

ACOUSTIC COHERENCE IMAGING THROUGH THE ATMOSPHERE

BY

JONATHAN WILLIAM BENSON

B.S., University of Illinois, 1992

M.S., University of Illinois, 1994

THESIS

**Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2000**

Urbana, Illinois

© Copyright by Jonathan William Benson, 2000

ACOUSTIC COHERENCE IMAGING THROUGH THE ATMOSPHERE

BY

JONATHAN WILLIAM BENSON

B.S., University of Illinois, 1992

M.S., University of Illinois, 1994

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2000

Urbana, Illinois

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
THE GRADUATE COLLEGE

APRIL 2000
(date)

WE HEREBY RECOMMEND THAT THE THESIS BY

JONATHAN WILLIAM BENSON

ENTITLED ACOUSTIC COHERENCE IMAGING

THROUGH THE ATMOSPHERE

BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR

THE DEGREE OF DOCTOR OF PHILOSOPHY

George W. Swenson, Jr.

Director of Thesis Research

D. M. King

Head of Department

Committee on Final Examination†

George W. Swenson, Jr.

Chairperson

A. J. Franke

John B.

Donald R. Jones

W. Brown

† Required for doctor's degree but not for master's.

ACOUSTIC COHERENCE IMAGING THROUGH THE ATMOSPHERE

Jonathan William Benson, Ph.D.
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, 2000
George W. Swenson, Jr., Advisor

The imaging of acoustic “scenes” through the atmosphere is a subject with little treatment in the literature. This thesis strives to set the state of the art in this field by presenting imaging methods and their performance in the turbulent atmosphere. Both narrowband interferometry and wideband techniques that can be solved using tomographic methods are presented. An extensive series of field experiments was conducted in which multiple sensor arrays of microphones recorded the signals received from loudspeaker sources outdoors under varying atmospheric conditions. Besides providing data for characterizing the imaging performance under different conditions, the field data also allowed for measurements to be made that characterize the atmosphere itself. Methods for enhancing image quality are discussed. It is shown that the popular self-calibration family of reconstruction methods from the radio astronomy community are not applicable to the acoustic imaging case. A new method, based on estimating the intensity distribution jointly with atmospheric distortion model parameters using multiple short integration period measurements, is derived. The performance of this method is encouraging.

DEDICATION

To my family and friends who have supported me throughout my studies, and especially to my wife Shelly for her never-ending love and support.

ACKNOWLEDGMENTS

First of all I would like to thank my advisor, Professor George W. Swenson, Jr., for his years of guidance and support. It has been a great experience working with him through the years. I would also like to thank my committee members, who have offered much advice and guidance, especially Professors Yoram Bresler and Steven Franke. I thank the staff at the U.S. Army Construction Engineering Research Laboratories, especially Dr. Larry Pater and Jeff Mifflin, for their generous support of my field experiment activities. This work has been supported in large part by contract DACA 88-95-C0010 from the U.S. Army Construction Engineering Research Laboratories.

TABLE OF CONTENTS

	Page
1. INTRODUCTION.....	1
1.1 Problem Statement	1
1.2 Previous Work on Acoustic Imaging in the Atmosphere.....	1
1.3 Organization	2
2. IMAGE FORMATION THEORY	4
2.1 Narrowband Interferometric Imaging.....	4
2.1.1 Derivation of the Van Cittert-Zernike theorem.....	4
2.1.2 Array design and image deconvolution.....	9
2.1.3 Multifrequency synthesis imaging.....	11
2.2 Wideband (2-D) Imaging	14
2.2.1 Tomographic solution method for wideband imaging	14
2.2.2 White noise imaging.....	17
3. FIELD EXPERIMENTS	20
3.1 Description of Experiments and Data	20
3.2 Characterization of the Atmosphere.....	26
3.2.1 A review of atmospheric turbulence theory	26
3.2.2 Coherence loss measurements.....	30
3.2.3 Computations of the wave structure function	31
3.2.4 Turbulence strength and scale parameter extraction	34
3.2.5 Propagation transfer function computations.....	35
4. IMAGING PERFORMANCE.....	43
4.1 Narrowband Imaging Performance	43
4.1.1 Non-turbulence-related performance issues	43
4.1.2 The effect of the atmosphere.....	52
4.1.3 Multifrequency synthesis imaging.....	63

4.2	Wideband Imaging Performance	65
4.2.1	Power spectrum imaging using tomographic solution methods.....	65
4.2.2	White noise imaging results	72
5.	IMAGE ENHANCEMENT	74
5.1	Self-Calibration and its Application to Acoustic Coherence Imaging	74
5.1.1	Phase closure and self-calibration	74
5.1.2	Applicability of phase closure to acoustic imaging.....	77
5.2	A Discussion of Atmospheric Distortion Models	82
5.3	An Estimation Problem Formulation for Image Recovery.....	86
5.3.1	Problem formulation	87
5.3.2	Derivation of the Cramér-Rao bounds	89
5.3.3	Cramér-Rao lower bound results and discussion	91
5.3.4	Initial solution attempts and results.....	95
6.	CONCLUSIONS	107
	REFERENCES.....	109
	VITA	114

1. INTRODUCTION

1.1 Problem Statement

The use of multiple sensor arrays to extract information about propagated waves in a medium has been widely studied for a variety of media and wave types. The fields of sonar image processing, radar, and radio astronomy have been studied for decades. The work described and proposed in this paper will strive to set the state of the art in the formation of images of acoustic wave fields in the atmosphere through the use of multiple sensor interferometric processes.

The problem to be explored in this thesis is that of producing an image of an acoustic scene using signals received from an array of microphones. This image would take the form of a map of the intensity as a function of angle of arrival, or more interestingly, a map of the power spectrum versus angle. For this thesis, a one-dimensional array will be assumed with imaging occurring in the horizontal plane near the ground. It will also be assumed that the acoustic sources being imaged are spatially incoherent and exist in the far field of the measurement array.

1.2 Previous Work on Acoustic Imaging in the Atmosphere

Acoustic imaging through the atmosphere has not been investigated in depth in the past, and only a handful of papers exist that describe this type of image formation. One of the first publications dealing with acoustic imaging in the atmosphere is by Billingsley and Kinns. They developed a rudimentary beamforming system for sound source localization on jet engines [1]. The images were produced via a simple delay-and-sum beamforming

technique using 14 microphones. A short time later, a group of French engineers published work that related ideas used in the radio astronomy community to the problem of acoustic imaging through the atmosphere [2]. Chiollaz et al. [3] and Escudié et al. [4] developed an acoustic imaging system based on the use of the two-dimensional Fourier transform of the space-time correlation functions of the elements of the array [2]-[4]. This method is equivalent in many ways to traditional interferometric techniques and is essentially the same as the wideband imaging technique described later in this thesis. They also introduce methods for dealing with moving objects. Much of the previous work in acoustic imaging through the atmosphere has been motivated by the measurement of the noise generated by moving vehicles, such as trains. Boone and Berkhout developed a similar system for the measurement of industrial noise [5]. Van Der Toorn et al. used this measurement system, which was named the SYNTACAN (SYNThetic Acoustic Antenna), to measure the spatial noise distributions of a high-speed train [6]. None of the work above included a thorough investigation of effects of the atmosphere on their measurements. This thesis will include a detailed discussion of the effects of the atmosphere on imaging performance. Also, some methods for removing these effects will be presented and investigated.

1.3 Organization

This thesis will strive to describe the state of the art in the formation of images of acoustic sources through the atmosphere by using multiple sensor interferometric processes. In Chapter 2, the theory for traditional narrowband interferometric imaging as well as wideband imaging will be presented. Next, in Chapter 3, field experiments conducted to support the imaging work are discussed. Also discussed are several measurement results that attempt to paint a picture of the atmosphere through which we are trying to produce images. In Chapter 4, performance characteristics for the narrowband imaging case, such as the variance associated with coherence measurements, will be discussed for situations both including and excluding the effects of the atmosphere. In addition, examples of the wideband imaging technique will be presented. In Chapter 5, methods for enhancing the performance of a narrowband imaging system under the influence of the atmosphere are

discussed. These methods include some used in the radio astronomy community and also another based on the joint estimation of an intensity image along with atmospheric model parameters. Chapter 6 contains a summary and the overall conclusions of the work completed in this thesis.

2. IMAGE FORMATION THEORY

In this chapter, the theory behind the formation of acoustic images in both the narrowband and wideband cases is presented. First is a derivation of the Van Cittert-Zernike theorem, the fundamental theorem of interferometric imaging. Then, a description of wideband imaging theory, including the case where the source distribution is assumed white, is presented.

2.1 Narrowband Interferometric Imaging

The radio astronomy community has been making images of celestial objects for decades using interferometry based on the Van Cittert-Zernike theorem. In this section, the theorem will be derived for the case of acoustic imaging. Also, discussions on array design, image deconvolution, and multifrequency synthesis will be presented.

2.1.1 Derivation of the Van Cittert-Zernike theorem

In this section, the widely used relationship known as the Van Cittert-Zernike theorem will be derived in the context of acoustic imaging. This theorem states that a coherence measurement on a given baseline in an array gives a sample of the Fourier transform of the intensity distribution of the source. This is true provided that the received signals are narrowband, that the source distribution is spatially incoherent, and that the source is stationary and is located in the far-field of the receiving array. The derivation here generally follows that given by Swenson and Mathur [7] and Thompson et al. [8].

The physical situation at hand is shown in Fig. 2.1. Let us assume that there is a source function $S(\xi, t)$ that represents the acoustic radiation of the source as a function of

angle and time, where ξ is the sine of the angle θ which is measured with respect to the normal to the array baseline. The acoustic waves then travel to the array sensors located at positions x_n and x_m . The signal received at the sensor located at position x_n is $y_n(t)$ and is given by

$$y_n(t) = \int_{-\infty}^{\infty} S(\xi, t - \frac{x_n}{c} \xi) d\xi \quad (2.1)$$

where c is the speed of sound. The source function is assumed to be zero outside the visible range

$$S(\xi, t) = 0, \quad |\xi| > 1. \quad (2.2)$$





Fig. 2.1: Diagram of the physical setup involved in the derivation of the Van Cittert-Zernike theorem.

Since $S(\xi, t)$ is assumed to be a stationary random process and therefore extends from infinite past to infinite future, it does not satisfy the condition for the existence of the Fourier transform. We can still use Fourier analysis if we consider the truncated function below.

$$S_T^r(\xi, t) \equiv \begin{cases} S(\xi, t) & |t| \leq T \\ 0 & |t| > T \end{cases} \quad (2.3)$$

In terms of this truncated function, we define the complex analytic signal

$$S_T(\xi, t) = S_T^r(\xi, t) + jS_T^i(\xi, t) \quad (2.4)$$

where the real and imaginary parts form a Hilbert transform pair. The real part represents the actual acoustic signal received. Even as the real part is truncated, the imaginary part is not necessarily so. We can now write the Fourier transform relationships

$$\begin{aligned} \hat{S}_T(\xi, f) &= \int_{-\infty}^{\infty} S_T(\xi, t) e^{-j2\pi ft} dt \\ S_T(\xi, t) &= \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) e^{j2\pi ft} df. \end{aligned} \quad (2.5)$$

A deeper treatment of the analytic signal concept can be found in [9].

Of great importance is the cross correlation between the source signals from two different directions at two different times, which is usually defined as

$$\Gamma_{12}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T S(\xi_1, t) S^*(\xi_2, t + \tau) dt \quad (2.6)$$

We now find it more convenient to use our truncated function to avoid Fourier transform existence problems. This leads to the definition of the *mutual coherence function* of the source

$$\Gamma(\xi_1, \xi_2, \tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-\infty}^{\infty} S_T(\xi_1, t) S_T^*(\xi_2, t - \tau) dt \equiv E[S(\xi_1, t) S^*(\xi_2, t - \tau)] \quad (2.7)$$

where E denotes the expected value operator. The mutual coherence function is a function of the time difference τ because of our stationarity assumption. With this definition let us continue and begin to analyze the situation shown in Fig. 2.1.

With reference to Fig. 2.1, we assume that the acoustic signals are received by omnidirectional microphones located at positions x_n and x_m . These signals, $y_n(t)$ and $y_m(t)$ are then filtered with a filter having frequency response $H(f)$. The outputs of the filters, $v_n(t)$

and $v_m(t)$, are then correlated, resulting in the output $r(\tau)$. The signals $y_n(t)$ and $y_m(t)$ can be written as

$$y_{n_r}(t) = \int_{-\infty}^{\infty} S_T(\xi, t - \frac{x_n}{c}\xi) d\xi \quad (2.8)$$

$$y_{m_r}(t) = \int_{-\infty}^{\infty} S_T(\xi, t - \frac{x_m}{c}\xi) d\xi \quad (2.9)$$

where from now on the truncated functions are assumed used. We can then also write the Fourier transforms of these signals as

$$\hat{y}_{n_r}(f) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) e^{-j2\pi f \frac{x_n}{c}\xi} d\xi \quad (2.10)$$

$$\hat{y}_{m_r}(f) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) e^{-j2\pi f \frac{x_m}{c}\xi} d\xi. \quad (2.11)$$

The Fourier transforms of the signals $v_n(t)$ and $v_m(t)$ can then be written as

$$\hat{v}_{n_r}(f) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) H(f) e^{-j2\pi f \frac{x_n}{c}\xi} d\xi \quad (2.12)$$

$$\hat{v}_{m_r}(f) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) H(f) e^{-j2\pi f \frac{x_m}{c}\xi} d\xi \quad (2.13)$$

where the convolution property of the Fourier transform was used [10]. We can then transform back into the time domain to obtain

$$v_{n_r}(t) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) H(f) e^{-j2\pi f \frac{x_n}{c}\xi} e^{j2\pi ft} d\xi dt \quad (2.14)$$

$$v_{m_r}(t) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) H(f) e^{-j2\pi f \frac{x_m}{c}\xi} e^{j2\pi ft} d\xi dt. \quad (2.15)$$

Following Fig. 2.1, the signals $v_n(t)$ and $v_m(t)$ are then correlated. Defining the cross correlation as

$$r(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-\infty}^{\infty} v_{m_r}(t) v_{n_r}^*(t - \tau) dt = E[v_{m_r}(t) v_{n_r}^*(t - \tau)] \quad (2.16)$$

the output of the correlator $r(\tau)$ can be written as

$$r_{mn}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-\infty}^{\infty} dt \int_{-\infty}^{\infty} df_1 \int_{-\infty}^{\infty} df_2 \int_{-\infty}^{\infty} d\xi_1 \int_{-\infty}^{\infty} d\xi_2 \hat{S}(\xi_1, f_1) \hat{S}^*(\xi_2, f_2) \times H(f_1) H^*(f_2) e^{-j2\pi f_1 \frac{x_m}{c} \xi_1} e^{j2\pi f_2 \frac{x_n}{c} \xi_2} e^{j2\pi f_1 t} e^{-j2\pi f_2 (t-\tau)}. \quad (2.17)$$

Note that the result of the t-integral is $\delta(f_1 - f_2)$. This allows us to evaluate one of the f -integrals, leading to the new expression

$$r_{mn}(\tau) = \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} d\xi_1 \int_{-\infty}^{\infty} d\xi_2 E[\hat{S}(\xi_1, f) \hat{S}^*(\xi_2, f)] |H(f)|^2 e^{-j2\pi f \frac{x_m}{c} \xi_1} e^{j2\pi f \frac{x_n}{c} \xi_2} e^{j2\pi f \tau}. \quad (2.18)$$

At this point we need to enforce our assumptions about the source function. We can write

$$E[\hat{S}(\xi_1, f) \hat{S}^*(\xi_2, f)] = \hat{\Gamma}(\xi_1, f) \delta(\xi_1 - \xi_2) \quad (2.19)$$

where $\hat{\Gamma}(\xi, f)$ is the temporal Fourier transform of the source autocorrelation function. The Dirac delta function follows from our requirement that the acoustic radiation is spatially incoherent. Using Eq. (2.19) in Eq. (2.18), we obtain

$$r_{mn}(\tau) = \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} d\xi \hat{\Gamma}(\xi, f) |H(f)|^2 e^{-j2\pi f \frac{(x_n - x_m)}{c} \xi} e^{j2\pi f \tau}. \quad (2.20)$$

Let us now assume that the filter is narrowband, centered at some frequency f_o with bandwidth Δf . Then, we can rewrite Eq. (2.20) as

$$r_{mn}(\tau) = \int_{f_o - \Delta f}^{f_o + \Delta f} df \int_{-\infty}^{\infty} d\xi \hat{\Gamma}(\xi, f) |H(f)|^2 e^{-j2\pi f \frac{(x_n - x_m)}{c} \xi} e^{j2\pi f \tau}. \quad (2.21)$$

If we make that assumption that Δf is very small and that $\hat{\Gamma}(\xi, f)$ does not change significantly over the range, then we obtain

$$r_{mn}(0) = \int_{-\infty}^{\infty} \hat{\Gamma}(\xi, f_o) e^{-j2\pi f_o \frac{(x_n - x_m)}{c} \xi} d\xi = \hat{\gamma}(u, f_o) \quad (2.22)$$

where we have evaluated the correlation at zero-lag, where $\hat{\Gamma}(\xi, f_o)$ is proportional to the intensity of the source, and where $\hat{\gamma}(u, f_o)$ is the spatial frequency spectrum of the mutual coherence function of the source. We define the spatial frequency u to be

$$u = \frac{(x_n - x_m)}{c} \xi = \frac{D_{mn}}{c} \xi \quad (2.23)$$

where D_{mn} is taken to be the distance between sensors m and n . Equation (2.22) is known as the Van Cittert-Zernike theorem. This relationship shows that the output of the correlator shown in Fig. 2.1 is a measurement of a sample of the spatial frequency spectrum of the mutual coherence function of the source. We can collect several samples of this function by correlating the signals of receivers that are separated by various distances D . These measurements are commonly referred to as the complex *visibilities* of the source. The source intensity distribution can be found via the inverse Fourier transform relation

$$\hat{\Gamma}(\xi, f_o) = \int_{-\infty}^{\infty} \hat{\gamma}(u, f_o) e^{j2\pi f_o u} du. \quad (2.24)$$

This is the basis for the interferometric image formation technique. Chapter 4 includes a discussion of performance issues related to this type of imaging, including the effect of finite bandwidth and the variance associated the correlation estimates.

2.1.2 Array design and image deconvolution

Given that the intensity distribution $P(\xi)$ has Fourier transform $\overline{P(u)}$ and that the spatial frequency sampling function S has unity value at every measured spatial frequency and is zero elsewhere, the measured visibilities for the monochromatic case can be expressed as

$$V(u) = S(u) \overline{P(u)}. \quad (2.25)$$

The resulting brightness map is given by the Fourier transform of Eq. (2.25) and can be expressed as

$$M(\xi) = P(\xi) * \overline{S(u)} = P(\xi) * B_0(\xi). \quad (2.26)$$

Equation (2.26) illustrates that the brightness map obtained is the convolution of the actual brightness with the beam B_0 . This beam is termed the “dirty” beam and the resultant map M is called the “dirty” map. Generally, the term *beam* is synonymous with the response of the imaging system to a point source located at 0° . We desire to find the actual brightness $P(\xi)$,

so a method of deconvolving $P(\xi)$ and $B_0(\xi)$ is needed. In general, one might deconvolve two such functions by moving to the Fourier domain and dividing $V(u)$ by $S(u)$ to find $\overline{P(u)}$ and thus $P(\xi)$. However, $S(u)$ most always has regions where the function is equal to zero, where no visibility measurements are made. Thus, dividing the Fourier transforms of the functions is not possible. Another method of deconvolution is needed. In the radio astronomy community, the most widely used algorithm for image deconvolution of this type is the method CLEAN.

The method CLEAN was developed by Högbom in 1974 and has since spawned many variants [11]. The basic idea of this nonlinear algorithm is to iteratively find point source components in the “dirty” map and then convolve these with a “clean” beam to produce a “clean” map. Even though the method is based on a collection of point sources, it does, in general, work on extended sources. The method can be summarized as follows:

1. Locate the maximum in the “dirty” map and determine its amplitude A .
2. Convolve the “dirty” beam with a point source at the location found in (1) of amplitude γA , where γ is known as the loop gain.
3. Subtract the result of (2) from the “dirty” map.
4. Repeat steps (1)-(3) until the residual (what is left of the “dirty” map after a number of subtractions as in (3)) is sufficiently small.
5. Convolve the collection of point sources found with a “clean” beam which is normally chosen to be Gaussian. This gives a smooth “clean” map.
6. Add the residual to the “clean” map.

CLEAN performs an interpolation-like function on the visibility data, attempting to fill in those areas in the spatial frequency domain that are unmeasured. As CLEAN is a nonlinear method, a complete mathematical analysis is not an easy task. However, Schwarz has given some important conclusions [12], [13]. Schwarz shows that if the following three conditions hold, then CLEAN will converge to a solution:

1. The “dirty” beam must be symmetric.
2. The Fourier transform of the “dirty” beam must be nonnegative.
3. There must be no spatial frequencies present in the “dirty” map which are not also present in the “dirty” beam.

If these conditions hold, which they almost always do in practice, and if the number of CLEAN components does not exceed the number of independent spatial frequencies

measured, then CLEAN converges to the solution that is the least squares fit of the Fourier transforms of the delta-function components to the measured visibilities. The method CLEAN is widely used and is also a component in many other image formation methods, such as self-calibration which is discussed in Chapter 5.

The choice of array will determine the spatial frequency sampling function $S(u)$, and therefore affects the quality of the image. One would like to have an array that uniformly covers a given range of spatial frequencies. One can consult [8] for a discussion of array design. In general, most interferometric imaging systems use a minimally redundant array that covers all multiples of a unit spacing up to the maximum extent of the array. The minimal redundancy requirement is more important for radio astronomy, where it allows for uniform coverage with the least number of elements, where each element can be of enormous cost. For our acoustic imaging case, we are usually content with an array that provides uniform coverage. Redundancies can be of use for certain aspects of image recovery. The resolution obtainable with a given array is proportional to the aperture size. The minimum baseline in the array gives the maximum view angle that can be achieved. The designs of the arrays used in this work can be found in Chapter 3.

2.1.3 Multifrequency synthesis imaging

Multifrequency synthesis is an extension to the monochromatic aperture synthesis theory outlined in the previous section. It was developed relatively recently in the radio astronomy field by Conway, Cornwell, and Wilkinson [14]. Gaps in the spatial frequency coverage of an interferometric array introduce artifacts into the resultant image. The multifrequency method serves to fill in these gaps and improve image quality. Though developed with the aim to improve the image quality possible with sparse interferometric arrays, this method does also show promise in when applied to the acoustic imaging problem.

As defined in Eq. (2.23), the spatial frequency variable u is proportional to the product of the frequency and the sensor separation. This relationship gives us a logical connection, or space-frequency equivalency, that one can use to increase the density of samples in the spatial frequency regime with a given array. This space-frequency

equivalency was discussed by Swenson and Mathur [15], and is the basis of multifrequency synthesis. By measurement of the visibilities associated with a given array of sensors at a series of several frequencies, one can sample a greater portion of the spatial frequency spectrum. When using multiple frequencies, one must be concerned about how the source emissions change with frequency. For most astronomical sources, Conway and Saul show that if frequencies are chosen within a $\pm 12.5\%$ range (with respect to the central analysis frequency) spectral errors are generally below the thermal noise [16]. The errors associated with spectral variation can be analyzed as follows [16], [17]. Let the intensity distribution as a function of direction and frequency be given by $P(\xi, f)$. The combination of visibilities measured at each frequency can be written as

$$V_T(u) = \sum_i S(u, f_i) \overline{G(f_i)^{-1} P(u, f_i)} \quad (2.27)$$

where the overline indicates the spatial Fourier transform, S is the sampling function of the visibility measurements (unity at spatial frequencies where a measurement is made), and G is an optional scaling factor that is a function of frequency. The intensity can be expressed as a Taylor series expansion with respect to the frequency dependent function $F(f) = f/f_o$ which gives

$$P(\xi, f) = P(\xi, f_o) + f_o \left. \frac{\partial P(\xi, f)}{\partial f} \right|_{f=f_o} \left(\frac{\Delta f}{f_o} \right) + \frac{1}{2} f_o^2 \left. \frac{\partial^2 P(\xi, f)}{\partial f^2} \right|_{f=f_o} \left(\frac{\Delta f}{f_o} \right)^2 + \dots \quad (2.28)$$

where $\Delta f = f - f_o$. If this expression is inserted into Eq. (2.13) and if we neglect G for now, we obtain the new expression for the combined visibilities:

$$V_T(u) = \overline{P(\xi, f_o)} \sum_i S(u, f_i) + f_o \left. \frac{\partial \overline{P(\xi, f)}}{\partial f} \right|_{f=f_o} \sum_i \left(\frac{\Delta f}{f_o} \right) S(u, f_i) + \frac{1}{2} f_o^2 \left. \frac{\partial^2 \overline{P(\xi, f)}}{\partial f^2} \right|_{f=f_o} \sum_i \left(\frac{\Delta f}{f_o} \right)^2 S(u, f_i) + \dots \quad (2.29)$$

Taking the Fourier transform of Eq. (2.14) we obtain the resultant composite intensity map

$$M_{comp} = P * B_0 + f_o \frac{\partial P}{\partial f} * B_1 + \frac{1}{2} f_o^2 \frac{\partial^2 P}{\partial f^2} * B_2 + \dots \quad (2.30)$$

where the derivatives are still evaluated at f_o and the generalized beams B_n are given by

$$B_n(\xi) = \sum_i S(u, f_i) \left(\frac{\Delta f}{f_o} \right)^n. \quad (2.31)$$

We can now discuss the spectral errors that result from the second and third terms of Eq. (2.30). The spectral beams given in Eq. (2.31) are shown in Fig. 2.2 for an example case with 18 frequencies located in the interval $\pm 12.5\%$ around a center frequency of 325 Hz. Note the relative amplitudes and the fact that the B_1 beam (and all odd beams) has no central peak. As long as the first-order derivative in Eq. (2.30) remains relatively small, the first-order spectral errors will be small, as the peaks of B_1 are 1/100 the peak in B_0 . If the intensity is dominated by a single bright component, then one can effectively remove the spectral errors by using the appropriate $G(f)$ function in Eq. (2.27) to make $\partial P / \partial f$ equal to zero. However, if there are multiple components with nearly equal brightness, this scheme will not work, and the first-order spectral errors will affect the image at a level of about 1% of the peak brightness. Second-order effects can also be effectively eliminated for the case of a dominant component with use of the G function. For those cases where there is no dominant component, the second-order effect can mainly be thought of as adding a small bias to the brightness due to the small central peak (1/160) and low sidelobe levels (1/2500) of B_2 .




Fig. 2.2: The spectral beams B_0 (a), B_1 (b), and B_2 (c) for the case of 18 frequencies in $325 \text{ Hz} \pm 12.5\%$.

2.2 Wideband (2-D) Imaging

2.2.1 Tomographic solution method for wideband imaging

This section concentrates on the theory involved in computing images of the power spectrum versus angle. As will be discussed, this two-dimensional image can be computed from wideband correlation data through a mathematically tomographic relationship. Early work on this type of imaging was done by Biraud et al. [2], Chiollaz et al. [3], and Escudié [4], as mentioned in Chapter 1.

The derivation of the imaging relation follows in a similar fashion to that in Section 2.1.1. The source function is once again assumed to follow the analytic signal formulation discussed earlier. However, the truncated function notation will be dropped in this case. The physical situation at hand is the same as that shown in Fig. 2.1, except that the narrowband

filter is absent in this case. We start by expressing the signals received at sensors located at positions x_n and x_m as

$$y_n(t) = \int_{-\infty}^{\infty} S\left(\xi, t - \frac{x_n}{c}\xi\right) d\xi \quad (2.32)$$

$$y_m(t) = \int_{-\infty}^{\infty} S\left(\xi, t - \frac{x_m}{c}\xi\right) d\xi. \quad (2.33)$$

Let us now examine the cross correlation of these signals, defined to be

$$r_{mn}(\tau) = E[y_m(t)y_n^*(t-\tau)] \quad (2.34)$$

where once again we enforce the stationarity condition. Plugging Eqs. (2.32)-(2.33) into Eq. (2.34) we obtain

$$r_{mn}(\tau) = \int_{-\infty}^{\infty} d\xi_1 \int_{-\infty}^{\infty} d\xi_2 E[S(\xi_1, t - \frac{x_m}{c}\xi_1)S^*(\xi_2, t - \frac{x_n}{c}\xi_2 - \tau)]. \quad (2.35)$$

As before, we need to specify a form for the mutual coherence function of the source. In this case we make the following assumption:

$$E[S(\xi_1, t_1)S^*(\xi_2, t_2)] = \Gamma(\xi_1, t_1 - t_2)\delta(\xi_1 - \xi_2) \quad (2.36)$$

where stationarity and spatial incoherency have been enforced. Applying this assumption to the correlation expression in Eq. (2.35), we obtain the relation

$$r_{mn}(\tau) = \int_{-\infty}^{\infty} \Gamma\left(\xi, \frac{(x_n - x_m)}{c}\xi + \tau\right) d\xi. \quad (2.37)$$

Equation (2.37) gives us the relationship between the wideband cross correlation of a pair of sensor signals and the mutual coherence function of the source. If one examines the form of Eq. (2.37), one sees that it closely resembles the relationship seen for a standard parallel beam tomographic imaging problem. The measured cross correlations are composed of line-integrals through the unknown mutual coherence function of the source. Figure 2.3 illustrates this relationship. The process of inverting this relationship has been well studied and there several different algorithms available to solve the problem. A good overview of tomographic imaging and the basic solution methods can be found in Haykin [18]. The projection slice theorem gives us one way to solve the problem. This theorem states that the one-dimensional Fourier transform of a projection taken at a given angle through the

unknown object or function is equal to a slice of the two-dimensional Fourier transform of the function at that same angle. By measuring the cross correlations of all sensor pairs in the array we are measuring the projections at different angles. Because the angle at which the projection is made is a function of the distance between sensor pairs, we see that we have a limited number and range of angles. This fact will limit the quality of the images we can recover. Once all the projections are measured, their Fourier transforms are computed and then arranged in the two-dimensional Fourier plane. Because the data are arranged in polar fashion in the Fourier plane, the data must be interpolated to lie on a Cartesian grid before the inverse Fourier transform can be applied to recover the unknown function. Results of this imaging process will be discussed in Chapter 4. Another common method of solving for the unknown object or function in tomographic imaging is the family of algebraic reconstruction techniques (ART). With ART, the unknown function is pixelated into a grid. Then, large linear equations are constructed by summing the pixel contributions for each point in the measured projections. This set of equations is usually very large and can only be solved using iterative methods [19], [20]. Results from this type of solution will also be shown in Chapter 4.




Fig. 2.3: Diagram illustrating the tomographic relationship between the mutual coherence function of the source and the array cross correlation measurements.

2.2.2 White noise imaging

Another wideband case to consider is the case where one assumes that the source radiation is everywhere white. We then would be solving for the intensity as a function of angle, similar to the narrowband case. Starting at Eq. (2.35), this time we make the following white-noise assumption on the nature of the source

$$E[S(\xi_1, t_1)S^*(\xi_2, t_2)] = P(\xi)\delta(\xi_1 - \xi_2)\delta(t_1 - t_2) \quad (2.38)$$

where $P(\xi)$ is the source intensity. Using this relation in Eq. (2.35), we obtain the simple relationship

$$r_{mn}(\tau) = P\left(-\frac{\tau c}{D_{mn}}\right). \quad (2.39)$$

Equation (2.39) tells us that to construct an image of the intensity of the source, we need only to arrange the cross correlation measurements. It is likely that for some positions in the intensity map, there will be more than one available coherence measurement. Therefore, we need to find the best way of combining these multiple measurements into a final image. To do this we can use linear estimation theory [21]. For each $\xi = -\tau c/D_{mn}$ where there are more than two measurements of $r_{mn}(\tau)$, we arrange those measurements into a column vector \vec{g} .

We can then construct the statement

$$\vec{g} = \mathbf{H}x + \vec{n} \quad (2.40)$$

where x is the actual value of the map, \mathbf{H} is a column vector of all ones, and \vec{n} is a noise vector representing the difference between the actual map value and each measurement. We assume that

$$\begin{aligned} E[\vec{n}] &= \vec{0} \\ E[\vec{n}\vec{n}^T] &= \mathbf{R}_n. \end{aligned} \quad (2.41)$$

Using the Gauss-Markov theorem we find that the minimum-variance, unbiased, linear estimate of x is

$$\hat{x} = (\mathbf{H}^T \mathbf{R}_n^{-1} \mathbf{H})^{-1} \mathbf{H}^T \mathbf{R}_n^{-1} \vec{g}. \quad (2.42)$$

The error in the estimate

$$\varepsilon_x = E[(x - \hat{x})^2] \quad (2.43)$$

can be found to be

$$\varepsilon_x = \left(\mathbf{H}^T \hat{\mathbf{R}}_n^{-1} \mathbf{H} \right)^{-1} \mathbf{H}^T \hat{\mathbf{R}}_n^{-1} \mathbf{R}_n \hat{\mathbf{R}}_n^{-1} \mathbf{H} \left(\mathbf{H}^T \hat{\mathbf{R}}_n^{-1} \mathbf{H} \right)^{-1}. \quad (2.44)$$

One might logically guess that simply averaging the measurements would be appropriate. In that case $\hat{\mathbf{R}}_n$ would be equal to the identity matrix. The general case would be $\hat{\mathbf{R}}_n = \mathbf{R}_n$. Expressions for the error in each case are shown below.

$$\begin{aligned} \hat{\mathbf{R}}_n = \mathbf{R}_n &\Rightarrow \varepsilon_x = \left(\mathbf{H}^T \mathbf{R}_n^{-1} \mathbf{H} \right)^{-1} = \left(\sum_{i,j} \mathbf{R}_n^{-1}(i,j) \right)^{-1} \\ \hat{\mathbf{R}}_n = \mathbf{I} &\Rightarrow \varepsilon_x = \left(\mathbf{H}^T \mathbf{H} \right)^{-1} \mathbf{H}^T \mathbf{R}_n \mathbf{H} \left(\mathbf{H}^T \mathbf{H} \right)^{-1} = \frac{1}{(\dim[\mathbf{R}_n])^2} \sum_{i,j} \mathbf{R}_n(i,j) \end{aligned} \quad (2.45)$$

These two errors can, in general, be different, giving motivation for finding \hat{x} . In order to compute the estimate, we need to compute \mathbf{R}_n . This can be done using the time series data from each sensor. Let us assume that we form the correlation estimate

$$\hat{r}(\tau_1, d_{mn}) = \frac{1}{M} \sum_{p=1}^M y_m(p) y_n(p - \tau_1) \quad (2.46)$$

where τ_1 is a discrete delay (based on the sampling rate used) and M is the total number of samples in the estimate. Also, let $y_n(p)$ denote the p^{th} sample of the signal received at a sensor located at position x_n . The signals are assumed to be real in this case. If we then collect all estimates for which τ/d is constant, we form the vector $\hat{\mathbf{g}}$, that has elements

$$\hat{r}(\tau_i, d_i) \quad \frac{\tau_i}{d_i} = \text{const.} \quad (2.47)$$

The vector \mathbf{g} , containing the true correlation values, can be formed in similar fashion. We then can express the correlation estimate (measurement) as the sum of the actual correlation plus noise:

$$\hat{r}(\tau, d) = r(\tau, d) + n. \quad (2.48)$$

This leads to the following expression for the covariance matrix of the noise:

$$\mathbf{R}_n = E[\hat{\mathbf{g}}\hat{\mathbf{g}}^T] - \mathbf{g}\mathbf{g}^T. \quad (2.49)$$

The received signals are white Gaussian noise since they are composed of a sum of source signals that are white. We then can write

$$E[y_n(p)y_n(q)] = \delta(p-q)N_o \quad (2.50)$$

where N_o is the total source flux. Now let us look at the expression for one element of the covariance matrix, given by

$$E[\hat{r}(\tau_1, d_{mn})\hat{r}(\tau_2, d_{ab})] - r(\tau_1, d_{mn})r(\tau_2, d_{ab}) \quad (2.51)$$

where we have

$$\frac{\tau_1}{d_{mn}} = \frac{\tau_2}{d_{ab}}. \quad (2.52)$$

This can be expanded using the fourth-order moment theorem for Gaussian processes [22].

The result consists of two terms as shown below.

$$\begin{aligned} & \frac{1}{M^2} \sum_p \sum_q E[y_j(p)y_l(q)]E[y_n(p-\tau_1)y_b(q-\tau_2)] \\ & + \frac{1}{M^2} \sum_p \sum_q E[y_m(p)y_b(q-\tau_2)]E[y_n(p-\tau_1)y_a(q)] \end{aligned} \quad (2.53)$$

The way in which the two terms in Eq. (2.53) are evaluated will depend strongly on the configuration of the array used. For each term one needs to consider several cases. These cases are determined largely by whether or not a sensor appears more than once among the four in the expression and, if so, which ones are equal. Using this approach, all elements of \mathbf{R}_n can be computed from the sensor data. That estimate of \mathbf{R}_n can then be used in Eq. (2.42) to find the optimal data for the given place in the reconstructed intensity map. Results of using this method are presented in Chapter 4.

3. FIELD EXPERIMENTS

In this chapter, field experiments used to collect real acoustic imaging data will be discussed. A series of completed experiments will be summarized. Measurements dealing with properties of the atmosphere that were made using the collected data will also be presented to provide a better understanding of the medium through which we are trying to image.

3.1 Description of Experiments and Data

A series of 20 acoustic imaging data collection experiments were conducted. The sites used were the Bondville and Monticello Road field sites, belonging to the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. The majority of the experiments took place at the Monticello site, which is located among cornfields in a rural area. The ground was flat and covered with grass of varying height. The layout of the experiments at the Monticello site is shown in Fig. 3.1. The small building, which was about 50 m behind the loudspeakers, should have no appreciable effect on the data, due to the fact the loudspeakers were at least marginally directional in the frequency range that was used. Large high-power loudspeakers broadcasting band-limited white noise were used as sources. The power spectrum of the noise signal input into the loudspeakers was 10 dB down at 100 and 5500 Hz. An array of eight microphones received the source emissions in the far field, and their signals were recorded on a multichannel digital recorder. The microphones consisted of Brüel and Kjær model 4145 1-in microphone cartridges, model 2639 preamplifiers and model 2204 power supplies. The digital recorder was a Sony PC-216. A sampling rate of 12 kHz was used for each channel. In addition to the array microphone outputs, the outputs of an anemometer and of a microphone placed near one of

the sources were also recorded for some experiments. Other environmental parameters were also kept, such as temperature, humidity, and the condition of the ground. For each experiment, there were multiple recordings, each with a different source configuration. Due to the fact that there were only two loudspeakers available, no more than two sources could operate at once. Table 3.1 summarizes the experiments, giving approximate conditions and configurations. In general, the sources were located from 0.7 to 1.5 m above the ground while the array was generally positioned 0.4 m above the ground.

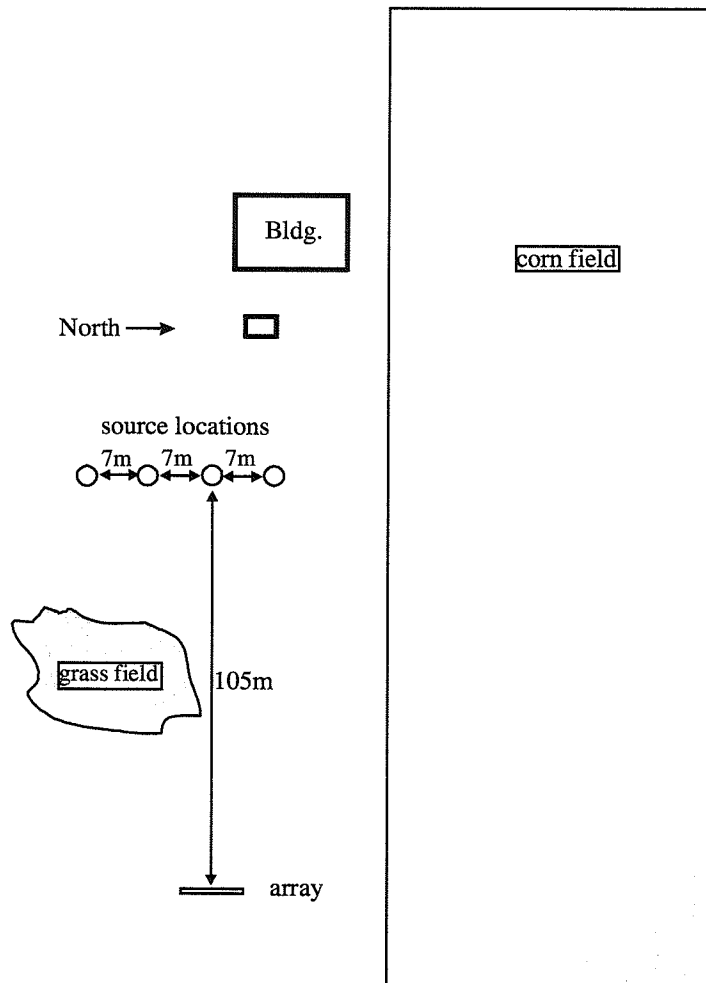


Fig. 3.1: Layout of the experiments at the Monticello Road field site.

There were three different arrays that were used in the experiments, referred to as types 1, 3, and 4. Types 1 and 3 were 10-m-long eight-element arrays with minimal redundancy and unit spacings of 0.4349 and 0.5 m, respectively. Array type 4 was also a minimally redundant array, but with a unit spacing of 0.2125 m and a total length of 4.25 m. All three arrays were complete in their spatial frequency coverage, meaning that their set of spatial frequency samples had no gaps in the Nyquist sense. The three arrays are shown in Fig. 3.2.




Fig. 3.2: The three arrays used in the experiments. For array 1 the base separation was 0.435 m and for arrays 3 and 4 it was 0.5 and 0.2125 m, respectively. Each array contains all multiples of the unit separation from 1 to 20. Arrays 3 and 4 are simply scaled versions of each other.

Table 3.1: Summary of completed experiments

Experiment Date	Time	Site	Number of Configurations	Wind Speed and Direction	Array Type Used	Weather	Temp. and Dew pt.
15SEP95	8-9am	Bondville	1	10 mph, E	1	partly cloudy	60
23OCT95	9-11am	Bondville	3	10-20 mph, SSE	1	partly cloudy	60
9FEB96	12-3pm	Bondville	12	5-10 mph, SW	1	mostly sunny	56
17APR96	8-11am	Bondville	12	10-15 mph, SE	1,3	mostly sunny	42-54
05JUN96	10am-1pm	Monticello	5	5-10 mph, SSE	3	mostly sunny	70-75
16JUL96	8am-12pm	Monticello	5	3-5 mph, WSW	3	partly sunny, humid	74-82
13AUG96	8am-12pm	Monticello	12	0-3 mph, NW	3	mostly sunny, humid	74-85
18OCT96	11am-2pm	Monticello	3	10-20 mph, NW	3	partly sunny	74
06MAR97	9am-1pm	Monticello	2	15-25 mph, WNW	3	Sunny	32
04JUN97	11am-3pm	Monticello	6	15-25 mph, ENE	4	mostly cloudy	70
18JUL97	10-11am	Monticello	3	0-5 mph, SW	4	mostly sunny	85
24JUL97	4-5pm	Monticello	3	0-5 mph, NE	4	mostly sunny	85, 69
25JUL97	10-11am	Monticello	3	7-8 mph, S	4	mostly sunny	79, 73
05AUG97	12-1pm	Monticello	3	6-12 mph, NNE	4	Sunny	73, 54
14AUG97	2-3pm	Monticello	3	10-12 mph, SSE	4	mostly cloudy	79, 65
21AUG97	2-3pm	Monticello	3	10-18 mph, NW	4	partly sunny	70, 57
27AUG97	9-10am	Monticello	3	1-5 mph, NW	4	heavy fog	72, 74
03SEP97	11am-12pm	Monticello	3	8-15 mph, NE	4	mostly sunny	63, 55
04SEP97	11am-12pm	Monticello	3	1-6 mph, NE	4	mostly sunny	65, 55
30OCT97	11am-12pm	Monticello	3	8-19 mph, S	4	mostly sunny	50, 35

Significant ground impedance effects have been observed in some of the data collected. The typical manifestation of this effect is the appearance of a significant amount of excess attenuation in a band of frequencies centered at about 500 Hz. This effect is to be expected with this type of measurement and is in agreement with similar measurements and theoretical calculations in the literature [23]-[25]. Figure 3.3, taken from [23], shows the predicted excess attenuation for various distances of propagation over grass for source and

receiver heights of 1.8 and 1.5 m, respectively. Figure 3.4 shows typical spectra of the transmitted and received signals for the experiments conducted for this thesis. The corresponding overall excess attenuation, based on the spectra in Fig. 3.4, is shown in Fig. 3.5. Generally, close agreement was found between the measured excess attenuation and that discussed in the literature. The presence of turbulence can significantly alter the amount of excess attenuation for a given propagation situation [24].






Fig. 3.3: Excess attenuation versus frequency and propagation distance for source and receiver heights of 1.5 and 1.8 m, respectively. Taken from [23].



(a)



(b)

Fig. 3.4: Typical measurements of the power spectra of the (a) source used in the experiments and (b) a signal received at the array. The source microphone was placed 1 m in front of the loudspeaker for the measurement.




Fig. 3.5: Excess attenuation computed from the spectra in Fig. 3.4 taking into account geometrical spreading and atmospheric absorption.

3.2 Characterization of the Atmosphere

In this section, a basic review of atmospheric turbulence will be given followed by a description of several atmospheric measurements that were made using the experimental data described above. The measurements include the calculation of coherence loss, the computation of wave structure functions, the extraction of turbulent strength and scale lengths from coherence data, and the computation of the time-varying propagation transfer function of the atmosphere.

3.2.1 A review of atmospheric turbulence theory

The atmosphere affects acoustic propagation in many ways. A discussion of some of these mechanisms can be found in [23], [26]. Besides the basic mechanisms, such as atmospheric absorption, the most destructive phenomenon for imaging is turbulence.

Turbulent air motion is generated by differential heating and cooling of the Earth due to the sun and the diurnal cycle. This generates large-scale inhomogeneities in the temperature of the air, which are broken into smaller inhomogeneities by the wind. The wind has complex motion and is generally random. Each of these randomly distributed inhomogeneities contains air having a characteristic temperature. As the index of refraction for sound is sensitive to temperature, we see that the atmosphere has a random index of refraction. The statistics of the index of refraction in the atmosphere are of great interest.

It was Kolmogorov who came up with the first comprehensive theories about the statistical nature of turbulence. Kolmogorov suggested that the large-scale inhomogeneities are continually broken down into smaller and smaller eddies until the point where the remaining kinetic energy is dissipated as heat [27]. According to this theory then, there is a range of scale sizes present at any given time in the atmosphere. Kolmogorov further suggested that the smaller-scale turbulent motions are homogeneous and isotropic. This gave rise to an expression for the power spectral density of the variation in the index of refraction, $\Phi_n(\vec{\kappa})$, where $\vec{\kappa}$ is the spatial wavenumber vector with orthogonal components κ_x , κ_y , and κ_z . The power spectrum $\Phi_n(\vec{\kappa})$ can be considered a measure of the abundance of turbulent eddies having scale sizes $l_x = 2\pi/\kappa_x$, $l_y = 2\pi/\kappa_y$, and $l_z = 2\pi/\kappa_z$. The spectrum is broken down into three regions, demarcated by the outer and inner turbulence scales, L_0 and l_0 . The shape and regions of the spectrum are shown in Fig. 3.6 and the characteristics of the regions are discussed in [28]. Near the ground, the outer scale L_0 can be approximated as being on the order of the height above the ground. The inner scale is on the order of a few millimeters. The region containing scales larger than the outer scale (small κ) is called the source or input range and is home to those large-scale eddies which break down according to Kolmogorov's theory. Turbulence in this region is not well understood and is generally thought to be inhomogeneous and anisotropic. Unfortunately, it is the eddies in this regime that most affect acoustic propagation and the imaging process [29]. The region between the outer and inner scales is called the inertial subrange, and is the regime over which Kolmogorov's expression for $\Phi_n(\vec{\kappa})$ (and those of others) is valid. The turbulence in this range is generally assumed to be homogeneous and isotropic. The inertial subrange is the

most well understood. The region of very small scales (large κ) is called the dissipation or viscous subrange. It is in this region where the very small eddies are dissipated as heat energy. For the inertial subrange, Kolmogorov theory predicts the following form for the spectrum:

$$\Phi_n(\bar{\kappa}) = 0.033 C_n^2 \kappa^{-11/3} \quad (3.1)$$

where C_n^2 is the structure constant of the index of refraction. The structure constant varies with local meteorological conditions and with height. Several models have been proposed, but a general consensus has not yet been met [30], [31]. For mathematical convenience, the three regions are sometimes combined into one relation with the form

$$\Phi_n(\bar{\kappa}) = \frac{0.033 C_n^2}{(\kappa^2 + \kappa_0^2)^{11/6}} e^{-\kappa^2/\kappa_m^2} \quad (3.2)$$

where $\kappa_0 = 2\pi/L_0$ and $\kappa_m = 5.92/l_0$. This form is often called the Von Karman spectrum [30]. It should be noted that even though Eq. (3.2) describes all three regions, its value in the source region should be taken as only approximate as the turbulence there depends strongly on local conditions. Similarly, the Kolmogorov spectrum in Eq. (3.1) should only be used for the inertial subrange.




Fig. 3.6: Diagram showing the general shape of the turbulence spectrum and the three subranges. Taken from [31].

It is the structure function of the fluctuating index of refraction that most clearly shows the influences of turbulence on the performance of an imaging system. Let us assume that the random index of refraction in the atmosphere can be broken down into the sum of the mean index and a randomly fluctuating index as

$$n(\vec{r}, t) = n_0 + n_1(\vec{r}, t) \quad (3.3)$$

where \vec{r} is the three-dimensional position vector and t is time. The *structure function* of the homogeneous index of refraction is defined to be

$$D_n(\vec{r}) = E[(n_1(\vec{r}_1) - n_1(\vec{r}_1 + \vec{r}))^2]. \quad (3.4)$$

One can express $D_n(\vec{r})$ as a function of the power spectral density, $\Phi_n(\vec{\kappa})$, as

$$D_n(\vec{r}) = 2 \int d\vec{\kappa} [1 - \cos(\vec{\kappa} \cdot \vec{r})] \Phi_n(\vec{\kappa}). \quad (3.5)$$

If we further assume that the index of refraction fluctuations are isotropic, the structure function then only depends on the scalar variable $r = |\vec{r}|$ and one can write

$$D_n(r) = 2 \int d\kappa \Phi_n(\kappa) \left[1 - \frac{\sin(\kappa r)}{\kappa r}\right]. \quad (3.6)$$

For straight-line propagation through atmospheric turbulence over a distance r , the log-amplitude and phase structure functions are commonly defined as

$$D\chi(r, \rho) = E[(\chi(r + \rho) - \chi(r))^2] \quad (3.7)$$

and

$$D\phi(r, \rho) = E[(\phi(r + \rho) - \phi(r))^2] \quad (3.8)$$

respectively, and where ρ is the separation along a plane perpendicular to the direction of propagation, χ is the log amplitude, and ϕ is the phase [32], [33]. The log-amplitude fluctuations typically suffer from saturation, wherein they do not increase with propagation distance beyond some characteristic distance [32]. The structure functions defined above can also be written as

$$\frac{1}{2} D_\chi(r, \rho) = E[\chi^2] - B_\chi(\rho) \quad (3.9)$$

$$\frac{1}{2} D_\phi(r, \rho) = E[\phi^2] - B_\phi(\rho) \quad (3.10)$$

where B_χ and B_ϕ are the covariances of the log-amplitude and phase fluctuations and where the mean square log-amplitude and phase fluctuations are given by

$$E[\chi^2] = E\left[\left(\ln\left(\frac{A}{A_0}\right)\right)^2\right] \quad (3.11)$$

$$E[\phi^2] = E[(\phi - \phi_0)^2]. \quad (3.12)$$

In the above expression A , A_0 , ϕ , and ϕ_0 are the amplitude and phase in the absence and presence of turbulence, respectively. The *wave structure function* $D(r, \rho)$ is defined as the sum of the log amplitude and phase structure functions. The mutual coherence function, defined to be the correlation function of the complex fields in a plane perpendicular to the direction of propagation, can be written as a function of the wave structure function as

$$\Gamma(r, \rho) = e^{-\frac{1}{2}D(r, \rho)} = e^{-\frac{1}{2}(D_A + D_\phi)}. \quad (3.13)$$

Measurements of the mutual coherence function, and thus of the wave structure function, were made using the data gathered in the experiments described in the last section. The results will be discussed in the next section.

3.2.2 Coherence loss measurements

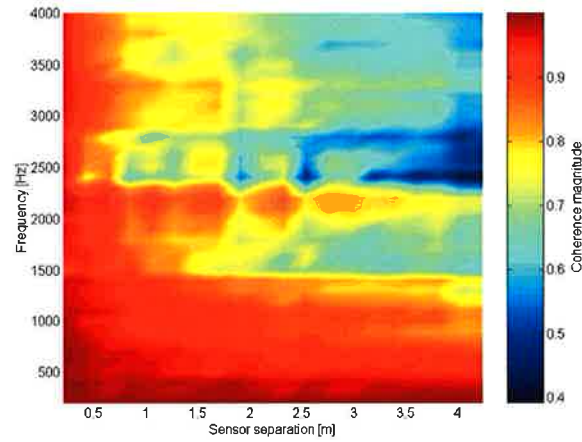
For all the datasets shown in Table 3.1, recordings were made with only one active source. With these data, one can measure the mutual coherence function in the array plane. These measurements can also be used to compute the wave structure function, which will be discussed in the next section. A Gaussian turbulence model can be fit to the coherence data in order to extract turbulence strength and scale length parameters. This will be discussed in Section 3.2.4.

The mutual coherence function, as shown in Eq. (3.13), can be computed from field data recordings where only one source was active. The mutual coherence function is constructed by measuring the visibilities on all baselines in the array. The magnitude of each visibility is a sample of the mutual coherence function. In the absence of turbulence, the wave impinging on the array should be perfectly coherent. However, with the addition of turbulence, the wave suffers a loss of coherence that is an increasing function of the baseline length. This coherence loss is important for the performance of an imaging system, as it reduces the resolution possible [2], [32]. The turbulence, by introducing the coherence loss on the longer baselines, essentially reduces the size of the array.


Representative examples of measurements of the mutual coherence function are shown in Fig. 3.7. The magnitude of the measured coherence as a function of sensor separation and frequency is shown for the three average wind speed cases of 2.4 , 10, and 14 mph in (a), (b), and (c) respectively. The coherence measurements were made using seven seconds of data over which the wind speed was relatively constant. As the wind speed increases, and thus the strength of the turbulence increases, ones sees significant loss of coherence at frequencies above a few hundred hertz. In Fig. 3.7(c) one sees little or no coherence above 300 Hz. These measurements agree well with those in the literature [32]-[34]. Figure 3.8 shows a comparison of the point source images that would result from the measured coherence data in Fig. 3.7(a), (c) at 800 Hz. The images were normalized to have unity peak power, and perfect phase measurements have been assumed. Note the slight widening of the mean beam response and the significant rise in sidelobe response. Aberrations in the measured phase would add additional distortion to the point source response.

3.2.3 Computations of the wave structure function


In this section, a discussion is presented on the results computing the wave structure function from coherence data like that discussed in the last section. The relationship between the coherence function and the wave structure function is given by Eq. (3.13). Representative results will be shown here. In Fig. 3.9, the computed wave structure functions at 800 Hz are shown for the three data sets shown in Fig. 3.7. The magnitude of the structure function increases with the wind speed, as one would expect. In the case of the 14-mph wind data, we are apparently seeing evidence of the outer scale, as the trace seems to have leveled off at about the 1-m scale. The outer scale would occur at a separation larger than those measured for the 10- and 2.4-mph cases. In Fig. 3.10, the wave structure functions for the 10-mph wind data are plotted as a function of frequency. Here we see the rollover associated with the outer scale in the 1.6- and 3.2-kHz traces. The computed structure functions verify the assumption that the acoustic waves are strongly affected by turbulent eddies in and near the outer scale.



(a)



(b)



(c)

Fig. 3.7: Examples of coherence measurements. The magnitude of the computed coherence (color coded according to the right of each example) is plotted versus sensor separation and frequency. The average wind speeds are (a) 2.4, (b) 10, and (c) 14 mph.




Fig. 3.8: Comparison of the resulting images based on the coherence measurements from Fig. 3.7 (a) and (c). Perfect phase measurement is assumed and the images have been normalized to have unity peak powers.

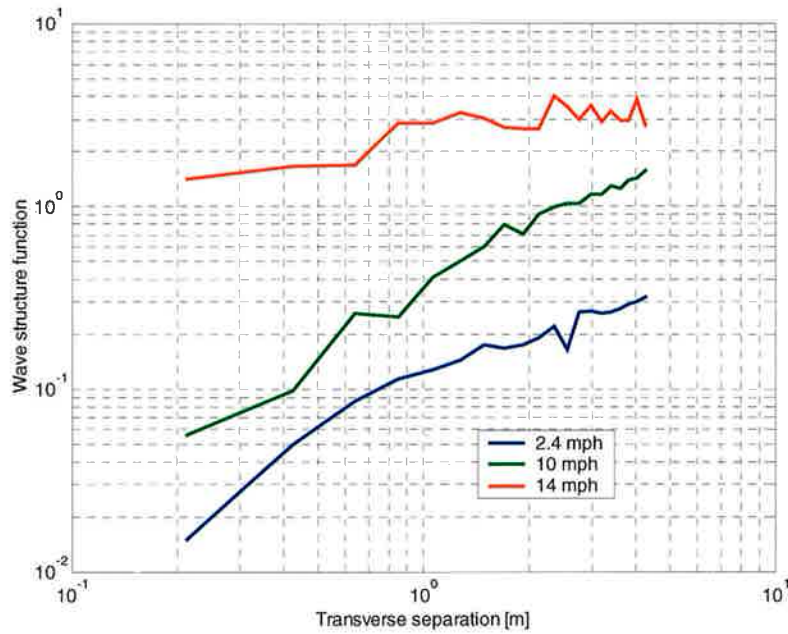


Fig. 3.9: Wave structure functions at 800 Hz, computed from the coherence data shown in Fig. 3.7 for average wind speeds of 2.4, 10, and 14 mph.

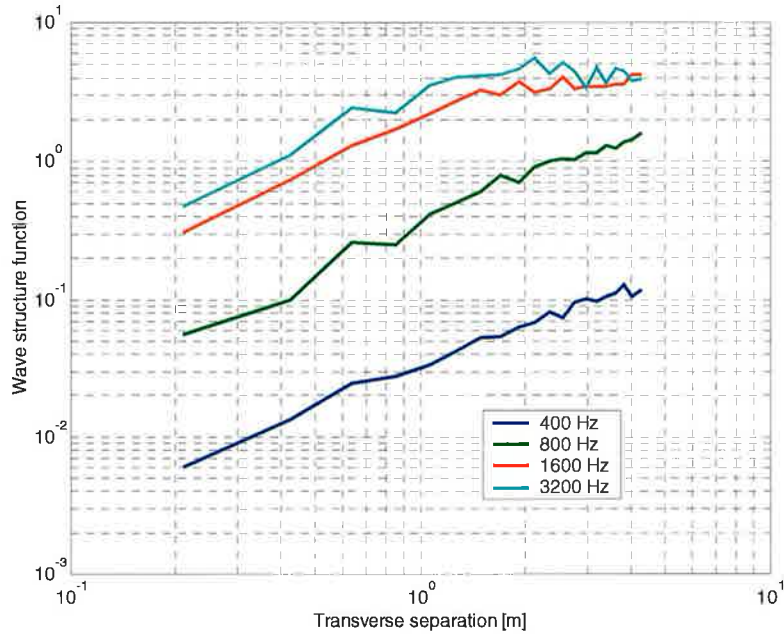


Fig. 3.10: Wave structure functions for the 10-mph wind case shown as a function of frequency.

3.2.4 Turbulence strength and scale parameter extraction

Using the coherence data mentioned in Section 3.2.2, turbulence strength and scale parameters can be estimated by applying a Gaussian turbulence model. Following the development in [32], the relatively simple theory of homogeneous and isotropic turbulence will be used. A Gaussian distribution is assumed for the spectrum of eddy sizes in the source or input subrange, since acoustic waves are most likely influenced by turbulent eddies having sizes equal to or larger than the so-called outer scale. This model implies that the fluctuating index is modeled as in Eq. (3.3). In this case, the spatial correlation of the refractive index will have the Gaussian form

$$E[n_1(\vec{r})n_1(\vec{r} + \rho)] = E[n_1^2]e^{-\rho^2/L^2} \quad (3.14)$$

where L is a measure of the scale of the turbulence and ρ is as was defined in Section 3.2.1.

Using Eqs. (3.9)-(3.13), along with the fact that for spherical waves

$$\frac{B_x(\rho)}{E[\chi^2]} = \frac{B_\phi(\rho)}{E[\phi^2]} = \frac{\Psi(\rho/L)}{\rho/L} \quad (3.15)$$

one obtains the following form for the mutual coherence function:

$$\Gamma(r, \rho) = \exp\left(-\sqrt{\pi} E[n_1^2] k^2 r L \{1 - [\Psi(\rho/L)/(\rho/L)]\}\right) \quad (3.16)$$

where k is the wave number, r is the propagation distance, and the function Ψ is given by

$$\Psi(\rho/L) = \int_0^{\rho/L} e^{-u^2} du. \quad (3.17)$$

The shape of this Gaussian spectrum is shown by the dotted line in Fig. 3.6. This model for the mutual coherence function was fit to the two-dimensional (frequency and space) measured coherence data discussed in Section 3.2.2. In order to account for the fact that the coherence does not fall to zero, but rather falls to a nonzero asymptote, two extra parameters were added to Eq. (3.17) to give

$$\Gamma'(r, \rho) = A_1 + A_2 \Gamma(r, \rho). \quad (3.18)$$

The least-squares parameter fits for each data segment were calculated. The results are shown in Figs. 3.11 and 3.12. The turbulence strength $E[n_1^2]$ is plotted in Fig. 3.11 as a function of the average wind speed for each of the data segments. The general trend displayed in Fig. 3.11 is that which one would expect, namely that the strength of the turbulence increases with increasing wind speed. The characteristic scale lengths extracted from the same coherence data, shown in Fig. 3.12, exhibit less organized behavior as a function of wind speed than the strength parameter. This can also be expected as the structure and size of eddies near the ground generally have a size on the order of the height above the ground. One would expect a slight decrease in the size of the eddies with increasing wind speed, and Fig. 3.12 does show a weak dependence.

3.2.5 Propagation transfer function calculations

Data collected with an additional microphone placed at the loudspeaker can be used to compute estimates of the time-varying propagation transfer function through the atmosphere using an adaptive filtering method. The method used is described in [35], but will be summarized here.




Fig. 3.11: Strength of turbulence $E[n_1^2]$ as calculated by fitting a Gaussian turbulence model to coherence data, versus average wind speed in mph.




Fig. 3.12: Characteristic scale length of the turbulence as extracted from coherence data by fitting a Gaussian turbulence model.

The sampled signals from the source microphone and a microphone in the array are input to an adaptive filter which is arranged in what is known as a system identification configuration. A good treatment of adaptive signal processing can be found in [36]. A block diagram of the system can be found in Fig. 3.13. The output of the adaptive filter is subtracted from the actual received signal to form the error signal. The error signal is fed back to the adaptive filter and is used to control the adaptation process. Several different adaptive filter algorithms were tested for this application, with the best results coming from the data-reusing least mean squares (DR LMS) and the recursive least square (RLS) lattice algorithms [37].




Fig. 3.13: Block diagram of the transfer function calculation system. The delay block Δ accounts for most of the propagation delay.

At any time instant, the transfer function is modeled as a linear time-invariant system. Accordingly, one must be aware of the limitations inherent to this type of representation, such as the inability to account for Doppler effects. The length of the finite impulse response (FIR) filter used will set the complexity of the transfer function that can be modeled.

The parameters of the adaptive algorithms should be chosen with care to ensure good performance. The length of the FIR filter must be chosen long enough to adequately represent the transfer function while not being so long as to lessen the ability of the adaptive algorithm to adjust the coefficients fast enough to keep up with the changing propagation conditions. A general guideline for choosing the maximum filter length is the fact that the transfer function should remain essentially unchanged over the period of time equivalent to

the length of the filter. The variability in the transfer function is related to the variability of the propagation conditions. Various filter lengths were tried and good results were obtained with 100 to 250 taps (using a 12-kHz sampling rate) depending on the specific propagation conditions (wind speed, etc.). The adaptation constant μ must also be chosen for good performance. The value of μ should be chosen according to guidelines used in the adaptive filtering literature [36]. The Δ delay block shown in Fig. 3.13 should be used to account for an expected minimum propagation delay. The adaptive filter can then account for the variations in the delay. The variation in the delay will be a function of the propagation distance and the wind speed fluctuations. Expected standard deviations of the propagation delay suggest that the filter lengths mentioned above are adequate to account for the variations [38].

A representative example of a computed transfer function is shown in Fig. 3.14. For this example, the average wind speed was 0.5 m/s. The data-reusing LMS algorithm was employed with a filter length of 250, a μ value of 0.7, and a constant delay Δ of 3550 samples using a sampling rate of 12 kHz. The filter coefficients representing the transfer function were averaged over intervals of 600 samples (1/20 s). Shown in Fig. 3.14 are the magnitude and phase of the average transfer function over 1 s of data. Note that one would expect 40 dB of attenuation due to geometrical spreading with the excess attenuation being caused by interactions with turbulence and the ground. In order to get a sense of the accuracy of the transfer functions computed using this method, it is illustrative to plot the estimated power spectrum of the received signal along with the power spectrum of the error signal. Just such a plot is shown in Fig. 3.15 for the same data that was used to compute Fig. 3.14. From Fig. 3.15 one sees that the model is accurate to about 2 kHz, and has marginal accuracy in certain frequency bands above that.

One can also view the time evolution of the computed transfer functions. Figure 3.16 shows two examples of the time evolution of transfer functions over 1 s. The magnitude of the transfer function is shown plotted as a function of frequency and time for average wind speeds of 0.5 and 4.5 m/s. One can easily see that the evolution of the transfer function for the low-wind-speed case is much more smooth than for the higher-wind-speed case where the evolution appears random, especially at higher frequencies. If the rate of change in the

transfer function becomes too great, the performance of this method may become poor. The point at which the performance degrades will be a function of the complexity and the coherence time of the transfer function.




Fig. 3.14: The average transfer function found using the adaptive method over 1 s of data having an average wind speed of 0.5 m/s. The magnitude (solid line) and phase (dashed line) are shown.






Fig. 3.15: Estimated power spectra of the received signal (solid line) and the error signal (dashed line) over the same one second data segment used for Fig. 3.14.



(a)



(b)

Fig. 3.16: Examples of the time evolution of computed transfer functions over 1 s of data having (a) 0.5 m/s and (b) 4.5 m/s average wind speeds. Shown is the magnitude of the transfer function as a function of time and the frequency range 200 to 2000 Hz.

These computed transfer functions can be used to compute estimates of the coherence time associated with the atmosphere. The coherence time associated with a process can be defined as [39]


$$\tau_c = \int_{-\infty}^{\infty} |\tilde{\Gamma}(\tau)|^2 d\tau \quad (3.19)$$

where $\tilde{\Gamma}(\tau)$ is the normalized autocorrelation function of the process given by


$$\tilde{\Gamma}(\tau) = \frac{\Gamma(\tau)}{\Gamma(0)}. \quad (3.20)$$

The coherence times found for computed transfer functions from data having various average wind speeds are shown in Fig. 3.17. Figure 3.17(a) shows the coherence time computed at 400 Hz and (b) shows the coherence times found by averaging over a range of frequencies from 200 to 2000 Hz. As one would expect, the coherence times drop with increasing wind speed. Due to the finite time window used in these calculations, the coherence times shown should be taken as being conservative estimates, meaning that they should generally be higher than the actual values. Often, the coherence time has been found to be less than 0.25 s for moderate wind conditions.

From the measurements discussed in this chapter, one gets a good picture of the medium through which we are attempting to image. Turbulence in the atmosphere, having sizes on the order of the outer scale, distorts the propagating acoustic waves. This effect becomes greater as the wind speed increases. The atmosphere causes significant and often crippling coherence loss in the arriving acoustic wavefronts. Also, the atmosphere exhibits coherence times that are often much less than 1 s. In the next chapter, a discussion of the performance of acoustic imaging will be presented, including the effect of turbulence. The discussion will focus on the classical narrowband imaging techniques.



(a)



(b)

Fig. 3.17: Coherence times extracted from the time evolutions of computed transfer functions versus average wind speed for (a) 400 Hz and (b) averaged over the frequency band 200 to 2000 Hz.

4. IMAGING PERFORMANCE

In this chapter, the performance of the imaging methods outlined in Chapter 2 will be discussed. First the performance of narrowband imaging will be presented. This discussion will include background on the effects of bandwidth, and the variance associated with the visibility measurements in the absence of turbulence or other distortions. Then, a discussion of the effect of turbulence on narrowband imaging will be given that will include the variance in visibility measurements and in images as a function of wind speed. Examples of multifrequency synthesis imaging will be presented. For wideband imaging, results from the tomographic imaging formulation will be presented along with a short description of performance issues. Lastly, results using the white noise imaging technique will be presented.

4.1 Narrowband Imaging Performance

In this section, the performance issues for narrowband imaging will be discussed. First a discussion of non-turbulence-related performance issues will be presented. Then, the field data described in Chapter 3 will be used to characterize the effect of the atmosphere on imaging performance. Finally, a short presentation of multifrequency imaging will be given.

4.1.1 Non-turbulence-related performance issues

Expressions for the variance or self-noise associated with the measurement of visibilities in interferometric imaging were derived by Kulkarni [40]. He gives expressions for the variances associated with the real and imaginary parts of the visibilities, and for the covariances of the real and imaginary parts of the visibilities measured on the same baseline and on different baselines. The derivations are quite complex, so only a summary discussion

will be given here. Since the visibility is a complex valued measurement, we must be careful how we calculate its variance. Kulkarni calculates the variance of the real and imaginary parts as well as the pseudo-variance, which is a measure of the variance of the complex quantity. The so called pseudo-variance for a visibility measured on a baseline composed of sensors m and n is given by

$$\sigma_{mn}^2 = \frac{(S + N)^2}{M} \quad (4.1)$$

where S is the total flux of the source given by

$$S = E[y_m(t)y_m^*(t)], \quad (4.2)$$

N is the power of any additive noise in the received signals, and M is the number of independent samples related to the product of the bandwidth and the integration time. In Eq. (4.2) we use the same definitions for the received signals as were used in Chapter 2. The variance of the real and imaginary parts of the same visibility are given by

$$V[r_{mn}^r] = \frac{1}{2M} \left((S + N)^2 + |R_{mn}|^2 \cos(2\phi_{mn}) \right) \quad (4.3)$$

$$V[r_{mn}^i] = \frac{1}{2M} \left((S + N)^2 - |R_{mn}|^2 \cos(2\phi_{mn}) \right), \quad (4.4)$$

where V is the variance operator and where $|R_{mn}|$ and ϕ_{mn} are the magnitude and phase of the expected value R_{mn} of the measured visibility r_{mn} . The superscripts i and r in Eqs. (4.3)-(4.4) denote the imaginary and real parts, respectively. In general, for the case of acoustic imaging, we can ignore the additive noise term. The pseudo-covariance between visibilities calculated on two different baselines, mn and ab , is given by

$$\mu_{mnab} = \frac{R_{ma}R_{nb}^*}{M}. \quad (4.5)$$

This expression holds even when there is a common sensor between the two baselines. From the above equations, we see that the pseudo-variance of the visibility measurements is independent of the particular baseline and of the structure of the source. However, the pseudo-covariance of two visibilities is highly dependent upon the source structure. Also, the variance of the real and imaginary parts of a visibility are dependent upon the source structure. With all the expressions we see that the variance falls as one over the number of

independent samples, which is an expected result. This also means that the variance falls with increasing bandwidth and/or with increasing integration time. This variance, which is due only to the fact that the visibility measurements are made using a finite bandwidth and integration time, can be significant in acoustic imaging. For example, for a typical bandwidth of 100 Hz, an integration time of 0.1 s, and a sampling rate of 12 kHz, there are only about 10 independent samples that make up the visibility estimate. By comparison, in radio astronomy, a typical bandwidth is 50 MHz, meaning that radio astronomers have about 500 000 times more independent samples in their visibility estimates.

At this point, the method used to compute visibilities from sensor time-series data should be presented. There are several available methods to compute the complex correlation between two signals [8], [41]. The method that was found to most closely approach the theoretical variance discussed above is the complex or quadrature correlator. The theory of sampling a band-pass signal is applied and is discussed in [42]. A block diagram of the correlator is shown in Fig. 4.1. First, the band-pass filter is applied to both of the signals, which are assumed to be real and sampled according to the Nyquist criterion. Then, each signal is split and is multiplied by $\sin(2\pi f_c t)$ and $\cos(2\pi f_c t)$ on each respective leg to form the quadrature representation of the signals. Each of these signals is then low-pass filtered to remove spectral copies and then downsampled by a factor related to the spectral width of the initial band-pass filter. The downsampling rate is approximately given by the ratio of the sampling rate to the bandwidth. Each pair of quadrature legs is summed to form a single complex signal for each original band-pass signal. The complex visibility is then found via the summation operation

$$r_{12}(0) = \sum_{k=1}^M v_1(k) v_2^*(k) \quad (4.6)$$

where M is the block length corresponding to the desired integration time period. The variance of visibilities calculated using this scheme come within 10% of the theoretical levels given by Kulkarni.

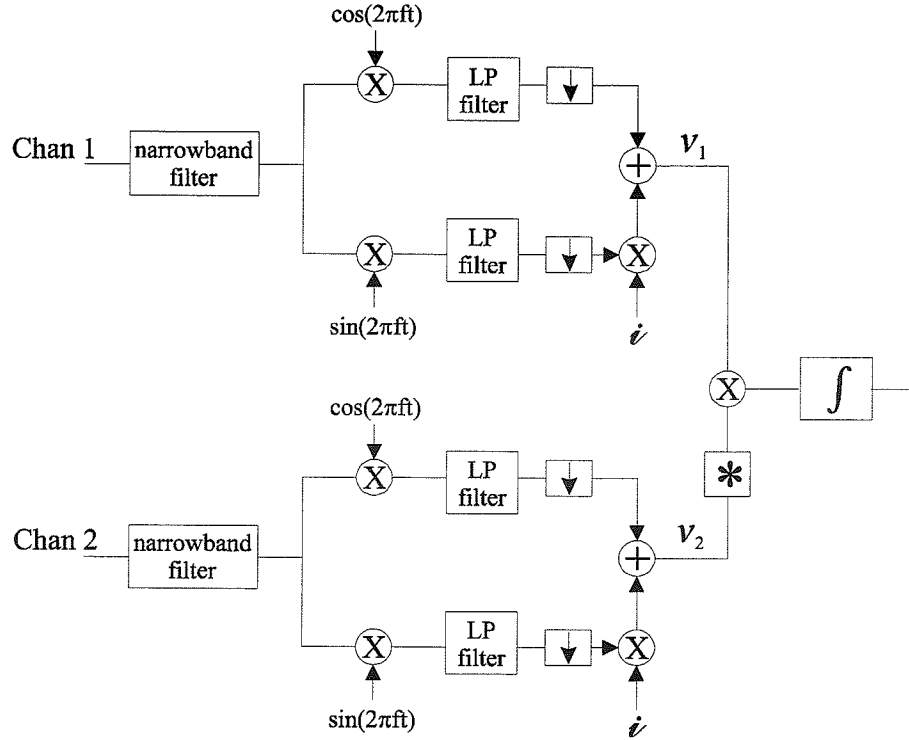


Fig. 4.1 Block diagram of the method used to compute complex visibilities from discrete-time samples of sensor outputs.

Presented in Fig. 4.2 are calculations of the variance of the complex visibilities as a function of integration time. These data were generated by creating synthetic sensor signals assuming that two point sources exist, one at each of the locations most used in the field experiments. This signal configuration is used here to allow for easy comparisons with similar results using the field data to be discussed in the next section. The visibilities were computed from synthetic sensor data using the quadrature method discussed above. The variance is shown normalized by the square of the total flux which is defined in Eq. (4.2). Once again, this is done to allow easy comparison to the variance values computed from the field data. The variance drops as the square of the integration time, as one would expect [43].

Another important performance issue is the variance of the intensity map associated with the visibilities. In Fig. 4.3, the average of the *signal-to-noise ratios* (SNRs) at each point source location and the overall image SNRs are shown as functions of the integration

time. The SNR is the ratio of the image intensity to the standard deviation of the intensity variations. As expected, the SNR values increase as the square root of the integration time. The difference between the two SNR measures can be explained by the fact that the expressions for the variance of the visibilities are highly dependent on the particular sensors involved in each baseline, which in turn causes the variance in the resulting image to be angle dependent. Note that the SNR at each point source would decrease if more sources were added. In Fig. 4.4, the mean image calculated from the synthetic data is plotted with the standard deviations associated with integration times of 0.1 and 3 s.




Fig. 4.2: Normalized visibility variance as a function of integration time. The variance is normalized by the square of the average power or total flux.




Fig. 4.3: The overall image SNR and the average SNR at the point source locations as a function of integration time.




Fig. 4.4: The mean image plotted with the standard deviations corresponding to integration times of 0.1 and 3 s. The image shown was made at 800 Hz using synthetic data. Similar results will occur at other frequencies.

In the discussion above we saw that the variance of the visibility measurements falls with increasing bandwidth. We now discuss the effect of increasing bandwidth on the accuracy of the measured visibilities. In the derivation of the Van Cittert-Zernike theorem in Chapter 2, we saw that a relationship between the source coherence function and the output of the correlator can be given by

$$r_{mn}(\tau) = \int_{f_0 - \Delta f}^{f_0 + \Delta f} df \int_{-\infty}^{\infty} d\xi \hat{\Gamma}(\xi, f) |H(f)|^2 e^{-j2\pi f \frac{(x_n - x_m)\xi}{c}} e^{j2\pi f \tau} \quad (4.7)$$

If we now assume that the bandwidth Δf is finite and that the filter response and the source do not change appreciably over the bandwidth, we can perform the integral over frequency to obtain

$$r_{mn}(0) = \beta \int_{-\infty}^{\infty} \hat{\Gamma}(\xi, f_0) \text{sinc}\left(\frac{D_{mn}\Delta f}{c} \xi\right) e^{-j2\pi \frac{D_{mn}f_0}{c} \xi} d\xi \quad (4.8)$$

where β is a constant made up of the value of the filter over the bandwidth multiplied by the bandwidth value itself. The sinc function is defined as $\text{sinc}(x)/x$. The difference between Eq. (2.22), which is the Van Cittert-Zernike theorem, and Eq. (4.8), is the addition of the sinc function in the integrand. This sinc function serves to degrade the accuracy of the visibility measurement by applying an attenuating window to the intensity distribution of the source. This attenuating sinc window is a function of the baseline length and the bandwidth. The attenuation mainly affects the intensity distribution at large angles, and is more severe for longer baseline lengths. Figure 4.5 illustrates this effect. Shown in Fig. 4.5 is the attenuation in image power as a function of angle and bandwidth. This shows that the bandwidth should be kept at about 100 Hz or less, unless one is sure that the source distribution is concentrated mainly in the center of the view. All relationships given in Figs. 4.2-4.5 were made using a center frequency of 800 Hz. Nearly identical results are obtained at other frequencies. Some small differences will result from inaccuracies in the design of bandpass filters for some frequencies.

Two other important aspects of the performance of narrowband acoustic imaging are the resolution and maximum view angle. The resolution is a function of the ratio of the operating wavelength to the size of the array. Figure 4.6 shows the relationship between the half-power beamwidth, which is being used as a measure of the resolution, and the ratio of

the wavelength to the size of the array. For most cases discussed in this thesis, the typical operating point would be where the ratio is approximately in the range from 0.08 to 0.2. The maximum view angle is a function of the ratio of the operating wavelength to the length of the shortest baseline in the array. This relationship is shown in Fig. 4.7. One must be aware that if the maximum view angle is less than a full 180° , source components outside the maximum view angle will cause aliasing. The range of ratios generally used in the work in this thesis is from 1.6 to 4.1.




Fig. 4.5: Attenuation in image power versus angle and bandwidth. As the bandwidth is increased, the attenuation becomes significant, especially at larger angles. Larger bandwidths can be used only with centrally grouped source distributions.

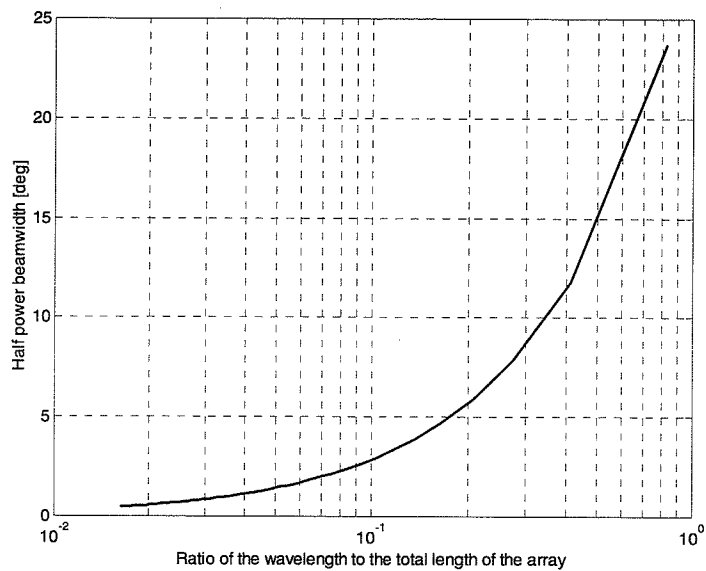


Fig. 4.6: The resolution that can be obtained as a function of the ratio of the wavelength to the total length of the array. The resolution is given in terms of the half-power beamwidth associated with a point source located at 0° .




Fig. 4.7: Maximum total view angle (centered at zero) possible as a function of the ratio of the wavelength to the length of the shortest baseline in the array.

4.1.2 The effect of the atmosphere

In this section, data compiled relating to the effects of the atmosphere on narrowband imaging will be presented. The variance of the measured visibilities will be given as a function of average wind speed. The SNR of the resulting intensity maps will also be presented as a function of the average wind speed. In both cases, the results will be compared to the results for atmospheric-distortion-free synthetic data given in the previous section. Several example images will be presented as well.

A total of 36 segments from the field experiment data were identified that have relatively constant wind speeds over their durations. The segments vary in length from 4 to 28.8 s. These segments have been used to compute imaging statistics such as the variance of the visibility measurements and image variations, as a function of the average wind speed. All the data presented in this section were prepared using an integration time of 0.1 s. This short time was chosen to ensure that the assumption that the atmosphere is essentially unchanged during the integration period is valid. This will be important for the maximum likelihood methods described in Chapter 5.

First, a measure of the amount of distortion present in the images produced from each segment was calculated. This distortion measure is based upon the total squared error existing between the images produced over the segment and a theoretical image produced with knowledge of the positions of the two sources and their relative strengths. The theoretical image is scaled according to the average power in the measured images and is translated in angle to best match the measured images. The translation accounts for any small tilt that may have been present in the setup of the array during each experiment. The distortion measure is also normalized to the length of the particular segment and to the average image power in the segment to ensure that those factors do not influence the distortion measure. Figure 4.8 shows this distortion measure as a function of the average wind speed for images made at both 800 and 400 Hz. One notices that for both frequencies, the distortion in the images increases with increasing wind speed. This is as one would expect, due to the fact that the strength of the turbulence increases with increasing wind speed as illustrated in Section 3.2.4. Also, we see that the distortion in the 400-Hz images is

less than for the 800-Hz images. This is in general agreement with the results of Section 3.2.2, where the structure functions implied that the turbulence has less effect at lower frequencies.

Next, the variance associated with the visibility measurements was calculated as a function of wind speed from the same set of data segments. The visibility variances for both 800 and 400 Hz are given in Fig. 4.9, where the variance is shown normalized to the square of the average power, as in the theoretical variance computations shown in Fig. 4.2. Once again we see the general trend of increasing variance with increasing wind speed. Like the trend exhibited in the distortion measure, we see that the 400-Hz variances are smaller than those at 800 Hz. When one compares the variances computed here with those shown in Fig. 4.2, it is clear that the atmosphere has a profound effect on the measurement of visibilities. The variance is increased over the theoretical value by a factor ranging from about 1.5 to 7.5 times for 400 Hz and from 3.5 to nearly 12 times for 800 Hz.




Fig. 4.8: A measure of the distortion present in the images as a function of average wind speed. The results for both 800- and 400-Hz images are shown.




Fig. 4.9 Normalized visibility variance as a function of average wind speed. The visibility variance is normalized by the square of the power for each data segment. The results presented here can be compared the theoretical results presented in Fig. 4.2.

As was done in Section 4.1.1 for the theoretical case, the SNR of the total image, and that associated with the point sources only, has been computed. The results are shown in Figs. 4.10 and 4.11. Figure 4.10 presents the total image SNR as a function of average wind speed for 800- and 400-Hz images. Here we see that the SNR decreases with increasing wind speed, as one would expect. In Fig. 4.11, the average SNR at the point source locations for 800- and 400-Hz images are plotted as a function of average wind speed. The SNR values shown in Figs. 4.10 and 4.11 are significantly lower than the theoretical values shown in Fig. 4.3. In many cases, especially at higher wind speeds, one is operating in the sub-unity SNR regime. Based on the results shown here, it can be reasonably stated that creating good images in the presence of even moderate wind is nearly impossible.




Fig. 4.10: Total image SNR as a function of average wind speed for 800 and 400 Hz.





Fig. 4.11: Average SNR at the point source locations versus average wind speed for 800 and 400 Hz.


An instructive way of viewing the variation in a series of images is to plot the images in cascade fashion as a function of time. Examples of these time-series images are shown in Figs. 4.12-4.14. A total of 7 s of data are shown in each example. Figure 4.12 shows two examples of time-series images at 800 and 400 Hz that correspond to an average wind speed of about 0.1 mph. For reference, the distortion measures for these two examples are 1.2 for 800 Hz, and 0.82 for 400 Hz. Figures 4.13 and 4.14 show examples where the wind speed is 6.4 and 17.6 mph, respectively. The distortion measures for these examples are 4.45, 0.53, 11.98, and 2.46, for Fig. 4.13(a)-(b) and Fig. 4.14(a)-(b), respectively. It is easy to view the random effects of the atmosphere with this type of picture. Following the trend seen in the rest of this section, the amount of distortion visible in the time-series images increases with increasing wind speed and frequency. At 17.6 mph, the 800-Hz image series amounts to little more than random noise.

To see the associated standard deviation of each of the example image series presented in Figs. 4.12-4.14, the plots in Figs. 4.15-4.17 were made in which the mean image over the entire 7 s of data and the corresponding standard deviation of the intensity are shown. Note that the intensity of each point source at each different wind speed and each different frequency is different. Those intensity ratios are listed in the figure captions. The standard deviation of the intensity overtakes the values of the mean intensity for the 800-Hz examples starting with the 6.4-mph wind speed, but not until the 17.6-mph speed for the 400-Hz examples. We see that relatively accurate images, in terms of the point source locations and relative intensities, are formed for the low-wind-speed case shown in Fig. 4.15. However, for the higher wind speeds shown in Figs. 4.16 and 4.17, the images do not generally portray the true physical situation. This is especially true at the highest wind speed, where the image at 800 Hz is extremely poor. It seems that in general, integration over 7 s is not long enough to produce accurate images in the atmosphere in conditions other than the most ideal. This also suggests that the atmosphere has a significant, large-scale, long-term effect on imaging. The effects of these very large-scale inhomogeneities have been found by others to be on the order of tens of seconds to minutes [44]. As an example, this effect can be seen in Fig. 4.14(a) where the left point source seems to essentially

disappear for several seconds and then reappear. This behavior is common to much of the data that were collected from the field experiments.




(a)

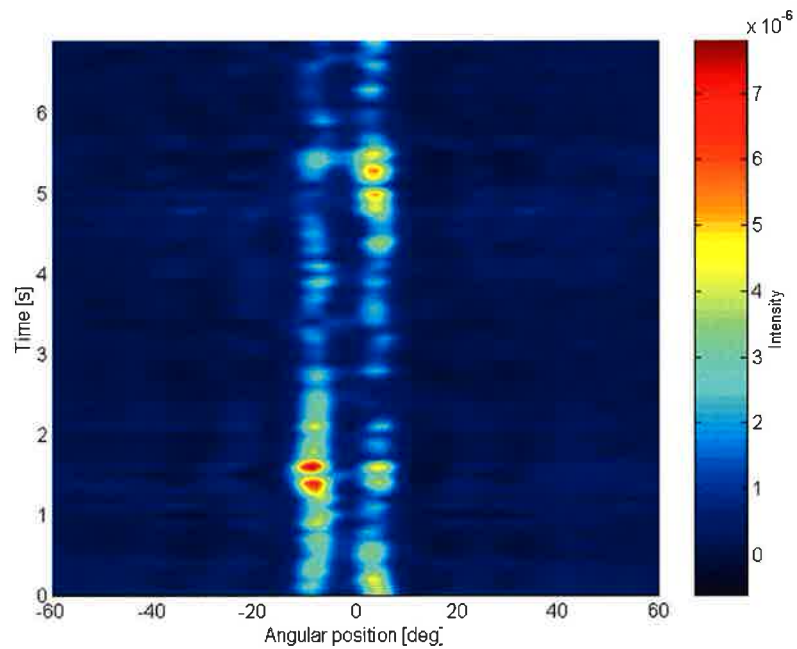


(b)

Fig. 4.12: Examples of time-series images for a low wind speed case for (a) 800 Hz and (b) 400 Hz. The wind speed was 0.1 mph. The distortion measure is 1.20 at 800 Hz and 0.28 at 400 Hz. The intensity is given by the color as indicated on the bar on the right. The images are composed of 70 0.1-s images. The CLEAN algorithm was applied to the images.




(a)




(b)

Fig. 4.13: Examples of time-series images for a medium wind speed case for (a) 800 Hz and (b) 400 Hz. The wind speed was 6.4 mph. The distortion measure is 4.45 at 800 Hz and 0.53 at 400 Hz. The intensity is given by the color as indicated on the bar on the right. The images are composed of 70 0.1-s images. The CLEAN algorithm was applied to the images.




(a)




(b)

Fig. 4.14: Examples of time-series images for a medium wind speed case for (a) 800 Hz and (b) 400 Hz. The wind speed was 17.6 mph. The distortion measure is 11.98 at 800 Hz and 2.46 at 400 Hz. The intensity is given by the color as indicated on the bar on the right. The images are composed of 70 0.1-s images. The CLEAN algorithm was applied to the images.




(a)




(b)

Fig. 4.15: The mean image (solid line) of the time-series data presented in Fig. 4.12 for (a) 800 Hz and (b) 400 Hz. The computed standard deviation of the intensity is shown with a dotted line. The relative intensities of the two point sources (ratio of left point source to the right point source) are (a) 1.32 and (b) 0.82. The CLEAN algorithm was applied to the images.




(a)




(b)

Fig. 4.16: The mean image (solid line) of the time-series data presented in Fig. 4.13 for (a) 800 Hz and (b) 400 Hz. The computed standard deviation of the intensity is shown with a dotted line. The relative intensities of the two point sources (ratio of left point source to the right point source) are (a) 4.14 and (b) 1.15. The CLEAN algorithm was applied to the images.



(a)



(b)

Fig. 4.17: The mean image (solid line) of the time-series data presented in Fig. 4.14 for (a) 800 Hz and (b) 400 Hz. The computed standard deviation of the intensity is shown with a dotted line. The relative intensities of the two point sources (ratio of left point source to the right point source) are (a) 0.52 and (b) 2.23. The CLEAN algorithm was applied to the images.

4.1.3 Multifrequency synthesis results

In Chapter 2, the theory related to the multifrequency imaging technique was presented. In this section, some results of using this method, and its general applicability to acoustic imaging, will be discussed.

First some comments on the how the technique was implemented. The process of computing the multifrequency image begins with the computation of the wideband correlation function for each baseline in the array. The Fourier transform of each correlation function is then taken. Each point in the resulting FFTs is a measurement of the complex visibility at that particular frequency. By using the notion of space-frequency equivalency mentioned during the theoretical development in Chapter 2, each measurement falling within a preselected band of frequencies can be placed into a single vector representing the Fourier transform of the source intensity distribution. This is done by creating a vector that contains all spatial frequencies from the dc term, up to a maximum value given by

$$u_{\max} = \frac{f_{\max} D_{\max}}{c} \quad (4.9)$$

where f_{\max} is the maximum frequency to be used and D_{\max} is the maximum baseline length of the array. Measurements that might fall into the same spatial frequency bin are averaged. Once this Fourier transform vector is filled with all the included measurements, the inverse Fourier transform is applied to obtain the intensity image.

As one adds more frequencies into the process, the effect on the resulting image is similar to that of increasing the integration time for the standard narrowband imaging process. Figure 4.18 shows the mean image obtained using synthetic data, like that described earlier, along with the associated standard deviation. This illustrates the fact that the SNR at the point source locations drops as the square root of the number of frequencies used. Also, as the maximum frequency is raised, the resolution obtained in the image is increased because the effective size of the array increases with increasing frequency. These seem like good properties; however, in practice there is little or no apparent advantage in using the multifrequency techniques on real acoustic imaging problems. The main problem is that the assumption of little or no change in the intensity of the source distribution over the band of

frequencies used in the process, does not generally hold. The method was applied to field data, but there was little success in significantly improving the image over using the standard narrowband quadrature method described in Section 4.1.1. In fact, often the image was much worse using multifrequency synthesis, most likely due to significant changes in the intensity as a function of frequency made even worse by the distorting atmosphere. A typical example of the best relative performance observed on field data is shown in Fig. 4.19. In this example a 0.5-s field data segment was used to produce both a standard 400-Hz narrowband image using the quadrature method, and an image using the multifrequency technique in which six frequencies in a band extending from 300 to 500 Hz were used. For these particular data, the ratio of the intensity of the left point source to that of the right should be approximately 0.6. The multifrequency image does improve on the accuracy of the relative intensity by a small amount.




Fig. 4.18: The relationship between the mean image computed using the multi-frequency technique, and the associated standard deviations for the cases where one, three, or six frequencies are used. The maximum frequency span is $\pm 12\%$ of 800 Hz. Synthetic data was used.




Fig. 4.19: A typical example comparing the standard narrowband quadrature image (solid line) to an image produced using the multifrequency technique. Both images were made with center frequency 400 Hz. The multifrequency process used six frequencies in the range from 300 to 500 Hz. The L/R intensity ratio for these data should be approximately 0.6.

4.2 Wideband Imaging Performance

In this section some results and performance issues related to the wideband imaging techniques discussed in Section 2.2.1 will be presented. First, the method of producing images of the power spectrum of the source distribution as a function of angle using a tomographic formulation will be discussed. Results using two different tomographic solution methods, direct Fourier inversion, and algebraic reconstruction, will be presented. Then, results of the white noise imaging technique outlined in Section 2.2.2 will be shown.

4.2.1 Power spectrum imaging using tomographic solution methods

In Section 2.2.1, Eq. (2.37) gives the relationship between the source coherence function and the array coherence function. As previously stated, this relationship is identical to that associated with a standard parallel beam tomography problem. Instead of using x-rays

and solving for the density of tissue inside a person's head, we can use the same solution methods to solve for the two-dimensional (2-D) source coherence function. Two methods will be discussed here. The first method is direct Fourier inversion based on the projection-slice theorem. The second is an algorithm that is a member of the general class of algebraic reconstruction techniques.

As discussed in Section 2.2.1, the projection-slice theorem states that the 1-D Fourier transform (FT) of the projections, which in this case are the cross-correlation functions of each baseline, are equal to radial slices of the 2-D FT of the unknown source coherence function at the same angles from which the projections were made. The angle of the projection in this case is proportional to the baseline length. One must be sure to also include the complementary projections corresponding to negative baseline lengths. The correlation functions for these projections are simply the delay-reversals of those for the positive angles. The 2-D Fourier transform of the source coherence function can thus be constructed by proper combination of the 1-D FTs of the cross-correlation functions. Since the 1-D Fourier transforms are aligned radially within the 2-D FT, a re-gridding to Cartesian coordinates is necessary in order to apply the standard 2-D inverse fast Fourier transform (FFT) to recover the source coherence function. A cubic interpolation routine called *griddata*, which is native to MATLAB, was used for this purpose [45]. One must be concerned with the possibility of aliasing when applying this reconstruction technique. As mentioned in Section 2.2.1, a relationship exists between the maximum frequency present in the image and the maximum angle that can be viewed in the image (see Fig. 4.7). Thus, in order to avoid aliasing in the final image, we must filter in both the frequency and spatial domains. Filtering in the frequency domain is straightforward. A filter is applied to the sensor data before correlation, or a window function is applied to the FT of the correlations. Spatial domain filtering consists of applying a window to the cross-correlation functions. The width of the window, in terms of correlation lag, is related to the maximum view angle by the simple relationship

$$\tau_{\max} = \frac{D\xi_{\max}}{c}, \quad (4.10)$$

where τ_{\max} is the magnitude of the maximum delay corresponding to the maximum view angle ξ_{\max} . Note that the maximum delay is a function of the baseline length D .

Computer code to compute these images was written in MATLAB. The code allows for arbitrary grid sizes to be used in the reconstruction, with cubic interpolation used to fill the 2-D FT grid using the radial 1-D FT information. Figure 4.20 shows an example image computed from synthetic data. This image was found by taking the 1-D FT along the autocorrelation lag direction in the 2-D reconstructed source coherence function. The synthetic data used were composed of four point sources. There is a low-pass source extending to 1500 Hz at 15°, a 1000-1500 Hz bandpass source at 5°, a 500-1000 Hz bandpass at -10°, and a low-pass source extending to 500 Hz at 20°. All four sources are clearly visible, along with some artifacts of the reconstruction. These artifacts are a result of issues such as imperfect spatial filtering and limited temporal resolution. The arching ridges are the effect of aliasing from imperfect filtering. Note that the resolution is a function of frequency, as one would expect. Due to this fact, the power spectrum level appears to fall at low frequencies as the energy is spread out. This is a predictable effect and can be removed. Shown in Fig. 4.21 is the 2-D coherence function, which shows the autocorrelation of the source as a function of angle.




Fig. 4.20: Example result of an image of the power spectrum versus angle using the direct Fourier inversion technique. The arbitrary power spectrum level is given by the color scale at the right. One second of synthetic data was used. Four sources exist: -20° with frequency band 0-500 Hz, -10° with 500-1000 Hz, 5° with 1000-1500 Hz, and 15° with 0-1500 Hz.




Fig. 4.21: The coherence function of the source as a function of angle from which Fig. 4.20 was computed. The arbitrary autocorrelation level is given by the color scale at the right.

The second technique to be discussed for recovering the 2-D source coherence function is a method called SIRT, or simultaneous iterative reconstruction technique [20], [46]. SIRT is a member of the algebraic reconstruction (ART) family of tomography methods. The general problem formulation is the same for all ART methods. The unknown object, in this case the source coherence function, is pixelated. Then, a large set of linear equations can be constructed by summing the contribution of each pixel to each point in each projection or cross-correlation function. The resulting set of linear equations can be very large, usually having several thousand equations and unknown pixels. The SIRT method is a simultaneous updating iterative solution method that produced good results when applied to the imaging problem at hand. The problem was generally set up by pixelating the source coherence function to form an unknown matrix of size 137×90 where the autocorrelation lag dimension received the greater number of pixels. This produces a total number of unknowns

equal to 12 330. A total of 12 341 measurements are available in the problem. This number of measurements is a function of the sampling rate used in the sensor data and the number of unique baseline lengths in the array. The number of baselines is related to the number of projections, or cross-correlation functions, and the sampling rate gives the number of relevant points in each projection. The cross-correlation functions for baselines having the same length are averaged for use. The very large matrix (12 330 x 12 341) of weights is computed ahead of time by determining, for each projection point, which pixels of the unknown function are involved. In this case, the width of the projection ray was taken into account, and the partial contributions of all pixels were computed. Because of memory constraints, this large matrix was compressed and stored in a multidimensional array for use in the iterative reconstruction algorithm. Figure 4.22 shows an example of the results obtained using the SIRT algorithm. Generally, about 100-200 iterations are required for convergence. Figure 4.22 shows the result of using the SIRT algorithm on the same data that was used to compute the image in Fig. 4.20 using the direct Fourier inversion technique. The image in Fig. 4.22 exhibits greater definition of the sources. For example, the low-pass source is much more visible. Artifacts similar to those found in Fig. 4.20 can be seen in Fig. 4.22 as well. Figure 4.23 shows the corresponding source autocorrelation function for the SIRT result. The autocorrelation structure is more visible here than in Fig. 4.21. The SIRT algorithm requires much greater computational requirement and execution time, but seems to produce higher quality images. Results using any of the vast number of available tomographic reconstruction algorithms may indeed offer performance that matches or exceeds that shown here. The aim here was to show two representative results. Figure 4.24 is an example of the SIRT solution method as applied to a segment of field data. The image shown here represents what should be considered a good result from field data. The wind speed for this data segment was 5.5 mph. Two seconds of data were used.




Fig. 4.22: Example result of an image of the power spectrum versus angle using the algebraic reconstruction technique, SIRT. The arbitrary power spectrum level is given by the color scale at the right. One second of synthetic data was used. Four sources exist: -20° with frequency band 0-500 Hz, -10° with 500-1000 Hz, 5° with 1000-1500 Hz, and 15° with 0-1500 Hz.




Fig. 4.23: The coherence function of the source as a function of angle from which Fig. 4.22 was computed. The arbitrary autocorrelation level is given by the color scale at the right.




Fig. 4.24: An example of a result from field data with wind speed equal to 5.5 mph. The arbitrary power spectrum level is given by the color scale at the right. Two seconds of data were used.

It should be mentioned at this point that one can produce images like those shown in this section by simply combining several narrowband images computed at different frequencies. An example result of this process is shown in Fig. 4.25. One must be careful to scale each image properly to account for the change in scale due to frequency. There should be an equivalence between the tomographic and piecemeal methods of producing 2-D images. That equivalence is not derived here, but it should be noted that any such equivalence may be destroyed when the distortion of the atmosphere is brought into the picture. This is due to induced distortions that are a function of frequency and position. The piecemeal images, like that shown in Fig. 4.25, are generally of superior quality, with fewer artifacts, than the tomographic-based results. The CLEAN algorithm can be applied to each narrowband image to further increase image quality. With some modifications in procedure, this can also be said for the tomographic case.




Fig. 4.25: An example result of an power spectrum image constructed from several narrowband images. The same synthetic data that was used to compute Figs. 4.20 and 4.21 was used here.

4.2.2 White noise imaging results

In Section 2.2.2, a method of producing images of source distributions that are composed of white noise elements was developed. The technique involves arranging wide-band cross-correlation measurements into a map of the intensity of the source distribution. For the many cases where multiple measurements correspond to the same point in the image, an optimal method of combining those measurements was derived based on linear estimation theory. Figure 4.26 shows a typical result of this method where synthetic data is used. The data contained five white noise point sources with varying intensities. The intensities were 1 at -20° , 0.5 at -15° , 2 at -5° , 1 at 8° , and 0.25 at 18° . The improvement in the image afforded by the optimal combination method is generally minor. Under certain conditions, the

improvement is more noticeable. Once such case is when there are sources located at opposite angles, such as $\pm 20^\circ$, for example.




Fig. 4.26: An example image from the white noise imaging technique. The synthetic data used contain five white noise point sources.

5. IMAGE ENHANCEMENT

We have seen in the preceding chapters that the atmosphere causes severe distortion in the signals received at an array of sensors, which results in errors in the measurements of the visibilities and subsequently in the images formed from those visibilities. In this chapter, methods for enhancing the quality and accuracy of acoustic images made through the atmosphere will be discussed. We will start with a discussion of the widely used and successful family of methods from the radio astronomy community called self-calibration. Evidence will be presented that shows that these techniques, which revolve around the concept of phase closure, cannot be applied to the acoustic imaging problem at hand. Next, an estimation problem based on multiple short-integration-period images, or snapshots, will be presented. Bounds will be derived for this problem and their implications discussed. Initial attempts at recovering solutions to the problem will also be presented.

5.1 Self-Calibration and Its Application to Acoustic Coherence Imaging

This section will present the concept of phase closure and the self-calibration family of methods that have been developed around it. These methods aim to enhance imaging performance in the presence of atmospheric distortions. Also, the application of these methods to the acoustic imaging problem at hand will be discussed.

5.1.1 Phase closure and self-calibration

Let us consider the following relationship between the visibility $V'_{jk}(t)$ as measured by sensors j and k and the corresponding actual visibility $V_{jk}(t)$:

$$V'_{jk}(t) = G_j(t)G_k^*(t)V_{jk}(t) + \varepsilon_{jk}(t). \quad (5.1)$$

In Eq. (5.1), $\varepsilon_{jk}(t)$ is the system noise contribution, and the complex gain factors $G_j(t)$ and $G_k(t)$ are lumped parameters that represent the effects of propagation and the instrumentation chain for each sensor. The time dependence of $V_{jk}(t)$ can be dropped if the measuring array does not move relative to the source. If we consider three array elements i , j , and k , ignore the noise term, and look at the phase part of Eq. (5.1), we find that

$$\Psi_{ijk}(t) = \angle V'_{ij}(t) + \angle V'_{jk}(t) + \angle V'_{ki}(t) = \angle V_{ij}(t) + \angle V_{jk}(t) + \angle V_{ki}(t) \quad (5.2)$$

because the phase of the gain factor of each sensor cancels out in the sum. Thus the sum of measured visibility phases around a triad of sensors yields a measurement that is independent of phase errors. This relationship was first noted by Jennison [47] and Ψ_{ijk} is called the closure phase after Rogers et al. [48]. In fact, closure phases can be formed for any closed polygon of sensors. In a similar manner, an amplitude closure relation can be constructed around a quartet of sensors

$$\frac{|V'_{ij}(t)||V'_{kl}(t)|}{|V'_{ik}(t)||V'_{jl}(t)|} = \frac{|V_{ij}(t)||V_{kl}(t)|}{|V_{ik}(t)||V_{jl}(t)|} \quad (5.3)$$

where the magnitudes of the gain factors all cancel [49]. The important underlying assumption at work here is that the atmospheric phase shifts at each sensor must remain constant over the angular extent of the source. This implies that the phase distorting medium, or screen, is located in very close proximity to the array. This assumption holds in many cases in radio astronomy where the field of view is small and the atmospheric effects are indeed located relatively close to the array--relatively in regard to the very distant sources. The implications of this assumption on the applicability of phase closure to acoustic imaging will be discussed in the next section.

For a nonredundant array of N sensors, there are $\frac{1}{2}(N-1)(N-2)$ independent closure phases available and $\frac{1}{2}N(N-3)$ closure amplitudes. The number of real numbers that are needed to completely describe all the complex visibilities that are measurable with the array is $N(N-1)$. Thus, the closure quantities cannot completely specify all the information needed.

With clever array design using redundancy, it is possible to determine all the complex visibilities to within an overall linear phase shift [50]. The measured visibility amplitudes can sometimes be used directly because they are often less affected by the atmosphere than the phases. Using that fact and array redundancy, one can recover all the potential information the array has to offer. This is especially true as the number of array elements grows. A number of related methods have been developed to accomplish this. They are all based on the closure relation. Some of these will be discussed next.

The closure relation defined above can be used to reconstruct maps which tend to be closer to the actual brightness of the source. One such method, called hybrid mapping, was first developed by Readhead and Wilkinson in 1978 [51], with numerous modifications following. In this class of procedure, a model map is iteratively refined using the closure relations following a scheme such as the following:

1. Decide upon an initial model.
2. Form a complete set of independent closure equations. Add to those a sufficient number of visibilities derived from the model such that the total number of relations equals the number of sensor spacings.
3. Solve these equations for the visibility corresponding to each baseline. Make a map using this data.
4. Apply the CLEAN algorithm (without using the residual) and use the resulting map as the new model map.
5. Test for convergence and repeat steps (2)-(4) as necessary.

When closure phases are used one loses the absolute position of the source, and when amplitude closure is applied, only relative brightness can be discerned. However, it is a relatively easy task to establish a brightness calibration, and often the visibility amplitude measurements are good enough to use without closure.

Closely related to hybrid mapping is the family of methods called self-calibration [52], [53]. The main thrust of self-calibration is to bypass the closure relations and let the complex gain factors of Eq. (5.1) be free parameters in a gradient optimization procedure. The optimization is carried out such that the weighted difference between the measured visibilities and model visibilities corrected by the gain factors is as small as possible. Usually, the squares of the residuals

$$C = \sum_t \sum_{j < k} W_{jk} \left| V'_{jk}(t) - G_j(t) G_k^*(t) V_{jk}(t) \right|^2 \quad (5.4)$$

is minimized, subject to $\angle G_r(t) = 0$ for some reference sensor and where W_{jk} are weights that are related to the inverse of the variance of the noise. The method can be summarized as follows:

1. Begin with an initial model as in hybrid mapping.
2. Find the visibilities at the observed sample points.
3. Solve for the complex gain factors by minimizing C using a gradient algorithm.
4. Correct the observed visibilities according to
$$V_{jk}(t) = G_j^{-1}(t) G_k^{*-1}(t) V'_{jk}(t)$$
5. Form a new model from V_{jk} using CLEAN.
6. Test for convergence and repeat steps (2)-(5) as needed.

Despite the fact that these two methods seem unfounded, they have been successfully and extensively used in the radio astronomy community.

5.1.2 Applicability of phase closure to acoustic imaging

The phase closure based methods described in the last section have been used very successfully in the radio astronomy community for decades. Unfortunately, in most cases it seems that the concept of phase closure cannot be applied to acoustic imaging through the atmosphere. There are two basic reasons why phase closure will not apply to acoustic imaging.

First, the assumption that the atmospheric phase distortions are constant across the angular extent of the source is violated. The extent of sources common in acoustic imaging may be very wide. The field data collected for this thesis include two point sources separated by about 12° , and there can easily be situations where larger angles will need to be imaged. Additionally, the distorting atmosphere is present over the entire propagation path. It is therefore inconceivable that the distortion effects would be independent of the position on the source. Evidence of this fact can be seen in Fig. 5.1, where the phase responses of transfer functions calculated using the method described in Section 3.2.5 are shown for three different sensors in the array. The phase has been compensated for geometrical delays and is shown

unwrapped for ease of viewing. For the two sensors located 0.22 m apart, their phase angles match relatively closely, with errors usually being less than about 0.5 rad. However, the phase of the transfer function for the sensor at 1.7 m behaves very differently than those of other two sensors. If we employ the concept of reciprocity, we see that this large difference in phase response between sensors separated by 1.7 m is analogous to the relative responses that one would see between source elements separated by about 1° . The wind speed associated with the data used in Fig. 5.1 was low at 3 mph. The requirement that the phase errors be independent of position on the source is equivalent to requiring that the distorting phase screen is located at the array. When this assumption is violated, even by a small amount, the performance of the self-calibration method suffers drastically. To illustrate this, a simulation was prepared where theoretical visibilities were distorted by a single phase screen located at varying distances from the array. The self-calibration method was then applied to the distorted data. The results of this simulation are shown in Fig. 5.2. Plotted in Fig. 5.2 are measures of the mean squared error in the images reconstructed with the self-calibration method as a function of the distance from the array to the phase screen. The simulation was done for two different point source separations. One was done at the 11.4° separation present in much of the field data, and one at half that, or 5.7° . One sees from Fig. 5.2 that for both point source configurations the image error rises quickly with phase screen distance, with the smaller source separation case rising more slowly. As evidenced by the saturation in the traces, the reconstructed images become essentially noise-like when the phase screen is moved far enough away from the array. For visual reference, examples of the reconstructed images for phase screen distances of zero and 0.2 m are shown in Fig. 5.3. Similar simulations by others have shown similar results, although very much more optimistic in nature [54]. In this case, the optimism was a result of a misinterpretation of the results. It should be noted that the self-calibration method works fine for the single-point-source case for all phase screen locations.




Fig. 5.1: An example illustrating the large difference between the phase responses of atmospheric transfer functions calculated at three different sensors over the same segment of data. The phase responses have been altered to remove the effect of the geometrical delay. The phase is shown unwrapped. The data used were associated with an average wind speed of 3 mph.




Fig. 5.2: The mean squared error in images reconstructed using the self-calibration method from data distorted by a phase screen located at increasing distances from the array. Two point sources were present. Results for angular separations of 11.4 and 5.7° are shown. The error was calculated using images normalized to unity maximum intensity.





Fig. 5.3: Examples of the solutions found using the self-calibration method for phase screen distances of zero (solid line) and 0.2 m (dotted line). The point source separation was 11.4° .

The second problem in applying the phase closure methods to acoustic imaging is the high variance associated with the visibility measurements. The phase closure based solution methods are very sensitive to noise in the visibility measurements. In radio astronomy, because they deal with much larger bandwidths and integration times, they can collect visibility measurements that have much less variance associated with them. As an example of this behavior, Fig. 5.4 shows the effect of measurement noise on the quality of reconstruction using the matrix equation approach. In this approach, the closure phase quantities, given by Eq. (5.2), are directly used to solve for the actual visibilities according to the matrix equation


$$\mathbf{A}\mathbf{v} = \mathbf{b} \quad (5.5)$$

where \mathbf{v} is the vector of the unknown actual visibility phases, \mathbf{b} is the vector of closure phase quantities, and \mathbf{A} is the matrix that relates the two. The matrix \mathbf{A} is rectangular, but is full rank assuming that care is taken in the design of the array [50]. The equation is solved using the pseudo-inverse. The condition number of the matrix is typically 500, making the solution sensitive to noise in the closure quantities. In Fig. 5.4(a), the reconstruction with no

measurement noise added is shown along with the image distorted by the phase screen. The reconstruction is perfect. In Fig. 5.4(b), the same distorted image is shown along with the reconstruction resulting when measurement noise corresponding to an integration time of 0.5 s is added. Note that the reconstruction in (b) is very poor. The self-calibration method behaves in a similar manner, but is slightly less sensitive to the measurement noise.



(a)



(b)

Fig. 5.4: Examples illustrating the effect of noise on closure phase solutions. The image as distorted by the phase screen is plotted with the reconstructed image for (a) no measurement noise added and (b) measurement noise corresponding to 0.5 s added.

5.2 A Discussion of Atmospheric Distortion Models

In this section, a discussion of the modeling of the atmosphere will be presented. In the last section, the concept of a phase screen was mentioned in the context of phase closure and self-calibration. Presented here will be a discussion of the form of the imaging equations now including the phase screen model of the atmosphere. The effect of the phase screen on the visibilities and resulting images will be discussed as well. It will be shown that the phase screen model is sufficient to capture most of the behavior exhibited in the field data. The phase screen model will be used extensively in the next section, where an estimation problem is set up to solve the image recovery problem.

Following the derivation of the Van Cittert-Zernike theorem in Section 2.1.1, we begin here by stating the form of the received signal at sensor m in the array as

$$y_m(t) = \int_{-\infty}^{\infty} S(\xi, t - \frac{x_m}{c}\xi - \phi_m(\xi)/2\pi f) d\xi \quad (5.6)$$

where $\phi_m(\xi)$ is phase delay, that is a function of angle, associated with sensor m . The frequency dependence of this shift is given explicitly in Eq. (5.6). In the same manner as before, we can then write the sensor signal after the narrowband filter as

$$v_m(t) = \int_{-\infty}^{\infty} \hat{S}_T(\xi, f) H(f) e^{-j2\pi f \frac{x_m}{c} \xi} e^{-j\phi_m(\xi)} e^{j2\pi f t} d\xi dt. \quad (5.7)$$

Adherence to the analytic signal formulation is still implied in the above equations.

Proceeding in the same way as before, we arrive at a new expression for the Van Cittert-Zernike theorem originally given in Eq. (2.22). The new expression becomes

$$r_{mn}(0) = \int_{-\infty}^{\infty} \hat{\Gamma}(\xi, f_o) e^{-j(\phi_n(\xi) - \phi_m(\xi))} e^{-j2\pi f_o \frac{(x_n - x_m)}{c} \xi} d\xi. \quad (5.8)$$

The addition of the complex exponential comprised of the phase errors will cause an error in the correlation or visibility measurement. An attenuation or gain function of the form $G_n(\xi)$ can be added for each sensor to account for amplitude fluctuations. This case is discussed later in this section.

If we make the assumption required by the phase closure based method, namely that the phase errors are independent of angle, we may pull the phase errors out of the integral in Eq. (5.8), which results in the expression

$$r_{mn}(0) = e^{j\phi_m} e^{-j\phi_n} \int_{-\infty}^{\infty} \hat{\Gamma}(\xi, f_o) e^{-j2\pi f_o \frac{(x_n - x_m)\xi}{c}} d\xi \quad (5.9)$$


that can be directly related to the form of Eq. (5.1), the basis for the phase closure techniques. As the phase screen is moved away from the array, the angle dependences of the phase errors come into play. As the phase screen is moved further and further away from the array, the sensor dependence becomes less significant than the angular dependence. In the limit where the phase screen is located at the source, the phase errors cancel out and the effect of the screen is removed.

For the phase screens used in the work to follow in the next section, the screen is represented in discrete fashion, with equally sized pixels in the lateral direction. Thus the position on the screen that affects a signal arriving at sensor n from direction ξ_0 is


$$x_s = x_n + R \tan(\sin^{-1}(\xi_0)) \quad (5.10)$$

where R is the distance from the array to the phase screen. The maximum extent of the phase screen is chosen according to the desired maximum view angle. In most cases in this work, $\xi = 0.8$ was chosen. Illustrating the effect of these phase screens on the resultant images, Fig. 5.5(a) shows the average image SNR at the two point source locations for a phase screen located 5 m away from the array and with a varying number of pixels. The synthetic data were created to simulate results from a 0.1 s integration time. The phase pixels were generated to be independent from one pixel to the next and uniformly distributed from zero to 2π . The SNR drops as the number of pixels is increased, but quickly reaches what appears to be an asymptotic SNR value of around 0.75. Looking back at Fig. 4.9, one notes that this value of SNR was associated with wind speeds in the neighborhood of 8 to 10 mph. Show in Fig. 5.5(b) is the normalized visibility variance corresponding to the SNR data shown in (a). Note that these variance values correspond to the same range of wind speeds, according to Fig. 4.11. The effect of moving the phase screen further away from the array is shown in Fig. 5.6, where the normalized visibility variance is plotted as a function of the range to the screen for a phase screen having 10 pixels. The variance drops down to the theoretical level

for a 0.1-s integration time. So it seems that a single phase screen can be used to simulate the behavior of the atmosphere for wind speeds up to about 10 mph.



(a)



(b)

Fig. 5.5: The (a) average SNR at the point source locations and (b) normalized visibility variance as a function of the number of phase screen pixels. Comparisons can be made with values found in Figs. 4.9 and 4.11 for field data. The phase screen was placed 5 m away from the array.




Fig. 5.6: The normalized visibility variance as a function of the distance from the array to the phase screen for 10 phase screen pixels.

Thus far we have discussed a screen that only affects the phase of an impinging wavefront. At this point, results will be discussed for the case where each pixel in the screen is allowed to have not only a phase shift, but also an amplitude gain factor. Equation (5.8) can now be rewritten as

$$r_{mn}(0) = \int_{-\infty}^{\infty} G_n(\xi) G_m(\xi) \hat{\Gamma}(\xi, f_o) e^{-j(\phi_n(\xi) - \phi_m(\xi))} e^{-j2\pi f_o \frac{(x_n - x_m)\xi}{c}} d\xi \quad (5.11)$$

where the gain factors G are a function of angle in the same manner as the phase shifts. Figure 5.7 shows the variance of the visibility measurements when a complex valued screen is located 5 m in front of the array. The phase shifts associated with each pixel were generated in the same manner as described above. The gain factors were taken to be Gaussian with a mean of unity and standard deviation of 0.4. The absolute value of any negative gain was used. The values obtained for the phase alone case are also plotted for easy comparison. Figure 5.7 shows that when gain factors are included, the visibility

variance is increased significantly from what one obtains using phases alone. With the gain factors, the visibility variance corresponds to about 12 mph wind speeds. Also important is the fact that the gain factors could be used to simulate the scintillation seen in images computed from actual field data. In Figs. 4.12-4.14, we see that the power in the image varies with time and that one or both point sources may fade in and out over time. This ability to model scintillation will become important in the next section, when an estimation problem is set up to reconstruct images from multiple short integration period images.




Fig. 5.7: The normalized visibility variance associated with both a phase only screen and a gain and phase screen. The gains were chosen to be Gaussian with mean unity and standard deviation 0.4. The screen in both cases was 5 m in front of the array.

5.3 An Estimation Problem Formulation for Image Recovery

In this section, a method for improving the quality of images obtained through the distorting atmosphere will be presented. This method is based on the idea of estimating a constant source intensity distribution through the use of several snapshots or short exposure

narrowband images. The exposure time is assumed short enough so that the atmosphere can be assumed frozen for each snapshot. The problem boils down to estimating a discrete version of the intensity distribution jointly with the phases associated with a phase screen located at some fixed position in front of the array. The phases for each snapshot are independent and are all estimated. First, more detail about the problem formulation will be provided. Then, a method for computing the Cramér-Rao bounds associated with the problem will be derived. The results of these bounds will then be discussed. Finally, some encouraging initial solutions to the problems of this type will be given.

5.3.1 Problem formulation

A diagram of the setup for this estimation problem is shown in Fig. 5.8. The source intensity distribution is pixelated as a function of angle, in the range $\xi = \pm 0.8$. The intensity is assumed to be constant with time. A phase screen is located at some range from the array and is composed of some number of pixels that are equally sized in the lateral dimension and which cover the lateral distance necessary to include the view angle mentioned above. The phase pixels are assumed to be independent from one snapshot to the next and within each snapshot, from one pixel to the next. The distance to the phase screen and the number of phase pixels can be chosen to suit the situation according to the discussion in Section 5.2.




Fig. 5.8: Diagram of the estimation problem setup.

A discrete approximation to the Van Cittert-Zernike relation is used. For the purposes of this estimation problem, the complex visibility measured on the baseline consisting of sensors m and n is given by

$$r_{mn} = \bar{v}_{mn} = \sum_{i=1}^L P_i e^{j\phi_n(\xi_i)} e^{-j\phi_m(\xi_i)} e^{j2\pi f_o \frac{D_{mn}}{c} \xi_i} \quad (5.12)$$

where L is the number of intensity pixels, P_i are the intensities, f_o is the frequency, D_{mn} is the baseline length, and $\phi_m(\xi_i)$ and $\phi_n(\xi_i)$ are the phase screen pixels, which are functions of angle, affecting sensors m and n . The phase screen pixel that affects a given sensor is found using Eq. (5.10). The pixel center closest to the value computed from Eq. (5.10) is used. The visibilities are assumed to be circularly Gaussian. This is somewhat of an approximation, as the real and imaginary parts of the visibilities have slightly different variances, depending on the source distribution [40]. The complex covariance matrix associated with the full set of visibility measurements for a given array can be constructed using the pseudo-covariances defined in Eq. (4.5). The visibility measurements are assumed independent between snapshots, such that the covariance matrix of all visibility measurements from all snapshots would be block diagonal in structure. Also note that if there are N_p phase screen pixels, only N_p-1 of those will actually be free variables in the estimation process. This is because if one writes all the other phases as the difference from a reference phase, the reference phase cancels from Eq. (5.12).

When picking the number of intensity and phase screen pixels, we must take into account the number of baselines in the measuring array, and the number of snapshots. The total number of unknowns must not exceed the number of measurements. For instance, if we have N_b baselines, N_p phase screen pixels, N_i intensity pixels, and S snapshots, the following inequality must hold for the problem to have a solution:

$$SN_b \geq N_i + S(N_p - 1). \quad (5.13)$$

Next, a method for computing the Cramér-Rao lower bounds for this estimation problem will be outlined.

5.3.2 Derivation of the Cramér-Rao bounds

The *Cramér-Rao lower bound* (CRLB) for an estimation problem gives the lower bound on the mean square error associated with each estimated parameter [55]. The CRLB is independent of the method used to solve the problem. Let $\boldsymbol{\theta}$ be the vector of unknown intensities and phase screen pixels. Since the estimator in this case is unbiased, the CRLB is given by the relation

$$E\left[(\theta_i - \hat{\theta}_i)^2\right] \geq [\mathbf{F}^{-1}]_{ii} \quad (5.14)$$

where $\hat{\theta}_i$ is the estimated value of the i^{th} element of $\boldsymbol{\theta}$, and where the right-hand side is equal to the i^{th} diagonal element of the inverse of the *Fisher information matrix* \mathbf{F} . The element at the j^{th} row and k^{th} column of the Fisher information matrix is given by

$$[\mathbf{F}]_{jk} = -E\left[\frac{\partial^2 L(\mathbf{v} | \boldsymbol{\theta})}{\partial \theta_j \partial \theta_k}\right] \quad (5.15)$$

where $L(\mathbf{v} | \boldsymbol{\theta})$ is the log-likelihood function of the vector of visibilities given the vector of unknowns. For the situation at hand, the likelihood function is given by the multivariate complex Gaussian probability density function

$$P(\mathbf{v} | \boldsymbol{\theta}) = \frac{1}{\det[\pi \mathbf{K}(\boldsymbol{\theta})]} \exp\left[-(\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta}))^H \mathbf{K}^{-1}(\boldsymbol{\theta})(\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta}))\right], \quad (5.16)$$

where $\bar{\mathbf{v}}(\boldsymbol{\theta})$ is the vector of the mean values of the visibilities computed according to the model given in Eq. (5.12), and $\mathbf{K}(\boldsymbol{\theta})$ is the covariance matrix associated with the mean or model visibilities. The symbol $\det[\]$ denotes the matrix determinant. The log-likelihood function is then given by

$$L(\mathbf{v} | \boldsymbol{\theta}) = -(\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta}))^H \mathbf{K}^{-1}(\boldsymbol{\theta})(\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta})) - \log(\det[\mathbf{K}(\boldsymbol{\theta})]) - N \log(\pi) \quad (5.17)$$

where N is the number of rows in $\mathbf{K}(\boldsymbol{\theta})$. In order to compute the Fisher matrix, Eq. (5.15) tells us that we need to take partial derivatives of Eq. (5.17) with respect to all of the parameters in $\boldsymbol{\theta}$. The diagonal entries of the inverse of the Fisher matrix are then equal to the CRLBs.

In the process of deriving a relation for Eq. (5.15), several identities will be needed. These identities are straightforward to derive and are listed below for reference.

$$\frac{\partial \mathbf{K}^{-1}}{\partial \theta_j} = -\mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial \theta_j} \mathbf{K}^{-1} \quad (5.18)$$

$$\frac{\partial}{\partial \theta_j} \{\log(\det[\mathbf{K}])\} = \text{tr} \left(\mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial \theta_j} \right) \quad (5.19)$$

$$\frac{\partial}{\partial \theta_j} \{\mathbf{x}^H \mathbf{A} \mathbf{x}\} = \text{Re} \left\{ \mathbf{x}^H \mathbf{A} \frac{\partial \mathbf{x}}{\partial \theta_j} \right\} + \mathbf{x}^H \frac{\partial \mathbf{A}}{\partial \theta_j} \mathbf{x} \quad (5.20)$$

$$E[\mathbf{x}^H \mathbf{A} \mathbf{x}] = \text{tr}(\mathbf{A} E[\mathbf{x}^H \mathbf{x}]) \quad (5.21)$$

The symbol $\text{tr}(\cdot)$ denotes the trace of the matrix. Using Eqs. (5.17)-(5.21), the expression for the j^{th} row and k^{th} column of the Fisher matrix becomes

$$[\mathbf{F}]_{jk} = 2 \text{Re} \left\{ \frac{\partial \bar{\mathbf{v}}^H(\boldsymbol{\theta})}{\partial \theta_j} \mathbf{K}^{-1}(\boldsymbol{\theta}) \frac{\partial \bar{\mathbf{v}}(\boldsymbol{\theta})}{\partial \theta_k} \right\} + \text{tr} \left(\mathbf{K}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{K}(\boldsymbol{\theta})}{\partial \theta_k} \mathbf{K}^{-1}(\boldsymbol{\theta}) \frac{\partial \mathbf{K}(\boldsymbol{\theta})}{\partial \theta_j} \right) \quad (5.22)$$

The entries in the vector of model visibilities $\bar{\mathbf{v}}$ are given by Eq. (5.12), and the entries in the covariance matrix \mathbf{K} are given by Eq. (4.5). Computation of the Fisher matrix is complex because of the nature of unknowns $\boldsymbol{\theta}$ and the form of Eq. (5.12). The derivatives of the model visibilities, which also appear in the expressions for the entries in the covariance matrix, differ depending on whether the parameter is an intensity or a phase. Recall that the phases are functions of angle. That dependence must be calculated and kept track of during the derivative calculations. In some instances, the phases in Eq. (5.8) or (5.12) may be from the same pixel and thus cancel out, making the derivative disappear. Also, one must note that one phase screen pixel in each snapshot must be held constant in order for the Fisher matrix to be full rank. This is because when all phase pixels are written as differences from that reference phase, the reference phase cancels from Eq. (5.8). In addition, when a complex gain screen is used, one must set one gain amongst all the snapshots, or one intensity pixel equal to a constant. Without this measure, there is an ambiguity in level that results in a rank deficient Fisher matrix. An example of this ambiguity can be shown if one divides all gains by one the gains which is equivalent to scaling the intensity distribution by

the square of that gain. Computer programs to compute the Fisher matrix were constructed in MATLAB. The results of those computations will be discussed in the next section.

5.3.3 Cramér-Rao lower bound results and discussion

In this section, the results of the CRLB computations for the image recovery estimation problem will be discussed. Results from both the phase screen and the complex gain screen cases will be presented. A large amount of computation time is required to compute the bounds. This is due in part to the use of MATLAB, which is somewhat slow in executing code with heavy use of looping, and partly due to the complex nature of the problem. Due to the computation time necessary, only a small number of the many possible parameter values were tried. However, enough were computed so that the behavior of the bounds could be adequately characterized.

To begin, one must note that this estimation problem requires that the screen be sufficiently excited by the source distribution. This means that if there are only a small number of point sources in the intensity distribution, the Fisher matrix becomes rank deficient. One way around this problem is to place very small intensities at all angle positions and then add in the larger intensities of interest. The small intensities can be at a level such that they are below the noise floor associated with the particular integration time of interest.

An intensity distribution containing two point sources of equal strength, along with intensity values at $1/1000^{\text{th}}$ of that value at all other angles, was used for many CRLB computations to test the effect of the different parameters. These parameters include the number of snapshots, the number of screen pixels, the range from the array to the screen, and the number of intensity pixels. The intensity distribution was normalized to have unity total flux. Changing the integration time of interest will have no effect on the bounds other than to shift the absolute level of the intensity bounds. Therefore, an integration time of 0.1 s was used for all results discussed here. The results of the CRLB computations for the phase screen case where the intensity distribution described above was used with a total of 20 intensity pixels, four snapshots, five phase pixels, and a range of 8 m are shown in Fig. 5.9.

The intensities correspond to parameter numbers 1-20, and the phases for all snapshots correspond to numbers 21-36. The phase screen pixels were chosen at random, according to a uniform distribution from 0 to 2π rad. The resultant standard deviation of the phase screen pixels averages about 0.1 rad. The standard deviation of the intensity values at the location of the two point sources is about $1/12^{\text{th}}$ the actual intensity value. Clearly these bounds are quite good and suggest that one can do very well when using these parameters. Several computations were made of the bounds for this case, all with different random phase screen pixel values. They showed that the intensity bounds were constant, but that the phase bounds changed by a small amount depending on the values chosen. The average standard deviation of the phase error remained at 0.1 rad.




Fig. 5.9: The square root of the CRLB for all parameters for the case where a phase screen was used. There were 20 intensities, 4 snapshots, 5 phase pixels, and a distance of 8 m to the screen. The intensity bounds correspond to parameter numbers 1-20 and the phase bounds to 21-36.

Still concentrating on the phase screen case and using the intensity distribution discussed above, computations to explore the effect of changing the number of snapshots were made. The CRLBs were computed for numbers of snapshots equal to 2, 4, 6, and 10. The effect of the number of snapshots is as one would expect. The resultant standard deviation of the intensity error decreased as the square root of the number of snapshots. The level of the phase error remained unchanged. When the distance from the array to the phase screen is changed, there appears to be no effect on the bounds, at least for distances ranging from 2 to 10 m. Changing the total number of intensity pixels also does not appear to have a significant effect on the bounds. Intensity pixel counts of 14, 20, 26, 32, and 38 were tried.

For some cases the phase pixel estimate bounds become very large. One such case is when the number of phase pixels is increased for the intensity profile described above. For 5 and 10 phase pixels the bound remains small with average standard deviations of 0.11 and 0.22 rad, respectively. However, when the number of phase screen pixels is increased to 13, the average standard deviation rises to 0.88 rad, and by the time 15 or 20 pixels are used it rises to about 1.6 rad. The bounds for the intensities remain essentially constant, seemingly unaffected by the increase in the number of phase pixels. It is not clear whether this behavior of the phase pixel bounds means that one would be unable to estimate the phases, or whether the error has increased to a point where the CRLB is no longer providing representative results [56]. This hypothesis can be tested when, in the next section, we attempt to obtain maximum likelihood solutions to the estimation problem. Another case where the phase errors become even larger is when the number of significant point sources is increased. In the case where there are three point sources, the average standard deviation of the phases is 0.55. But when four point sources are present, the standard deviation rises to 2.58, and with six it jumps to 33.6. The bounds on the intensities remain reasonable, being in line with what one would expect as the number of sources increases.

It was mentioned above that the SNR associated with the CRLB for a phase screen for the two point source case was approximately 12. The theoretical SNR value corresponding to an integration time of 0.4 s given in Fig. 4.3, for the standard synthesis imaging case, is approximately 5.5. When we compare these two numbers we see that the CRLB SNR is a little more than twice what we would normally expect using synthesis

imaging. Additionally, one can compute the CRLB for the case where there is no phase screen in the model. In this case the associated SNR at the two point source locations is the same as is calculated for the two point source case with a phase screen present.

Cramer-Rao lower bounds were also computed for the complex gain screen case. In computing these bounds, one gain factor was set equal to unity per the discussion in the last section. This gain factor must be one of the central pixels in the screen in order for the bounds to be reasonable in level. The performance indicated by the CRLBs for the complex gain screen is lower than for the phase screen case. This should be expected, as the model here is of increased complexity. An example CRLB result for the same setup as was used in Fig. 5.9 is shown for the complex gain screen case in Fig. 5.10. The intensity, phase, and gain bounds are indicated on the graph. The third gain in the first snapshot was fixed. The SNR equivalent to the intensity bounds at the point source locations is about 8.5 which is lower than for the phase screen case, but still above what one would expect with standard synthesis imaging. The phase bounds always appear as being large. It is expected that this is the same phenomenon that occurs for some configurations for the phase screen case and therefore they are not believed to accurately represent the bounds. The bounds for the gain parameters indicate that the bounds for the outer pixels in the screen are high while those for the center pixels are much lower. This behavior is likely due to the fact that the central pixels appear much more often in the visibility computations and thus there is more information available with which to estimate them. The bounds for the complex gain screen behave in a manner similar to that for the phase screen case as a function of the parameters mentioned above.

It seems that in many situations, the CRLBs for this estimation problem are favorable. Whether these bounds can be achieved is uncertain. If an estimator is unbiased and consistent, then the maximum likelihood approach can achieve the CRLB [57]. Although our estimator is unbiased, or at least asymptotically so, it is not consistent. The estimator is not consistent because the estimation error will never decrease to zero, no matter how many snapshots are used. This is because each snapshot carries with it an additional set of unknown screen pixels that need to be estimated. In the next section, a discussion of

attempts that have been made to solve the estimation problem using maximum likelihood and least-squares approaches will be presented.




Fig. 5.10: The square root of the CRLB for all parameters for the case where a complex gain screen was used. There were 20 intensities, 4 snapshots, 5 phase pixels, and a distance of 8 m to the screen. The intensity bounds correspond to parameter numbers 1-20, the phase bounds to 21-36, and the amplitudes to 37-55.

5.3.4 Initial solution attempts and results

Because the bounds discussed in the last section were, in general, very encouraging, some initial attempts have been made to obtain maximum likelihood and least-squares solutions to the estimation problem. This section will present these initial solutions and discuss the method's applicability to field data.

First, attempts were made to solve the estimation problem on synthetic data sets that were constructed using the single phase screen model. The negative of the log-likelihood

function given in Eq. (5.16) was minimized subject to the constraint that the intensity pixels be nonnegative. A constrained optimization routine available within MATLAB was used to complete the optimization. This routine, which uses a sequential quadratic programming (SQP) method, was used as a quick alternative to implementing a similar method from scratch [58]. The initial condition for the intensity distribution was obtained by first applying the CLEAN algorithm to each image in the data set, averaging those images, and then resampling the average image to conform to the discrete intensity distribution model being used. It was found that the best results were obtained when the initial intensity distribution was scaled down by an order of magnitude or more. It seems that if the model visibilities and the measured visibilities start close in magnitude, the optimization routine tended to continually increase the model intensities. Because the first term of the log-likelihood function is relatively immune to changes in the total flux, the value of the negative of the log-likelihood function can remain the same or even decrease when the intensities go off toward infinity. The determinant term of the log-likelihood function would eventually stop the increase in the intensities. The phase screen pixel values were either initialized to be all zero, or randomly distributed over the interval 0 to 2π rad. A typical result of this optimization is shown in Fig. 5.11, where the normalized intensity distributions of the optimization result and the initial condition are shown. Note that the initial condition represents the best quality image obtained for the set of data using the traditional direct method of imaging based on the Van Cittert-Zenike theorem. The case shown in Fig. 5.11 corresponds to the case for which the CRLB is shown in Fig. 5.9. Shown in Fig. 5.12 are the phase screen pixel values corresponding to the synthetic data and those found through the optimization process. The average phase error in Fig. 5.12 is 0.7 rad.

Better performance was obtained using methods not based on the log-likelihood function. A weighted least-squares cost function provided the best results. This cost function takes the form

$$C = (\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta}))^H \mathbf{W} (\mathbf{v} - \bar{\mathbf{v}}(\boldsymbol{\theta})) \quad (5.23)$$

where \mathbf{W} is a diagonal matrix that can be used to weight each visibility in the cost function. As will be shown, the weighting becomes important when field data is used. To find the minimum of Eq. (5.23) with respect to the unknown intensities and phases, two methods

were tried. The first method tried was the previously mentioned MATLAB routine. This produced results on synthetic data sets that were far superior to those obtained using the log-likelihood approach. Even better results were obtained when yet another optimization routine was applied. With this method, the cost function was minimized one variable at a time. All the unknowns were cycled through until the change in the cost function value was less than a specified small value. The optimization was constrained by requiring that the intensity values be nonnegative. Figure 5.13 shows the result of using this routine, along with the weighted least-squares cost function, on the same synthetic data set that was used for the results shown in Fig. 5.11. The actual and estimated phase screen values are shown in Fig. 5.14. The average phase error here is 0.53 rad.




Fig. 5.11: The result of applying the maximum likelihood method to 4 snapshots of synthetic data distorted by a phase screen containing 5 pixels located 8 m from the array. The initial image, representing the result of traditional interferometry is shown with the dotted line, and the result of the optimization is indicated with the solid line.




Fig. 5.12: The phase screen pixel values corresponding to the solution shown in Fig. 5.11. The model (dotted) and resulting from the optimization process (solid) are shown. The average phase error is 0.7 rad.




Fig. 5.13: The solution (solid line) found using the least-squares approach, for the same synthetic data as used for Fig. 5.11.




Fig. 5.14: The model and computed phase screen pixel values corresponding to the solution shown in Fig. 5.13. The average phase error is 0.53 rad.

In the last section, several combinations of intensity distributions and numbers of phase screen pixels produced phase bounds that were very large. The optimized result for one of those cases is found in Fig. 5.15. In this case, there were five phase screen pixels per snapshot and four point sources. The average standard deviation for the phases was 2.58 rad according to the CRLB computations. The average phase error for the solution shown in Fig. 5.15 is 0.56 rad, which is significantly lower than the bound. Therefore, it seems reasonable that the conjecture involving the inaccuracy of the CRLB in these cases is valid.

An example result for the case where a complex gain screen is used is shown in Fig. 5.16. In this case a screen with 5 pixels, 4 snapshots, and a range of 8 m was used. The result shows degraded performance compared to the phase screen case, as was predicted by the bounds. The average phase error for this case was 0.65 rad. The error in the amplitude factors was large, especially for the outer pixels. This shows the importance of phase in the reconstruction as indicated by the relatively good reconstruction possible with poor amplitude information. The amplitude factors will be more important when we attempt to

deal with the scintillation seen in field data. The least-squares cost function and the round-robin optimization method were again used for this example.




Fig. 5.15: The traditional interferometric image (dotted line) and the least-squares solution (solid) for a synthetic data case with four point sources. The average phase screen pixel error is 0.56 rad.




Fig. 5.16: Solution (solid line) sound using the least-squares method for synthetic data produced using a complex gain screen. The corresponding error in the screen phases was 0.65 rad.

Two simple cases were investigated to see what happens when the model used to recover the image is less complex than the model used to create the data. The first case is where the number of phase screen pixels is reduced during the recovery. The case studied here is where 13 pixels are used to recover a set of images distorted by 15 pixels. The second is where one phase screen is used to recover an image distorted by two phase screens. Here, a single phase screen located at 5 m with 5 pixels is used to recover an image distorted by two phase screens located at 5 and 10 m with 5 pixels each. The results for these two cases are found in Figs. 5.17 and 5.18, respectively. Both figures show some improvement in image quality over the traditional method. As one would assume that the atmosphere will not be fully modeled by the screen parameter approach, these two results give one some hope that the method can be successfully applied to real field data.




Fig. 5.17: The result of attempting to recover an image from a synthetic data set distorted with 15 phase screen pixels using only 13 pixels. The two point sources are located at intensity pixel numbers 9 and 12.




Fig. 5.18: The result of attempting to recover an image using a single phase screen with 5 pixels located at 5 m when the data set was distorted using two phase screens, at 5 and 10 m, with 5 pixels each. The two point sources are located at intensity pixel numbers 9 and 12.


Several sets of field data were used to test the effectiveness of this recovery method on real world data. All the results discussed here were computed using the least-squares cost function minimized using the alternate iterative method that was mentioned above. In general the results varied greatly. In many cases, real improvements in image quality were obtained. In other cases, little or no improvement resulted. In general the best results were obtained when a complex valued screen was used. It is thought that the gains allow the model to better compensate for the significant changes in the total flux from one snapshot to the next as well as for variations with angle. The weighting, in the form of the matrix W in Eq. (5.22), that tended to produce the best results was where the inverse of the square of the total flux measurement for each snapshot weighted the visibilities of the corresponding snapshot. This was applied in order to make sure that the one or two most energetic snapshots did not dominate the solution. One might argue that the most energetic snapshots may well be the most important; however, experience has shown that the results are generally

better with the weighting in place. One might also apply a weighting equal to the inverse of the square of the magnitude of each individual visibility. This makes all visibilities equally important. It was found that this method did not generally produce results that were as good as with the flux weighting. The measured data for the examples shown in this section consisted of a varying number of 0.1-s snapshots taken from the same field data that was discussed in Chapter 3. In general, the snapshots were chosen from a larger set of data in such a way that the individual snapshots were not consecutive in time. The time gap between them was generally 0.3 to 1 s. This was done to obtain greater variety in the data. The best results are generally obtained when the phase pixel values are initialized to zero. The initial intensities, as long as they are not too large in reference to the measured visibilities, do not seem to influence the final solution.


An example of a successful recovery is shown in Fig. 5.19. The data consisted of six snapshots at 800 Hz with an associated wind speed of 11.7 mph. Two solutions are shown in Fig. 5.19. Figure 5.19(a) shows the result when a five-pixel phase screen is used, while (b) shows the result when a five-pixel complex valued screen is used. Both screens were located at 5 m. The left point source should have an intensity of approximately 0.8, which is nearly correctly given by the solution in (b). There were 25 intensity pixels. This number was chosen because it closely matches the theoretical resolution obtainable with the array and frequency used. Adding more intensity pixels sometimes tended to produce spurious peaks in the resulting image. Using fewer pixels gave a smeared representation. There also seems to be an optimal location for the screen. Placing the screen further away or closer to the array generally produced less favorable results.

Two additional examples of image enhancement are shown in Figs. 5.20 and 5.21. Again, each data set consisted of six snapshots. The associated wind speeds were 5.7 and 9.7 mph, respectively. In Fig. 5.20 the right point source should have an intensity of 0.76, and in Fig. 5.21, the left point source should have a level of 0.54. Five complex valued screen pixels were used for the case in Fig. 5.20, and seven for Fig. 5.21. In both cases 25 intensity pixels were again used. The screens were located at 2 and 5 m, respectively. In most cases, a good solution was possible with careful choice of the number of screen pixels and the distance from the screen to the array. At this point, no real guidelines are available

for the choice of these parameters. Therefore, human intervention in the process is necessary.



(a)



(b)

Fig. 5.19: Successful examples of image recovery from field data using (a) a phase screen with 5 pixels located at 5 m, and (b) a complex valued screen with 5 pixels located at 5 m. The average wind speed associated with the data was 11.7 mph. The left point source should have an intensity of 0.79. There were 25 intensity pixels.




Fig. 5.20: Another successful recovery example. A complex valued screen at 2 m with 5 pixels was used. The associated wind speed was 5.7 mph. The right point source should have an intensity of 0.76. There were 25 intensity pixels.




Fig. 5.21: Another successful recovery example. A complex valued screen at 5 m with 7 pixels was used. The associated wind speed was 9.7 mph. The left point source should have an intensity of 0.54. There were 25 intensity pixels.

Although the screen parameter estimation method seems promising for application to real outdoor acoustic imaging, the results shown in this section are only preliminary. There is much more testing that can be done and many possible avenues that could be followed for further performance gains. The method of optimization is one area where further work could provide performance benefits. The maximum likelihood method was abandoned here because of difficulties in its interaction with the optimization routine used. It is possible that other optimization routines might fare better. The maximum likelihood solution model contains information, in the form of the covariance matrix, that the least-squares method does not. It is therefore possible that with the right optimization algorithm, the maximum likelihood approach could perform as well as or better than least-squares. For an image recovery algorithm like the one presented here to have reliable application to real image data, much more testing is needed. Guidelines on the choice of parameters such as the number and placement of phase screen pixels would be very important. For example, when a larger number of screen pixels are used, the level of error given by the cost function always goes down. However, the quality of the solution, in terms of matching the real physical case, does not improve, but rather degrades. This behavior is analogous to the case of fitting a polynomial to a discrete set of data points. When too many degrees of freedom as given by a large-order polynomial are allowed, the fitting error can become very small, but the function may have wild oscillations in the areas between samples. So, one challenge with this image recovery method is to find that optimal number of pixels, or degrees of freedom, that is appropriate for the specific data set in question. Two methods that may help in this respect involve letting the number of screen pixels become a variable of the optimization. First, we could add a term to the cost function in Eq. (5.23) that would be an increasing function of the number of pixels. This would penalize solutions with larger number of pixels. Another, more promising method is cross-validation. With this method, the idea is to leave out data points, or visibilities in this case, one at a time and to choose the number of pixels for which the missing data points are best predicted by the remainder of the data [59]-[60]. In addition to the parameter choice issues, the effect of initial conditions on the result of the optimization and the possible existence of multiple solutions must be more fully investigated.

6. CONCLUSIONS

In this thesis, the state of the art for acoustic imaging through the atmosphere has been described. An overview of imaging techniques and methods was given. These methods include both the classical narrowband interferometric process and the wideband tomographic solution methods. Examples of the performance of the wideband methods on both synthetic and field data were given. Good performance was found using two different methods borrowed from the tomographic imaging community. In addition to examples, a detailed discussion of the performance issues related to the narrowband technique was presented.

Central to the performance of narrowband acoustic interferometry through the atmosphere is the effect of the atmosphere. In addition to providing imaging data, the field experiment data also allowed for the measurement of several important physical characteristics of the atmosphere. Among these were the wave structure function, coherence loss, and the propagation transfer function. From these measurements we see that the atmosphere is an unfriendly place to image acoustic sources. The atmosphere has a very short coherence time, on the order of 0.5 s or less. Acoustic wavefronts suffer from substantial coherence loss as they propagate through the atmosphere. Using the field data, a compilation of data were computed that show the effect of the atmosphere on the variance of visibility measurements and on the quality and signal-to-noise ratio of the resulting images. These data show that the visibility variance can be increased by a factor of ten as compared to the theoretical. Image SNR values also suffer greatly, often becoming less than unity, even for modest wind speeds. In order to produce more accurate images, long integration times are required. Large-scale effects of the

atmosphere require that the integration time be on the order of minutes. Even with these images, the resolution damping effect of coherence loss is still present.

Methods for enhancing the quality of the images produced through the atmosphere were investigated. First, the widely used and successful self-calibration family of methods from the radio astronomy community were considered. It was shown that these methods, which are based on the concept of phase closure, cannot be applied to the case of acoustic imaging through the atmosphere. This is because the atmospheric distortion is a function of source coordinates, and because the methods are sensitive to measurement noise which is much more pronounced here than in radio astronomy where they deal with much higher bandwidths. As an alternative to self-calibration, an estimation problem approach was taken in which a discrete version of the source intensity is estimated along with the parameters of a distorting screen representing the atmosphere. Several short integration period images or snapshots are used, where the screen parameters are allowed to vary for each snapshot. Performance bounds, in the form of the Cramér-Rao lower bound, were calculated. These bounds were encouraging and so maximum likelihood and least squares techniques were employed to solve several image reconstructions, both with synthetic and field data. When applied to field data, it was necessary to make the distorting screen complex valued, which helped to account for significant variations in the power associated with the visibility measurements from one snapshot to the next. The least-squares approach was found to be the most successful, but perhaps only because of the optimization routine used. One can think of this method, especially the case where a complex valued screen is used with the least-squares cost function, as an angle-dependent version of the self-calibration method. Several examples of significant gains in image quality were presented. As was discussed, there is much work that can be done to further investigate this image enhancement procedure.

REFERENCES

- [1] J. Billingsley and R. Kinns, "The acoustic telescope," *Journal of Sound and Vibration*, vol. 48, no. 4, pp. 485-510, 1976.
- [2] Y. Biraud, B. Escudié, A. Hellion, and J. Munier, "Formation of acoustic images from spatial coherence functions: Processing by superresolution techniques," in *Image Formation From Coherence Functions in Astronomy*, vol. 76, *Proceedings*, C. van Schooneveld, Ed., Dordrecht: D. Reidel Publishing Company, 1979, pp. 299-313.
- [3] M. Chiollaz, B. Escudié, Y. Biraud, and G. Pachiaudi, "Interferometric acoustic imaging, joint representation and image deconvolution," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, pp. 1485-1488, 1982.
- [4] B. Escudié, M. Chiollaz, and E. Parent de Curzon, "Interferometric acoustic imaging of railway noise," *Journal of Sound and Vibration*, vol. 120, no. 2, pp. 303-310, 1988.
- [5] M. M. Boone and A.J. Berkhout, "Theory and applications of a high-resolution synthetic acoustic antenna for industrial noise measurement," *Noise Control Engineering Journal*, vol. 23, no. 2, pp. 60-68, Sept.-Oct. 1984.
- [6] J.D. Van Der Toorn, H. Hendriks, and T.C. Van Den Dool, "Measuring TGV source strength with SYNTACAN," *Journal of Sound and Vibration*, vol. 193, no. 1, pp. 113-121, 1996.
- [7] G.W. Swenson, Jr., and N.C. Mathur, "The interferometer in radio astronomy," *Proceedings of the IEEE*, vol. 56, no. 12, pp. 2114-2130, Dec. 1968.
- [8] A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr., *Interferometry and Synthesis in Radio Astronomy*. Malabar: Krieger Publishing Company, 1991.
- [9] R.N. Bracewell, *The Fourier Transform and its Applications*. New York: McGraw-Hill, 1978.
- [10] R.A. Gabel and R.A. Roberts, *Signals and Linear Systems*. New York: John Wiley and Sons, 1987.

- [11] J.A. Högbom, "Aperture synthesis with a non-regular distribution of interferometric baselines," *Astron. Astrophys. Supp.*, vol. 15, pp. 417-426, 1974.
- [12] U.J. Schwarz, "Mathematical-statistical description of the iterative beam removing technique (method CLEAN)," *Astron. Astrophys.*, vol. 65, pp. 345-356, 1978.
- [13] U.J. Schwarz, "The method "CLEAN" - Use, misuse, and variations," in *Image Formation from Coherence Functions in Astronomy*, C. van Schooneveld, Ed., Dordrecht: D. Reidel Publishing Company, 1979, pp. 261-275.
- [14] J.E. Conway, T.J. Cornwell, and P.N. Wilkinson, "Multi-frequency synthesis: A new technique in radio interferometric imaging," *Mon. Not. R. astr. Soc.*, vol. 246, pp. 490-509, 1990.
- [15] G.W. Swenson, Jr., and N.C. Mathur, "On the space-frequency equivalence of a correlator interferometer," *Radio Sci.*, vol. 4, pp. 69-71, 1969.
- [16] J.E. Conway and R.J. Sault, "Multi-frequency synthesis," *Astron. Soc. Pac. Conf. Ser.*, vol. 82, pp. 309-325, 1995.
- [17] R.J. Sault and M.H. Wieringa, "Multi-frequency synthesis techniques in radio interferometric imaging," *Astron. Astrophys. Suppl. Ser.*, vol. 108, pp. 585-594, 1994.
- [18] S. Haykin, Ed., *Array Signal Processing*. Englewood Cliffs: Prentice-Hall, 1985, pp. 351-428.
- [19] C. C. Paige and M. A. Saunders, "LSQR: An algorithm for sparse linear equations and sparse least squares," *ACM Transactions on Mathematical Software*, vol. 8, no. 1, March 1982, pp. 43-71.
- [20] A.H. Andersen and A.C. Kak, "Simultaneous algebraic reconstruction technique: A new implementation of the ART algorithm," *Ultrasonic Imaging*, vol. 6, Jan. 1984.
- [21] H. Stark and J. W. Woods, *Probability, Random Processes, and Estimation Theory for Engineers*. Englewood Cliffs: Prentice-Hall, 1994, pp. 292-296.
- [22] J. B. Thomas, *An Introduction to Statistical Communication Theory*. New York: John Wiley and Sons, 1969, pp. 63-64.
- [23] J.E. Piercy, T.F.W Embleton, and L.C. Sutherland, "Review of noise propagation in the atmosphere," *J. Acoust. Soc. Am.*, vol. 61, no. 6, pp. 1403-1418, 1977.
- [24] G.A. Daigle, "Effects of atmospheric turbulence on the interference of sound waves above a finite impedance boundary," *J. Acoust. Soc. Am.*, vol. 65, no.1, pp. 45-49, 1979.

- [25] G.A. Daigle, "Correlation of the phase and amplitude fluctuations between direct and ground-reflected sound," *J. Acoust. Soc. Am.*, vol. 68, no. 1, pp. 297-302, July 1980.
- [26] T.F.W. Embleton, "Sound propagation outdoors: Improved prediction schemes for the 80's," *Noise Control Engineering Journal*, vol. 18, no. 1, pp. 30-39, 1982.
- [27] A.N. Kolmogorov, "The local structure of turbulence in incompressible viscous fluids for very large Reynolds' numbers," in *Turbulence, Classic Papers on Statistical Theory*, S.K. Friedlander and L. Topper, Ed., New York: Wiley-Interscience, 1961, pp. 151-155.
- [28] A. Ishimaru, *Wave Propagation and Scattering in Random Media*. Vol. 2, New York: Academic Press, 1978, pp. 336-338.
- [29] G.A. Daigle, T.F.W. Embleton, and J.E. Piercy, "Propagation of sound in the presence of gradients and turbulence near the ground," *J. Acoust. Soc. Am.*, vol. 79, no. 2, March 1986.
- [30] M.C. Roggemann and B. Welsh, *Imaging Through Turbulence*. Boca Raton: CRC Press, 1996, pp. 60-62.
- [31] D.K. Wilson et al., "Experimental determination of the effective structure-function constant parameter for atmospheric turbulence," *J. Acoust. Soc. Am.*, vol. 105, no. 2, pt. 1, pp. 912-914, 1999.
- [32] G.A. Daigle, J.E. Piercy, and T.F.W. Embleton, "Line-of-sight propagation through atmospheric turbulence near the ground," *J. Acoust. Soc. Am.*, vol. 74, no. 5, pp. 1505-1513, 1983.
- [33] H.E. Bass, L.N. Bolen, R. Raspet, W. McBride, and J. Noble, "Acoustic propagation through a turbulent atmosphere: Experimental characterization," *J. Acoust. Soc. Am.*, vol. 90, no. 6, pp. 3307-3313, 1991.
- [34] D.K. Wilson, "Calculated coherence and extinction of sound waves propagating through anisotropic, shear-induced turbulent velocity fluctuations," *J. Acoust. Soc. Am.*, vol. 105, no. 2, pt. 1, pp. 658-671, 1999.
- [35] J.W. Benson, "Atmospheric propagation transfer function determination using an adaptive filter," *J. Acoust. Soc. Am.*, vol. 104, no. 6, pp. 3681-3684, 1998.
- [36] B. Widrow and S.D. Stearns, *Adaptive Signal Processing*. Englewood Cliffs: Prentice-Hall, 1985.
- [37] W. K Jenkins, A.W. Hull, J.C. Strait, B.A. Schnauffer and X. Li, *Advanced Concepts in Adaptive Signal Processing*. Boston: Kluwer Academic Publishers, 1996, pp. 54-68.

- [38] D.K. Wilson and D.W. Thomson, "Acoustic propagation through anisotropic, surface-layer turbulence," *J. Acoust. Soc. Am.*, vol. 96, no. 2, pt. 1, pp. 1080-1095, 1994.
- [39] J.W. Goodman, *Statistical Optics*. New York: John Wiley & Sons, 1985, p. 167.
- [40] S.R. Kulkarni, "Self-noise in interferometers: Radio and infrared," *The Astronomical Journal*, vol. 98, no. 3, pp. 1112-1130, 1989.
- [41] R.L. Mickelson and G.W. Swenson, Jr., "A comparison of two correlation schemes," *IEEE Transactions on Instrumentation and Measurement*, vol. 40, no. 5, pp. 816-819, Oct. 1991.
- [42] J.G. Proakis, C.M. Rader, F. Ling, and C.L. Nikias, *Advanced Digital Signal Processing*. New York: Macmillan, 1992, pp. 41-44.
- [43] G.M. Jenkins and D.G. Watts, *Spectral Analysis and Its Applications*. San Francisco: Holden-Day, Inc., 1968.
- [44] J.M. Noble, H.E. Bass, and R. Raspet, "The effect of large-scale atmospheric inhomogeneities on acoustic propagation," *J. Acoust. Soc. Am.*, vol. 92, no. 2, pt. 1, pp. 1040-1044, 1992.
- [45] The Mathworks Inc., *MATLAB Function Reference, Volume 1: Language*, 1999.
- [46] A. van der Sluis and H.A. van der Vorst, "Numerical solution of large, sparse linear algebraic systems arising from tomographic problems," in *Seismic Tomography*, G. Nolet, Ed., Dordrecht: D. Reidel Publishing Company, 1987.
- [47] R. Jennison, "A phase sensitive interferometer technique for the measurement of the Fourier transforms of spatial brightness distributions," *Mon. Not. R. Astron. Soc.*, vol. 118, pp. 176-284, 1958.
- [48] A.E.E. Rogers, H.F. Hinteregger, A.R. Whitney, C.C. Counselman, I.I. Shapiro, J.J. Wittels, W.K. Klemperer, W.W. Warnock, T.A. Clark, L.K. Hutton, G.E. Marnandino, B.O. Ronnang, and O.E.H. Rydbeck, "The structure of radio sources 3C 273B and 3C 84 deduced from the "closure" phases and visibility amplitudes observed with three-element interferometers," *The Astrophysical Journal*, vol. 193, pp. 293-301, 1974.
- [49] R.Q. Twiss, A.W.L. Carter, and A.G. Little, "Brightness distribution over some strong radio sources at 1427 Mc/s," *Observatory*, vol. 80, pp. 153-159, 1960.
- [50] D.H. Rogstad, "A technique for measuring visibility phase with an optical interferometer in the presence of atmospheric seeing," *Applied Optics*, vol. 7, no. 4, pp. 585-588, 1968.

- [51] A.C.S. Readhead and P.N. Wilkinson, "The mapping of compact radio sources from VLBI data," *The Astrophysical Journal*, vol. 223, pp. 25-36, 1978.
- [52] T.J. Cornwell and P.N. Wilkinson, "A new method of making maps with unstable radio interferometers," *Mon. Not. R. astr. Soc.*, vol. 196, pp. 1067-1086, 1981.
- [53] T.J. Pearson and A.C.S. Readhead, "Image formation by self-calibration in radio astronomy," *Ann. Rev. Astron. Astrophys.*, vol. 22, pp. 97-130, 1984.
- [54] K.A. Marsh, J.M. Richardson, and J.P. Martin, "Application of the phase closure technique to passive acoustic imaging through inhomogeneous media," in *Acoustical Imaging, Proceedings of the 14th International Symposium*, New York: Plenum Press, 1985, pp. 133-140.
- [55] H.V. Poor, *An Introduction to Signal Detection and Estimation*. New York: Springer-Verlag, 1994.
- [56] H.L. Van Trees, *Detection, Estimation, and Modulation Theory, Part 1: Detection, Estimation, and Linear Modulation Theory*. New York: John Wiley and Sons, 1968, pp. 63-72.
- [57] D.H. Johnson and D.E. Dudgeon, *Array Signal Processing, Concepts and Techniques*. Englewood Cliffs: Prentice Hall, 1993, pp. 267-268.
- [58] The Mathworks Inc., *MATLAB Optimization Toolbox User's Guide*, 1999.
- [59] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215-223, May 1979.
- [60] R. Whatmough, "Applying generalized cross-validation to image restoration," *ICASSP-94, 1994 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 5, pp. 453-5, 1994.

VITA

Jonathan William Benson was born on December 25, 1969, in Chicago Heights, Illinois. He received a bachelor's degree in electrical engineering in May of 1992 from the University of Illinois at Urbana-Champaign. He began graduate studies immediately thereafter and received the M.S. degree in January of 1994, also from the University of Illinois. His thesis was entitled "A low-frequency directional microphone system," and his research was directed by Professor George W. Swenson, Jr. He then began work for the U.S. Army Construction Engineering Research Laboratories in Champaign, Illinois. In this position he continued to work in the fields of acoustics and signal processing through the development of noise monitoring equipment and participation in acoustic propagation field experiments. He returned to the University of Illinois in the fall of 1995 to begin work toward the Ph.D. His thesis research was once again under the direction of Professor Swenson and his thesis was entitled "Acoustic coherence imaging through the atmosphere." His Ph.D. was granted in May 2000. He currently resides in Champaign, Illinois, with his wife Michelle.