Remote Acoustic Sensing of Vibrating Structures for Structural Health Monitoring

by

Tyler Flynn

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in the University of Michigan
2019

Doctoral Committee:

Professor David Dowling, Chair
Assistant Professor Bogdan Popa
Assistant Professor Serife Tol
Professor Nickolas Vlahopoulos
Professor Kon-Well Wang

Tyler J. Flynn

tjayflyn@umich.edu

t.jayflynn@gmail.com

ORCID iD: 0000-0002-3126-5015

Dedication

The following document – and the efforts toward it – were made possible with the support of friends, family, and colleagues. This thesis is dedicated to them, and to you, the reader.

Acknowledgements

Many have played a role in supporting this thesis work – certainly more than can be enumerated here. Nevertheless, here is an earnest attempt:

My advisor, Prof. David Dowling, was an indispensable resource. His exceptional research acumen, willingness to let me pursue interesting problems, and his much appreciated open-door policy were all instrumental in shaping the content, quality, and accuracy of this document. This thesis has been further revised and enriched with feedback received from my doctoral committee members, Asst. Prof. Bogdan Popa, Asst. Prof. Serife Tol, Prof. Nickolas Vlahopoulos, and former Department Chair, Prof. Kon-Well Wang. Discussions with Prof. Karl Grosh also yielded valuable structural acoustics insight and were much appreciated.

Funding is, of course, a necessity in the research sciences. I am very thankful for the several organizations that deemed it prudent to support me financially throughout my graduate work. The U.S. Navy supported this work through multiple channels. Several of the projects I worked on – along with a number of undergraduate and Master's students I had the pleasure of working with – were motivated and funded through the Naval Engineering Education Consortium. I was also very fortunate to be awarded a National Defense Science and Engineering Graduate fellowship through the Office of Naval Research. I must also thank the numerous points of contact and project managers I have interacted with at the Carderock Naval Surface Warfare Center including Natasha Chang, Jonathan Forest, Jason Smoker, Alexi Titovich, and Kuangcheng Wu. I would also like to acknowledge the Acoustical Society of America for various travel subsidies allowing me to present

my research all over the nation, and the University of Michigan Rackham Graduate School for funding various research related endeavors.

During my time at Michigan I have been lucky to participate in several student organizations and numerous extracurricular activities. From ice cream socials to weekend outreach workshops, the friends and colleagues I have worked with have truly been a pleasure. I appreciate the many Michigan students I served with as a part of the Mechanical Engineering Graduate Council, Michigan Acoustics, and the Graduate Student Advisory Committee, as well as those students from other universities that I served with on the Acoustical Society of America Student Council.

Finally, I must thank my family and friends in Michigan, back in Kentucky, and elsewhere. Be it lengthy (and occasionally relevant) discourses with labmates, or holiday get-togethers with way too much food, or the music festivals where we definitely didn't bring enough water – these experiences have made the last five years an amazing time. I am grateful my marvelous mother, Renee, for easily being my #1 fan – without her I would (literally and figuratively) never have been able to complete this journey. And finally, to my wonderful partner Hannah: I am so lucky to have met you (and Francis, Eugene, and Leo) along this journey. Your support means everything, and I can't wait to begin our next chapter!

Table of Contents

List of Figures

List of Tables

Abstract

Key words: structural acoustics, phased arrays, signal processing, vibration, remote sensing, structural health monitoring

Acoustic radiation from a vibrating mechanical structure subject to narrow or broadband forcing is inherently dependent on the structure's condition, material, geometry, and boundary conditions. In many applications – such as those in naval, automotive, and civil engineering – acoustic recordings of a structure of interest may be easily and affordably obtained using a remote passive receiver or an array of receivers. As a result, there is great value in structural health monitoring techniques that can yield insight into a structure's condition using such conveniently accessible data sets, even though these methods bring forth the added complexity of acoustic propagation in the environment, source-receiver geometry, exogenous acoustic sources, etc. This dissertation details the development and evaluation of remote acoustic sensing techniques for the detection, localization, and classification of damage and/or mechanical changes in a vibrating structure. These techniques are based on comparisons of current test recordings with known baseline measurements in concert with both conventional and advanced array signal processing methods to achieve accurate and robust remote sensing performance.

Results are divided primarily into three sections reflecting the three classical components of structural health monitoring: Detection, Localization, and Classification of damage or changes. An experimental apparatus was constructed for base excitation of 30-cm square plates, wherein various forms of mechanical changes (i.e. synthetic damage) were applied to plates of differing material properties and geometry in order to test and evaluate the proposed methods. *Detection* of synthesized damage including cuts, fastener faults,

delamination, and localized added mass are achieved via baseline comparisons of remote array recordings of the plates' acoustic signatures. Cuts in metal plates as small as 1.3 cm and delamination areas of 6 cm^2 in composite fiber plates are detectable with >99% certainty using remote acoustic measurements captured meters from the plates. The addition of a blind deconvolution algorithm is shown to improve robustness of the method in practical environments with considerable unknown reverberation, for signal-to-reverberation ratios as low as -13 dB. A method for *localization* of mechanical changes/damage is presented using an innovative array-based coherent background subtraction technique that combines an existing high-resolution beamforming method with an additional subarray averaging step, ultimately enabling localization of small changes from a distributed acoustic source. Results are presented for the localization of cuts, fastener faults, and delamination in non-compact vibrating plates. Finally, a data-driven method is presented for the *classification* of cuts, fastener faults, and corrosion in vibrating plates via remote acoustic recordings. Using Monte Carlo finite element simulated training data for various damage cases, several classification strategies are evaluated and compared. Damage classification rates in excess of 96% are achieved using simulated training data and these results are corroborated with actual acoustic recordings of damaged plates.

Chapter I: Background & Introduction

1.1 Background

1.1.1 Vibroacoustic SHM

The presence of damage in a mechanical structure often changes its vibrational characteristics. For example, resonant frequencies may shift and mode shapes may be distorted in a damaged structure relative to a healthy counterpart [1, 2]. Damage may also affect structural integrity, stiffness, damping, and vibration-energy loss mechanisms such that resonant peak frequency locations, and widths may change. In severe cases, damage may even cause nonlinearities to arise. In 1969, Lifshitz & Rotem published one of the first articles motivating the use of vibroacoustic changes for damage detection (specifically the detection of reinforcement unbonding in composite materials) [3]. Since then, a litany of structural health monitoring (SHM) techniques have been developed which utilize measurable vibroacoustic changes to illuminate properties of potential faults in a mechanical system [4, 5, 6, 7]. The process of SHM is divided into four pillars [4, 8]: Detection – *does damage exist?*, Localization – *if so, where in the structure?*, Classification – *what type of damage is it?*, and Evaluation – *how severe is the damage and what is the system prognosis?* As implied in Figure 1.1, each pillar is necessary to obtain a complete picture of a structure's health. Moreover, there generally also exists some amount of overlap in the quantification of the four pillars, the extent of which

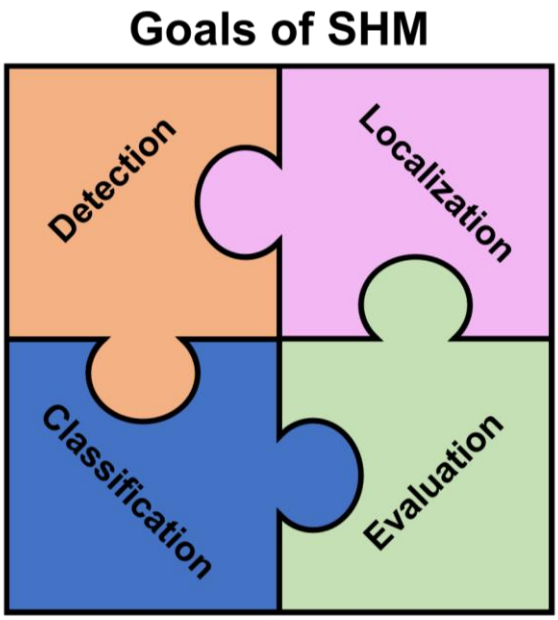

Figure 1.1: Fundamental pillars of Structural Health Monitoring (SHM).

depends heavily on application and approach. Understanding these relationships is key to fully recognizing the utility of a SHM technique for a given system. Since the implementation of guidelines for evaluation is highly specific to a given system and its application, the findings of the following thesis are generally limited to investigating the *detection*, *localization*, and *classification* capabilities of remote acoustic sensing.

There exist many approaches for vibroacoustic-based SHM. Direct, discrete measurements of acceleration and strain are by far the most common data types for monitoring structural health [5]. These approaches have the benefit of simplicity and affordability, however they may ultimately require a large number of reliable long-life transducers for sufficient coverage of large robust structures. This, in turn, can make scalability challenging [9].

For structures immersed in a fluid medium, changes in vibrational behavior will result in alterations to radiated sound that can be passively measured [10, 11]. In acoustics, three popular methods applied toward SHM are nearfield acoustic holography [12, 13, 11], guided waves [14,

15], and acoustic emissions [16, 17]. Nearfield acoustic holography (NAH) leverages the evanescent wave field measured near a surface to image acoustic sources in a three-dimensional space. NAH is attractive because it allows for high resolution spatial reconstruction of vibrating surfaces. However, by definition, it requires nearfield measurements, and thusly is not *remote* (see Section 1.2), although it is *non-contacting*. Furthermore, many receivers (on the order of hundreds) are typically required to achieve desirable resolution of surface sources. Guided wave SHM involves using surface mounted transducers to excite and record Lamb or Rayleigh waves that travel through a structure. These structural waves, typically in the ultrasound range, scatter from structural defects and the time of flight characteristics of the scattered waves can be used to detect and localize the defects [18]. As with vibrational measurements, acoustic emissions techniques are not remote and require transducers mounted to the structure, proximal to any damage to be detected. Some studies have produced positive results using acoustic emissions techniques with air coupled ultrasonic transducers placed up to 25 mm from a vibrating surface [19]. Acoustic emission monitoring is similar to guided wave SHM, however instead of using an external excitation source, transient surface waves are generated via the rapid release of strain energy local to a defect. These localized emissions are then monitored using surface mounted receiving transducers, and damage can be evaluated by counting the number of times the emissions exceed some predetermined threshold [16]. Despite a number of practical applications for these methods, none of them is truly remote from the vibrating structure. The requirement that transducers be mounted on, in, or near the test structure presents operational challenges; installation may be difficult, in situ measurement could be infeasible, and transducers aging with a structure may exhibit their own health problems indistinguishable from actual damage in the structure of interest. Furthermore, neglecting these drawbacks, direct

contact of transducers and test structures requires physical coupling which can ultimately alter the behavior of the structure.

1.1.2 Remote acoustic methods for SHM

Remote acoustic sensing of structural damage has been investigated in past studies, but to a lesser extent than the aforementioned contacting or nearfield techniques. For many years “tap” (or "gage") tests have been used, with various levels of rigor, to detect perceptible damage-induced changes in resonant frequencies of railway wheels and aeronautical structures [9]. Such gage tests, while essentially free to perform, suffer from variability depending on the intensity and location of the tap, the loading and mounting of the structure, and the skill and perceptiveness of the tester, among other factors. More rigorous approaches toward remote acoustic sensing of structural damage have been investigated for rotating systems, particularly regarding bearing health [13]. It is known that bearings, when nearing their end of life, sometimes exhibit characteristic “screeching” that can be detected with monitoring microphones. Time-frequency analysis of bearing sounds have been used to monitor and predict bearing failure, and similar studies have been done with full gearbox systems [20]. Ultimately, remote acoustic NDE requires accurate signal identification and robust processing which can be difficult in real-world environments due to unknown environmental factors like noise, propagation uncertainty, multipath propagation, and unknown reverberation [21]. Unknown reverberation and multipath present a particular difficulty as the desired vibration-induced signal becomes convolved with the acoustic environment's Green’s function which is generally unknown, and conventional techniques are, generally speaking, incapable of inverting this convolution to determine the original signal. And though truncation in the time domain is often

a means of sidestepping multipath for impulsive signals, this is infeasible for long duration structural excitation and signals, which is challenging since these types of excitation and signals are commonly of interest for passive NDE, necessitating additional signal processing methods to ensure robustness in such circumstances.

The second goal of a remote structural health analysis is to determine *where* the potential damage is located, i.e. *localization*. Damage localization is typically treated as an inverse problem, wherein measured vibrational or acoustic data of some sort are used to infer damage location parameters [5]. Within the scope of *contacting* vibration-based SHM, localization is approached primarily in two fashions. If the system under investigation permits (i.e. if it is simple), then analytical or numerical models can be exploited in conjunction with sensor measurements to estimate defect locations. Though model-based methods have the convenience of relatively simple signal processing and predictability, their primary limitation is in the unavailability of practical models for real-world systems. This is further complicated since model accuracy for a system typically worsens with system age, which is precisely when accurate SHM is most needed. The second, and increasingly more common approach, involves solving an inverse scattering problem between propagating surface waves and the damage location. This can sometimes be done relatively easily using simple time-of-flight estimates and triangulation [14] although more advanced algorithms exist for improved performance [18].

Attacking the problem of localization with *remote* sensing typically necessitates the use of a receiver array to map an acoustic source level (i.e. energy) over the structural surface of the interest. This procedure may be performed when the structure is known to be healthy (baseline) and at some later time when the condition is unknown or of interest (or, perhaps known to be defective). In theory, differences between these two source level maps will

correspond to changes presumed to arise from damage. Several such methods have been proposed, using deconvolution-based approaches such as DAMAS [22], or variants of the classic CLEAN algorithm [23, 24]. Unfortunately, these methods are only accurate under the assumption that the modeled acoustic sources are all mutually incoherent (statistically uncorrelated). While such an assumption is permissible in certain aeroacoustic applications, it is too restrictive for the purposes of detecting small changes in an extended vibrating and radiating structure. This is because each the vibrationally active regions on a vibrating structure are intrinsically coherent, at any chosen frequency, due to the coupled and modal nature of the structural response. As such, an effective method for remote localization of damage must ensure that vibrational-source coherence is correctly treated, while also maintaining sufficiently high-resolution to spatially resolve small structural-response changes.

Damage *classification*, or damage *identification*, is another common structural health monitoring goal that is prominent in the literature. Classification is challenging because the exact form in which any particular type of damage will affect structural response is difficult to predict, especially when the type, severity, and location of the damage is not known *a priori*. Worse yet, different types of damage may affect structural response in very similar ways, making distinction between damage types even more challenging. For these reasons, damage classification is typically approached using statistical, data-driven methods, the most popular modern method being machine learning. Such data-driven methods apply pattern-recognition concepts to develop statistical models with simulated or measured damage scenarios spanning the range of potential damage cases. From these scenarios, data *features* are then extracted and analyzed using either regression analysis or group classification algorithms. At any rate, the ability to effectively classify structural damage is heavily dependent on the character of

available measurement data, the damage scenarios used for data-driven-method training, and the extracted features' ability to discriminate damage types.

Remote acoustic sensing for the classification structural damage has not been well-explored in the literature. Some researchers have used remote microphone arrays to glean information about crank, piston, and clutch faults in motorcycle [25] and helicopter [26] transmissions. The use of machine learning has become somewhat common in the bearing community where it has been applied to acoustic signals for automated classification of various types of bearing faults [27, 28]. Vibrational and acoustic sensor data have also been combined for damage classification in composite plates [29]. The primary difference between using remote acoustic data, versus more conventional vibrational data (which is more extensively discussed in the literature [30, 31]), is the inclusion of the acoustic propagation in the local environment, and the increased difficulty of obtaining high-fidelity structural data. Fortunately, outside the realm of SHM, the use of remote acoustic measurements for classification has been applied to many problems including the automatic classification of road vehicles [32], agricultural sorting [33], and musical instruments [34]. These findings indicate that the extension of remote acoustic sensing to structural damage classification is feasible, although appropriate features and classification schemes must be considered to extract sufficient information from the available remotely recorded acoustic data.

Compositing these detection, localization, and classification goals together, a pathway is presented for structural health monitoring using remote acoustic sensing. As is described in this thesis, there are a number of challenges in using passive acoustic measurements for nondestructive evaluation; however, many of these concerns are addressed through the implementation of advanced signal processing methods. Still, the minimum detectable damage

size of a passive acoustic technique is generally larger than for active techniques (such as ultrasonic methods [19] or distributed vibrational sensors [7]). The benefit of remote acoustic SHM is, of course, that it is remote. The reduced cost of remote acoustic measurements and the lack of dependence on expensive, bespoke transducers, allows remote acoustic SHM methods to fill a gap in applications where other methods are impractical or even impossible.

1.2 On Remoteness

Remote sensing has been variously defined [35, 36, 37] and the term possesses different meanings in different contexts. One of the more concise definitions of remote sensing is provided by Fischer, *et al.* [38]

"Remote sensing is … the art or science of telling something about an object without touching it."

This definition of remote sensing – and most others – qualifies "remote" as essentially being equivalent to "non-contacting". In practice, however, there may be very significant differences between remote sensing methods which require some additional clarity and explanation.

Throughout this thesis the term "remote" is generally used to indicate that a measurement has been collected in the *acoustic farfield* of a radiating structure, or at least, at multiple acoustic wavelengths from the radiating structure. The acoustic farfield defines a region in which waveforms incident from some acoustic source of finite-extent appear to be locally planar (i.e. the wavefronts exhibit negligible curvature). The acoustic farfield stands in contrast to the acoustic nearfield, which is typically characterized by complex interference patterns due to energy propagating from separate regions of the source in different directions. For a single receiver, the farfield can be quantified by the Fraunhofer diffraction distance [39]

$$R_{Fr} \sim D^2/\lambda \tag{1.1}$$

where D is the characteristic length scale of the source, and λ is the acoustic wavelength [40]. The constant of proportionality for R_{Fr} varies from reference to reference, though a value of 1/4 is perhaps most common [39]. When the receiver range R is comfortably greater than R_{Fr}, the receiver is said to be in the source's farfield. Intuitively, then the Fraunhofer diffraction distance is such that the difference in path length between opposing ends (or edges) of the source is significantly less than one wavelength (subsequently, R_{Fr} is also a function of frequency). In applications where an array of multiple, spatially distributed acoustic receivers is used to record radiation from a sound source, the farfield is typically defined in terms of the characteristic length scale of the array, L

$$R_{Fr} \sim L^2/\lambda \tag{1.2}$$

The methods outlined in Chapter 3 and Chapter 5 are applicable in the acoustic farfield as defined by (1.2). In essence, they were developed for applications necessitating that measurements be collected at great distances from a structure of interest. The techniques

described. and the results presented in Chapter 4, for the acoustic localization of structural damage, are somewhat more constrained. These techniques require an array proximal to a structure, such that different regions of the structure can be angularly resolved by the recording array. The strict limits are formulated in Chapter 4, but in general measurements can still be collected at ranges on the order of multiple source lengths, D. The techniques developed in this thesis for acoustical structural health monitoring are – in all senses of the word – remote, but do require the structure to vibrate with sufficient vigor to radiate sound.

Chapter II: Methods and Tools

All techniques described in this thesis were developed in MATLAB with in-house written code. No commercial packages were used in generating any of the reported findings with the exception of the CVX convex optimization toolkit [41] for MATLAB (used throughout Chapter 4). The commercial ANSYS Workbench finite element analysis software was used to benchmark MATLAB code implemented in Chapter 5.

2.1 Array signal processing

Array signal processing encompasses a broad range of techniques used in the fields of acoustics [42, 43], radar [44, 45], medicine [46], astronomy [47], geophysics [48], communications [49], among others [50, 51]. At its most basic, array signal processing is measurement of a quantity of interest in both space and time to leverage more information or better performance than is achievable from sampling a single point in space. In applications this is typically done using a spatial distribution – or *array* – of sensors which record some state variable as a function of time (though in some applications arrays are synthesized from just a single sensor that is moved around a field [43]). Common applications of array signal processing include detection, localization, and tracking of targets, source identification, and estimation of environmental properties through measurement inversion techniques.

In acoustics, microphones (for in-air measurements) and hydrophones (for aqueous measurements) are the sensors of choice, with the generic term *receiver* often used to specify

both. Typically, acoustic pressure $p(t)$ is recorded as a function of time at all receivers, with the $p_j(t)$ signifying the recording from the j^{th} receiver of an M-element array. In classical linear acoustics [52], the acoustic pressure at time t and position vector $\boldsymbol{x}$ obeys the wave equation

$$\frac{\partial^2 p(\boldsymbol{x},t)}{\partial t^2} - c^2\nabla^2 p(\boldsymbol{x},t) = q(\boldsymbol{x},t) \tag{2.1}$$

where c signifies the speed of sound in the environment (which may depend on time and, more commonly, space) and $q(\boldsymbol{x},t)$ is some source pressure. In addition to satisfying (2.1), the acoustic pressure must also satisfy corresponding initial and boundary conditions to be fully defined. Solutions to (2.1) can be constructed for well-behaved $q(\boldsymbol{x},t)$ by leveraging the linearity of the wave-equation using a Green's function, $g(\boldsymbol{x}|\boldsymbol{x}_0,t)$, which satisfies

$$\frac{\partial^2 g(\boldsymbol{x}|\boldsymbol{x}_0,t)}{\partial t^2} - c^2\nabla^2 g(\boldsymbol{x}|\boldsymbol{x}_0,t) = \delta(\boldsymbol{x}-\boldsymbol{x}_0)\delta(t) \tag{2.2}$$

where $\delta(\boldsymbol{x}-\boldsymbol{x}_0)$ and $\delta(t)$ are the Dirac delta functions, in addition to the initial and boundary conditions. Then, the solution of (2.1) is produced via a convolution operation [53]

$$p(\boldsymbol{x},t) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(\boldsymbol{x}|\boldsymbol{x}_0,t-\tau)q(\boldsymbol{x}_0,\tau)d\tau d\boldsymbol{x} \tag{2.3}$$

Often, however, analysis of acoustics problems begins by recasting (2.1) from the space-time domain into to space-frequency domain via the Fourier transform pair

$$P(\boldsymbol{x},\omega) = \mathcal{F}\big(p(\boldsymbol{x},t)\big) \equiv \int_{-\infty}^{\infty} p(\boldsymbol{x},t)e^{i\omega t}dt$$

$$p(\boldsymbol{x},t) = \mathcal{F}^{-1}\big(P(\boldsymbol{x},\omega)\big) \equiv \frac{1}{2\pi}\int_{-\infty}^{\infty} P(\boldsymbol{x},\omega)e^{-i\omega t}d\omega \quad \text{(2.4a,b)}$$

yielding the Helmholtz equation

$$\nabla^2 P(\boldsymbol{x},\omega) + k^2 P(\boldsymbol{x},\omega) = -Q(\boldsymbol{x},\omega) \quad (2.5)$$

where $\omega = 2\pi f$ is the angular frequency and $k \equiv \omega/c = 2\pi/\lambda$ is the wavenumber. The Helmholtz formulation has the convenience of transforming the time convolution operation in (2.3) into an easier-to-implement multiplication due to the convolution theorem [53]. In the frequency domain the pressure is given by

$$P(\boldsymbol{x},\omega) = \int_{-\infty}^{\infty} G(\boldsymbol{x}|\boldsymbol{x}_0,\omega)Q(\boldsymbol{x},\omega)\,d\boldsymbol{x} \quad (2.6)$$

where $G(\boldsymbol{x}|\boldsymbol{x}_0,\omega) = \mathcal{F}\big(g(\boldsymbol{x}|\boldsymbol{x}_0,t)\big)$ satisfies $\nabla^2 G(\boldsymbol{x}|\boldsymbol{x}_0,\omega) + k^2 G(\boldsymbol{x}|\boldsymbol{x}_0,\omega) = -\delta(\boldsymbol{x}-\boldsymbol{x}_0)$. In practice, $G(\boldsymbol{x}|\boldsymbol{x}_0,\omega)$ is often a known or modeled function. One of the most fundamental Green's functions corresponds to a source in an unbounded free-space (no surfaces) with uniform sound speed and is given by

$$G(\boldsymbol{x}|\boldsymbol{x}_0,\omega) = \frac{e^{ik|\boldsymbol{x}-\boldsymbol{x}_0|}}{4\pi|\boldsymbol{x}-\boldsymbol{x}_0|} \quad (2.7)$$

If the sources are sufficiently far away from the array ('sufficiently' being defined as $L^2/4\lambda R \ll 1$ where L is the characteristic array length, λ the wavelength, and R the source-to-array distance), the magnitude factor in the denominator of (2.7) may be neglected, and a Fraunhofer approximation [54] can be used for the phase. Rather than the sources emanating from points

in space they can then be treated as arriving from angles of inclination θ and azimuth ϕ with Green's function

$$G(\boldsymbol{x}|\theta,\phi\,;\omega) = e^{ik(x\cos\theta\cos\phi\ +\ y\cos\theta\sin\phi)} \tag{2.8}$$

Further, in the case of uniform array geometry wherein the receivers are placed with equal spacing the Green's function can be further reduced. A common design is the uniform line array (ULA), which consists of equally spaced receivers along a line. For a ULA with receiver spacing d, the position of the m^{th} receiver is $\boldsymbol{x}_m = (0,(m-1)d)$, yielding a simplified Green's function

$$G(\boldsymbol{x}_m|\theta;\omega) = e^{ik(m-1)d\cos\theta} \tag{2.9}$$

It is notable that the ULA is only capable of resolving the angles of inclination, θ, due to the axially symmetry of its linear geometry.

If sufficient information is known about an environment, more complex Green's function can be constructed, but in applications where the boundary conditions are complex, numerical methods such a rays [40], boundary elements [55], or parabolic equation techniques [56] are typically employed.

2.1.1 Beamforming

Beamforming is a method of focusing an array of receivers toward a particular point or direction. This is typically done to either apply a gain to signals impinging from a predetermined direction, or to determine from which directions signals have originated. Beamforming is often referred to as *spatial filtering* due to its mathematical duality to temporal filtering [44]. At its

core, beamforming exploits the measured phase between spatially separated receivers to induce constructive interference (i.e. positive gain) in directions where signals are arriving (or anticipated), and destructive interference (i.e. negative gain) where signals are not.

Beamforming can be performed in either the time-domain or the frequency-domain. In acoustics, recorded signals are typically *wideband*, meaning that phase of the impinging signals varies significantly over the array length, or *aperture* (this is in contrast to *narrowband* signals, typically used in radar, where incident waveforms are typically modulated into a much higher carrier frequency [45]). As a result, it is generally more convenient to perform wideband beamforming in the frequency domain. Thus, for wideband signals, the output of an M-element receiver array consists of a data matrix $\hat{\boldsymbol{r}}(\omega)$ whose rows are the Fourier transforms of the M sensor recordings.

$$\hat{\boldsymbol{r}}(\omega) = [\hat{r}_1(\omega)\ \hat{r}_2(\omega)\ \cdots\ \hat{r}_M(\omega)\ ]^T \tag{2.10}$$

A model for the array output can be constructed using the K unknown source waveforms $\hat{s}_k(\omega)$ and arrival locations θ_k (note, θ_k may be an arrival angle, an (x, y) coordinate pair, or some other location parameters), in addition the modeled Green's functions $G_{mk}(\omega) \equiv G(\boldsymbol{x}_m|\boldsymbol{x}_k, \omega)$ and exogenous noise, $\hat{n}_m(\omega)$

$$\hat{\boldsymbol{r}}_m(\omega) = \sum_{k=1}^{K} G_{mk}(\omega)\hat{s}_k(\omega) + \hat{n}_m(\omega) \tag{2.11}$$

which yields the vector representation of the array output

$$\hat{\boldsymbol{r}}(\omega) = \mathbf{G}(\omega)\hat{\mathbf{s}}(\omega) + \hat{\boldsymbol{n}}(\omega) \tag{2.12}$$

where the M x K matrix $\mathbf{G}(\omega)$ is sometimes referred to as the *propagation matrix* or *array manifold*, the rows of $\hat{\boldsymbol{s}}(\omega)$ are the Fourier transforms of the K unknown source waveforms, and $\hat{\boldsymbol{n}}(\omega)$ are the Fourier transforms of the additive noise at each of the M receivers. Generally, the noise is treated as a zero-mean variable, uncorrelated at each receiver such that $E\{\hat{n}_i(\omega)\hat{n}_j^*(\omega)\} = \sigma^2\delta_{ij}$ where σ^2 is the *a priori* unknown noise variance (i.e. the 'amplitude' of the noise). The most common value for quantifying the noise magnitude in a measurement is the *signal to noise ratio* (SNR), reported in decibels

$$\text{SNR} \equiv 10\log\left(\frac{\sigma^2}{|\boldsymbol{s}|^2}\right) \tag{2.13}$$

In acoustics, SNRs above 15 dB are considered high (very little noise), whereas SNRs of -10 dB or below are considered to be significantly corrupted by noise.

It will also be convenient to define the steering vector for a single look direction k

$$\boldsymbol{g}_k(\omega) = [G_{1k}(\omega)\ \ G_{2k}(\omega)\ \cdots\ G_{Mk}(\omega)]^T \tag{2.14}$$

For example, in the case of a farfield uniform line array, recalling (2.9) the steering vector directing the array in the θ_k direction is

$$\boldsymbol{g}_k(\theta_k;\omega) = \left[1\ \ e^{i\frac{\omega}{c}d\cos\theta_k}\ \ \cdots\ \ e^{i\frac{\omega}{c}(M-1)d\cos\theta_k}\right]^T \tag{2.15}$$

For beamforming operations, it is often convenient to define the cross-spectral density matrix (CSDM) which is the M x M complex matrix of cross-correlations between the receivers.

$$\boldsymbol{\Gamma}(\omega) = E\left(\hat{\boldsymbol{r}}(\omega)\hat{\boldsymbol{r}}^H(\omega)\right) \tag{2.16}$$

$(\ \ )^H$ indicates the conjugate transpose and the expectation operator $E()$ indicate that the presence of noise (among other uncertainties), limits true evaluation of the CSDM. In practice, the CSDM can be better approximated by *snapshot averaging*, wherein the CSDM is computed for N multiple measurements of $\hat{r}(\omega)$

$$\mathbf{\Gamma}(\omega) \approx \frac{1}{N}\sum_{i=1}^{N} \hat{\boldsymbol{r}}(\omega)\, \hat{\boldsymbol{r}}^H(\omega)|_{snapshot\ i} \tag{2.17}$$

2.1.1.1 Conventional Beamformer

The simplest and most fundamental beamformer is the conventional beamformer (CBF) [57], also dubbed the Bartlett or delay-and-sum beamformer. As indicated by the latter of its names, the CBF picks a look direction, applies individual time delays to each receiver corresponding to the look direction and modeled Green's function, then sums the delayed data. In the frequency domain, the time delay is achieved via multiplication with a complex weight vector

$$\boldsymbol{w}_{CBF} = \frac{\boldsymbol{g}_k}{\sqrt{\boldsymbol{g}_k^H \boldsymbol{g}_k}} \tag{2.18}$$

where the subscript indicates the weight vector steers to the k^{th} source direction. The summed beamformer output is given by

$$y(\omega) = \sum_{i=1}^{M} w_i^*(\omega)\hat{r}_i(\omega) \ = \ \boldsymbol{w}^H(\omega)\hat{\boldsymbol{r}}(\omega) \tag{2.19}$$

The output power, being the real-valued function that is typically plotted against source location, is defined

$$P(\omega) = |y(\omega)|^2 = \boldsymbol{w}^H \hat{\boldsymbol{r}}\hat{\boldsymbol{r}}^H \boldsymbol{w} = \boldsymbol{w}^H \boldsymbol{\Gamma} \boldsymbol{w} \tag{2.20}$$

where $\boldsymbol{\Gamma}$ is the CSDM defined in (2.17).

CBF enjoys frequent use due to its simplicity and robustness, but it suffers from poor resolution (i.e. the ability to distinguish two sources arriving from nearby locations), due to its large beamwidth. Since its inception, many more beamforming methods have been developed which improve on conventional beamforming in a variety of metrics.

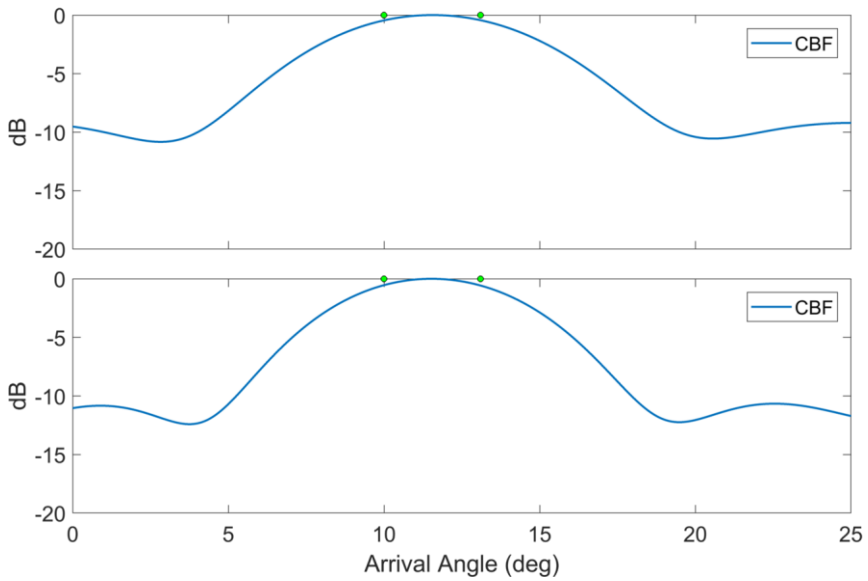

Figure 2.1: Conventional beamformer output for two nearby planar sources (green dots, arriving from 10° and 13° at a 16-element ULA with $kd = \pi$ and 0 dB SNR). In the top figure, the two sources are completely uncorrelated (incoherent, e.g. noise) with CSDMs averaged over 100 snapshots. In the bottom figure, the two sources are fully correlated (coherent).

2.1.1.2 MVDR

The Minimum Variance Distortionless Response (MVDR) beamformer, also called Capon's beamformer [58], is a variant of CBF which can give better source resolution. To do

this, the weights applied in MVDR are optimized such that the power is minimized, but the gain in the look direction is constrained to unity (the maximum)

$$\min_{\boldsymbol{w}} P(\boldsymbol{w}) \quad \text{subject to } \boldsymbol{w}^H \boldsymbol{G_k} \tag{2.21}$$

where $\boldsymbol{G_k} = [\boldsymbol{G}_{1k}\ \boldsymbol{G}_{2k} \cdots \boldsymbol{G}_{Mk}]^T$ is the steering vector pointing the array in the k^{th} direction. This optimization yield a unique steering vector

$$\boldsymbol{w}_{MVDR} = \frac{\boldsymbol{\Gamma}^{-1}\boldsymbol{g}_k}{\boldsymbol{g}_k^H \boldsymbol{\Gamma}^{-1}\boldsymbol{g}_k} \tag{2.22}$$

that, when plugged into (2.20) yields the MVDR 'spectrum'

$$P_{MVDR}(\omega) = \frac{1}{\boldsymbol{g}_k^H \boldsymbol{\Gamma}^{-1}\boldsymbol{g}_k} \tag{2.23}$$

which can be computed directly over all possible source locations/directions of interest. MVDR benefits from better resolution than CBF, due its use of the measured data to adaptively suppress interference, but it suffers when incoming sources are coherent (i.e. sources impinging from different directions are shifted/scaled multiples of one another), and when the CSDM is nearly singular (such as in a very *low* noise environment with fewer than *M* sources).

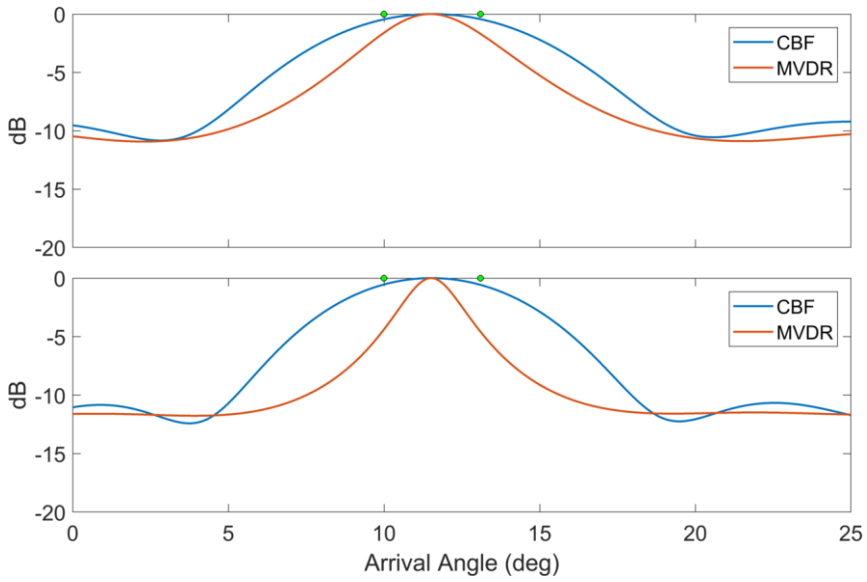

Figure 2.2: MVDR output for two nearby planar sources (green dots, arriving from 10° and 13° at a 16-element ULA with $kd = \pi$ and 0 dB SNR). In the top figure, the two sources are uncorrelated noise using 100 snapshots. In the bottom figure, the two sources are fully correlated.

2.1.1.3 MUSIC

The Multiple Signal Classification [59] (MUSIC) beamformer is the most popular of a class of algorithms known as subspace methods. These methods are based on eigenvalue decomposition of the CSDM and subsequent partitioning of the eigenvectors into a signal subspace $\mathbf{U}_s$ of rank K, and a noise subspace $\mathbf{U}_n$ of rank $M - K$, where eigenvectors of $\mathbf{U}_s$ correspond to the *largest* K eigenvalues of the CSDM. The subspace partitioning requires either *a priori* knowledge of the number of sources, K, though in low noise environments K may be estimated by the number of 'non-small' eigenvalues. In either case, once $\mathbf{U}_n$ can be computed, the 'pseudospectrum' for MUSIC is given by [59]

$$P_{MUSIC}(\omega) = \frac{\boldsymbol{g}_k^H \boldsymbol{g}_k}{\boldsymbol{g}_k^H \mathbf{U}_n \boldsymbol{g}_k} \tag{2.24}$$

P_{MUSIC} typically features very sharp peaks at the true arrival directions, with better resolution than CBF and MVDR. This behavior is intuitively explained by considering (2.11); the steering vectors $\boldsymbol{g}_k$ corresponding to the true source directions are in the signal subspace $\mathbf{U}_s$. It follows that the denominator, $\boldsymbol{g}_k^H \mathbf{U}_n \boldsymbol{g}_k$, must be very small for those $\boldsymbol{g}_k$ not in the noise subspace, resulting in the sharp peaks.

MUSIC is very common in the literature, due to its high resolution for what is a computationally cheap eigenvalue decomposition. However, similarly to MVDR, MUSIC suffers when coherent signals are present. This is due to a degeneracy in the signal subspace eigenvectors, resulting in $rank(\mathbf{U}_s) < K$. Some techniques do exist for circumventing this limitation, including subarray smoothing [60], which is discussed further in Chapter 4.

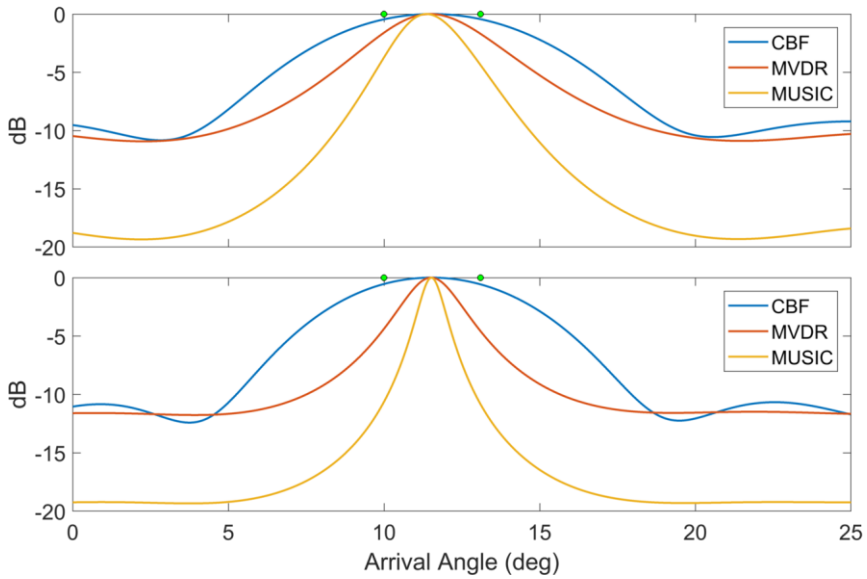

Figure 2.3: MUSIC output for two nearby planar sources (green dots, arriving from 10° and 13° at a 16-element ULA with $kd = \pi$ and 0 dB SNR). In the top figure, the two sources are uncorrelated noise using 100 snapshots. In the bottom figure, the two sources are fully correlated.

2.1.1.4 MIN-NORM

The MIN-NORM algorithm is a variant of MUSIC and part of a class of methods known as *weighted MUSIC* beamformers. These beamformers possess pseudospectra of the form

$$P_{WM}(\omega) = \frac{\boldsymbol{g}_k^H \boldsymbol{g}_k}{\boldsymbol{g}_k^H \mathbf{U}_n \mathbf{W} \mathbf{U}_n \boldsymbol{g}_k} \tag{2.25}$$

where $\mathbf{W}$ is a matrix of weights selected to improve upon MUSIC in some respect. In the case of MIN-NORM, $\mathbf{W}$ takes the form [61]

$$\mathbf{W} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \tag{2.26}$$

This choice of $\mathbf{W}$ essentially has the effect of reducing the noise subspace to the single noise eigenvector with the smallest eigenvector, ultimately yielding a smaller inner product in the denominator of (2.25) and, subsequently, peaks that are somewhat sharper than MUSIC. Despite its moderate benefits, MIN-NORM also features the same difficulties of rank deficiency as MUSIC is the case of coherent sources, and likewise exhibits poor performance in cases of multipath or mutually correlated sources.

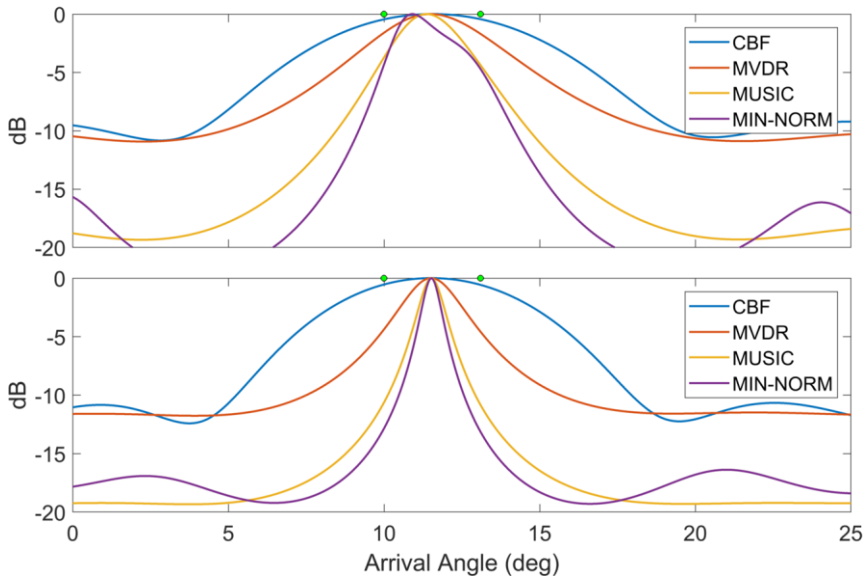

Figure 2.4: MIN-NORM output for two nearby planar sources (green dots, arriving from 10° and 13° at a 16-element ULA with $kd = \pi$ and 0 dB SNR). In the top figure, the two sources are uncorrelated noise using 100 snapshots. In the bottom figure, the two sources are fully correlated. For uncorrelated sources (top), MIN-NORM begins to indicate that two source are present.

2.1.1.5 MLE

The Maximum Likelihood Estimator (MLE) beamformer leverages the parametric model in (2.12) along with the statistical concept of joint probabilities between receiver recordings, to reduce beamforming to a multivariable optimization problem [44]. Referring back to (2.12), presuming that the exogenous noise term $\hat{\boldsymbol{n}}(\omega)$ is normally distributed with zero-mean and *a priori* unknown standard deviation σ, then the probability of obtaining a particular measurement $\hat{\boldsymbol{r}}(\omega)$ takes the form of a M-variate gaussian distribution

$$\text{Prob}\big(\hat{\boldsymbol{r}}(\omega)\big) = \frac{1}{(\pi\sigma^2)^M} e^{\|\hat{\boldsymbol{r}}(\omega) - \mathbf{G}(\omega)\hat{\boldsymbol{s}}(\omega)\|/\sigma^2} \tag{2.27}$$

which is dependent on several unknown quantities, including the signal waveforms $\hat{\boldsymbol{s}}(\omega)$, the noise amplitude σ^2, and the source directions $\boldsymbol{\theta}$, which we ultimately would like to know. The crux of MLE is that the statistically best estimate of these parameters maximizes (2.35). An equivalent objective, is to *minimize* the negative log-likelihood (e.g. the logarithm of (2.35)). This is equivalent to minimizing the function

$$l(\boldsymbol{\theta}, \hat{\boldsymbol{s}}(\omega), \sigma^2) = M \log \sigma^2 + \frac{1}{\sigma^2 N} \sum_{i=1}^{N} \|\hat{\boldsymbol{r}}(\omega) - \mathbf{G}(\omega)\hat{\boldsymbol{s}}(\omega)\| \tag{2.28}$$

over its arguments. It can be shown [62] that the best estimate for the unknown source locations is given via the optimization

$$\boldsymbol{\theta_{est}} = \arg\left\{\min_{\theta} \mathrm{Tr}\{(\mathbf{I} - \mathbf{G}(\mathbf{G}^H\mathbf{G})^{-1}\mathbf{G}^H)\boldsymbol{\Gamma}\}\right\} \tag{2.29}$$

where, again, $\mathbf{G}(\omega, \boldsymbol{\theta})$, contains the model information (Green's functions) for the possible source locations, and $\boldsymbol{\Gamma}(\omega)$ is the measured CSDM.

The MLE beamformer has two useful properties. First, with sufficient number of snapshots (samples) it converges to the theoretical lower bound (the Cramer-Rao Lower Bound) for uncertainty in $\boldsymbol{\theta_{est}}$ [62]. Second, the MLE formulation is insensitive to mutually coherent sources, as is shown in Figure 2.5. Despite its in-some-sense optimal performance, the MLE beamformer has a few drawbacks that has limited its use in the remote acoustic sensing literature. The primary difficulty is that the minimization in (2.37) is costly. The cost function is not well behaved, including large shallow canyons that seize up convergence for common gradient decent approaches. Other optimization techniques, such as genetic algorithms and

particle methods, can yield better results, but the computational complexity becomes very large as a greater number of sources are considered.

A secondary challenge of the MLE beamformer, is that the number of sources must either be known *a priori*, or the beamformer must multiple times for all different possible numbers of sources. One method for estimating the number of sources is to determine how many "large" eigenvalues exist in the decomposition of the CSDM. This approach, however, suffers at low SNRs and also for coherent sources (where the eigenvectors are degenerate, resulting in fewer 'large' eigenvalues than actual sources). Computational costs increase further when multiple frequencies (i.e. broadband sources) are considered, in which case the MLE would be computed at each desired frequency bin. Nevertheless, when it is applicable, the MLE beamformer can be an effective tool for localization and signal identification.

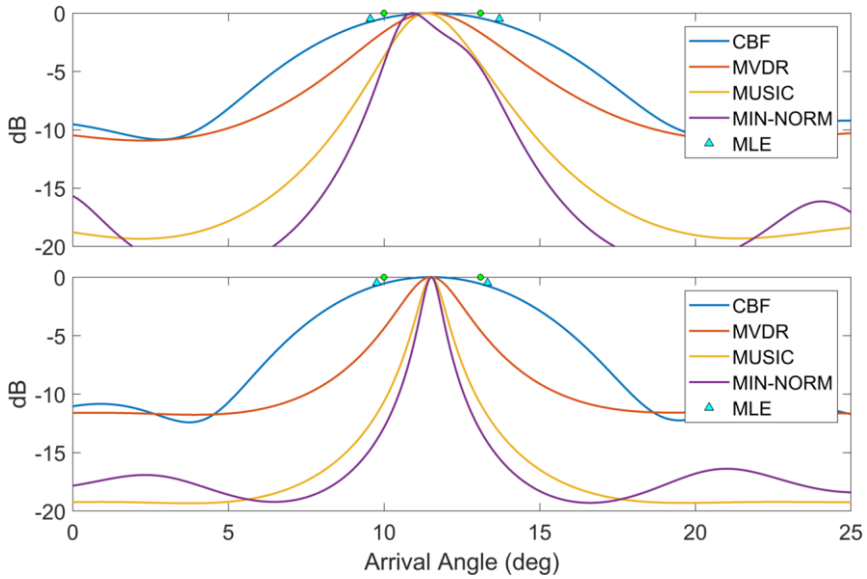

Figure 2.5: MLE output (blue triangles) for two nearby planar sources (green dots, arriving from 10° and 13° at a 16-element ULA with $kd = \pi$ and 0 dB SNR). In the top figure, the two sources are uncorrelated noise using 100 snapshots. In the bottom figure, the two sources are fully correlated. Note, the MLE beamformer is capable of localizing both sources, regardless of their coherence. However, this comes at the cost of a computationally demanding multivariable minimization task (a genetic algorithm was used to produce the results above).

2.1.1.6 SEM

The Spectral Estimation Method (SEM) is a more recent beamformer that can achieve high-resolution localization of sources via a convex optimization [63]. SEM utilizes the structure of the CSDM from (2.16)

$$\mathbf{\Gamma}(\omega) = E\left(\hat{\boldsymbol{r}}(\omega)\hat{\boldsymbol{r}}^H(\omega)\right) = \mathbf{G}\hat{\boldsymbol{s}}\hat{\boldsymbol{s}}^H\mathbf{G}^{\boldsymbol{H}} \tag{2.30}$$

the matrix $\hat{\boldsymbol{s}}\hat{\boldsymbol{s}}^H$ is named the source cross-spectral density matrix and its elements consist of the frequency domain cross-correlations of the source signals. In practice $\hat{\boldsymbol{s}}\hat{\boldsymbol{s}}^H$ is not known, since the source waveforms $\hat{\boldsymbol{s}}(\omega)$ are generally unknown. However, if all the sources can be assumed mutually incoherent (a reasonable assumption for noisy sources such as a turbulent flow field) then the source cross-spectral density matrix can be reduced to

$$E\left(\hat{\boldsymbol{s}}(\omega)\hat{\boldsymbol{s}}^H(\omega)\right) = E\left(\hat{s}_i(\omega)\hat{s}_j{}^H(\omega)\right) = \hat{s}_i(\omega)\hat{s}_j{}^H(\omega)\delta_{ij} = S_i(\omega) \tag{2.31}$$

where δ_{ij} is the Kronecker delta function and S_i is the power spectrum of the i^{th} source which is a positive definite quantity. In other words, for mutually incoherent sources, $\hat{\boldsymbol{s}}\hat{\boldsymbol{s}}^H$ reduces to a matrix with only diagonal components which represent the *a priori* unknown, but positive definite source strengths of the K sources.

The objective of SEM is then to discretize the domain of all possible source locations, then perform the minimization

$$\min_{S_k}\left(\left|\boldsymbol{\Gamma}_{m,n}^{meas}(\omega)-\sum_{l}^{K}G_{m,k}S_kG_{n,k}^{*}\right|^2\right) \tag{2.32}$$

$$\text{such that} \quad S_k > 0$$

yielding an estimate S_k for the source strength at all discretized locations. Given the positivity constraint on S_k, (2.32) is a convex optimization problem, meaning that at any point in the search space, the gradient of the cost function points in the direction of a single global minimum. This is convenient because many methods have recently been developed for quickly solving convex optimization problems [64, 65, 41]. The result of minimizing (2.32) is an estimate of source strength that features nearly discrete peaks at, or very near, the true source locations as shown in Figure 2.6.

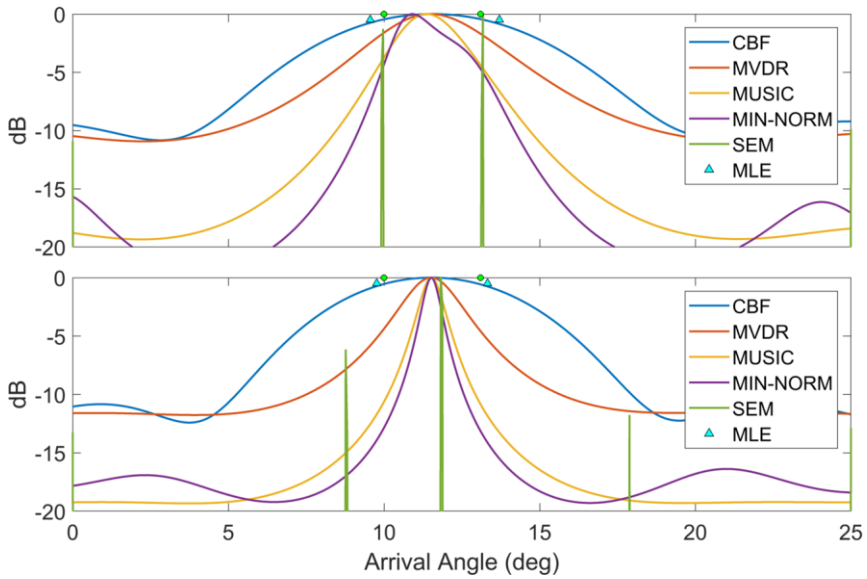

Figure 2.6: The convex optimization produce performed in SEM yields very sharp "peaks", which successfully resolves the 3° separated incoherent sources (top), but fails for two coherent sources (bottom).

The convex behavior of SEM is illustrated in Figure 2.7. A two-dimensional surface plots the error $\left|\boldsymbol{\Gamma}_{m,n}^{meas}(\omega) - \sum_{l}^{K} G_{m,k} S_k G_{n,k}^{*}\right|$ as a function of the source strength in the direction of Source 1, and the source strength in the direction of Source 2. The strength of Source 1 was set to unity, while the strength of Source 2 was set to 0.5. It is clear that the minimum (which is a global minimum for positive source strengths) is coincident with source strengths of 1 and 0.5, respectively. In practice, the locations of the sources are not known in advance, so the domain of possible source locations is discretized into J bins. Then, SEM involves a J-dimensional convex minimization. Throughout this thesis, the CVX toolbox [41] for MATLAB is used to solve these convex problems in which the dimensionality J runs anywhere from ~100 for one-dimensional localization problems, to ~10,000 for two-dimensional localization problems. To speed up performance, the proprietary Gurobi solver [66] was implemented within the CVX toolbox.

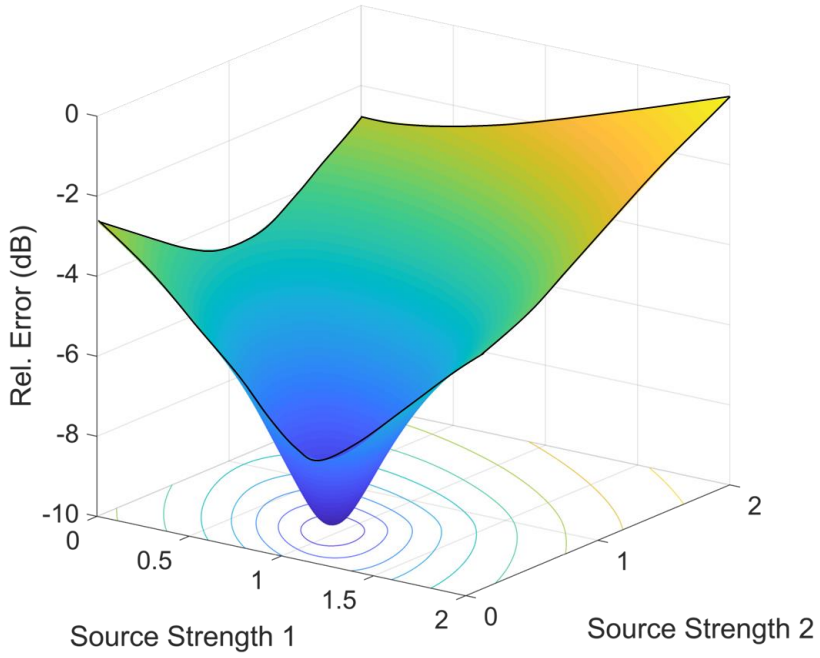

Figure 2.7: A two-dimensional error surface depicts the characteristic convex behavior of the SEM minimization. Above, only two dimensions are plotted (corresponding to the source strengths in the direction of Source 1 and Source 2, respectively). In this example Source 1 has an amplitude of 1 and Source 2 has an amplitude of 0.5, which is the location of the minima.

It must be noted that the benefits of SEM are still limited to sources which are mutually incoherent. In Chapter 4, it is shown that by exploiting translational symmetry in certain array geometries, SEM (as well as MUSIC and MIN-NORM) can be extended to applications involving numerous highly, and even perfectly, correlated sources.

2.1.2 Synthetic Time Reversal

Signal estimation is another common application of array signal processing. Source identification is the process of estimating the true waveforms of one or more broadcast sources in a propagation environment and it is an important task wireless communications both in air [49] and underwater [67]. A common problem encountered is these fields is that of a single source arriving at a receiver array from multiple directions at multiple times due to reflections from surfaces [21] or from refraction [68]. This effect is known as *multipath* due to the multiple propagation paths connecting the source to the receiver array. Multipath is particularly challenging because the multiple arrivals may interfere, distorting the recorded signal, and the arrivals are, by definition, fully correlated, which introduces additional challenges. Multipath may be problematic even when only a single reflection in present, for instance, when making underwater recordings near the ocean surface. In the other limit, when hundreds or even thousands of paths impinge upon the receiving array, multipath is often termed *reverberation*. Reverberation is a household phenomenon, characterized by a diffuse field of sound where the model of individual arrivals is traded for a statistical framework [40]. In certain scenarios,

acoustic recordings in a reflecting environment can be considered as the sum of several early-time, individual arrivals (multipath), followed by a longer duration of diffuse reverberation.

Signal estimation in the presence of multipath may be improved when with the use array signal processing. If an M-element array is used to capture the individual receiver recordings $r_m(t)$, where m serves as the receiver index, then the simplest scheme for estimating the true broadcast signal $s(t)$ is an incoherent average of the recordings

$$s_{incoh}(t) = \sum_{m=1}^{M} r_m(t) \tag{2.33}$$

Unfortunately, this average in no way accounts for the environmental effects, and is usually ineffective at improving the signal estimate (however, (2.33) may help suppress incoherent noise effects). If the incoherent average is performed in the frequency domain

$$\hat{s}_{incoh}(\omega) = \sum_{m=1}^{M} \hat{r}_m(\omega) \tag{2.34}$$

Recalling (2.11) (and neglecting the noise term), $\hat{r}_m = G_{mk}\hat{s}$ where G_{mk} is the *unknown* Green's function from the source to the m^{th} receiver, accounting for multipath. It follows from (2.34)

$$\hat{s}_{incoh}(\omega) = \sum_{m=1}^{M} G_{mk}(\omega)\hat{s}(\omega) = \hat{s}(\omega)\sum_{m=1}^{M} G_{mk}(\omega) \tag{2.35}$$

For $\hat{s}_{incoh}(\omega)$ to accurately approximate the true source waveform (up to a time-delay and normalization constant), then $\sum_{m=1}^{M} G_{mk}(\omega)$ must be proportional to $e^{ia\omega}$, for some real

constant a. Unfortunately, this condition is generally not met and as such, $\hat{s}_{incoh}(\omega)$ is typically a poor estimator of $\hat{s}(\omega)$.

The beamforming techniques discussed in Section 2.1.1 can also be beneficial for source estimation. Using a beamformer of choice to first determine the various source-to-array paths, the array can be steered in the direction of one of these paths, effectively applying a negative gain to all other (interfering) paths. The beamformed estimate of the source signal is

$$\hat{s}_{BF}(\omega) = \boldsymbol{w}^H(\theta_m; \omega)\hat{\boldsymbol{r}}(\omega) \tag{2.36}$$

where $\boldsymbol{w}^H(\theta_m; \omega)$ are the beamformed weights used to steer in the direction of the θ_m path (where θ_m may be an angle, set of angles, spatial coordinates, etc.). In practice, conventional beamforming is a common beamformer for this practice, due to its robustness to coherent effect.

Synthetic Time Reversal (STR) is a technique for signal estimation in shallow water channels where significant amounts of multipath are present due to reflection from the sea surface and the seafloor [69]. STR is specifically a method of *blind deconvolution*: **deconvolution** in that the impulse response of the water channel is deconvolved from the receiver recordings $r_m(t)$ to approximate $s(t)$, and **blind** because this operation is done with no prior knowledge of the channel's impulse response *or* the signal waveform.

STR accomplishes blind deconvolution by first estimating the M source-to-receiver Green's functions, $G_{mk}(\omega)$, then using these to estimated Green's functions to deconvolve each receiver output. Mathematically, the STR estimate of the source signal is [70]

$$\hat{s}_{STR}(\omega) = \sum_{m=1}^{M} \tilde{G}_{mk}^{*}(\omega)\hat{r}_m(\omega)$$
$$= \sum_{m=1}^{M} \left\{ \frac{\hat{r}_m(\omega)e^{-i\alpha(\omega)}}{\sqrt{\sum_{l=1}^{M}|\hat{r}_m(\omega)|^2}} \right\}^{*} \hat{r}_m(\omega) \tag{2.37}$$

where $\alpha(\omega)$ is a phase correction factor of the form

$$\alpha(\omega) = \arg\left(\sum_{m=1}^{M} w_m(\omega)\hat{r}_m(\omega) \right) \tag{2.38}$$

If the weights $w_m(\omega)$ can be chosen such that $\arg(\sum_{m=1}^{M} w_m(\omega)G_{mk}(\omega))$ is, at most, linear in ω, then $s_{STR}(t)$ will be equivalent to $s(t)$ up to an arbitrary time shift and normalization constant (meaning that STR is insufficient for determining source range). It can be shown [71] that the beamformer weights from (2.18) are sufficient to satisfy this condition.

Though other blind deconvolution techniques exist [72] [73], STR is particularly convenient because it requires essentially no more processing than a typical beamforming computation. To showcase the performance of STR, a simulation was performed wherein five plane waves impinged upon a 15-element uniform line array (d = 22.4 cm) from 6.7°, 16.3°, –12.3°, –23.3°, and 26.8° with time delays of 0, 5, 10, 20, and 25 ms and relative amplitudes of 0.9, 1.2, 0.5, 0.2, and 0.6, respectively (these values were not arbitrary, rather they are based on the CAPEX09 dataset taken in Lake Washington off the coast of Seattle [70, 74]). The true signal waveform, broadcast along each path, was a linear chirp from 0-5000 Hz, with a 0.1 second duration. The chirps frequency content of the chirps were enveloped such that the frequencies 2100, 2800, and 3700 Hz were bandpassed (effectively mimicking the response of a system with resonances centered at these three frequencies).

Figure 2.8 compares the performance of the three aforementioned signal estimation methods for the simulated multipath scenario detailed above. The top row shows the true signal waveform in both the time and frequency domains. The second row shows the estimated signal waveform using a rudimentary incoherent average across the 15 receivers. The second resonance is split in the time domain waveform, and the magnitude of the Fourier transform indicates that much of the energy at the higher frequencies has been lost through interference. The third row showcases the estimated signal waveform using the CBF weights to steer in the direction of the reference ray. Though beamforming does a better job at reconstructing the signal, it is clear from the Fourier transform that interference effects are still present in the estimate. Finally, the last row shows the signal waveform estimate using STR. The time-domain reconstruction is very close to the actual broadcast signal (the normalized cross-correlation is 98%), and the frequency-domain plot indicates that STR yield the best estimate of the signal energy as a function of frequency.

These simulation results are promising for structural health monitoring applications because the availability of an accurate estimate of a *true* signal waveform is necessary for detecting *small* changes in said waveform (changes that might be due to damage). STR provides a computationally affordable method of signal estimation that can be leveraged to improve robustness of acoustic techniques in reverberant environments.

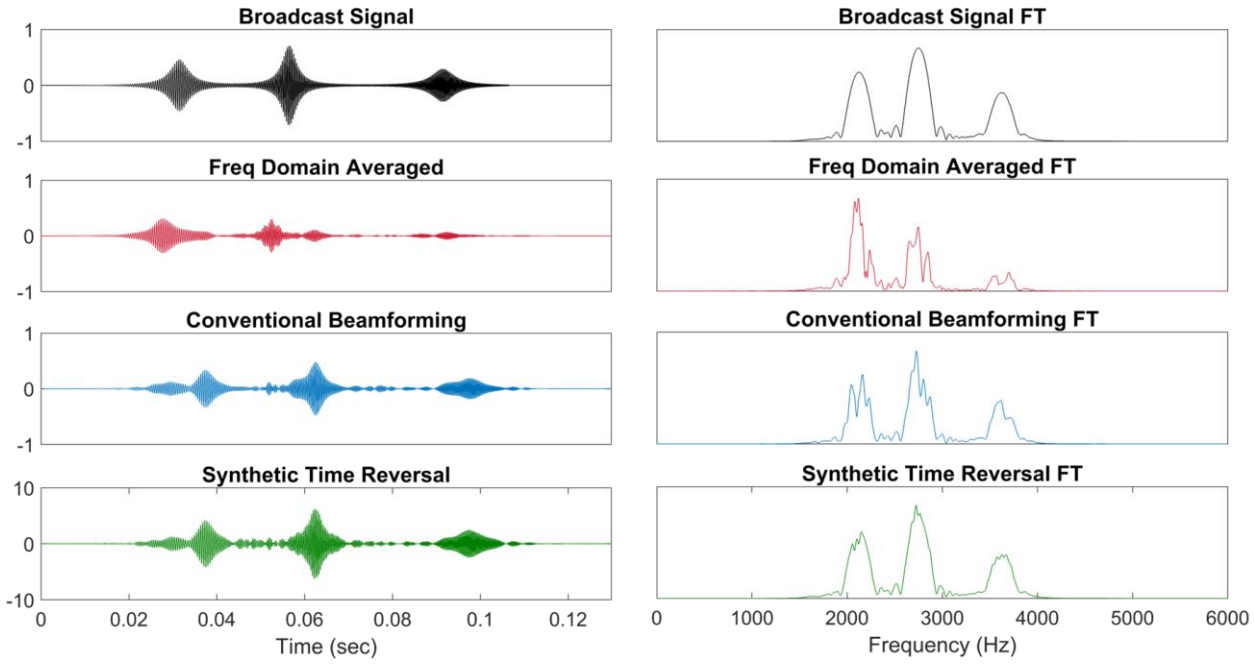

Figure 2.8: Simulation results comparing multiple signal estimation techniques applied to a multipath environment. In black, the top row shows the true source waveform (left) and magnitude Fourier transform (right). In red is shown the signal estimates using an incoherent average across the 15 receivers. In blue is shown signal estimates using a conventional beamformer to steer in the direction of a reference ray. The bottom green curves show the best signal estimates, using STR, with the same beamformer weights as the blue curves above.

2.2 Experimental apparatus

The following experimental tools were frequently used to test and validate the methods developed in this thesis and as such it is beneficial to describe them here.

2.2.1 Microphone array

16 PCB 130E20 electret microphones were used to collect all acoustic measurements in this thesis, unless otherwise stated. Each microphone was calibrated to flat frequency response ($\pm$2 dB), from 20-20,000 kHz. The microphones were used in a variety of array geometries, the

most common being a 2-inch spaced linear array mounted either vertically or horizontally, as shown in Figure 2.9.

Figure 2.9: Linear array of 16 PCB 130E20 electret microphones used in this thesis. This horizontal configuration with 2-inch spacing was one of a multiple of array configurations available.

For two-dimensional source localization (see Chapter 4, Section 3), the microphones were configured in a Cartesian planar geometry, shown in Figure 2.10. For several experiments, a planar array of greater size than 4x4 was required. To accomplish this, the base 4x4 array was affixed to a two-axis automated gantry and used to take an acoustic measurement. The 4x4 array was then translated in the x-axis and used to take a second, nominally identical acoustic measurement. This was repeated two more times for a total of four 4x4 measurements, and the resulting data were 'stitched' to synthesize a single 8x8 array measurement. This stitching procedure was found to be effective for acoustic signals reproducible from trial-to-trial, such as deterministic chirps and sinusoids.

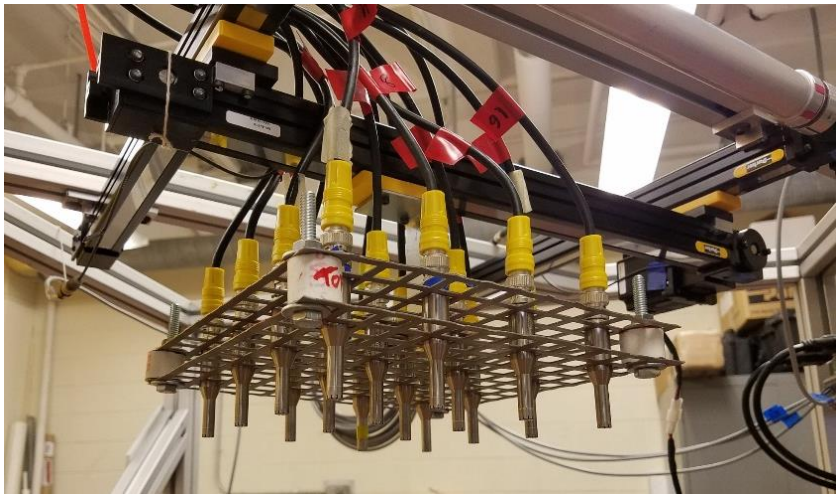

Figure 2.10: 4x4 Cartesian array of PCB 130E20 electret microphones. The 4x4 was affixed to a 2-axis automated gantry for precision translation, allowing synthesis of larger arrays up to 8x8 elements in size.

Figure 2.11 provides the schematic data flow used in all reported acoustic and vibroacoustic measurements in this document. A computer was used to forward output/drive signals to a National Instruments PXIe-6368 DAQ. These signals were converted to usable voltages which were then forwarded to and a source. The two primary sources included home audio speakers for acoustic tests and a plate vibration apparatus (see Section 2.2.2 Plate vibration rig) for vibroacoustic tests. Acoustic responses from the source(s) were simultaneously captured at the aforementioned microphone receiver arrays. The microphone outputs were then passed through a signal conditioner, a set of Krohn-Hite 3364 analog bandpass filters, routing back to the DAQ for final storage in the computer. Each individual microphone was calibrated to ±2 dB within their 20-20,000 Hz operating band [75]. The microphones also exhibit a nearly omnidirectional response within the operating band, and so individual receiver transfer function were neglected (assumed unity) for all subsequent array processing methods.

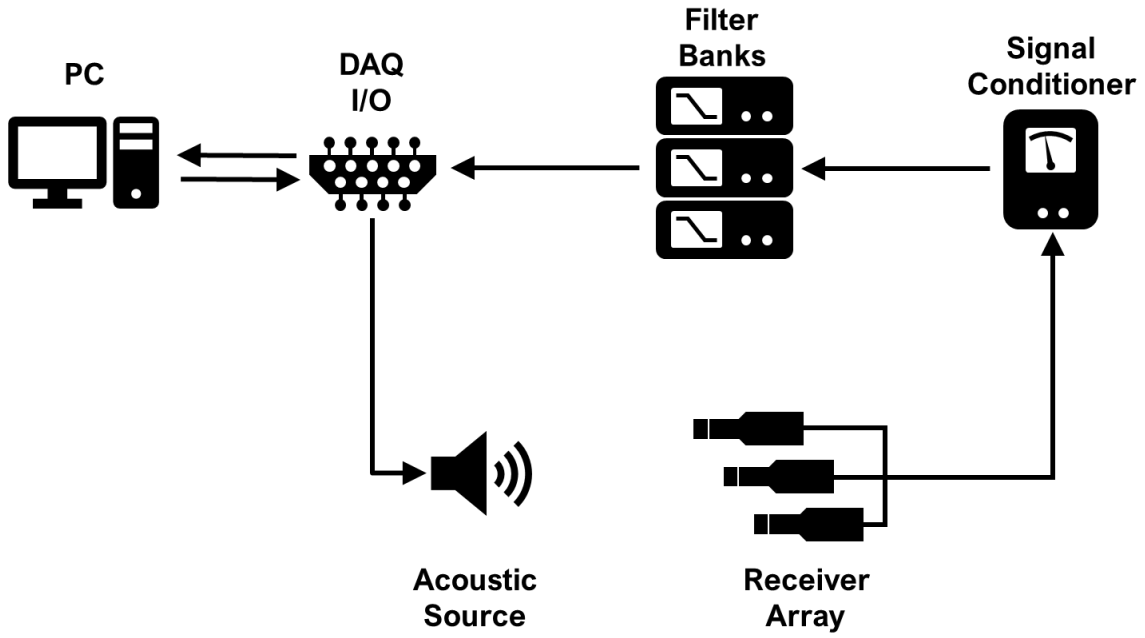

Figure 2.11: Data flow schematic for the acoustic and vibroacoustic measurements reported in this thesis.

2.2.2 Plate vibration rig

An apparatus was designed and built for controlled, precise, and repeatable vibration of clamped square plates. This plate vibration rig consists of a rigid MIC-6 aluminum base, on which plates of nominal size 15” x 15” could be clamped. A 1/8” steel flange above the test plates allowed for continuous, distributed clamping along the plate perimeter, with the clamping force applied from 16 over-center toggle clamps evenly distributed around the plate. Each toggle clamp was rated for a clamping force of approximately 200 lbs. A schematic of the plate vibration rig is shown in Figure 2.12. The effective vibrating region of the clamped plates is 12” x 12”.

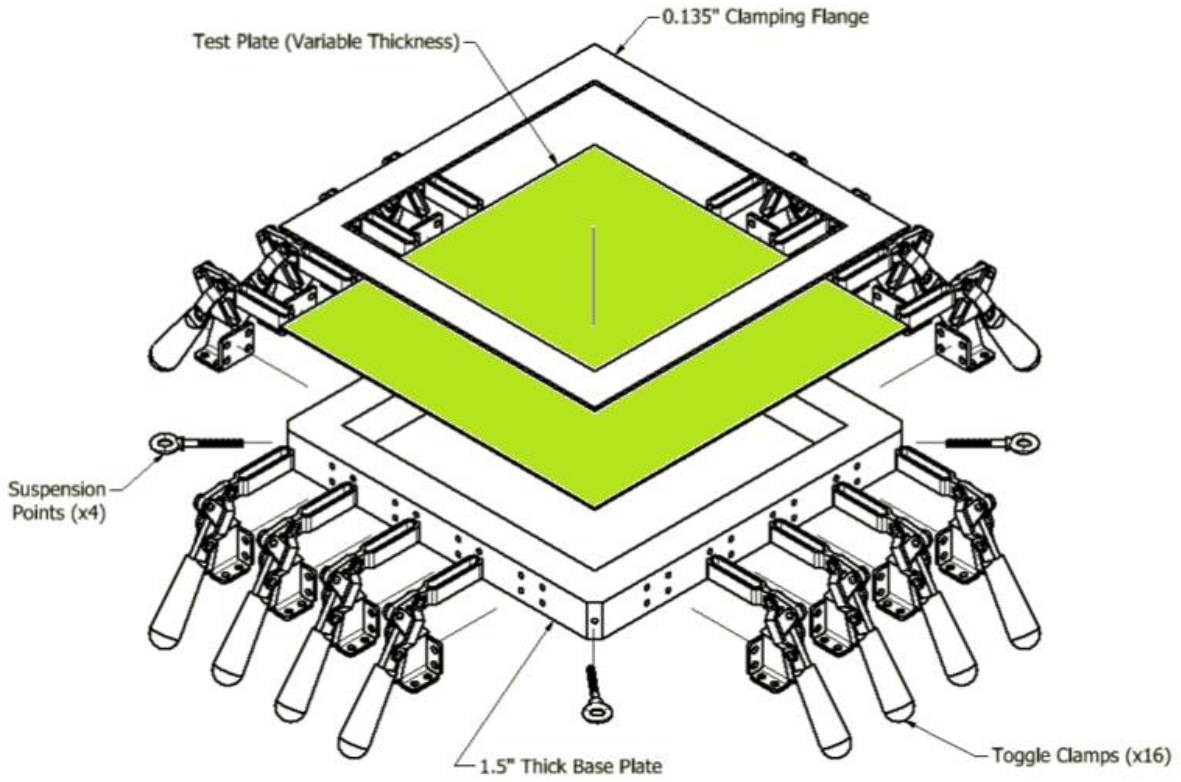

Figure 2.12: Exploded view of the plate vibration apparatus. The test plate is highlighted in light green.

Eye bolts mounted into each corner of the MIC-6 base allowed the system to be suspended from four bungee cords. A Modal Shop K2007E01 series electrodynamic shaker was used to drive the rigid base from below, shown at the bottom right of Figure 2.13. The energy input from the shaker base-excited the clamped test plates, exciting many vibrational modes without requiring direct connection with the plate. The shaker was rated for a nominal load limit of 7 lbf, and an operational bandwidth of DC-9000 Hz. A PCB 208C03 force couple, in-line with the shaker, was used to simultaneously record input forcing, so that frequency response functions of various outputs could later be computed. The geometry of the system was selected to generate several dozen vibrational plate modes of a 1/16" aluminum (or steel) plate within the 100-2000 Hz frequency band. The corresponding coincidence frequency [10] is

approximately 4500 Hz, though lower frequency modes still radiate significantly into the farfield. Throughout this thesis, however, several different plate thicknesses, materials, and excitation bandwidths have been investigated. The modularity of the system also enabled the controlled application of real and synthetic damage, as plates could either be quickly exchanged by toggling the clamps, or by simply applying damage to plates already mounted into the rig.

Figure 2.13: The plate vibration rig and microphone array (left). A close-up of a mounted 1/16" aluminum plate is shown (top right). An electrodynamic shaker (bottom right) drives the system from the rigid base, thereby base-exciting the test plate. A force couple, in-line with the shaker, allows for simultaneous measurement of excitation forcing into the system.

2.2.1 Laser Doppler vibrometer

An OMS LP-01 single-point laser Doppler vibrometer (LDV) was used to measure plate deflection as a function of frequency, material, condition, excitation, etc. The LDV provided non-contacting insight into the vibrational response of the various test plates, which complimented and informed subsequent acoustic measurements. As depicted in Figure 2.14, the LDV operational theory exploits the Doppler shift that a moving structure applies to a near-IR laser beam. Isolating this effect allows for the direct measurement of plate velocity at the inspection point (and in the direction of the laser), which may then be integrated to determine plate displacement. Additional displacement frequency response functions can be computed with simultaneous measurement of the input forcing, obtained via the in-line force couple.

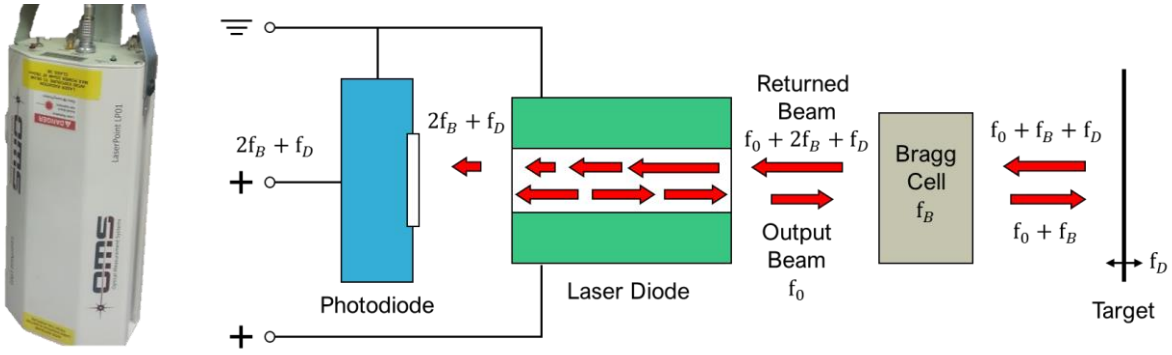

Figure 2.14: Schematic of the physical mechanisms behind a laser Doppler vibrometer.

To achieve two-dimensional measurements of plate velocity, displacement, acceleration, etc., the single-point LDV was mounted to an automated two-axis gantry. This allowed the LDV to raster along the plate, collecting an ensemble of point measurements which were later stitched into two-dimensional plots. Figure 2.15 shows an image of the scanning LDV system, facing down toward the vibrating plate rig. Figure 2.16 compares LDV measured mode shapes of a 12 x 12 x 1/16 inch aluminum plate against finite element simulations using

ANSYS Workbench. Very good agreement is exhibited in mode shape, and modal frequencies for the measured plate were found to be approximately 8% below what is predicted using finite elements or analytical techniques [76]. It is likely that this slight reduction in the modal frequencies is due to the aluminum plate materials being cold-rolled (thereby altering its stiffness properties from that of bulk aluminum, which was used in simulation), and also possibly due to the toggle clamps not providing a perfectly clamped boundary, potentially making the effective size of the plate slight larger than the presumed 12 x 12 inches.

Figure 2.15: Laser Doppler vibrometer setup. The single point OMS LP-01 LDV is mounted to a two-axis automated gantry with 46 x 46 cm area of motion, allowing for full surface scans of the test plate apparatus.

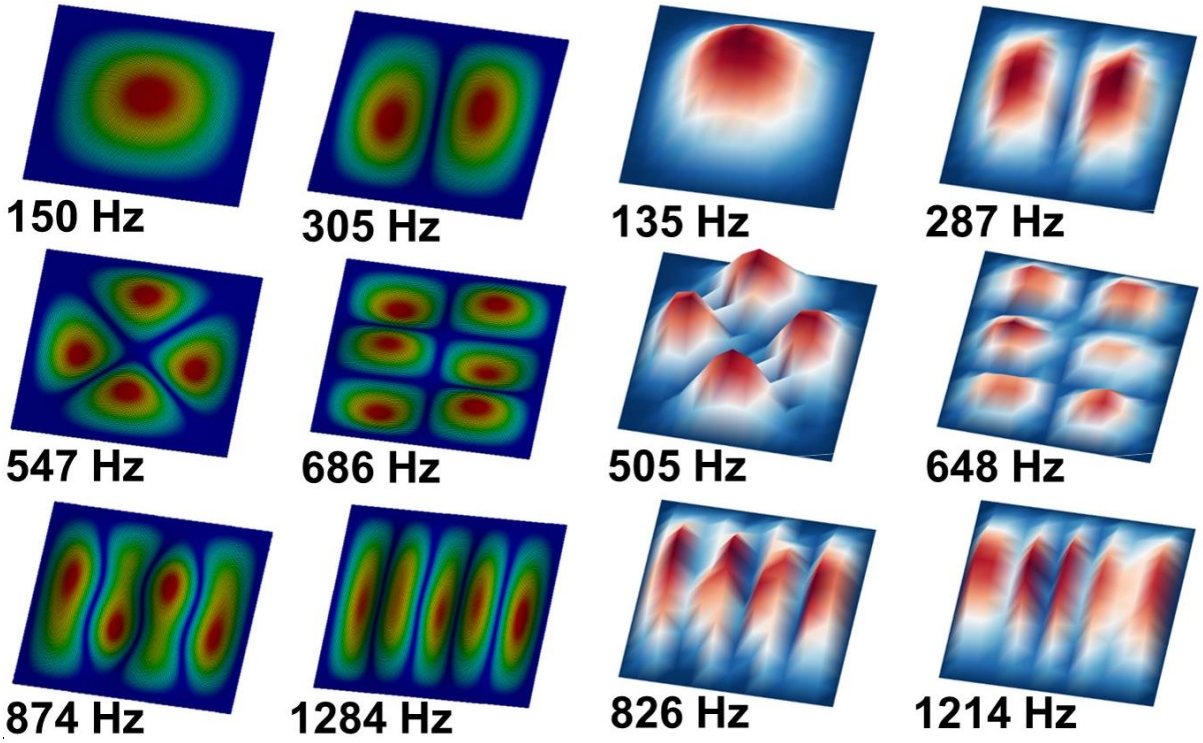

Figure 2.16: Comparison of six select mode shapes of an undamaged 30x30x0.16 cm clamped aluminum plate, using finite element methods (ANSYS 15.0) (left), and from LDV measurements (right). The ~8% reduction in measured resonance frequencies is likely due to the modulus of elasticity of the cold-rolled aluminum plates being less that of bulk aluminum.

2.3 Finite element analysis

To better predict vibroacoustic results, validate measurements, and design various experimental apparatus, finite element methods were used to numerically model the structural responses of plates. A plate bending finite element analysis (FEA) was developed in MATLAB to allow for better data flow and user-control than proprietary software packages.

Following is a brief description of the finite element formulation for plate bending. For thin isotropic plates of thickness h, elastic modulus E, Poisson's ratio ν, and areal density ρ, undergoing transverse motion with no damping, the governing partial differential equation (using Kirchhoff-Love theory) is [76]:

$$D\left(\frac{\partial^4 w}{\partial x^4}+\frac{\partial^4 w}{\partial x^2 \partial y^2}+\frac{\partial^4 w}{\partial y^4}\right)-\rho\ddot{w}=\ q(x,y) \tag{2.39}$$

where $D = Eh^3/12(1-\nu^2)$ is the plate flexural rigidity, $q(x,y)$ is the prescribed transverse pressure loading, and $w(x,y)$ is the unknown transverse displacement. Boundary conditions must also be prescribed at the plate boundary, with the most common forms being 'clamped' (prescribed zero displacement and zero slope), 'simply-supported' (prescribed zero displacement and zero moment), and 'free' (prescribed zero moment and zero shear). Closed form analytical solutions are generally only available for simple geometries and boundary conditions; for rectangular plates this requires at least two opposing edges be fully simply-supported [76].

The goal of FEA is to reduce the complex governing PDE into a finite system of linear equations of the form

$$\mathbf{M}\ddot{\mathbf{d}}+\mathbf{C}\dot{\mathbf{d}}+\mathbf{K}\mathbf{d}=\mathbf{f} \tag{2.40}$$

where $\mathbf{d}$ is a vector of generalized displacements, $\mathbf{M}$ is the system's *global mass matrix*, $\mathbf{C}$ is the *global damping matrix*, $\mathbf{K}$ is the *global stiffness matrix,* and $\mathbf{f}$ is a vector of generalized forces acting on generalized displacements. In this thesis, plates were modeled using simple 4-node quadrilateral elements. Only transverse plate bending was considered (no in-plane stresses), and as such, each node possessed three degrees of freedom (DOF): transverse displacement, *x*-axis rotation, and *y*-axis rotation. In total, each element possessed 12 DOF, depicted in Figure 2.17.

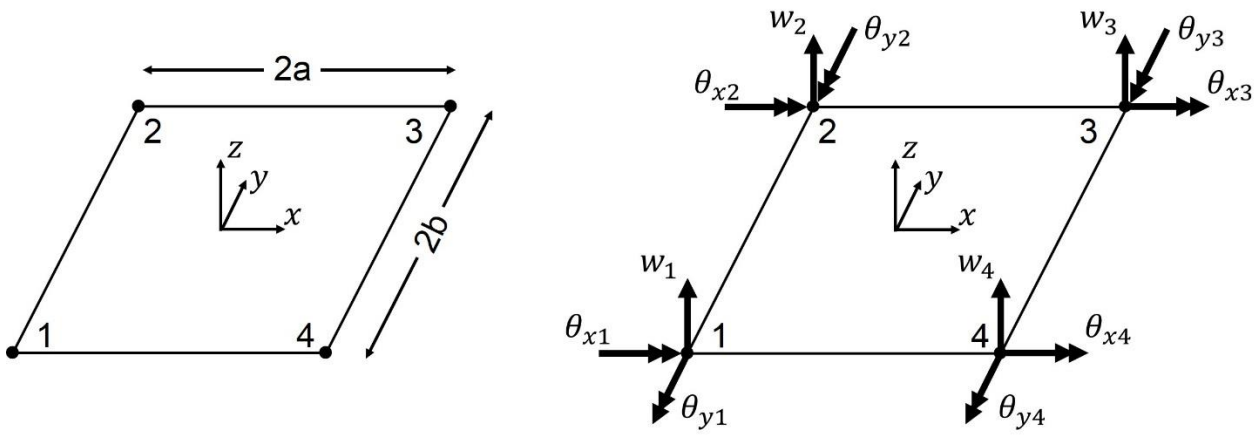

Figure 2.17: Generic 4-node, 12 degree of freedom plate bending element.

Derivations for the quadrilateral bending element mass and stiffness matrices, **m** and **k** along with the applied load vector **f**, can be found in numerous undergraduate and graduate finite element method texts [77, 78, 79]. Simply put, since each element contains 12 DOF, any polynomial with 12 undetermined coefficients can be used to model the transverse displacement $w(x, y)$ within the element. In this work displacement takes the form

$$w(x,y) = a_1 + a_2x + a_3y + a_4x^2 + a_5xy + a_6y^2 + a_7x^3 + a_8x^2y + a_9xy^2 + a_{10}y^3 + a_{11}x^3y + a_{12}xy^3 \tag{2.41}$$

It is noted that this polynomial is not complete to fourth order, but this selection of terms generally yields the best performance for plate bending problems [77].

To ensure that the plate bending FE code was performing properly a benchmark test was conducted against ANSYS Workbench 15.0. Both programs were used to model the center-plate displacement for a 12 x 12 x 1/16 inch clamped aluminum plate, with an loss factor applied loss factor $\eta = 0.008$, such that $E \rightarrow E(1 + i\eta)$. In both scenarios the plate was harmonically excited with a uniform pressure of 14.5 psi from 0-2000 Hz. The displacement

curves from both programs are plotted in Figure 2.18. Similarly, a comparison of predicted modal responses at 874 Hz is provided in Figure 2.19. Both programs are in good agreement, within the 0-2000 Hz bandwidth. Additional finite element results for plates with various types and severities of damage are reported in Chapter 5.

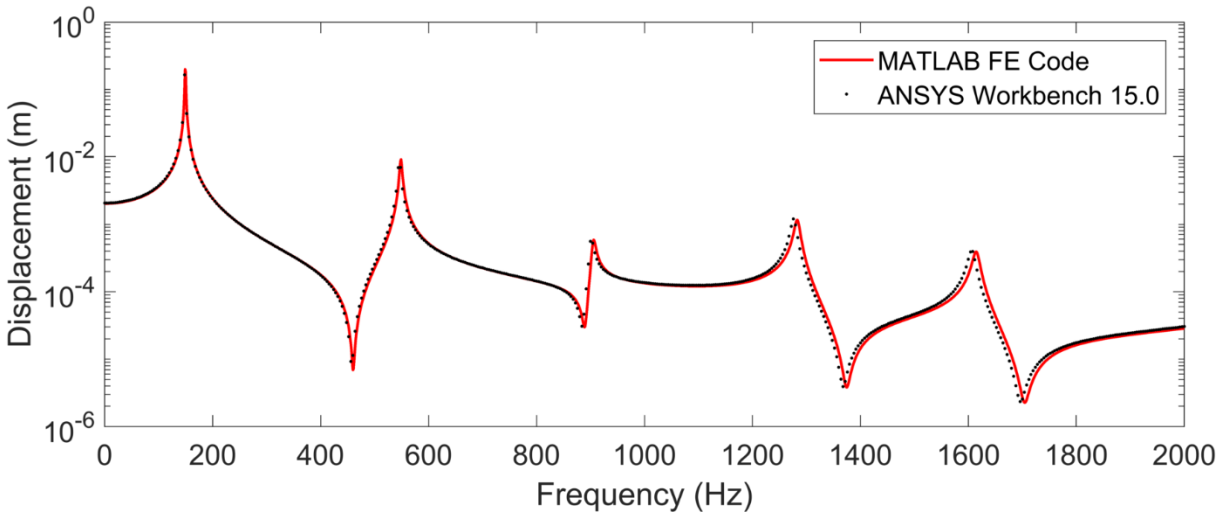

Figure 2.18: Benchmark test for MATLAB plate bending code against ANSYS Workbench 15.0. The two methods are in good agreement, both in resonance peak locations and amplitudes.

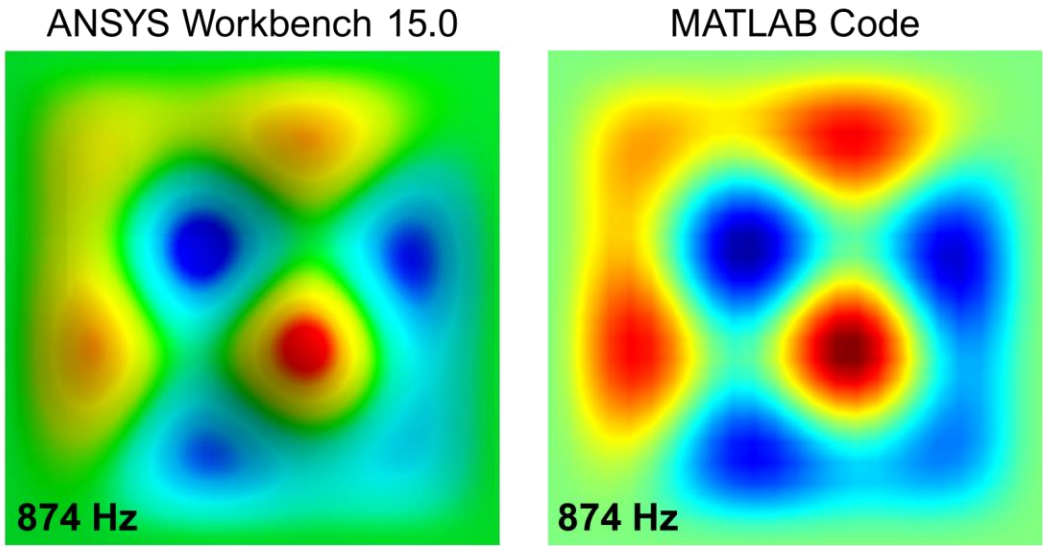

Figure 2.19: Benchmark comparison of the harmonic response of a clamped plate at 874 Hz. The ANSYS Workbench 15.0 result (left) and the MATLAB code (right) agree quite well.

Chapter III: Detection of Damage

3.1 Detection in reverberant environments

The presence of damage in a mechanical structure often changes its vibrational characteristics. For example, resonant frequencies may shift, and mode shapes may be distorted in a damaged structure relative to a healthy counterpart [1,2]. Damage may also affect structural stiffness, damping, and vibration-energy loss mechanisms, so resonant peak widths may change, and nonlinearities may even arise. A schematic example of some common effects of damage on a structure's frequency response is shown below in Figure 3.1.

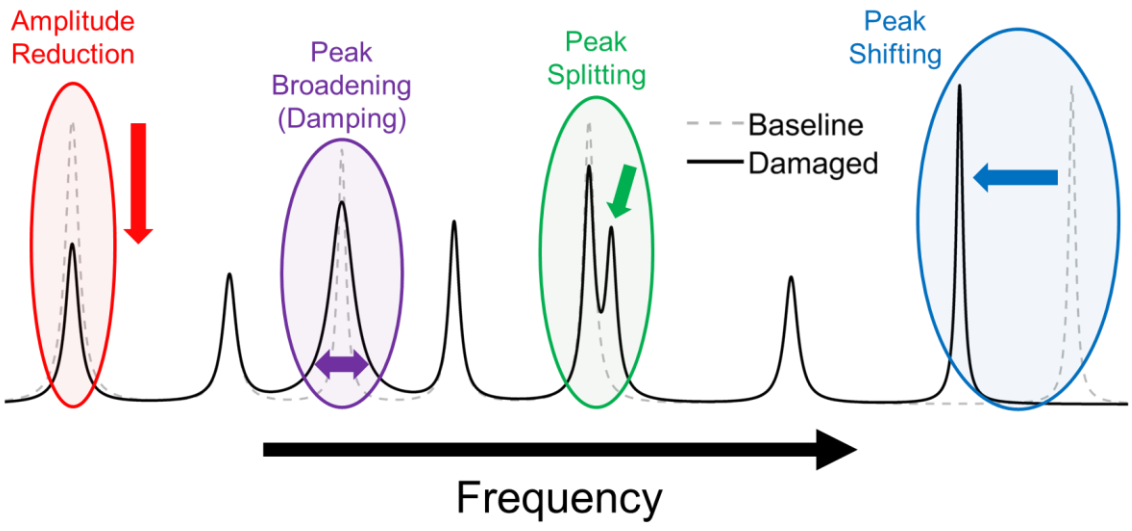

Figure 3.1: Schematic illustrating some of the common effects of damage on a generic frequency response function. The baseline (healthy) frequency response function is shown as a gray dashed line, while the nominally damaged frequency response function is in solid black.

For structures immersed in a fluid medium, such changes in vibrational behavior will result in alterations to radiated sound that can be passively measured [3]. Several nondestructive evaluation (NDE) techniques exist which utilize measurable vibroacoustic changes for defect detection [4-6]. Many studies have investigated use of discrete measurements of acceleration and strain to monitor structural health. These approaches have the benefit of simplicity and affordability, however they may ultimately require large numbers of transducers for sufficient coverage for large structures which can make scalability challenging. In acoustics, two popular methods of NDE are nearfield acoustic holography [7,8] and acoustic emissions [9,10]. Nearfield acoustic holography (NAH) leverages the evanescent wave field measured near a surface to image acoustic sources in a three-dimensional space. NAH is attractive because it allows for high resolution spatial reconstruction of vibrating surfaces. However, by definition, it requires nearfield measurements, and thusly is not *remote*, although it is *non-contacting*. Furthermore, many receivers (more than 100) are typically required for good results. Acoustic emissions involve using surface mounted transducers to excite and record Lamb or Rayleigh waves that travel through a structure. These structural waves, typically in the ultrasound range, scatter from structural defects and the time of flight characteristics of the scattered waves can be used to detect and localize the defects. As with vibrational measurements, acoustic emissions techniques are not remote and require transducers mounted to the structure, proximal to any damage to be detected. Some studies have produced positive results using acoustic emissions techniques with air coupled ultrasonic transducers placed up to 25 mm from a vibrating surface [11]. Such measurements are non-contacting, but they are not remote measurements. For each of these methods the requirement of transducers being in or near the test structure presents operational challenges; installation may be difficult, in situ measurement could be infeasible,

and transducers aging with a structure may exhibit their own health problems indistinguishable from actual damage in the structure of interest. Furthermore, neglecting these drawbacks, direct contact of transducers and test structures requires physical coupling which can alter the behavior of the structure.

Remote acoustic sensing of structural damage has been investigated in past studies, but to a lesser extent than the aforementioned techniques. For many years “tap” (aka "gage") tests have been used, with various levels of rigor, to detect perceptible damage-induced changes in resonant frequencies of railway wheels and aeronautical structures [12]. Such gage tests, while essentially free to perform, suffer from variability depending on the level and location of the tap, the loading and mounting of the structure, and, of course, the skill and perceptiveness of the tester. More rigorous approaches toward remote acoustic sensing of structural damage have been investigated for rotating systems, particularly regarding bearing health [8]. It is known that bearings, when nearing their end of life, sometimes exhibit characteristic “screeching” that can be detected with monitoring microphones. Time-frequency analysis of bearing sounds have been used to monitor and predict bearing failure, and similar studies have been done with full gearbox systems [13]. Ultimately, remote acoustic NDE requires accurate signal identification which can be difficult in real-world environments due to unknown environmental factors like noise, propagation uncertainty, and multipath and reverberation [14]. Reverberation and multipath presents a particular difficulty as the desired broadcast signal becomes convolved with the acoustic environment's Green’s function which is generally unknown, and conventional techniques are incapable of inverting this convolution to determine the original signal. And, though truncation in the time domain is often a means of sidestepping multipath

for impulsive signals, this is infeasible for long duration signals and such signals are commonly of interest for passive NDE.

In this paper, remote acoustic array measurements of sound radiated by a vibrating plate structure in a reverberant environment are used in concert with array signal processing techniques to detect small mechanical changes to the mounting of the vibrating plate. Measurements were conducted in a laboratory setting which provided a suitable unknown reverberant environment. To accurately identify the acoustic signature of the plate structure in the presence of strong reverberation, array recordings are blindly deconvolved using synthetic time reversal [15,16]. With this reconstruction technique, a baseline comparison approach is employed for detection of mechanical changes, using time domain cross correlations as a quantitative metric. The performance of the technique for varying levels of artificially imposed defects is analyzed using statistical classification techniques and results are presented using conventional receiver operating characteristic (ROC) curves.

3.1.1 Experimental methods

The primary experimental set up used throughout this thesis is shown in Figure 3.2. The radiating test structure consists of a 30 x 30 x 0.3 cm square aluminum plate clamped along all sides to a nominally-rigid aluminum base by 16 toggle clamps with nominally identical loading. The plate and its base are suspended 43 cm above the laboratory floor via elastic cables. The plate's nominally-rigid test base is directly coupled to a 30-N electrodynamic shaker through a force coupler. Eight-second duration linear frequency sweeps from 100-2000 Hz were created on a computer and sent through a digital-to-analog converter to the shaker to excite vibrations of the test plate. The shaker was connected directly to the thick aluminum base onto which the

plate was clamped. This base-excitation configuration enabled sufficient excitation of all vibrational modes within the excitation bandwidth, without requiring direct coupling to the plate. The vibrating plate's resulting acoustic signature was then radiated into the laboratory environment and recorded with a vertical array of 15 5.08-cm-spaced microphones placed 1.0, 2.0, or 3.0 meters from the vibrating plate. The microphone-array signals were then conditioned, band-passed filtered between 50 and 5,000 Hz, passed through a DAQ system, sampled at 20 kHz per channel, and stored on a laboratory computer system for later analysis. The bandwidth of the input forcing was chosen to encompass approximately the first dozen vibrational modes of the test plate.

Figure 3.2: Experimental setup with plate apparatus and 15 receiver vertical array at 1m range in the test environment (left), detailed view of clamped plate system (top right), and 7 lbf shaker providing base excitation of the plate (bottom right).

The first type of artificial damage was a boundary-type defect. This was achieved by disengaging one or more clamps at the plate edge to mimic a fastener or weld failure. As shown in Figure 3.3, three test cases were investigated: unclamping an entire side, unclamping a single clamp near the middle of a plate edge, and unclamping a single clamp near the plate corner. Acoustic measurements of the fully-clamped plate served as the baseline data against which the test cases were compared to detect the presence of the mechanical changes.

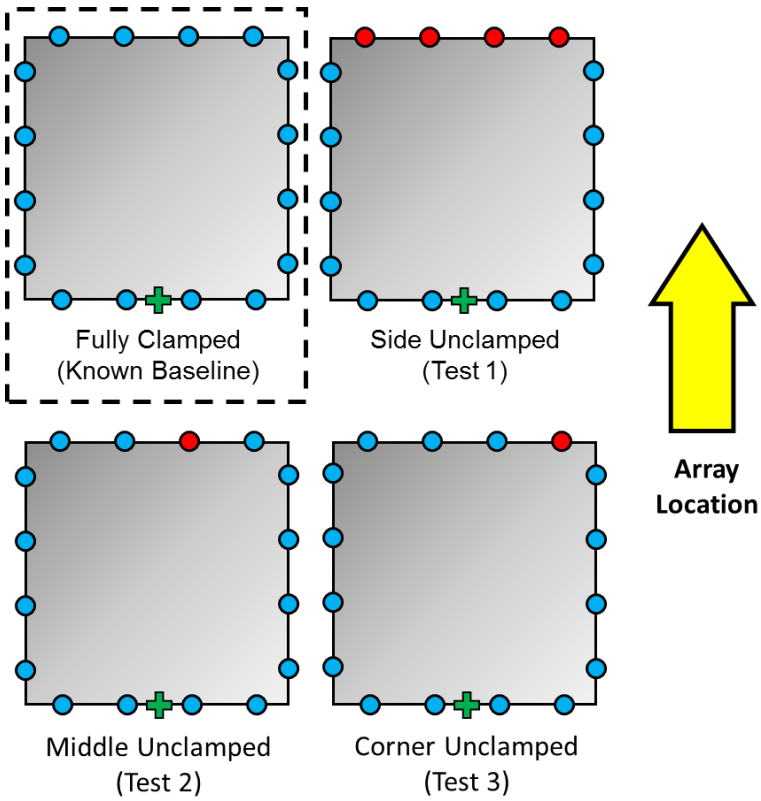

Figure 3.3: Clamping configurations of the 30 x 30 x 0.16 cm aluminum plate for the three test cases and for the baseline. **Blue** circles represent fully engaged clamps and **red** circles represent fully disengaged clamps. The **green** plus sign (+) indicates the location of input forcing into the rigid base to which the test plate was clamped.

The test environment consisted of an acoustically-untreated laboratory with approximately 50-m^2 of floor area, a ceramic tile floor, painted cinder-block walls, a painted

cement ceiling, a variety of furniture and equipment, and numerous utility fixtures which together provided a complex and reverberant environment with a T_{60} reverberation time of about 0.5 sec. For these experiments, the amount of reverberation was quantified by the signal-to-reverberation ratio (SRR) defined by

$$SRR \equiv 10 \log_{10} \left(\frac{SE_{\text{Direct}}}{SE_{\text{Reverb}}} \right) \tag{3.1}$$

where SE_{Direct} and SE_{Reverb} are the signal energies carried by the direct plate-to-array path and all other plate-to-array paths, respectively. The *SRR* is a function of plate position in the room (fixed for all measurements) and the receiving array location. For each array location, the *SRR* was measured by temporarily replacing the plate apparatus with a small loudspeaker that broadcast an 8-second-duration 100-2000 Hz linear chirp that was recorded by a single receiver at the array center. A separate reference measurement of this loudspeaker broadcast signal allowed for matched filtering of the recorded broadcast to reveal the laboratory environment's impulse response, $h(t)$, in the signal bandwidth. This additional filtering step, detailed in Figure 3.4, was required because the broadcast signal was far longer than the reverberation time of the test environment, so temporal truncation and/or windowing could not be used to isolate the direct path signal. However, the direct path arrival was always clearly evident in $h(t)$; thus, SE_{Direct} and SE_{Reverb} were calculated from the time integrals of $h^2(t)$ for the first arrival and for all subsequent arrivals, respectively. The measured *SRR* values at plate-to-array distances of 1.0, 2.0, and 3.0 meters were –7, –11, and –13 dB, respectively.

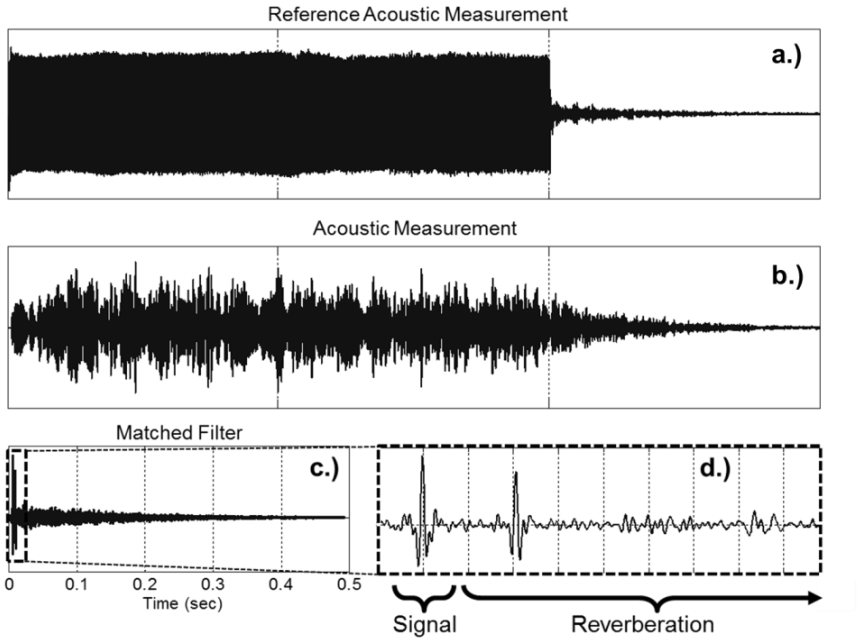

Figure 3.4: SRR measurement procedure. An 8-second-duration source waveform a) was broadcast into the reverberant laboratory and recorded using a single receiver at the array center b) from which a matched filter output was computed c). The SRR is defined as the ratio of the energy of the signal (direct path) to all later arrivals d).

3.1.2 Array signal processing schemes

Though the vibrating plate generally does not radiate sound omnidirectionally, for simplicity it was treated as a remote point source with a unique time domain signature $s(t)$. Three different techniques were used to quantify the vibrating plate's time-domain acoustic signature, $s(t)$, from each experimental trial. The first and simplest technique was to use the recorded time history, $r_m(t)$, from a single receiver (m = 11) without additional processing, i.e. $s_1(t) = r_{11}(t)$. The second technique was conventional time-domain beamforming using all 15 $r_m(t)$ to produce a single beamformed time-domain signature, $s_{BF}(t)$. For this effort, the steering direction was determined from the recorded data and typically pointed from the array center to the plate's center. The third technique was STR applied to the array recordings to

recover a reverberation-suppressed time-domain signature, $s_{STR}(t)$. Here, the STR reference direction was the same as the beamforming steering direction. Results from these three techniques are provided and compared in the next section to show how damage detection robustness may be improved with better suppression of the effects of reverberation.

Conventional time-domain beamforming is a simple technique for fully or partially isolating the sound that arrives at a receiving array from a particular look direction or location [44]. It was implemented using all M = 15 $r_m(t)$ to produce $s_{BF}(t)$, the Fourier transform pair of $\hat{s}_{BF}(\omega)$, defined in Equation (3.2). The weighting vectors $w_m(\omega)$ are computed using the typical spherical wave approach of searching through a range of possible (x, y) source locations, applying the predicted location-dependent phase shifts to each frequency-domain receiver recording $\hat{r}_m(\omega)$, then calculating the magnitude of the sum of the phased responses. $w_m(\omega)$ is then defined as the vector of phases corresponding to the potential source location which yields the largest magnitude output.

$$\hat{s}_{BF}(\omega) = \sum_{m=1}^{M} w_m(\omega)\hat{r}_m(\omega) \tag{3.2}$$

Synthetic time reversal (STR) is a blind deconvolution technique developed for estimating source-broadcast waveforms from array recordings made after the source signal has been corrupted by propagation through an unknown reverberant environment. STR was first proposed for multi-mode underwater sound channels [69], but it has also been successfully applied in multiple-ray-path environments [70]. A more detailed description of STR is given in Section 2.1.2 Synthetic Time Reversal. STR is used here as a more sophisticated means of

reverberation suppression than conventional beamforming. Effectively, STR uses only the recorded array waveforms and beamformer weights $W_j(\omega)$ to estimate impulse responses for each of the receivers. These impulse responses can then be used to synthetically time-reverse or back-propagate the array recordings to their source location, yielding an estimate, $s_{STR}(t)$, of the source's original signal. Equation (3.3) provides the final formula for computing the Fourier transform of $s_{STR}(t)$, $\hat{s}_{STR}(\omega)$, using the same receiver recordings and weighting vectors applied in Equation (3.2).

$$\hat{s}_{BF}(\omega) = \sum_{m=1}^{M} \left\{ \frac{\hat{r}_m(\omega) e^{-i \arg\left(\sum_{l=1}^{M} w_l(\omega) \hat{r}_l(\omega)\right)}}{\sqrt{\sum_{k=1}^{M} |\hat{r}_k(\omega)|^2}} \right\}^* \hat{r}_m(\omega) \tag{3.3}$$

Sample results for the three signal processing techniques are provided in Figure 3.5 for the acoustic signature of a small loudspeaker broadcasting a 1-second-duration 1.0-to-4.0 kHz chirp that was recorded with the 15-microphone array in the reverberant laboratory environment. The panels on the left side of Figure 3.5 show time-domain waveforms while the panels on the right side of the figure show the Fourier transform amplitudes of the various waveforms. Figure 3.5 a) and b) show the reference waveform measured by a single microphone when the loudspeaker-to-microphone distance is 20 cm and the effects of reverberation are nearly negligible. Figure 3.5 c) and d) show the single-microphone acoustic signature, $s_{single}(t)$, when the loudspeaker-to-microphone distance is 3.0 m. Here, the effects of reverberation are more prevalent, and the cross correlation of this signal with the reference signal is 67.9%. Figure 3.5 e) and f) show the conventionally-beamformed acoustic signature, $s_{BF}(t)$, when the loudspeaker-to-array-center distance is 3.0 m. Here, the effects of reverberation are partially suppressed compared to the

single microphone result, and the cross correlation of this signal with reference signal is higher: 91.9%. Figure 3.5g) and h) show the STR signature, $s_{STR}(t)$, when the loudspeaker-to-array-center distance is again 3.0 m. Here, the effects of reverberation are better suppressed than when conventional beamforming is used, and the cross correlation of this signal with the reference signal is even higher: 94.3%. These results show that array-based measurements processed with the blind deconvolution algorithm best suppress the detrimental effects of reverberation. Furthermore, these results for a three-dimensional reverberant environment are in agreement with previous findings for STR in a two-dimensional water channel [70].

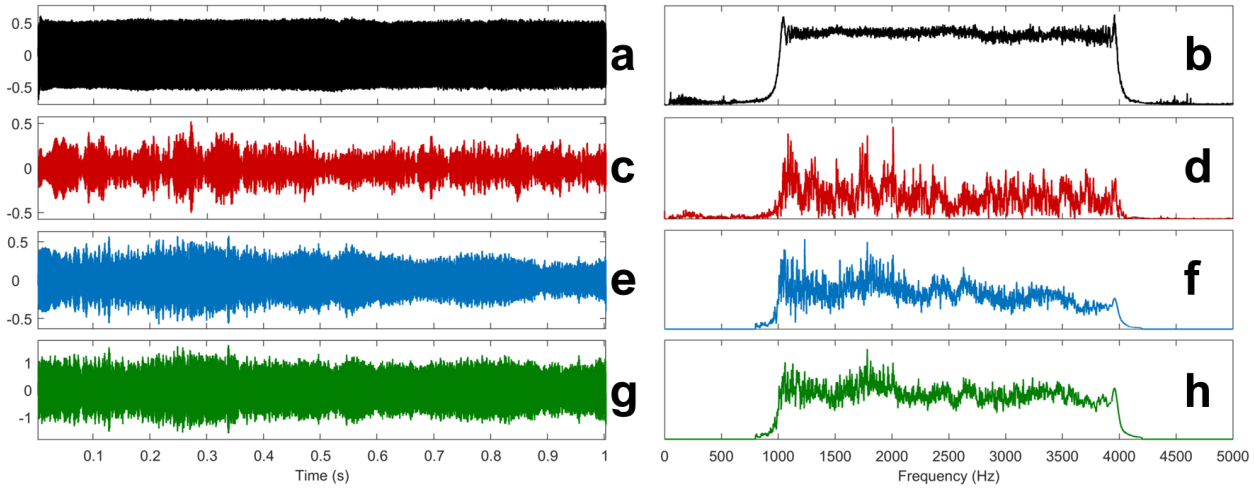

Figure 3.5: Comparisons of signal estimation methods for a source broadcasting a 1-4kHz linear chirp in a reverberant room. a,b) Reference measurement collected at 8 inches from source with corresponding FT magnitude, c,d) single unprocessed receiver with 67.9% correlation to reference with FT magnitude, e,f) conventional spherical wave beamforming (SWBF) with 91.9% correlation with FT magnitude, and g,h) STR reconstruction showing 94.3% correlation with the reference measurement and its FT magnitude.

3.1.2 Results and analysis

Detection of mechanical changes in the test plate was accomplished using the statistics of time domain cross correlations between remote array-measured acoustic signatures of the

vibrating plate for a given test case and known baseline measurements. Figure 3.6 shows a sample set of array recordings, $r_m(t)$ for $1 \leq j \leq 15$, for one baseline experiment. Here, the first and fifteenth microphones are 86 cm and 163 cm from the laboratory floor, respectively. The temporal variations in signal amplitude are primarily caused by the driving frequency moving through radiating vibration modes of the plate. The differences in signal envelope between microphones are caused by non-uniform radiation from the various plate modes and by spatial variations in the reverberant field. The experiment was repeated eight times (a *batch*) for one fully-clamped-plate baseline test, one fully-clamped *defect-free* test (essentially a null test that is equivalent to the baseline), and each of three partially-unclamped-plate test cases for a total of five batches (40 experiments). These five batches of measurements were collected at the plate-center-to-array-center ranges of 1.0, 2.0 and 3.0 m. Thus, the complete experimental data set included results from 120 experiments. Defect detection results for single-receiver, conventionally beamformed, and STR acoustic signatures are provided here to show how the detrimental effects of reverberation can be mitigated.

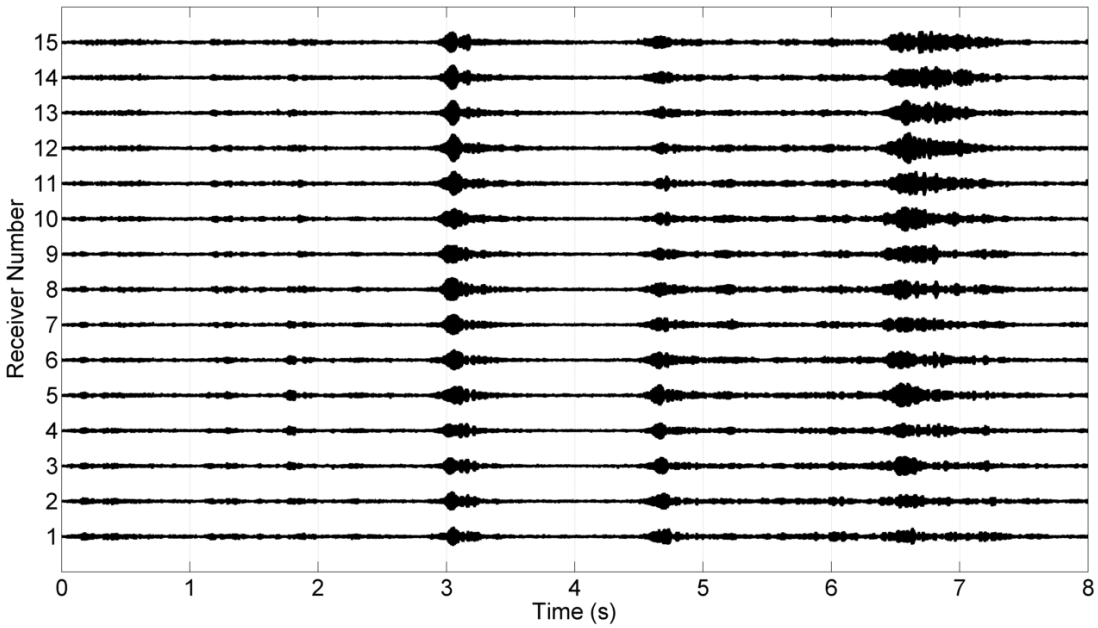

Figure 3.6: Example waterfall plot for the baseline measurement of plate radiation at 1 meter range due to an 8 second 100-2000 Hz base excitation.

The binary detection test considered here was accomplished statistically using the procedure shown in Figure 3.7 and described here. This method effectively captures and quantifies changes in the acoustic signature caused by frequency shifting of the plate's modes, changes in the plate's radiation pattern, and sound generated locally from the plate's defect. For each signal-reconstruction technique (single receiver, conventional beamforming, and STR), every pair-wise combination of reconstructed signals was cross-correlated between *one* of the baseline batches and the three test case batches plus the *other* baseline batch. This ultimately yielded four sets of 64 cross correlation coefficients (one set for each baseline-to-test case, and one for the baseline-to-baseline case). These 64 correlation coefficient samples were then binned into histograms and used to statistically quantify detection performance. To reduce the effect of statistical fluctuations, Gaussians (mean μ, standard deviation σ) were fit to each of the four histograms.

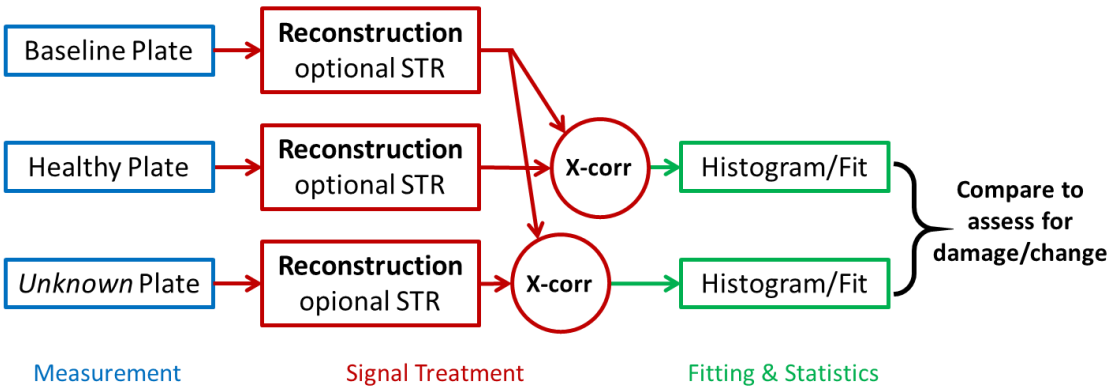

Figure 3.7: Procedural flow chart for the statistical evaluation of structural health using time domain cross-correlation of acoustic signatures determined from baseline and test measurements. This can be extended to an arbitrary number of test cases. The three test cases shown in Figure 3.3 were considered in this study.

A set of sample histograms of correlation coefficients for a single receiver (no. 11) at 1 meter range is shown in Figure 3.8. The histograms of cross correlations for an entire side unclamped and for one near-middle clamp undone are centered at 0.54 and 0.75, respectively, and are fully distinguished from the other test and baseline cases. Thus, the presence of such mechanical changes is readily detected with a single receiver. However, the histograms of cross correlations for one near-corner clamp undone and for the second baseline batch are overlapping. Therefore, these two conditions cannot be fully distinguished using a single receiver.

A common technique for quantitatively evaluating such a binary detector is to plot the calculated probability of true detection P_D versus the probability of false alarm P_{FA} for a given set of measurements with *a priori* known baseline and non-baseline groups. Both P_D and P_{FA} are plotted as a function of the user defined *detection metric*, in this case the cross-correlation with the first baseline, such that measured cross correlation values above and below the detection metric are labeled *detection* and *not-detection*, respectively. The locus of such points

is called a receiver operating characteristic (ROC) curve and, provided a sufficient number of samples are collected, the ROC curve allows the user to select an appropriate detection metric threshold, typically based on some predefined cost function of type I (false positive) and type II (false negative) errors [80]. Additionally, assuming the detection metric is normally distributed and sampled sufficiently to estimate (fit) the parent distribution of both the baseline and non-baseline case, P_D and P_{FA} for the corresponding "fitted" ROC curve can be shown to be

$$P_D = \frac{1}{2}\text{erfc}\left(\frac{1}{\sqrt{2}\sigma_{test}}\left(T - |\mu_{base} - \mu_{test}|\right)\right) \tag{3.4}$$

$$P_{FA} = \frac{1}{2}\text{erfc}\left(\frac{T}{\sqrt{2}\sigma_{base}}\right) \tag{3.5}$$

where $\text{erfc}(x)$ is the complementary error function, μ_{base} and μ_{test} are the baseline and test-case means, σ_{base} and σ_{test} are the baseline and test-case standard deviations, and T is the variable detection threshold which parameterizes the curve

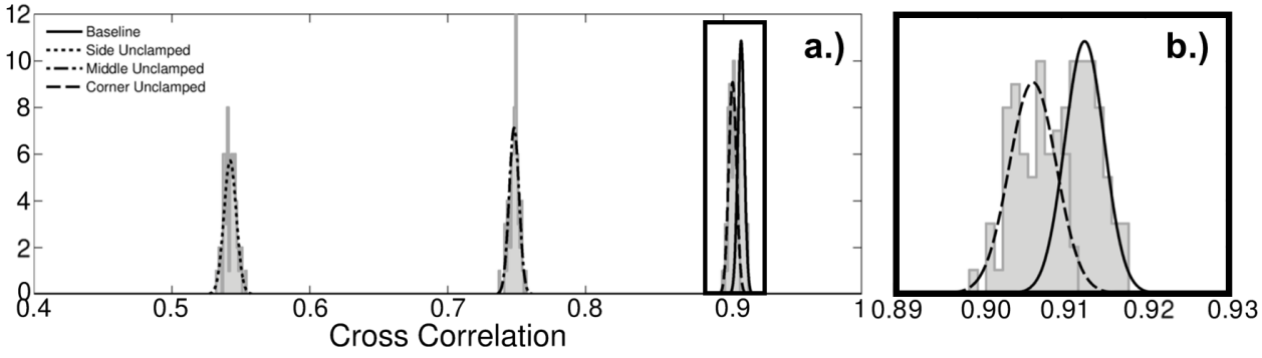

Figure 3.8: a) Binned histograms of cross-correlation coefficients for the three test cases and for the baseline case using a *single unprocessed receiver* (#11). Solid and dashed curves are the Gaussian fits for each test case. b) A detailed view shows overlapping PDFs between the baseline and disengaged corner clamp cases resulting in limited detectability using no processing.

It is evident from Figure 3.8 that the full side unclamped and the disengaged near-center clamp test cases are easily distinguished from the baseline. However, the single disengaged near-corner clamp case is significantly more difficult to distinguish from the baseline. It follows that the probability of making a type I (false alarm) error is far greater, given a desired minimum rate of type II (missed detection) errors. This result is exemplified in Figure 3.9 which shows the corresponding fitted ROC curves for the test cases binned in Figure 3.8.

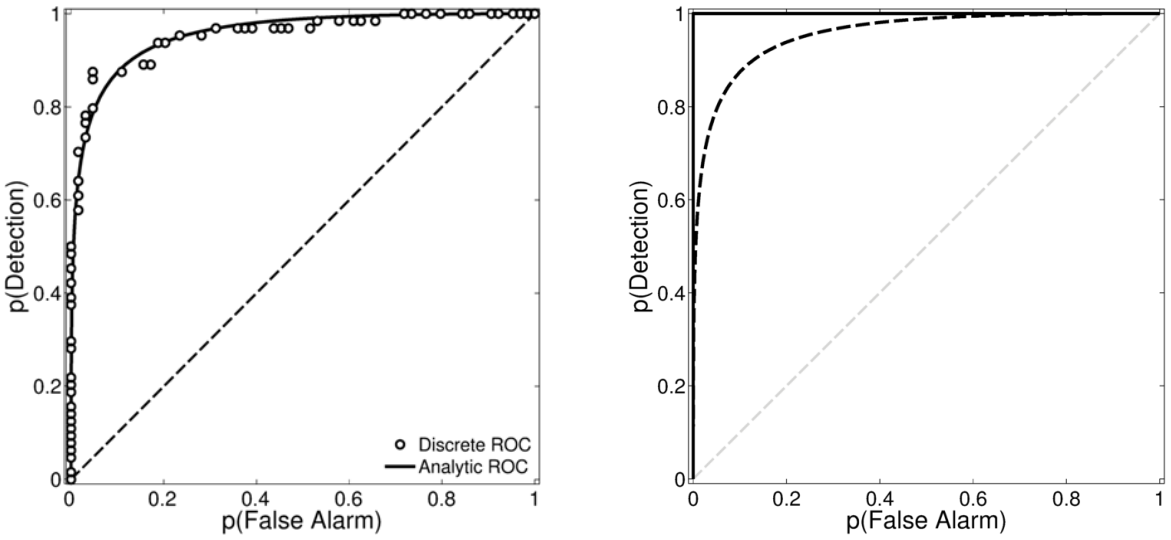

Figure 3.9: *Left*: Discrete, calculated points for the side-unclampled ROC curve. The solid line is the associated analytical fit. *Right*: Corresponding ROC curves for the three test cases using an unprocessed single receiver (#11) at a range of 1 meter. The black dashed line corresponds to the corner-unclamped case, whereas both the center-unclamped and side-unclamped cases overlap top right corner, indicating good detection.

Using the same array recordings at 1 meter range, the detection process was repeated, but this time using STR to construct the acoustic signature of the plate. Conventional spherical wave beamforming (SWBF) was also used for comparison. Figure 3.10 shows that, with STR, there is substantially more separation between the correlation coefficients even though they are

lower. This is reflected in the ROC curves which are located optimally in the upper left corner of the detection domain, meaning that very high rates of accurate detection are possible for negligible rates of false alarm. Conventional SWBF, on the other hand, does not perform as well as STR, and actually performs worse than simply using an unprocessed receiver.

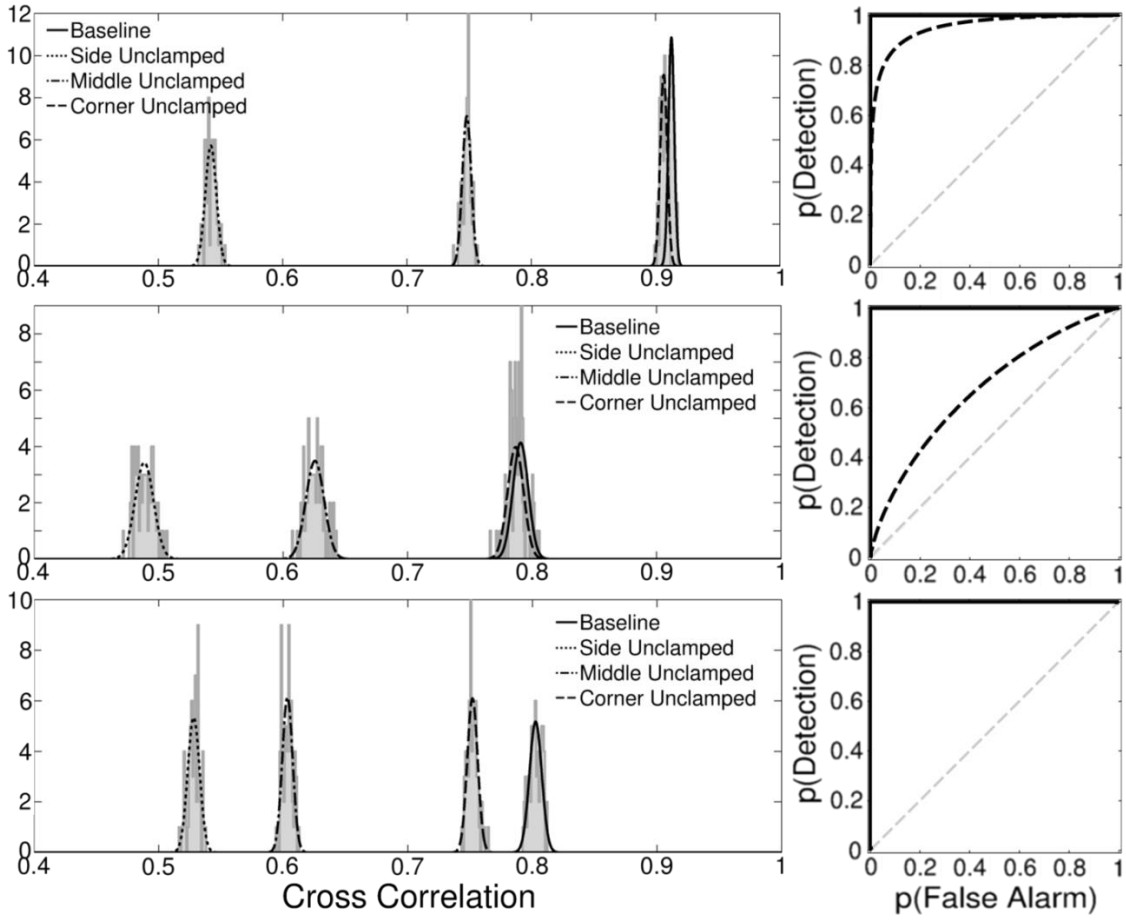

Figure 3.10: Cross-correlation histograms and corresponding ROC curves for an acoustic array at 1 meter range using no reconstruction (top), conventional SWBF (center), and STR (bottom).

ROC curves can be further compacted into a single scalar value by computing the deficit area (DA) between the actual ROC curve and the ideal ROC curve (a vertical line at $P_{FA} = 0$ and a horizontal line at $P_D = 1$). Assuming the associated costs of type I and type II errors are

the same between tests, the DA provides a convenient, compact comparison metric between classifiers with low value of DA being best. The largest possible DA is 1/2 and this corresponds the dashed diagonal line (sometimes referred to as the coin-flipping line) shown in the right-side panels of Figure 3.10. DA values from the current experiments are tabulated in Table 3.1 for all three clamping configurations, at the three plate-array ranges, using a single receiver, conventional SWBF, and STR. The full side unclamped test case was essentially perfectly detectible by all techniques regardless of range or reverberation correction. The single disengaged clamp near the plate-side's center was also nicely detectable, but at larger ranges it was found to be slightly less so. The near-corner unclamped case was not well detected using either SWBF or a single unprocessed receiver, particularly at ranges of more than 1 meter. However, using STR, even the unclamped corner case was detected almost perfectly for all three plate-array ranges.

Table 3.1: Deficit area (DA) of the ROC curve for each test case at 1 m, 2 m, and 3 m array ranges. Results are compared using STR, SWBF, and no signal processing (i.e. using a single unprocessed receiver). Cases with the worst detection performance (DA > 0.1) are boxed in red.

		Unprocessed Receiver			**Conventional Beamforming**			**Synthetic Time Reversal**		
Range	SRR	Full Side	Center	Corner	Full Side	Center	Corner	Full Side	Center	Corner
1 m	-7 dB	0.00	0.00	0.10	0.00	0.00	0.33	0.00	0.00	2.5e-12
2 m	-11 dB	0.00	3.1e-15	0.19	0.00	5.9e-3	0.41	0.00	2.2e-16	2.2e-16
3 m	-13 dB	0.00	2.4e-16	0.45	0.00	3.7e-3	0.38	0.00	0.00	6.6e-12

Noting that it is often difficult to place a test structure in exactly the same location or orientation that it had been in for a baseline recording, an additional experiment was conducted by recording the test measurements with the array at 3 meters from the plate (SRR = –13 dB),

with *only the first* batch of baseline measurements recorded at the same range but with a 0.2 m shift perpendicular to the original plate-array direction. In this scenario, geometric discrepancies can cause significant differences in the array recordings, even between the two baseline batches, due to the reverberant characteristics of the laboratory. With this change in baseline location, detection performance was significantly reduced using both the unprocessed single receiver and conventional SWBF methods as is shown in Figure 3.11. Here, neither the disengaged near-corner clamp case nor the disengaged near-center clamp case were well detected, and even in the case of the fully unclamped side, the unprocessed single receiver was beginning to struggle. On the other hand, with the additional step of STR reconstruction, all test cases were well detected with DA < 0.03. This result agrees with the expectation that STR yields an estimate of the broadcast signal that is more independent of the environment's reverberant propagation, so its results are thereby less sensitive to changes in the receiving array location. Without such a reconstruction step, differences between test and baseline measurements are indeed found, however these differences are not the result of changes in the test structure but are instead primarily dependent on changes in the acoustic connections between source and receivers in the reverberant environment which leads to a higher rate of false alarm.

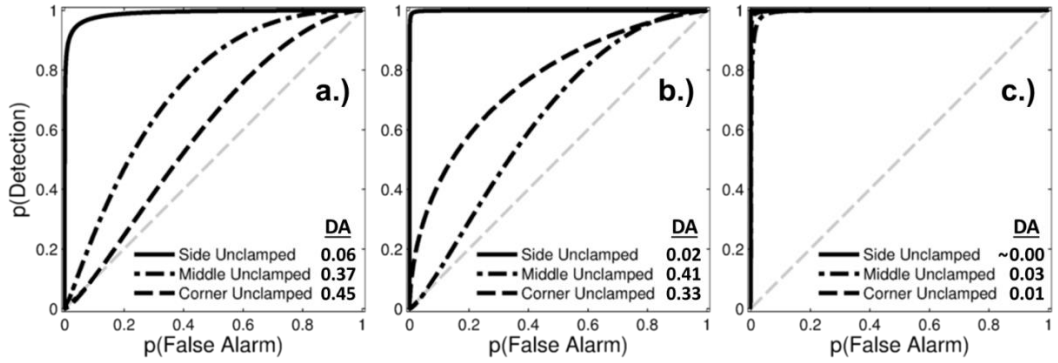

Figure 3.11. ROC curves for test measurements collected at a range of 3 meters using a baseline collected at 0.2 meters away from the test measurements, a) using no signal reconstruction (receiver #11), b) using SWBF, and c) using STR.

3.2 Detection with stochastic input forcing

When a mechanical structure is damaged or otherwise mechanically changed, its vibrational characteristics are altered [1, 30]. If the structure resides in a fluid medium, such alterations in vibratory behavior result in changes to the sound radiated by the structure [10]. Thus, even at remote distances, recordings of this radiated sound can be used to detect and assess mechanical changes of the structure. Moreover, if *baseline* measurements of a known *healthy* structure can be collected, subsequent *test* measurements may be compared to these baseline recordings for the purposes of structural health monitoring.

Such remote acoustic detection provides certain benefits over contacting or near-contacting methods of damage detection such as vibration monitoring [5, 6, 7], guided wave testing [14, 15], acoustic emissions [16, 17], and nearfield acoustic holography [11, 12]. The drawbacks associated with requiring transducers to be near or in contact with a test structure can include: transducer installation difficulties, infeasibility of *in-situ* measurements, monitoring and replacement of aging and failing transducers, and unexpected effects of physical

coupling between transducers and test structures. However, fully remote acoustic monitoring circumvents many of these issues while also allowing a single (more-easily maintained) receiver array to be used for monitoring multiple independent systems.

Sound radiated from a vibrating structure that is recorded remotely depends on the vibratory forcing applied to the structure, the structure's frequency response function, and the characteristics of the acoustic propagation between the structure and the receiver. A primary impediment to remote acoustic sensing in real acoustic environments is the required discrimination between changes in the structure's frequency response and those resulting from changes in the other two factors (vibratory forcing, and structure-to-receiver propagation). Such a task may be difficult because of unknown vibratory forcing and environmental factors including noise, environmental uncertainty, and multipath propagation [67].

For a linear dynamic system, the recorded response in the frequency domain can be modeled as the product of the input forcing signal(s), the system response, and the source-to-receiver transfer function for the acoustic environment. Thus, the desired detection and identification of changes in the system response may be masked by changes in the input forcing or acoustic transfer function. As a result, some *a priori* knowledge of the input forcing must be available in order to accurately ascribe detected changes to actual structural changes. Fortunately, this requirement is often admissible in many practical applications, as discussed below. In this paper, two input forcing cases are considered, one where the frequency content *and* phase of the input forcing is held constant, and one where only the frequency content of the input forcing is held constant (i.e. bandlimited noise). Effects of the character of the input forcing on detection performance are analyzed for three separate detection strategies.

Attention, then, must also be given to the acoustic transfer function of the environment which is generally unknown. Though there are numerous types of propagation effects, signal corruption from unknown and variable multipath presents a particularly difficult challenge for three reasons: 1) it is not readily suppressed with multiple measurements, 2) the Green's function for a typical three-dimensional testing environment with multiple reflecting surfaces is generally unknown and may vary over sufficiently long time periods [81], and 3) multipath propagation can be very sensitive to even small changes (fractions of an acoustic wavelength) in receiver-to-source geometry [82], which is problematic when comparing test recordings to baseline measurements which may have been collected days, months, or even years earlier. In this paper, the problem of unknown multipath is tackled using the synthetic time reversal (STR) *blind*-deconvolution algorithm [70] which requires no information about the acoustic environment or the source waveform. STR leverages the acoustic field directionality deduced from transducer array recordings to estimate the true broadcast waveform in the absence of multipath interference.

Even with corrections for variations in the vibratory-excitation waveform and for the environment's unknown multipath characteristics, the task of comparing *test* data to *baseline* recordings and detecting changes remains. A recent study yielded successful detection of clamping changes using time-domain cross-correlations between test and baseline signatures for a vibrating plate undergoing repeatable swept-frequency excitations [83]. However, in many, if not most, applications, the input forcing to a system is not perfectly repeatable and does not exhibit the convenient characteristics of a frequency sweep. For large and/or complex systems with many load paths and multiple sources of vibration (for example a ship with a large engine, secondary machinery, and hydrodynamic loading), the total vibratory loading is

complicated and can often result in radiated sound more closely resembling broadband noise than a deterministic signal such as a frequency sweep [84, 85]. As such, it is desirable that an acoustic detection method for structural health monitoring provides accurate detection performance, even if the input forcing (and resulting response) is stochastic.

Herein, cuts of various lengths in a vibrating plate are detected using remote acoustic array-recorded data collected in an unknown reverberant environment. The STR blind deconvolution algorithm is used to estimate the plate's radiated acoustic signature in this reverberant environment. Three detection metrics – time-domain cross-correlations, power spectrum comparisons, and frequency response function comparisons – are then used to quantitatively compare baseline radiated sound recordings from *healthy* vibrating plates to equivalent recordings from potentially changed (or *damaged*) plates. Detection performance is compared between the three metrics for both stochastic and frequency-sweep forcing. Performance robustness to source-array geometry changes in the reverberant environment is also addressed. Thus, the work reported here extends the prior results involving boundary-clamping changes and frequency-sweep forcing [83] to cuts in the vibrating plate, random structural excitation, and additional detection metrics.

3.2.1 Experimental methods

The experimental work was completed in an ordinary laboratory without acoustic wall or ceiling treatments. Figure 3.12 shows the apparatus. The acoustically radiating test structure consists of a 30 cm x 30 cm x 1.6 mm aluminum plate clamped along all sides to a nominally-rigid aluminum mounting base (top right of Figure 3.12). For deterministic excitation, eight-second-duration linear frequency sweeps from 100 to 2000 Hz were created on a computer and

sent through a PXIe-6368 DAQ (National Instruments, Austin, TX) to a K2007E01 series electrodynamic shaker (The Modal Shop, Cincinnati, OH) providing input forcing to the plate's mounting base (bottom right of Figure 3.12). For stochastic excitation, eight seconds of 100-to-2000-Hz band-limited noise were input via the shaker to the same drive point. The actual input forcing, $f(t)$, delivered to the mounting base was recorded at the drive point using a 208C03 inline force transducer (PCB, Depew, NY) (see Fig. 1.) and was used for one of the detection metrics. Sound from the resulting plate vibration radiated into the laboratory environment and was recorded using a 15-element vertical array with d = 5.1 cm spacing between the 130E20 microphones (PCB) for a total aperture of L = 71.1 cm ($kd \approx$ 1.0, $kL \approx$13 for a nominal center frequency of 1000 Hz and a sound speed of 344 m/s). The distance from the plate's center to the array's center was 132 cm, as shown at the left of Figure 3.12. The array element recordings, $r_j(t)$, were analog band-passed filtered between 50 Hz and 5 kHz, digitized at a 50 kHz sample rate per channel with the same DAQ system, and stored on the laboratory computer for subsequent processing. The input forcing bandwidth was chosen to encompass approximately the first dozen vibrational modes of the test plate.

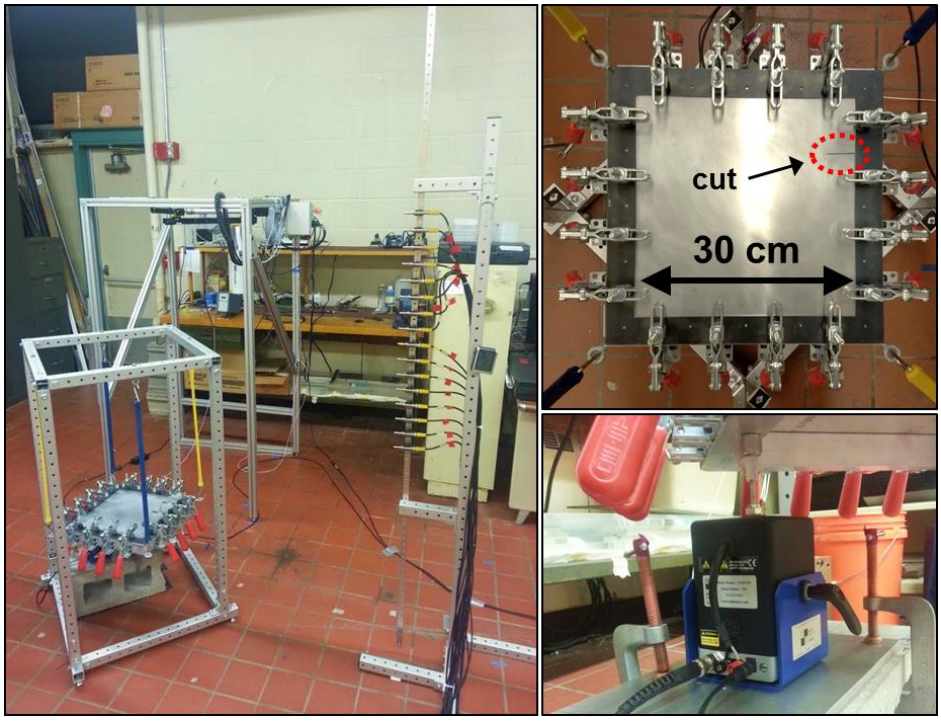

Figure 3.12: Experimental setup with plate apparatus and 15 receiver vertical line array (5.1 cm spacing) in the test environment (left), top down view of a 30-cm square aluminum plate with a 38 mm cut (top right), and 30 N shaker providing base excitation to the plate (bottom right). [Color on-line].

The test plate was mechanically changed by introducing cuts perpendicular to an edge of the plate (see upper right panel of Fig. 1). Seven test cases were considered: no cut (baseline/healthy), and six cut lengths ranging from 13 mm to 76 mm in 13 mm increments (±2 mm cut-length uncertainty). Figure 3.13 shows an example of the effects of a 38 mm cut on the vibrational response of the plate for the (1,4) plate vibration mode, measured using an LP-01 laser Doppler vibrometer (LDV) (OMS Corporation, Laguna Hills, CA). The (1,4) mode exhibits changes in the vibratory mode shape localized near the cut in addition to the global effect of a reduced resonant frequency, both of which aid in acoustic detection. Such vibrational changes are expected for cuts on a resonating structure [5].

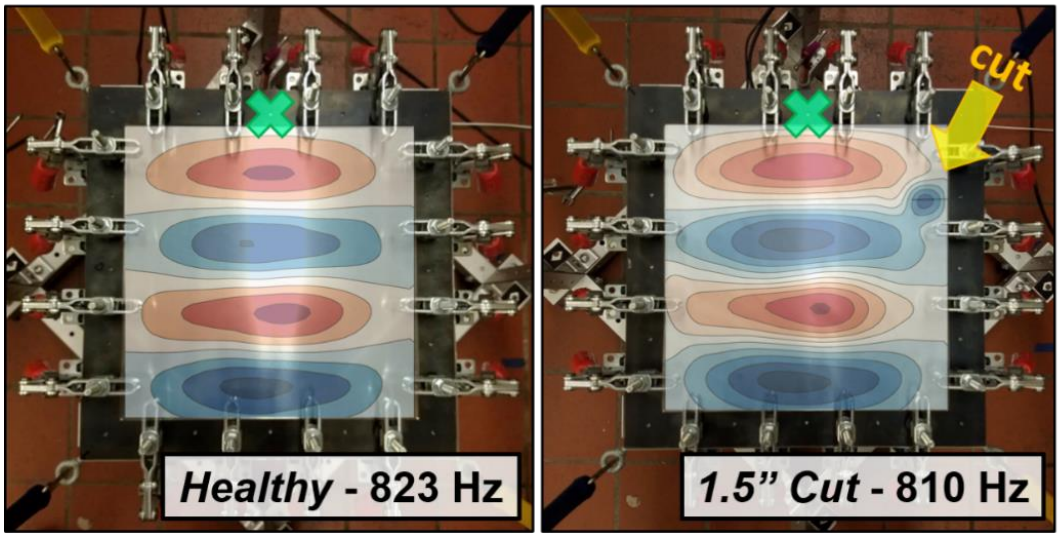

Figure 3.13: Effects of 38 mm cut on the (4,1) plate mode. The plate is base excited underneath at the green cross. The 38 mm cut exhibits localized changes in the mode shape in addition to shifting the resonant frequency. Similar results were found at other modes.

Acoustic measurements were made in a roughly 50 sq-m laboratory which provided a reverberant environment with a T_{60} time of approximately 0.5 sec. The extent of reverberation is quantified by the signal-to-reverberation ratio (SRR) defined by

$$SRR \equiv 10 \log_{10} \left(\frac{SE_{\text{Direct}}}{SE_{\text{Reverb}}} \right) \tag{3.6}$$

where SE_{Direct} and SE_{Reverb} are the signal energies that travelled between the plate and the array on the direct path and all-other paths, respectively. SRR depends on the relative source-receiver-room geometry and was estimated to be –7 dB for the geometry described above (i.e. multipath signal energy was five times greater than direct path signal energy). The measurement of SRR was accomplished by substituting a small loudspeaker for the plate apparatus and

broadcasting a known 0.1-to-2 kHz bandwidth signal which was then matched filtered to identify direct and subsequent arrival paths, times, and amplitudes [83].

To correct the recorded signals for multipath interference the synthetic time reversal (STR) blind-deconvolution algorithm [69] was applied to the recorded array data. Using only the array recordings and known receiver locations, STR provides estimates of the source-to-receiver impulse responses and the original signal waveform broadcast from the source. In particular, the Fourier transform of the STR-estimated signal waveform is [70]

$$\hat{s}_{est}(\omega) = \sum_{m=1}^{M} \left\{ \frac{\hat{r}_m(\omega) e^{-i \arg\left(\sum_{l=1}^{M} w_l(\omega) \hat{r}_l(\omega)\right)}}{\sqrt{\sum_{k=1}^{M} |\hat{r}_k(\omega)|^2}} \right\}^* \hat{r}_m(\omega) \tag{3.7}$$

where M is the number of receivers, $\hat{r}_m(\omega)$ is Fourier transform of the m^{th} microphone's recording $r_m(t)$, and $w_l(\omega)$ is a weighting vector, identical to that which would steer a conventional beamformer in the direction of maximum gain, typically along the direct path between the source and the array. For its benefit, STR requires no more information (and negligibly more computation) than that of classical delay-and-sum beamforming. For the purposes of this study, STR collapses the 15 microphone recordings to a single signal from which the multipath propagation characteristics of the source-to-array geometry and the laboratory environment have been at least partially removed, as recommended in a prior remote structural health monitoring study [83].

3.2.2 Detection metrics & procedure

Using STR-reconstructed signals for baseline and test cases, three metrics were considered for detecting structural changes. A sample of each detection metric was generated

from one test measurement and one baseline measurement, resulting in a scalar value between 0 and 1. To statistically evaluate detection performance batches of sixteen acoustic array measurements were collected for: the defect-free baseline case, a second defect-free *test* case (zero length cut), and six additional *test* cases, one for each non-trivial cut length, for a total of eight batches (128 individual experiments). For each of the seven test batches, every pair-wise combination of STR-reconstructed signals was compared to those from the *baseline* batch using the three metrics outlined below. Here, results for the baseline and zero-length-cut batches serve to quantify the outcome for the null hypothesis for damage detection and are necessary for statistical evaluation of detection performance for the six non-zero-length-cut cases. This procedure yielded seven sets of 16 x 16 = 256 metric samples for each of the three detection metrics. These samples were binned into histograms, fitted to Gaussian distributions, and used to statistically quantify detection performance. This whole procedure was completed once for stochastic noise input forcing and once for frequency-sweep input forcing. In total, 15 x 16 x 8 x 2 = 3,840 total receiver measurements were collected for this effort. A flowchart of the detection procedure is provided in Figure 3.14.

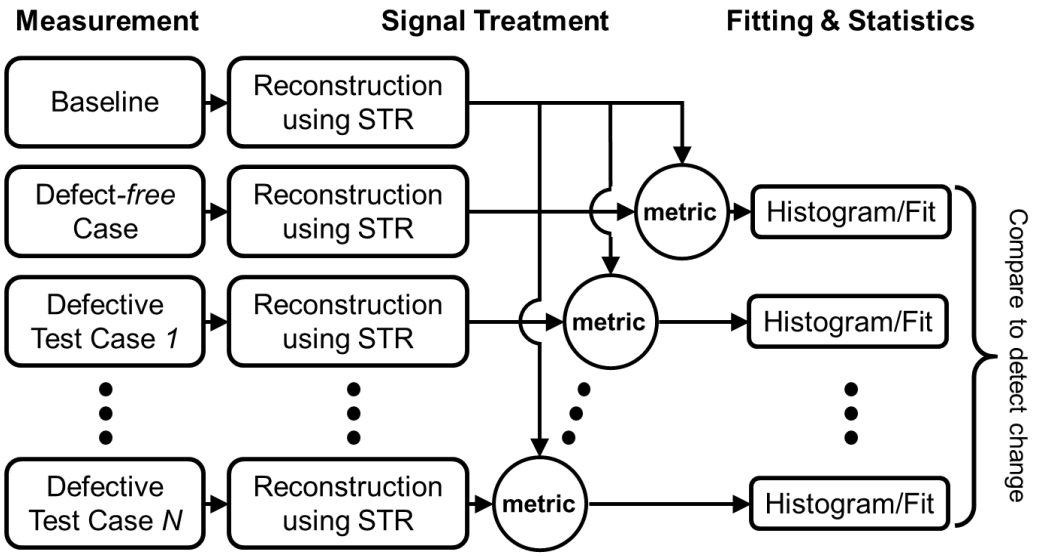

Figure 3.14: Procedural flow chart for the statistical evaluation of structural health using a comparative metric between remote acoustic test and baseline measurements. This can be extended to an arbitrary number of test cases, though seven were considered in this study.

Three detection metrics, denoted by Δ, were considered in this study, each possessing its own strengths and weaknesses. The first detection metric is the time domain cross correlation coefficient between baseline, $b(t)$, and test, $s(t)$, measurements, and is defined by

$$\Delta_{TD} = \max_{\tau} \frac{\int_{-\infty}^{\infty} b(t)s(t+\tau)dt}{\sqrt{\int_{-\infty}^{\infty} b(t)^2 dt}\sqrt{\int_{-\infty}^{\infty} s(t)^2 dt}}. \tag{3.8}$$

Here, $b(t)$ and $s(t)$ were typically obtained from an inverse Fourier transform of Equation (3.7), and Δ_{TD} of ±1 specifies perfect correlation and anti-correlation between $b(t)$ and $s(t)$, respectively, while Δ_{TD} near zero indicates little or no correlation between $b(t)$ and $s(t)$. Although this metric is relatively easy to interpret, it provides low (near-zero) values between the baseline and all test cases when the vibratory forcing is unmeasured and random.

The second detection metric is the zero-shifted cross correlation coefficient between the power spectral density (PSD) of the baseline and test measurements and is defined by:

$$\Delta_{PSD} = \frac{\int_{-\infty}^{\infty} \left|\hat{b}(\omega)\right|^2 |\hat{s}(\omega)|^2 d\omega}{\sqrt{\int_{-\infty}^{\infty} \left|\hat{b}(\omega)\right|^4 d\omega} \sqrt{\int_{-\infty}^{\infty} |\hat{s}(\omega)|^4 d\omega}}. \tag{3.9}$$

where $\hat{b}(\omega)$ and $\hat{s}(\omega)$ are the Fourier transforms of $b(t)$ and $s(t)$, respectively, and were typically obtained directly from (3.7). This metric is bounded between zero and unity, with Δ_{PSD}-values near unity indicating a close match between the PSD of the baseline and test measurements. Given that changes in stiffness or mass distribution due to damage in a structure are likely to result in changes in frequency response [4], and that such changes are often captured in the power spectrum of a signal, then a decline in Δ_{PSD} away from unity should indicate the existence of structural change(s). In addition, this metric is less dependent on the waveform of the vibratory forcing and achieves near-unity values for all baseline and zero-length-cut test comparisons in this study. However, all integrands in (3.9) are positive definite, so Δ_{PSD}-values are always positive, even when $b(t)$ and $s(t)$ are uncorrelated ($\Delta_{TD} = 0$). Thus, the overall dynamic range of Δ_{PSD} may be small. And, Δ_{PSD} is only a reliable metric when the frequency dependence of the vibratory forcing is consistent between test and baseline measurements. To achieve sufficient similarity of the PSDs, each measured spectrum is smoothed using a boxcar average with a 5 Hz window before the Δ_{PSD} metrics are computed. The selection of a 5 Hz smoothing window for the 8.5 second measurements is further discussed in Section 3.2.3b Conditions for Input Forcing Repeatability.

The third detection metric is the zero-shifted cross correlation coefficient between acoustic frequency response functions (FRFs) from test and baseline measurements, and is defined by

$$\Delta_{FRF} = \frac{\int_{-\infty}^{\infty} FRF_b(\omega) \cdot FRF_s^*(\omega) d\omega}{\sqrt{\int_{-\infty}^{\infty} FRF_b^2(\omega) d\omega} \sqrt{\int_{-\infty}^{\infty} FRF_s^2(\omega) d\omega}} \tag{3.10}$$

where

$$FRF_b(\omega) \equiv \frac{\hat{b}(\omega)}{\hat{f}(\omega)} \quad \text{and} \quad FRF_s(\omega) \equiv \frac{\hat{s}(\omega)}{\hat{f}(\omega)} \tag{3.11a,b}$$

are the FRFs computed for the baseline and test measurements, respectively, and $\hat{f}(\omega)$ is the Fourier transform of the vibratory forcing function $f(t)$. FRFs are convenient functions as they contain the same beneficial structural frequency-response information as the power spectra. However, unlike the PSD of $b(t)$ or $s(t)$, they are independent of variations in the input forcing, in the absence of noise when $\hat{f}(\omega)$ is non-zero in the bandwidth of interest. However, use of FRFs and (3.10) requires knowledge of $f(t)$ that may not be available in many structural health monitoring applications. As with the Δ_{PSD} metric, the measured FRFs are smoothed with a 5 Hz width boxcar average before computing Δ_{FRF}.

3.2.3 Results & analysis for detection of cuts

Following the procedure outlined in Figure 3.14, cut detection, and its dependence on cut length, was determined for both stochastic input forcing and deterministic frequency-sweep input forcing. Here the detectability of the 13 to 76 mm cuts was determined by binning the

computed baseline-to-test metrics into histograms, fitting a Gaussian distribution (with mean μ, and standard deviation σ), and comparing this baseline-to-test distribution with the equivalent baseline-to-zero-length-cut distribution. The resulting histograms and Gaussian fits for stochastic input forcing are shown in the three parts of Figure 3.16. for the three metrics Δ_{TD} (a), Δ_{PSD} (b), and Δ_{FRF} (c) for seven different test cases (one for zero cut length and six for non-zero cut lengths). In all three parts of Figure 3.16, the jagged curves are the measured histograms while the smooth bell-shaped curves are fitted Gaussian distributions. Here, horizontal and vertical axes have been adjusted in each panel to best portray measured distributions.

Using the data shown in Figure 3.15, the detection performance of the each metric can be assessed by considering the overlap of the baseline-to-zero-length-cut distribution, which is shown in black and presents the statistics of 'nothing to detect', with the various baseline-to-non-zero-length-cut distributions, which are shown in colors and present the statistics of 'something to detect'. In practice, an operator would use data like these to select a *threshold* for the chosen detection metric; measured metrics *below* that threshold would correspond to the presence of damage, and metrics *above* the threshold would correspond to no damage. For example, using the Δ_{PSD} metric (Figure 3.16b), an operator may select a threshold value for the PSD inner product between an unknown signature and baseline signature to be just above 0.98. Since the histograms are well separated, this threshold would yield a high probability that any measured Δ_{PSD} below 0.98 is likely due to a mechanical change, at least for damage severity equal to or greater than a 13 mm cut.

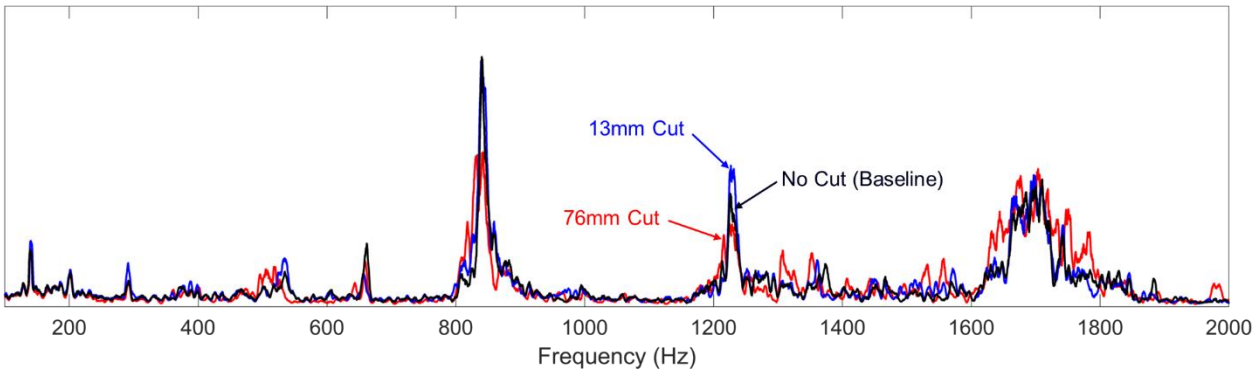

Figure 3.15: Acoustic FRF magnitudes (arbitrary units) for baseline, 13 mm cut, and 76 mm cut cases vs. frequency (Hz). Each FRF was computed using the center-most receiver output (receiver 9) with stochastic input forcing. A 5 Hz boxcar smoothing operation was applied in each case.

With this in mind, some observations on detection performance can be easily inferred from Figure 3.16. Firstly, the Δ_{TD} metric completely fails in the presence of stochastic input forcing. In this case, the average of all Δ_{TD} values is approximately 0.07 and with little or no separation of the various distributions by cut length. This is not a surprising result; cross-correlation between two signals is highly-dependent on the similarity of their phase structure. In the case of noisy input forcing, the phase structure of the forcing is effectively random, resulting in random phases for the acoustic measurements and very low cross-correlations. Detection metrics in the time domain (essentially all of which require some knowledge of the underlying phase structure [7]) are unsuccessful when applied to stochastically forced systems. In comparison, the use of frequency domain metrics Δ_{PSD} and Δ_{FRF} both exhibit very good detection performance (good separation between baseline-to-zero-length-cut distributions and those of baseline-to-non-zero-length-cut distributions). These two metrics, in contrast to Δ_{TD}, rely only on *magnitudes* in the frequency domain thereby making them independent to the random phases associated with stochastic forcing.

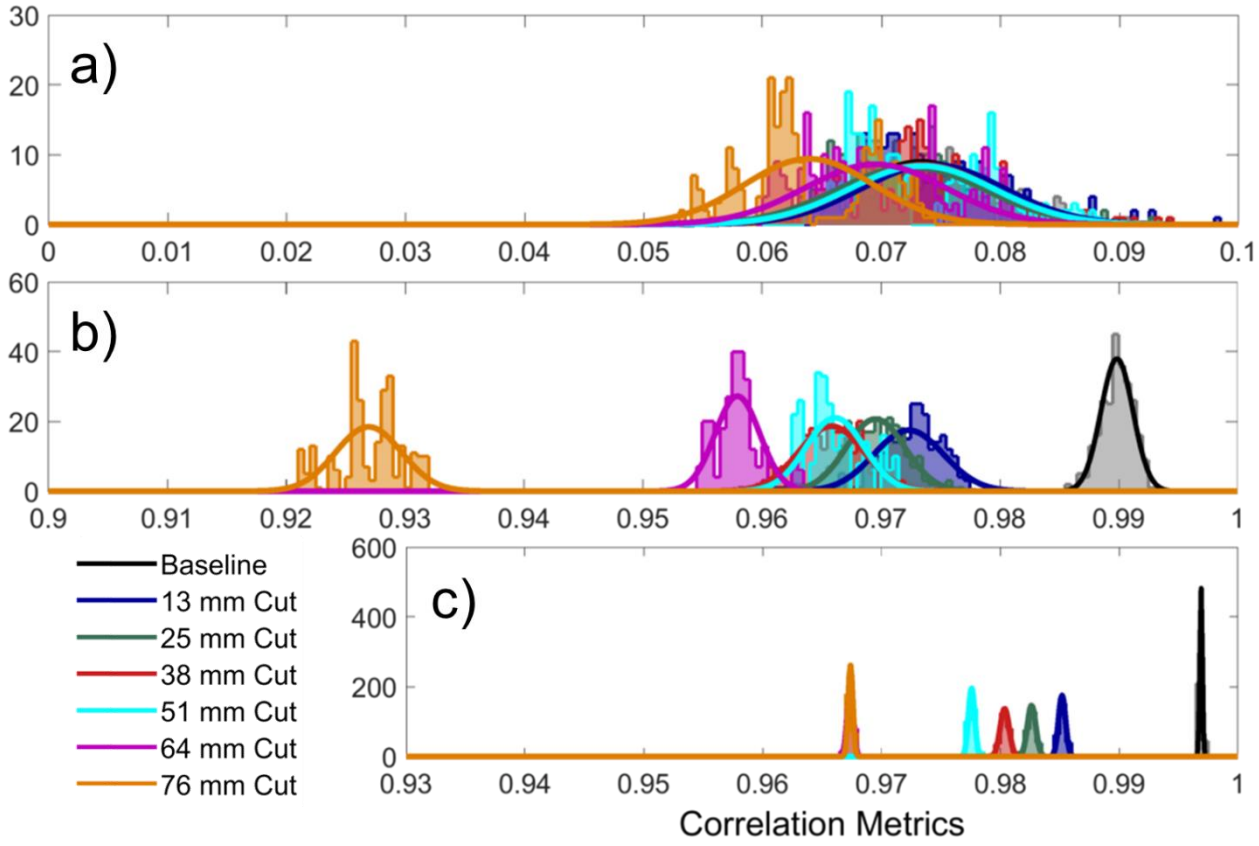

Figure 3.16: Detection histograms and fitted gaussian profiles from noisy input forcing using the (a) Δ_{TD} metric, (b) Δ_{PSD} metric (middle), and (c) Δ_{FRF} metric (bottom). Both x- and y-axis scales are equal in all three plots. Good separation between the gray-black histogram and the other histograms indicates good damage detection performance for the Δ_{PSD} and Δ_{FRF} metrics. Note that the horizontal axis range has been shifted in the top image to show the very low values of Δ_{TD} ~ 0.07.

For comparison, detection performance was also investigated using repeatable frequency sweep input forcing over the same 0.1-to-2 kHz bandwidth. The corresponding detection histograms are similar to those in Figure 3.16 (and therefore not shown) except that the Δ_{TD} metric performed much better due to the repeated phase structure between successive experiments.

The detection histograms shown in Figure 3.16 can be further quantified using a detection index, d, for each cut length, defined as the ratio of the difference in the means between baseline-to-zero-length-cut (hereafter denoted 'baseline') and baseline-to-non-zero-length-cut

(hereafter denoted 'test') histograms, and the geometric-mean standard deviation of the two distributions [80]. Explicitly, d is given by

$$d = \frac{\mu_b - \mu_s}{\sqrt{\sigma_b \sigma_s}} \tag{3.12}$$

where μ_b and μ_s are the means of the fitted-Gaussian distributions and σ_b , σ_s are the respective standard deviations. Intuitively, the detection index is a scalar indicator of how well separated a test histogram is from the baseline histogram. Values of $d \gg 1$ indicate that the histograms are well separated, i.e. detection performance is good. Detection indices are tabulated Table 3.2 for each cut length and each metric investigated. It is clear that while Δ_{TD} is an insufficient metric for stochastic forcing, it performs comparably to Δ_{PSD} in the case of repeatable frequency sweep forcing. Regardless of forcing character, Δ_{FRF} performed the best of all three metrics due to very low standard deviation of in the resulting histograms; though the cost of these gains is the necessity of direct knowledge of the input forcing, which may be unobtainable in many applications.

		Frequency Sweep Forcing			Noise Forcing		
		TD	***PSD***	***FRF***	***TD***	***PSD***	***FRF***
Cut Length (mm)	13	7.8	9.4	39.8	-0.052	7.1	37.0
	25	8.7	10.7	47.3	0.046	8.2	41.3
	38	15.3	13.1	52.7	0.098	12.6	47.6
	51	22.5	14.1	81.1	-0.085	12.1	61.0
	64	81.0	19.5	101.1	0.21	28.0	93.3
	76	80.8	32.4	144.1	0.71	34.5	104.3

Table 3.2: Detection indices for each of the cases considered. Values highlighted in pink are low detection index values, $d < 1$, which indicate poor detection performance.

3.2.3a Robustness to Multipath Variation

All results above were obtained using the STR algorithm to correct for unknown multipath in the reverberant laboratory (SRR = -7 dB). However, unlike the tests above, a structure of interest may not be in the exact same position for baseline and test measurements that are separated by days, weeks, months, or even longer. In reality, uncertainties in the source-receiver geometry will exist and these uncertainties can be detrimental in a reverberant environment [81]. In an effort to quantify the robustness of the proposed remote acoustic detection method to geometric uncertainties, two additional sets of *baseline* array recordings were collected with the array shifted laterally from its nominal position by 51 mm and 102 mm (angular shifts of roughly 3° and 5°, respectively, see Figure 3.17). Calculation of the detection metrics was again performed as described in Figure 3.14, however with the offset baseline recordings in place of the 'ideal' baseline recordings.

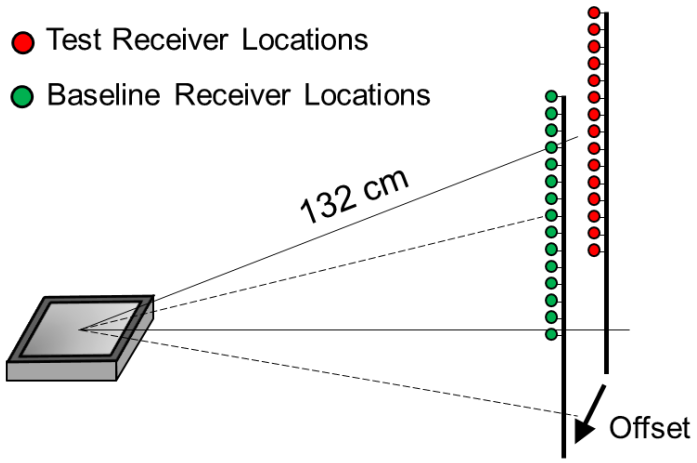

Figure 3.17: Schematic showing the baseline offset used in to investigate detection robustness to multipath.

Figure 3.18 shows detection results using the Δ_{PSD} metric for three levels of baseline offset using random forcing, obtained using no signal reconstruction by directly processing the

recording from the arbitrarily selected 12th receiver, and by using STR. Detection performance is depicted via Receiver Operating Characteristic (ROC) curves which are a common tool used for evaluating binary detection schemes [80]. For each length of cut, the ROC curve is the locus of points specifying the probability of false alarm (undesired, on the horizontal axis) and the probability of true detection (desired, on the vertical axis) as the detection threshold is changed. ROC curves exhibiting good detection performance reach the top left corner (high probability of detection for low probability of false alarm), while those from poor detectors fall closer to the dashed diagonal. Though the unprocessed case performed very well when the source-array geometry was identical between baseline and test measurements (i.e. 0 mm offset), it shows signs of weakness with only 51 mm of offset and fails in detecting all but the most severe level of damage when the baseline is offset by 102 mm. For comparison, corresponding ROC curves were also produced using signals reconstructed with STR to correct for multipath due to the applied geometric errors. For each level of offset, detection with STR was essentially perfect resulting in ideal ROC curves as shown on the bottom left of Figure 3.18. The observed performance improvements indicate both the importance of considering multipath interference and the robustness of the proposed method to changes in it. Since STR applied to the 15-element array resulted in optimal detection, an important question is then: how many receivers are necessary for this performance? To address this question, the entire detection procedure using STR was run again; however, a reduced number of receivers N were used in the computations (receiver spacing was maintained at 51 mm). Results of this analysis are shown in the bottom half of Figure 3.18 for the 13 and 38 mm cut cases, at both the 51 and 102 mm array offset distances. In each case it was discovered that a $N = 4$ element array was necessary for effectively perfect detection, but for $N < 4$ detection performance was somewhat reduced,

ultimately approaching (but still beating) the performance of using a single unprocessed receiver. These results are logical outcomes of the fact that STR requires array weights (typically derived from a beamforming operation [70]), and sufficient array resolution to distinguish signal arrival directions for success. However, the estimation of such weights and the resolution of signal arrival directions is worsened as the number of receivers – and subsequent array length – are reduced.

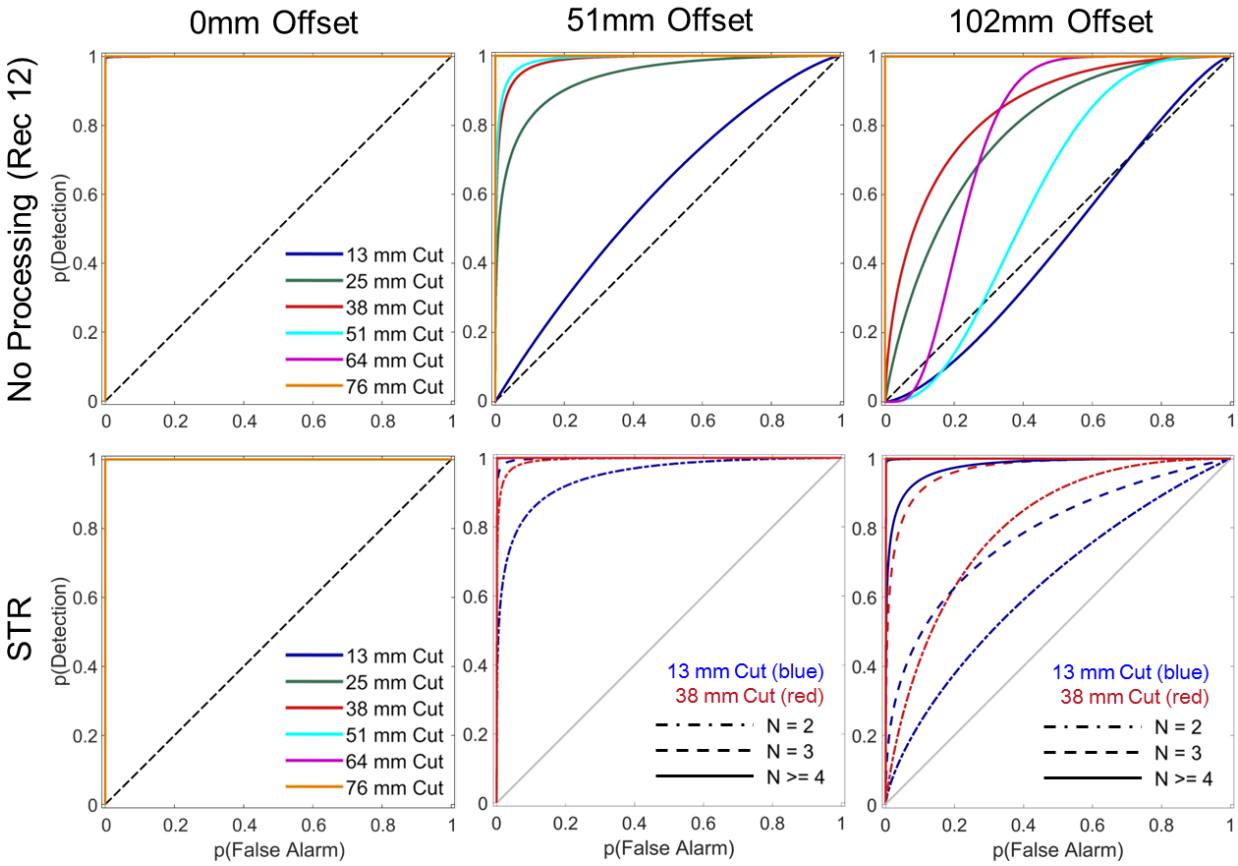

Figure 3.18: ROC curves for varying levels of baseline-recording offset (0, 51, and 102 mm) using a single receiver recording with no processing. It is clear that even for small amounts of geometric uncertainty, detection performance suffers greatly. Corresponding ROC curves were produced using STR to correct for geometric uncertainty, and in each case the ROC curves depicted essentially *perfect* detection, visually identical to the plot for 0 mm offset above. All ROC curves were computed using the Δ_{PSD} detection metric.

It is common to reduce ROC curve results like those shown in Figure 3.18 to single scalar values for ease of comparison [86]. One such value is the Deficit Area (DA) which is the integral of the area between a given ROC curve, and the line corresponding to an ideal probability of detection of 1. An ideal detector exhibits a DA of 0, while a very poor detector exhibits a DA of roughly 0.5 (statistically equivalent to coin flipping). DA values corresponding to the ROC curves in Figure 3.18 are provided in Table 3.3. It is evident that without processing the receiver recordings to account for the presence of multipath, the remote acoustic detection method is highly sensitive to geometric variations. Fortunately, the inclusion of the blind-deconvolution algorithm yields negligible performance impact for all cases tested.

		No Processing			**Synthetic Time Reveral**		
		Baseline Offset (mm)			Baseline Offset (mm)		
		0	*51*	*102*	*0*	*51*	*102*
Cut Length (mm)	*13*	0.02	0.41	0.52	0	0	0.0004
	25	0.0002	0.12	0.26	0	0	0
	38	0	0.04	0.22	0	0	0
	51	0	0.03	0.38	0	0	0
	64	0	0.01	0.27	0	0	0
	76	0	0	0	0	0	0

Table 3.3: Deficit Area (DA) values for the associated ROC curves in Figure 3.18. Pink cells highlight DA > 0.1, which implies poor performance. Without multipath and reverberation compensation (left half of the table), detection performance is poor for moderate baseline offsets, while detection using STR was found to be robust in all cases.

3.2.3b Conditions for Input Forcing Repeatability

For the presented baseline comparison techniques to be successful, certain information about the input forcing must be known. If a direct measurement of the input forcing is unavailable, a sufficient condition for accurate comparison is that the PSD of the input forcing is identical – or at least sufficiently constant – throughout successive measurements. Input forcing repeatability can be quantified by computing the pairwise PSD correlation coefficients (Eq. (3.9) from a batch of successive input forcing measurements, then computing the mean of this set. This mean PSD input forcing correlation serves as a ceiling for the *acoustic* correlations discussed in Section 3.2.3a Robustness to Multipath Variation. As such, a sufficient value for the mean input forcing correlation should exceed the maximum correlation derived from the acoustic data. For example, from the baseline Δ_{PSD} histograms shown in Fig. 4b), an acceptable mean input forcing correlation should exceed roughly 0.99.

The PSD of noisy signals in general do not exhibit correlations of ~0.99, though smoothing of the PSDs can increase the correlations to such values. The implicit tradeoff is that too much smoothing erases valuable spectral structure which discriminates the presence of damage (and the resulting radiated sound changes) from the lack thereof. The goal then is to determine experimental parameters which allow smoothing of the statistically *random* variations associated with a noise excitation, without discarding the *bias* variations in the acoustic recordings associated with changes in structural health.

To quantify the influence of smoothing and input-excitation consistency, a brief analysis was performed to ensure input forcing was sufficiently repeatable in the above experiments. A batch of 16 noisy input forcing measurements (100-2000 Hz bandlimited, 8.5 second durations)

were processed into PSDs then pairwise correlated and averaged. The time-domain signals were then slightly truncated and the mean correlation was again computed, and so on, until the signals were reduced to 0 seconds in duration. This process was then repeated with the inclusion of a boxcar smoothing operation applied to the PSDs pre-correlation. Boxcar windows of 2, 5, 10, and 20 Hz were investigated. Figure 3.19 shows the mean input correlations as a function of signal duration and smoothing window width. It is evident that if no smoothing is applied, correlations are less than 0.8, well below the 0.99 required for detection as per Figure 3.16b. However, the inclusion of smoothing allows the 0.99 target to be achieved, and with a 5 Hz smoothing window, a sample duration of approximately 6 seconds is required (see inset panel in Figure 3.19). Acceptably high mean correlation values may be achieved with even shorter time durations; however, the fine structure of the PSDs – which are important for distinguishing damage – may be lost, reducing detection performance. A compromise of a 5 Hz smoothing window with 8.5 second signals was used for the results presented throughout this paper, though optimal settings would in general be application dependent.

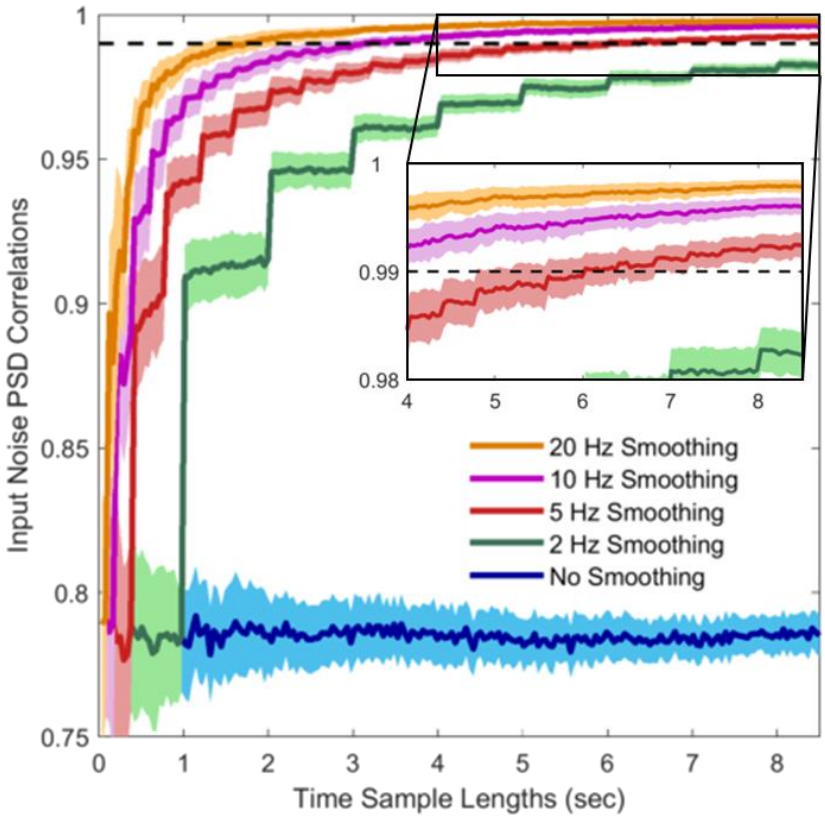

Figure 3.19: Input forcing correlations as a function of signal duration and PSD box-car smoothing filter width (in Hz). Mean correlations are shown as dark lines, with single standard deviation spread shown in lighter shades. The inclusion of PSD smoothing allows the target correlation of 0.99 to be achieved (dashed line). With 5 Hz smoothing, a time duration of roughly 6 seconds is required – less than the 8.5 seconds used in the experiments reported here. The 'stairstep' jumps occur when the time duration is increased such that an additional frequency bin falls within the smoothing window width, increasing the total number of bins being averaged.

3.2.4 Summary and conclusions

An investigation has been conducted to determine if remote acoustic sensing with an array of 15 microphones is a viable method for detecting cuts with lengths from 13 mm to 76 mm in a vibrating base-excited, clamped 30-cm-by-30-cm-by-1.6-mm plate in an unknown reverberant environment at a signal-to-reverberation ratio of –7 dB when baseline recordings of the uncut plate are available. Results were collected for six different cut lengths, for

deterministic and random vibratory excitation in a 100-Hz-to-2-kHz bandwidth, and for three different detection metrics: time-domain (TD), power-spectral-density (PSD), and frequency-response-function (FRF) cross-correlations. The vibratory forcing excited acoustic radiation from several plate-vibration modes. The collected detection results come from thousands of microphone recordings and are based on using synthetic time reversal (STR) to reduce the 15-microphone-array recordings to a single signal that is at least partially corrected for reverberation-induced signal distortion. Additional tests that included plate-to-array geometrical changes confirmed the utility of STR for this application. Detection results are presented in terms of the detection index, ROC curves, and ROC-curve deficit area.

This research effort leads to four conclusions. First, when array recordings can be made and baseline recordings are available, then remote acoustic detection of cuts as small as 13 mm in a 30-cm-square plate is possible in an environment with significant unknown multipath propagation and reverberation. The approach followed here has also been found effective for detecting boundary-clamping defects [17] with repeatable frequency-sweep excitation. The present work extends that detection success to plate cuts and to both repeatable and noisy (not repeatable) vibratory excitation. Second, when the vibratory excitation is not repeatable, the TD cross-correlation detection metric fails. Although this conclusion might be readily anticipated from the wealth of prior work in the structural health monitoring realm, it is stated here for completeness. The remote acoustic techniques utilized here do not overcome this limitation. Third, when the vibratory excitation is not repeatable, both PSD and FRF cross-correlations metrics can be successful with FRF-based detection being the better of the two. However, FRF-based detection requires measurement or knowledge of the vibratory excitation waveform and might therefore be more difficult to implement. By comparison, the PSD metric merely requires

the vibratory excitation's PSD to have the same frequency dependence during the baseline and test measurements; the actual excitation waveforms can differ. And fourth, geometrical robustness attributed to use of STR in the prior boundary-clamping-defect study persists in the current findings even though the current work primarily involves random (not repeatable) structural excitation. This conclusion is promising for remote detection applications in fully or partially enclosed environments, where boundary reflections and echoes are common, and the time elapsed between the collection of baseline and test recordings is long enough so that the source-to-receiver geometry and environmental characteristics are likely to have varied because the original (baseline) experimental configuration is only approximately known or cannot be precisely duplicated.

Though the FRF metric outperformed the PSD metric, they were both fully effective for detecting cuts of the severity investigated, even with noisy input excitation. Though FRF detection methods would be expected to outperform PSD methods in general, they come with the cost of requiring some information about the input excitation. This could either be measured directly, or estimated via some model or known parameters (e.g. knowledge of the speed of a ship may allow one to infer the rate and loads due to the engines). Further, requiring measurement of the input excitation is the antithesis of this *remote* acoustic SHM endeavor. Benefits and drawbacks of the tested metrics are organized in Table 3.4. Due its adequate performance and lack of a requirement to know input forcing, the PSD metric was chosen as the best metric to apply in general. Nevertheless, there are certain to exist scenarios where other metrics – or perhaps functions of multiple metrics – may be needed. One immediate such example in that of completely *nonstationary* noisy inputs, wherein the PSD of the noise is a

reasonably fast function of time such that the PSD metric would be unsuited to make appropriate comparisons between baseline and test cases.

	Benefits	Limitations
Δ_{TD}	Minimal processing. Only requires array recordings Potential benefits for transient signals	Very sensitive to changes in phase/random phase signals
Δ_{PSD}	Ease of implementation Robust to phase changes Only requires array recordings	Sensitive to changes in PSD of input forcing
Δ_{FRF}	Most robust to changes in input forcing Marginal performance improvement to Δ_{PSD}	Requires knowledge of input forcing or additional measurements

Table 3.4: Cost-benefit analysis table for the three detection metrics considered. A PSD based metric was determined the best compromise due to the performance improvements compared to time-domain cross-correlation, and the ease of measurement compared to FRF based metrics. If the input forcing is expected to be highly non-stationary FRF based methods may be necessary.

3.3 Detection of delamination in composite materials

Composite materials are becoming more accessible, affordable, and ubiquitous in many engineering applications due to their high strength-to-weight ratios, improved stealth characteristics, and ability to conform to complex geometries [87]. One of the most common types of composite materials in use today are fiber-reinforced composite laminates [88]. Fiber-reinforced laminates consist of multiple layers – or lamina – of woven fiber (sometimes woven directionally, other times randomly) which are bonded (or *laminated*) using a matrix material, typical epoxy resin. Depending on the relative ordering and orientation of the fiber layers, various directionally dependent stiffnesses may be obtained in a laminate depending on design requirements. Despite these benefits, a major concern for the accelerated implementation of

fiber-reinforced composites is the problem of delamination [89]. Delamination occurs when adjacent fiber layers become unbonded, creating a region of low-stiffness which is prone to spreading and may ultimately lead to catastrophic system failure. One of the challenges of detecting delamination is that it is typically visually obstructed due to being one or more layers below the surface. As such, the vast majority of detection strategies for delamination in the literature utilize contacting vibration-based methods discussed in Section 1.1.1, such as guided wave methods [15] and accelerometer measurements of changes in modal characteristics [89]. There are few, if any, results in the literature for remote acoustic detection of delamination.

The motivation for investigating remote acoustic detection of delamination is twofold: 1.) delamination is an additional failure mode of interest to test, and 2.) since the fabrication of composite plates leads to inherent structural uncertainties, such an experiment provides insight into the performance of the technique when used with non-ideal baseline measurements. Four 12”x12”x0.06” composite plates were fabricated using a common vacuum bagging method [90] with biaxial weave carbon fiber sheets and epoxy resin. Each plate consisted of 6 layers of carbon fiber, each with the principle fiber directions running parallel with the two square axes. The number of layers per plate was selected such that the modal density (in frequency) of the composite plates would be comparable to that of the previously tested aluminum plates. Two of the four plates were a fully bonded in a typical fashion, one of which served as the *baseline* plate and the other as the *known healthy* plate, each with nominally identical vibroacoustic properties. Of the remaining two plates, thin layers of aluminum foil were carefully placed between the medial layers during the layup process, as to prevent bonding within a localized region, thereby approximating delamination. One test plate featured a 3” square delaminated

region, shown in Figure 3.20, and the final test plate featured a 1” square delaminated region (centered at the same location as the 3” square region of the third plate).

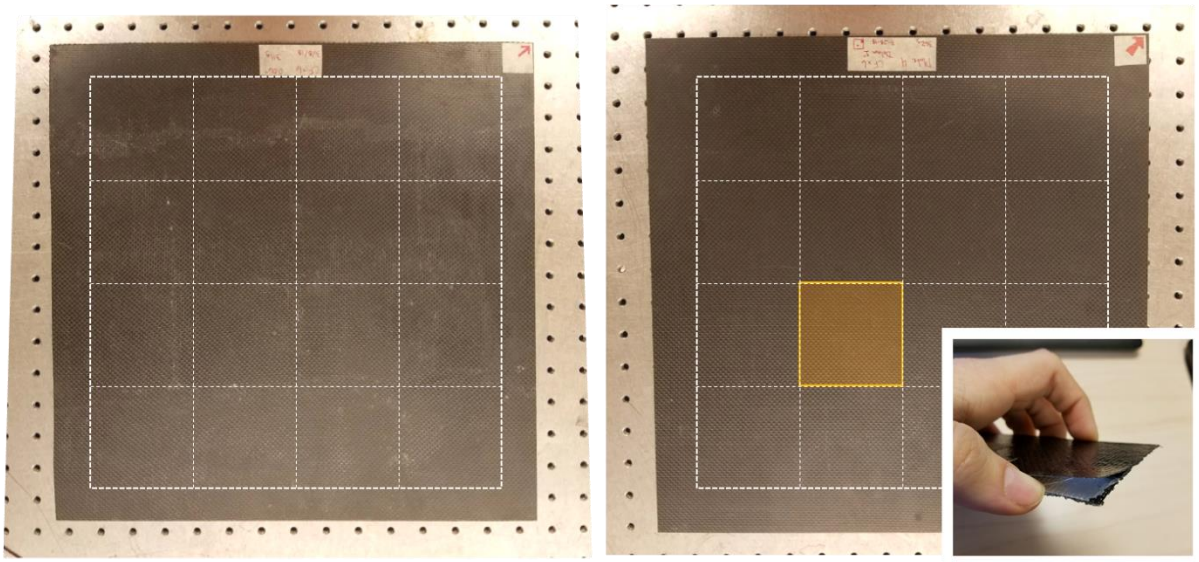

Figure 3.20: Six layer biaxial-weave carbon fiber plate with no-delamination (*left*), six layer biaxial-weave carbon fiber plate with 3” square delamination in highlighted region (*right*), test coupon showing synthesized delamination (*inset*).

Effectively the same methodology outlined in Figure 3.7 was applied for the delamination detection tests. 0.1-2 kHz 8-second linearly swept chirps were used to excite the plate and the acoustic response was recorded with a 15-element vertical microphone array with 2” spacing. Based on the results presented in Section 3.2 Detection with stochastic input forcing, comparisons between Power Spectral Densities were used to quantify changes in plate response. In addition to acoustic measurements, vibrational measurements were collected with an LDV to provide insight into mode shape changes due to the presence of delamination. Comparisons of measured mode shapes for delaminated and fully laminated plates are provided in Figure 3.21. The presence of delamination has some visually discernable effect on the lower plate modes, but it is not extreme. Note that the (4,2) mode of the baseline plate switched to a

(2,4) mode with delamination. In addition to slight changes in the mode shapes, significant changes in the modal frequencies was also observed in the LDV measurements. The presence of the 3" square delamination region reduced resonance frequencies by roughly 2-4% for the modes inspected, which is agrees with the qualitative prediction of delamination behaving as a localized reduction in stiffness.

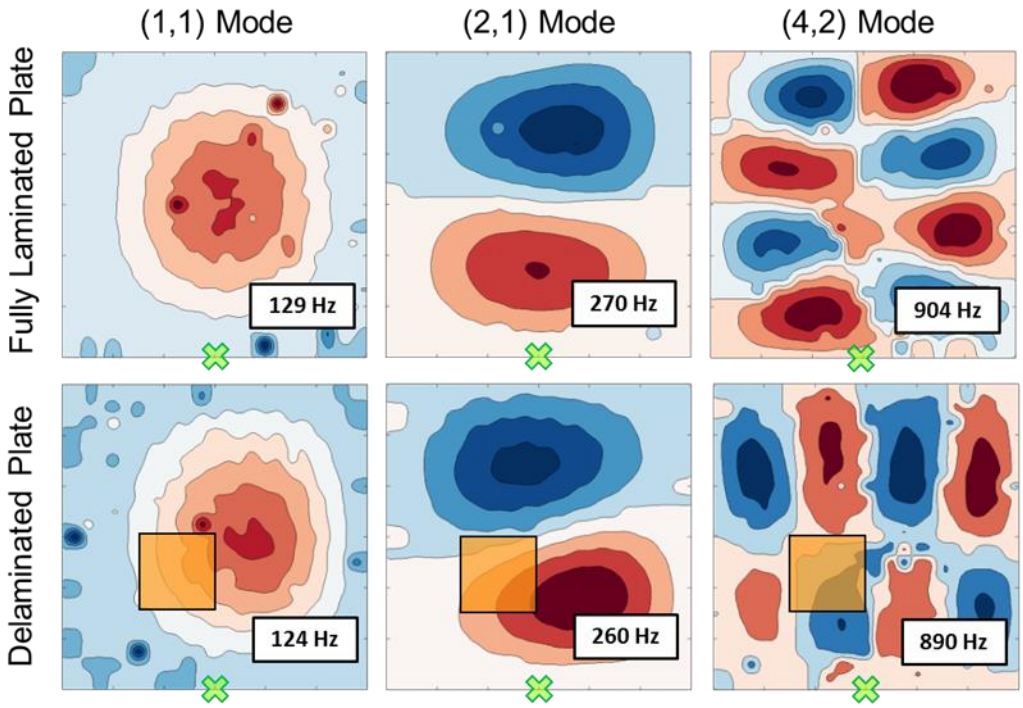

Figure 3.21: Comparison of measured mode shapes for three vibrational plate modes. Fully laminated (baseline) carbon fiber plate (*top*) and carbon fiber plate with a 3" square delamination region, highlighted (*bottom*).

Detection results are quantified as before, using detection histograms of the PSD metric between the test plates and the baseline plate. From Figure 3.22 it is clear that both the 3" and 1" delamination regions are easily distinguishable, based on their separation from the baseline-to-known healthy plate histogram. Further, the 3" delaminated region was noticeably "easier" to detect (i.e. greater separation distance) than the 1" delaminated region, which is to

be expected as it is roughly 9 times greater in severity. An additional red histogram, shown in Figure 3.22, indicates the result of taping two 1” square pieces of aluminum foil to the top and bottom of the known healthy plate. This was done to verify that the results of the delamination detection were not due the added *mass* of the aluminum foil (but rather the reduction in stiffness associated with unbonding). Though some separation is present between the added foil case and the baseline case (which could indicate an ability to detect small localized changes in mass due to the foil), the stark difference between the red and gold histograms implies that the observed change in power spectral response for the 1” square delaminated plate (compared to the baseline) was likely *not* due to the added mass of the foil. These results indicate that remote acoustic sensing is feasible for detecting the presence delamination in composite plates, even when imperfections due to fabrication errors exist in the baseline and healthy plates.

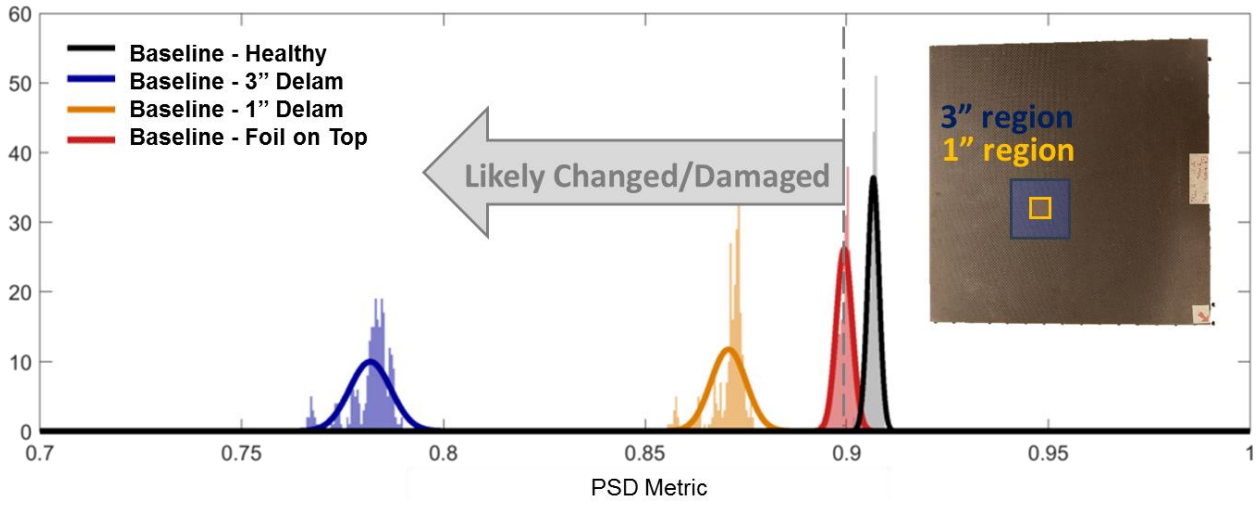

Figure 3.22. Detection histograms for localized delamination. Schematic representations of the 3” and 1” square delaminated regions are shown at the right-hand side of the plot. Good separation of the histograms indicates good detection performance for both the 3” and 1” delamination using remote acoustic sensing.

Chapter IV: Localization of Damage

4.1 Introduction

There are few examples of *remote* acoustic sensing used explicitly for damage *localization* in the literature. Nearfield Acoustic Holography has seen application in damage localization [12], however since it is inherently *nearfield* it does not practically fall under the umbrella of *remote* acoustic. Remote measurements of bearing noise have been used to estimate where on a system (even where on a single bearing) a fault may be located [91], however this is due primarily to the simple geometry of bearings which lend themselves well to models. For remote localization in systems wherein *a priori* models are either unavailable or infeasible, beamforming with a receiver array is a potential recourse. However, beamforming with the intent of localizing small changes in a vibrating body is a difficult task for the following two reasons:

1.) *Small defects should result in small changes to the radiated acoustic field*. In essence, for reasonably small damage the goal is to beamform accurately enough to detect a change over a small spatial extent. This is further complicated due to the spatial filtering effect of a beamformer, any finite-size array exhibits an aperture function (in some cases referred to as a Point Spread Function) that will convolve itself with the

source field [52], effectively smearing out the small (desired) changes with the rest of the large (undesired) background field. This then calls for the use of a high resolution beamformer, of which there are fortunately many [44].

2.) *The majority of the acoustic radiation emanating from the structure of interest is both undesired and coherent.* This is because distributed structures – like a plate – are phase-locked, meaning different regions will generally oscillate in or out-of-phase, albeit at different amplitudes. This coherent field, heretofore referred to as a *Coherent Background* causes mathematical problems for most of the common high resolution high resolution beamformers such as MUSIC [59] and MVDR [92].

One common method for mitigating the first of these effects is to implement a beamformer deconvolution algorithm such as CLEAN [23], or DAMAS [22]. Both CLEAN and DAMAS/DAMAS2 are iterative computational techniques which aim to reconstruct the 'true' non-smeared source field by successively removing smeared peaks and replacing them by sharp 'true' peaks. CLEAN – which originally came out of the astronomy community for 'cleaning up' telescope images [93] – has found moderate application in acoustics, but the two main drawbacks of deconvolution algorithms like these are that they require accurate knowledge of the array's aperture function, and they can suffer in noisy environments [22].

For the second challenge, associated with the coherent background field, a technique sometimes referred to as spatial smoothing (but in this document referred to as subarray averaging) has been shown to partially circumvent the coherence issue with certain high

resolution beamformers [60]. Subarray averaging is an operation which can be applied directly to uniformly spaced line arrays, wherein the array is segmented into smaller, overlapping *subarrays*, then these subarrays are used to produce a modified cross spectral matrix (CSM) which is simply the average of the CSMs over all subarrays. A schematic of this procedure is shown in Figure 4.1. Mathematically, this is justified in the case of MUSIC by constructing a new CSM which has full rank (unlike the original CSM which features at least two degenerate eigenvectors when two or more coherent sources are present) [60].

Another somewhat newer high-resolution beamformer, titled the Spectral Estimation Method (SEM), has certain properties making it useful for this problem. SEM, introduced by [63] is a CSM based beamformer which implements a convex optimization for the minimization of the L_2-normed difference of a measured CSM and a modeled CSM. The result of this approach is a somewhat robust beamformer with very high resolution, due to the optimization step. An additional benefit of SEM is that, due to the way it is formulated, it can also be extended to reject noise from spatially distributed background noise sources, so long as the noise power spectral density is time invariant and a reference measurement solely of the noise is available. Blacodon titled this extension to SEM, SEM with Additive Noise (SEMWAN) [94]. It is crucial to note, however, that SEMWAN is only directly applicable when applied to background *noise* sources, i.e. unwanted sound sources that are mutually *incoherent*. Out of the box, SEMWAN does not apply to *coherent* acoustic backgrounds.

Let $r_m(t)$ be the time domain recording at the m^{th} receiver with Fourier transform denoted $\hat{r}_m(\omega)$. By definition, the cross spectral density matrix (CSDM), $\Gamma(\omega)$ is given

$$\Gamma_{\mathrm{mn}}(\omega) \equiv \hat{r}_m(\omega)\hat{r}_n^*(\omega) \tag{4.1}$$

It is also known that the recording $\hat{r}_m(\omega)$ is related to the acoustic sources and the environmental Green's function

$$\hat{r}_m(\omega) = \sum_k^K G_{m,\,k}(\omega)\hat{s}_k(\omega) \tag{4.2}$$

where $\hat{s}_k(\omega)$ is one of the K unknown source waveforms, k is simply an index standing in for a potential source location, and $G_{m,\,k}(\omega)$ is the Green's function from location k to location m. Then, $\Gamma(\omega)$ can be expressed in terms of Green's functions and the unknown source waveforms.

$$\Gamma_{mn}(\omega) = \left(\sum_k^K G_{m,k}\hat{s}_k\right)\left(\sum_l^K G_{n,l}^*\hat{s}_l^*\right) = \sum_k^K\sum_l^K G_{m,k}\hat{s}_k\hat{s}_l^*G_{n,l}^* \tag{4.3}$$

If the source waveforms are assumed to be *incoherent* then, by definition,

$$E\{\,\hat{s}_k\hat{s}_l^*\,\} = \hat{s}_k\hat{s}_l^*\delta_{ik} = \hat{s}_k\hat{s}_k^* \equiv S_k \tag{4.4}$$

where $E\{\cdot\}$ is the expectation value operator, and S_k is the positive semi-definite, real-valued power spectrum density (PSD) of $\hat{s}_k(\omega)$, or in other words, the spatial distribution of the sources, which we desire. Then

$$\Gamma_{mn}(\omega) = \sum_k^K G_{m,k}S_kG_{n,k}^* \tag{4.5}$$

If a reasonable estimate can be made for the Green's function (i.e. free-field, plane wave, etc.), then, provided some measured CSDM, $\Gamma_{m,n}^{meas}(\omega)$ is available, S_k can be estimated by minimizing the following L_2-norm.

$$\min_{S_k}\left(\left|\Gamma_{m,n}^{meas}(\omega)-\sum_{l}^{K}G_{m,k}S_kG_{n,k}^{*}\right|^2\right) \tag{4.6}$$

This specific optimization problem is convex for positive definite S_k. It can be performed in MATLAB using any of a number of toolboxes or third-party software, such as the popular CVX package (used throughout this chapter) [41]. Further, through this formulation it is immediately evident that, provided a reference measurement of a static noise field is available, say $\Gamma_{m,n}^{ref}(\omega)$, then the noise field can effectively be subtracted out in the CSDM domain, allowing for significant SNR improvements if localization of signals buried in noise is the goal. This is effectively the form of SEMWAN

$$\min_{S_k}\left(\left|\Gamma_{m,n}^{meas}(\omega)-\Gamma_{m,n}^{ref}(\omega)-\sum_{l}^{K}G_{m,k}S_kG_{n,k}^{*}\right|^2\right) \tag{4.7}$$

Refocusing on the original goal of localization, this technique cannot be immediately extended for the subtraction of the coherent background, exactly because for coherent sources

$$E\{\,\hat{s}_k\hat{s}_l^{*}\,\}\neq S_k \tag{4.8}$$

Instead we can exploit subarray averaging. Subarray averaging is the process of segmenting the array into smaller, overlapping *subarrays*, then using these subarrays to produce a modified

CSDM which is simply the average of the subarray CSDMs. Figure 4.1 shows a schematic of the subarray averaging procedure.

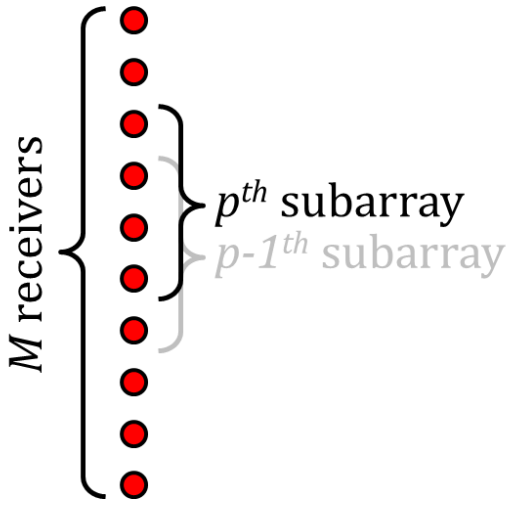

Figure 4.1: Schematic for implementation of subarray averaging on a uniform line array of receivers.

It can be exactly shown for plane waves impinging on a uniform line array, that, for an array dividing into P subarrays the, averaged CSDM $\Gamma_{m,n}^{avg}$ is such that

$$\lim_{P\to\infty} \Gamma_{m,n}^{avg} = \sum_{k}^{K} G_{m,k} S_k G_{n,k}^{*} \tag{4.9}$$

which is exactly the form of the CSDM for mutually incoherent sources in Equation (4.5). This subarray averaging approach now allows for the application of SEMWAN, but with references of coherent backgrounds instead of incoherent noise backgrounds. As such this extension is referred to as SEM with Coherent Background Removal (SEMCBR). SEMCBR can also be shown to work with spherical, nearfield Green's functions. In practice, the performance of SEMCBR depends on several parameters, some of which are user selected, such as the number

of subarrays to divide the full array into. It is also worth noting that while more subarrays increase the number of averaged CSDMs, thus better cancelling the coherent cross terms as shown in Equation (4.9), the benefits do not continue indefinitely. This is because as the number of subarrays increases, the length – or aperture size – of each subarray decreases, thereby decreasing resolution [52]. There is indeed a 'sweet spot' for the number of subarrays to apply.

The following sections detail the use of SEMCBR for the localization of damage in vibrating square plates, first using a one-dimensional uniform line array of 15 receivers, then using a two-dimensional planar Cartesian array of 64 receivers.

4.2 1-D localization

A proof of concept simulation was done to showcase the capabilities of SEMCBR. In this simulation, two sources were placed on a line collinear to a 15-element uniform linear array at a range of 60 receiver spacings (3 m). The lateral coordinate of one source was located at -0.15 m from array center, and the other source was located at +0.30 m. The spherically radiating sources are such that the wavenumber ($k = 2\pi/\lambda$) multiplied by the receiver spacing was 3. The baseline condition consisted of both sources emitting with equal amplitude. For the 'damaged' condition, the +0.30 m source was reduced in amplitude by 50%. Beamformer outputs showing the baseline and damaged conditions are shown at the top and bottom of Figure 4.2.

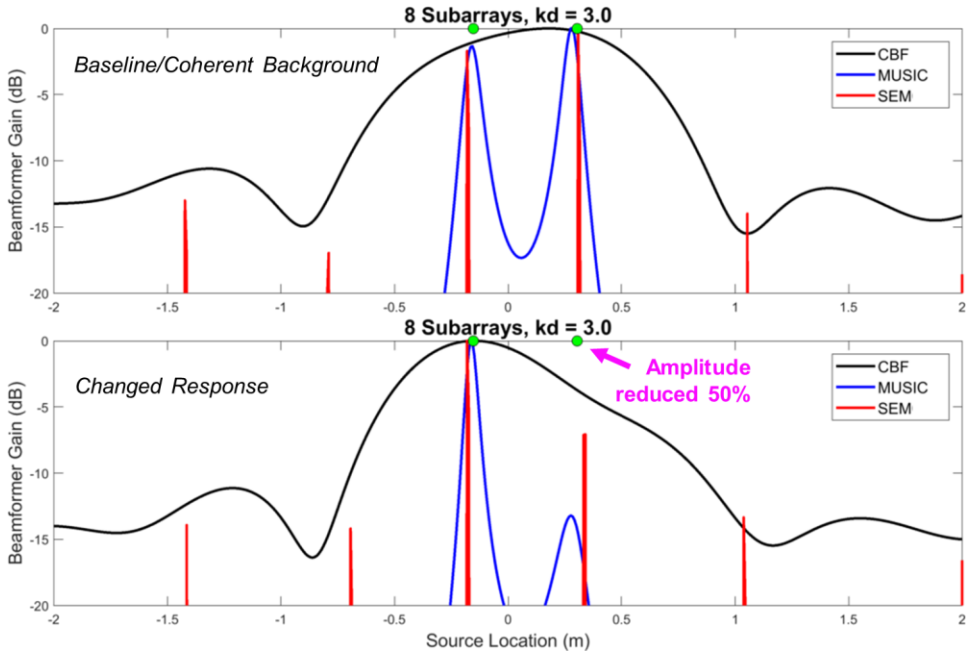

Figure 4.2: (*Top)* The beamformer output for two equal amplitude monopoles situated at -0.15 and +0.30 m on an axis parallel to the line array, separated by 3 m. (*Bottom*) Same simulation, but the source at +0.30 m is reduced in amplitude by 50%.

Figure 4.3 details the results of SEMCBR for increasing numbers of subarrays. It is clear that with no subarray averaging (effectively using the full array) the background subtraction fails. With four subarrays, performance improves and for eight subarrays, the peak is highly accurate with a peak-to-sidelobe ratio of greater than 20 dB, thus validating the method.

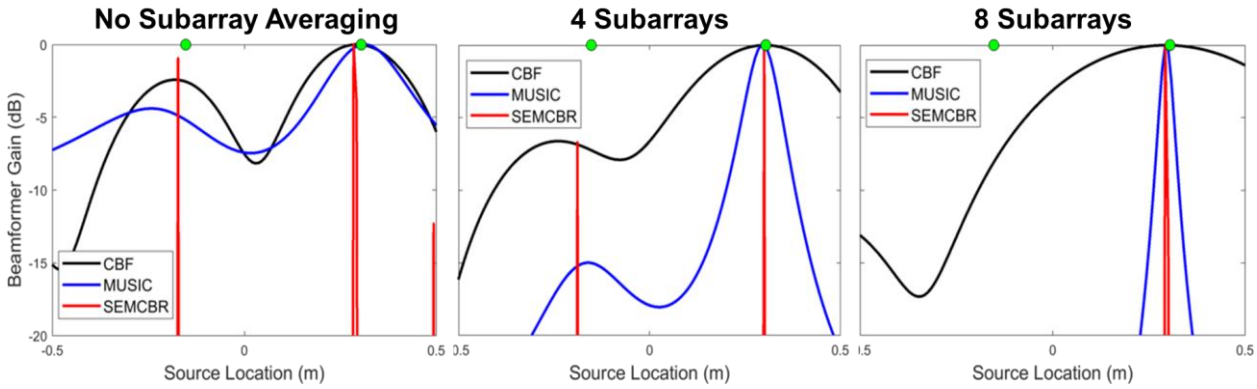

Figure 4.3: SEMCBR results for the problem statement in Figure 4.2. Subtraction performance improves as the number of subarrays is increased. When 8 subarrays are used, the peak to sidelobe ratio exceeds 20 dB.

4.2.1 One-dimensional Localization – Experiment

SEMCBR was tested experimentally on vibrating plates featuring through cuts of various length made with a band saw. A 1.5" cut was placed into a 12 x 12 x 1/16" cm aluminum plate which was clamped into the plate vibration apparatus detailed above. A uniform linear array of 15 microphones with 1" receiver spacing was placed parallel to the plane of the plate, 9.5" above the cut, as shown in Figure 4.4. For comparison, comparable NAH experiments have applied standoff distances of less than 1 inch [12, 95]. A one-dimensionally constrained spherical wave Green's function was implemented such that the beamformer would search for the 'acoustic change' just over the line-of-interrogation, shown in magenta in Figure 4.4.

Acoustic measurements of 0.1-4 kHz input chirps were collected for both a baseline, not-cut plate, in addition to the plate with a 1.5" cut. SEMCBR with P = 10 subarrays was run for each 5 Hz increment across the 0.1-4 kHz bandwidth. Results of the analysis are shown in Figure 4.5. The colormap in the background of Figure 4.5b is simply the conventionally beamformed source mapping which indicates the vibrationally active frequency bands as well

as some of the spatial character of the various plate modes near the line-of-interrogation. At each inspection frequency (in 5 Hz intervals) the output of the SEMCBR algorithm was computed, multiplied by the linear value of the FFT magnitude in Figure 4.5c (such that localizations in vibrationally active bands are weighted more that localizations in inactive bands) and finally binned into the histogram shown in Figure 4.5d. The final declaration for the cut location is determined from the maximum value of the histogram, which only exhibited ~0.5" error, and this is considered a successful proof of concept. The experiment was repeated with the plate rotated by 90° such that the cut was in line with a more active region of the plate. These results, shown in Figure 4.6, are very similar. The cut is again successfully located, on average, by leveraging many measurements throughout the bandwidth. The distribution is, however, more spread than the first experiment, indicating greater uncertainty. This is likely due to the change in relative position of the array (see the left side of Figure 4.6), which picks up a greater amount of sound from the highly-active center of the plate, effectively drowning out the small change due to the cut.

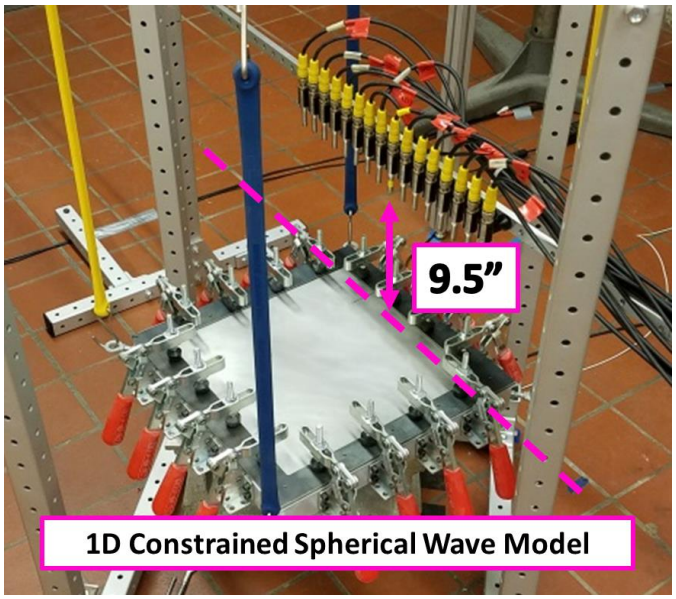

Figure 4.4: SEMCBR experimental validation for localization of a cut.

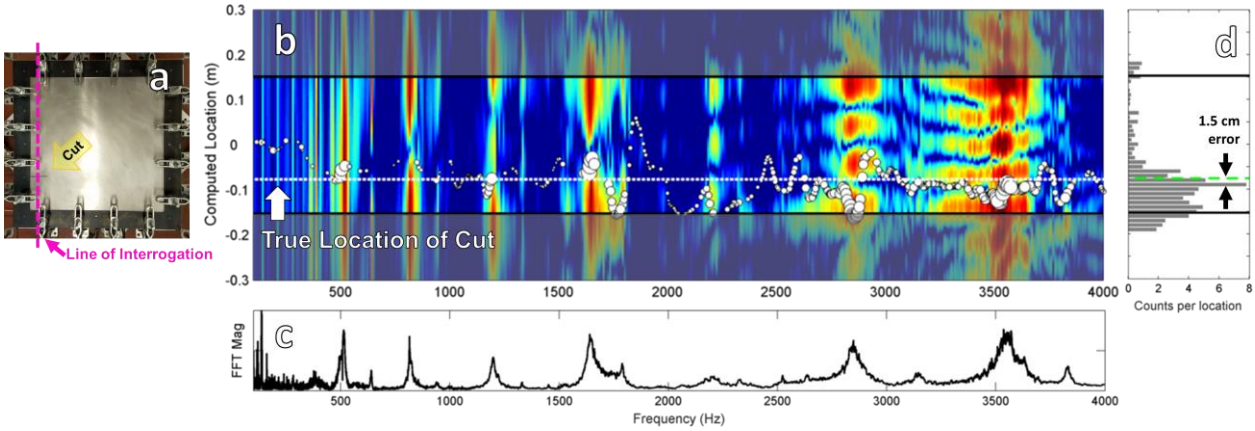

Figure 4.5: Results for the SEMCBR localization of a cut on an aluminum plate. Plate and cut geometry are shown in a.) Part b) indicates the true location of the cut, along with SEMCBR estimates calculated at every 5 Hz across the bandwidth. Overlap in the white circles (SEMCBR outputs, weighted by FFT magnitude at that frequency) indicate accurate localization. Localization results were ultimate binned into the histogram in part d) which shows only 0.5" error in the computed and actual cut locations (over a 24" search space).

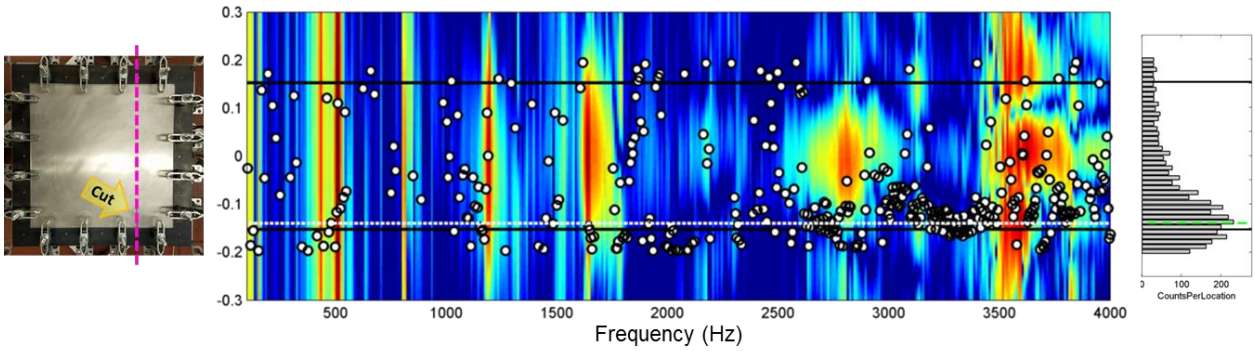

Figure 4.6: Results for the SEMCBR for the same experiment as shown in Figure 4.5, but with the plate rotated by 90 degrees, placing the cut in a different location. The cut is again successfully localized using the 0-4000 Hz bandwidth; however, the uncertainty is somewhat greater than for the first case.

4.3 2-D localization

Two-dimensional maps of acoustic energy are generally more informative than one-dimensional maps. Not only can a greater extent of a structure be investigated using a two-dimensional map, but interpretation of the results is more straight-forward. Fortunately, the mathematical formulation for 1-D SEMCBR can be readily extended to 2-D, provided a Cartesian array of two-dimensional extent is available.

4.3.1 Spectral Estimation Method

Given an M element array with receiver locations $\boldsymbol{x}_m$impinged upon by K monopole sources, each represented by generic time signatures $s_i(t)$, the Fourier transform of the measured signal at each receiver (in the absence of noise) is modelled as

$$\hat{r}_m(\omega) = \sum_{i}^{K} \hat{G}_{m,i}(\omega)\hat{s}_i(\omega) \tag{4.10}$$

where $\hat{G}_{m,i}(\omega)$ denotes an appropriate Green's function from the i^{th} source to receiver location $\boldsymbol{x}_m$. The goal of acoustic localization is to estimate the signal energies $S_i(\omega) = |\hat{s}_i(\omega)|^2$ originating from each potential source location within some domain. To do this, typically a cross spectral density matrix (CSDM) $\Gamma(\omega)$ is computed, consisting of cross-correlations between measured receiver data

$$\Gamma_{mn}^{meas}(\omega) = \hat{r}_m(\omega)\hat{r}_n^*(\omega) \tag{4.11}$$

If the measured signals obey the form described in Equation (4.10), then $\Gamma(\omega)$ may be modelled as

$$\Gamma_{mn}^{mod} = \sum_{i,j=1}^{K} \hat{G}_{m,i} \hat{s}_i \hat{s}_j^* \hat{G}_{n,j}^* \tag{4.12}$$

where the explicit dependence on ω is omitted for brevity. In scenarios where the monopole source are *mutually incoherent*, $\mathrm{E}\{\hat{s}_i \hat{s}_j^*\} = S_i \delta_{ij}$, and the modelled CSDM can be reduced to

$$E\{\Gamma_{mn}^{mod}\} = \sum_{i=1}^{K} \hat{G}_{m,i} S_i \hat{G}_{n,i}^* \tag{4.13}$$

which is achieved in practice by averaging CSDMs over a sufficient number of snapshots. S_i can then be solved for via minimization of the difference between the modelled and measured cross spectral density matrices, i.e.

$$\min_{S_i} \sum_{m,n=1}^{M} \left| \Gamma_{mn}^{meas} - \sum_{i=1}^{K} \hat{G}_{m,i} S_i \hat{G}_{n,i}^* \right|^2 \tag{4.14}$$

with the constraint that $S_i \geq 0$. This method is known as the spectral estimation method (SEM) and can yield very high-resolution estimates for incoherent sources. Further, Equation (4.14) has an additional convenience in that any available reference measurements $\hat{r}_m^{ref}(\omega)$ of a noisy background can be subtracted before the minimization procedure, i.e.

$$\min_{S_i} \sum_{m,n=1}^{M} \left| \Gamma_{mn}^{meas} - \Gamma_{mn}^{ref} - \sum_{i=1}^{K} \hat{G}_{m,i} S_i \hat{G}_{n,i}^* \right|^2 \tag{4.15}$$

for $\Gamma_{mn}^{ref} = \hat{r}_m^{ref}(\omega) \hat{r}_n^{ref\,*}(\omega)$. This modified technique, titled SEM with additive noise (SEMWAN) then allows for high-resolution localization of noisy signals amid a much louder

noise background. However, in many applications the assumption of mutually coherent sources is not valid, and in these situations the performance of SEMWAN expectedly suffers.

4.3.2 Subarray averaging

To extend the above high-resolution localization approach to problems involving coherent sources, we implement a technique known as subarray averaging, which has been previously used to improve the performance of the MUSIC algorithm for coherent sources. Subarray averaging is a four-part procedure consisting of 1) dividing a regularly patterned array into smaller, possibly overlapping subarrays, 2) computing a measured CSDM for each subarray, 3) averaging each of these CSDMs to form a 'smoothed' CSDM which can then be 4) used as an input into a beamforming method of choice. Figure 4.7 is a schematic representation of a 5x5 Cartesian receiver array divided into 3x3 subarrays.

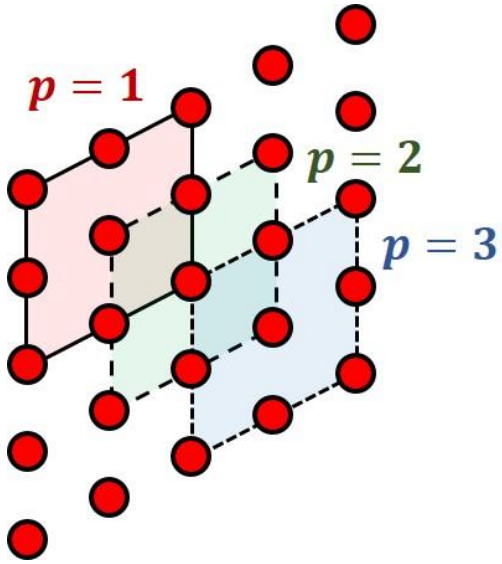

Figure 4.7: Schematic of a possible subarray configuration for a uniformly distributed Cartesian planar array.

The purpose of subarray averaging is to produce a surrogate CSDM with beneficial properties, i.e. one with a form identical to Equation (4.13). Equivalently stated, subarray averaging eliminates the terms in Equation (4.12) for which $i \neq j$, forcing the CSDM for a general distribution of coherent sources to approach that of a system of mutually incoherent sources. In the case of a uniform line array with far-field sources, the subarray averaged CSDM $\bar{\Gamma}_{mn}$ can be shown to be (Appendix A)

$$\bar{\Gamma}_{mn} = \sum_{k,l=1}^{K} e^{i\frac{\omega}{c}d(m\sin\theta_k - n\sin\theta_l)} \hat{s}_k \hat{s}_l^* \, \eta_{kl} \tag{4.16}$$

for receiver spacing d, sound speed c, and incident arrival angles θ_k. The cross-term factor η_{kl} is a function of P that is small for $k \neq l$ and unity for $k = l$. Explicitly,

$$\eta_{kl} = \frac{\sin\left(\frac{\omega d}{2c} P(\sin\theta_k - \sin\theta_l)\right)}{\sin\left(\frac{\omega d}{2c}(\sin\theta_k - \sin\theta_l)\right)} \tag{4.17}$$

Figure 4.8 contains plots of η_{kl} for two sources separated by various incident angles averaged over 15 subarrays. Points below the horizontal dashed line indicate values of $\eta_{kl} \leq 0.2$ wherein subarray averaging has been found to perform well. The vertical dashed line marks the $kd = \omega d/c = \pi$ boundary, above which aliasing becomes a concern. Increased source separation has the effect of averaging out the cross terms more quickly, thereby requiring fewer subarrays for accurate localization of coherent sources. Closer sources require either more subarrays or higher frequencies to successfully mitigate cross-terms than for more separated sources. Further, for $k = l$, the source separation can be thought of as zero, and thus the cross-term factor is unity, as expected.

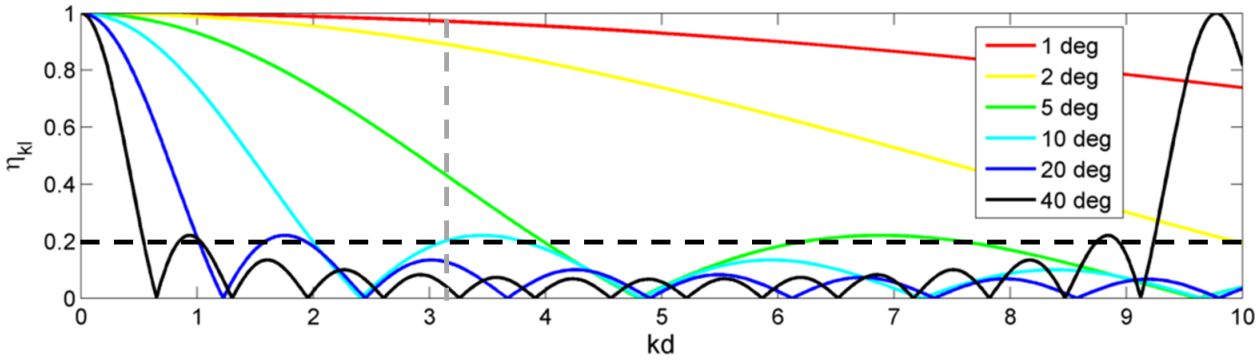

Figure 4.8: Cross-term factor for two sources separated by various incident angles with 15 subarrays averaged. Points below the horizontal dashed line ($\eta_{kl} \leq 0.2$) indicate correspond to good localization performance. The vertical dashed line corresponds to $kd = \pi$, above which grating lobes may appear.

Subarray averaging can be immediately extended to sources that are not in the farfield, provided some geometric constraints are satisfied. In the case of a uniform line array, if the sources can be assumed to all lie approximately along a line parallel to the array separated by *a priori* known distance D, then the appropriate CSDM computation then becomes a *weighted* average over subarray CSDMs (Appendix B)

$$\bar{\Gamma}_{mn} = \frac{1}{P}\sum_{p=1}^{P} e^{-i\frac{\omega d}{c D}(y_m - y_n)(p-1)}\,\Gamma_{mn}^{p} \tag{4.18}$$

where y_m is the vertical coordinate of the m^{th} receiver. Intuitively, the weighting is due to the fact that each subarray has a different center coordinate relative to the acoustic sources, and the resulting phase differences must be taken into account. The weights themselves are only functions of known geometry and are easily implemented into the subarray averaging method. Extension of the method to 2D Cartesian planar arrays is also straightforward, however subarrays then take the form of squares or rectangles.

Since subarray averaging is used to generate a surrogate CSDM of the form Equation (4.13), it is worth emphasizing that the benefits of the SEMWAN formulation are simply extended to coherent source cases when subarray averaging is used. Namely, the same minimization procedure can be used to solve for the source powers S_i, and coherent backgrounds can be directly subtracted out using reference recordings (Equation (4.15)). These combination of these two steps, henceforth referred to as SEMCBR (Spectral Estimation Method with Coherent Background Subtraction), enables the localization of even small changes in an otherwise loud, distributed backdrop of coherent sources.

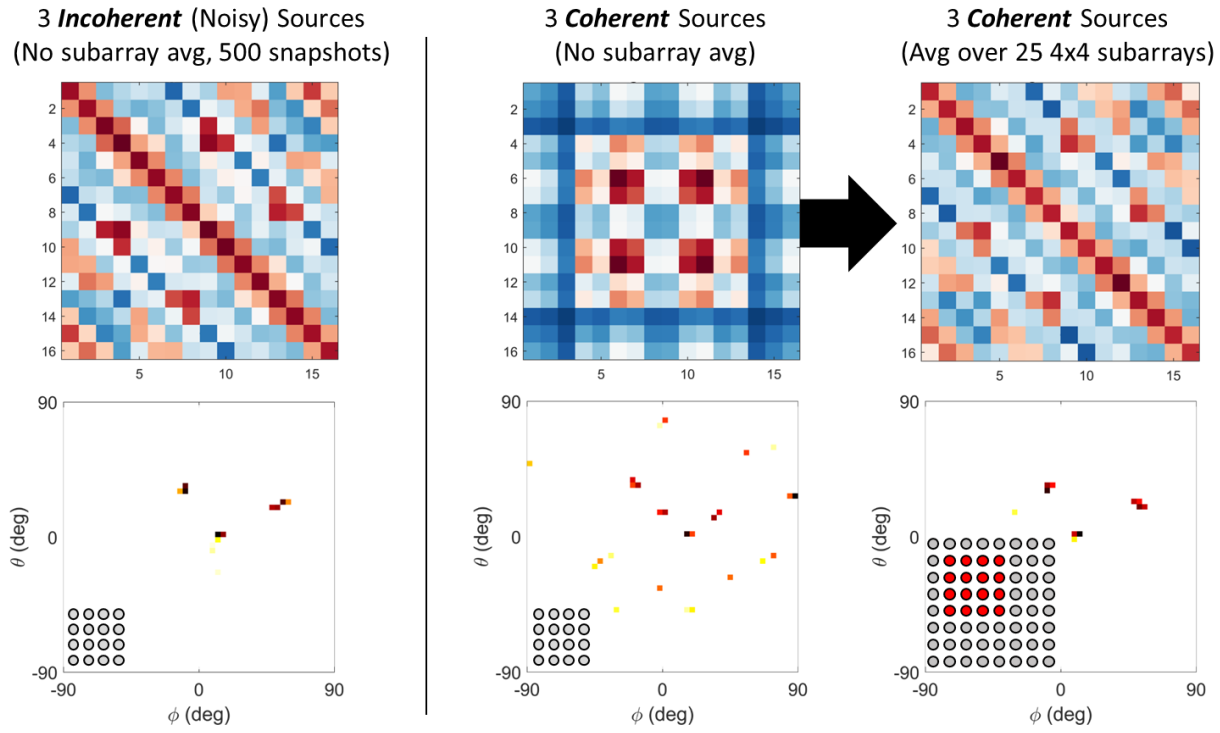

Figure 4.9: The effect of subarray smoothing on CSDM structure. Each case corresponds to three simulated farfield sources impinging on a 2D cartesian array ($kd = \pi$). When all three sources are incoherent, the CSDM has a generally banded structure (top left), resulting in effective source mapping using SEM (bottom left). The same operation applied to three coherent sources (middle) fails, and the CSDM structure is of noticeably different character than for incoherent sources (due to cross terms). However, applying subarray averaging with 25 4x4 subarrays (right), the CSDM structure approaches that of the incoherent case and the source map (bottom right) is again accurate.

4.3.3 Validation experiment

A simple proof-of-concept experiment was conducted to validate SEMCBR. An 8x8 Cartesian array of PCB 130E20 microphones (Depew, NY) with 6 cm uniform spacing was placed parallel to the floor of an untreated laboratory. Gaussian enveloped sinusoids were broadcast from three home-audio speakers on the floor 1.2 m below the array, pointed upward toward the receivers as shown in Figure 4.10. The three sources simultaneously broadcast $f =$ 5000 Hz pulses ($kd = 2\pi fd/c =$ 5.5) with 0.1 s durations. As only 16 microphones were available, measurements from an 8x8 array were generated by stitching together recordings from a smaller, shifted 4x4 array (see Figure 4.10). For precision, the array was translated between runs via a two-axis motorized gantry with 0.1 mm precision. Microphone recordings were collected using a PXIe-6368 DAQ (National Instruments, Austin, TX) at a 50 kHz simultaneous sampling rate.

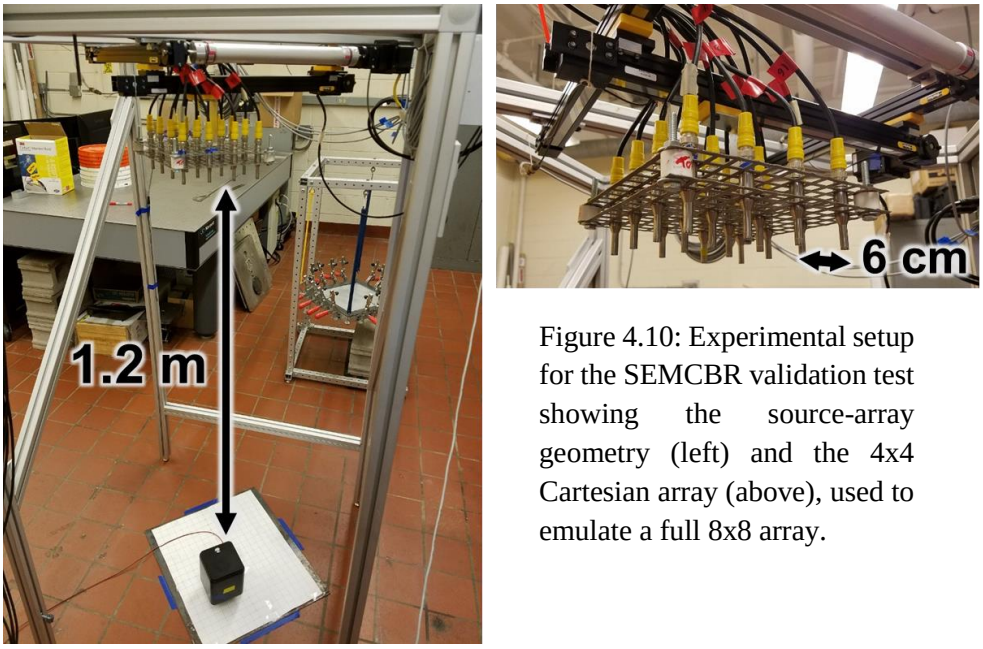

Figure 4.10: Experimental setup for the SEMCBR validation test showing the source-array geometry (left) and the 4x4 Cartesian array (above), used to emulate a full 8x8 array.

The source locations were estimated using both SEM and a conventional delay-and-sum beamformer (CBF) for comparison. In both cases, a spherical wave Green's function, $G_{km}(\omega) = e^{i\frac{\omega}{c}R_{km}}/R_{km}$ was used to model the nominally monopole, non-farfield sources (in actuality, the Fraunhofer number was $L^2/2\lambda R \sim 1.1$, given array length L and wavelength λ). The outputs of SEM are the estimated source strengths as a function of x and y (Equation (4.14)), while the CBF output is given by [44]

$$\mathrm{BF}_{\mathrm{conv}}(x,y) = a^H(x,y)\,\Gamma_{mn}^{meas}(\omega)a(x,y) \tag{4.19}$$

where $a(x,y) = [G_{1k}\ G_{2k}\ \ldots G_{Mk}]^T$, k serves as an indexing variable for the gridded (x,y) coordinates which must be searched over, and $(\cdot)^H$ represents the conjugate transpose.

Figure 4.11 shows the localization results of the two techniques as a function of number of subarrays averaged. While CBF can reasonably distinguish the sources when the full 8x8 array is used, its resolution degrades as the array is divided into subarrays, due primarily to the aperture reduction. On the other hand, SEM performs only moderately well when the full 8x8 array is used, and performs best when 25 4x4 subarrays are averaged, localizing all three point sources precisely. This improvement in accuracy supports the notion that the subarray averaging procedure reduces the effects of the source covariance cross-terms that otherwise degrade the performance of SEM.

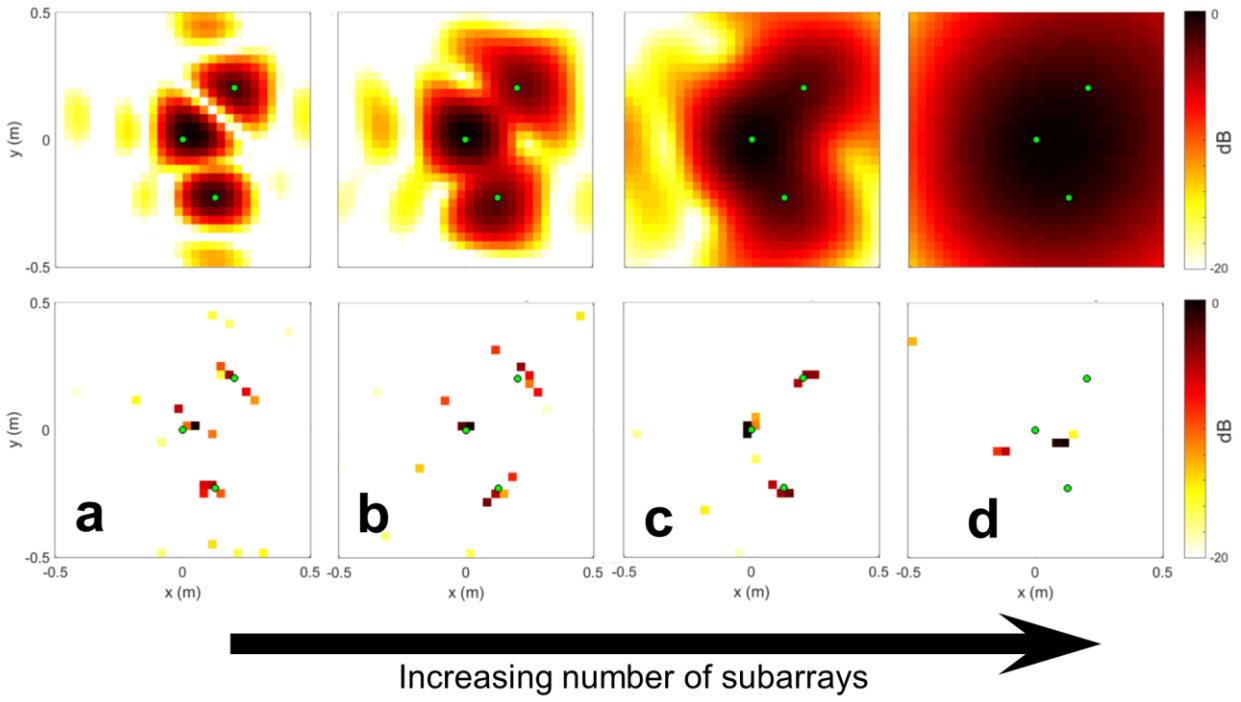

Figure 4.11: Localization of three coherent sources ($\boldsymbol{kd} = 5.5$) with conventional beamforming (left) and SEM (right) using a) 1 8x8 subarray (i.e. no averaging), b) 9 6x6 subarrays, c) 25 4x4 subarrays, and d) 49 2x2 subarrays.

Having shown the benefit of subarray averaging for highly resolved localization of coherent sources, attention was turned to the coherent background subtraction problem; namely, extending SEM to SEMCBR. To test SEMCBR, another set of acoustic measurement were collected, nominally identical to the previously detailed three-source setup, however one of the sources was decreased in amplitude by 20% (the other two sources remaining unchanged). Both localization methods were repeated, this time with the addition of the CSDM subtraction step of Equation (4.15) where Γ^{meas} serves as the CSDM of the original data, and Γ^{ref} as the CSDM of the data with a reduced source. 25 subarrays of 4x4 receivers were averaged to generate the CSDMs for the results shown in Figure 4.12. SEMCBR successfully localized the small reduction in amplitude of the changed source with high precision (Figure 4.12c). Though

several spurious points were resolved from directions other than the changed source, these peaks were more than 8 dB below the maximum. Conventional beamforming was also generally successful in estimating the direction of the changed source; however, it suffers from a much larger lobe width in addition to a broad sidelobe only 3 dB below the main lobe. Further, both results agree with simulation using identical parameters, shown in Figure 4.12d, further validating that SEMCBR is capable of localizing a small *change* in a loud, coherent sound field, with high precision.

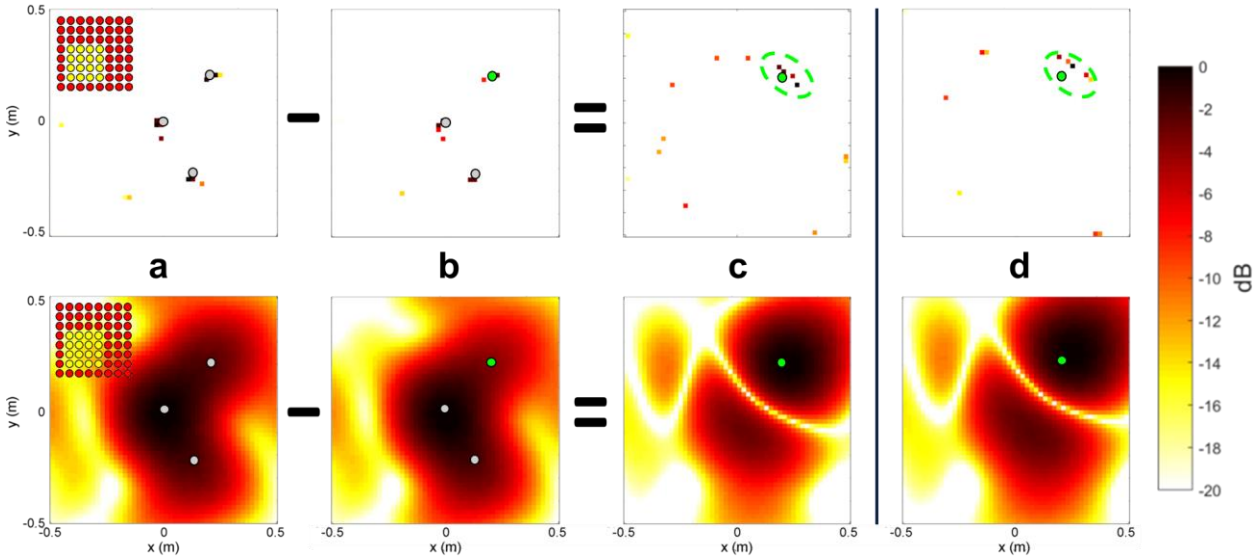

Figure 4.12: Localization of a small change in a sound field of three coherent sources using SEMCBR (top row) and conventional beamforming with a reference subtraction (bottom row). Three sources with equal amplitudes were recorded (a) then rerecorded, with one of the sources reduced in amplitude by 20% (indicated by a green dot in (b), (c), (d)). Experimental results for localization of the *changed* source are shown in (c) and agree well with simulation (d). For each figure, averaging of 25 4x4 subarrays across the full 8x8 array was used.

4.3.4 Localizing damage in vibrating plates

Having outlined and established the SEMCBR method, several experiments were conducted to apply the method toward the structural health monitoring of vibrating plates. An experimental apparatus, shown at the left of Figure 4.13, was used to clamp 30 cm x 30 cm plates to a nominally-rigid aluminum mounting base. The plates were then base-excited with frequency-swept input forcing from 100 to 6000 Hz, generated by a K2007E01 series electrodynamic shaker (The Modal Shop, Cincinnati, OH) attached to the plate's mounting base (see Figure 4.13, bottom right). The acoustic response of the plates was recorded with the same array configuration detailed in Section 4.3.3 ($kd = 3.2$ for the 3 kHz center frequency), using the same PXIe-6368 DAQ sampling at 50 kHz. The distance from plate center to array center was 1 m. The input forcing bandwidth was chosen to encompass approximately the first dozen unique vibrational modes of the test plates. In all following results, size 4x4 subarrays were averaged over the full 8x8 array to mitigate coherent cross terms.

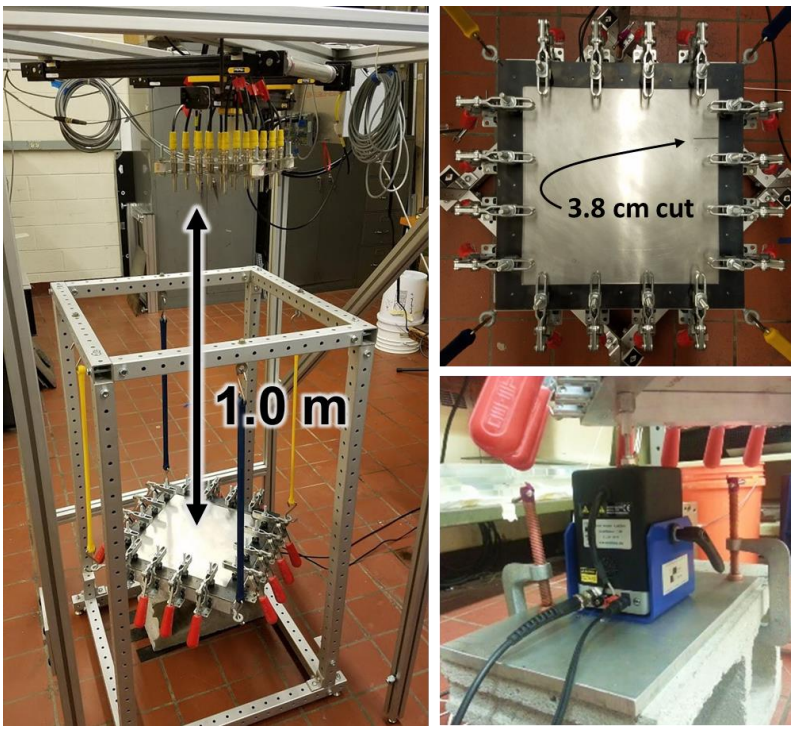

Figure 4.13: Experimental apparatus for remote acoustic evaluation of vibrating 30 cm x 30 cm plates. The array plate configuration (left) is used to measure the acoustic response of damaged plates (top-right) that are base-excited with a 30 N electrodynamic shaker (bottom-right).

4.3.4.a Cuts in an aluminum plate

The first damage case consists of through cuts in a 30 x 30 x 0.16 cm aluminum plate. A 3.8 cm cut was made perpendicular to the edge of the plate with a Dremel tool. The presence of the cut affects the plate's vibroacoustic response, both via changes in mode shape and through (downward) shifts in resonant frequencies. Figure 4.14 showcases laser Doppler vibrometer measurements of the effect of the edge cut on the vibrational response of the plate at its (4,1) mode. Though the spatial structure of the mode is similar along the majority of the plate's surface, a localized change in vibrational behavior is observed near the cut, and similar

behavior is present in other modes. It is these vibrational – and subsequent acoustical – changes near the damage that present an opportunity for remote acoustic localization.

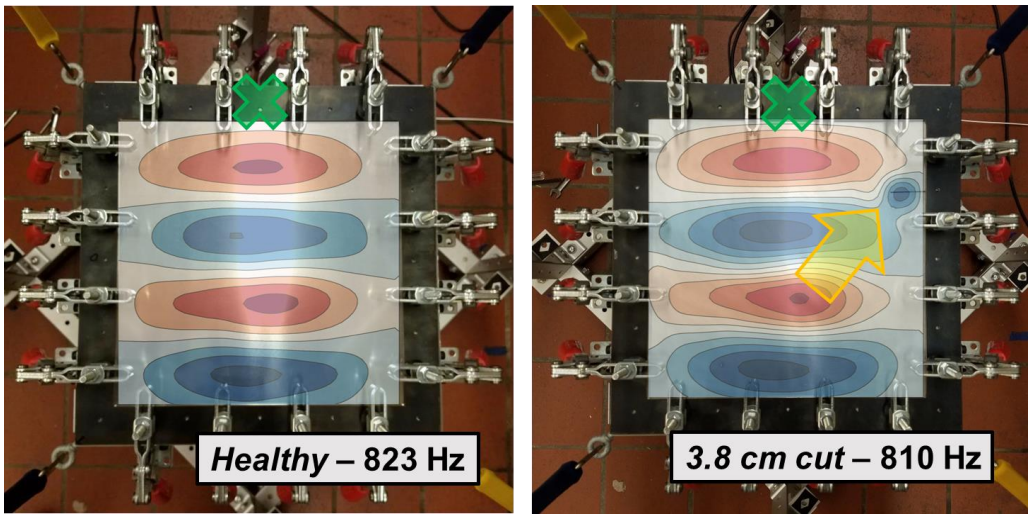

Figure 4.14: Laser Doppler vibrometer measurements of (4,1) vibrational modes of a 30 x 30 x 0.16 cm aluminum plate undamaged (left) and with a 3.8 cm cut (right, indicated by the yellow arrow). In addition to modal changes, the damage also shifts the resonant frequency down by 1.5%.

SEMCBR was performed for plate data using the same procedure discussed in Section 4.3.3. Referring to Equation (4.15), the cut-plate data were used to construct Γ^{meas}, while the healthy-plate data were used for Γ^{ref}. SEMCBR and conventional beamforming, as stated, are applied to a single frequency bin of the CSDMs. This yields a single localization estimate which can be trivially extended to multiple frequencies throughout the bandwidth of the recorded signals. Averaging localization estimates over multiple frequencies serves to improve performance since more information (in this case, more modes) are being incorporated. However, computing all available frequency bins quickly becomes computationally intensive for conventional beamforming and especially SEMCBR, which requires a convex optimization. Instead, frequencies were selected by first identifying the resonances of the array data, then

selecting frequencies distributed around these high-energy regions. This helps to avoid processing data in modally 'inactive' frequency bands that offer little information. Figure 4.15 shows the frequency power spectra for all 64 receivers in for both healthy and damaged plates. Average spectra for each plate are also plotted and comparison between the two curves shows moderate-but-noticeable downshifting of the resonances due to the 3.8 cm cut. A peak finding algorithm was used to identify 11 resonance peaks within the bandwidth, and for each peak 11 frequency bins were selected with uniform 5 Hz spacing yielding a total of 121 frequency bins to process.

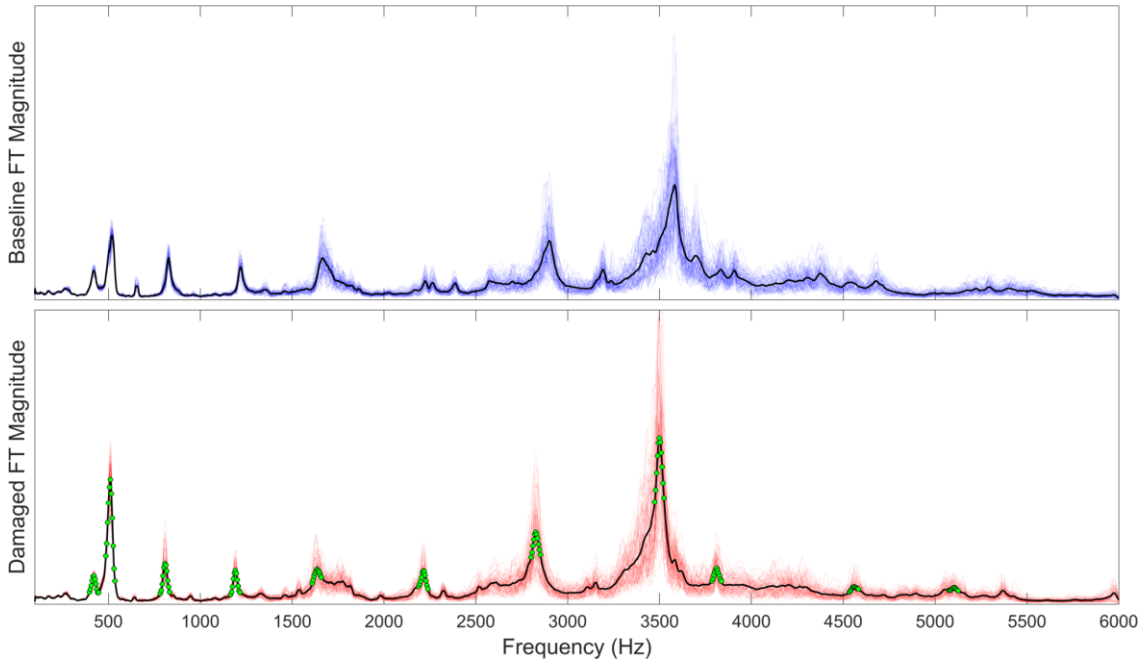

Figure 4.15: Acoustic power spectra for the baseline (top) and 3.8 cm edge-cut plate (bottom). Black curves signify the average spectra across all 64 receivers. The presence of the cut results in downward shifts of the resonant frequencies. Green dots, distributed about the peaks, indicate the frequency bins at which SEMCBR and conventional beamforming were computed.

The bandwidth-averaged damage localization results for a 3.8 cm edge cut are shown in Figure 4.16. Though both conventional beamforming and SEMCBR are distributed around the

cut, SEMCBR features a much tighter grouping than conventional beamforming. The resolution of each method was quantified by comparing the 1-standard deviation spread of each method, about the centroid. For the conventional beamformer output, this corresponds to $20\log\left(e^{-\frac{1}{2}}\right) = -4.3$ dB and for SEMCBR, the weighted RMS spread of the source level map was used. The spread contours are depicted as dashed blue lines in Figure 4.16. For the 3.8 cm cut, the spread of the conventional beamforming method was found to be approximately 20 cm, or roughly 2/3 the plate length. SEMCBR, however, exhibits a spread of approximately 3 cm, approximately equal to the actual cut length. In other words, SEMCBR achieves a resolution over 6 times greater than that of conventional methods, though both methods do point in the correct location of the cut.

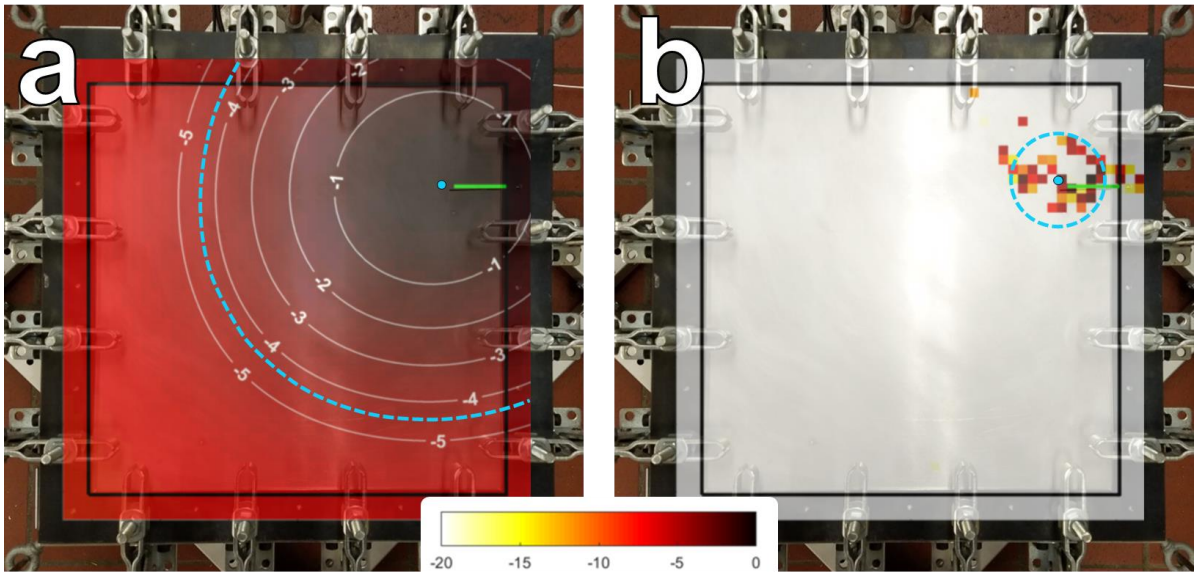

Figure 4.16: Remote acoustic localization of a 3.8 cm through cut near the plate edge using background subtraction with a conventional beamformer (a) and SEMCBR (b). Green lines indicate the true cut location and dashed-blue lines indicate the standard deviation spread of each method.

An additional cut test was conducted for a cut made interior to the plate. In this case a 3.8 cm horizontal cut was made in the upper left quadrant of the plate. The same procedure was used, including frequency bandwidth averaging over 121 frequency bins. Conventional beamforming and SEMCBR results are compared in Figure 4.17. Again, while both SEMCBR and conventional beamforming are centered about the cut, the SEMCBR estimates of the *changed* source energy are much more tightly distributed (by a factor of ~6) around the 3.8 cm cut.

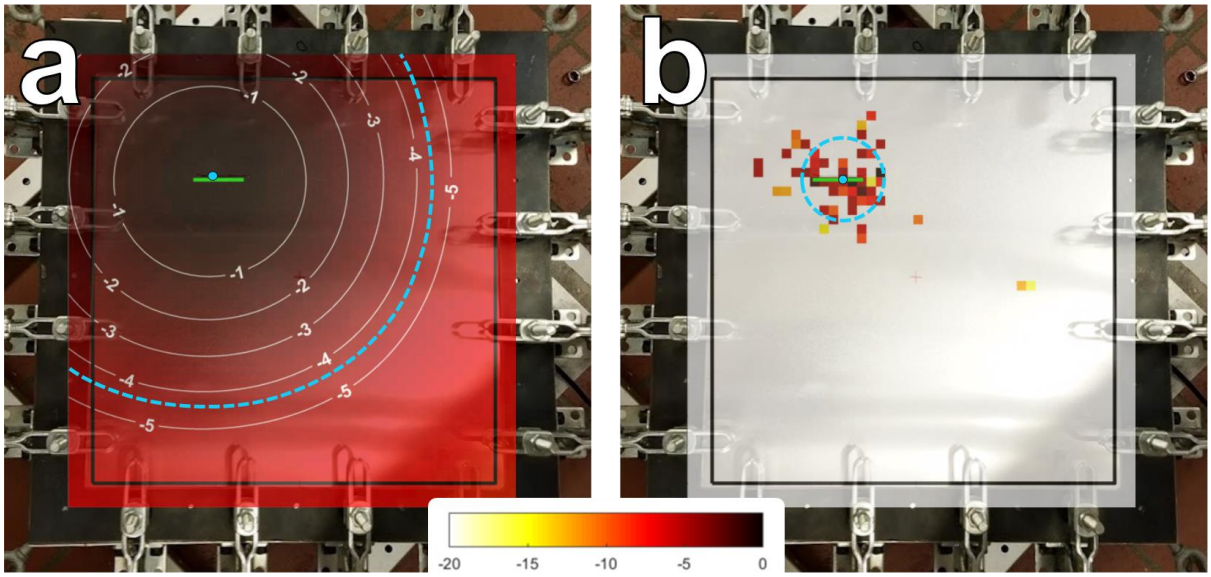

Figure 4.17: Remote acoustic localization of a 3.8 cm through cut, center to the plate using background subtraction with a conventional beamformer (a) and SEMCBR (b). Green lines indicate the true cut location and dashed-blue lines indicate the standard deviation spread of each method.

4.3.4.b Changes in boundary condition of an aluminum plate

Another common form of damage in many plate-like structures is weld or fastener faults [96]. Either the failure of a weld or the loss of a fastener can result in significant changes to the boundary conditions of a structure, thereby altering its vibroacoustic response [97]. To approximate this effect experimentally, a single clamp was released from the nominally

clamped 30 x 30 x 0.16 cm aluminum plate. A baseline recording of the fully clamped plate was collected, then a second recording was collected after releasing a clamp located near the center of one side of the plate. Using the same procedure as in Section 4.3.4a, acoustic localization of the unclamped region was performed, with a comparison of conventional beamforming and SEMCBR shown in Figure 4.18. As with the cut cases, both methods pointed in the direction of the released clamp. Conventional beamforming performed roughly equivalent as with the cuts, however SEMCBR featured a significantly less tight set of source estimates (though still a factor of 5 smaller spread than conventional). A proposed explanation for this result is that the change in the boundary conditions results in a more distributed change to the vibrational mode shapes than for the cuts. In this case, SEMCBR is still localizing the change in the acoustic source (which are associated with changes in the vibrational response), though the extent of these changes is less localized than the unclamped region.

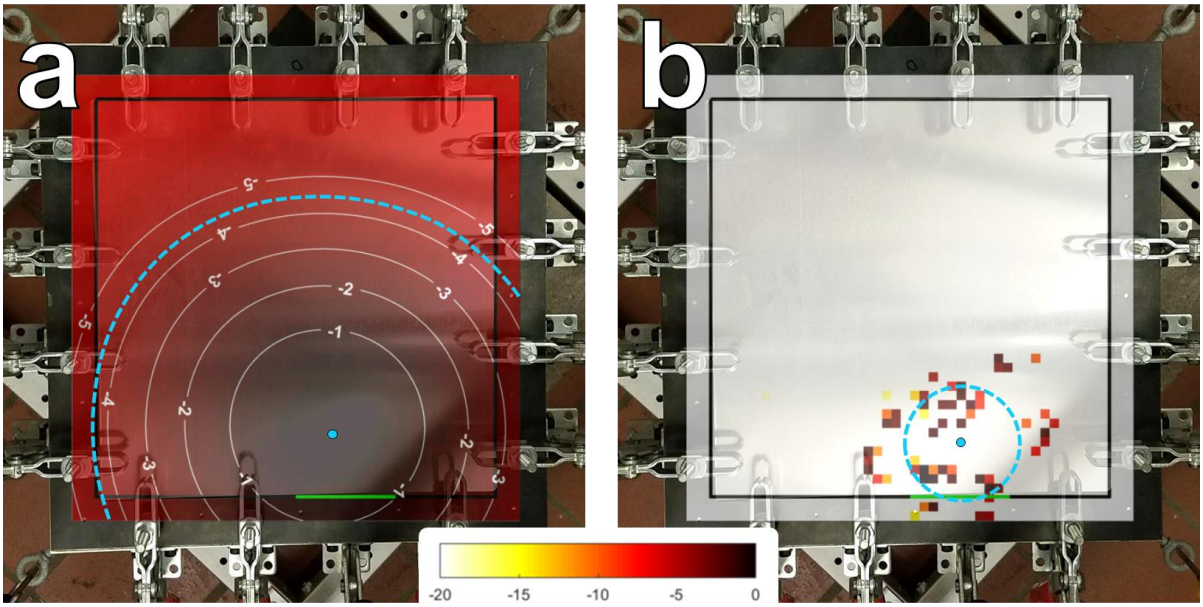

Figure 4.18: Remote acoustic localization of a partially unclamped plate using background subtraction with a conventional beamformer (a) and SEMCBR (b). Green lines indicate the nominally unclamped section of the perimeter and dashed-blue lines indicate the standard deviation spread of each method.

4.3.4.c Changes in boundary condition of an aluminum plate

A final test was conducted to determine the localizability of delamination in a composite plate. Composite materials, and in particular composite laminates, have become increasingly popular in both industrial and consumer applications over the last several decades due to their high strength-to-weight ratios [98]. Composite laminates are a subset of composite materials that consist of multiple layers of fiber laminae (often woven), encased in a resin matrix. Lamina material and orientation, weave patterns, and numbers of layers all serve as design parameter to allow specific structural constraints such as yield strength and rigidity to be achieved with modest mass. However, when the bonding matrix between layers detaches, an effect commonly titled *delamination*, a significant reduction in material strength may occur [99]. Delamination, being a subsurface effect, is also usually non-visual, meaning that effective detection and localization strategies are desirable to treat delamination damage before it can grow.

Two carbon fiber composite plates were constructed using a conventional vacuum bagging process [100]. Each plate consisted of 6-layers of bi-axial weave carbon fiber, encased in resin, with the principle fiber direction running parallel to the square axes. The number of layers was chosen to approximately match the rigidity of aluminum to achieve a similar modal density in the excitation bandwidth. In the damaged plate, thin 7.5 x 7.5 cm layers of aluminum foil were carefully placed between the medial layers during the layup process, as to prevent bonding within a localized region, thereby approximating partial delamination. After curing and trimming, both composite plates were nominally 0.3 x 0.3 x 0.15 cm in size.

Both the baseline and partially delaminated plates were tested akin to the previous sections Results are provided in Figure 4.19; as was seen with the cut and unclamping test, both methods successfully localize around the delaminated region, however SEMCBR yielded a

spread of approximately 12 cm, in comparison to the conventional beamformer spread of 23 cm. These results are promising because, as is seen in Figure 4.19a, there is no apparent visual indication of the delaminated region.

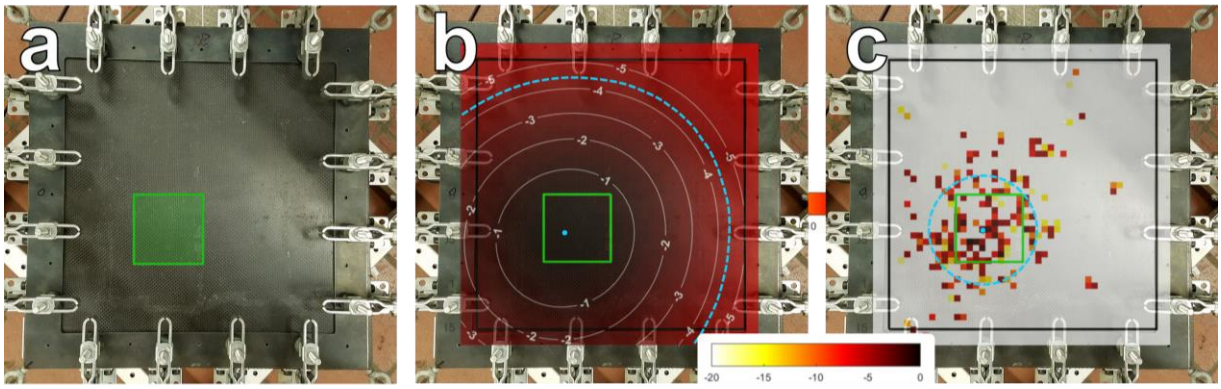

Figure 4.19: Remote acoustic localization of a partially delaminated carbon fiber plate (a) using background subtraction with a conventional beamformer (b) and SEMCBR (c). Green squares indicate the 7.5 x 7.5 cm delaminated region and dashed-blue lines indicate the spread of each method.

4.3.5 Vibrational dependence of localization

It must be noted that, though the goal of SEMCBR in the previous section was to localize the damage site, in actuality it is only the *vibrational change* due to the damage that is being localized. In theory, the vibrational change in a structure is localized at or near the damage site, and so they are effectively equivalent. In particular, the localization of fastener failure (boundary damage) shown in Figure 4.18 did not point at the edge, but rather slightly inward toward the plate center. This is sensible because the boundary, where the change is located, is not vibrationally active, and thereby casts no sound. However, the region of the plate near the unclamped edge *is* vibrationally active, as well as changed relative to baseline, explaining the achieved results.

This is illustrated in Figure 4.20 using simulated plate response. The baseline plate response at 450 Hz (shown at center) exhibits the expected 2x2 mode, while a 3-inch-cut plate (left) vibrates with a similar, but slightly difference mode shape. The difference of the two mode shapes yields a "residual" that is maximized near the cut, but is also non-zero at other regions of the plate. However, by averaging this vibrational residual over the 0-2000 Hz frequency band, the average difference is clearly localized to the cut, shown at the left of Figure 4.21. A similar analysis was performed for a boundary-type defect, shown at the right if Figure 4.21. As expected, boundary damage results in vibrational changes localized not *at* the damage, but rather *near* to the damage (since the boundary itself does not vibrate).

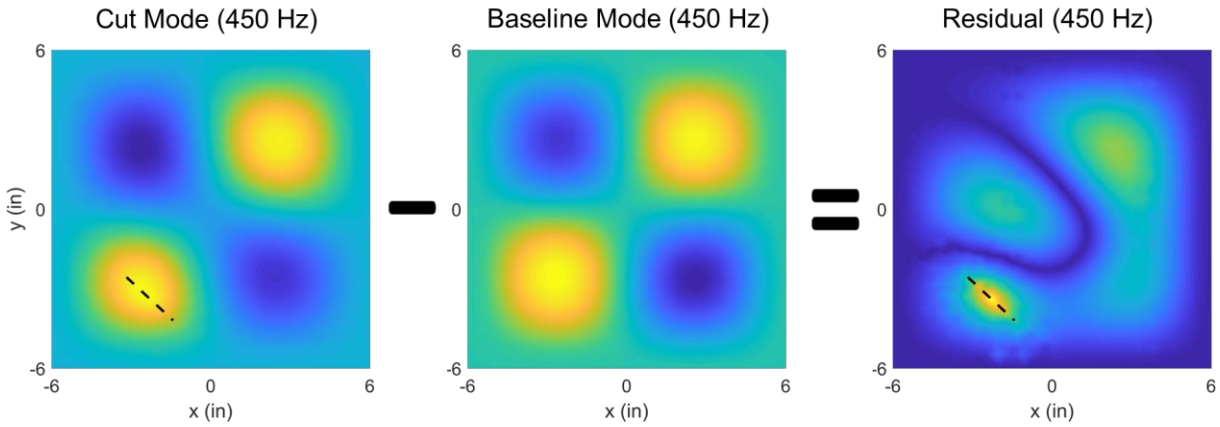

Figure 4.20: Subtraction of the baseline 2x2 vibrational mode from the cut mode yields a residual local to the cut, though other regions of the plate are also changed.

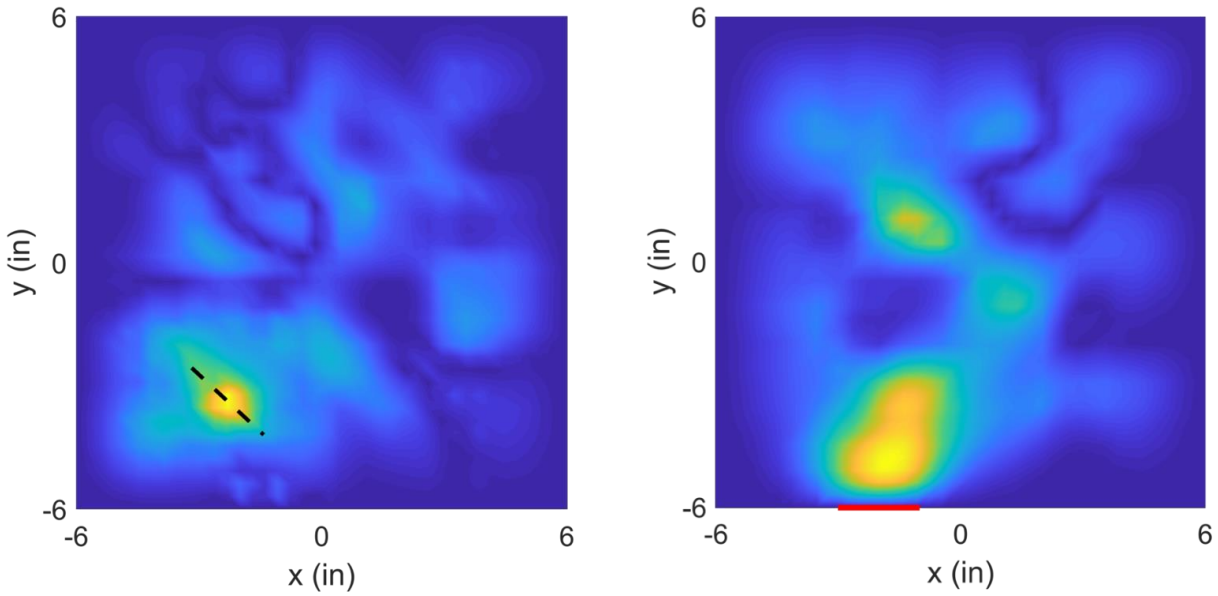

Figure 4.21: Averaging the vibrational residuals over all modes from 0-2000 Hz, yields an average vibrational coincident with the damage in the case of a cut (left), and slightly interior to the actual location of a fastener failure (right).

From Figure 4.21 it can be seen that the simulated average *vibrational* change agrees well with the measured *acoustic* localization results shown in Figure 4.17 and Figure 4.18. This indicates that SEMCBR does indeed localize the vibrational change rather than the damage itself. For more complex structures it may be possible that the vibrational changes due to the damage are *not* localized to the damage itself, and in this scenario SEMCBR may be inadequate for damage localization. However, in the case of the square, vibrating plate, the average vibrational change was consistently coincident (or nearly coincident) with the each of the damage type investigate.

4.3.6 Summary and conclusions

The Spectral Estimation Method with Coherent Background Removal (SEMCBR) was presented for the high-resolution acoustic localization small changes in a coherent field. SEMCBR is an extension of the existing the Spectral Estimation Method with Additive Noise [101], by incorporating a subarray averaging step before the CSDM subtraction. This allows for a small change in an otherwise loud and complex coherent acoustic field to be isolated and localized, provided a background measurement of the field, without the change, is available. SEMCBR was validated first using a planar Cartesian 8x8 array to successfully localize a 20% decrement in the amplitude one of three coherent sources, using a 4x4 subarray averaging step. Comparisons to conventional beamforming (with CSDM subtraction) showed a significant improvement in source resolution for SEMCBR.

SEMCBR was then employed for the localization of damage in 30 x 30 cm base-excited plates, again using a 4x4 subarray averaging step. Using SEMCBR and conventional beamforming (with CSDM subtraction), through-cuts and partial unclamping of aluminum plates were successfully localized. Delamination of a carbon-fiber composite plate was also achieved. In all cases, conventional beamforming was generally successful in pointing toward the correct damage sites, however the beams were not well resolved, with spreads on the order of the plate length for all tests. Comparatively, SEMCBR was found to achieve roughly a factor of 6 increase in damage resolution, yielding a better estimate of the spatial extent of the damage. In addition to its applicability to for structural health monitoring, SEMCBR could also provide a useful tool to the underwater acoustics community, where multipath is a common concern, or in radar applications to counteract jammers and coherent interferers.

4.4 Derivations for SEMCBR

4.4.1 Effect of subarray averaging on source cross-spectra terms

In the case of a uniform line array with receiver spacing d, divided into P subarrays, the m^{th} receiver output of the p^{th} subarray (where, p represents the subarray increment from the centermost subarray) is

$$\hat{r}_m^p(\omega) = \sum_{k=1}^{K} e^{i\frac{\omega}{c}d(m+p)\sin\theta_k}\hat{s}_k(\omega) \tag{4.20}$$

where the far-field assumption has been made for the K sources arriving from incident angles θ_k and c is sound speed. The resulting averaged CSDM is then

$$\bar{\Gamma}_{mn} \equiv \frac{1}{P}\sum_{p=1}^{P} \hat{r}_m^p \hat{r}_n^{p*} = \frac{1}{P}\sum_{p=-\frac{P-1}{2}}^{\frac{P-1}{2}}\sum_{k,l=1}^{K} e^{i\frac{\omega}{c}d\left((m+p-1)\sin\theta_k-(n+p)\sin\theta_l\right)}\hat{s}_k\hat{s}_l^* \tag{4.21}$$

$$\bar{\Gamma}_{mn} = \sum_{k,l=1}^{K} e^{i\frac{\omega}{c}d(m\sin\theta_k-n\sin\theta_l)}\hat{s}_k\hat{s}_l^* \frac{1}{P}\sum_{p=-\frac{P-1}{2}}^{\frac{P-1}{2}} e^{i\frac{\omega}{c}dp(\sin\theta_k-\sin\theta_l)} \tag{4.22}$$

The summation term over the P subarrays is analogous to that derived for the beampattern of a finite uniform line array and may be shown to be [102]

$$\eta_{kl} \equiv \frac{1}{P}\sum_{p=-\frac{P-1}{2}}^{\frac{P-1}{2}} e^{i\frac{\omega}{c}d(p-1)(\sin\theta_k-\sin\theta_l)} = \frac{\sin\left(\frac{\omega d}{2c}P(\sin\theta_k-\sin\theta_l)\right)}{\sin\left(\frac{\omega d}{2c}(\sin\theta_k-\sin\theta_l)\right)} \tag{4.23}$$

4.4.2 Subarray averaging correction phase for non-farfield sources

When subarray smoothing is used to localize *non*-farfield sources, a phase correction must be applied to each subarray CSDM. These phases counteract the effect of shifting the center location of each subarray relative to the non-planar incident sources (i.e. the global position of each subarray must be considered when localizing proximal sources). Consider the case shown in Figure 4.22, wherein a spherical-wave source impinges on an M-element ULA.

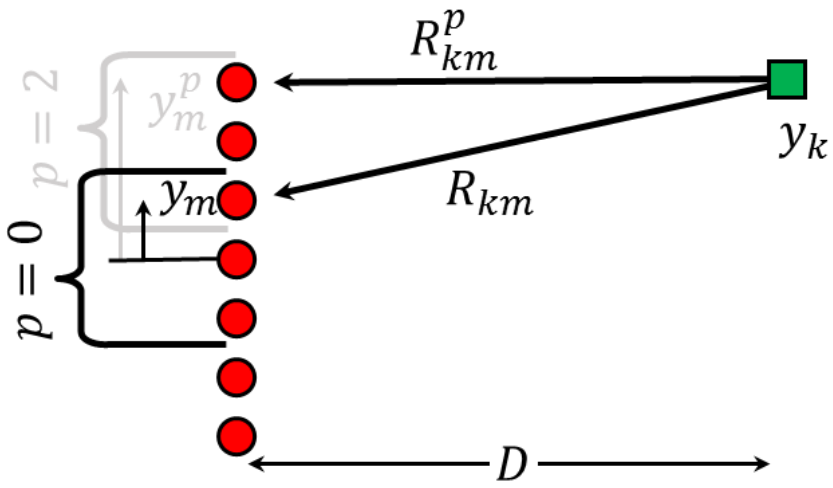

Figure 4.22: Schematic for a spherical wave source impinging on a uniform line array. The distance of the source as seen by the $\boldsymbol{m^{th}}$ element of the central ($\boldsymbol{p = 0}$) subarray is different than the distance seen by the $\boldsymbol{m^{th}}$ element of the $\boldsymbol{p = 2}$ subarray, resulting in a phase difference that must be accounted for.

The distance R_{mk} from the source at y_k to the m^{th} receiver of the center-most subarray (centered at the origin of the source plane) is given by $R_{mk} = \sqrt{(y_m - y_k)^2 + D^2}$. Setting R_{mk}^{p} as the distance from the source to the m^{th} receiver of the p^{th} subarray, and noting, for a ULA

with receiver spacing d, $y_m^p = y_m + (p-1)d$, the difference in source distance as seen between the two subarrays is

$$\begin{aligned} R_{mk}^p - R_{mk} &\equiv \Delta R_{mk}^p \\ &= \sqrt{(y_m + (p-1)d - y_k)^2 + D^2} - \sqrt{(y_m - y_k)^2 + D^2} \end{aligned} \tag{4.24}$$

Making the assumption $D \gg |y_m - y_k|$ (which is less restrictive than assuming the sources are properly in the farfield) we can truncate the Taylor expansion for the roots yielding

$$\Delta R_{mk}^p \approx \frac{1}{2D}(2(p-1)dy_m + (p-1)^2 d^2 - 2y_k(p-1)d) \tag{4.25}$$

at first glance it appears any correction involving this term is stymied by the unknown source location(s) y_k. However, the modeled CSDM follows the form $\Gamma_{mn} = \hat{G}_{m,i}\hat{s}_i\hat{s}_j^*\hat{G}_{n,j}^*$ and, for spherical wave sources this yields a phasor of the form $\exp(i\omega/c(R_{mk} - R_{nk}))$. Considering the p^{th} subarray, this phasor can be expanded into

$$e^{i\frac{\omega}{c}(R_{mk}^p - R_{nk}^p)} = e^{i\frac{\omega}{c}(R_{mk} - \Delta R_{mk}^p - (R_n - \Delta R_{nk}^p))} = e^{i\frac{\omega}{c}(R_{mk} - R_{nk})} e^{-i\frac{\omega}{c}(\Delta R_{mk}^p - \Delta R_{nk}^p)} \tag{4.26}$$

where the $\exp(-i\omega/c(\Delta R_{mk}^p - \Delta R_{nk}^p))$ is the 'shifting' we aim to correct for, for each subarray. Fortunately we find

$$\Delta R_{mk}^p - \Delta R_{nk}^p \approx \frac{d}{D}(p-1)(y_m - y_n) \tag{4.27}$$

which is independent of the unknown source location(s). Therefore, in computing the subarray averaged CSDM for non-farfield sources, a correction phase of

$$\Gamma_{mn}^{p} \rightarrow \Gamma_{mn}^{p} e^{i\frac{\omega d}{c D}(p-1)(y_m - y_n)} \tag{4.28}$$

must be applied to the elements of each subarray CSDM. This formulation is easily extended to 2D Cartesian arrays,

$$\Gamma_{mn}^{p} \rightarrow \Gamma_{mn}^{p} e^{i\frac{\omega}{cD}(d_x(p_x-1)(x_m - x_n) + d_y(p_y-1)(y_m - y_n))} \tag{4.29}$$

where d_x , d_y correspond to the x- and y-direction receiver spacings, and p_x , p_y correspond to the subarray increments of the p^{th} subarray along the x- and y-directions, respectively. It is also worth noting while the derivation above was restricted to $D \gg |y_m - y_k|$ (essentially requiring sources be significantly more distant from the array than they are spread parallel to the array), stable localization results were achieved even for $D \sim |y_m - y_k|$.

Chapter V: Classification of Damage

5.1 Introduction

Vibrating structures are common in numerous applications and when such a structure is in contact with a fluid (e.g. air or water) these vibrations commonly result in radiated sound [10]. As discussed in previous chapters, knowledge of the vibroacoustic behavior of a structure is useful for remotely interfacing with it *in situ*, and for performing diagnostics (e.g. *how has the acoustic signature changed and what does that tell us?*). For the purposes of structural health monitoring, these vibroacoustic changes may be effective indicators of structural damage [4, 7], which, having been identified, can be accommodated or repaired. This chapter investigates the feasibility of using data-driven approaches for *classifying* multiple types of damage in vibrating plates using vibrational and/or remote acoustic measurements. More specifically, many finite element simulations of damaged plates were used to train classifiers and subsequently distinguish damage types based solely on changes to acoustic (and/or displacement) frequency response functions, with 'changes' quantified relative to an *a priori* known healthy baseline structural response. With this data-driven approach, classification between modeled cut, corrosion, and fastener failure (i.e. weld) -type damage to square plates was performed, yielding correct classification rates in excess of 96% using modeled acoustic responses. Experimental data support the simulated findings.

5.2 Motivation for data-driven approaches

The process of structural health monitoring (SHM) for damage inspection is typically divided into four steps: *Detection* (is damage present?), *Localization* (if so, where?), *Classification* (what type of damage, and how severe?), and *Evaluation/Treatment* (what is the prognosis or next step(s)?) [4]. Chapters 3 and 4 detailed several methods for the first two of these steps – detection and localization – in regard to cuts, fastener failures, and delamination in vibrating plates using remote acoustic sensing methods. In this chapter, the third step of SHM, *damage classification* is tackled, again with an emphasis on remote acoustic methods.

Damage classification is often a challenging task due to the many ways in which a system can become damaged. Because damage typically manifests itself in complex changes to the system's vibroacoustic behavior, the exact form in which any particular type of damage will affect structural response is difficult to predict. Though in theory a given type of damage will yield a deterministic change to some output response (e.g. a frequency response curve), in practice, solving such an inverse problem is infeasible because the location and severity of the damage are typically unknowns, and the accuracy of any structural vibration model will be limited. Even when the localization techniques described in the previous chapter are applied, the search space of possible damage configurations may be unmanageably large and the needed model accuracy may be out of reach. Furthermore, applying analytical models becomes even more challenging when the *type* of damage is *a priori* unknown and when structural models become more complex, and potentially less accurate.

In lieu of appropriate analytical or computational models, the damage classification task is often approached using data-driven methods [30]. Data-driven methods are appealing because the difficulties of solving potentially-intractable inverse problems is replaced with the

more achievable task of generating training data for a ‘damaged’ system [5, 103]. With training data in hand, any of a number of common classification approaches can be implemented, with the number of assignable damage classes (or types) dependent on the number of damage classes included in the training data set.

5.3 Classification methodology

The canonical clamped, square vibrating plate was selected as the system of interest. This simple structure was chosen 1) for ease of modeling, 2) to avoid unnecessary complexities associated with more complicated structures while maintaining some level of practical application, and 3) because an experimental setup for monitoring vibrating plates was previously developed (see Section 2.2 Experimental apparatus), allowing for empirical validation of the simulation results. In this investigation, three distinct damage classes were considered for the vibrating plate: cuts, corrosion, and fastener failures (e.g. weld/bolt failures).

The data-driven damage classification method discussed in this report consists of the following four steps:

1) **Generate Data** – ‘damaged’ plate datasets were generated using Monte Carlo finite element simulations for the three damage cases. For each damage case, 200-250 simulations were run, varying the location, severity (size), and type of damage present. To model acoustic responses, the plate was treated as existing in an infinite baffle and the Rayleigh integral [40] was evaluated at user-selected receiver locations. See Section 2.3 Finite element analysis for more information about the modeling process.

2) **Feature Extraction** – comparisons between the damaged-plate frequency responses and the baseline-plate frequency response were used as features (i.e. inputs fed into a classifier which are ultimately used to distinguish between different damage classes) for classification. Specifically, a dilational correlation [104] was applied to each damage case, quantifying how 'deformed' and 'stretched' each test-case frequency response function is compared the baseline case.
3) **Training Classifier** – 70% of the generated data was used to train a classification scheme based on the extracted features. Several classification schemes were implemented, included K Nearest Neighbors (KNN) [103]., Support Vector Machines (SVM) [105], Neural Nets (NN) [106], and others [107, 108].
4) **Testing Classifier** – The remaining 30% of data was used to test the performance of the trained classifier, quantified using the Correct Classification Rate (CCR), defined as the percent of test data samples that were successfully classified to their true damage type. To obtain an average value for the CCR (and to avoid lucky/unlucky partitionings of training/test data), steps 3 and 4 were performed 50 times for each classifier, with random permutations of training/test data selected each time. The average CCR of these 50 partitionings is reported.

Figure 5.1 shows the finite element mesh for the clamped, square aluminum plate, which is driven at an arbitrary, but consistent, location throughout all simulations. Also included is a frequency response curve for the RMS displacement of the plate from 0 to 2000 Hz. Ultimately, it is the changes to this frequency response curve that were quantified and subsequently utilized

for damage classification (or in other words, each type of damage affects the plate's frequency response in complex – but distinguishable – ways).

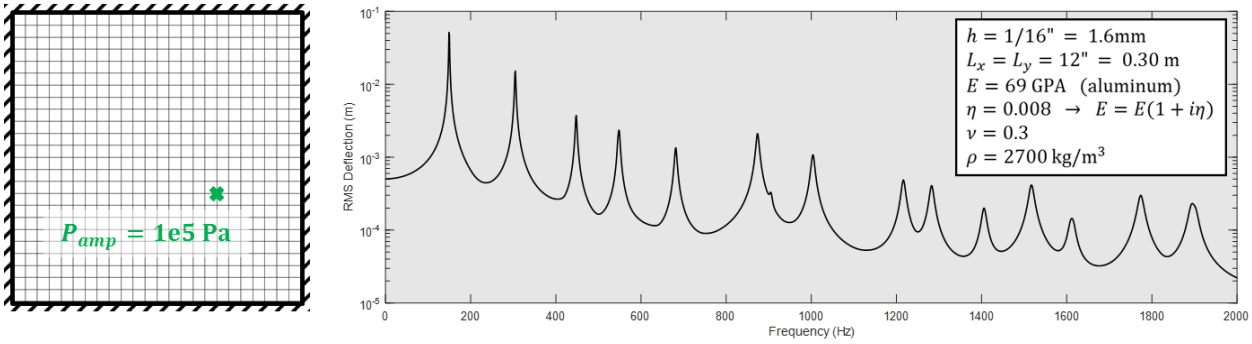

Figure 5.1: FE model of a clamped, square aluminum plate showing: the forcing drive point (left), and the resulting frequency response of RMS displacement for the baseline (undamaged) plate (right).

Three types, or classes, of damage were considered: cuts, corrosion, and boundary failures. The first of these, cuts, was modeled by assuming straight through-cuts wherein the length, center, and angle of the cuts were randomly varied throughout 250 training cases. 50 cases each were computed for cuts of 8%, 12%, 16%, 20%, and 24% of the plate side length (12 in). Figure 5.2 gives an example of a vibrational mode shape in the presence of a large through-cut, alongside 50 RMS displacement frequency response for the smallest and largest cut test cases. It is notable that changes to the RMS frequency response are both more extreme and more complex at higher frequencies (especially for frequencies above 1 kHz). Figure 5.3 shows the averaged frequency response curves for the 50 simulations for each of the five difference cut lengths. It is notable that increased damage severity (in this case, increased cut length) results not only in greater frequency response variation relative to baseline, but also greater intra-class variation (i.e. variation between all simulation of a 3" cut).

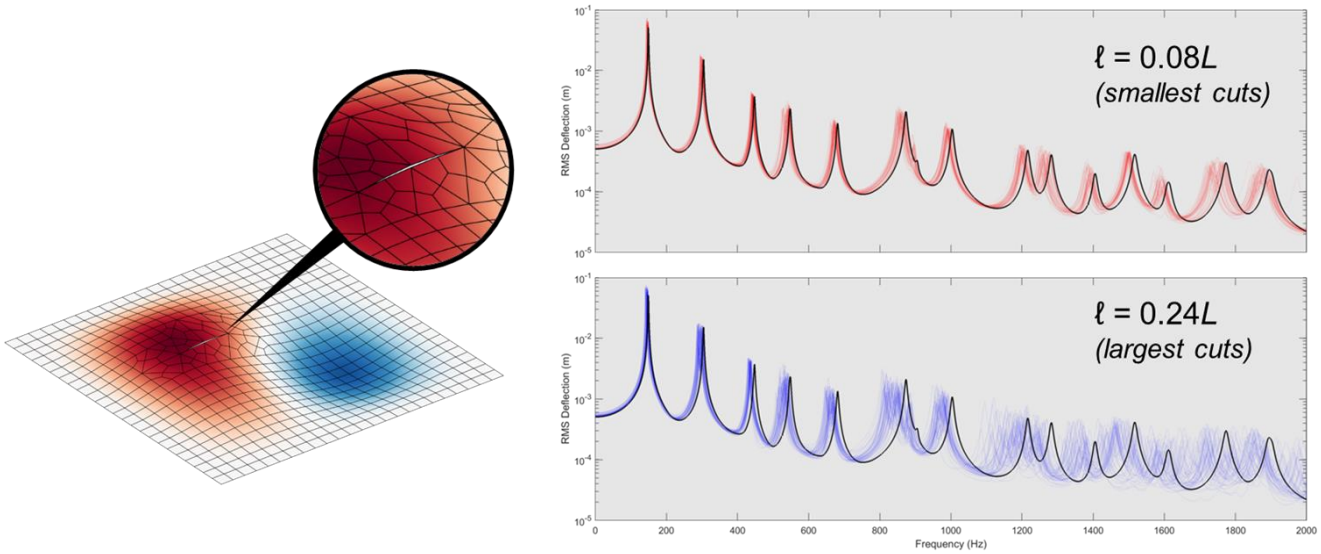

Figure 5.2: Example of a steady state mode shape (1000 Hz) in the presence of a 2 inch cut (left). Frequency responses of RMS displacement for all 50 cases of the smallest cut size (1 inch, red) and the largest cut size (3 inches, blue) (right). Baseline frequency responses are shown in black.

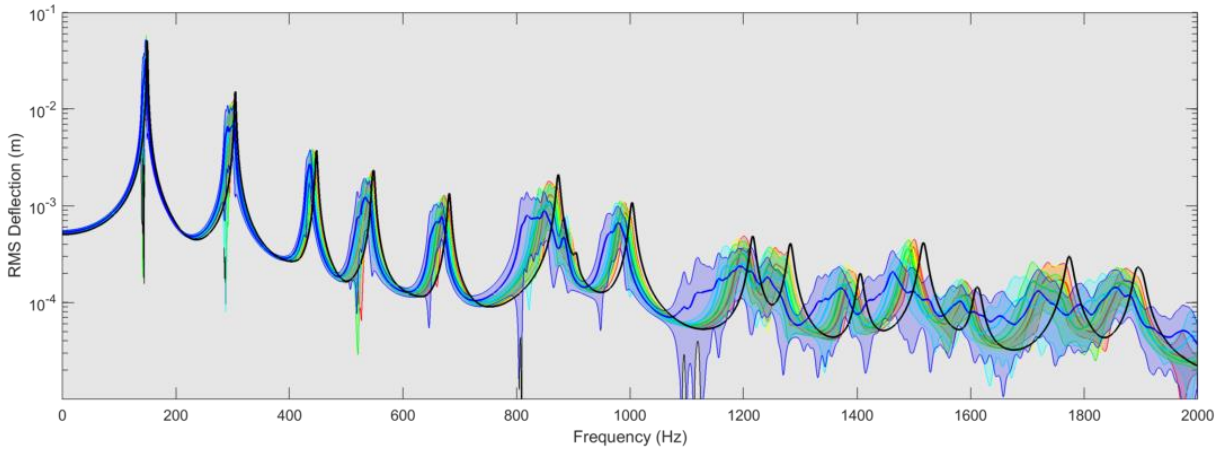

Figure 5.3: Frequency response curves averaged over 50 simulations for each of the five cut length cases (from 1" cuts (red) to 3" cuts (blue)). The filled regions indicate ± 1 standard deviation bounds. Increased damage severity not only results in greater variation from baseline, but also greater intra-class variation (larger spread).

Corrosion was modeled via a smooth but spatially-varying reduction in thickness of the plate. This phenomenological approach to corrosion modeling is common in the literature [109,

110]. More specifically, corrosion was modeled as a Gaussian reduction in thickness, parameterized by a center-point (the location of greatest thickness reduction), an amplitude, and a spread size (e.g. the standard deviation of the corrosion region, equal in all directions). 200 simulations were run for various corrosion cases, the severity of each case quantified by the percentage of plate material removed which ranged from 0.01% to 1.6% across all cases. Figure 5.4 shows the RMS displacement frequency response curves for all 200 corrosion simulations. Comparing against Figure 5.2, it is apparent that the corrosion causes a more subtle 'compression' effect on the frequency response relative to baseline.

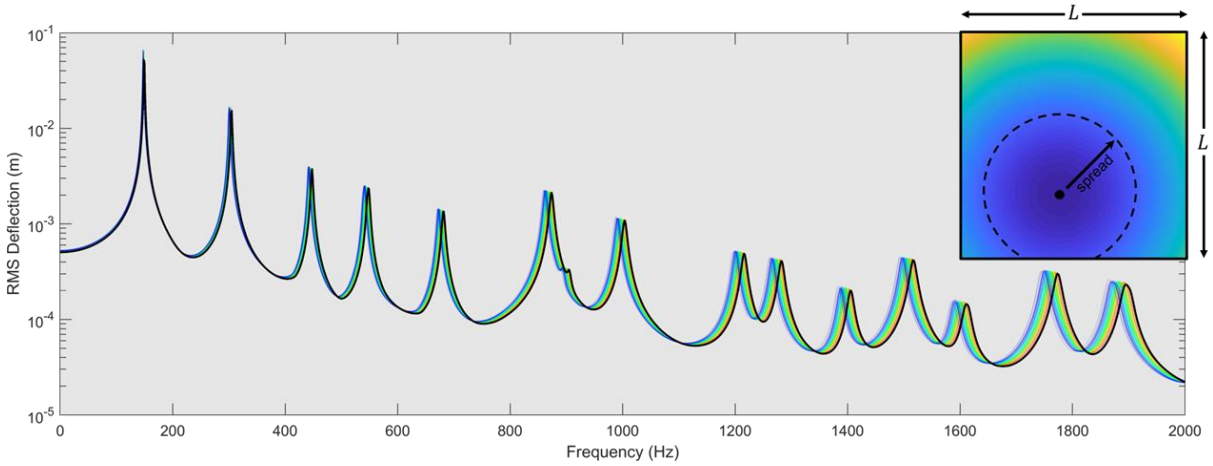

Figure 5.4: Frequency response curves of RMS displacement for all 200 corrosion cases. The red lines indicate little corrosion (0.01% of material removed) and the blue lines indicate the greatest corrosion (around 1.6% removed). A schematic of the Gaussian spread of the corrosion thickness reduction is shown in the top right inset.

Fastener failures, such as broken bolts and weld failures, were modeled as a localized change in the plate boundary conditions. Specifically, for each of 200 cases, a small fraction of the nominally-clamped perimeter was modified to be either simply-supported or free. The location and extent of this damaged-boundary region was randomized over the 200 samples, with half assigned to simply-supported and half assigned to free type constraints. The effects

of fastener damage can be very noticeable, especially near the plate's edges as shown in Figure 5.5 (note that the extent of damage shown here is illustrative and much greater than what was applied in the training data). Figure 5.6 displays the RMS displacement frequency response curves for the 200 simply-supported and free type boundary condition changes.

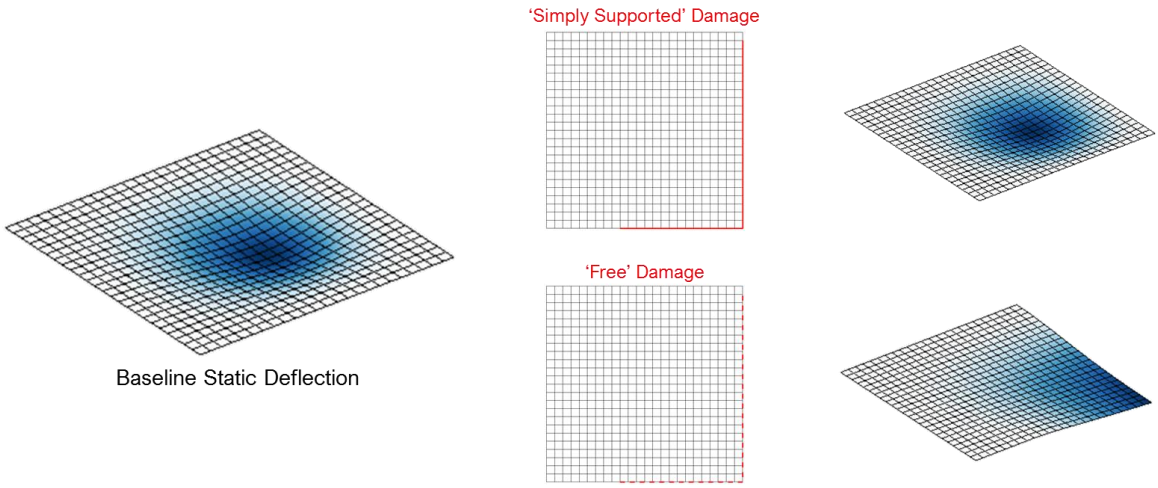

Figure 5.5: Illustrative example of the effects of severe fastener damage. At the left is the static displacement of the undamaged baseline plate due to point loading (at zero Hz). At the top right is the displacement of the same plate with approximately 30% of the perimeter simply-supported (shown as red lines on the mesh). At the bottom right is the displacement for the same extent of damage, but modeled instead with the more extreme free boundary condition.

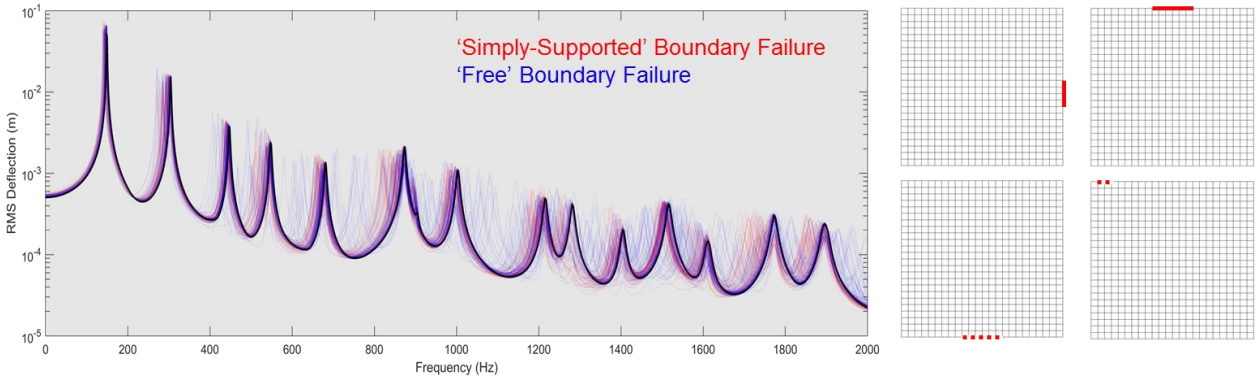

Figure 5.6: RMS displacement frequency response for 200 cases of fastener failure-type damage (left). Red and blue curves correspond to simply-supported and free edge modeled damage, respectively. Meshes from four cases indicating the general extent of boundary failure damage (right).

In order to train a classifier on the damaged data sets, meaningful features must be extracted from the frequency response curves. On one hand, there is no shortage of potential features that may be used to train a classifier [e.g. maximum value of the frequency responses, kurtosis of the frequency responses over the bandwidth, absolute (L_1) difference between test case and baseline frequency responses], though an ideal feature set is one that maximizes the inter-class variation in the feature space (to better separate/distinguish between classes) while minimizing the number of features (to avoid overfitting, especially when the size of the training data set is limited) [103]. The features used in this study were extracted using the dilational correlation function [104]

$$D_{xy}(\lambda) = \frac{\int x(\lambda f)\, y(f)\, df}{\sqrt{\int x^2(f) df \int y^2(f) df}} \tag{5.1}$$

The dilational correlation is analogous to the more conventional cross correlation, except the shifting property is exchanged for a stretching/scaling factor, λ. For each test case, the dilation correlation is computed between the log frequency response of that case and the log frequency response of the baseline plate. For each computation, two features are extracted, the maximum value of D_{xy} and the corresponding stretch factor λ_{max} at that maximum. Intuitively, the dilational correlation quantifies how similar a damaged plate frequency response curve is to a stretched or compressed version of the baseline frequency response curve. For instance, in the case of uniform thickness reduction it can be shown analytically that the frequency response curve is has a constant scaling compared to a full-thickness baseline (all peaks shift downward in frequency by equivalent ratios). In this scenario, the maximum dilation correlation would be almost exactly 1, and the value of λ_{max} would be dependent on the magnitude of the thickness reduction. As is seen in Figure 5.2 and Figure 5.6, cut and boundary failure-type damage clearly

do not result in scaled versions of the baseline frequency response, and it is the goal of the extracted dilational correlation features to identify and distinguish trends in how the damage affects the frequency response curves for the various damage classes. Figure 5.7 shows an example of the dilational correlation output between a damage case and baseline.

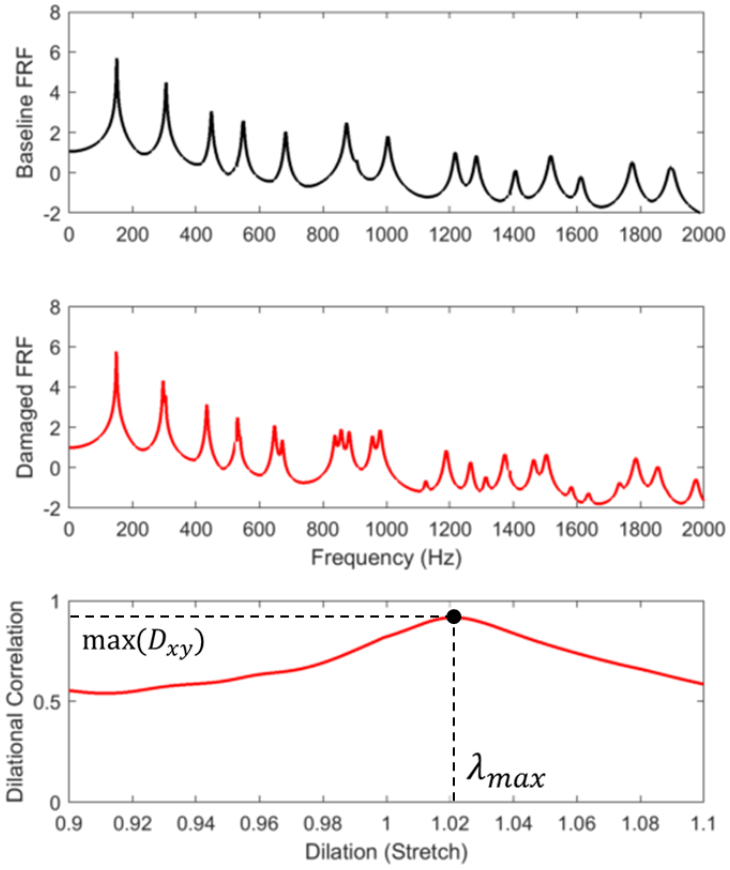

Figure 5.7: Example of dilational correlation function (bottom) between baseline (top) and 2-inch cut case (middle).

Finally, to investigate the feasibility of utilizing remote acoustic measurements of plate response, the plate was treated as existing in an infinite baffle and a Rayleigh integral was computed to estimate the acoustic amplitude response at particular receiver points in the 3D

field above the plate. The propagation medium was treated as air and fluid loading and damping effects were neglected in the analysis.

5.4 Simulation results

Figure 5.8 showcases the 2D feature space for the three damage cases considered. The features were extracted using the RMS displacement frequency responses, so no acoustic coupling is considered in Figure 5.8 (rather, in practice, the RMS displacement on the surface of the plate would need to be measured across the frequency band). Even by visual inspection, there is clear

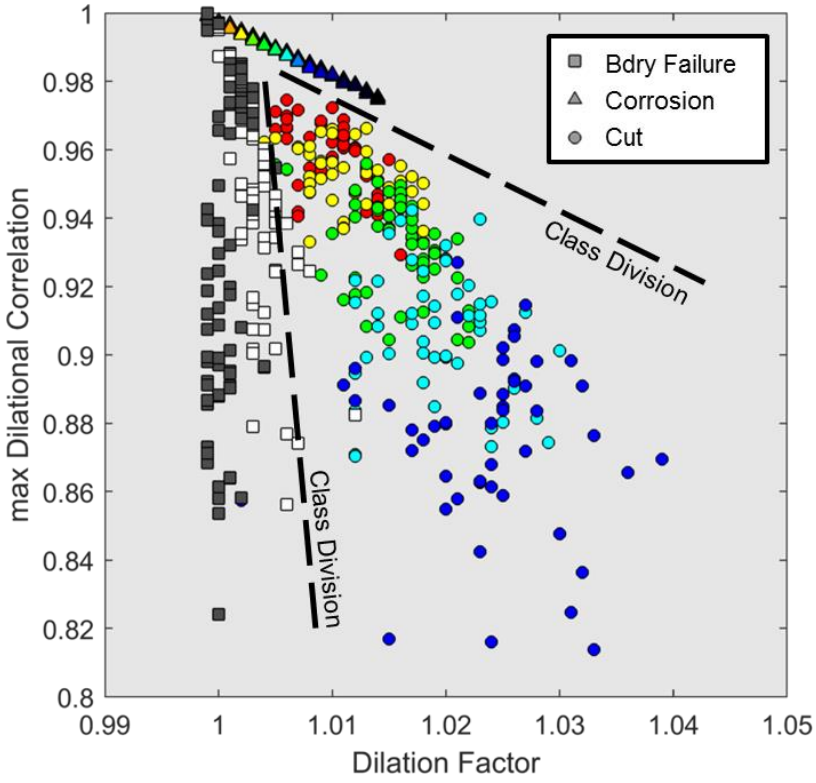

Figure 5.8: Feature space for the three damage classes using RMS displacement data. Colors increase from red to blue as damage severity (e.g. length of cut) increases. Qualitative decision boundaries are included for insight.

separation between the three damage classes. To quantify how well the data can be classified, an classification strategy must be selected. Myriad popular classification algorithms exist, many of which are easily implemented with a standard MATLAB package. The algorithms investigated in this study include:

- ***K Nearest Neighbors (KNN)*** – Points in the feature space are classified by popular vote of the 'closest' neighboring K points. For this study, $K = 5$, though this hyperparameter can be tuned to the data [103].
- ***Discriminant Analysis*** – The classes are assumed to be normally distributed, and boundaries are derived minimizing an L_2-norm. Linear Discriminant Analysis (LDA) and Quadratic Discriminant Analysis (QDA) fit linear and quadratic boundaries, respectively [107].
- ***Decision Trees (DT)*** – Decision boundaries are determined via iterative perpendicular partitioning of the feature space [108].
- ***Support Vector Machines (SVM)*** – Straight lines which maximize the margin of separation between classes are solved for, typically with better performance than LDA among other methods [105]. The use of kernel functions enables the feature space to be mapped (and later remapped) allowing for good separation with nonlinear boundaries. The kernel function considered here include *linear* (the unitary kernel function), *polynomial*, and *radial basis function* (RBF, sometimes referred to as a Gaussian kernel function). SVMs also typically feature one hyperparameter which can be tuned to the data (for instance, the spread of the RBF kernel) [111] .
- ***Neural Networks (NN)*** – A network of weight functions are tuned to training data to output a classification probability for later test data points (the class with the greatest

probability being assigned) [106]. The NN architecture used herein is *shallow*, meaning it features only a single *hidden layer* of nodes. The number of nodes in the hidden layer serves as a hyperparameter that can be tuned to the data, although 10 nodes were used in this analysis to avoid overfitting to the training data.

Figure 5.9 shows partitioned feature spaces for six of the classification strategies employed. The classifiers were trained with 70% of the generated data, then evaluated with the remaining 30% of test data. The Correct Classification Rate (CCR) is defined as the percentage of all of the test data that were correctly assigned to the correct damage type. Generally, all of the methods showed good performance, with the possible exception of LDA (which is expected since discriminant analysis is based on the errant assumption that the classes are normally distributed in the feature space). The more adaptable classifiers tend to perform very well, with CCRs above 96%.

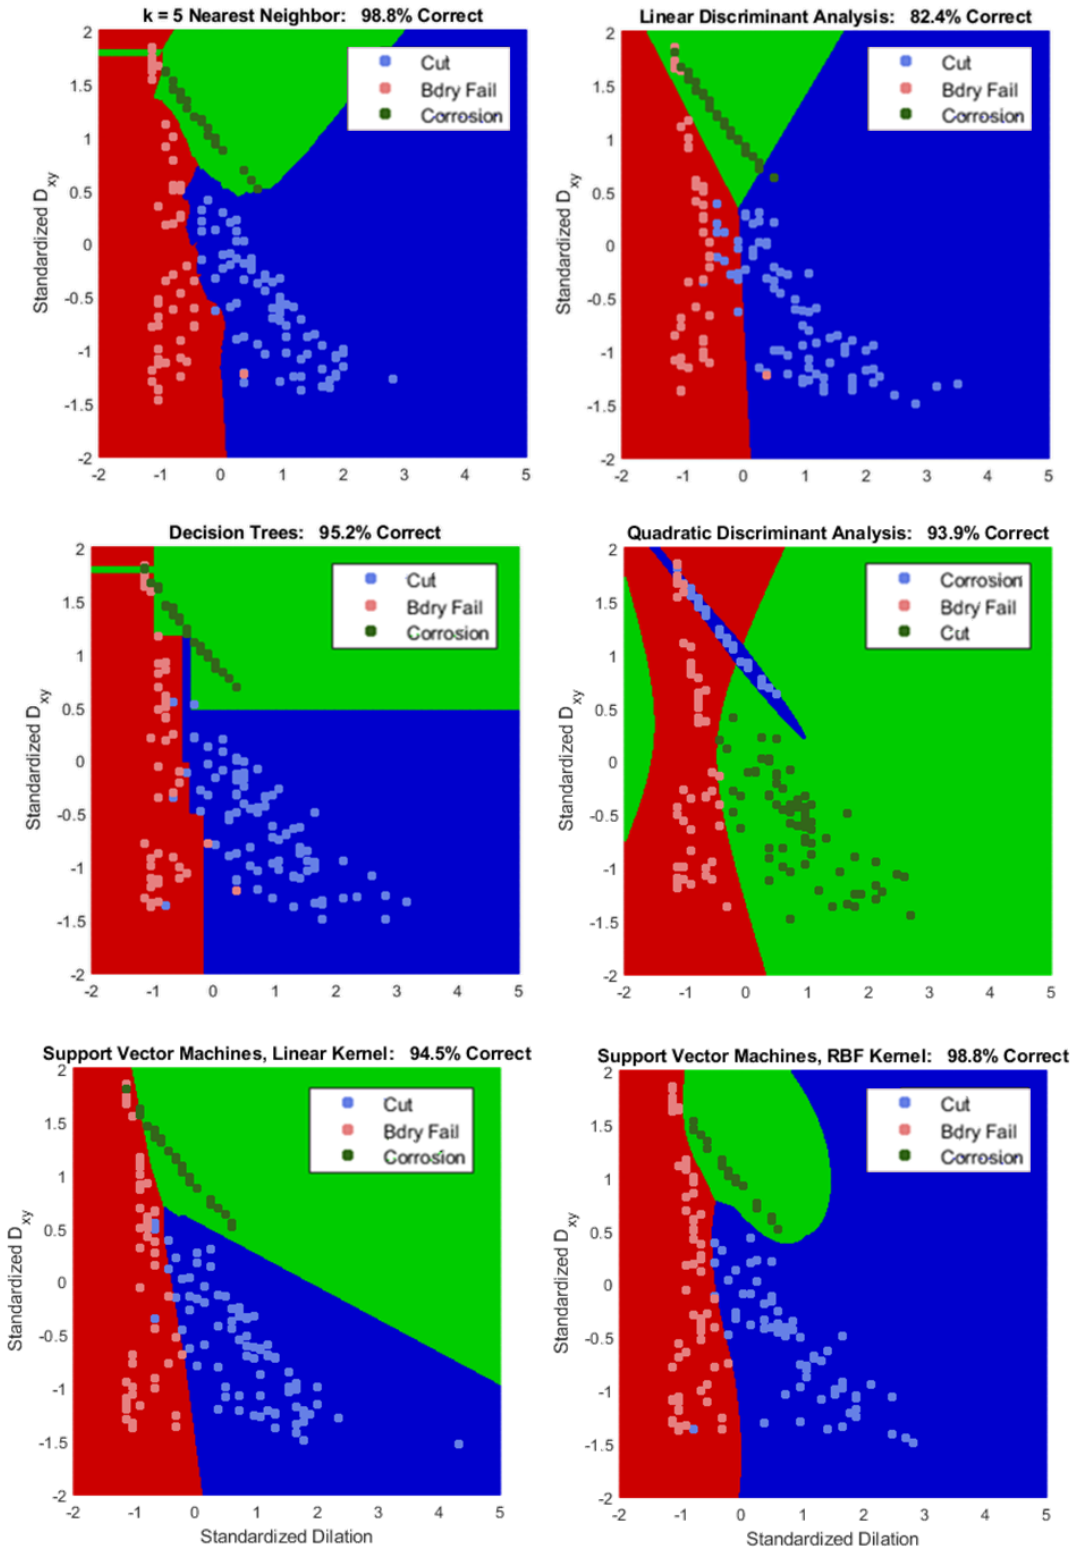

Figure 5.9: Partitioned feature space for six of the classification algorithms investigated. The correct classification rate of the test data is shown in the heading of each plot. With the exception of LDA, all methods resulted in >90% correct classification of the test data.

The training-then-test sequence was performed a total of 50 times for each classifier, with random permutations of training/test data, and the mean of each resulting CCR was computed to avoid outlying samples. The average classification performance of each classifier investigated, using RMS displacement data, are displayed in Table 5.1. The four best performing classifiers were determined to be KNN, Polynomial SVM, Radial Basis Function (RBF) SVM, and NN.

Table 5.1: Average Correct Classification Rate of each classifier investigated. The best performing classifiers are highlighted.

Classifier	Rate
KNN	97.5 %
LDA	84.3 %
QDA	93.5 %
Decision Trees	96.2 %
SVM – Linear	96.8 %
SVM – Poly	97.7 %
SVM – RBF	97.5 %
Neural Net	97.8 %

As previously mentioned, some of the classifiers listed in Table 5.1 possess hyperparameters, which are, in essence, additional user-defined 'knobs' that affect classifier performance. Intuitively, hyperparameters allow a user to 'tune' a classifier to maximize class separation without overfitting to the training data [103]. Hyperparameter sweeps were performed to optimize the hyperparameters and achieve optimal CCR. Figure 5.10 shows the hyperparameter sweeps for the K-Nearest Neighbor and Support Vector Machine classifiers. The CCR was only moderately dependent on values of K ranging from 1 to 20, with best performance for K = 5. For each of the SVM classifiers, the optimal hyperparameter value was

found to be approximately unity for linear, polynomial, and RBF kernels and polynomial-SVM generally outperformed linear and RBF SVMs, albeit slightly.

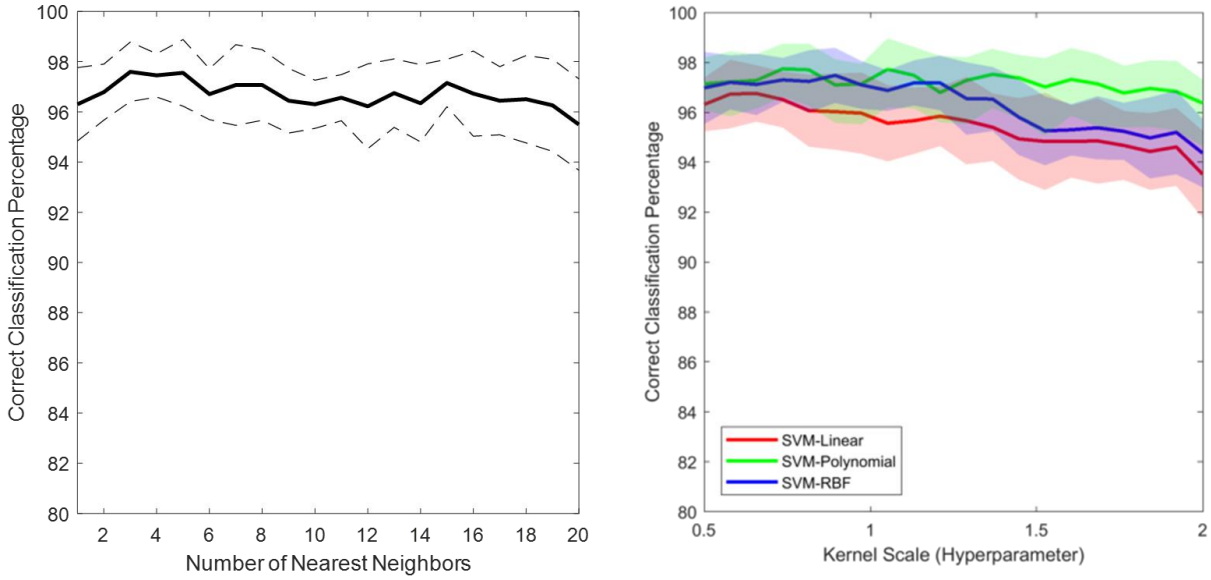

Figure 5.10: Hyperparameter sweeps for KNN (left) and SVM (right) classifiers. Classification performance was only subtly dependent on the hyperparameter values, with the optimal $K = 5$ and the SVM hyperparameters ~1.

It is important to note that, up to this point, all classification results have been predicated on knowing the RMS displacement of the plate as a function of frequency, which necessitates vibrational measurements over the entire surface of the plate. In practice such a measurement is likely challenging and/or expensive, or even impossible. A simpler, more obtainable measurement is that of a single acoustic receiver at some distance from the vibrating structure. To test the utility of an acoustic approach, a Rayleigh integral was used to model the acoustic response of the many test cases, and the subsequent frequency responses were computed for 74 receiver locations equispaced on a hemisphere centered at the plate. For each receiver location, the following classification approach was performed using a KNN ($K = 5$) classifier, and the

average CCR is plotted for each receiver location in Figure 5.11, for receiver ranges of 0.25 m (nearfield), 1.0 m, and 10 m (farfield).

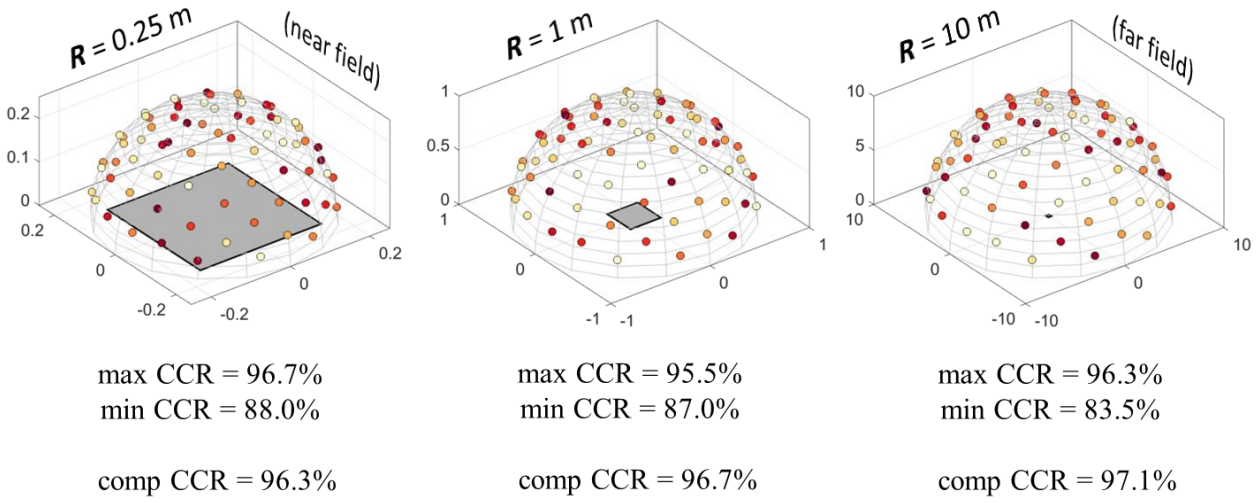

Figure 5.11: Remote acoustic classification results at various receiver positions relative to the plate. Results are shown for receiver ranges of 0.25 m (left), 1.0 m (center), and 10 m (right). The lightest dots correspond to 83% CCR and the darkest points correspond to 98% CCR. The 'comp CRR' results are for random batches of 10 composited receiver locations.

Referencing Figure 5.11, two results are immediate; on average, a remote acoustic receiver performs slightly worse that the RMS displacement measurements (CCR ~ 93% vs CCR ~98%), and there is no easily identifiable spatial trend for the CCR across the receiver locations. In general, it appears that any single, arbitrary receiver location is likely to perform worse that using RMS displacement data, with particularly disadvantageous receiver locations yielding CCR ~ 84% (at a range of 10 m). However, when data from a batch of receiver locations is used to train the classifier, performance is found to increase. Particularly, several random batches of 10 receiver locations were used to train the classifier, resulting in a composite CCR that was found to consistently perform above 96%, close to that achievable

with displacement data. This shows promise that by using several random receiver locations, robust and accurate classification is achievable with remote acoustic measurements.

5.5 Experimental validation

An experimental test apparatus, shown in Figure 5.13, was used to validate the simulated results. A clamped 12”x12”x1/16” aluminum plate was base-excited from below by an electrodynamic shaker outputting a 16 second linear-swept chirp from 100-2000 Hz. A laser Doppler vibrometer was used to measure the RMS deflection of the plate as a function of excitation frequency from 100-2000 Hz. This was done for two cases, the undamaged baseline plate, and a plate with a 3.8 cm (12.5% plate length) cut perpendicular to one of the sides of the plate. The RMS deflection curves are shown in Figure 5.12. The presence of the cut noticeably alters the frequency response, generally shifting the resonance peaks to lower frequencies. The measured response of the undamaged plate also agrees well with the simulated behavior within the excitation band

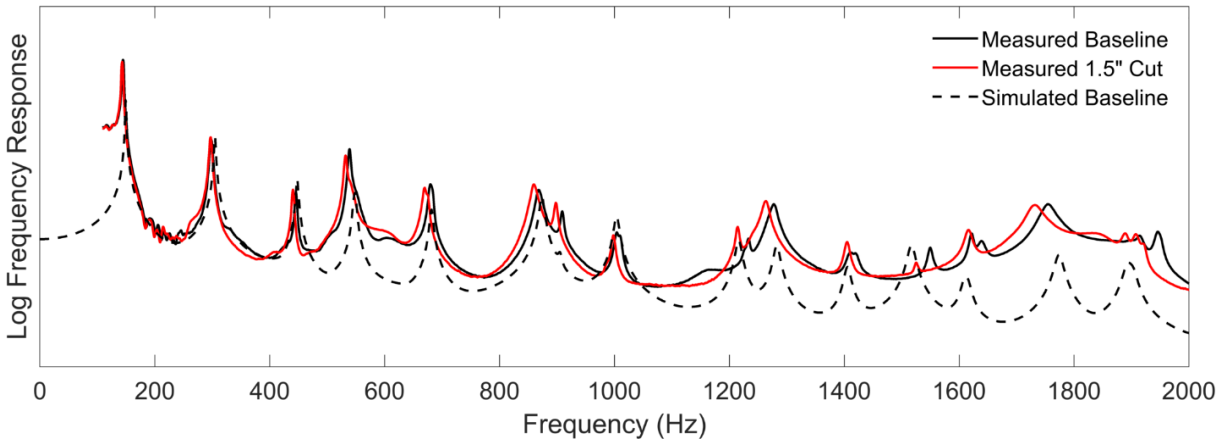

Figure 5.12: RMS deflection frequency response curves for the undamaged baseline plate (black) and a plate with a cut (red). The presence of the cut changes the peak shapes and locations,

generally by shifting them to lower frequencies. The simulated response of the undamaged plate is shown via the dashed line, and is in nominal agreement with the measured responses.

Sound cast from the plate was measured at the center receiver of a 15-element uniform line array at a range of 1 m and height of 0.4 m relative to plate center. The response of a healthy, undamaged (baseline) plate was first measured at the receiver. Then damage was added in the form of a 0.5” cut extending from one side of the plate and new acoustic measurements were taken. This procedure was repeated five more times, increasing the cut length in half inch increments to a maximum length of 3.0”. The top right inset of Figure 5.13 shows the test plate with a 1.5” through cut. In total, six damage cases and one baseline case were investigated.

Figure 5.13: Experimental test apparatus for the remote acoustic measurement of vibrating plates. The nominally clamped test plate (top right) is excited from below with an electrodynamic shaker (bottom right). An array of microphones (left) placed roughly 1 m from the plates were used to record radiated sound.

The magnitude frequency responses of the acoustic measurements were computed for the cut and baseline cases. Features were extracted using the same dilation correlation approach previously discussed, correlating each cut case with the single baseline case. A 1 Hz smoothing window (16 samples) was applied to the frequency response data to reduce fluctuations (1 Hz was selected based on the 1 Hz sampling resolution of the simulated data).

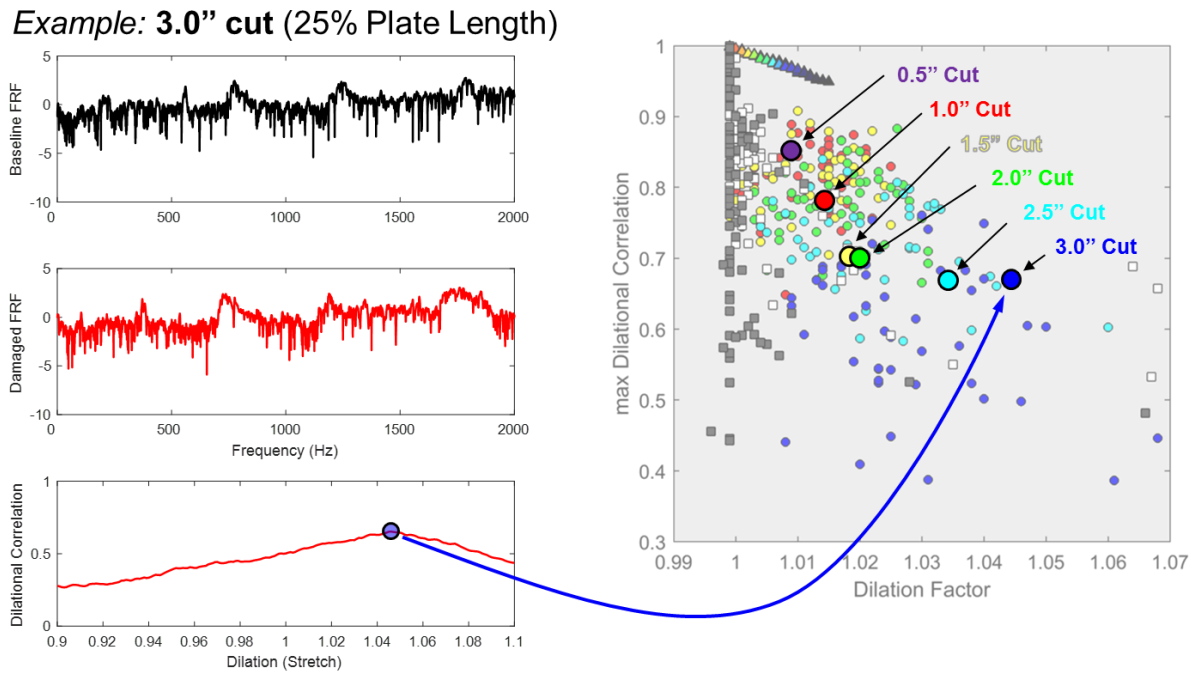

Figure 5.14: Empirical results from the feature extraction procedure for various length cuts using remote acoustic measurements. An example of the dilation correlation between the baseline and 3" cut case are shown (left). The feature set corresponding the six empirical cut cases is superimposed over the simulated feature space corresponding to a single receiver in the same relative location as in the measurements (right).

Figure 5.14 shows an example dilation correlation curve between measured baseline plate, and the same plate with a 3" cut. It is notable that, both in simulation and experiment, dilation correlation values are generally lower for acoustic data than for RMS displacement

data. The extracted feature set for all six measured cut cases is also provided in Figure 5.14, superimposed onto the feature space derived from the *simulated* acoustic response data at a theoretical receiver with the same geometry. The measured cases fall well within the grouping of cut cases, indicating good agreement between simulation and measurement. Further, the color of the circular markers (for both simulated and measured cases) corresponds to the cut length (e.g. 1.0" cut cases are red, 2.0" cut cases are green), and again, good agreement can be seen between simulated and measured cut cases in regards to the effect of cut length. Additional dilational correlation curves for other cut lengths are included in Figure 5.15

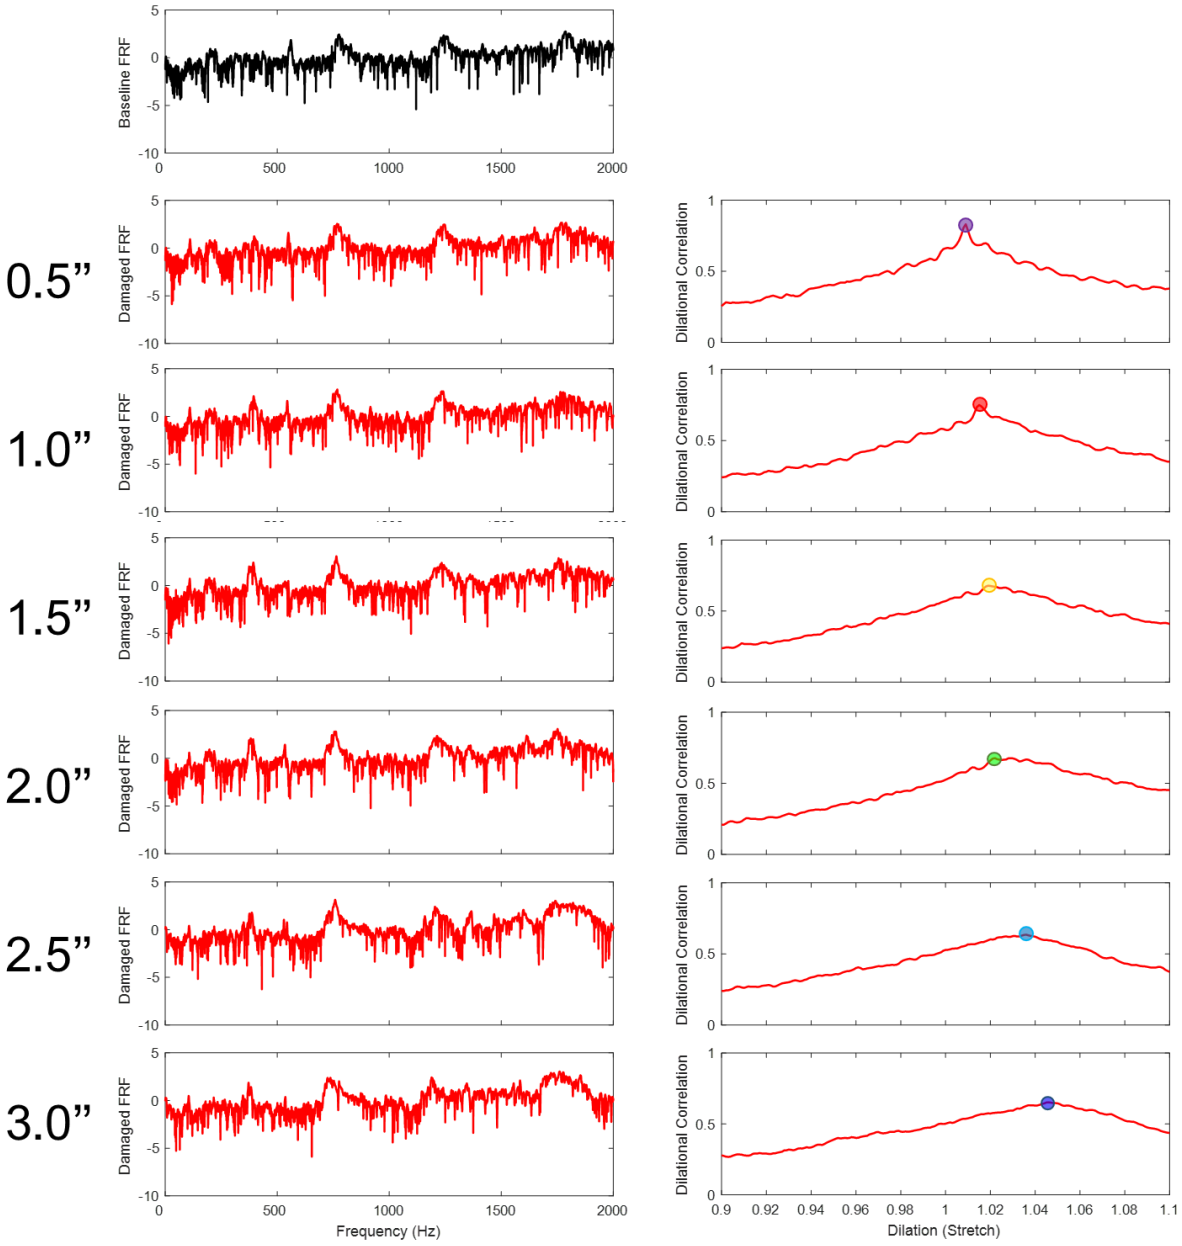

Figure 5.15: Dilation correlation feature extraction using a single remote acoustic receiver data at 3 meter range. At top left, is the baseline frequency response in black. Below, in red, are the frequency responses corresponding to a plate with a 0.5-3.0” cut. To the right are the corresponding dilation correlation curves, with the peak correlation value indicated with a colored marker, corresponding to Figure 5.14.

Additional acoustic measurements were conducted for various levels of fastener failure damage. Fastener failures were experimentally approximated by releasing one or more of the sixteen toggle clamps seen in Figure 5.13. Three cases were considered in order of increasing severity: 1) a single toggle clamp removed at a corner, 2) a single toggle clamp removed in the middle of a side, and 3) an entire side of four toggle clamps is removed. Acoustic recordings were collected at a 3 m range from the plate. The dilational correlation curves for these three cases are shown in Figure 5.16. A 1 Hz smoothing window was applied to the frequency response curves.

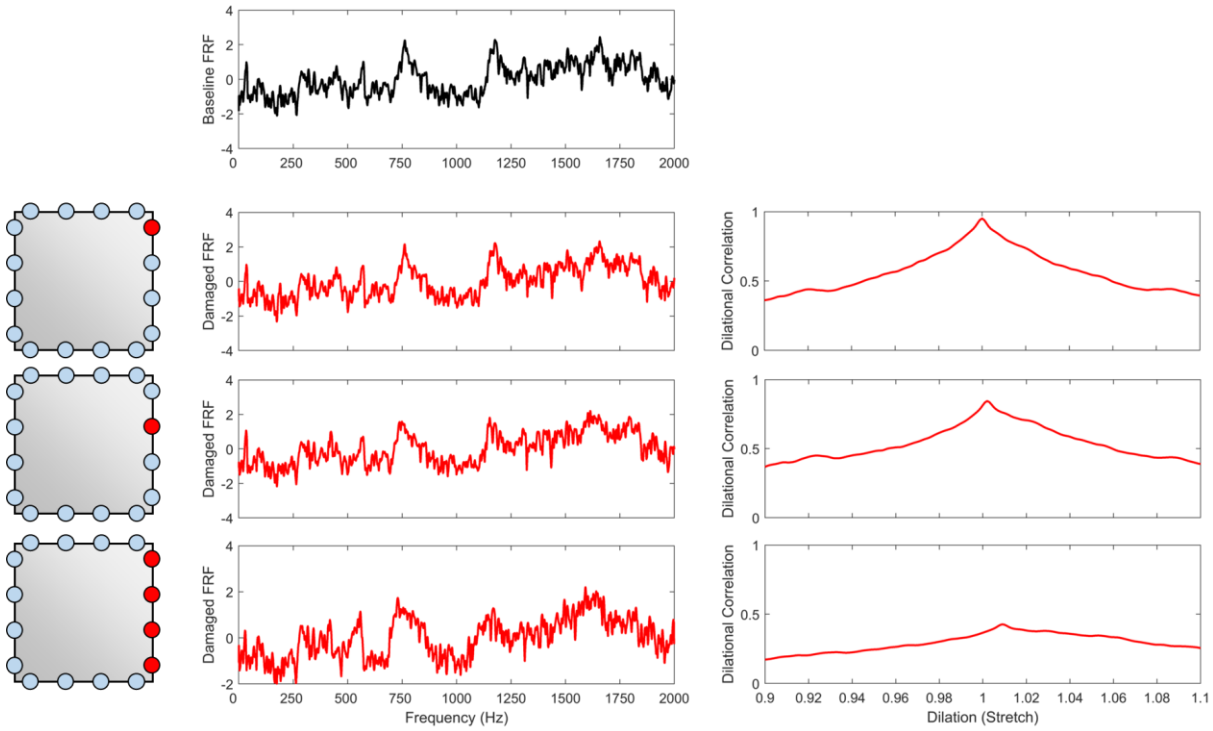

Figure 5.16: Dilation correlation curves for the measured acoustic response of a vibrating plate with synthetic fastener failure. Schematics at the left indicate which clamp(s) were removed (red dots) and which clamps remained in place (blue dots).

Figure 5.17 shows the complete feature space including experimental results for both cuts and fastener-failures (unclamping). Even with a limited number of experimental data points, it is clear that the empirical results agree well with the simulated cases, overlapping strongly with their respective groups. These results are promising for applications involving more complex structures or more realistically modeled damage, assuming that sufficiently accurate finite element models are available for both.

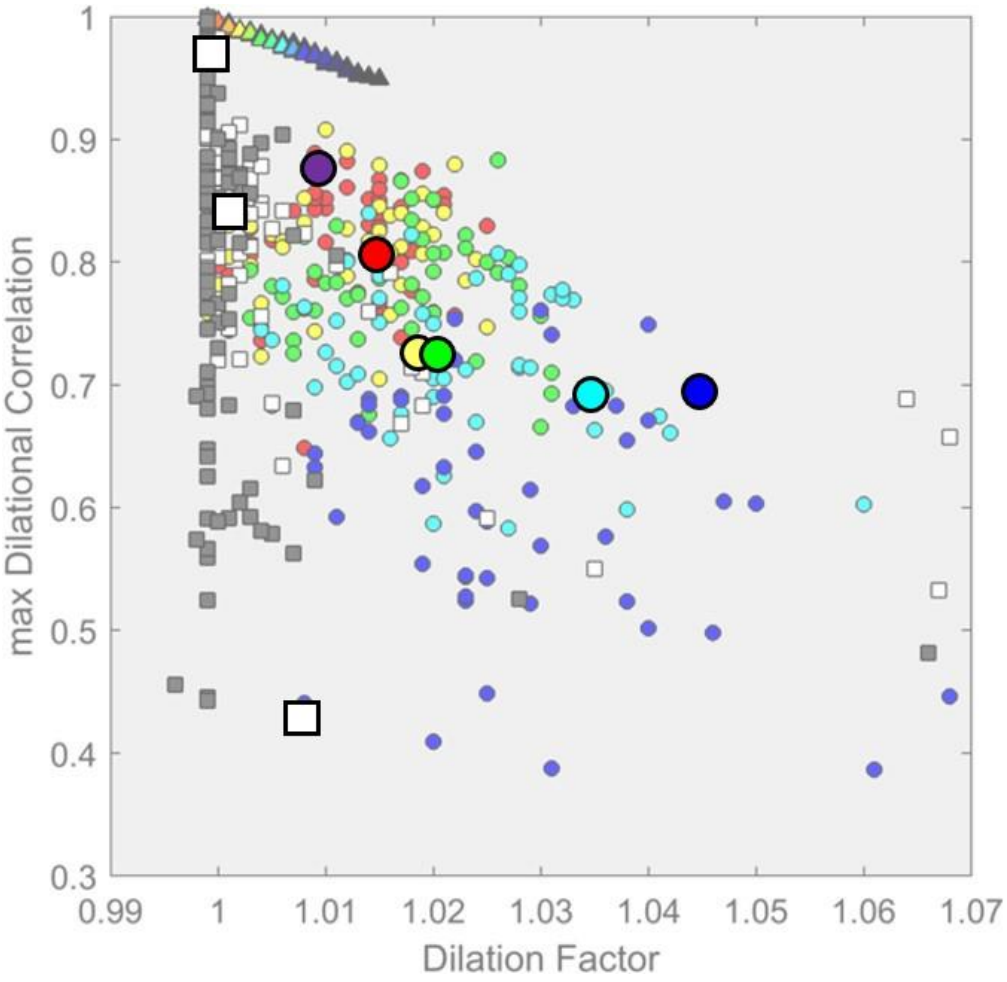

Figure 5.17: Acoustic response dilational correlation feature space, superimposed with experimental data points for cuts (large circles) and fastener failures (white squares). Both sets of experimental data agree well with the simulated training data.

With a populated feature space, and promising overlap between experimental and measured acoustic data, an preliminary attempt was made at extending the classification method to address the question of damage severity, specifically "given this recorded change in the acoustic frequency response relative to baseline, what is estimated *length* of a classified cut?". To address this question, which encroaches on the *Evaluation* component of SHM outlined in Figure 1.1, a regression analysis was performed the 2D feature space, using all of the 250 simulated cut cases, with dilation and maximum dilation correlation serving as the two inputs, and cut length serving as the regression output. Using Bayesian Linear Regression [112], a model was fit to the training data, along with associated uncertainty. Figure 5.18 shows the 95% confidence interval (±2 standard deviations) of the fitted model which contains the six measured damage cases (0.5-3.0 inch cuts, with 0.5 inch intervals).

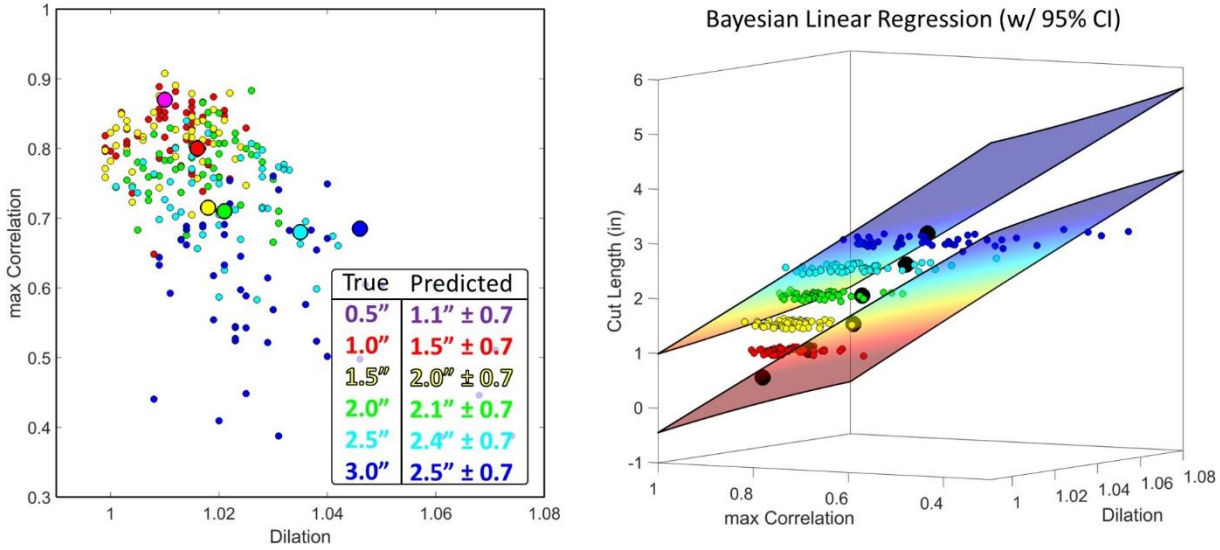

Figure 5.18: Regression analysis for extracting cut length from classification features. The feature space (left) is populated with 250 simulated cut cases (small circles), in addition to 6 measured cut cases (large circles) spanning from 0.5-3.0 inches in half-inch increments. The 95% confidence interval of the linear fit is shown (right), which contains the six measured damage cases (black circles), indicating that extracting cut size (estimating damage severity) is feasible within the 0.5-3.0 cut length range.

Notably, cut cases were only simulated for cut sizes down to 1.0 inch, so the 0.5 inch measured case is actually an extrapolation from the training data. Similar regression results were achieved with the use of nonlinear parametric fits, SVM regression, and neural nets. These results are promising, but the current 2D feature space is likely insufficient for accurate and precise estimation of damage severity (i.e. estimation with low uncertainty). This is die to the significant amount of overlap between the distinct damage cases, particularly for small damage severity (the color overlap at the left of Figure 5.18). Improved classification features, along with a broader range of simulated damage severity, would likely serve to improve the estimation of damage severity using remote acoustic measurements.

5.6 Conclusions

A feasibility study was conducted for the remote acoustic classification of damage in vibrating plates. Three classes of damage were investigated – cuts, corrosion, and boundary failures – and vibroacoustic data, generated from Monte Carlo finite element simulations for each class of damage, was used to train and evaluate classifiers. Features input into the classifiers were extracted using a dilational correlation between the frequency responses of each damage case and a 'healthy' simulated baseline case. When purely vibrational data were considered (i.e. RMS displacement of the plate vs frequency), test data were classified with roughly 97% accuracy using easily implementable KNN, SVM, and NN classification strategies. When remote acoustic data were used, classification rates were generally lower (~93%) and dependent on receiver location, with no clear trend as to which locations are

optimal. However, it was found that training with data from roughly 10 randomly placed receivers was sufficient for robust classification rates of 96%, commensurate with that achieved using more difficult-to-obtain displacement data. Preliminary experiments with various length cuts and amounts of unclamping in a 12”x12”x1/16” aluminum plate agreed well with simulations, which is promising for the applicability of data-driven damage classification methods.

Chapter VI: Summary & Conclusions

6.1 Summary

Damage in structures is often detrimental to their longevity, performance, and safety; and proper monitoring of damage is of extreme importance in many fields. Fortunately, the presence of damage generally results in measurable changes to a structure's vibroacoustic behavior which can be exploited for structural health monitoring (SHM). Though many methods exist for SHM the vast majority necessitate either: 1) that the structure is outfitted with dedicated sensors which serve to continuously monitor the structure actively or passively, or 2) a secondary system is employed which scans/interfaces with the structure, typically in close proximity. The drawback of the first set of methods is that dedicated sensors may be costly, they may interfere with the functionality or couple with the structure itself, and, for decades-long structural lifetimes, it may be challenging to distinguish faults in the structure from faults in the sensors. Drawbacks of the second set of methods, using deployable scanning systems, is that these systems also generally require close proximity to a test structure which may be difficult, time consuming, or costly to accomplish *in situ*.

In this thesis, an assortment of techniques was developed for remote acoustic sensing of structural changes (damage) in the canonical structure of a vibrating plate. Passive remote acoustic sensing (i.e. listening) for changes in a structure's acoustic response, can be performed at a distance with a relatively simple array of acoustic transducers that can be maintained

independently of the test structure. These beneficial properties could serve to reduce maintenance costs and increase measurement feasibility and reliability. For example, the use of a hydrophone array could be used to monitor sound radiated from a vessel leaving, then later returning, to a harbor; and significant changes to the acoustic signature of said vessel could be used to generate an alert recommending inspection and/or maintenance. However, remote acoustic methods are accompanied by some difficulties including complexities associated with acoustic propagation (e.g. multipath and reverberation), exogenous sources of noise, and the *a priori*-unknown sensitivity of the radiated acoustic field to structural damage. This thesis addresses these topics and is chapter-by-chapter summarized in the following paragraphs.

Chapter 2 details the methods and tools used throughout this thesis. Several array signal processing methods are outlined, which are used throughout later chapters. The experimental tools including the test apparatus, various sensors, and data acquisition devices are described. A discussion of the finite element techniques and acoustic coupling models used for simulating the vibroacoustic behaviors of damaged plates is also included.

Chapter 3 discusses the *detection* of damage in vibrating plates. A baseline-comparison method using remote receiver measurements is used to detect damage with statistically-quantified certainty. Of primary interest is the effectiveness of various detection metrics, the dependence on the character of the input forcing, and robustness to geometric error (i.e. receiver placement error) in unknown reverberant environments.

Chapter 4 discusses the *localization* of damage in vibrating plates. A high-resolution acoustic beamformer was developed (as an extension to an existing beamformer) to localize *small changes* in an otherwise loud, coherent acoustic field via a baseline subtraction. This

technique is evaluated through simulation and experiment and is applied for the localization of several damage types on actual plates.

Chapter 5 discusses the *classification* of damage in vibrating plates. A data-driven approach is used for the classification of various damage types using Monte Carlo finite element simulations of damaged plates as training data. Physics-based features sensitive to the different damage types are exploited for classification, and classification performance is investigated for different classification schemes. Additional comparisons between vibration-based classification and acoustic-based classification are also included.

Though the scope of this thesis was limited to the remote acoustic monitoring of vibrating square plates, the underlying techniques can theoretically be generalized to more specific, practical structures. As none of the methods summarized above are inherently model-based, they should be extendable to other structures provided: 1) the structures repeatably radiate sound while in operation, 2) the presence of damage affects the vibroacoustic behavior, and 3) the frequency bandwidth of the vibration encompasses several structural modes. Fortunately these constraints are met in a number of real-world systems. For instance, naval ships and submarines are known to possess measureable and consistent acoustic signatures [113, 114], and, due to the plate-like nature of their hulls, the vibroacoustic behaviors of these structures are sensitive to damage. Moreover, these structures radiate sound from <10 Hz, well into the kiloHertz regime, exciting many modes. Though testing and validation experiments would need to be performed for these and other structures, the scaling-up of these methods is likely feasible.

6.2 Conclusions

This thesis work has yielded the following unique conclusions relevant for remote acoustic sensing of structural health:

6.1 – ***Remote acoustic structural health monitoring can be studied with a modular test apparatus for broadband base-excitation of a 0.30-m-by-0.30-m square plate with clamped edges.***

Detailed in Section 2.2.2 Plate vibration rig and depicted in Figure 2.13, an experimental apparatus was constructed to investigate the vibrational and acoustic effects of various forms of damage on square plates. With the use of microphone arrays (Figure 2.9) and a laser Doppler vibrometer (Figure 2.15), precise measurements of changes in vibrational modes and the acoustic radiated signature were measured for several damage cases of common interest. Some of these findings include the localized vibrational effects of cuts on an aluminum plate (Figure 4.14), and delamination in a carbon fiber composite plate (Figure 3.21).

6.2 – ***Physically synthesized damage to a square vibrating plate, including through cuts, fastener failures, and delamination, can be remotely detected.***

Using the detection procedure detailed in Figure 3.14, remote acoustic detection of damage in vibrating plates was achieved by statistically comparing baseline (healthy) recorded measurements with potentailly damaged plate measurements. Several metrics

were investigated to quantify similarity-to-baseline, and the most well-rounded metric was found to be a correlation of power spectra (see Table 3.4). This procedure, shown to be effective for both deterministic and stochastic input forcing, was used to detect cuts as small as 1.3 cm in a 30-cm-square aluminum plate with greater than 99% probablity of detection and less than 1% probabilityof false alarm (see Figure 3.16 and Section 3.2.3 *Results & analysis for detection of cuts*). Similar findings were reported for synthesized fasterner/boundary failures (Section 3.1.2 *Results and analysis*) and delamination in composite plates (Section 3.3 Detection of delamination in composite materials). In the case of negligible reverberation (and/or perfectly repeatable source-receiver geometry between baseline and test measurements), greater than 99% detection was achievable using only a single acoustic receiver.

6.3 – ***Remote acoustic detection of damage can be made more robust in reverberant environments by incorporating a blind deconvolution step into the detection procedure.***

Unprocessed recordings made in reverberant environments were found to be corrupted by multipath propagation, leading to poor detection performance even for slight changes in source-to-receiver geometry (or in the environmental geometry) between baseline and test measurements. Implementation of the Synthetic Time Reversal (STR) algorithm (Section 2.1.2 Synthetic Time Reversal) for blind deconvolution was found to significantly improve the robustness of the remote acoustic detection procedure. For a receiver array 3 meters from the plate, a 108 mm shift of the array between baseline and test measurements resulted in poor detection performance of a 1.3 cm cut on par with a 50/50

coin flip (see Figure 3.18). With STR, the probability of detection for the same cut was improved to greater than 99%. Furthermore, it was found that as few as four equispaced receivers were required for STR to be effective (see Figure 3.18).

6.4 – ***A method was developed for constructing composite plates with known built-in delamination and such delamination is remotely detectable using the techniques described in the prior two conclusions.***

A method of synthesizing delamination in composite carbon fiber plates was presented (Section 3.3 Detection of delamination in composite materials) and remote acoustic detection of this damage was achieved for delamination patch sizes as small as 2.5 x 2.5 cm in a 30 cm square plate. Delamination-type damage is of particular noteworthiness because it is both common and visually indetectable since if occurs below a structure's surface.

6.5 – ***An existing method of high-resolution acoustic beamforming originally limited to incoherent acoustic source can be extended to distributed (spatially extended) coherent acoustic sources.***

A pre-existing acoustic high-resolution beamformer that was developed for use with incoherent noise-like sources was extended in applicability to coherent sources by implementing a subarray averaging step into the formulation (see Section 4.3.2 *Subarray averaging*). In addition to the structural health monitoring applications discussed herein,

this contribution is relevant in fields of study where source localization is plagued by coherent effects, such as multipath in underwater acoustic scenarios, or jamming/interference in radar applications.

6.6 – ***High-resolution localization of* small changes *in a distributed coherent acoustic source is possible using a novel technique, the Spectral Estimation Method with Coherent Background Removal (SEMCBR), developed here for this purpose.***

SEMCBR was developed and tested; it combines the beamformer of Conclusion 6.5 with an additional baseline subtraction step, allowing for the localization of small acoustic changes in an otherwise loud spatially-extended coherent sound source. This method was evaluated through simulation and experiment (see Figure 4.12 and Section 4.3.3 *Validation experiment*) and could be potentially relevant for applications involving the detection or localization of small changes in machines, automobiles, loudspeakers, etc.

6.7 –***Synthesized damage including through cuts, fastener failures, and delamination in broadband excited plates can be localized with the SEMCBR method, described in Conclusion 6.6, with improved resolution compared to the conventional method.***

Using a baseline subtraction technique for a localization of small changes in the response of various damaged plates, cuts (Figure 4.17), fastener failures (Figure 4.18), and delamination (Figure 4.19) were all successfully localized using SEMCBR. In all cases

SEMCBR yielded better resolution, and a more robust localization output than conventional methods.

6.8 – ***Simulation-based data-driven damage classification using vibration signatures and/or remote acoustic measurements is possible for vibroacoustic systems.***

Using Monte Carlo finite element simulations of a system with randomized, labeled damage, various classifiers were trained and evaluated based on the modeled vibrational or acoustic sensor responses. A physically-meaningful feature set (defined in Section 5.3 Classification methodology) was implemented which yielded 97% correct classification rates of simulated cut, corrosion, and fastener failure damage types. Results indicate that direct vibrational measurements yield better classification performance than remote acoustic measurements, however, when multiple acoustic receivers (i.e. an array) is used, the classification performance is nearly equal for both techniques. Additionally, comparisons between measured recordings of cut plates showed good agreement with simulated outcomes (Figure 5.14).

References

[1] D. Balageas, C.-P. Fritzen and A. Güemes, Structural health monitoring, John Wiley & Sons, 2010.

[2] E. P. Carden and P. Fanning, "Vibration Based Condition Monitoring: A Review," *Structural Health Monitoring,* pp. 355-377, 2004.

[3] J. Lifshitz and A. Rotem, "Determination of reinforcement unbonding of composites by a vibration technique," *Journal of Composite Materials,* vol. 3, no. 3, pp. 412-423, 1969.

[4] C. R. Farrar and K. Worden, "An introduction to structural health monitoring," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences,* vol. 365, no. 1851, pp. 303-315, 2007.

[5] S. W. Doebling, C. R. Farrar, M. B. Prime and D. W. Shevitz, "Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: a literature review," *Los Alamos National Laboratory report, LA-13070-MS,* 1996.

[6] O. S. Salawu, "Detection of Structural Damage through Changes in Frequency: A Review," *Engineering Structures,* pp. 718-723, 1997.

[7] D. Montalvao, N. M. M. Maia and A. M. R. Ribeiro, "A review of vibration-based structural health monitoring with special emphasis on composite materials," *Shock and vibration digest,* pp. 295-324, 2006.

[8] A. Rytter, "Vibration based inspection of civil engineering structures.," *Ph.D. Dissertation, Aalborg University, Denmark,* 1993.

[9] P. Chang, A. Flatau and S. Liu, "Review Paper: Health Monitoring of Civil Infrastructure," *Structural Health Monitoring,* vol. 2, no. 3, pp. 257-267, 2003.

[10] M. C. Junger and D. Feit, Sound, Structures, and Their Interaction, MIT Press, 1986.

[11] E. Williams, Fourier acoustics: sound radiation and nearfield acoustical holography, 1 ed., Elsevier, 1999.

[12] J. D. W. E. G. a. L. Y. Maynard, "Nearfield Acoustic Holography: I. Theory of Generalized Holography and the Development of NAH," *The Journal of the Acoustical Society of America,* pp. 1395-1413, 1985.

[13] W. Lu, W. Jiang, H. Wu and J. Hou, "A fault diagnosis scheme of rolling element bearing based on near-field acoustic holography and gray level co-occurrence matrix," *Journal of Sound and Vibration,* pp. 3663-3674, 2012.

[14] A. Raghavan and C. E. S. Cesnik, "Review of Guided-Wave Structural Health Monitoring," *Shock and Vibration Digest,* pp. 91-116, 2007.

[15] S. S. Kessler, S. M. Spearing and C. Soutis, "Damage Detection in Composite Materials Using Lamb Wave Methods," *Smart Materials and Structures,* pp. 269-278, 2002.

[16] A. Nair and C. Cai, "Acoustic emission monitoring of bridges: Review and case studies," *Engineering structures,* vol. 32, no. 6, pp. 1704-1714, 2010.

[17] M. Wevers and K. Lambrighs, "Applications of acoustic emission for SHM: A review," in *Encyclopedia of Structural Health Monitoring*, 2009.

[18] A. Croxford, P. Wilcox, B. Drinkwater and G. Konstantinidis, "November. Strategies for guided-wave structural health monitoring," *Proceedings of the Royal Society of*

London A: Mathematical, Physical and Engineering Sciences, vol. 463, no. 2087, pp. 2961-2981, 2007.

[19] D. W. H. D. A. Schindel, "Applications of micromachined capacitance transducers in air-coupled ultrasonics and nondestructive evaluation," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control,* vol. 42, no. 1, pp. 51-58, 1995.

[20] K. Zhu, Y. San Wong and G. S. Hong, "Wavelet Analysis of Sensor Signals for Tool Condition Monitoring: A Review and Some New Results," *International Journal of Machine Tools and Manufacture,* vol. 49, no. 7, pp. 537-553, 2009.

[21] M. Stojanovic and J. Preisig, "Underwater Acoustic Communication Channels: Propagation Models and Statistical Characterization," *IEEE Communications Magazine,* vol. 47, no. 1, pp. 84-89, 2009.

[22] K. Ehrenfried and L. Koop, "Comparison of iterative deconvolution algorithms for the mapping of acoustic sources," *AIAA Journal,* vol. 45, no. 7, pp. 1584-1595, 2007.

[23] P. Sijtsma, "CLEAN based on spatial source coherence," *International journal of aeroacoustics,* vol. 6, no. 4, pp. 357-374, 2007.

[24] Z. &. Y. Y. Chu, "Comparison of deconvolution methods for the visualization of acoustic sources based on cross-spectral imaging function beamforming," *Mechanical Systems and Signal Processing,* vol. 48, no. 1-2, pp. 40-422, 2014.

[25] B. S. Anami, V. B. Pagi and S. M. Magi, "Wavelet-based acoustic analysis for determining health condition of motorized two-wheelers," *Applied Acoustics,* vol. 72, no. 7, pp. 464-469, 2011.

[26] P. D. Samuel and D. J. Pines, "A review of vibration-based techniques for helicopter transmission diagnostics," *Journal of Sound and Vibration,* vol. 282, no. 1-2, pp. 475-508, 2005.

[27] N. Zuber and R. Bajrić, "Application of artificial neural networks and principal component analysis on vibration signals for automated fault classification of roller element bearings," *Eksploatacja i Niezawodność,* vol. 18, no. 2, 2016.

[28] H. Kumar, T. R. Kumar, M. Amarnath and V. Sugumaran, "Fault diagnosis of antifriction bearings through sound signals using support vector machine," *Journal of Vibroengineering,* vol. 14, no. 4, 2012.

[29] S. Das, A. N. Srivastava and A. Chattopadhyay, "Classification of damage signatures in composite plates using one-class SVMs," in *IEEE Aerospace Conference*, Big Sky, MT, 2007.

[30] C. R. Farrar, S. W. Doebling and D. A. Nix, "Vibration–based structural damage identification," *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences,* vol. 359, no. 1778, pp. 131-149, 2001.

[31] J. E. Stephens and J. T. Yao, "Damage assessment using response measurements," *Journal of Structural Engineering,* vol. 113, no. 4, pp. 787-801, 1987.

[32] X. Valero Gonzalez and F. Alías, "Automatic classification of road vehicles considering their pass-by acoustic signature," in *Proceedings of Meetings on Acoustics ICA2013*, Montreal, CA, 2013.

[33] H. Kalkan, N. F. Ince, A. H. Tewfik, Y. Yardimci and T. Pearson, "Classification of hazelnut kernels by using impact acoustic time-frequency patterns," in *EURASIP Journal on Advances in Signal Processing*, 2008.

[34] P. Herrera-Boyer, G. Peeters and S. Dubnov, "Automatic classification of musical instrument sounds," *Journal of New Music Research,* vol. 32, no. 1, pp. 3-21, 2003.

[35] J. B. Campbell and R. H. Wynne, Introduction to remote sensing, Guilford Press, 2011.

[36] D. R. Dowling and K. G. Sabra, "Acoustic remote sensing," *Annual Review of Fluid Mechanics,* vol. 47, pp. 221-243, 2015.

[37] C. Elachi and J. J. Van Zyl, Introduction to the physics and techniques of remote sensing, John Wiley & Sons, 2006.

[38] W. A. Fischer, W. R. Hemphill and A. Kover, "Progress in remote sensing," *Photogrammetria,* vol. 32, no. 2, pp. 33-72, 1976.

[39] K. G. Foote, "Discriminating between the nearfield and the farfield of acoustic transducers," *The Journal of the Acoustical Society of America,* vol. 136, no. 4, pp. 1511-1517, 2014.

[40] A. Pierce, Acoustics: An Intoduction to Its Physical Principles and Applications, Springer, 1989.

[41] M. Grant, S. Boyd and Y. Ye, "cvx users' guide," Stanford University, 2009.

[42] R. J. Urick, Principles of underwater sound for engineers., Tata McGraw-Hill Education, 1967.

[43] M. A. Ainslie, Principles of sonar performance modelling, Berlin: Springer, 2010.

[44] H. Krim and M. Viberg, "Two Decades of llrray Signal Processing Research," *IEEE signal processing magazine,* vol. 13, no. 4, pp. 67-94, 1996.

[45] M. I. Skolnik, Radar Handbook, 1970.

[46] T. L. Szabo, Diagnostic ultrasound imaging: inside out, Academic Press, 2004.

[47] A. R. Thompson, J. M. Moran and G. W. Swenson, Interferometry and synthesis in radio astronomy, New York: Wiley, 1986.

[48] P. E. Green Jr, E. J. Kelly Jr and M. J. Levin, "A comparison of seismic array processing methods," *Geophysical Journal International,* vol. 11, no. 1, pp. 67-84, 1966.

[49] H. G. Bölcskei, P. C. B. D. and A. J. Van der Veen, Space-time wireless systems: from array processing to MIMO communications, Cambridge University Press, 2006.

[50] S. Haykin, Array Signal Processing, Englewood Cliffs, NJ: Prentice-Hall, 1985.

[51] T. J. Johnson, R. L. Brown, D. E. Adams and M. Schiefer, "Distributed structural health monitoring with a smart sensor array," *Mechanical Systems and Signal Processing,* vol. 18, no. 3, pp. 555-572, 2004.

[52] L. E. Kinsler, A. R. Frey, A. B. Coppens and J. V. Sanders, Fundamentals of Acoustics, 4th ed., New York, NY: Wiley, 2000.

[53] W. A. Strauss, Partial differential equations: An introduction, Wiley, 2007.

[54] J. W. Goodman, Introduction to Fourier optics, Roberts and Company Publishers, 2005.

[55] L. C. Wrobel, The boundary element method, applications in thermo-fluids and acoustics, John Wiley & Sons, 2002.

[56] V. A. Fock, Fock, Vladimir Aleksandrovitch. Electromagnetic diffraction and propagation problems, Pergamon Press, 1965.

[57] M. S. Bartlett, "Periodogram analysis and continuous spectra," *Biometrika,* vol. 37, no. 1, pp. 1-16, 1950.

[58] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE,* vol. 57, no. 8, pp. 1408-1418, 1969.

[59] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation,* vol. 34, no. 3, pp. 276-280, 1986.

[60] T. Shan, M. Wax and T. Kailath, "On spatial smoothing for direction-of-arrival estimation of coherent signals," *Shan, Tie-Jun, Mati Wax, and Thomas Kailath. "On spatial smoothing for direction-of-arrival estimation of coherent signals." IEEE*

Transactions on Acoustics, Speech, and Signal Processing, vol. 33, no. 4, pp. 806-811, 1985.

[61] F. Li and R. J. Vaccaro, "Analysis of min-norm and MUSIC with arbitrary array geometry," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 26, no. 6, pp. 976-985, 1990.

[62] J. Li and P. Stoica, Robust adaptive beamforming, Hoboken, NJ: John Wiley, 2006.

[63] D. Blacodon and G. Elias, "Level estimation of extended acoustic sources using a parametric method," *Journal of Aircraft*, vol. 41, no. 6, pp. 1360-1369, 2004.

[64] S. Boyd and L. Vandenberghe, Convex optimization, Cambridge University Press, 2004.

[65] A. Ben-Tal and A. Nemirovski, "Robust convex optimization," *Mathematics of operations research*, vol. 23, no. 4, pp. 769-805, 1998.

[66] "Gurobi optimizer reference manual," GUROBI OPTIMIZATION, 2015. [Online]. Available: http://www.gurobi.com.

[67] J. Preisig, "Acoustic propagation considerations for underwater acoustic communications network development," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 11, no. 4, pp. 2-10, 2007.

[68] W. Munk, P. Worcester and C. Wunsch, Ocean acoustic tomography, Cambridge university press, 2009.

[69] K. G. Sabra and D. R. Dowling, "Blind Deconvolution in Ocean Waveguides Using Artificial Time Reversal," *The Journal of the Acoustical Society of America*, vol. 116, no. 1, pp. 262-271, 2004.

[70] S. H. Abadi, D. Rouseff and D. R. Dowling, "Blind Deconvolution for Robust Signal Estimation and Approximate Source Localization," *The Journal of the Acoustical Society of America*, vol. 131, no. 4, pp. 2599-2610, 2012.

[71] K. G. S. H. C. &. D. D. R. Sabra, "Ray-based blind deconvolution in ocean sound channels," *The Journal of the Acoustical Society of America,* vol. 127, no. 2, pp. EL42-EL47, 2010.

[72] S. C. Douglas, "Blind separation of acoustic signals," in *Microphone Arrays*, Heidelberg, Springer Berlin, 2001, pp. 355-380.

[73] R. Aichner, H. Buchner, F. Yan and W. Kellermann, "A real-time blind source separation scheme and its application to reverberant and noisy acoustic environments," *Signal Processing,* vol. 86, no. 6, pp. 1260-1277, 2006.

[74] D. Rouseff, Z. Wang, S. Zhou and L. Zhou, "The 2009 Cooperative Array Performance Experiment (CAPEx09) comparing vector-and pressure-sensor arrays," in *European Conference on Underwater Acoustics* , Istanbul, Turkey, 2010.

[75] P. Piezotronics, "Model 130E20 ICP® Electret Array Microphone Installation and Operating Manual," PCB Group, Depew, NY, 2014.

[76] A. W. Leissa, "The free vibration of rectangular plates," *Journal of sound and vibration,* vol. 31, no. 3, pp. 257-293, 1973.

[77] D. L. Logan, A first course in the finite element method, Cengage Learning, 2011.

[78] O. C. Zienkiewicz, R. L. Taylor, P. Nithiarasu and J. Z. Zhu, The finite element method, London: McGraw-Hill, 1977.

[79] S. S. Rao, The finite element method in engineering, Butterworth-Heinemann, 2017.

[80] R. J. Urick and G. C. Gaunaurd, "Detection of Fluctuating Sonar Targets, Technical Report NOLTR-72-47," Naval Ordnance Lab, White Oak, MD, 1972.

[81] J. A. Catipovic, " Performance limitations in underwater acoustic telemetry," *IEEE Journal of Oceanic Engineering,* vol. 15, no. 3, pp. 205-216, 1990.

[82] I. F. Akyildiz, D. Pompili and T. Melodia, "Underwater acoustic sensor networks: research challenges," *Ad hoc networks,* vol. 3, no. 3, pp. 257-279, 2005.

[83] T. Flynn and D. Dowling, "Remote acoustic detection of mechanical changes in a vibrating plate in an unknown reverberant environment," *The Journal of the Acoustical Society of America,* vol. 143, no. 2, pp. 1093-1101, 2018.

[84] W. Wang, J. Liou, R. Sutton and B. Dobson, "Machine vibration induced underwater acoustic radiation," *Journal of Marine Science and Technology,* vol. 8, no. 1, pp. 30-40, 2000.

[85] B. Petersson and B. Gibbs, "Towards a structure-borne sound source characterization," *Applied Acoustics,* vol. 61, no. 3, pp. 325-343, 2000.

[86] E. R. DeLong, D. M. DeLong and D. L. Clarke-Pearson, "Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach," *Biometrics,* vol. 44, no. 3, pp. 837-845, 1988.

[87] A. Mouritz, E. Gellert, P. Burchill and K. Challis, "Review of advanced composite structures for naval ships and submarines.," *Composite structures,* vol. 53, no. 1, pp. 21-42, 2001.

[88] P. Mallick, Fiber-reinforced composites: materials, manufacturing, and design, 1 ed., CRC Press, 2007.

[89] S. Valdes and C. Soutis, "Delamination detection in composite laminates from variations of their modal characteristics.," *Journal of sound and vibration,* vol. 228, no. 1, pp. 1-9, 1999.

[90] W. Seemann, "Seemann, William. "Unitary vacuum bag for forming fiber reinforced composite articles," U.S. Patent, 1994.

[91] A. Widodo, E. Kim, J. Son, B. Yang, A. Tan, D. Gu, B. Choi and J. Mathew, "Fault diagnosis of low speed bearing based on relevance vector machine and support vector machine," *Expert systems with applications,* vol. 36, no. 3, pp. 7252-7261, 2009.

[92] M. Zoltowski, "On the performance analysis of the MVDR beamformer in the presence of correlated interference," *IEEE Transactions on Acoustics, Speech, and Signal Processing,* vol. 36, no. 6, pp. 945-947, 1988.

[93] J. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astronomy and Astrophysics Supplement Series,* vol. 15, p. 417, 1974.

[94] D. Blacodon, "Spectral estimation method for noisy data using a noise reference," *Applied Acoustics,* vol. 72, no. 1, pp. 11-21, 2011.

[95] E. Williams, H. Dardy and K. Washburn, "Generalized nearfield acoustical holography for cylindrical geometry: Theory and experiment," *The Journal of the Acoustical Society of America,* vol. 81, no. 2, pp. 389-407, 1987.

[96] D. Adams, Health monitoring of structural materials and components: methods with applications, London, UK: John Wiley & Sons, 2007.

[97] T. J. Flynn and D. R. Dowling, "Remote acoustic detection of mechanical changes in a vibrating plate in an unknown reverberant environment," *The Journal of the Acoustical Society of America,* vol. 143, no. 2, pp. 1093-1101, 2018.

[98] R. M. Jones, Mechanics of composite materials, CRC Press, 2014.

[99] E. J. &. R. J. N. Barbero, "Modeling of delamination in composite laminates using a layer-wise plate theory," *International Journal of Solids and Structures,* vol. 28, no. 3, pp. 373-388, 1991.

[100] W. Seemann, "Unitary vacuum bag for forming fiber reinforced composite articles". U.S. Patent 5,316,462, 31 May 1994.

[101] D. Blacodon, "Spectral estimation method for noisy data using a noise reference," *Applied Acoustics,* vol. 72, no. 1, pp. 11-21, 2011.

[102] D. H. Johnson and D. E. Dudgeon, Array signal processing: concepts and techniques, Englewood Cliffs: PTR Prentice Hall, 1993.

[103] J. Friedman, T. Hastie and R. Tibshirani, The elements of statistical learning, New York City, NY: Springer, 2001.

[104] M. R. Belmont, "Dilation correlation functions and their applications," *IEE Proceedings F,* vol. 137, no. 5, pp. 345-350, 1990.

[105] H. Drucker, C. J. Burges, L. Kaufman, A. J. Smola and V. Vapnik, "Support vector regression machines," *Advances in neural information processing systems,* pp. 155-161, 1997.

[106] H. N. Mhaskar and T. Poggio, "Deep vs. shallow networks: An approximation theory perspective," *Analysis and Applications,* vol. 14, no. 6, pp. 829-848, 2016.

[107] W. R. Klecka, G. R. Iversen and W. R. Klecka, Discriminant analysis, Sage, 1980.

[108] J. R. Quinlan, "Induction of decision trees," *Machine Learning,* vol. 1, no. 1, pp. 81-106, 1986.

[109] R. E. Melchers, "Modeling of marine immersion corrosion for mild and low-alloy steels—Part 1: Phenomenological model," *Corrosion,* vol. 59, no. 4, pp. 319-334, 2003.

[110] J. K. Paik, A. K. Thayamballi, Y. I. Park and J. S. Hwang, "A time-dependent corrosion wastage model for seawater ballast tank structures of ships," *Corrosion Science,* vol. 46, no. 2, pp. 471-486, 2004.

[111] J. A. Suykens and J. Vandewalle, "Least squares support vector machine classifiers," *Neural processing letters,* vol. 9, no. 3, pp. 293-300, 1999.

[112] J. M. Marin and C. Robert, Bayesian core: a practical approach to computational Bayesian statistics, Springer Science & Business Media, 2007.

[113] M. Rauer and K. O'Neill, "R/V Neil Armstrong (AGOR-27) Acoustic Trial," Carderock, West Bethesda, MD, 2016.

[114] A. Averbuch, V. Zheludev, P. Neittaanmäki, P. Wartiainen, K. Huoman and K. Janson, "Acoustic detection and classification of river boats," *Applied Acoustics,* vol. 72, no. 1, pp. 22-34, 2011.

[115] P. M. Morse and K. U. Ingard, Theoretical Acoustics, 2nd ed., Princeton, NJ: Princeton University Press, 1987.

[116] S. D. Fassois and J. S. Sakellariou, "Time-series methods for fault detection and identification in vibrating structures," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences,* pp. 411-448, 2007.

[117] P. Poozesh, K. Aizawa, C. Niezrecki, J. Baqersad, M. Inalpolat and G. Heilmann, "Structural health monitoring of wind turbine blades using acoustic microphone array," *Structural Health Monitoring,* vol. 16, no. 4, pp. 471-485, 2017.

[118] T. J. Flynn and D. Dowling, "Remote acoustic sensing of boundary defects in a plate in a reverberant environment," London, UK, 2017.

[119] H. Sohn and C. R. Farrar, "Damage Diagnosis Using Time Series Analysis of Vibration Signals," *Smart Materials and Structures,* pp. 446-451, 2001.