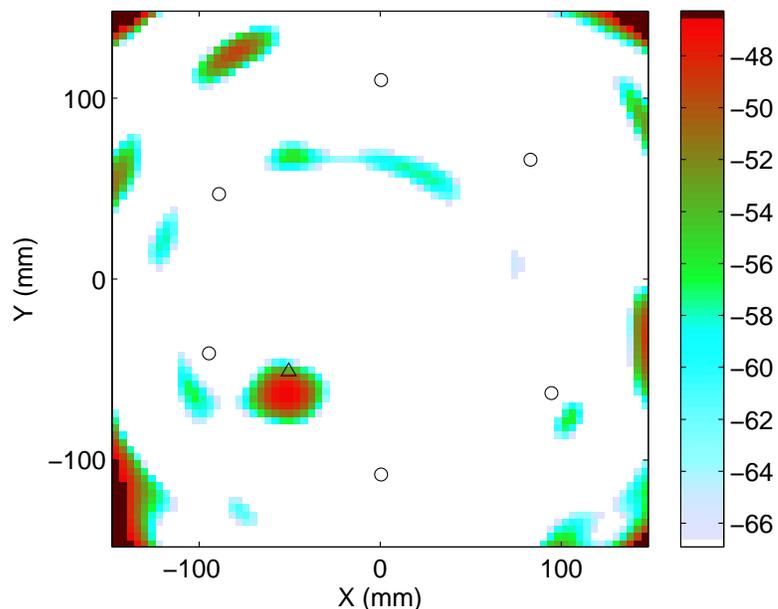

Figure 28: Result for experiment three (small plate) using sparse reconstruction with signal envelopes, where dispersion is not modeled in the dictionary. The image is shown on a 20 dB scale and normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.

Figure 29: Result for experiment three (small plate) using block-sparse reconstruction, shown on a 20 dB scale, normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.
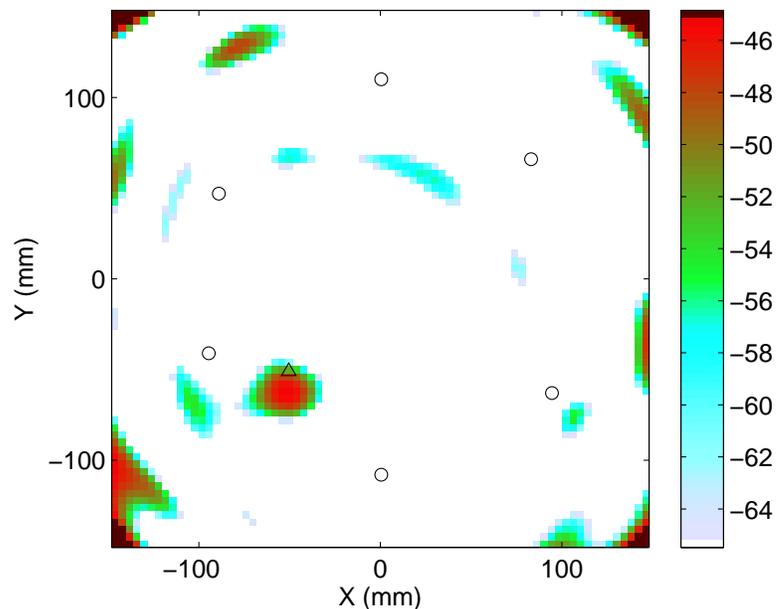

Figure 30: Result for experiment three (small plate) using block-sparse reconstruction, where dispersion is not modeled in the dictionary. The image is shown on a 20 dB scale and normalized to the largest pixel value near the scatterer. Circles denote transducer locations, and the triangle denotes the location of the notch.

contrast, the envelope from experiment three has large regions of overlapping echoes; it is impossible to discern at first glance which part of the signal is due to the scatterer. (In fact, the scatterer is responsible for the small wavepacket with the peak at 75 µs.) Both signals are truncated and amplitude-scaled to fit on the same axes.

One particularly noticeable phenomenon is the bright artifacts at the corners of the plate. As explained in Section 4.3.2.3, these can result from the solver attempting to match later portions of the signal. Additionally, if baseline subtraction is imperfect, there will be some residual due to signals bouncing at plate boundaries. These edge reflections appear as scatterers, and an algorithm has no way of differentiating this sort of behavior from that of an actual scatterer. The DAS image does not suffer from the specific issue of "corner lighting"; however, the DAS image is much worse overall.

The performance of imaging in an actual structure is likely to fall somewhere between experiments two and three. A realistic structure for this application would not have free edges like the small plate of experiment three, but instead would have more clutter at smaller amplitudes due to geometrical reflectors such as fastener holes, stiffeners, and welds.

Table 8: Summarized results of experiment three (small plate)

	Dispersion Curves			Group Velocity		
	RF Analytic	Envelope[a]	Block	Envelope-GV[a]	Block-GV	DAS
Projection energy fraction	0.073	0.061	0.077	0.057	0.077	N/A
Scatterer spot size[b]	64[c] mm^2	64 mm^2	352 mm^2	192 mm^2	272 mm^2	3216[d] mm^2
Largest artifact (entire plate)	2.18 dB	16.61 dB	7.71 dB	15.53 dB	10.85 dB	0.41 dB
Largest artifact (excluding edges)[e]	−2.87 dB	−3.28 dB	−2.83 dB	−1.92 dB	−2.71 dB	−0.89 dB
Largest artifact (within aperture)[f]	−2.87 dB	−6.47 dB	−8.93 dB	−4.95 dB	−11.31 dB	−1.68 dB

a Without noise compensation.
b Total area of connected pixels within 3 dB of the scatterer's maximum pixel value.
c Includes nearby, but non-adjacent, pixels.
d Does not include area of an adjacent artifact.
e Excludes any artifact whose spot size includes a pixel at the plate edge.
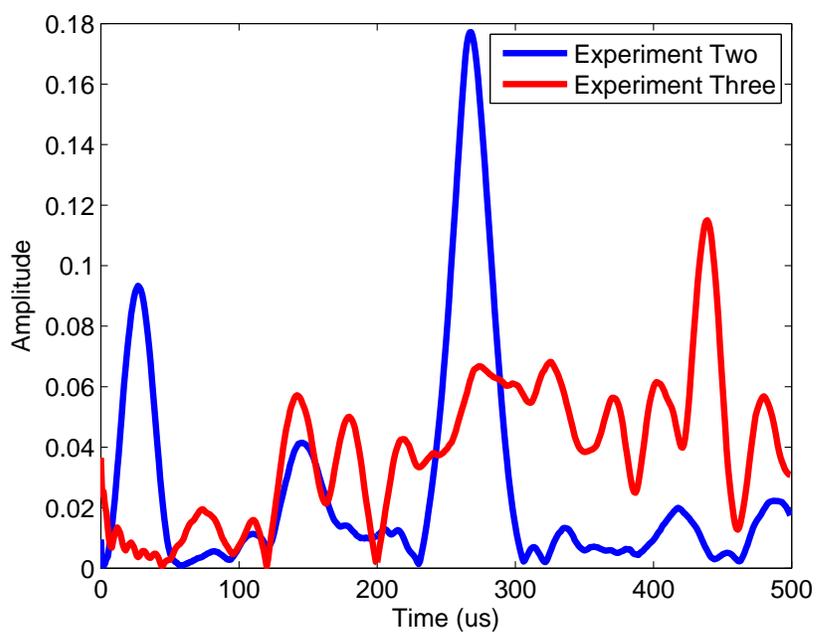f Includes any artifact that is more than 5 mm inside the convex hull of the transducer locations.

Figure 31: Comparison of differential signal envelopes for one pair in experiments two and three. The signals are truncated at 500 µs and normalized to fit on the same scale.

6.4 Experiment Four

The final experiment was performed on an 8-layer cross-ply carbon fiber composite plate with dimensions 460 mm × 460 mm × 2.5 mm. Six 0.5 mm thick, 7 mm diameter transducers were affixed at the locations shown in Table 9. Transducer #6 was later found to be faulty and was not used, reducing the total number of transducer pairs to ten.

Dispersion curves were unavailable for this material; instead, the group velocity was computed from experimental data using the slant-stack Radon transform [118], with no applied time offset. Damage was simulated with small rare earth magnets that were affixed to both sides of the plate. The acquisition system was identical to previous experiments (Figure 11), with the exception of the number of attached transducers. The entirety of the plate, -230 mm $\leq x, y \leq 230$ mm was designated as the area of interest. After exciting and postprocessing with a broadband pulse [117], it was found that the wavefield is nearly single-mode at 75 kHz for a pseudo-A_0 mode that is almost isotropic; this frequency was selected as the center frequency for the excitation, which was selected to be a 5-cycle, Hann-windowed toneburst. The dictionaries again had a resolution of 4 mm. Even though the plate was relatively small, the greater material attenuation in the composite compared to aluminum helped mitigate the effect of unmodeled edge reflections. Sparse reconstruction with raw signals performed similarly to experiment three and is not shown.

Table 9: Transducer placement for experiment four

Transducer	X (mm)	Y (mm)
#1	42.5	142.0
#2	-118.5	105.0
#3	-165.0	-24.0
#4	-56.0	-142.5
#5	97.0	-125.0
#6 (defective)	144.0	-6.5

6.4.1 Delay-and-Sum

Once again, a DAS image was generated for comparison; results are shown in Figure 32. The scatterer has a large spot size and two high-value artifacts above and below it, to the right. A very noticeable artifact in the shape of an ellipse is present with foci at the two rightmost transducers; this is likely caused by imperfect baseline subtraction for that pair for the portions of the signals that correspond to at least one edge reflection.
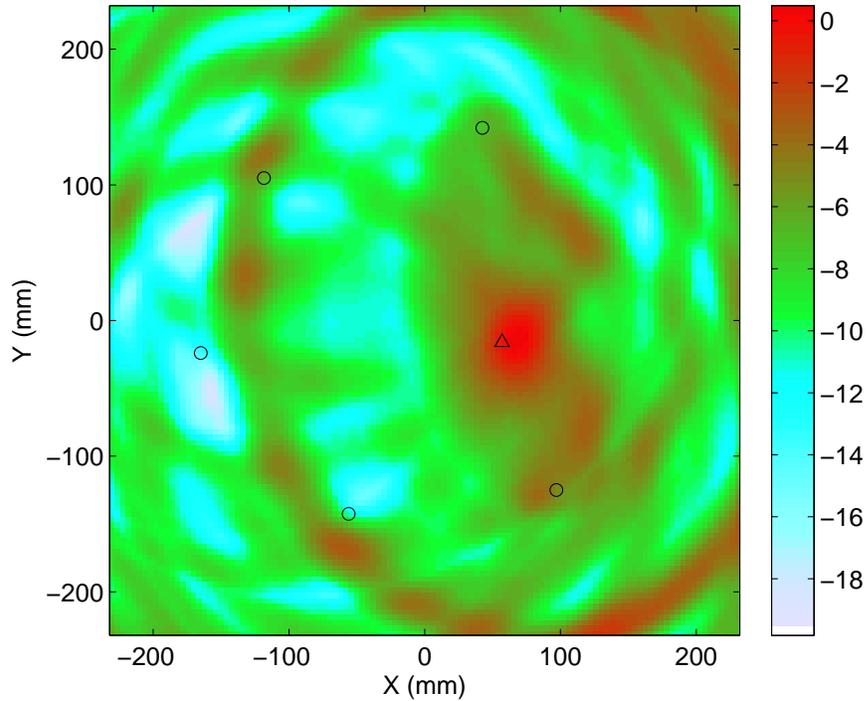

Figure 32: Result for experiment four (composite plate) using delay-and-sum imaging on signal envelopes, shown on a 20 dB scale, normalized to the largest pixel value. Circles denote transducer locations, and the triangle denotes the location of the magnets.

6.4.2 Sparse Reconstruction with Signal Envelopes with Unknown Dispersion Curves

The projection energy an envelope-detected dictionary on signals without noise compensation was $E_\parallel = 0.61$. Imaging results using $\sigma' = 0.5\sigma'_{\max}$ are shown in Figure 33.

97

The scatterer is well localized, with only a small number of artifacts away from plate edges. The largest such artifact is 6.76 dB below the pixel intensity of the scatterer. The artifact at the top edge of the plate on the right, which lies on the ellipse of artifacts in Figure 32, is 1.52 dB above the magnitude of the scatterer.

For comparison, when noise-compensated signals are used, $E_\parallel = 0.65$, and results are shown in Figure 34 for $\sigma' = 0.5\sigma'_{\max}$. The artifact at (118 mm, 58 mm) is 4.79 dB below the scatterer, and the largest artifact at the plate edge is 4.06 dB above the scatterer. Although the spot sizes are reduced, the noise compensation has adversely affected the image in terms of artifact level, possibly due to poor estimation of the true noise level.

6.4.3   Block-Sparse Reconstruction with Unknown Dispersion Curves

An image was also generated using block-sparse reconstruction, where $E_\parallel = 0.68$ using the block dictionary. Figure 35 shows the results using $\sigma' = 0.5\sigma'_{\max}$. The image is nearly identical to Figure 33, though the amplitude of the scatterer is somewhat increased compared to that of the artifacts, with the largest internal artifact 6.91 dB below the scatterer, and the largest artifact at an edge 0.19 dB above the scatterer.

6.4.4   Discussion

Results of experiment four are summarized in Table 10. In this experiment, only results using group velocity were available, and very little difference is observed between envelope and block-sparse methods, with artifacts in the same locations and similar spot sizes. Because of the small size of the plate and only a rough estimate of propagation behavior, a small number of artifacts are apparent outside of the aperture. Additionally, artifacts are present at the plate boundary; these types of artifacts are also present in experiment three and are due to edge reflections and imperfect baseline subtraction. Interestingly, the corresponding artifacts in the DAS image are smaller in magnitude than those in the sparsity-based images; however, DAS suffers
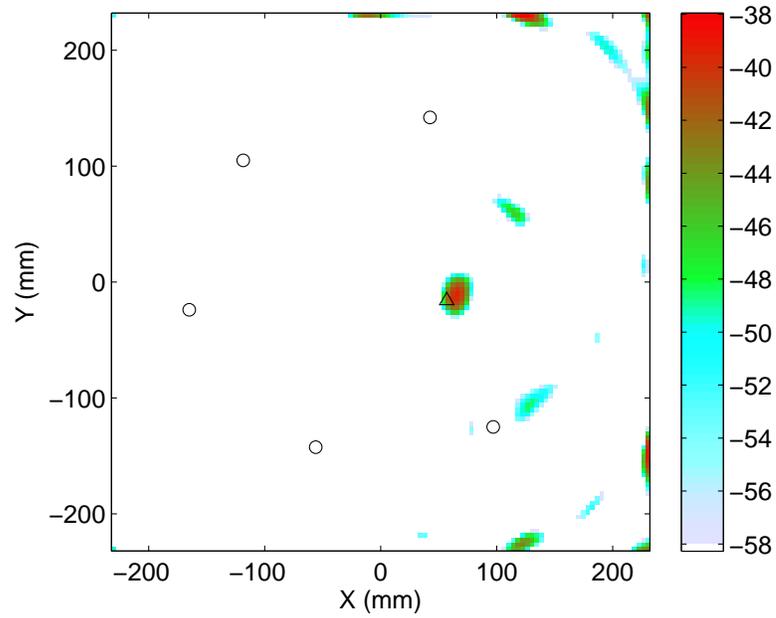
Figure 33: Result for experiment four (composite plate) using sparse reconstruction with signal envelopes that were not compensated for noise, shown on a 20 dB scale, normalized to the largest pixel value. Dispersion is not modeled in the dictionary. Circles denote transducer locations, and the triangle denotes the location of the magnets.
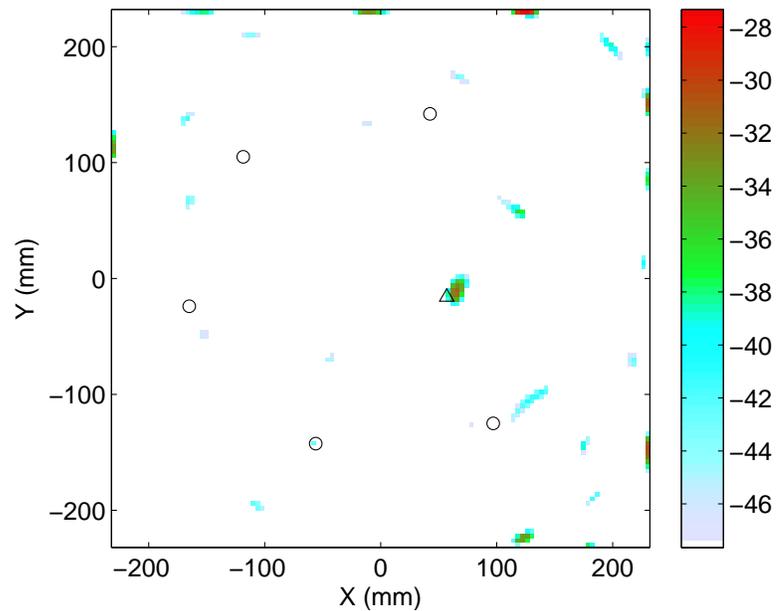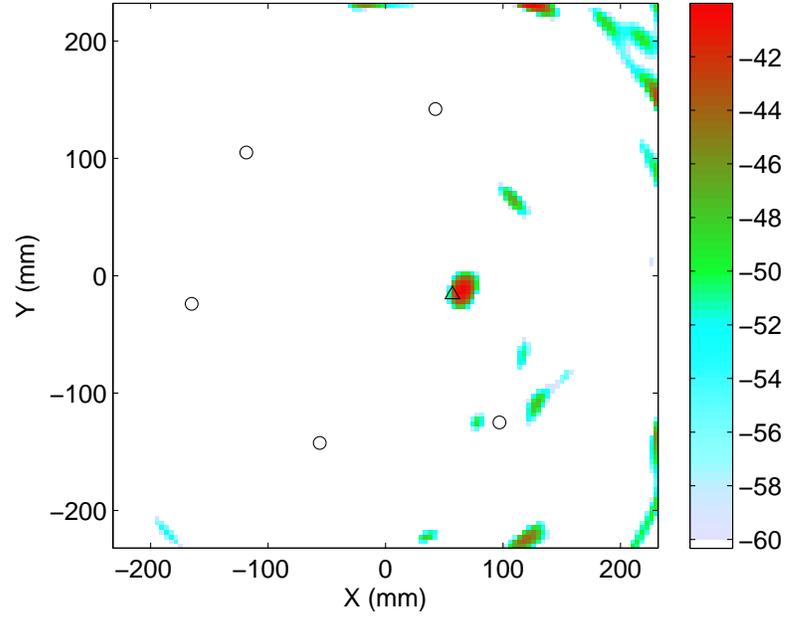

Figure 34: Result for experiment four (composite plate) using sparse reconstruction with signal envelopes and noise compensation, shown on a 20 dB scale, normalized to the largest pixel value. Dispersion is not modeled in the dictionary. Circles denote transducer locations, and the triangle denotes the location of the magnets.

Figure 35: Result for experiment four (composite plate) using block-sparse recon-
struction, shown on a 20 dB scale, normalized to the largest pixel value. Dispersion is
not modeled in the dictionary. Circles denote transducer locations, and the triangle
denotes the location of the magnets.

from multiple large artifacts in the center of the plate.

A surprising result is reduction in quality from noise compensation; unlike ex-
periments one and two, the image is not improved by this operation. The cause is
unclear; it may be due to the relatively low noise levels, which made the true noise
level difficult to estimate. Even this image, however, is superior to the DAS image
overall in terms of the spot size and amplitude of artifacts in the center of the plate.

Table 10: Summarized results of experiment four (composite plate)

	Group Velocity			
	Envelope-GV	Envelope-GV (N.C.[a])	Block-GV	DAS
Projection energy fraction	0.61	0.65	0.68	N/A
Scatterer spot size[b]	256 mm^2	144 mm^2	192 mm^2	2128 mm^2
Largest artifact (entire plate)	1.52 dB	4.06 dB	−0.19 dB	−2.15 dB
Largest artifact (excluding edges)[c]	−6.77 dB	−4.80 dB	−6.91 dB	−3.67 dB

[a] Noise compensated.
[b] Total area of connected pixels within 3 dB of the scatterer's maximum pixel value.
[c] Excludes any artifact whose spot size includes a pixel at the plate edge.

## 6.5 *Effect of Decreasing Numbers of Transducer Pairs*

To analyze the effects of varying numbers of transducers, an analysis of data from experiments one and two was re-run, first using only six of the available transducers, and then using only four, for both the non-block case with signal envelopes and for the block-sparse case. The smaller number of transducers was simulated by zeroing out data: for the non-block method, the portions of the residual signal vector and dictionary corresponding to unused transducer pairs were zeroed; for the block method, only the residual vector needed to be zeroed.
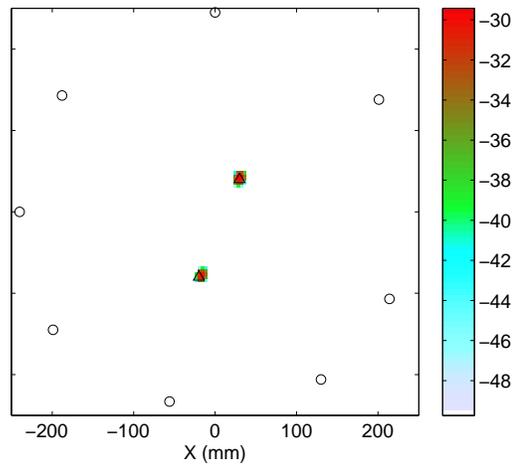
6.5.1 Experiment One

Figure 36 shows results for experiment one with eight, six, and four transducers. For the envelope method, only results without noise compensation are shown; results are similar with noise compensation. With six or eight transducers, both methods are successful at localizing the scatterers. With only four transducers, the envelope method succeeds, but one scatterer's pixel value is reduced; in contrast, the block-sparse method produces a large artifact in the center of the plate. This highlights a weakness of the block-sparse method; as discussed in Section 5.6.2, some of the dimensions in the subspaces that correspond to each pixel are highly correlated. As the number of dimensions decreases, the likelihood of this becoming an issue with reconstruction increases. In the case of this experiment with four transducers, the theoretical residual signal lies in a 12-dimensional subspace (six transducer pairs and two different scatterers). This subspace happens to be highly correlated with the artifact location in five of its six dimensions. Figure 37 shows the cosines of the principal angles (i.e., correlations) between the subspaces $\mathcal{A}_{\text{scatterers}} = \text{span}\left(\begin{bmatrix} \hat{\mathbf{A}}_{(32,40)} & \hat{\mathbf{A}}_{(-20,-80)} \end{bmatrix}\right)$ and $\mathcal{A}_{\text{artifact}} = \text{span}\left(\hat{\mathbf{A}}_{(36,-48)}\right)$; the former represents the two pixel locations that are closest to the scatterers, and the latter is the location of the largest-valued artifact pixel. The figure shows the dimensional correlations when restricted to dimensions

present in the dictionaries generated using eight, six, and four transducers (28, 15, and 6 transducer pairs, respectively). With eight transducers, 12 of the dimensions have a correlation of less than 0.75; with six transducers, 5 dimensions do; with four transducers, only one of the six dimensions has a correlation below 0.9. This explains why the artifact occurs at that particular location; other pixels are not nearly so correlated. For example, the subspace corresponding to the (arbitrarily selected) pixel at $(132, -8)$ is only highly-correlated with $\mathcal{A}_{\text{scatterers}}$ in two of its six dimensions for the four-transducer case.

The phenomenon of correlated artifacts is caused by the transducer configuration. Recall that each transducer pair detects time-of-flight elliptically, where the transducers are the foci of a set of ellipses that correspond to increasing propagation times. For the roughly square configuration of the four transducers selected in Figure 36(f), most of the ellipses that contain the two sites of damage cross the central area of the plate. With a large number of pairs, this is not a major issue, since there are enough uncorrelated dimensions to allow detection. However, with only six pairs total, this configuration can lead to artifacts if flaws are in the center of the aperture. This can be addressed by varying the transducer geometry; for example, Figure 38 shows results with a different configuration of four transducers. No artifacts are present at the 20 dB level.
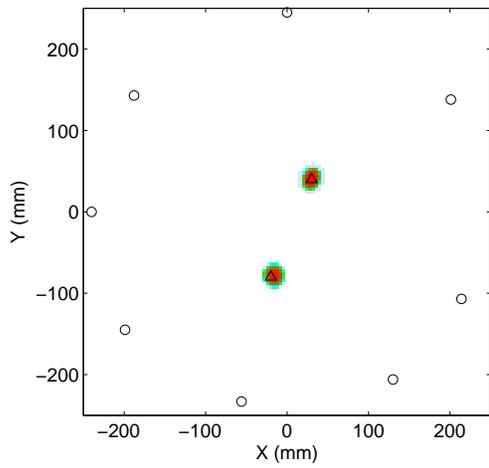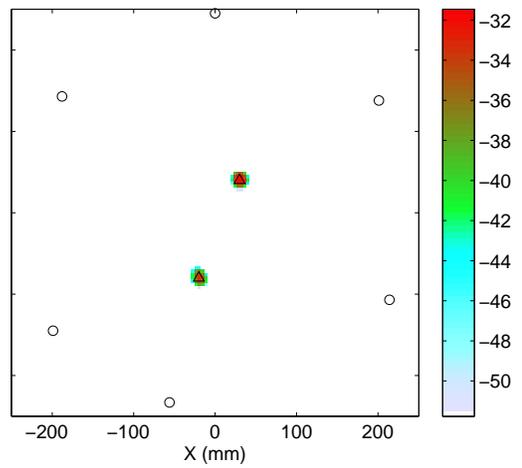
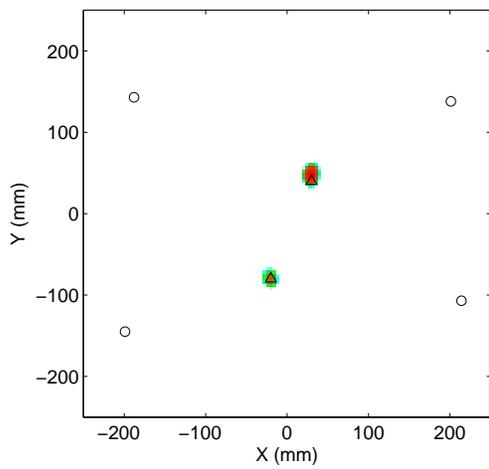(a) Envelope method, 8 transducers (28 pairs)   (b) Block method, 8 transducers (28 pairs)
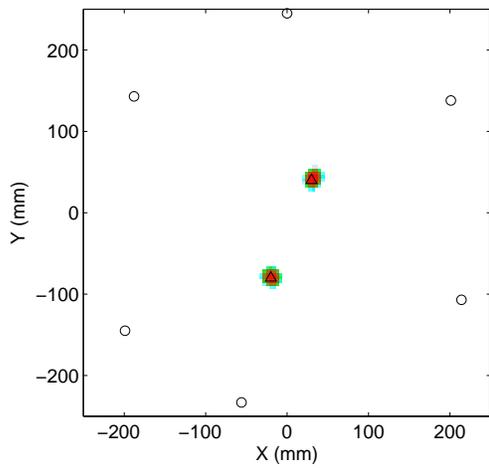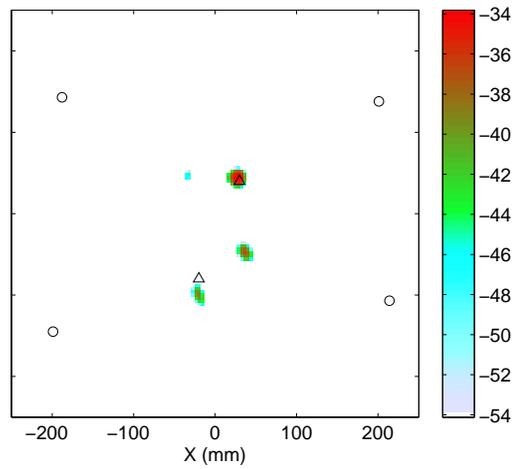
(c) Envelope method, 6 transducers (15 pairs)   (d) Block method, 6 transducers (15 pairs)

(e) Envelope method, 4 transducers (6 pairs)   (f) Block method, 4 transducers (6 pairs)

Figure 36: Results for experiment one with varying numbers of transducers. Circles denote active transducer locations and triangles denote locations of the rods.
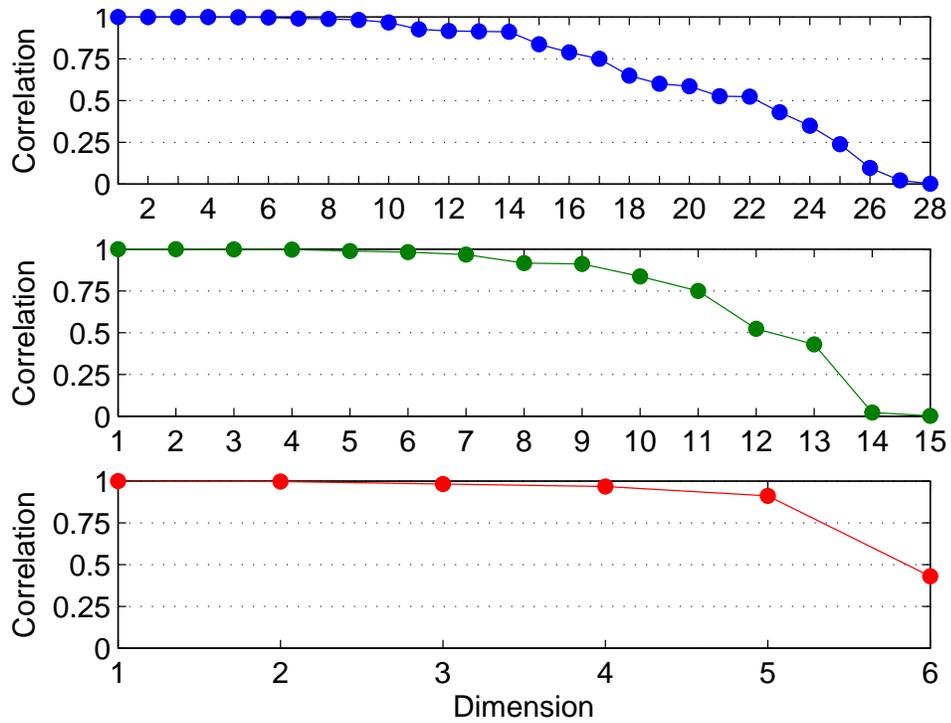
Figure 37: Principal correlations between the dictionary blocks corresponding to scatterer locations and the dictionary block corresponding to the artifact in Figure 36(f), for eight (top), six (middle), and four (bottom) transducer pairs.
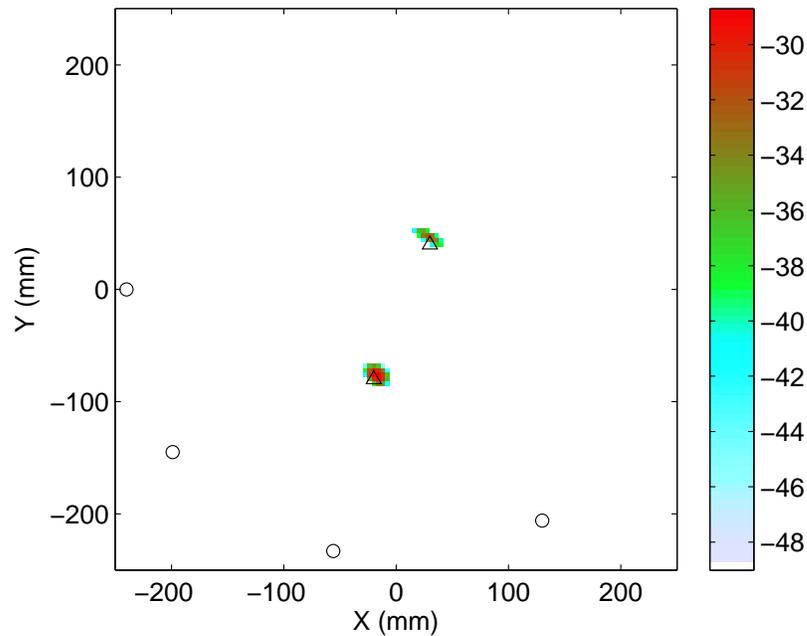

Figure 38: Results for experiment one for block-sparse imaging with four transducers in an alternate configuration.
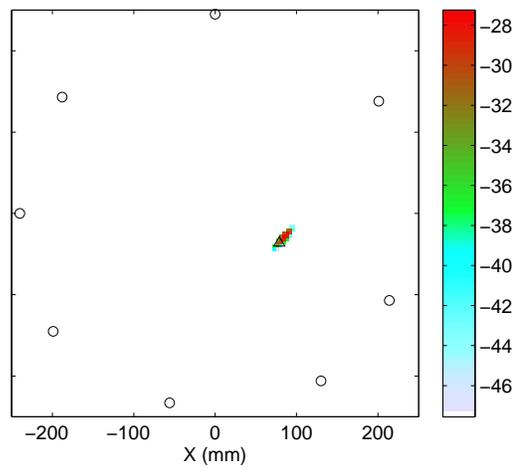
6.5.2 Experiment Two

Figure 39 shows results for experiment two with eight, six, and four transducers, using the same procedure as was used for the previous section. Since there is only one scatterer present, the block-sparse method does not suffer the same problems with artifacts as it did in the previous section. On the other hand, sparse reconstruction with signal envelopes suffers a modest reduction in image quality as the number of transducers is reduced, due to dictionary mismatch (since the scatterer is directional). For the four-transducer case, two high-valued artifacts are located near the two transducers that excite waves that hit the notch's broadside (upper-left and lower-right). These pixel locations contribute disproportionately to the signal that travels between these two transducers and are selected by the algorithm, since the scattering between this pair has the highest amplitude.

6.5.3 Discussion

These cases demonstrate how each algorithm degrades with decreasing numbers of transducers. The block-sparse method develops severe artifacts with four transducers in experiment one (two glued-on steel rods), but is only slightly degraded by using four transducers in experiment two (one hand-cut notch). This suggests that as the number of transducers decreases, the block-sparse reconstruction method degrades as the number of scatterers increases; the union of the subspaces due to these scatterers grows in dimension and becomes more and more correlated with certain subspaces on the plate, which increases the likelihood of artifacts. In contrast, sparse reconstruction with signal envelopes is almost completely unaffected by a reduction of transducers in experiment one; the only noticeable effect is a reduction in amplitude of the bottom scatterer. In experiment two, however, the spot size of the scatterer increases as transducers are removed, and artifacts develop at or near the transducer locations. As discussed in Section 5.6.1, the envelope method can have difficulty with multiple
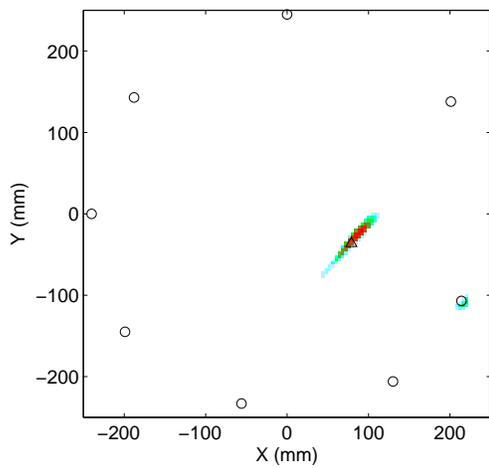
(a) Envelope method, 8 transducers (28 pairs)   (b) Block method, 8 transducers (28 pairs)
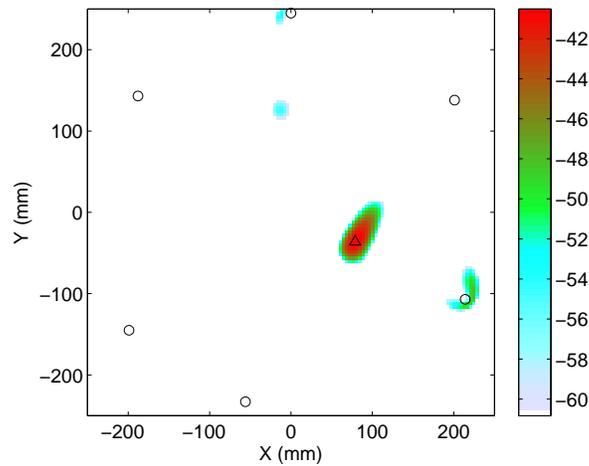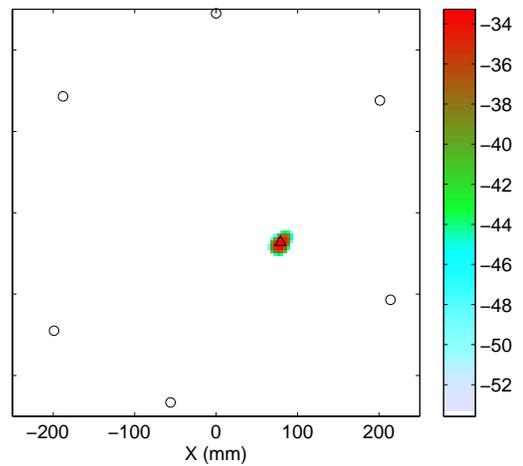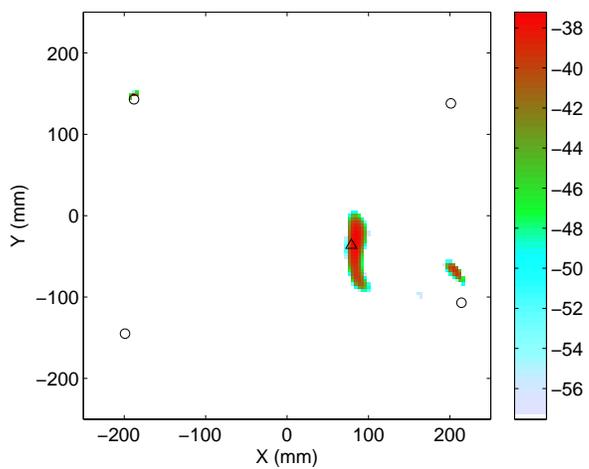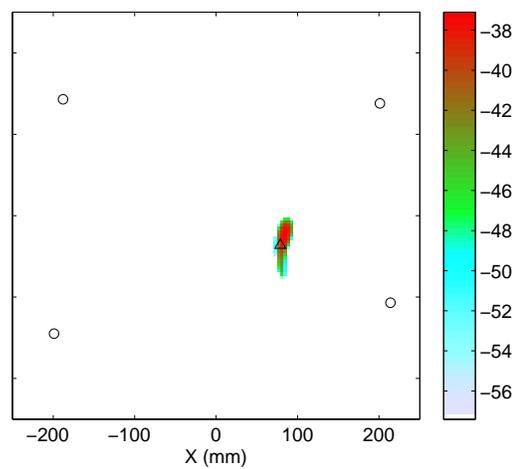
(c) Envelope method, 6 transducers (15 pairs)   (d) Block method, 6 transducers (15 pairs)

(e) Envelope method, 4 transducers (6 pairs)   (f) Block method, 4 transducers (6 pairs)

Figure 39: Results for experiment two with varying numbers of transducers. Circles denote active transducer locations and the triangle denotes the location of the notch.

scatterers, masking weaker-detected scatterers due to the nonlinear effects of envelope detection; this effect did not appear to be significantly affected by decreasing the number of transducers, although it may not appear until more than two scatterers are present. The noticeable progression of scatterer spot size in experiment two suggests that the effects of scatterer mismatch increase as the number of transducers is decreased for the envelope method.

One clear implication is that the quality of images generated with both methods increases as the number of transducers increases, especially when increasing from four to six transducers; the image quality is less substantially affected when increasing from six to eight transducers. In situations with only a small number of transducers, the different characteristics of the algorithms might make one method more favorable than the other, depending on anticipated flaw types and numbers.

6.6   Scatterer Characterization

One advantage of the block-sparse method is the ability to characterize damage by using the scattering coefficients from the solution. To demonstrate, block-sparse reconstruction was performed on the residual measurement set from experiment one (glued-on steel rods) for only a single glued-on rod, and the residual measurement set from experiment two (hand-cut notch). For both sets of data, image coefficient denormalization was performed, as described in Section 4.4.2. For each case, the result of block-sparse reconstruction is a set of 28 coefficients for every pixel location – one for each transducer pair. These coefficients were summed element-wise for all pixels in the scatterer's spot size to create a single set of 28 complex coefficients. The magnitudes of these coefficients were then circularly interpolated in one dimension with cubic splines to create eight scattering curves; each has a fixed incoming angle and is a function of outgoing angle.

Figure 40 shows the curves generated for the data from experiment one from

transducers 1, 3, and 5. Because the glued-on rod has a circular profile, the three scattering patterns theoretically should be identical other than their rotation; however, differences in transducer responses (which were not taken into account) and the large gaps between interpolation points can also have an effect on results. Nevertheless, the general characteristics of near-omni-directional scattering are apparent. The interpolated curves generated with the data from experiment two, using transducers 1, 3, and 5, are shown in Figure 41. The directional nature of the scatterer for this experiment is easily observed, with clear nulls in the curves around the end-on incident angles for the notch. Additionally, the scattering pattern from transducer 3 (red) has very low amplitude. Again, these results do not incorporate differences in transducer transfer functions, and the scattering curves are only sampled at seven angles, which could "straddle" a peak or valley in the response curves.

The clear differences between the scattering curves for experiment one (Figure 40) and two (Figure 41) demonstrate that, at the very least, a rudimentary characterization or classification algorithm should be able to discriminate between the two types of damage.

Figure 40: Interpolated scattering curves for experiment one (with only one attached rod) using block-sparse reconstruction and image denormalization. The colored arrows show incoming angle, and the matching curves show scattered amplitude as a function of angle. The white circles along the curves denote interpolation points.
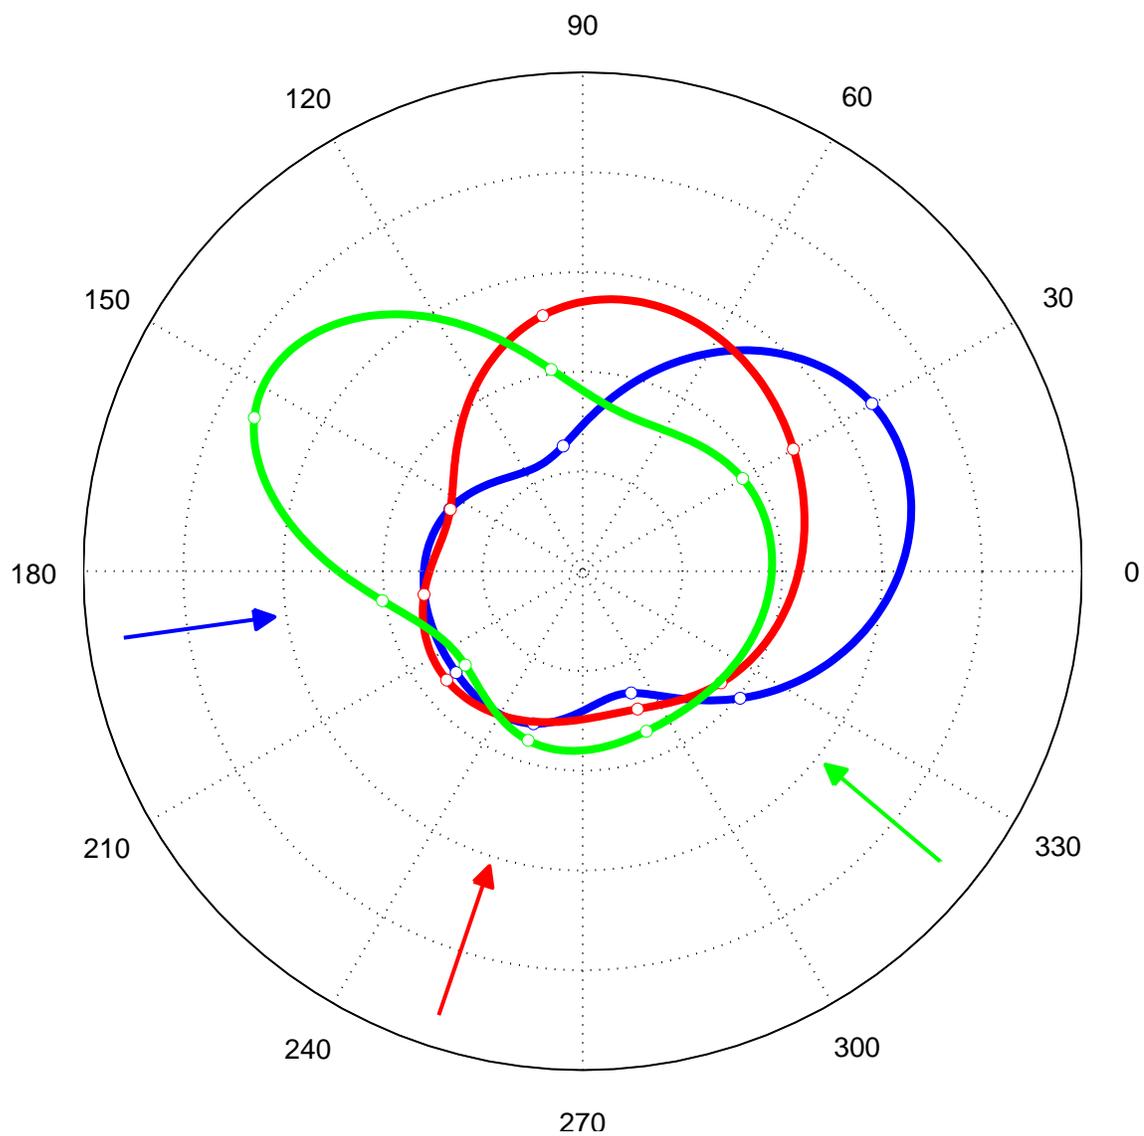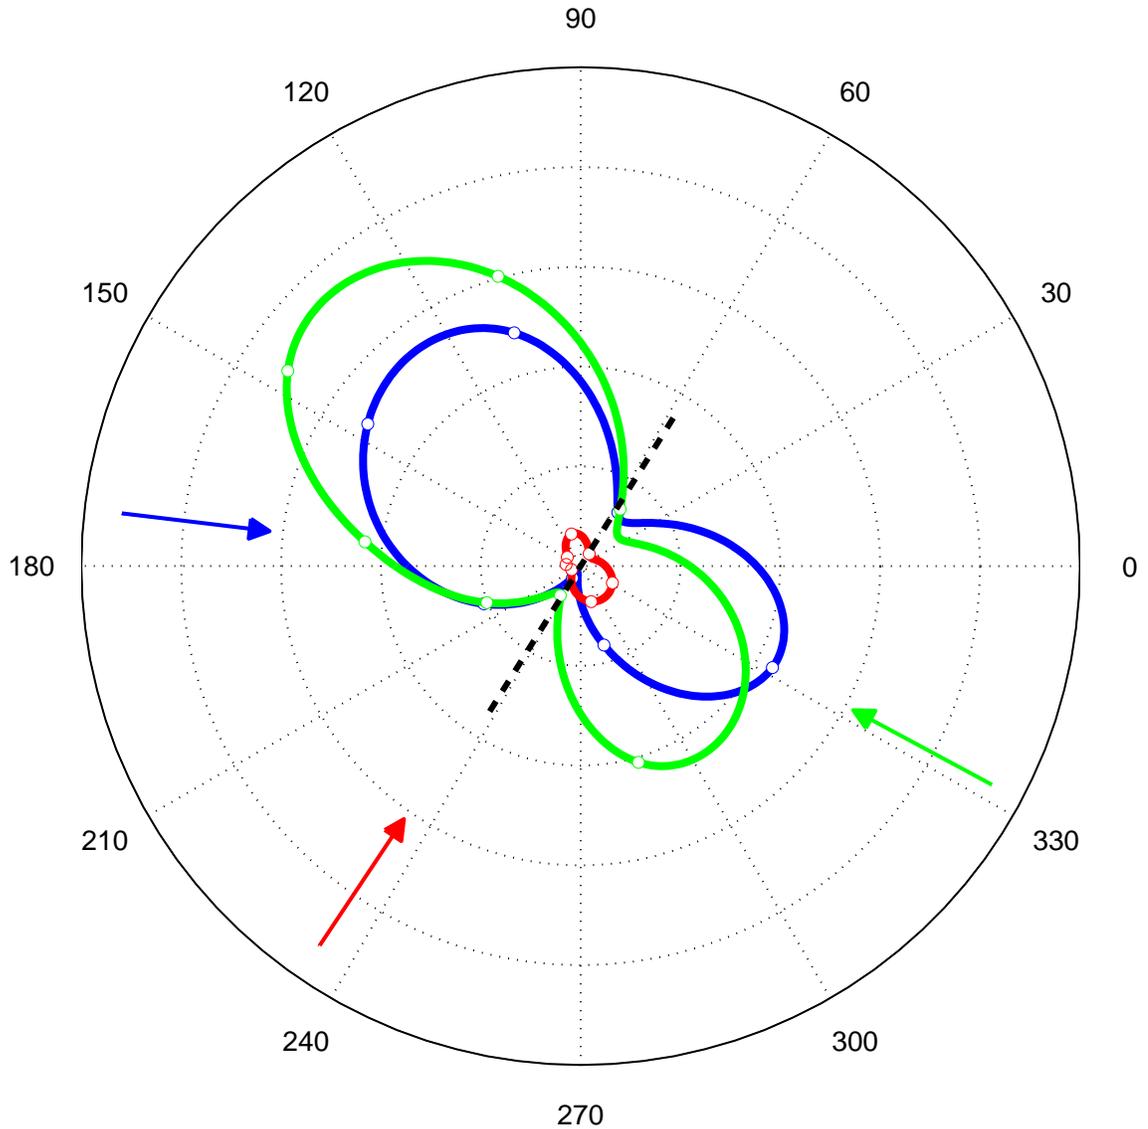
Figure 41: Interpolated scattering curves for experiment two (hand-cut notch) using block-sparse reconstruction and image denormalization. The colored arrows show incoming angle, and the matching curves show scattered amplitude as a function of angle. The white circles along the curves denote interpolation points. The black dashed line indicates the orientation of the notch.

6.7 Comparison of Methods

Table 11 compares the relative performance of sparse reconstruction with signal analytic representations, sparse reconstruction with signal envelopes, block-sparse reconstruction, and delay-and-sum imaging; this comparison is based on theory as well as simulated and experimental data. The performance of the envelope and block methods is shown for both dispersive and nondispersive dictionaries. Sparse reconstruction with real-valued signals is strictly inferior to sparse reconstruction with analytic representations and is not shown. The table serves as an outline of the overall trends exhibited by each method; actual performance depends on many factors, including the problem geometry, number of transducers, and noise levels, and a particular methods may perform better or worse than expected under various circumstances. A few observations do not fit in the table:

- When using a small number of transducers, images generated with the envelope method are degraded if the scatterer does not match the dictionary.

- When using a small number of transducers, images generated with the block-sparse method are likely to contain high-magnitude artifacts if multiple scatterers are present.

The general trend of the table is that sparse reconstruction with signal envelopes and block-sparse reconstruction are the two most successful methods, with block-sparse reconstruction tending to produce larger spot sizes of scatterers and artifacts, but the artifacts have smaller magnitudes than those present in images generated using sparse reconstruction with signal envelopes. Both methods exhibit overall image superiority to DAS, which produces images with large spot sizes and is very sensitive to interference, and sparse reconstruction with raw signals, which can fail outright due to its low robustness and overly-sensitive model.

Table 11: Comparison of imaging methods

	Dispersion Curves			Group Velocity		
	RF Analytic	Envelope	Block	Envelope-GV	Block-GV	DAS
Preparation required	dictionary[a]	dictionary[a]	dictionary[a]	dictionary[a]	dictionary[a]	calibration
Algorithm speed	minutes	minutes	minutes	minutes	minutes	seconds
Spot sizes[b] (matched scatterer)	v. small[c]	v. small	small	small	moderate	v. large
Spot sizes[b] (mismatched scatterer)	v. small[c]	small	small	moderate	moderate	v. large
Effect of interference on spot sizes[b]	extreme[d]	small	moderate	moderate	large	small[e]
Artifact severity[f] (matched scatterer)	moderate	small	small	moderate	small	large[g]
Artifact severity[f] (mismatched scatterer)	moderate	moderate	small	large	small	large[g]
Effect of interference on artifacts[f]	v. large	large	moderate	large	moderate	v. large

a Dictionary generation requires time and disk space.
b Refers to the spot sizes of both scatterers and artifacts.
c Spot size is typically one or two pixels; however, detection can fail outright in the presence of multiple scatterers.
d This method can fail when subject to even a moderate amount of interference.
e Spot sizes are already very poor for this method.
f Refers to number and magnitude of artifacts.
g Artifacts may be obscured by the large spot sizes of scatterers.

CHAPTER VII

SUMMARY AND CONCLUSIONS

7.1   Improving Localization via the Assumption of Sparsity

Most existing Lamb wave SHM imaging methods are based on delay-and-sum techniques. These are conceptually simple, but can have large spot sizes and many artifacts. These algorithms do not exploit the powerful nature of sparsity, though the problem is well-suited to that assumption for working structures. By assuming single-scattering behavior (which is an assumption that DAS already uses), responses to multiple sites of damage are linearized and can be represented by corresponding vectors (or blocks) in a redundant dictionary. The resultant linear equation can be solved by a sparse (or block-sparse) solver, as described in Sections 4.4 and 5.5. For most cases, where scattering behavior is unknown, there are two possible approaches, each with its strengths and weaknesses.

Sparse reconstruction using signal envelopes is described in Section 4.3.2 and uses envelope detection to remove phase information from signals before performing imaging. Envelope detection creates a nonlinearity that can have detrimental effects on scatterers that are close together, as demonstrated in Figure 9(a). Additionally, the resultant dictionary does not incorporate information about scatterer directionality, which can cause artifacts near transducer locations; the number and magnitudes of such artifacts are significantly increased when using small numbers of transducers. However, its performance is somewhat faster and its images tend to have smaller spot sizes than the block-sparse method (though it is still slower than DAS).

Block-sparse reconstruction uses a multidimensional model to match scatterers with phase shifts or directionality. Unlike the envelope method, it does not require

114

an *a priori* assumption of scatterer behavior or a nonlinear operation to correct for phase mismatch; instead block-sparse reconstruction allows any scatterer to have varying amplitude and phase between each transducer pair. Additionally, it can be used to characterize scattering behavior and possibly to characterize defects. Its more permissive model, however, can allow more artifacts and larger spot sizes than the envelope method, especially with a small number of transducers and multiple scatterers; it is somewhat slower as well.

The resultant images for both methods are quite similar overall, so the selection of which algorithm to use depends on the type and severity of anticipated damage as well as the trade-off between low artifact levels and small spot sizes. Both types of imaging outperform DAS imaging and are relatively insensitive to model mismatch; they yield superior results even when using only the group velocity.

One negative trait these algorithms share is that they require a precomputed dictionary. This large matrix can take several minutes to generate, even on a very fast computer, and if dictionary projection is performed, more time is required to compute its projection operator (or its approximation). These matrices consume storage space as well, and while they can be compressed somewhat via orthogonal transforms, sparse matrix representations, and downsampling, the file sizes of these matrices can become quite large for a large area of interest or a finely-sampled pixel grid.

7.2 Contributions

7.2.1 Formulation as Sparse Reconstruction

The major contribution of this work is the formulation and implementation of Lamb wave detection and localization as both sparse and block-sparse reconstructions, as described in Chapters 4 and 5. This new approach to imaging produces results that are overall superior to existing DAS methods.

For sparse reconstruction, the central concept is the redundant dictionary of location-based residual signal components using either real-valued or analytic representations. This dictionary requires *a priori* knowledge of scatterer behavior, which is often unrealistic; a modified method using signal envelopes is also presented, along with an explanation of benefits and drawbacks. An analysis of the coherent nature of the dictionary and the relation of its column space to the residual signals is also performed. The use of denormalization to compensate for the normalization of the dictionary matrix is described, and some general considerations are discussed.

Block-sparse reconstruction allows a multidimensional scattering model that requires no *a priori* information. The fundamental concept presented is the extension of the standard sparse reconstruction formulation to a block-sparse problem, including the idea of a scatterer having one or more distinct coefficients per transducer pair. An analysis of the coherence between dictionary blocks was performed, and computational considerations were compared to those of the non-block methods.

7.2.2   Noise Analysis of Envelope-Detected Signals

An analysis of envelope detection and the effects of Rician noise on sparse reconstruction with signal envelopes is presented in Sections 4.3.2.2 and 4.3.2.3. This analysis includes a method of estimating the true mean of noisy portions of the signal, the "corner lighting" effect of images generated with noisy signals, a simple compensation procedure, and the use of simulations to verify the effectiveness of noise compensation.

7.2.3   Use of Nondispersive Dictionaries

One vital question that arose during this research was that of robustness. How sensitive are such methods to the model assumptions? For any method to be viable, it must have some tolerance to model mismatch. Robustness is addressed here by using dictionaries that are nondispersive; i.e., they use only the group velocity of the propagation mode. In cases where dispersion curves are unavailable, it is important

116

that dictionaries can be used that do not require knowledge of dispersion curves. All experiments performed evaluated the performance of nondispersive dictionaries; these dictionaries were outperformed by those that incorporated dispersion curves, but still successfully localized scatterers. This success indicates that the sparsity-based methods are tolerant of at least some propagation model mismatch, demonstrating their potential for use in "real-world" applications.

7.2.4 Extraction of Scattering Patterns

Another contribution is a method of extracting scattering patterns from block-sparse image coefficients. A relation between the multidimensional pixel values and scattering amplitude and phase is briefly discussed in Section 5.5; a simple one-dimensional interpolation operation is performed on data from the images for experiments one and two to generate approximate scattering curves. Experiment one used glued-on steel masses, while experiment two featured a hand-cut notch; the extracted scattering curves for the two scatterers confirm their expected differences.

7.2.5 Experimental Verification

An important contribution of this work is the use of both presented methods on data from a variety of experiments along with comparisons to the corresponding delay-and-sum images. The first two experiments used an ideal plate with two different types of scatterers to show "best-case" behavior. The next experiment featured a small plate with many edge reflections to evaluate feasibility of the methods in a more difficult environment. An experiment was performed on a composite plate to demonstrate robustness to the propagation model, which is difficult to estimate for this material. Finally, images from the first two experiments were generated using differing numbers of transducers to determine the detrimental effects of limited amounts of data.

7.3 Potential for Future Work

Presented here is a mathematical framework and initial laboratory experiments for sparsity-based imaging methods. The next step for these algorithms is larger-scale experiments on more realistic structures. Any novel application gains acceptance only through extensive testing in a wide variety of scenarios. The aircraft industry is a multi-billion dollar sector of the economy; they are rightly conservative about accepting any new technology. Numerous experiments comprise the first step to gaining such acceptance.

The most important issue that further research must address is the general problem of model mismatch. The methods presented are somewhat robust to mismatch in the propagation model as well as unpredicted interference from edge reflections and poor baseline subtraction; however, additional work is required to determine a more precise relation between mismatch severity and image quality for various types of model mismatch. A particularly valuable result would be quantification of the effect of poor baseline subtraction, which can cause interference that, to a sparse algorithm, is indistinguishable from damage.

One potential direction for future research is the possible application to damage characterization, a capability that few current sparse array methods possess. Initial examples of the potential of characterization are shown for the first two experiments, demonstrating that scattering behavior and directionality can be determined, at least to some degree. The extent to which these data can be used, however, is currently unknown and beyond the scope of this present work. Experiments could be performed to see if reconstruction coefficients might be usable to, for example, track the growth of cracks, in addition to discriminating between cracks and benign scatterers.

Another possibility for future investigation is improved dictionary generation. It might be able to generate an adaptive dictionary on the fly that can use measured

Lamb wave parameters to produce superior images. This would also have the advantage of requiring far less disk space, which could make the proposed methods more feasible for implementation on certain embedded hardware.

7.4 Conclusions

7.4.1 Sparsity Assumption

The central concept of this research is that the assumption of damage sparsity can be used to generate Lamb wave images of scatterers. The analysis and experiments presented demonstrate that sparsity-based methods are viable and generally behave as expected.

7.4.2 Performance

Experimental results show that images generated via sparse reconstruction of signal envelopes and block-sparse reconstruction are superior to conventional delay-and-sum images when an appropriate model is provided. In particular, sparse and block-sparse images have significantly smaller spot sizes; fewer artifacts; and lower artifact magnitudes, with the exception of artifacts at strongly reflecting plate edges.

7.4.3 Model Mismatch

The success of imaging using only group velocities demonstrates that both sparse reconstruction of signal envelopes and block-sparse reconstruction are tolerant to some degree of model mismatch. Additionally, the block-sparse imaging method is not susceptible to amplitude or phase mismatch of the scatterer; sparse reconstruction with signal envelopes uses a scattering model, but the primary effect of scatterer mismatch is an increase in the number of artifacts, usually present at or near the transducer locations.

7.4.4 Additional Advantages

The sparse and block-sparse methods presented have additional benefits over conventional delay-and-sum imaging. The availability of a precomputed dictionary allows the use of projection to quantify the level of model mismatch in residual signals; this can indicate, for example, poor baseline subtraction, high levels of noise or interference, or a problem with the SHM system. The use of BPDN offers a trade-off between sparsity and reconstruction fidelity, which allows flexibility at the cost of requiring some degree of user input. Finally, the block-sparse method allows scattering patterns to be readily extracted from imaging results.

7.4.5 Concluding Remarks

Many hurdles stand between this document and the widespread adoption of these techniques. The process will be gradual and will require additional research in cooperation with industry, as well as advances in computational resources. However, it is the belief of this author that a day will come that these techniques, specifically, the ability to incorporate the assumption that damage is sparse, will be used to keep critical structures safe and reduce the need for expensive and time-consuming manual inspections.

REFERENCES

[1] H. Lamb, "On waves in an elastic plate," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 93, no. 648, pp. 114–128, 1917.

[2] D. C. Worlton, "Ultrasonic testing with Lamb waves," tech. rep., General Electric Co., Hanford Atomic Products Operation, Richland, Wash., 1956.

[3] J. L. Rose *et al.*, "A baseline and vision of ultrasonic guided wave inspection potential," *Journal of Pressure Vessel Technology*, vol. 124, no. 3, pp. 273–282, 2002.

[4] D. N. Alleyne and P. Cawley, "The interaction of Lamb waves with defects," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 3, pp. 381–397, 1992.

[5] J. L. Rose, A. Pilarski, and J. J. Ditri, "An approach to guided wave mode selection for inspection of laminated plate," *Journal of Reinforced Plastics and Composites*, vol. 12, no. 5, pp. 536–544, 1993.

[6] N. Guo and P. Cawley, "The interaction of Lamb waves with delaminations in composite laminates," *The Journal of the Acoustical Society of America*, vol. 94, pp. 2240–2246, 1993.

[7] W. Wright, D. Hutchins, D. Jansen, and D. Schindel, "Air-coupled Lamb wave tomography," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 1, pp. 53–59, 1997.

[8] J. C. P. McKeon and M. K. Hinders, "Parallel projection and crosshole Lamb wave contact scanning tomography," *The Journal of the Acoustical Society of America*, vol. 106, pp. 2568–2577, 1999.

[9] P. D. Wilcox, "Omni-directional guided wave transducer arrays for the rapid inspection of large areas of plate structures," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 50, no. 6, pp. 699–709, 2003.

[10] A. S. Purekar, D. J. Pines, S. Sundararaman, and D. E. Adams, "Directional piezoelectric phased array filters for detecting damage in isotropic plates," *Smart Materials and Structures*, vol. 13, no. 4, pp. 838–850, 2004.

[11] P. Fromme, P. D. Wilcox, M. J. S. Lowe, and P. Cawley, "On the development and testing of a guided ultrasonic wave array for structural integrity monitoring," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 53, no. 4, pp. 777–785, 2006.

[12] C. H. Wang, J. T. Rose, and F. K. Chang, "A synthetic time-reversal imaging method for structural health monitoring," *Smart Materials and Structures*, vol. 13, no. 2, pp. 415–423, 2004.

[13] T. E. Michaels and J. E. Michaels, "Sparse ultrasonic transducer array for structural health monitoring," in *AIP Conference Proceedings*, vol. 700, pp. 1468–1475, 2004.

[14] J. E. Michaels and T. E. Michaels, "Guided wave signal processing and image fusion for *in situ* damage localization in plates," *Wave Motion*, vol. 44, no. 6, pp. 482–492, 2007.

[15] X. Zhao, H. Gao, G. Zhang, B. Ayhan, F. Yan, C. Kwan, and J. L. Rose, "Active health monitoring of an aircraft wing with embedded piezoelectric sensor/actuator network: I. defect detection, localization and growth monitoring," *Smart Materials and Structures*, vol. 16, no. 4, pp. 1208–1217, 2007.

[16] R. Sicard, J. Goyette, and D. Zellouf, "A SAFT algorithm for Lamb wave imaging of isotropic plate-like structures," *Ultrasonics*, vol. 39, no. 7, pp. 487–494, 2002.

[17] R. Sicard, A. Chahbaz, and J. Goyette, "Guided Lamb waves and L-SAFT processing technique for enhanced detection and imaging of corrosion defects in plates with small depth-to wavelength ratio," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 51, no. 10, pp. 1287–1297, 2004.

[18] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[19] V. Giurgiutiu, "Tuned Lamb wave excitation and detection with piezoelectric wafer active sensors for structural health monitoring," *Journal of Intelligent Material Systems and Structures*, vol. 16, no. 4, pp. 291–305, 2005.

[20] S. J. Lee, J. E. Michaels, H. Sohn, and T. E. Michaels, "Piezoelectric transducer diagnostics via linear reciprocity for guided wave structural health monitoring," *AIAA journal*, vol. 49, no. 3, pp. 621–629, 2011.

[21] U.S. Federal Aviation Administration, *Aviation Maintenance Technician Handbook: General 2008*. JL Aviation LLC, 2008.

[22] P. J. Shull, *Nondestructive Evaluation: Theory, Techniques, and Applications*, vol. 142. CRC, 2002.

[23] U.S. Code of Federal Regulations, "14 C.F.R. part 91," `http://www.ecfr.gov/cgi-bin/text-idx?node=14:2.0.1.3.10`, Nov. 2013.

[24] D. Inaudi, M. Bolster, R. Deblois, C. French, A. Phipps, J. Sebasky, and K. Western, "Structural health monitoring system for the new I-35W St. Anthony Falls Bridge," in *4th International Conference on Structural Health Monitoring of Intelligent Infrastructure*, p. 10, 2009.

[25] R. J. Ditchburn, S. K. Burke, and C. M. Scala, "NDT of welds: state of the art," *NDT & E International*, vol. 29, no. 2, pp. 111–117, 1996.

[26] B. W. Drinkwater and P. D. Wilcox, "Ultrasonic arrays for non-destructive evaluation: A review," *NDT & E International*, vol. 39, no. 7, pp. 525–541, 2006.

[27] M. J. S. Lowe, D. N. Alleyne, and P. Cawley, "Defect detection in pipes using guided waves," *Ultrasonics*, vol. 36, no. 1, pp. 147–154, 1998.

[28] D. N. Alleyne and P. Cawley, "The excitation of Lamb waves in pipes using dry-coupled piezoelectric transducers," *Journal of Nondestructive Evaluation*, vol. 15, no. 1, pp. 11–20, 1996.

[29] J. D. Achenbach, *Wave Propagation in Elastic Solids*, vol. 16. North-Holland Amsterdam, 1973.

[30] D. C. Worlton, "Experimental confirmation of Lamb waves at megacycle frequencies," *Journal of Applied Physics*, vol. 32, no. 6, pp. 967–971, 1961.

[31] I. A. Viktorov, *Rayleigh and Lamb Waves: Physical Theory and Applications*, vol. 147. Plenum Press New York, 1967.

[32] V. Giurgiutiu, *Structural Health Monitoring with Piezoelectric Wafer Active Sensors*. Elsevier, 2008.

[33] P. Wilcox, M. Lowe, and P. Cawley, "Lamb and SH wave transducer arrays for the inspection of large areas of thick plates," in *AIP Conference Proceedings*, vol. 509, pp. 1049–1056, 2000.

[34] B. Pavlakovic, M. Lowe, D. Alleyne, and P. Cawley, "Disperse: a general purpose program for creating dispersion curves," *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 16, pp. 185–192, 1997.

[35] Vallen Systeme GmbH, "Vallen Dispersion," Available as part of the AGU-Vallen Wavelet software package, `http://www.vallen.de/products/software/wavelet`, 2013.

[36] J. Zhang, B. W. Drinkwater, and P. D. Wilcox, "Defect characterization using an ultrasonic array to measure the scattering coefficient matrix," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 55, no. 10, pp. 2254–2265, 2008.

[37] Y. Cho, D. D. Hongerholt, and J. L. Rose, "Lamb wave scattering analysis for reflector characterization," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 1, pp. 44–52, 1997.

[38] O. Diligent, T. Grahn, A. Boström, P. Cawley, and M. J. S. Lowe, "The low-frequency reflection and scattering of the S_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*, vol. 112, pp. 2589–2601, 2002.

[39] M. J. S. Lowe and O. Diligent, "Low-frequency reflection characteristics of the S_0 Lamb wave from a rectangular notch in a plate," *The Journal of the Acoustical Society of America*, vol. 111, pp. 64–74, 2002.

[40] M. J. S. Lowe, P. Cawley, J. Y. Kao, and O. Diligent, "The low frequency reflection characteristics of the fundamental antisymmetric Lamb wave a_0 from a

rectangular notch in a plate," *The Journal of the Acoustical Society of America*, vol. 112, pp. 2612–2622, 2002.

[41] J. M. Galán and R. Abascal, "Numerical simulation of Lamb wave scattering in semi-infinite plates," *International Journal for Numerical Methods in Engineering*, vol. 53, no. 5, pp. 1145–1173, 2002.

[42] A. J. Croxford, P. D. Wilcox, G. Konstantinidis, and B. W. Drinkwater, "Strategies for overcoming the effect of temperature on guided wave structural health monitoring," in *Proceedings of SPIE*, vol. 6532, p. 65321T, 2007.

[43] J. E. Michaels, "Detection, localization and characterization of damage in plates with an *in situ* array of spatially distributed ultrasonic sensors," *Smart Materials and Structures*, vol. 17, no. 3, p. 035035, 2008.

[44] T. Clarke, P. Cawley, P. D. Wilcox, and A. J. Croxford, "Evaluation of the damage detection capability of a sparse-array guided-wave SHM system applied to a complex structure under varying thermal conditions," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 56, no. 12, pp. 2666–2678, 2009.

[45] E. Moulin, J. Assaad, C. Delebarre, H. Kaczmarek, and D. Balageas, "Piezoelectric transducer embedded in a composite plate: application to Lamb wave generation," *Journal of Applied Physics*, vol. 82, no. 5, pp. 2049–2055, 1997.

[46] J. S. Hall and J. E. Michaels, "Analysis of distributed sparse array configurations for guided wave imaging applications," in *AIP Conference Proceedings*, vol. 1335, pp. 859–866, 2011.

[47] X. Chen, J. E. Michaels, and T. E. Michaels, "Design of distributed sparse arrays for Lamb wave SHM based upon estimated scattering matrices," in *AIP Conference Proceedings*, vol. 1581, pp. 248–255, 2014.

[48] K. F. Braun, "Electrical oscillations and wireless telegraphy," *Nobel Lectures*, vol. 11, pp. 226–245, December 1909.

[49] D. H. Johnson and D. E. Dudgeon, *Array Signal Processing: Concepts and Techniques*. Prentice Hall, 1993.

[50] L. F. Brown, "Large-area ultrasonic gap profile measurements using two-dimensional PVDF," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 42, no. 1, pp. 120–126, 1995.

[51] L. Azar, Y. Shi, and S.-C. Wooh, "Beam focusing behavior of linear phased arrays," *NDT & E International*, vol. 33, no. 3, pp. 189–198, 2000.

[52] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "Post-processing of the full matrix of ultrasonic transmit–receive array data for non-destructive evaluation," *NDT & E International*, vol. 38, no. 8, pp. 701–711, 2005.

[53] P. D. Wilcox, C. Holmes, and B. W. Drinkwater, "Advanced reflector characterization with ultrasonic phased arrays in NDE applications," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 54, no. 8, pp. 1541–1550, 2007.

[54] L. Yu and V. Giurgiutiu, "In-situ optimized PWAS phased arrays for Lamb wave structural health monitoring," *Journal of Mechanics of Materials and Structures*, vol. 2, no. 3, pp. 459–487, 2007.

[55] T. Stepinski, "Ultrasonic nondestructive inspection of solid objects," in *AIP Conference Proceedings*, vol. 1433, pp. 11–20, 2012.

[56] M. Spies, H. Rieder, A. Dillhöfer, V. Schmitz, and W. Müller, "Synthetic aperture focusing and time-of-flight diffraction ultrasonic imaging–past and

present," *Journal of Nondestructive Evaluation*, vol. 31, no. 4, pp. 310–323, 2012.

[57] J. A. Jensen, S. I. Nikolov, K. L. Gammelmark, and M. H. Pedersen, "Synthetic aperture ultrasound imaging," *Ultrasonics*, vol. 44, pp. e5–e15, 2006.

[58] A. J. Croxford, J. Moll, P. D. Wilcox, and J. E. Michaels, "Efficient temperature compensation strategies for guided wave structural health monitoring," *Ultrasonics*, vol. 50, no. 4, pp. 517–528, 2010.

[59] J. E. Michaels, A. C. Cobb, and T. E. Michaels, "A comparison of feature-based classifiers for ultrasonic structural health monitoring," in *Proceedings of SPIE*, vol. 5394.

[60] Y. Lu and J. E. Michaels, "A methodology for structural health monitoring with diffuse ultrasonic waves in the presence of temperature variations," *Ultrasonics*, vol. 43, no. 9, pp. 717–731, 2005.

[61] S. J. Lee, N. Gandhi, J. E. Michaels, and T. E. Michaels, "Comparison of the effects of applied loads and temperature variations on guided wave propagation," in *AIP Conference Proceedings*, vol. 1335, pp. 175–182, 2011.

[62] N. Gandhi, J. E. Michaels, and S. J. Lee, "Acoustoelastic Lamb wave propagation in biaxially stressed plates," *The Journal of the Acoustical Society of America*, vol. 132, pp. 1284–1293, 2012.

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

[64] G. Konstantinidis, B. W. Drinkwater, and P. D. Wilcox, "The temperature stability of guided wave structural health monitoring systems," *Smart Materials and Structures*, vol. 15, no. 4, pp. 967–976, 2006.

[65] Z. Lu, S. J. Lee, J. E. Michaels, and T. E. Michaels, "On the optimization of temperature compensation for guided wave structural health monitoring," in *AIP Conference Proceedings*, vol. 1211, pp. 1860–1867, 2010.

[66] J. E. Michaels, R.-J. Liou, J. P. Zutty, and T. E. Michaels, "On the interpolation of ultrasonic guided wave signals," in *AIP Conference Proceedings*, vol. 1430, pp. 679–686, 2012.

[67] S. J. Lee, J. E. Michaels, X. Chen, and T. E. Michaels, "Fatigue crack monitoring via load-differential guided wave methods," in *AIP Conference Proceedings*, vol. 1430, pp. 1575–1582, 2012.

[68] A. Mal, S. Banerjee, and F. Ricci, "An automated damage identification technique based on vibration and wave propagation data," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 365, no. 1851, pp. 479–491, 2007.

[69] J. C. Dodson, J. R. Foley, and D. J. Inman, "Structural health monitoring of an impulsively loaded structure using wave-propagation based instantaneous baseline," in *Proceedings of SPIE*, vol. 7650, p. 765038, International Society for Optics and Photonics, 2010.

[70] H. T. Banks, D. J. Inman, D. J. Leo, and Y. Wang, "An experimentally validated damage detection theory in smart structures," *Journal of Sound and Vibration*, vol. 191, no. 5, pp. 859–880, 1996.

[71] H. Sohn, D. Dutta, and Y.-K. An, "Temperature independent damage detection in plates using redundant signal measurements," *Journal of Nondestructive Evaluation*, vol. 30, no. 2, pp. 106–116, 2011.

[72] B. Xu and V. Giurgiutiu, "Single mode tuning effects on Lamb wave time reversal with piezoelectric wafer active sensors for structural health monitoring," *Journal of Nondestructive Evaluation*, vol. 26, no. 2-4, pp. 123–134, 2007.

[73] J. S. Hall and J. E. Michaels, "Minimum variance ultrasonic imaging applied to an in situ sparse guided wave array," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 57, no. 10, pp. 2311–2323, 2010.

[74] E. B. Flynn, M. D. Todd, A. J. Croxford, B. W. Drinkwater, and P. D. Wilcox, "Enhanced detection through low-order stochastic modeling for guided-wave structural health monitoring," *Structural Health Monitoring*, vol. 11, no. 2, pp. 149–160, 2012.

[75] J. S. Hall and J. E. Michaels, "Multi-path guided wave imaging for *in situ* monitoring of complex structures," in *AIP Conference Proceedings*, vol. 1511, pp. 707–713, 2013.

[76] E. V. Malyarenko and M. K. Hinders, "Ultrasonic Lamb wave diffraction tomography," *Ultrasonics*, vol. 39, no. 4, pp. 269–281, 2001.

[77] K. R. Leonard, E. V. Malyarenko, and M. K. Hinders, "Ultrasonic Lamb wave tomography," *Inverse Problems*, vol. 18, no. 6, pp. 1795–1808, 2002.

[78] P. Belanger and P. Cawley, "Feasibility of low frequency straight-ray guided wave tomography," *NDT & E International*, vol. 42, no. 2, pp. 113–119, 2009.

[79] P. Belanger, P. Cawley, and F. Simonetti, "Guided wave diffraction tomography within the Born approximation," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 57, no. 6, pp. 1405–1418, 2010.

[80] T. R. Hay, R. L. Royer, H. Gao, X. Zhao, and J. L. Rose, "A comparison of embedded sensor Lamb wave ultrasonic tomography approaches for material loss detection," *Smart Materials and Structures*, vol. 15, no. 4, pp. 946–951, 2006.

[81] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[82] M. Fornasier and H. Rauhut, "Compressive sensing," *Handbook of Mathematical Methods in Imaging*, vol. 1, pp. 187–229, 2011.

[83] F. Santosa and W. W. Symes, "Reconstruction of blocky impedance profiles from normal-incidence reflection seismograms which are band-limited and miscalibrated," *Wave Motion*, vol. 10, no. 3, pp. 209–230, 1988.

[84] A. Skodras, C. Christopoulos, and T. Ebrahimi, "The JPEG 2000 still image compression standard," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 36–58, 2001.

[85] A. P. Lobo and P. C. Loizou, "Voiced/unvoiced speech discrimination in noise using Gabor atomic decomposition," in *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1, pp. I–820, IEEE, 2003.

[86] A. C. Gurbuz, J. H. McClellan, and W. R. Scott Jr., "Compressive sensing of underground structures using GPR," *Digital Signal Processing*, vol. 22, no. 1, pp. 66–73, 2012.

[87] F. Boßmann, G. Plonka, T. Peter, O. Nemitz, and T. Schmitte, "Sparse deconvolution methods for ultrasonic NDT," *Journal of Nondestructive Evaluation*, vol. 31, no. 3, pp. 225–244, 2012.

[88] A. Perelli, T. Ianni, A. Marzani, L. Marchi, and G. Masetti, "Model-based compressive sensing for damage localization in Lamb wave inspection," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 60, no. 10, 2013.

[89] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[90] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.

[91] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.

[92] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[93] E. J. Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians: Madrid, August 22-30, 2006: invited lectures*, pp. 1433–1452, 2006.

[94] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.

[95] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *1993 Conference Record of The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers*, pp. 40–44, IEEE, 1993.

[96] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.

[97] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proceedings of the National Academy of Sciences*, vol. 106, no. 45, pp. 18914–18919, 2009.

[98] L. I. Rudin and S. Osher, "Total variation based image restoration with free local constraints," in *Proceedings of the 1994 IEEE International Conference on Image Processing*, vol. 1, pp. 31–35, IEEE, 1994.

[99] "Sparse- and low-rank approximation wiki," `http://www.ugcs.caltech.edu/~srbecker/wiki/Main_Page`, 2013.

[100] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Transactions on Information Theory*, vol. 55, no. 11, pp. 5302–5316, 2009.

[101] Y. C. Eldar, P. Kuppinger, and H. Bolcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Transactions on Signal Processing*, vol. 58, no. 6, pp. 3042–3054, 2010.

[102] Internation Telecom Union, "ITU-T Recommendation Q.23: Technical features of push-button telephone sets," `http://www.itu.int/rec/T-REC-Q.23-198811-I/en`, Nov. 2013.

[103] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction," June 2007. http://www.cs.ubc.ca/labs/scl/spgl1.

[104] R. M. Levine, J. E. Michaels, and S. J. Lee, "Guided wave localization of damage via sparse reconstruction," in *AIP Conference Proceedings*, vol. 1430, pp. 647–654, 2012.

[105] R. M. Levine and J. E. Michaels, "Model-based imaging of damage with Lamb waves via sparse reconstruction," *The Journal of the Acoustical Society of America*, vol. 133, pp. 1525–1534, 2013.

[106] B. L. Sturm and M. G. Christensen, "Comparison of orthogonal matching pursuit implementations," in *2012 Proceedings of the 20th European Signal Processing Conference*, pp. 220–224, IEEE, 2012.

[107] J. Huang, C. R. Berger, S. Zhou, and J. Huang, "Comparison of basis pursuit algorithms for sparse channel estimation in underwater acoustic OFDM," in *OCEANS 2010 IEEE Sydney*, pp. 1–6, IEEE, 2010.

[108] S. O. Rice, "Mathematical analysis of random noise-conclusion," *Bell Systems Technical Journal*, vol. 24, pp. 46–156, 1945.

[109] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.

[110] B. Dawes, D. Abrahams, and R. Rivera, "Boost C++ libraries," `http://www.boost.org/`.

[111] E. Jones, T. Oliphant, and P. Peterson, "SciPy: Open source scientific tools for Python," `http://www.scipy.org/`, 2013.

[112] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Communications on Pure and Applied Mathematics*, vol. 41, no. 7, pp. 909–996, 1988.

[113] R. M. Levine and J. E. Michaels, "Block-sparse Lamb wave structural health monitoring using generic scattering models," in *AIP Conference Proceedings*, vol. 1581, pp. 232–239, 2014.

[114] R. M. Levine and J. E. Michaels, "Block-sparse reconstruction and imaging for Lamb wave structural health monitoring," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, accepted for publication.

[115] Å. Björck and G. H. Golub, "Numerical methods for computing angles between linear subspaces," *Mathematics of Computation*, vol. 27, no. 123, pp. 579–594, 1973.

[116] D. R. Hardoon, S. Szedmak, and J. Shawe-Taylor, "Canonical correlation analysis: An overview with application to learning methods," *Neural Computation*, vol. 16, no. 12, pp. 2639–2664, 2004.

[117] J. E. Michaels, S. J. Lee, A. J. Croxford, and P. D. Wilcox, "Chirp excitation of ultrasonic guided waves," *Ultrasonics*, vol. 53, no. 1, pp. 265–270, 2013.

[118] A. Dawson, J. E. Michaels, T. E. Michaels, and R. M. Levine, "Analysis of angle-beam wavefield measurements," in *AIP Conference Proceedings*, vol. 1581, pp. 1716–1723, 2014.