

Chapter 1

Introduction

There are many practical problems in which it is required to detect and characterise hidden structures or remote objects by virtue of the scattered acoustic or electromagnetic fields they generate. It remains an open question, however, as to which reconstruction algorithms offer the best images for a given set of field measurements. Classical beamforming, while conceptually simple, does not perform well in situations where data are limited, and unfortunately, many existing and prospective applications for ultrasonic imaging involve limited data.

This thesis is primarily concerned with the application of a general, though computationally intensive method of image reconstruction to ultrasonic imaging, with particular emphasis on three-dimensional (3-D) problems in the non-destructive evaluation (NDE) of materials. Known as the maximum entropy method, and henceforth referred to as MaxEnt, it has previously been applied successfully to many problems involving incomplete data [1-7], although applications in ultrasonic imaging have been limited by perceived incompatibilities between complex-valued imagery and the positivity constraint imposed by MaxEnt theory. In addition, due to their complexity, MaxEnt calculations have traditionally been the province of supercomputer-class machines [8]. Until very recently, therefore, applications to real-world problems have been beyond the reach of most researchers.

To add to these problems, there was initially considerable confusion and controversy regarding the rationale and limitations of MaxEnt. As with any newly developed scientific theory dealing with fundamental questions, much of the early reasoning behind MaxEnt was based more on philosophy than mathematics. Sensing vagueness in the philosophy, many researchers were reluctant to use MaxEnt in place of conventional methods of data analysis without solid justification.

Today, the philosophical and mathematical underpinnings of MaxEnt, namely Bayesian probability and decision theory, are far more developed and less controversial than they were in the 1960s and '70s. From a pragmatic standpoint, it is also more widely appreciated that MaxEnt *does* yield excellent results in many practical circumstances, so the search for new applications and the development of old ones are continuing enterprises.

Although the *raison d'être* of this thesis relates primarily to the field of NDE, the principles developed are applicable to a variety of inverse scattering problems involving homogeneous media and planar apertures. Only reflection mode imaging is considered, as this enables comprehensive treatment of the commonly used NDE geometries while retaining links with other applications such as medical and underwater imaging. In reflection mode imaging, data must be acquired from one face of an object, and this is a limitation encountered often in practice. Although the equations describing acoustic and electromagnetic propagation are sufficiently similar for the methods developed here to be applicable to some microwave imaging problems, further consideration of these problems is beyond the scope of this thesis.

1.1 Entropy and Inference

Mathematical problems in the real world usually fall into two classes, or at least reflect the characteristics of these classes to varying degrees. Firstly, there are those problems for which descriptive equations can be formulated and solved according to the rules of algebra. These are problems of *deduction*, and for a particular set of assumptions and approximations, they have only one uniquely correct solution. A second class of problems exists, however, which defy straightforward calculation, regardless of the numerical accuracy available. These problems must instead be viewed as problems of *inference*, and the vast majority of practical *inverse problems*; that is, problems involving indirectly observed phenomena, belong to this class.

The distinction between problems of deduction and problems of inference is primarily due to a mathematical pathology known as non-uniqueness. In cases where the mapping from image space to data space is incomplete; that is, whenever elements of the image structure are not transmitted in the data, the solution to the corresponding inverse problem is non-unique. As a further consequence of this non-uniqueness, the solution that appears optimal in the sense of agreeing with the data

most closely may depend as much on noise as on the data itself. For these reasons, a robust and consistent mode of reasoning is required whereby *uncertainty* is correctly handled. Until recently, it was unclear what this mode of reasoning could possibly be, although the technological imperatives of real-world problems gave rise to a collection of *regularisation* techniques in its absence [9-11]. Each of these had the *ad hoc* justification of enforcing desirable characteristics of *smoothness* and *stability* in solutions, yet no single technique emerged as universally applicable or successful.

In the late 1940s, amidst considerable progress in the field of information theory, the importance of *entropy* as a fundamental measure of uncertainty was established by Shannon [12]. Whereas previously, entropy had been afforded only a thermodynamic interpretation, its significance in information theory was far more profound. Given a set of discrete, independent and exhaustive probabilities p_i ($i = 1, \dots, N$), the *Shannon entropy* is given by¹

$$S = - \sum_{i=1}^N p_i \ln p_i \quad (1.1)$$

where the p_i can be associated with any ensemble of outcomes, meaning the results of scientific experiments, or even the intensities of individual pixels which collectively comprise an image. In a probabilistic sense, S measures the uncertainty associated with any such ensemble, which is also a negative measure of its information content. It must be noted though, that as defined here, both information and uncertainty are strictly mathematical constructs, having only abstract connections to what human beings regard as knowledge and ignorance.

Jaynes [13] was one of the first to see the possibilities of incorporating this new abstract measure into the old regularisation formalism. In the absence of complete information, he reasoned, inferences drawn from probability distributions of maximal entropy (maximum uncertainty) should be favoured over all others, and hence the maximum entropy principle found its modern incarnation².

In the history of signal processing, the introduction of maximum entropy methods, starting with the auto-regressive power spectral estimate due to Burg [16], marked a departure from established thinking. Whereas previously, signal

¹ Natural logarithms are assumed unless otherwise specified. This has little bearing on the definition of entropy.

²Actually, the principle had been elucidated several times throughout history, but its significance was not appreciated in the orthodox sciences until the advent of high-speed computing [14, 15].

processing had essentially involved the manipulation of electrical analogues of physical quantities, the availability of fast digital computers allowed abstract concepts such as entropy to be brought to bear on practical problems. In recent years, MaxEnt has evolved considerably, attaining the status of an axiomatically justified, self-consistent framework for inference, applicable to any *positive, additive distribution* (PAD). Optical images fit the description of such distributions well, as image intensity is always positive and additive. In other areas, however, the application of MaxEnt has not been without controversy.

In nuclear magnetic resonance (NMR) spectroscopy for example, free induction decay (FID) waveforms are now routinely processed via MaxEnt [17-19], yielding spectra of superior resolution and signal-to-noise ratio (SNR) relative to Fourier transform processing. This is despite the fact that FID waveforms are coherent by nature, and their spectra can possess arbitrary phase, depending on the mode in which nuclei are excited. Given the arbitrary phases of the spin excitations and the interference effects associated with coherent wave phenomena, it is difficult to imagine how NMR spectra can be associated with positive, additive distributions, but this has not impeded the effectiveness of MaxEnt as a form of regularisation. Given the success of MaxEnt in NMR, it is a natural step to apply it to other problems that involve coherent waves - such as ultrasonic inverse scattering.

1.2 Inverse Scattering

Ultrasonic waves in the megahertz frequency range propagate with low levels of attenuation in many optically opaque media of interest to science, industry and medicine. Inhomogeneities in the mechanical parameters of these media, namely elasticity and density, give rise to *scattered* waves, which then convey information regarding their shape and physical properties.

Unfortunately, due to multiple scattering and other second-order phenomena, the exact equations governing scattering are non-linear [20, 21]. This makes the treatment of either *forward* or *inverse* scattering problems extremely complex unless first-order solutions, excluding all but primary scattering, are employed. Fortunately, such approximations yield acceptable results in many practical situations, and in this sense, the forward problem of calculating scattered fields from known incident fields and scatterers can be regarded as mathematically well posed.

The same cannot be said, however, for inverse scattering problems, in which it is required to compute images of scatterers from measurements of scattered fields. Even with first-order approximations, these problems demand considerable mathematical subtlety for their satisfactory solution. Particularly in circumstances where the measurements are noisy or incomplete, classical image reconstruction theory can fail to yield images of acceptable resolution and clarity. It is important, therefore, to explore alternative approaches to solving inverse scattering problems.

1.3 Non-Destructive Evaluation

NDE is an important technology that supports the life management of components and structures subject to stress or harsh environments. Ultrasound responds to the mechanical and elastic parameters of materials and therefore yields information directly related to the strength and remaining life of structures [22]. For this reason, and due to the inherent safety and ease with which it can be generated and detected, ultrasound can be considered one of the most widely used modalities in NDE [23].

Compared with ionising radiation, ultrasound possesses other major advantages in terms of sensitivity to important types of defects. As X-rays and γ -rays respond essentially to density variations through enhanced or attenuated transmission, certain metallurgical defects not associated with density variations, and cracks not aligned with the incident rays, can easily pass undetected³. In contrast, ultrasonic waves will generally undergo both specular reflection and diffraction from the boundaries of such defects. Even in the event that the specular reflections are directed completely away from the receiving aperture, the diffracted waves will radiate more or less isotropically, and be detectable over a wide range of orientations.

Figure 1.1 depicts a common geometry for data acquisition in NDE, wherein a single broadband transducer collects data over a one-dimensional (1-D) synthetic aperture A-B. Digitised radio frequency (RF) waveforms proportional to the scattered field are known as A-scans (Amplitude-scans) and comprise the basic data unit of broadband ultrasonic imaging. To aid in localising scatterers, it is also important that scattered wavefields are sampled spatially. This can be achieved as

³ The DC10 crash in Sioux City, Iowa, USA in 1989 was attributed to a "hard- α " defect in an engine turbine blade, which precipitated crack growth and catastrophic failure of the engine. The characteristic of such a defect is that it has the same density as the surrounding material, but is significantly harder and hence more brittle [24].

illustrated in Figure 1.2 by acquiring multiple A-scans at successive locations. The resulting "B-scan" displays alternate light and dark shades corresponding to the positive and negative cycles of the underlying A-scan waveforms.

Echo signals corresponding to diffracted wave-fronts from the ends of the crack depicted in Figure 1.1 are observed as a pair of hyperbolic curves, reflecting the variation in arrival times across the aperture. Due to the orientation of the hypothetical crack in this instance, specular reflections are not evident in Figure 1.2, however B-scans generally contain a mixture of reflected and diffracted waves. This makes systematic estimates of defect size, shape and orientation extremely difficult without the application of some kind of image reconstruction algorithm.

Beyond simply detecting the presence of material defects, ultrasonic imaging techniques are capable of discriminating between cracks, which may pose a threat to structural integrity, and more benign defects such as gas bubbles and inclusions, which should be monitored periodically. High-resolution ultrasonic image reconstruction is therefore steadily gaining importance in quantitative NDE.

In analogy with optical imaging, the simplest and most direct route to forming images from ultrasound data is via an acoustic lens. Sound waves originating from the focal zone of the lens are then automatically compensated for their arrival times and summed coherently to yield a single reflectivity estimate at a point. Unfortunately, a transducer equipped with such a lens must be mechanically scanned to form an image, and depending on the depth of field of the lens, only a narrow range of depths will be correctly focussed. For these reasons among others, the use of physical lenses is often impractical and undesirable in ultrasonic imaging [22]. There has, therefore, been considerable interest in methods of *synthetically* focussing or reconstructing ultrasonic images from discrete A-scan measurements.




Figure 1.1: Typical data acquisition geometry for ultrasonic NDE.




Figure 1.2: Simulated B-scan corresponding to Figure 1.1.

1.4 A Review of Classical Beamforming

1.4.1 The Time-Domain Approach

The oldest and most rudimentary method of achieving synthetic focussing with discrete arrays of transmitters and receivers is classical *beamforming*, otherwise known as "delay-and-sum" beamforming, or simply the synthetic aperture focussing technique (SAFT) [25-28]. The basic idea is to impose delays on transmitted and/or received signals such that either the sound energy of the transmitter or the spatial sensitivity of the receiver is concentrated in a particular location or direction. In each case, the desired signals interfere constructively, while unwanted signals combine incoherently. This is completely analogous to the operation of a physical lens.

Figure 1.3 shows how time-domain beamforming can be implemented in hardware. This is the approach usually taken in medical phased-array systems, which are able to provide near-ideal cross-sectional images of soft tissue by dynamically steering and focussing pulsed ultrasonic beams in two dimensions.




Figure 1.3: Schematic representation of time-domain beamforming.

For simplicity, Figure 1.3 is indicative of plane wave transmission and reception, which relates primarily to far field imaging. In the more general near field case, the focussing delays must dynamically match the curvature of wavefronts over the range of depths to be imaged, and this greatly complicates implementation. It is still useful to review far field beamforming however, as many of the performance limiting aspects in this regime are characteristic of linear beamforming in general. Specifically, it is of interest to characterise linear beamforming in terms of its inherent point spread function (PSF), which is a function of both the array beam pattern, and the bandwidth of the transmitted pulses.

Mathematically, time-domain beamforming is symmetrical in respect to the way beams are synthesised in transmission and reception. In transmission, elementary waves physically coalesce into a single wavefront by virtue of the elastic linearity of the medium. In reception, time-domain beamforming is simply an analogue of this linearity, with the summations implemented via linear amplifier circuitry or digital arithmetic. This thesis is mainly concerned with focussing in reception, due to the relative freedom with which received signals can be processed, and the significant gains in image quality that stand to be made by applying advanced signal processing in a *virtual* computational space. This contrasts with the case of beamforming in transmission, where relatively few opportunities exist for enhancing the fundamental physics of propagation in *physical* media.

1.4.2 The Fourier-Domain Approach

While time-domain beamforming is conceptually simple and relatively easily implemented in hardware, many aspects of its performance are far from optimal. To gain insights into these deficiencies, it is useful to resort to the Fourier-domain, wherein the entire beamforming process can be viewed equivalently in terms of standard linear systems theory. Of course, in addition to providing an alternative interpretation for ultrasonic image reconstruction, the Fourier-domain has been profitably exploited in its own right for many years as a means of efficiently implementing beamforming and image reconstruction algorithms [26, 27].

A prerequisite to beamforming in the Fourier-domain is that all time-dependent fields are first subjected to *temporal* Fourier transformation. This results in a set of complex-valued time-harmonic fields, which can be individually processed in the

spatial Fourier-domain and then recombined to yield broadband solutions if required. Figure 1.4 shows the arrangement of the classic direction of arrival (DOA) estimation problem, where a time-harmonic plane wave from an unknown direction is incident on an equally spaced linear array of N isotropic receivers⁴.




Figure 1.4: The classic direction of arrival (DOA) estimation problem.

For the case depicted in Figure 1.4, simultaneous coherent samples⁵ from the array would yield samples of a second wave with a *spatial frequency* given by

$$f_x = \frac{1}{\lambda} \sin \theta \quad (\text{m}^{-1}) \quad (-90^\circ < \theta < 90^\circ) \quad (1.2)$$

where λ is the wavelength of the original plane wave in the medium and θ is its angle of incidence in the sense indicated. According to the Fourier optics interpretation, time-harmonic waves can be decomposed into *angular spectra* of plane waves via Fourier transformation [29]. Subjecting the complex spatial samples to a discrete spectral estimation procedure would therefore ideally yield a single spectral peak, proportional in amplitude to the incident plane wave. The performance of the fast Fourier transform (FFT) in achieving this expected response is shown in Figure 1.5 for a plane wave incident at $\theta = 0^\circ$ on an array of N elements arranged at the critical spacing of $\lambda/2$. To illustrate the effect of aperture size on the results, FFT spectra have been computed for $N=8, 16, 32$ and 64 .

⁴ It will be assumed for the present that the array elements are equally sensitive in all directions.

⁵ Referring to the measurement of both real and imaginary parts of a complex signal.




Figure 1.5: FFT spectra corresponding to a plane wave incident at $\theta = 0^\circ$ on a $\lambda/2$ sampled N -element linear array⁶. (a) $N=8$; (b) $N=16$; (c) $N=32$; (d) $N=64$.

The resolution of an imaging system is a fundamental parameter indicative of the scale at which structure can be retrieved from images. Of particular interest in Figure 1.5 are the widths of the main lobes of the beam patterns, as these are indicative of the lateral (angular) resolution obtained. For coherent imaging systems, resolution is difficult to define exactly, due in part to its dependence on the relative phases of the sources imaged [29, 30]. Nevertheless, the angular separation between the main beam and the first zero of the $\sin(x)/x$ (*sinc function*) response is found to be a useful indicator of overall resolving power. This resolution, corresponding to the *Rayleigh* resolution for rectilinear apertures, is given by

$$\theta_R = \sin^{-1}\left(\frac{\lambda}{A}\right), \quad (1.3)$$

or expressed in Cartesian co-ordinates as a function of the range z [29],

⁶ As the ultimate aim here is to reconstruct images, linear amplitude scales have been used.

$$d_R = \frac{\lambda z}{A} \quad (\text{m}). \quad (1.4)$$

From Equation (1.4), it might appear that for a sufficiently large aperture A , that an arbitrarily high resolution could be achieved. From Equation (1.2), however, it is apparent that wavefields originating in the far field are limited to spatial frequencies in the range $\pm 1/\lambda$, which places a lower bound of $\lambda/2$ on the spacing at which distinct scatterers can be resolved [31]⁷. Despite this second limitation, it is usually the F-number of an array (z/A) that determines its resolution.

Compared with many other types of imaging, the resolution obtainable from ultrasonic imaging is severely limited by long wavelengths and small apertures [32–34]. Whereas in absolute terms, the aperture sizes available in ultrasonic imaging are comparable or even sometimes larger than those available in optical imaging, ultrasonic wavelengths are typically three orders of magnitude longer than in the optical case. For a given aperture size, therefore, image resolution can be expected to be three orders of magnitude poorer with ultrasound. This problem is highlighted in Table 1.1, which lists in descending order, the typical aperture sizes (in wavelengths) associated with various optical and ultrasonic imaging systems.

Imaging System		Aperture Size
Optical ⁸	Two inch optical telescope	$1 \times 10^5 \lambda$
	Camera lens in bright sunlight	$1 \times 10^4 \lambda$
	Human eye in bright sunlight	$3 \times 10^3 \lambda$
Ultrasonic	50 cm underwater imaging array at 1 MHz	$3 \times 10^2 \lambda$
	5 cm medical array imaging in tissue at 7.5 MHz	$2 \times 10^2 \lambda$
	5 cm NDE array imaging in steel at 5 MHz	$4 \times 10^1 \lambda$

Table 1.1: A comparison of typical aperture sizes of various imaging systems.

⁷ This neglects the presence of evanescent waves, which will be discussed further in Chapter 2.

⁸ Data from Tsao and Steinberg [35].

In view of the data in Table 1.1, it is apparent that resolution is a commodity in especially short supply in ultrasonic NDE, where (longitudinal) sound velocities in metals are of the order of 6000 ms^{-1} , and the available apertures span only a few tens of wavelengths. While it is generally desirable to use the highest frequencies possible to optimise resolution, the attenuation of ultrasound increases with at least the square of frequency [22], and the resulting degradation in signal-to-noise ratio (SNR) sets an upper limit to the useful range of frequencies. This is particularly true of materials with coarse grain structures at the wavelength scale [36].

Super-resolution, meaning any improvement in resolution beyond the Rayleigh limit of Equation (1.3), has been the subject of intensive research in ultrasonics and other imaging related fields in recent years. Unfortunately, despite some promising results [31, 33, 37], there have been very few examples of true, stable super-resolution in practice, and many researchers are perhaps justifiably sceptical as to the prospects of achieving it in the future. Rihaczek and Hershkowitz [30] state their position bluntly: "despite significant efforts in this area, true superresolution methods have not worked in a practical sense for the past thirty years, and we do not expect them to work in the next thirty years."

In addition to aperture and wavelength limited resolution, other defects identifiable from Figure 1.5 are the side-lobes, which are characteristic of beamforming with truncated data. When it is of interest to resolve the shape of extended objects, or detect weakly scattering objects in the presence of strongly scattering ones, side-lobes pose major problems [22]. Whereas the amplitude of the sinc function response is approximately 13 dB down from the main lobe at the peak of the first side-lobe, constructive interference can lead to total obliteration of image structure at amplitudes well above this.

A crude method of controlling the side-lobe levels in linear beamforming (again borrowed from spectral estimation) is to *apodise* the array samples - tapering them smoothly to near zero at the extremities of the aperture. By eliminating sharp transitions between the measured samples and the unknown (assumed zero) samples outside the aperture, the spectral response, and hence the side-lobe distribution, becomes less oscillatory. As an example, Figure 1.6 was computed from the same parameters as Figure 1.5, except that a Gaussian apodising function with a standard deviation of $N/4$ was applied to the data prior to FFT processing.




Figure 1.6: FFT spectra of Gaussian-apodised, N -element critically-sampled array data for a plane wave incident at $\theta = 0^\circ$. (a) $N=8$; (b) $N=16$; (c) $N=32$; (d) $N=64$.

The most obvious differences between Figures 1.5 and 1.6 are that the peak side-lobe levels have decreased by approximately an order of magnitude in the latter. This has occurred, however, at the expense of increasing the 6-dB width of the main lobes by approximately 50% in each case. This result must be expected, as by diminishing the contribution of data near the ends of the aperture, the effective aperture size has been reduced. In situations where spatial data samples are easily obtained, as in synthetic aperture imaging from a single transducer, apodisation may represent an effective means of side-lobe control. If spatial samples are obtained from a physical array of receivers however, discarding, or partially discarding a significant proportion of them is highly undesirable.

In regard to apodisation and beamforming in general, the advice of E.T. Jaynes [38] is clear: "Once one is committed to using a computer to analyse the data, the high-resolution imaging problem is completely changed. What one has tried to do in the past by fancy optical and mechanical engineering (apodizing, image stabilizing) can be done better, and at a small fraction of the cost by the computer."

1.4.3 Three-Dimensional Beamforming

To this point, linear beamforming has been discussed in the context of one-dimensional arrays, which are capable of imaging in a single plane only. In many applications, including NDE, there is often no guarantee that the objects of interest lie in the reconstructed plane, or even that they are two-dimensional. For scatterers with unknown shape and orientation then, the best solution is to carry out full three-dimensional image reconstruction, which requires data to be acquired in two dimensions (2-D). Conceptually, this represents no difficulty, as beamforming theory extends naturally to 2-D, as illustrated by Figure 1.7. Technically, however, 3-D imaging presents huge complications, as providing N^2 as opposed to N elements in a 2-D array is out of the question for all but the smallest arrays, and mechanically scanning a transducer over a 2-D aperture is very time-consuming.




Figure 1.7: A 3-D wavefront incident on a 2-D array.

It is well known that the number of elements in an array can be reduced to a small fraction of the number required to completely fill the aperture without significantly affecting the shape of the main beam [39]. Sparsely sampled apertures, therefore, offer significantly improved resolution relative to filled arrays, however, this improved resolution comes at the expense of additional high-amplitude side-lobes [40]. Two extremes in this trade-off are depicted in Figure 1.8 for filled and sparse-random configurations consisting of 64 elements each.




Figure 1.8: Sampling patterns and associated far field beam patterns for 2-D arrays consisting of 64 elements arranged on a 64×64 critically sampled grid. (a) 8×8 filled array; (b) Beam pattern for (a); (c) Random array; (d) Beam pattern for (c).

The chief advantage of random element configurations is that the aperture size can be increased and hence the resolution improved without the high-amplitude grating lobes typically associated with uniformly spaced under-sampled arrays [41]. According to Equation (1.3), the angular resolution of the uniform array in Figure 1.8(a) in either the x - z or y - z planes would be approximately 14.5° , whereas that of the sparse array in Figure 1.8(c) would be approximately 1.8° . Unfortunately, it is also observed from Figure 1.8(d) that the average side-lobe level of the sparse-random array is significantly higher. This high side-lobe level would substantially degrade the dynamic range of images obtained, and given the small number of array elements, apodisation would not yield any significant improvement [35].

1.5 MaxEnt versus Classical Beamforming

From the preceding discussion, it is evident that the deficiencies of classical beamforming manifest whenever the quality of the data is compromised. Lateral image resolution is thus directly related to aperture size, while side-lobe levels are related to the aperture sampling pattern. In broadband imaging systems, which will be the subject of further discussion in Chapter 3, range-resolution is related to the bandwidth of the interrogating waveform, which is also limited in practice by medium-dependent attenuation and the resonant characteristics of ultrasonic transducers [42, 43]. In the context of the earlier discussion of inverse problems, all of these defects can be attributed to the simplistic manner in which classical beamforming deals with incomplete data. Additional problems related to noise have not been mentioned here, but will be addressed in subsequent chapters.

As demonstrated in many successful applications to date [2, 44-47], MaxEnt represents a fundamentally different approach to image reconstruction. By framing inverse scattering problems in terms of regularisation theory and invoking the maximum entropy principle, seemingly unrelated problems of poor resolution from small apertures and image distortion from missing samples will be seen to take on the more general character of problems of inference from incomplete information.

1.6 Thesis Contributions

The primary contribution of this thesis is in advancing a new method of reconstructing ultrasonic images that is more robust against the effects of incomplete data than classical beamforming and other conventional methods. The maximum entropy principle itself is by no means new, and several researchers to date have noted that complex-valued data at least appear to benefit from MaxEnt processing. This work extends all previous results, however, in that it verifies the benefits of MaxEnt signal processing in ultrasonic image reconstruction while demonstrating tractable algorithms for solving large-scale multi-dimensional problems.

The secondary contribution of this thesis is in successfully demonstrating novel Fourier-domain sampling concepts for 2-D array imaging. While the array constructed was very simple and inexpensive, it enabled surprisingly good 3-D images to be obtained by virtue of the MaxEnt reconstruction algorithm.

Other specific aspects of this work that are claimed as novel are listed below.

- A matrix operator formalism for linearised inverse scattering that specifically includes projection, backprojection and sparse sampling operators.
- The application of a composite image MaxEnt algorithm to complex-valued image reconstruction, proving that there is more than one way of reconstructing complex-valued images while adhering to the MaxEnt formalism.
- The adaptation of the Cambridge entropy metric for positive images to complex images, superseding the requirement for composite image approaches.
- The achievement of true, stable super-resolution in both range and cross-range dimensions of ultrasonic images under realistic experimental conditions.
- The demonstration of side-lobe clutter suppression with filled and random arrays, and grating lobe suppression in the case of uniformly under-sampled arrays.
- The adaptation of the coherent CLEAN algorithm due to Tsao and Steinberg [35] to broadband ultrasonic imaging in two and three dimensions.

1.7 Thesis Structure

In approximate proportion to the time spent during research, this thesis is split into five main sections, comprising Chapters 2 to 6. Following from this introduction, Chapter 2 deals with the modelling of the underlying physics of propagation and scattering, and the derivation of a matrix model for time-harmonic scattering. The cases considered are limited to impenetrable and penetrable scattering, as these are of generic importance in NDE. Although the linear equations developed to describe the forward scattering problem are quite simple, they are demonstrated to yield accurate results in comparison with a well-known analytic solution.

In Chapter 3, mathematical concepts applicable to the solution of ill-posed image reconstruction problems are reviewed, with particular emphasis on the forward scattering model proposed in Chapter 2. The traditional "least-squares" framework

for image reconstruction is introduced as a prelude to the unifying concepts of Bayesian probability theory, of which MaxEnt is identified as a special case. To allow comparisons with subsequent MaxEnt algorithms, a small number of additional algorithms, including the non-linear CLEAN algorithm from radio astronomy [35, 48] are applied to solving a simulated 3-D image reconstruction problem.

Chapter 4 deals with the development of practical MaxEnt algorithms for ultrasonic image reconstruction. Starting with the basic subspace optimisation strategy developed by the “Cambridge group” [2, 44-46, 49-51], algorithms are developed for broadband image reconstruction from sparsely sampled aperture data. Two methods are proposed; the former of which is based on a multi-channel approach, which is similar in some respects to methods previously used for phasing complex-valued NMR spectra [17]. The second method is a stabilised adaptation of the Cambridge entropy metric to the complex domain that is computationally simpler than the multi-channel approach, though yielding identical results.

Chapter 5 applies the above MaxEnt algorithms to data obtained in synthetic aperture experiments. The effects of sparse aperture sampling on image quality are investigated relative to the matched filter and CLEAN approaches discussed in Chapter 3. As noted earlier, the general thrust of this work is towards solving 3-D problems, although the preponderance of 2-D imaging techniques in use today necessitated the inclusion of at least one 2-D experiment. Due to the difficulty of obtaining suitable solid test specimens containing well-characterised scatterers, all of the 3-D imaging experiments were conducted in immersion. The scatterers imaged in the 2-D experiment were side-drilled holes in an aluminium plate.

Chapter 6 is a detailed case study of the design and construction of a two-dimensional polyvinylidene fluoride (PVDF) array transducer based on novel Fourier-domain sampling concepts and MaxEnt signal processing. The array offers advantages over conventional point-sampled arrays in terms of construction simplicity and signal-to-noise ratio (SNR). A simple differential multiplexer design based on CMOS analogue technology is presented, whereby the array can be quickly scanned without moving parts. Three-dimensional images of several test objects obtained via the array are presented and discussed.

Chapter 7 summarises the findings of the thesis and discusses future research directions, including methods for improving the robustness of MaxEnt algorithms against potential errors in modelling propagation phenomena.

Chapter 2

Physics and Modelling


The aim of this chapter is to review the underlying physics of acoustic scattering and to translate the continuous domain acoustic equations into the discrete domain. The scattering models developed are linear and accurate to first-order; representing the scalar behaviour of homogeneous media, as opposed to their full vector behaviour. The necessary assumptions and approximations are stated explicitly, and the salient mathematical characteristics impacting on the inverse problem are highlighted.

The major outcome of this chapter is a matrix observation model that incorporates arbitrary insonification and sparse aperture sampling. To conclude the chapter, this observation model is used to simulate time-domain scattering in a simple canonical test case drawn from the field of NDE: the flat-bottom hole. Close agreement between the numerical and analytical solutions confirms the physical accuracy of the model to the level of the approximations made.

2.1 Mathematical Preliminaries

2.1.1 Co-ordinate System and Terminology

For clarity of explanation, the co-ordinate system to be employed will first be described in three-dimensions, as illustrated in Figure 2.1, and then generalised to m dimensions in subsequent equations. The image reconstruction problems to be addressed will be limited to two and three-dimensions ($m = 2,3$). In all cases, the x - y plane is assumed to be coincident with the free surface of an isotropic and lossless medium contained in the volume V_0 and bounded by the composite mathematical surface $S_1 + S_2$. The position vector $\mathbf{r} = (x_1, \dots, x_m)$ points to arbitrary co-ordinates within V_0 and the unit normal $\hat{\mathbf{n}}$ is directed outwards from V_0 , as shown below.



As the orientation of this work is towards *imaging* rather than *detecting* scatterers, it will be assumed that their locations are known approximately before any imaging experiments are conducted. A *region of interest* (ROI) can therefore be defined within the medium, as illustrated in Figure 2.2, for subsequent reconstruction from field measurements made in the aperture.

In order for the forward scattering problem to be properly modelled, it is important that the ROI contain *all* scatterers illuminated by the incident field. Fields in the receiving aperture contributed from outside the ROI constitute a modelling error, as the sources of these fields cannot be mapped to their correct locations during image reconstruction. The exception to this rule is that A-scan signals can be time-gated to select any prescribed depth range, thereby removing contributions from scatterers above and below the ROI. The range parameter R will be used to specify the offset between the aperture plane and the ROI, as illustrated in Figure 2.2.




Figure 2.2: The region of interest in relation to the aperture plane.

The transmitting and receiving apertures will be assumed to act independently, though for simplicity, each will be restricted to lie in the x - y plane. This restriction allows efficient subdivision of V_0 into rectilinear elements (pixels or voxels) in the discrete modelling phase, which is discussed later in this chapter. At this stage, the distribution of transmitting and receiving elements will remain undefined; however, it will be assumed that ultrasonic fields are transmitted and received over sets of coordinates labelled ξ and ξ' respectively.

2.1.2 Fourier Transform Definitions

As discussed in section 1.4.2, time-domain scattering problems can be formulated and solved as linear combinations of time-harmonic scattering problems. The prerequisite for dealing with time-harmonic problems is that all time-dependent fields must first be transformed into an equivalent set of time-harmonic fields. In the continuous domain, Fourier amplitude spectra $F(\omega)$ and time-domain pressure amplitudes $p(t)$ are related by the following pair of *temporal* Fourier transforms, stated in terms of the radian frequency ω and time t :

$$F(\omega) = \int_{-\infty}^{+\infty} p(t) \exp(j\omega t) dt \quad (2.1)$$

$$p(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} F(\omega) \exp(-j\omega t) d\omega \quad (2.2)$$

In defining *spatial* Fourier transforms, it is customary in the physics literature to reverse the sign convention used above so that the m -dimensional spatial spectra $F(\mathbf{k})$ and the time-harmonic pressure fields $p(\mathbf{r})$ are related by¹

$$F(\mathbf{k}) = \int_{-\infty}^{+\infty} p(\mathbf{r}) \exp(-j\mathbf{k} \cdot \mathbf{r}) d^m \mathbf{r} \quad (2.3)$$

$$p(\mathbf{r}) = \frac{1}{(2\pi)^m} \int_{-\infty}^{+\infty} F(\mathbf{k}) \exp(j\mathbf{k} \cdot \mathbf{r}) d^m \mathbf{k} \quad (2.4)$$

where the radian wavenumbers associated with individual co-ordinate directions are contained in the m -dimensional vector

$$\mathbf{k} = (k_1, \dots, k_m). \quad (2.5)$$

The background wavenumber in the medium is given by

$$k = \|\mathbf{k}\| = \frac{2\pi}{\lambda} = \frac{\omega}{c} \quad (\text{rad m}^{-1}) \quad (2.6)$$

where c is the velocity of wave propagation.

2.1.3 Scalar Wave Equations

In liquids and soft biological tissues, only longitudinal waves are able to propagate. In solids however, the situation is more complex, with a variety of wave modes possible. From the perspective of imaging, the most important of these are the longitudinal and transverse wave modes.

Although scattering models and even some inversion schemes have been formulated for multi-mode scattering [53, 54], longitudinal and transverse waves are treated here as uncoupled scalar fields. In practice, the two are easily separated by virtue of their widely differing velocities, and one or the other can be selected as input to image reconstruction algorithms depending on the particular application².

Mathematically, the separation of longitudinal and transverse wave modes is achieved through the Helmholtz decomposition [21, 22] of arbitrary vector displacement fields in terms of scalar and vector potentials respectively, according to

$$\mathbf{u}(\mathbf{r}, t) = \nabla \Phi(\mathbf{r}, t) + \nabla \times \vec{\Psi}(\mathbf{r}, t) \quad \nabla \cdot \vec{\Psi}(\mathbf{r}, t) = 0, \quad (2.5)$$

¹ The justification for this is that the directions of wave propagation are then automatically positive [21, 52].

where the del operator ∇ is defined in three dimensions to be

$$\nabla = \frac{\partial}{\partial x} \hat{\mathbf{i}} + \frac{\partial}{\partial y} \hat{\mathbf{j}} + \frac{\partial}{\partial z} \hat{\mathbf{k}}. \quad (2.6)$$

Given this decomposition, the homogeneous forms of the time-dependent wave equations for Φ and $\bar{\Psi}$ can be written in essentially the same form [21],

$$\left. \begin{aligned} \nabla^2 \Phi - \frac{1}{c_l^2} \frac{\partial^2 \Phi}{\partial t^2} &= 0 \\ \nabla^2 \bar{\Psi} - \frac{1}{c_t^2} \frac{\partial^2 \bar{\Psi}}{\partial t^2} &= 0 \end{aligned} \right\} \quad (2.7)$$

where the Laplacian operator ∇^2 is defined in three dimensions to be

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}. \quad (2.8)$$

In cases where the fields associated with the vector potential $\bar{\Psi}$ do not undergo significant depolarisation, it is acceptable to treat the Cartesian components of $\bar{\Psi}$ as scalar fields in a similar manner to Φ [21]. Inversion algorithms developed for longitudinal waves are then equally applicable to transverse waves after substituting the transverse wave velocity c_t , which is usually of the order of 60 % of the longitudinal velocity c_l in solids [22]. After Fourier transformation, either of Equations (2.7) can be expressed in the form

$$[\nabla^2 + k^2] \Phi(\mathbf{r}, \omega) = 0, \quad (2.9)$$

which is the basic Helmholtz wave equation governing time-harmonic scalar potentials in homogeneous media. As most experiments are carried out with piezoelectric transducers that respond directly to acoustic pressure, the equations to follow will be stated in terms of the time-harmonic pressure $p(\mathbf{r}, \omega)$, which in liquids is related to the time-harmonic scalar potential $\Phi(\mathbf{r}, \omega)$ via [22]

$$p(\mathbf{r}, \omega) = \omega^2 \rho \Phi(\mathbf{r}, \omega) \quad (\text{Nm}^{-2}). \quad (2.10)$$

This relation also holds approximately for isotropic solids.

² Reference [55] gives some reasons why transverse waves might be preferred in NDE.

2.1.4 Free-Space Green's Functions

As the intention here is to model propagation and scattering phenomena linearly, it is important to identify the impulse response of the medium, or in the terminology of image processing, the *point-spread-function* (PSF) that maps image space to data space. Green's functions model the spatial responses of media to point sources of excitation, and can thus be identified as point spread functions.

Mathematically, point sources of ultrasonic energy can be represented by Dirac delta functions of the form $\delta(\xi' - \mathbf{r})$, with $\mathbf{r}$ taken as the *source point* and ξ' as the *observation point* [20]. Substituting $\delta(\xi' - \mathbf{r})$ for the right hand side of Equation (2.9) (with a negative sign to simplify subsequent expressions) gives

$$[\nabla^2 + k^2] g_m(\xi' | \mathbf{r}, \omega) = -\delta(\xi' - \mathbf{r}), \quad (2.11)$$

where the solution $g_m(\xi' | \mathbf{r}, \omega)$ is known as the free-space scalar Green's function in m dimensions. The notation $\cdot | \cdot$ implies that $g_m(\xi' | \mathbf{r}, \omega)$ is generally dependent on both ξ' and $\mathbf{r}$, however, in all of the cases to be discussed, $g_m(\xi' | \mathbf{r}, \omega)$ can be taken purely a function of the norm $r = |\xi' - \mathbf{r}|$ [52, 56]. In image processing terminology, this property is known as shift-invariance, and it is a pre-requisite to the application of efficient numerical techniques such as fast convolution [57, 58].

Bleistein [52] gives the following general expression for $g_m(r, \omega)$:

$$g_m(r, \omega) = \frac{j}{4} \left[\frac{\omega}{2\pi r c} \right]^p H_p^{(1)}(\omega r / c) \quad p = (m-2)/2, \quad (2.12)$$

where $H_p^{(1)}$ is a p th order Hankel function of the first kind. Hankel functions with half integer values of p are expressible in terms of complex exponentials and powers, which are more convenient for numerical computation. For the case of $m = 2$ ($p = 0$), an asymptotic approximation is adequate for problems of the type considered here in which scatterers are located an appreciable number of wavelengths from the aperture ($r \gg \lambda$) [59, 60]. Table 2.1 summarises the exponential forms for $g_m(r, \omega)$, $m = (2, 3)$ used in numerical computations. While both expressions possess singularities at $r = 0$, these can be safely ignored provided $r > \lambda$.

Dimensions (m)	Green's Function	Condition
2	$g_2(r, \omega) \cong \frac{1}{2} \sqrt{\frac{jc}{2\pi\omega r}} \exp(j\omega r/c)$	$\omega r/c \gg 1$ $(r \gg \lambda)$
3	$g_3(r, \omega) = \frac{\exp(j\omega r/c)}{4\pi r}$	—————

Table 2.1: Exponential forms for the free-space scalar Green's functions.

2.2 First-Order Scattering Models

The scalar Green's functions just discussed form the basis of a linear, shift-invariant mapping between secondary sources of ultrasonic energy and scattered wavefields. It is now important to quantify the interactions between incident fields and scatterers that are responsible for the secondary sources in the first instance. Unfortunately, the situation is complicated by the fact that scatterers manifest in several different varieties. In anticipation of the types of scatterers encountered in NDE, the classes of *penetrable* and *impenetrable* scatterers will now be considered.

The necessity for different treatments for these classes arises from the different domains of validity for the assumptions required to linearise the scattering equations. Weakly inhomogeneous scatterers are penetrable in the sense that a proportion of the incident energy is allowed to propagate relatively unperturbed. Other types of scatterers, however, such as open cracks and voids, act more like perfect mirrors for ultrasonic waves, and are thus classified as impenetrable.

2.2.1 Scattering by Penetrable Objects

To develop a quantitative model for penetrable scattering, it is necessary to augment Equation (2.9) to its inhomogeneous form by adding a secondary source term. To first-order accuracy, the equations for elasticity and continuity of density lead to the following expression for the secondary source $o(\mathbf{r}, \omega)$ associated with a weakly scattering inhomogeneity in terms of the total pressure $p(\mathbf{r}, \omega)$ [20, 56, 61, 62]:

$$o(\mathbf{r}, \omega) = k^2 \gamma_\kappa(\mathbf{r}) p(\mathbf{r}, \omega) - \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega)]. \quad (2.13)$$

The two components of $o(\mathbf{r}, \omega)$ are *field dependent* monopole and dipole secondary source distributions [20, 62], which are functions of the normalised deviations from the background compressibility κ_0 and density ρ_0 , given respectively by [56, 61-63]

$$\gamma_\kappa(\mathbf{r}) = \begin{cases} \frac{\kappa(\mathbf{r}) - \kappa_0}{\kappa_0} & \mathbf{r} \in V_c \\ 0 & \mathbf{r} \notin V_c \end{cases} \quad (2.14)$$

and

$$\gamma_\rho(\mathbf{r}) = \begin{cases} \frac{\rho(\mathbf{r}) - \rho_0}{\rho_0} & \mathbf{r} \in V_c \\ 0 & \mathbf{r} \notin V_c \end{cases} \quad (2.15)$$

The distinction between these sources can be important, as dipole sources display an inherent directivity by virtue of the gradient operator in Equation (2.13). This makes the scattered field amplitudes dependent on the relative orientations of the transmitter and receiver, and affects the accuracy with which such secondary sources fit the approximation of shift-invariance already established for free-space Green's functions. In the detailed derivation in Appendix A, the directivity of the dipole sources is neglected, with the assumption that the incident and scattered fields propagate along nearly anti-parallel vectors. For co-located transmitting and receiving apertures, this assumption is valid, at least beyond the extreme near field, and both components of the secondary source distribution can be lumped into the single *reflectivity function* $f(\mathbf{r})$. In view of Equation (2.13), Equation (2.9) can be modified to the following inhomogeneous Helmholtz equation

$$[\nabla^2 + k^2] p(\mathbf{r}, \omega) = -o(\mathbf{r}, \omega), \quad (2.16)$$

where the total pressure $p(\mathbf{r}, \omega)$ is the sum of the incident and scattered pressures,

$$p(\mathbf{r}, \omega) = p_i(\mathbf{r}, \omega) + p_s(\mathbf{r}, \omega). \quad (2.17)$$

Through Equation (2.17), the intrinsic non-linearity of the scattering phenomenon is seen to be associated with the dependence of the scattered pressure field on the *total pressure*, which is itself the sum of incident and scattered pressures. In order to proceed further with Equation (2.17), some approximation is required to

linearise this dependence. The Born approximation achieves this aim by replacing $p(\mathbf{r}, \omega)$ with $p_i(\mathbf{r}, \omega)$; that is, it is assumed that the incident field is not significantly perturbed by the scatterer, as illustrated in Figure 2.3.




Figure 2.3: Penetrable scattering under the Born Approximation.

With further manipulation of Equation (2.17), as detailed in Appendix A, the scattered pressure field in the receiving aperture $p_s(\xi', \omega)$ can be related to the *reflectivity function* $f(\mathbf{r}) = \gamma_\kappa(\mathbf{r}) - \gamma_\rho(\mathbf{r})$ through the convolution integral

$$p_s(\xi', \omega) \cong k^2 \int_{V_0} f(\mathbf{r}) p_i(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \quad (2.18)$$

Using Equation (2.18), it is possible to compute $p_s(\xi', \omega)$, provided $f(\mathbf{r})$ and $p_i(\mathbf{r}, \omega)$ are known in advance, or by means of mathematical inversion, compute $f(\mathbf{r})$ from knowledge of $p_s(\xi', \omega)$ and $p_i(\mathbf{r}, \omega)$. Equation (2.18) is not *exact*, however, and the inherent approximations and assumptions are summarised below:

- Mode conversion effects have been neglected.
- Multiple scattering effects have been neglected.
- The directivities of dipole secondary sources have been neglected.
- Plane-wave approximations for $\nabla g(r, \omega)$ and $\nabla p_i(\mathbf{r}, \omega)$ have been employed.
- $p(\mathbf{r}, \omega)$ has been set to $p_i(\mathbf{r}, \omega)$ in accordance with the Born approximation.

As an alternative to the Born approximation, the Rytov approximation [56], which is based on an expansion of $p(\mathbf{r}, \omega)$ in terms of its complex phase, can also be used in some circumstances. As noted by Norton and Linzer [64] and Ishimaru [65], however, the Born approximation is the most applicable in relation to reflection mode problems, whereas the Rytov approximation is better suited to transmission mode problems. Additional justification for excluding the Rytov approximation in this instance relates to its requirement for unwrapped phases [61, 66], which would be difficult to calculate for the sparsely sampled apertures to be considered later in this chapter.

2.2.2 Scattering by Impenetrable Objects

The most typical scatterer encountered in NDE is the crack, or void of arbitrary shape. Provided the time-domain displacement $\mathbf{u}(\mathbf{r}, t)$ associated with the total pressure field $p(\mathbf{r}, t)$ is always less than the width of the scatterer, the characteristic surface S_c can be assumed to represent a Dirichlet (stress-free) boundary condition at all frequencies [20, 21, 67], whence

$$p(\mathbf{r}, \omega) = 0 \quad \mathbf{r} \in S_c. \quad (2.19)$$

In contrast, the acoustic impedance of some scatterers, given by the product of their density and acoustic velocity, may be sufficiently higher than that of the surrounding medium for them to be regarded as perfectly rigid. In such cases, Neumann boundary conditions apply [20, 21], whence it can be assumed

$$\nabla p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' = 0 \quad \mathbf{r} \in S_c. \quad (2.20)$$

In either of these cases, the reasoning by which impenetrable scatterers must be considered is different from that of the preceding analysis, chiefly because weak scattering approximations are inappropriate. Physical optics theory, in conjunction with Kirchhoff theory, however, offers an alternative linearising approximation [52, 68, 69]. Under this approximation, scatterers are assumed to behave as collections of discrete facets, and the scattering contribution of each facet is calculated via plane wave reflection theory as though it were part of an infinite surface. Individual scattering contributions are then summed coherently to arrive at the overall scattered field. Implicit in this approach is the notion that impenetrable scatterers can be separated into illuminated and shadow regions as illustrated in Figure 2.4.




Figure 2.4: Impenetrable scattering under the physical optics approximation.

According to Kirchhoff theory, the total wavefield *gradient* associated with the illuminated portion of a stress-free boundary is simply twice the value that would exist in the absence of the boundary [20, 22, 52], or

$$\nabla p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' \cong 2\nabla p_i(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' \quad \mathbf{r} \in S_c. \quad (2.21)$$

Conversely, in the case of a perfectly rigid scatterer, it is the pressure *amplitude* at the boundary that is doubled, giving

$$p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' \cong 2p_i(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' \quad \mathbf{r} \in S_c. \quad (2.22)$$

In either case, the total field in the shadow zone and its derivatives are taken to be zero. These approximations hold adequately for high frequencies, and provided other second order effects (listed below) are neglected, a linear scattering model results. Following the detailed derivation in Appendix B, the approximate scattering formula for both stress-free and rigid boundaries is found to be

$$p_s(\xi', \omega) \cong 2jk \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) f(\mathbf{r}) p_i(\mathbf{r}, \omega) d^m \mathbf{r}, \quad (2.23)$$

where the reflectivity function $f(\mathbf{r})$ has been substituted for the illuminated portion of the *singular function* to arrive at a similar notation to that used in Equation (2.18). It is important to note that although the reflection coefficient has been discussed as though it were necessarily a real number, in reality it depends on the surface admittance of the scatterer, which can be real, imaginary or complex [20].

As in the previous derivation of penetrable scattering, several approximations and assumptions have been made to arrive at Equation (2.23). The most significant of these are summarised below, with further details available in Appendix B.

- Multiple scattering has been neglected.
- Mode conversion effects have been neglected.
- Scatterers have been assumed smooth and convex.
- The total field intensity in the shadow zone has been assumed to vanish.
- Plane-wave approximations for $\nabla g(r, \omega)$ and $\nabla p_i(\mathbf{r}, \omega)$ have been employed.
- The surfaces of scatterers have been assumed not to interact; that is, $\mathbf{u}(t) < w$.
- The obliquity factors associated with individual facets have been neglected.
- Scattering boundaries have been assumed either stress-free or perfectly rigid.

While plane-wave reflection from infinite surfaces might seem unrepresentative of practical cases involving finite sized scatterers, it is found in practice that the combination of physical optics theory and Kirchhoff theory accurately models both reflected and diffracted waves [52]. This will be demonstrated through a model validation experiment at the end of this chapter.

2.2.3 Generalised Scattering Formula

Having derived scattering equations for the two important limiting cases, it can be seen that no one equation is generally applicable to modelling ultrasonic scattering. In particular, penetrable scattering is seen to have a frequency-squared dependence, while impenetrable scattering is only linearly dependent on frequency. Unfortunately, these two extremes also imply a continuum of intermediate cases, in which neither Equation (2.18) nor (2.23) would necessarily be accurate. Nevertheless, by expressing this frequency dependence as a *spectral function* $K(\omega)$ [60], Equations (2.18) and (2.23) can be re-written in the generalised form

$$p_s(\xi', \omega) \cong K(\omega) \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) f(\mathbf{r}) p_i(\mathbf{r}, \omega) d^m \mathbf{r}, \quad (2.24)$$

where for penetrable scatterers under the Born approximation,

$$K(\omega) = k^2, \quad (2.25)$$

and for impenetrable scatterers under the physical optics approximation

$$K(\omega) = 2jk. \quad (2.26)$$

The basic structure of Equation (2.24) is that of a Fredholm integral equation of the first kind [70, 71], with the reflectivity function $f(\mathbf{r})$ spatially modulated by the incident field $p_i(\mathbf{r}, \omega)$ prior to being convolved with the Green's function $g_m(|\xi' - \mathbf{r}|, \omega)$. The next two sections are concerned with highlighting the particular aspects of this equation that are responsible for the mathematical structure of first-order scattering and the information content of scattered fields.

2.2.4 Generalised Projection Theorem

Though the discrete observation model to be defined shortly will be tailored to the requirements of fast convolution in the Fourier domain, the characteristics of the mapping implied by Equation (2.24) will first be analysed in the spatial domain. In one dimension, the convolution between two functions $h(x)$ and $f(x)$ is defined as

$$d(x) = h(x) * f(x) = \int_{-\infty}^{+\infty} h(x - \alpha) f(\alpha) d\alpha, \quad (2.27)$$

where the bounds of the integral cater for functions with hypothetically infinite support. In real problems, it is usually only necessary to evaluate $d(x)$ over some finite domain ξ' (in this case the receiving aperture) and so the domains of $h(x)$ and $f(x)$ can also be restricted. In fact, the scope here has been even further limited by the fact that only planar and linear apertures are of interest. Moreover, in the m -dimensional co-ordinate system defined at the beginning of this chapter, the aperture co-ordinate in dimension m is always zero. Hence, in this dimension, Equation (2.27) can be rewritten

$$d(0) = \delta_x(0) [h(x) * f(x)] = \int_{-\infty}^{+\infty} h(-\alpha) f(\alpha) d\alpha. \quad (2.28)$$

The intermediate expression in Equation (2.28) corresponds to evaluating $d(x)$ via a full convolution, and then sampling it with a delta function. For reasons soon to be apparent, this approach would never be used in practice, but it has been illustrated in Figure 2.5 to show the sequence of operations that would be required if

no better option were available. Noteworthy from Figure 2.5 are the requirements for zero padding the ROI by a factor of two in all dimensions, and the corresponding expansion of the region of support for the Green's function. These operations prevent wrap-around contamination from periodic extensions of the ROI [70].




Figure 2.5: Data flow for 2-D field calculation via 3-D convolution.

In addition, to ensure that the result of the convolution is correctly referenced to the origin, the Green's functions must be shifted to "wrap-around" format prior to convolution. For clarity, this operation has been shown implemented in the spatial domain, however, in practice, the same result can be achieved more efficiently by applying the Fourier domain translation theorem [57, 58].

Aside from the $O(N^m \log N)$ operations required by m -dimensional fast Fourier transforms (FFTs), zero padding places severe demands on memory, especially in three-dimensions. Fortunately, the right hand side of Equation (2.28) presents an immediate shortcut, as the only operations required in dimension m are seen to be the

multiplication of $f(x)$ by an *inverted* copy of $h(x)$, and integration of the result. The implication of this is that to calculate scattered fields over planar and linear apertures, convolutions are only required in the aperture dimensions. Furthermore, the integration in Equation (2.28) implies that the m -dimensional inverse Fourier transform required by the convolution theorem can be replaced by a single inverse Fourier transform in $m-1$ dimensions *after the integration*. Utilising FFTs, as illustrated in Figure 2.6, the corresponding calculation in the discrete domain requires only $O(N^{m-1} \log N)$ operations, which is a significant improvement.




Figure 2.6: Data flow for 2-D field calculation via 2-D convolution.

In Figure 2.6, the integration of the 2-D Fourier transforms in the third dimension is referred to as a projection operation. This is in keeping with the terminology used in straight ray tomography, where the *projection-slice* or *Fourier-slice* theorem forms the basis of many algorithms for reconstructing images of objects from parallel projection data [56, 72].

The projection-slice theorem states that the Fourier transform of a function's projection (made with straight rays) is equal to a slice of the function's Fourier transform taken parallel to the projection plane. The corresponding theorem for diffracting waves is known as the *diffraction-slice* or *Fourier-diffraction* theorem. This theorem states that the Fourier transform of a function's projection (made with cylindrically or spherically spreading waves) is equal to a slice of the function's Fourier transform taken on circular or spherical contours respectively [56, 73]. By equating the Fourier transform of a sum to a sum of Fourier transforms, it can be seen that Figure 2.6 represents a generalisation of these theorems, referred to as the *generalised projection theorem* [61, 74, 75]. The only difference between dealing with straight rays and diffracting rays is the requirement for pre-multiplying the Fourier transform of the object, taken in the dimensions parallel to the aperture, by a similarly transformed Green's function; the so-called *propagation transfer function*.

2.2.5 Propagation Transfer Functions

As noted in Section 2.1.4, Green's functions can be interpreted as point-spread functions in an imaging context. It is useful, therefore, to examine their Fourier domain properties to ascertain their impact on the information content of wavefield data. In Section 2.2.4, it was also noted that Green's functions only needed to be Fourier transformed in the aperture dimensions in order for scattered fields to be calculated over linear and planar apertures. In three-dimensions ($m = 3$), the *propagation transfer function* $G_3(k_x, k_y, z, \omega)$ is obtained by Fourier transforming $g_3(r, \omega)$ in the x and y dimensions. According to Equation (2.3) then,

$$G_3(k_x, k_y, z, \omega) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} g_3(r, \omega) \exp[-j(k_x x + k_y y)] dx dy. \quad (2.29)$$

This integral can be expedited with the Weyl identity, which decomposes Green's functions into an equivalent spectrum of plane waves. Again for $m = 3$, this identity takes the form [76]

$$g_3(r, \omega) = \frac{j}{2(2\pi)^2} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \frac{\exp[j(k'_x x + k'_y y + K_z |z|)]}{K_z} dk'_x dk'_y, \quad (2.30)$$

where

$$K_z = \sqrt{k^2 - (k'_x)^2 - (k'_y)^2} . \quad (2.31)$$

Substituting equation (2.30) into equation (2.29) then gives [62]

$$G_3(k_x, k_y, z, \omega) = \frac{j}{2(2\pi)^2} \times \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \exp[j(k'_x - k_x)x + j(k'_y - k_y)y] \frac{\exp[jK_z|z|]}{K_z} dk'_x dk'_y dx dy, \quad (2.32)$$

which can be simplified by substituting Fourier-domain delta functions for the complex exponentials in x and y , giving

$$G_3(k_x, k_y, z, \omega) = \frac{j}{2(2\pi)^2} \times \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \frac{\exp[jK_z|z|]}{K_z} (2\pi)^2 \delta(k'_x - k_x, k'_y - k_y) dk'_x dk'_y . \quad (2.33)$$

Invoking the sifting properties of these delta functions then leads to the expression

$$G_3(k_x, k_y, z, \omega) = \frac{j}{2} \frac{\exp[j|z|\sqrt{k^2 - k_x^2 - k_y^2}]}{\sqrt{k^2 - k_x^2 - k_y^2}} . \quad (2.34)$$

Figure 2.7 is a plot of $|G_3(k_x, k_y, z, \omega)|$ for $z = \lambda$, which among its salient features includes a ridge of singularity on the *radiation circle* $k^2 = k_x^2 + k_y^2$ (rendered as vertical delta functions). This singularity defines a spatial frequency band-pass characteristic that rises towards a sharp (though not infinitely sharp) cut-off. From Equation (2.34) it can be observed that spatial frequencies outside the radiation circle are attenuated exponentially, meaning that beyond a few wavelengths, the *evanescent waves* conveying these frequencies are virtually impossible to detect [22, 29]. Provided $r \gg \lambda$, therefore, it can be safely assumed that the spatial frequency response of the propagation transfer function is perfectly band-limited; that is,

$$G_3(k_x, k_y, z, \omega) = \left\{ \begin{array}{ll} \frac{j}{2} \frac{\exp[j|z|\sqrt{k^2 - k_x^2 - k_y^2}]}{\sqrt{k^2 - k_x^2 - k_y^2}} & k_x^2 + k_y^2 < k^2 \\ 0 & k_x^2 + k_y^2 \geq k^2 \end{array} \right\} \quad (2.35)$$




Figure 2.7: A plot of the function $\left| G_3\left(\frac{k_x}{k}, \frac{k_y}{k}, z, \omega\right) \right|$ for $z = \lambda$.

As the bounds of integration applied in the derivation of Equation (2.34) were infinite, Figure 2.7 is actually only indicative of the spatial spectrum received by an *infinite* 2-D aperture ($m = 3$). With finite apertures, the singular behaviour of the spatial spectrum would be less evident, as would be the central dip in the frequency response. Similar frequency-domain behaviour is observed for $m = 2$.

In terms of the appearance of reconstructed images, Figure 2.7 has several implications. Firstly, the band-limited response results in the resolution limitation of all conventional ultrasonic imaging systems. To extend this resolution effectively requires an extrapolation of spatial frequencies beyond the cutoff. Secondly, the severity of the band-pass characteristic is partly responsible for the ringing artefacts, or Gibbs oscillations often observed in ultrasonic images. These oscillations are exacerbated by the side-lobe artefacts mentioned in the introduction, and the combined effect is to increase the background clutter in reconstructed images. The Fourier domain characteristics of the propagation transfer function are therefore a major reason for the overall ill-posedness of inverse scattering problems, which will be analysed from a more mathematical perspective in Chapter 3.

2.3 Computational Aspects

To this point, the equations describing propagation and scattering have been derived and characterised in the continuous domain. The remainder of this chapter is concerned with translating these continuous domain equations into the discrete domain and validating the accuracy of the resulting matrix model against analytical solutions.

2.3.1 Data Structures and Organisation

Figure 2.8 illustrates the two main imaging geometries of interest. In each case, the functions of transmission and reception may be performed by the same transducer, or more generally, by separate transducers, as shown. The continuous m -dimensional image $f(\mathbf{r})$ associated with scatterers in the region of interest (ROI) is approximated by an m -dimensional array of discrete reflectivity samples arranged in a Cartesian grid. In ultrasonic image reconstruction, due to the sizes of the matrix multiplications otherwise necessary, the fast Fourier transform (FFT) is used to perform discrete convolutions. For this reason, only regular Cartesian grids with dimensions n_i restricted to powers of two are considered³. Each image, therefore, consists of a number of voxels or pixels given by

$$N = \prod_{i=1}^m n_i . \quad (2.36)$$

The leading $m - 1$ dimensions of the Cartesian grid define the aperture. This terminology has only computational significance, however, as although the corresponding memory locations are used to accumulate aperture data during computations, the *physical* aperture is generally displaced from the ROI by a range R , as discussed in Section 2.1.1. The total number of aperture elements is given by

$$A = \prod_{i=1}^{m-1} n_i . \quad (2.37)$$

Due to sparse sampling, not all of the aperture locations necessarily represent data, and so the number of data D is generally defined such that $D \leq A$.

³ Mixed radix FFT algorithms might have obviated this restriction, but were not implemented in this study.




Figure 2.8: Two imaging geometries and their corresponding computational grids. (a) 2-D imaging from a 1-D (linear) aperture; (b) 2-D computational grid; (c) 3-D imaging from a 2-D (planar) aperture; (d) 3-D computational grid.

Aside from multi-dimensional FFTs, most of the algorithms to be presented involve simply multiplication, addition and summation of vectors. These operations can be carried out independently of the mapping between physical co-ordinates and vector addresses, and are thus well suited to modern pipelined and superscalar computer architectures. When the physical mapping of the data structures is important, however, as when storing and retrieving data, it is convenient to use the column-oriented FORTRAN addressing scheme as opposed to the row-oriented ANSI C convention for addressing multi-dimensional arrays.

The reason for this concerns the way data are typically accessed during computations. If aperture data occupying the first $m-1$ dimensions are to be read from an m -dimensional array, it is straightforward to access these dimensions

without regard for the total number of dimensions. Additionally, aperture data can be accessed sequentially, without regard to the total vector length. According to the ANSI C convention, this would require either storing the dimensions in reverse order (supplying m as an additional function argument and then counting back from the last dimension), or addressing the data out of sequence.

2.3.2 Approximation of Green's Functions

In Section 2.1.4, expressions were derived for the free-space Green's functions $g_m(r, \omega)$. To facilitate the calculation of scattered fields numerically, these continuous domain functions must be approximated over a finite domain.

Veronesi and Maynard [77] discussed the options for approximating Green's functions by sampling them in either the spatial or Fourier-domains. They concluded that piecewise-constant approximations in the spatial-domain were adequate, especially beyond the near field, although the errors associated with sampling their analytical Fourier transforms tended to be larger, due to the singularities noted in Section 2.2.5. In this work, only spatially sampled Green's functions were used, and all Fourier transforms were computed numerically via FFT algorithms.

Figures 2.9 (a) and (b) illustrate Green's functions computed in 2-D and 3-D respectively. In view of the zero padding requirements of the ROI, the number of complex elements comprising each Green's function is given by

$$P = 2^{(m-1)} N, \quad (2.39)$$


where N was previously defined as the number of elements in the unpadded ROI. As only waves propagating *towards* the aperture are of interest, and because Equation (2.28) requires an inverted convolution kernel in dimension m , only the lower portion of each Green's function is required. The distance by which the origin is displaced from the aperture is simply the range offset R defined in Section 2.1.1.

In view of the discussion in Section 2.2.5, Green's functions can be regarded as sharply band-limited, and hence the maximum sample spacing is determined by the Whittaker-Shannon sampling theorem [29, 57] such that


$$\delta_i < \frac{\lambda_{\min}}{2} \quad i = (1, \dots, m-1), \quad (2.40)$$

where $\lambda_{\min}$ corresponds to the minimum wavelength (maximum frequency).

As convolutions, and hence Fourier transforms, are only required in the leading $m - 1$ dimensions, the Whittaker-Shannon criteria do not apply in dimension m (the range dimension). In order to avoid the ambiguity that might result from variable grid spacings, however, all of the images presented in subsequent chapters are reconstructed on grids with the same spacing in each dimension.



(a)



(b)

Figure 2.9: Real parts of scalar Green's functions computed over finite domains. (a) 2-D Green's function computed on a 64×32 grid; (b) 3-D Green's function computed on a $64 \times 64 \times 32$ grid. Grid spacing (δ_i) = 1.0 mm. Wavelength (λ) = 6.4 mm. $R = 10$ mm.

2.3.3 Incident Field Modelling

The previous two sections dealt with discrete approximations for images and Green's functions that allowed the numerical modelling of ultrasonic scattering. The remaining field quantity to be approximated is the incident pressure field $p_i(\mathbf{r}, \omega)$. In this section, various modes of insonification are considered, corresponding to those that have been used experimentally in subsequent chapters. At this stage, only time-harmonic fields are considered, from which broadband fields can be easily synthesised.

The simplest assumption to make in relation to the incident field is that it can be approximated by a plane wave propagating perpendicularly to the aperture, as illustrated in Figure 2.10 (a). Indeed, both preceding scattering derivations included plane-wave approximations for $p_i(\mathbf{r}, \omega)$, and would therefore be most accurate for this case. Given the Fourier transform conventions adopted in Section 2.1.2, the time-harmonic components of such a plane wave would have the mathematical form

$$p(\mathbf{r}, \omega) = F(\omega) \exp(jkx_m - j\omega t), \quad (2.41)$$

where $F(\omega)$ represents the temporal spectrum of the transmitted waveform. Equation (2.41) is particularly simple in the spatial Fourier domain, as it varies only in dimension m . In all other dimensions, its spatial Fourier transform is equal to a Dirac delta function, weighted by $F(\omega)$ and an appropriate complex phase factor depending on the depth into the medium x_m and time t . These characteristics will be seen to play an important role in expediting numerical calculations of forward and inverse scattering in later chapters. Approximating incident fields via plane waves propagating perpendicularly to the aperture is therefore desirable from a mathematical standpoint. Of course, perfect plane waves imply either infinite transmitting apertures or propagation over infinite distances, and so they can only be approximated in practice.

Approximate plane wave insonification is typically associated with *bistatic* data-acquisition geometries in which there is a single stationary transmitter in the far field and multiple receivers⁴. By dispensing with focussing in transmission and adopting an entirely computed focussing approach, considerable savings in system complexity can be achieved. Plane wave insonification and bistatic imaging geometries are of interest here by virtue of the array imaging experiments to be discussed in Chapter 6.

⁴ Or by reciprocity, multiple transmitters and a single stationary receiver.




Figure 2.10: Discrete approximations for time-harmonic incident fields. (a) Plane-wave; (b) Spherical wave; (c) Shaded spherical wave; (d) Field from square aperture. Grid spacing (δ_i) = 1.0 mm. Wavelength (λ) = 6.4 mm. Aperture offset (R) = 10 mm.

The next most common assumption regarding incident fields is that they can be associated with geometric line and point sources in 2-D and 3-D respectively. One example of a spherical field from a 3-D point source is illustrated in Figure 2.10 (b). In *monostatic* data acquisition, a small transducer, or pair of such transducers, is scanned over the aperture plane, providing wide-angle coverage [78-80]. Monostatic data-acquisition is of interest here due to its simplicity, and suitability for both mechanically scanned and array based imaging systems. It also forms the basis of simulations in Chapter 3 and synthetic aperture imaging experiments in Chapter 4.

As noted in Section 2.1.4, the fields resulting from point sources are simply free-space Green's functions centred on the source⁵. This explains the similarity between Figures 2.9 (b) and 2.10 (b). In particular, if the transmitter and receiver are coincident in their aperture dimensions ($\xi = \xi'$), the transmitted field becomes shift-invariant with respect to the receiver, giving

$$p_i(\mathbf{r}, \omega) \cong g_m(|\mathbf{r} - \xi|, \omega) = g_m(|\xi' - \mathbf{r}|, \omega). \quad (2.42)$$

Two-way propagation between the transmitter and receiver may then be accounted for by simply squaring the Green's function, simplifying Equation (2.24) to

$$p_s(\xi', \omega) \cong K(\omega) F^2(\omega) \int_{V_0} g_m^2(|\xi' - \mathbf{r}|, \omega) f(\mathbf{r}) d^m \mathbf{r}, \quad (2.43)$$

where $F^2(\omega)$ is the two-way temporal frequency response of the transducer. It should be noted that as a result of squaring the Green's function, the critical sampling rate for monostatic data acquisition is effectively doubled from that given by Equation (2.40); that is, the maximum sample spacing allowable δ_i becomes $\lambda_{\min}/4$.

In practice, the beam patterns associated with small transducers (and the focussed transducers used to emulate them) can never be perfectly isotropic, and allowances for this must be made. In the monostatic synthetic aperture imaging experiments reported in Chapter 5, cylindrically and spherically focussed transducers were employed to produce wide-angle beams in 2-D and 3-D respectively. In each case, it was found that the beam patterns were well approximated by Gaussian functions of the form

$$b(\theta) = \exp\left[\frac{4 \ln(0.5) \theta^2}{w_{6dB}^2}\right], \quad (2.44)$$

where w_{6dB} corresponds to the 6 dB angular beam width. Figure 2.10 (c) illustrates one such Gaussian shaded field computed for $w_{6dB} = 30^\circ$. Incorporating two-way beam patterns of this type into Equation (2.44) yields

$$p_s(\xi', \omega) \cong K(\omega) F^2(\omega) \int_{V_0} [b^2(\theta) g_m^2(|\xi' - \mathbf{r}|, \omega)] f(\mathbf{r}) d^m \mathbf{r}. \quad (2.45)$$

Strictly speaking, Equation (2.44) only applies at a single frequency, as beam patterns are generally frequency dependent [81]. For wide bandwidths, and/or

⁵ This ignores the boundary conditions in the aperture plane, which is acceptable provided $r \gg \lambda$ [22, 78].

arbitrarily shaped transmitting apertures, beam patterns must be calculated numerically at each frequency of interest. At several wavelengths from the transmitting aperture, the Huygens-Fresnel form of the Rayleigh-Sommerfeld formula for a baffled piston source can be approximated by [20, 22, 29]

$$p_i(\mathbf{r}, \omega) = -j\omega\rho \int_{\xi} v_z(\xi, \omega) g_m(|\mathbf{r} - \xi|, \omega) da, \quad (2.46)$$

where $v_z(\xi, \omega)$ is the temporal spectrum of the normal velocity of the transducer, ρ is the density of the medium, ξ is the surface of the transmitting aperture, and $g_m(|\mathbf{r} - \xi|, \omega)$ is the same free-space Green's function discussed earlier. Figure 2.10 (d) illustrates the time-harmonic field calculated from Equation (2.46) for a square transmitting aperture. This field is seen to be reasonably well approximated by the plane wave discussed earlier: a fact that will be exploited in Chapter 6.

2.3.4 Matrix Formulation

With discrete approximations for the major field quantities now in place, a matrix model can be formulated to describe the forward scattering and data acquisition processes. To allow the use of compact matrix notation, the elements of the computational grid are lexicographically ordered into column vectors. The unknown (and in most cases unknowable) *true* image of the scatterers is defined to be the vector $\mathbf{f}$, and all estimates of this vector are denoted $\hat{\mathbf{f}}$. When required, it will be implicitly assumed that all such vectors represent appropriately zero padded data structures so that the results of convolutions are linear rather than circular.

As noted in Section 2.2.3, incident fields have a modulating influence on secondary sources - influencing them only in amplitude and phase. Computationally, this modulation represents point-wise multiplication between two data structures; the first being the image of the scatterers as "seen" by the aperture, and the second being the incident field computed throughout the ROI. The corresponding diagonal matrix operator acting on $\mathbf{f}$ contains samples of the incident field, and is denoted by $\mathbf{E}$.

In regard to the lexicographical ordering of the Green's functions, the aim is to establish both Fourier and matrix interpretations for the partial convolution discussed in Section 2.2.4. The Fourier domain interpretation is straightforward, with zero-padded images and Green's functions Fourier transformed in their leading $m-1$

dimensions and multiplied together. In a matrix context, convolutions are usually associated with block-circulant matrices operating on lexicographically ordered vectors. In this instance, due to the modified form of the convolutions, the block-circulant matrices are replaced by block-diagonal matrices, with each diagonal block possessing either a circulant or block-circulant structure. The Green's function operators, therefore, are diagonalised by discrete Fourier transforms (DFTs) in $m - 1$ dimensions. To simplify subsequent notation, the spectral functions introduced in Section 2.2.3 are absorbed into the discrete representations of the Green's functions, and the resulting semi-block-circulant matrices are denoted by $\mathbf{H}$.

Following modulation by the incident field operator $\mathbf{E}$, and partial convolution with the Green's function operator $\mathbf{H}$, it remains to project the resulting m -dimensional lexicographically ordered vector into $m - 1$ dimensions. Whereas the equivalent matrix operation is difficult to express generally, the action of the projection operator, denoted $\mathbf{P}$, in one dimension ($m = 1$) can be expressed as

$$\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_N \end{bmatrix} = \sum_i^N a_i. \quad (2.47)$$

In m dimensions, $\mathbf{P}$ takes the form of an $A \times N$ matrix operator consisting of n_m identity matrices arranged row-wise, where n_m was previously defined to be the number of discrete depth layers in the computational grid.

Having defined the matrix operators that together describe the mapping between image and data space, two additional components are required in order to represent realistic experimental conditions. The first of these is the vector $\mathbf{n}$ of complex additive Gaussian noise terms with uniform variance $2\sigma^2$. The second is the sampling operator $\mathbf{S}$, which accounts for the fact that in many practical circumstances, it is not possible to sample the aperture completely; that is,

$$\mathbf{S} = \text{diag}[s_i] \quad (i = 1, \dots, A), \quad (2.48)$$

where $s_i = 1 + j0$ for each of the D values that correspond to actual measurements in the aperture, and $s_i = 0 + j0$ otherwise. Combining all of the above matrix operators and vectors results in the single matrix equation

$$\mathbf{d} = \mathbf{S}(\mathbf{P}\mathbf{H}\mathbf{e}\mathbf{f} + \mathbf{n}), \quad (2.49)$$

which constitutes the discrete *forward* scattering model used throughout this thesis. In particular, it must be emphasised that the diagonal operators $\mathbf{E}$ and $\mathbf{S}$, and the Green's function and projection operators $\mathbf{H}$ and $\mathbf{P}$ are only *formally* matrices, as in reality they are all prohibitively large to store. In practice, diagonal operators are stored and manipulated as vectors and only the diagonal forms (the DFTs of the first column) of each block-circulant matrix are involved in calculations.

2.4 Model Validation

To jointly validate the computer programs written to implement the incident field calculations and the numerical scattering model, a canonical scattering arrangement was selected for which accurate analytic solutions were available. In NDE, the flat-bottom hole represents such a standard test case, and for this reason, analytic expressions have been derived to enable the calculation of scattered fields for comparison and standardisation of defect detection sensitivity.

Figure 2.11 illustrates the classic configuration of coaxial circular transducer and flat-bottom hole, separated by a distance z_i . The sequence of events commences with some time-domain velocity excitation applied to the circular piston transducer and ultrasonic waves propagating towards the end of the hole from which they are then scattered back towards the transducer. The received pressure wave, averaged over the entire circular aperture of the transducer, is denoted $\langle p \rangle$.




Figure 2.11: Hypothetical configuration of circular transducer and flat-bottom hole.

For this arrangement, Schmerr and Sedov [82] derived expressions for the spatially averaged time-harmonic pressure received by the transducer, in terms of its initial time-harmonic normal velocity excitation v_0 . For a given wavenumber k , the average pressure $\langle p \rangle$ received by a transducer of diameter a , located as illustrated with respect to a flat-bottom hole of diameter b is given by

$$\begin{aligned} \langle p \rangle = & -(\rho v_0 c b^2 / a^2) \{ \exp(jkz_i) [\exp(jkz_i) - 2 \cos \alpha J_1(kb \sin \alpha) \\ & \cdot \exp(jk\sqrt{(a^2 + z_i^2)}) / (kb \sin \alpha)] - \exp(jk\sqrt{(a^2 + z_i^2)}) \\ & \cdot [2 J_1(kb \sin \alpha) \exp(jkz_i) / (kb \sin \alpha) - \cos \alpha (J_0^2(kb \sin \alpha) \\ & + J_1^2(kb \sin \alpha)) \exp(jk\sqrt{(a^2 + z_i^2)})] \} \end{aligned} \quad (2.50)$$

where

$$\alpha = \sin^{-1} \left[a / \sqrt{(a^2 + z_i^2)} \right] \quad (2.51)$$

and the J_0 and J_1 are Bessel functions of the first kind, of order 0 and 1 respectively.

To compare the analytical solution above with the numerical model derived earlier in this chapter, time-dependent fields were computed by summing time-harmonic fields, given by Equations (2.49) and (2.50), over the bandwidth of the pulse shown in Figure 2.12. This Gaussian pulse possessed a frequency spectrum centred on 10 MHz, and a 6dB bandwidth of 7 MHz. The propagation velocity in the medium was assumed to be 5960 ms^{-1} , corresponding to mild steel [36]. For a given time duration $\tau = z_i / c$ of the computed response, the corresponding maximum frequency increment used for the evaluation of $\langle p(t) \rangle$ was $1/\tau$.




Figure 2.12: (a) Gaussian weighted velocity pulse; (b) Normalised spectrum.

Figure 2.13 illustrates the discrete models used for the transducer and flat-bottom hole. The diameter of the transducer was set to 12.8 mm and the diameter of the flat-bottom hole was set to 3.6 mm. Both the transducer and flat-bottom hole were modelled on a regular 64×64 grid with a spacing of 0.2 mm. Each model was generated via the GNU image manipulation package, which provided automatic anti-aliasing to improve the approximation to a circular geometry in each case [83].





Figure 2.13: Discrete approximations for the circular transducer and flat-bottom hole. (a) 12.8 mm ϕ transducer; (b) 3.6 mm flat-bottom hole (0.2 mm per division).


To facilitate the calculation of time-domain responses via the matrix model, the discrete model for the flat-bottom hole in Figure 2.13 (b) was input to the scattering evaluation program as a PGM format 8-bit graphics file. The particular program used allowed the option of selecting either Equation (2.25) or (2.26) as the spectral function in the scattering model. As the flat-bottom hole was clearly impenetrable, the latter, corresponding to Kirchhoff theory was selected. In lexicographical form, the image in Figure 2.13 (b) assumed the role of the image vector $\mathbf{f}$.

The incident fields from the circular transducer were calculated according to Equation (2.46) via a separate computer program, which used Figure 2.13 (a) as the model for the transducer aperture ξ . In this instance, as the problem involved two flat geometries at a known distance, it was permissible to formulate the matrix model with $m = 2$. The aperture offset R then provided the means of setting the transducer - hole spacing z_i . Finally, once the scattered fields were computed, the data were spatially sampled and averaged by the $\mathbf{S}$ operator, which had previously been initialised with the same PGM image file used to model the transmitting aperture.


Figure 2.14 is a collection of time-domain responses calculated from both the analytic and numerical models. To highlight the changes in the received pressure with the distance z_i , responses were calculated for $z_i = 5, 10, 20$ and 50 mm.




(a)



(b)



(c)



(d)

Figure 2.14: Comparisons between analytical flat-bottom hole responses (left), and numerical responses (right). $z_i =$ (a) 5 mm; (b) 10 mm; (c) 20 mm; (d) 50 mm.

These responses can be seen to consist of high amplitude reflected signals, followed by weaker signals arising from edge diffraction [82]. The specular signals appear phase inverted with respect to the excitation pulse in the near field, due to the Dirichlet boundary conditions of the hole. Their phases shift to 90° into the far field, with the weaker diffracted signals merging with the main reflected signal.

The agreement between these two sets of responses is quite close. In particular, the agreement in amplitude and phase of the edge diffracted signals gives confidence that this aspect of scattering has been well modelled, despite not being explicitly taken into account by the underlying Kirchhoff theory.

In some respects, the complexity of the calculation required to produce the numerical results in Figure 2.14 (30 minutes on a 200 MHz R10000 CPU) is unjustifiable in the light of Equation (2.50). The primary aim of this exercise, however, which was to validate the theory and computer programs to be used in subsequent inverse scattering formulations, has been achieved.

2.5 Summary

In this chapter, approximate formulae were derived for the two important cases of penetrable and impenetrable ultrasonic scattering, and from these, a generalised scattering formula was obtained. The mapping between image space and data space was then analysed in terms of both the mathematical operations required to implement it and the spatial frequency behaviour of the propagation transfer function. The resolution limitations of conventional ultrasonic imaging systems were seen to result from severe restrictions on spatial frequency bandwidth. Correspondingly, the ill-posedness of inverse scattering problems is seen to be associated with the need to extrapolate beyond a finite set of spatial frequencies.

Each of the terms from the generalised scattering formula was then translated into the discrete domain prior to establishing a matrix observation model for scattered fields measured over linear and planar apertures. This matrix model was subsequently tested for a simple case to verify the performance of both the theory and the associated computer programs. While the matrix model itself does not enhance the interpretation of the physics, it does facilitate more rigorous investigation of the inverse problem, which is the subject of Chapter 3.

Chapter 3

A Survey of Image

Reconstruction Concepts

Early this century, Hadamard [84] first observed that the solutions to certain integral equations were nonunique and unstable; nonunique in the sense that different solutions satisfied the same data, and unstable in the sense that even small perturbations in the data resulted in unbounded fluctuations. Since that time, such symptoms of *ill-posedness* have been identified in many practical inverse problems. Rather than implying any defect in the reasoning by which these problems have been formulated, ill-posedness refers to the difficulty of recovering information that has been suppressed by mathematical (usually integral) transformations.

This chapter constitutes a survey of image reconstruction concepts applicable to solving ill-posed inverse scattering problems formulated according to the linearised model derived in Chapter 2. In addition to the MaxEnt approach, which is the main thrust of this thesis, several other regularised and unregularised approaches are examined, including inverse filtering, matched filtering, minimum-norm least-squares and constrained least-squares. Due to its competition with MaxEnt in the field of radio-astronomy, the non-linear CLEAN algorithm is also evaluated.

The purpose of evaluating methods other than MaxEnt is firstly to identify the problems associated with simple solutions, particularly in relation to sparse data, and secondly to provide baseline solutions against which the results of MaxEnt can be compared. As MaxEnt is by far the most complicated approach discussed, it is important that its benefits are demonstrated unambiguously and independently of other factors, such as the physical accuracy of the scattering model. To achieve such

a demonstration, numerical experiments are conducted in which both point-like and locally smooth scatterers are reconstructed from simulated broadband data. The algorithmic details of how the new approach to MaxEnt is implemented, as well as the results of real experiments will be left to subsequent chapters.

3.1 Unregularised Image Reconstruction

3.1.1 Inverse Filtering

To solve inverse scattering problems, it is required to infer images of scatterers from noisy and sparsely sampled field data. The first requirement of prospective solutions is that they satisfy the forward scattering model to within the expected error, and so some measure of this agreement must be formulated. The following analysis commences with the time-harmonic observation model for forward scattering derived in Chapter 2, given by

$$\mathbf{d} = \mathbf{S}(\mathbf{P}\mathbf{H}\mathbf{E}\mathbf{f} + \mathbf{n}), \quad (3.1)$$

where

- $\mathbf{d}$ is a data vector consisting of a maximum of A aperture samples.
- $\mathbf{S}$ is a binary-valued diagonal operator sampling D field data from $\mathbf{d}$.
- $\mathbf{P}$ is a projection operator mapping m dimensions to $m - 1$ dimensions.
- $\mathbf{H}$ is a propagation operator mapping secondary sources to field data.
- $\mathbf{E}$ is an incident field operator mapping scatterers to secondary sources.
- $\mathbf{n}$ is a vector of A complex Gaussian deviates with variance $2\sigma^2$.

In view of the assumption that the noise is Gaussian, it is convenient to compute the data error as a single Chi-squared statistic. At a single temporal frequency, this statistic can be expressed in terms of the above notation as

$$\chi^2 = \frac{1}{2\sigma^2} \left\| (\mathbf{S}\mathbf{P}\mathbf{H}\mathbf{E}\hat{\mathbf{f}} - \mathbf{d}) \right\|^2, \quad (3.2)$$

where the “true image” $\mathbf{f}$ has been replaced by the image estimate $\hat{\mathbf{f}}$. From this expression, it might appear that there could be no further input to the image reconstruction problem other than the data $\mathbf{d}$, which would imply that the optimal image could be estimated in all cases by simply minimising χ^2 .

Unfortunately, unconstrained error minimisation is problematic for two reasons. Firstly is the problem of non-uniqueness, which even in the hypothetical case of noiseless data, results in an infinite number of solutions equally able to minimise χ^2 . Rather than a single optimal solution, therefore, only a *feasible set* of solutions can be identified. Secondly, the image selected by straightforward error minimisation is likely to be dominated by noise, which is amplified in the process of reconstructing weakly transmitted image features.

Both of these problems can be viewed in a mathematical context by differentiating Equation (3.2) with respect to $\hat{\mathbf{f}}$ and setting the resulting expression to zero in an attempt to find a stationary point. Solving the resulting equation for $\hat{\mathbf{f}}$ would then be expected to yield a least-squares (minimum variance) image estimate. Proceeding along these lines via elementary matrix calculus gives¹

$$[\nabla \chi^2]_i = \frac{\partial(\chi^2)}{\partial f_i} = \frac{1}{\sigma^2} \mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} (\mathbf{S} \mathbf{P} \mathbf{H} \hat{\mathbf{f}} - \mathbf{d}) \quad (i = 1, 2, \dots, N), \quad (3.3)$$

where N is the number of complex image elements in the region of interest (ROI), and $\mathbf{E}^*$, $\mathbf{H}^{*T}$ and $\mathbf{P}^T$ are the conjugates, conjugate transposes and transposes of the respective matrix operators. Equating $\nabla \chi^2$ to zero and invoking the simplification $\mathbf{S} \mathbf{S} = \mathbf{S}$ for the binary sampling operators then gives

$$\mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H} \hat{\mathbf{f}} = \mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d}. \quad (3.4)$$

To solve this equation for the least-squares image estimate $\hat{\mathbf{f}}$, the $N \times N$ matrix $\mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}$ must be inverted. As this matrix takes the general form $\mathbf{A}^{*T} \mathbf{A}$, it is symmetric with real and nonnegative eigenvalues. This follows from the singular value decomposition (SVD) of any $A \times N$ matrix $\mathbf{A}$ according to

$$\mathbf{A} = \mathbf{U} \text{diag}[w_1, w_2, \dots, w_N] \mathbf{V}^{*T}, \quad (3.5)$$

where $\mathbf{U}$ is an $A \times N$ column-orthogonal matrix, $\mathbf{V}$ is an $N \times N$ orthogonal matrix, and the w_i are known as singular values, of which at most A will be nonzero [70]. By virtue of the column-orthogonality of $\mathbf{U}$,

$$\mathbf{A}^{*T} \mathbf{A} = \mathbf{V} \text{diag}[w_1^2, w_2^2, \dots, w_N^2] \mathbf{V}^{*T}, \quad (3.6)$$

¹ This usage of the ∇ symbol supersedes its previous use in Chapter 2 as a continuous-domain gradient operator.

from which it can be seen that the eigenvalues $\lambda_i = w_i^2$ ($i = 1, 2, \dots, N$) are real and non-negative. The *rank* of $\mathbf{A}$, therefore, which is the number of non-zero singular values, is also the rank of $\mathbf{A}^{*T}\mathbf{A}$, and the corresponding columns of $\mathbf{V}$ are the set of orthonormal basis vectors that constitute its *range*. There is also a complementary set of vectors that constitutes the *null-space* of $\mathbf{A}^{*T}\mathbf{A}$. These vectors correspond to zero singular values, or those less than some threshold determined by either the signal-to-noise ratio (SNR) of the data $\mathbf{d}$ or the available numerical precision.

When $\mathbf{A}^{*T}\mathbf{A}$ is block-circulant, its orthonormal eigenvectors are simply the basis vectors of the discrete Fourier transform (DFT) in m dimensions [57, 58, 85]. Provided the matrix $\mathbf{E}^*\mathbf{H}^{*T}\mathbf{P}^T\mathbf{S}\mathbf{P}\mathbf{H}\mathbf{E}$ is block-circulant, therefore, the orthogonal decomposition

$$\mathbf{E}^*\mathbf{H}^{*T}\mathbf{P}^T\mathbf{S}\mathbf{P}\mathbf{H}\mathbf{E} = \mathbf{W} \text{diag}[\lambda_1, \lambda_2, \dots, \lambda_N] \mathbf{W}^{*T} \quad (3.7)$$

can be carried out via DFTs, where the matrices $\mathbf{W}$ and $\mathbf{W}^{*T}$ are forward and inverse DFT matrices respectively, and the eigenvalues λ_i are DFT coefficients.

As noted in Section 2.3.4, the character of the matrix observation model derived in Chapter 2 is semi-block-circulant; that is, block-circulant in $m-1$ dimensions rather than m . Because of its semi-block-circulant structure, the Green's function matrix $\mathbf{H}$ can be diagonalised by DFTs in $m-1$ dimensions according to its m -dimensional lexicographical ordering. In order for Equation (3.7) to have a simple interpretation in the Fourier domain, the operators $\mathbf{E}$, $\mathbf{P}$ and $\mathbf{S}$ must also exhibit this semi-block-circulant structure.

In addition to being a diagonal matrix, the incident field operator $\mathbf{E}$ is semi-block-circulant under two restrictive conditions. Firstly, as noted in Section 2.3.3, if the incident field is a plane wave propagating perpendicular to the aperture, then its complex amplitude at any given depth will be uniform, as illustrated in Figure 2.10 (a). In that case, each diagonal block of $\mathbf{E}$ is a complex constant times an $A \times A$ identity matrix, where A was defined in Section 2.3.1 to be the number of discrete elements (both sampled and unsampled) in the aperture plane. Under these conditions, $\mathbf{E}$ is semi-block-circulant with rank N . Secondly, if data acquisition is monostatic, as discussed in Section 2.3.3, then the incident field is shift-invariant with respect to the receiver, and $\mathbf{E}$ is again semi-block-circulant.

In Section 2.3.4, the projection operator $\mathbf{P}$ was illustrated in one dimension ($m=1$) as a unit row vector operating on a column vector, with the result of its application being a single datum in the aperture plane. In the same way, the operation of the transposed projection operator $\mathbf{P}^T$ can be expressed as

$$\begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^T [a] = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} [a] = \begin{bmatrix} a \\ a \\ \vdots \\ a \end{bmatrix}, \quad (3.8)$$

where the unit row vector has become a unit column vector operating on a single datum a so as to replicate it at every layer of the computational grid. In the terminology of tomography, the datum a is said to have been *backprojected*. Beyond $m=1$, the matrix operator $\mathbf{P}^T$ is a juxtaposition of n_m identity matrices arranged column-wise, where n_m is the number of depth layers in the computational grid. In all cases, both $\mathbf{P}$ and $\mathbf{P}^T$ are semi-block-circulant matrices with rank A .

In Figure 3.1, the operations of projection and backprojection have been illustrated in three dimensions ($m=3$). As $\mathbf{P}$ and $\mathbf{P}^T$ invariably follow or precede Fourier transformations (through convolutions with $\mathbf{H}$ and $\mathbf{H}^{*T}$ respectively), Fourier transforms rather than spatial functions have been shown projected *to* and backprojected *from* the aperture plane. This corresponds to the computational situation where it is always more efficient to Fourier transform a sum of functions than it is to sum individually computed Fourier transforms.




Figure 3.1: The relationship between projection and its transpose (backprojection).

In order to express the solution to Equation (3.4) in a simple closed form, the spatial-domain sampling operator $\mathbf{S}$ is also required to be semi-block-circulant. The necessary condition for this to be true is that the aperture is completely sampled, whereby $\mathbf{S}$ then takes the form of an $A \times A$ identity matrix.

Given the above conditions relating to the mode of insonification and the sampling of the aperture, it is a simple matter to calculate the diagonal form of $\mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H} \mathbf{E}$ via successive multiplications, projections and fast convolutions. Using the orthogonal properties of the Fourier transform matrix $\mathbf{W}$ in Equation (3.7), the solution to Equation (3.4) can be *expressed*, but not necessarily *computed* as

$$\hat{\mathbf{f}} = \mathbf{W}^{*T} \text{diag} \left[\frac{1}{\lambda_i} \right] \mathbf{W} \mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d} \quad (i = 1, 2, \dots, N). \quad (3.9)$$

Extension of this solution to broadband problems follows from Equation (3.2), whereby considering K discrete frequencies, the Chi-squared error is modified to

$$\chi^2 = \frac{1}{2\sigma^2} \sum_{k=1}^K \left\| \left(\mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} - \mathbf{d}_k \right) \right\|^2. \quad (3.10)$$

In addition to a separate data vector $\mathbf{d}_k$, each frequency requires a separate Green's function operator $\mathbf{H}_k$ and incident field operator $\mathbf{E}_k$. A different noise variance σ_k^2 could also be used for each frequency, however in this work, the power spectrum of the noise is assumed to be flat within the range of frequencies considered. By redefining the orthogonal decomposition in Equation (3.7) to be

$$\sum_{k=1}^K \left(\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \right) = \mathbf{W} \text{diag} \left[\lambda'_1, \lambda'_2, \dots, \lambda'_N \right] \mathbf{W}^{*T}, \quad (3.11)$$

the broadband least squares image estimate becomes

$$\hat{\mathbf{f}} = \mathbf{W}^{*T} \text{diag} \left[\frac{1}{\lambda'_i} \right] \mathbf{W} \sum_{k=1}^K \left(\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d}_k \right) \quad (i = 1, 2, \dots, N). \quad (3.12)$$

Both Equations (3.9) and (3.12) have simple interpretations in the Fourier domain whereby the inverse eigenvalues serve as inverse filter coefficients equalising the spatial frequency response of the reconstructed images.

Complete equalisation is impossible, however, due to the occurrence of zero eigenvalues. These are inevitable, firstly because the propagation transfer function to which the eigenvalue spectrum of $\mathbf{H}$ is a discrete approximation is band-limited, and

so the λ_i will be non-zero only up to some index $r < N$ where r is the rank of $\mathbf{H}$. Zero eigenvalues correspond to spatial frequencies outside the pass-band of the propagation transfer function (the range of $\mathbf{H}$) and as such are not transmitted in the data. Similarly, the projection operator $\mathbf{P}$ involves a loss of information, as it represents a mapping from m dimensions to $m - 1$ dimensions. Furthermore, it can be seen that any linear combination of eigenvectors from the null-space of $\mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H} \mathbf{E}$ can be added to $\hat{\mathbf{f}}$ without affecting the equality in Equation (3.4). This indicates that the data are insufficient to determine a unique solution².

In regard to the stability of solutions obtained via either Equations (3.9) or (3.12), substituting Equation (3.1) for the data vector $\mathbf{d}$ into the latter gives

$$\hat{\mathbf{f}} = \mathbf{W}^{*T} \text{diag} \left[\frac{1}{\lambda_i'} \right] \mathbf{W} \sum_{k=1}^K \left[\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} (\mathbf{P} \mathbf{H}_k^{*T} \mathbf{E}_k^* \mathbf{f} + \mathbf{n}) \right], \quad (3.13)$$

from which it can be seen that the noise will be amplified along with any weakly transmitted eigenvectors. This leads to problems when inverse filtering is used to obtain solutions directly, as the presence of even moderate noise results in oscillations of such large amplitude that image detail can be totally obscured.

From the above discussion, it is evident that nonuniqueness and instability are inherent aspects of inverse filtered (unconstrained least-squares) image reconstruction, even in the event that the aperture data are completely sampled. When sparsely sampled apertures are considered, image reconstruction becomes even more ill-posed, and the composite matrix $\mathbf{E}^* \mathbf{H}^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H} \mathbf{E}$ can no longer be diagonalised by Fourier transformation, meaning that iterative approaches to minimising χ^2 must be employed, as discussed later in this chapter.

3.1.2 Matched Filtering

Equation (3.13) shows that it is possible in principle to estimate broadband images of scatterers in a least-squares sense by inverse filtering, although the uniqueness and stability of the resulting images are not assured. A simple means of avoiding instability is to instead compute the broadband *matched filtered* solution

² The issue of nonuniqueness in inverse scattering has been debated extensively. See for example [86-89].

$$\hat{\mathbf{f}} = \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d}_k, \quad (3.14)$$

which is intrinsically stable, as it involves only transpositions, backprojections and multiplications, as opposed to matrix inverses. Whereas inverse filtering achieves optimal resolution from linear signal processing under conditions of high SNR, matched filtering can be shown to achieve optimal peak SNR in the presence of both random and non-random interference [22, 90].

Equation (3.14) is the basis of many practical approaches to inverse scattering, unifying time-domain synthetic aperture focussing (SAFT) techniques and Fourier domain techniques such as Rayleigh Sommerfeld holography as mathematically equivalent implementations of classical beamforming [60, 63, 91-95]. Whereas broadband signals can be shifted in the time-domain to establish the required focussing delays, the equivalent operation in the Fourier domain is convolution with the complex conjugate of the appropriate free-space Green's function [21, 93]. Taking the ratio of the propagation transfer function given by Equation (2.34) at zero range relative to that at range z , the corresponding *backpropagation filter* commonly employed in bistatic acoustic holography is given for the case $m = 3$ by [29, 92, 93]

$$\begin{aligned} \frac{p(k_x, k_y, 0, \omega)}{p(k_x, k_y, z, \omega)} &= \frac{G_3(k_x, k_y, 0, \omega)}{G_3(k_x, k_y, z, \omega)} \\ &= \begin{cases} \exp[-j|z|\sqrt{k^2 - k_x^2 - k_y^2}] & k_x^2 + k_y^2 < k^2 \\ 0 & k_x^2 + k_y^2 \geq k^2 \end{cases}, \end{aligned} \quad (3.15)$$

where as before, the evanescent waves have been discarded. Similarly, for the monostatic case, the backpropagation filter is found to be [68, 96]

$$\frac{p(k_x, k_y, 0, \omega)}{p(k_x, k_y, z, \omega)} = \begin{cases} \exp\left[-j2|z|\sqrt{k^2 - \frac{k_x^2}{4} - \frac{k_y^2}{4}}\right] & k_x^2 + k_y^2 < 4k^2 \\ 0 & k_x^2 + k_y^2 \geq 4k^2 \end{cases} \quad (3.16)$$

Strictly speaking, matched filtering does not yield quantitative images of *scatterers* in terms of their frequency-independent scalar reflectivity. The images obtained, however, do constitute reconstructions of the *fields* $p(\mathbf{r}, \omega)$, which converge to the secondary sources on the surfaces of the scatterers [63, 87, 93].

As mathematically equivalent implementations of classical beamforming, Fourier-domain acoustic holography and time-domain synthetic aperture focussing suffer the same limitations in relation to resolution and side-lobes as discussed in Chapter 1. Moreover, when apertures are sparsely sampled, neither direct inverse filtering nor matched filtering is optimal in any sense, and optimal image reconstruction then requires *regularised* iterative algorithms as discussed next.

3.2 Regularised Image Reconstruction

In the previous section, the prescription for calculating least-squares image estimates was seen to be incomplete due to the ill-posedness of the underlying inverse problem. Despite being more stable than inverse filtering, matched filtering faces similar resolution limitations and image degradation due to sparse aperture sampling, and for this reason, alternative approaches need to be considered.

In response to the need to solve practical inverse problems numerically, Tikhonov and Arsenin [11] introduced the concept of *regularisation*, which essentially involves transforming ill-posed problems into closely related well-posed problems. Mathematically, this can be done by augmenting the least-squares inversion constraint with additional constraints reflecting desirable properties of mathematical regularity, such as uniqueness and stability (usually interpreted as smoothness). The function to be minimised then takes the variational form

$$J(\hat{\mathbf{f}}) = \chi^2 + \alpha Q(\hat{\mathbf{f}}). \quad (3.17)$$

where α is a Lagrange multiplier intended to mediate between the influence of the “goodness of fit” constraint (chi-squared) and the *regularising functional* Q .

By minimising $J(\hat{\mathbf{f}})$, both the uniqueness and stability of $\hat{\mathbf{f}}$ are guaranteed provided Q and α are appropriately chosen. Many choices for the functional Q have been suggested in the image processing literature [9, 11, 97-106], ranging from those simply intended to ensure the uniqueness of solutions to those capable of introducing non-trivial *prior information*. The first class of regularising functionals to be considered is the class of linear discrete functionals of the form

$$Q(\hat{\mathbf{f}}) = \|\mathbf{R}\hat{\mathbf{f}}\|^2, \quad (3.18)$$

where $\mathbf{R}$ is formally an $N^2 \times N^2$ matrix known as a *regularising operator* [85, 107, 108]. To ensure the tractability of the minimisations involved, the assumption of shift-invariance previously made in relation to the observation model will also be made for the regularising operator. $\mathbf{R}$ is therefore required to be semi-block-circulant such that it can be diagonalised by discrete Fourier transforms (DFTs).

3.2.1 Minimum-Norm Least-Squares

From the perspective of uniqueness, an effective choice of regularising operator $\mathbf{R}$ is the identity matrix $\mathbf{I}$. The joint function to be minimised then takes the form

$$J(\hat{\mathbf{f}}) = \chi^2 + \alpha \|\hat{\mathbf{f}}\|^2, \quad (3.19)$$

which leads to the minimum-norm least-squares (MNLS) or minimum energy solution. As with the inverse and matched filtered image estimates, the MNLS image estimate can be expressed analytically provided the operators $\mathbf{S}$ and $\mathbf{E}_k$ are semi-block-circulant. Following a similar derivation to that given for the unregularised least-squares solution; calculating $\nabla J(\hat{\mathbf{f}})$ and equating it to zero to find a stationary point, the broadband MNLS image is found to be

$$\hat{\mathbf{f}} = \mathbf{W}^{*T} \text{diag} \left[\frac{1}{\lambda'_i + \alpha 2\sigma^2} \right] \mathbf{W} \sum_{k=1}^K (\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d}_k), \quad (3.20)$$

where the eigenvalues λ'_i are those in Equation (3.11). From this equation, it can be seen that the equalising effect of the inverse eigenvalues has been stabilised at a threshold determined by both the Lagrange multiplier α and the noise variance $2\sigma^2$.

In acoustic imaging, the MNLS constraint is a standard feature of many iterative reconstruction schemes, added simply to ensure the uniqueness of the solutions obtained. The physical significance of this constraint is that *nonradiating sources* [86], *invisible objects* [73] or *ghost objects* [75], which correspond to the null-space of the observation model and have no projection in the field measurements; that is,

$$\mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} = 0 \quad \forall k, \quad (3.21)$$

are explicitly excluded from the image estimate $\hat{\mathbf{f}}$. In fact, under conditions of complete aperture sampling, it is observed that backpropagation solves the *inverse*

source problem in a MNLS sense by virtue of the exclusion of evanescent waves. For fully sampled apertures, therefore, SAFT, holography and related classical beamforming approaches represent inverse scattering solutions closely associated with the MNLS solution. Ladas and Devaney [75] and Devaney [61] have demonstrated that similar conclusions apply in limited view diffraction tomography by comparing the results of iterative and non-iterative reconstruction algorithms.

Recent examples of iterative approaches to MNLS regularisation include the work of Wang and Chew [109], and Lin and Chew [110]. In each of these cases, the MNLS criterion was combined with sophisticated non-linear approaches to inverse scattering. In a sense, the approach to be taken here is complementary, in that it addresses linearised inverse scattering problems with more sophisticated forms of regularisation; in particular, MaxEnt, which is non-linear³.

Before leaving the topic of MNLS regularisation, it can be noted that the minimisation of $J(\hat{\mathbf{f}})$ in Equation (3.19) is a standard quadratic problem, and virtually any multivariate minimisation algorithm based on gradient descent is satisfactory for its solution. Reviews of several iterative algorithms applicable to image reconstruction can be found in [85, 107]. The specific algorithm used to enable comparisons between quadratic regularisation and MaxEnt later in this chapter was the *Numerical Recipes in C* version of the Polak-Ribiere multidimensional conjugate-gradient algorithm **frprmn** [70]. Although the details of this standard algorithm will not be discussed here, the method of selecting α , which is common to all regularised algorithms compared in this chapter, will be discussed in detail in Section 3.2.5.

3.2.2 Constrained Least-Squares

In the previous section, the minimum-norm constraint was discussed as a means of ensuring the uniqueness of reconstructed images. In this section, additional criteria are examined for suppressing artefacts commonly observed in acoustic imagery.

In inverse scattering problems involving homogeneous media, such as those in NDE and underwater imaging, images are usually dominated by relatively high contrast targets set against low-amplitude backgrounds [111]. Whether apertures are

³ Wang and Chew [109] have also suggested MaxEnt as a regularisation approach in non-linear inverse scattering.

sparsely sampled or not, image artefacts arise in such circumstances from the combination of the band-limited spatial frequency response of the propagation transfer function and the band-limited temporal frequency responses of transducers. Particularly common artefacts are those of oscillation or "ringing" around sharp discontinuities in image intensity in addition to side-lobe clutter and speckle, though the latter is less of a problem in broadband imaging [111].

In the framework of Tikhonov regularisation, one approach to suppressing oscillation and clutter is to enforce smoothness by minimising the curvature (second partial derivative) of the image [57, 58, 70, 112]. In image processing, the second partial derivative, or Laplacian operator, can be calculated in the discrete domain by differencing neighbouring pixels according to the weights illustrated in Figure 3.2.




Figure 3.2: Discrete approximations to the Laplacian operator. (a) $m = 2$; (b) $m = 3$.

The use of finite difference approximations to the Laplacian operator in so-called constrained least-squares (CLS) regularisation was first proposed by Phillips [113], and later demonstrated in practical image restoration / reconstruction applications by Hunt [112]. By virtue of the shift-invariance of the Laplacian, the matrix $\mathbf{R}$ corresponding to its discrete approximation is block-circulant in m dimensions and can be efficiently convolved with lexicographically ordered image vectors via fast convolution. The gradient $\nabla Q(\hat{\mathbf{f}})$ can be similarly computed as

$$\nabla Q(\hat{\mathbf{f}}) = 2\mathbf{R}^{*T}\mathbf{R}\hat{\mathbf{f}}. \quad (3.22)$$

A minor difficulty arises in relation to expressing the CLS image estimate $\hat{\mathbf{f}}$ in a closed form comparable to Equation (3.20). This is because the Laplacian is an m -

dimensional operator, whereas the observation model used throughout this thesis is only semi-block-circulant; that is, block-circulant in $m-1$ dimensions. The eigenvectors associated with block-circulant and semi-block-circulant matrices cannot be associated with the same orthogonal decomposition.

The solution to this problem is either to formulate a closed form expression for the CLS image, regularised in only $m-1$ dimensions, or forgo the closed form interpretation of the solution altogether and compute the fully regularised image iteratively. Adopting the former approach, the semi-regularised CLS image, which, in any case, is applicable only to fully sampled apertures, is given by

$$\hat{\mathbf{f}} = \mathbf{W}^{*T} \text{diag} \left[\frac{1}{\lambda'_i + \alpha 2\sigma^2 \mu_i} \right] \mathbf{W} \sum_{k=1}^K (\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S}^T \mathbf{d}_k), \quad (3.23)$$

where the μ_i ($i=1,2,\dots,N$) are the eigenvalues of $\mathbf{R}^{*T}\mathbf{R}$ computed via $m-1$ dimensional DFTs. In the numerical simulations given in Section 3.5, the iterative approach to CLS regularisation is used exclusively, whereby the images obtained are fully regularised in three dimensions. It can be noted, however, that Equation (3.23) is closely related to the classic Wiener or least-mean-squares (LMS) filter formulation [57, 58, 85, 107]. In terms of the $m-1$ dimensional wavenumber vector $\mathbf{k}$, the eigenvalues μ_i can be written directly as [2]

$$\mu_i = \|\mathbf{k}\|^4, \quad (3.24)$$

so that Equation (3.23) enforces smoothness by simply suppressing high spatial-frequencies. This contrasts with the usual stochastic Wiener filter approach, in which the filter response depends on the signal-to-noise ratio of the data as a function of spatial-frequency. As Wiener filtering can be subsumed into the general deterministic framework of Tikhonov regularisation, it will not be discussed further.

3.2.3 Edge Preserving Regularisation

At the simplest level, regularisation can be viewed as a means of reducing the effects of zero eigenvalues in the integral operators of forward problems that otherwise become infinite when solving inverse problems [114]. A deficiency in this interpretation is that it does not acknowledge the additional improvements in solutions that can accompany more sophisticated approaches to regularisation.

In the Tikhonov framework, CLS regularisation was seen to enforce both uniqueness and stability in images by suppressing high spatial frequencies. The flaw in this approach is that regions in images containing sharp transitions in intensity that correspond to edges are smoothed in addition to homogeneous regions. As the underlying normal equations to be solved in reconstructing quadratically regularised images are linear, ultrasonic images reconstructed in this way are limited to linear combinations of the measured eigenvectors [70]. With band-limited ultrasonic data, the resolution obtainable is therefore limited to that determined by the Rayleigh criterion defined in Chapter 1.

Deterministic "edge-preserving" regularisation [102-104, 115-119] is based on a generalisation of the Tikhonov CLS formulation whereby the regularising influence of the modulus of the image gradient $|\nabla f|$ is controlled adaptively via a function $\varphi(t)$ in addition to a Lagrange multiplier α . Equation (3.17) thus becomes

$$J(\hat{\mathbf{f}}) = \chi^2 + \alpha \int_{ROI} \varphi[|\nabla f|] dr. \quad (3.25)$$

For the case $\varphi(t) = t^2$, Equation (3.25) is equivalent to CLS regularisation. If $\varphi(t)$ is non-quadratic, however, the normal equations to be solved in minimising $J(\hat{\mathbf{f}})$ become non-linear, and the smoothing behaviour of the right-hand regularising term can be controlled; that is decreased, in the vicinity of edges, where $|\nabla f|$ is large. This is particularly important when reconstructing images from ultrasonic data, as it enables a degree of image smoothing without significant resolution loss.

Unfortunately, by altering Equation (3.25) from its quadratic form, two major difficulties are introduced. Firstly, when the underlying normal equations become non-linear, there are no analytic solutions comparable to those obtained for MNLS and CLS regularisation; images must always be computed iteratively. Secondly, the property of convexity⁴ that ensures that the function $J(\hat{\mathbf{f}})$ has only a single minimum is lost, depending on the exact form of $\varphi(t)$, and so algorithms

⁴ A function $\varphi(t)$ is convex in the interval (a,b) if $\varphi(\eta t_1 + (1-\eta)t_2) \leq \eta\varphi(t_1) + (1-\eta)\varphi(t_2)$ for all $t_1, t_2 \in (a,b)$ and proportions $0 \leq \eta \leq 1$. $\varphi(t)$ is strictly convex if the equality holds only for $\eta = 0$ and $\eta = 1$. Some examples of strictly convex functions are x^2 , e^x , $\log(1/x)$ for $x > 0$, and $x \log x$ for $x > 0$.

designed to minimise $J(\hat{\mathbf{f}})$ must contend with both local and global minima [70, 103, 119]. As most gradient descent algorithms easily become stuck in local minima, computationally expensive algorithms are usually required to perform non-convex optimisations [117, 120-123].

In the context of computed image reconstruction from sparse ultrasonic data, very few applications of edge-preserving regularisation have been reported. This is undoubtedly due to the computational complexity associated with non-convex optimisations [124] and the lack of proven convergence of the associated algorithms for underdetermined problems [103]. Although edge-preserving regularisation is an advance over quadratic regularisation in terms of noise suppression, it does not specifically address the issue of missing information. This requires the introduction of prior information via more powerful Bayesian probabilistic formulations.

3.2.4 Bayesian Image Reconstruction

To this point, regularised image reconstruction has been viewed deterministically. An alternative interpretation stems from Bayesian probability, which has provided the framework for many recent advances in inverse problem theory [7, 105, 106, 125, 126]. Until recently, many researchers did not acknowledge the validity or necessity of the Bayesian approach to solving inverse problems, or may have believed that it was in some way controversial [127]. In image processing for example, Klug and Crowther [128] define a reliable image reconstruction scheme as "one which extracts maximum information from the data while preventing spoiling of the results by effects which are beyond resolution with the limited number of views available." While this statement is indisputable, they go on to state that "if the object is general and the projections contain noise, no mathematical or computational manipulation can make up for the lack of information in the data."

The reason for this somewhat pessimistic outlook is rooted in the orthodox *frequentist* interpretation of probability, which denies the role of prior information in problems of inference. Bayesian probability theory, on the other hand, always assumes the existence of some form of prior information, however general, and it is through such *priors* that information can actually be added to images during restoration or reconstruction. The fact remains, however, that successfully

introducing prior information to resolve uncertainty in problems of inference is difficult in practice, for while Bayesian probability theory is powerful, it often results in computationally unwieldy problems [129].

The main distinction between frequentist and Bayesian interpretations of probability is that the former are based on the average long-run behaviour of repeated experiments, whereas the latter deal strictly with *degrees of belief* assigned numerical values between zero and one. This latter interpretation was actually the one adopted by Bernoulli [130] early in the development of probability theory, who's law of large numbers was intended to imply a link (but not necessarily an equality) between relative frequencies and probabilities⁵. In addition to further codifying these concepts, Bayes [132] derived the important relationships between *prior* (before data) probabilities and *posterior* (after data) probabilities that were subsequently rediscovered and applied by Laplace [133]. In particular, the theorem for which Bayesian probability theory is known; Bayes' Theorem, is usually stated as

$$p(H|DI) = p(H|I) \frac{p(D|HI)}{p(D|I)}, \quad (3.26)$$

where

- $p(H|DI)$ is the posterior probability of the hypothesis H conditioned on both the data D and the background information I regarding the problem.
- $p(H|I)$ is the prior probability of H given I .
- $p(D|HI)$ is the *likelihood* of D given both H and I .
- $p(D|I)$ is the probability of D given I , also known as the *evidence*.⁶

Making the background information I implicit for clarity, Bayes' Theorem can be restated in the notation of discrete image reconstruction as

$$p(\hat{\mathbf{f}}|\mathbf{d}) = p(\hat{\mathbf{f}}) \frac{p(\mathbf{d}|\hat{\mathbf{f}})}{p(\mathbf{d})}, \quad (3.27)$$

where

⁵ See Jaynes [14, 131] and Loredó [15] for historical insights into the recent renaissance in Bayesian probability.

⁶ In certain Bayesian problems involving parametric modelling, the evidence is used for model selection [134].

- $p(\hat{\mathbf{f}}|\mathbf{d})$ is the posterior probability of the image $\hat{\mathbf{f}}$ given the data $\mathbf{d}$.
- $p(\hat{\mathbf{f}})$ is the prior probability of the image $\hat{\mathbf{f}}$.
- $p(\mathbf{d}|\hat{\mathbf{f}})$ is the likelihood of the data given the image $\hat{\mathbf{f}}$.
- $p(\mathbf{d})$ is the evidence.

As the evidence term $p(\mathbf{d})$, can usually be taken as a constant for a given data set [105, 135], Equation (3.27) can be rewritten

$$p(\hat{\mathbf{f}}|\mathbf{d}) \propto p(\mathbf{d}|\hat{\mathbf{f}}) p(\hat{\mathbf{f}}), \quad (3.28)$$

which is the classic Bayesian formulation for image reconstruction and restoration, where the posterior probability is proportional to the likelihood times the prior. To make the likelihood and prior additive rather than multiplicative, it is customary to deal with the logarithms of probabilities such that (3.28) becomes

$$\ln[p(\hat{\mathbf{f}}|\mathbf{d})] \propto \ln[p(\mathbf{d}|\hat{\mathbf{f}})] + \ln[p(\hat{\mathbf{f}})]. \quad (3.29)$$

There is an immediate and appealing interpretation for Equation (3.29) in view of the discussion of Tikhonov regularisation earlier in this chapter. Firstly, with the assumption of Gaussian noise, the likelihood $p(\mathbf{d}|\hat{\mathbf{f}})$ can be calculated directly as

$$p(\mathbf{d}|\hat{\mathbf{f}}) = \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^M \exp\left(-\frac{1}{2} \chi^2 \right), \quad (3.30)$$

where M is the number of complex data. Substituting this expression into Equation (3.29) and placing the Lagrange multiplier α before the log-prior gives

$$\ln[p(\hat{\mathbf{f}}|\mathbf{d})] \propto \alpha \ln[p(\hat{\mathbf{f}})] - \chi^2, \quad (3.31)$$

whereby $\ln[p(\hat{\mathbf{f}})]$ is seen to have a role similar to that of the regularising functional $Q(\hat{\mathbf{f}})$ in Equation (3.17). Selecting the mode of the posterior probability by maximising Equation (3.31) yields the maximum *a posteriori* (MAP) image estimate

$$\hat{\mathbf{f}}_{MAP} = \arg \max_{\hat{\mathbf{f}}} \left(\alpha \ln[p(\hat{\mathbf{f}})] - \chi^2 \right), \quad (3.32)$$

or more closely according to the Tikhonov formulation,

$$\hat{\mathbf{f}}_{MAP} = \arg \min_{\mathbf{f}} \left(\chi^2 - \alpha \ln[p(\hat{\mathbf{f}})] \right). \quad (3.33)$$

As previously, the Lagrange multiplier α is required to mediate between the influence of the chi-squared constraint and the regularising operator. The role of the latter, however, is now wider than implied by the earlier discussion, as $p(\hat{\mathbf{f}})$ can embody *probabilistic* prior information about the image $\mathbf{f}$ ⁷. Moreover, these probabilities are not restricted to describing the long-run *statistical* behaviour of *random variables*, but are open to the full Bayesian interpretation whereby non-statistical preferences to certain types of images can also be incorporated. Loredó [15] refers to this aspect of Bayesian probability as a kind of "quantitative epistemology" for the numerical encoding of one's state of knowledge.

One early criticism of Bayesian probability was that although it adequately described how to incorporate prior probabilities in calculating posterior probabilities, it did not describe how to *assign* these prior probabilities in the first instance. This question is still an open one and subject to intensive research. Aside from the MaxEnt prior, which has strong deterministic and probabilistic justifications, as described later in this chapter, various other image priors have been applied to general image processing.

One particularly successful family of image priors has been that based on Markov random fields (MRFs) [105, 106, 115]. In MRF priors, the conditional probabilities (intensities) of pixels are dependent on their configuration within *local* neighbourhood systems or *cliques*. With suitable non-linear weighting of these dependencies, it is possible to encode general prior knowledge such as [105]:

- Straight running edges tend to continue running straight.
- Edges do not usually end blindly.
- Corners are most often right-angled.
- Pixels separated by edges are usually uncorrelated.

⁷ In fact, the stochastic Wiener filter mentioned in Section 3.2.2 also yields a kind of MAP estimate in which the *a priori* assumptions *imposed* on the problem are that the statistics of the image and the data are jointly Gaussian [107, 125].

The optimisation problem posed by Equation (3.33) has essentially the same form as that posed by Equation (3.25) for deterministic edge-preserving regularisation. In the MRF context, however, $\varphi(t)$ takes the form of a potential function for a Gibbs energy [106, 135-137], which allows considerably richer *a priori* structure to be encoded than is possible via deterministic regularisers. Unfortunately, this level of flexibility is also associated with extreme non-convexity. Modified Gauss-Newton and conjugate gradient algorithms fail to converge in any sense, and steepest descent algorithms are extremely slow. While stochastic relaxation algorithms based on simulated annealing (SA) require extreme amounts of computation to reach a global minimum [103], they offer one of the most general and tractable routes to solving large-scale Bayesian problems involving MRF priors.

As with deterministic regularisation, published accounts of the application of MRF priors to computed ultrasonic imaging are scarce, largely due to the complexity of the optimisations involved. This situation may not change in the near future, as it has recently been noted that there are additional convergence problems associated with SA algorithms when PSFs display significant non-local characteristics [127, 138] (time-harmonic Green's functions are effectively infinite!). In such cases, simulated annealing algorithms may not yield optimal results in a finite amount of time.

3.2.5 The Discrepancy Principle

At the beginning of this chapter, the measure of the agreement between the forward scattering model and the experimental data was established according to Equation (3.2). Low values of χ^2 indicate close agreement, whereas high values indicate disagreement. The function of the regularising parameter α , therefore, is to arbitrate between the data $\mathbf{d}$ and the regularising operator Q so that the scattering model approximates the true signal with minimal noise amplification.

As $Q(\hat{\mathbf{f}})$ is a function of $\hat{\mathbf{f}}$, the selection of α is best based on the actual appearance of the image, however such a selection procedure is difficult to automate in an image reconstruction algorithm. A simpler approach, based on the *discrepancy principle* [10, 11, 46] uses the statistical behaviour of the noise, which can usually be quantified according to its variance.

According to the discrepancy principle, the statistical expectation of χ^2 is equal to the number of measurements M , which in the time-harmonic case is simply the number of aperture samples D . In the broadband case, D must be multiplied by the number of discrete frequencies K , such that $M = DK$. In minimising $J(\hat{\mathbf{f}})$, therefore, the value of α is adjusted until χ^2 achieves this expected value, and the set of all images for which this condition is achievable is the set of *feasible* solutions. This contrasts with inverse filtering, where χ^2 is optimised to achieve its global minimum, with all the associated symptoms of ill-posedness noted previously.

There are two disadvantages associated with the use of the discrepancy principle in selecting α . The first of these is that it requires *a priori* knowledge of the noise variance. In fact, there are alternatives to the discrepancy principle based variously on Bayesian probability, cross-validation and generalised maximum likelihood for which no such *a priori* knowledge is required [10, 116, 139, 140]. Unfortunately, the complexity of these methods rules them out of consideration in this instance.

The second problem is that, except in cases of extremely low signal-to-noise ratio, it is well known that the discrepancy principle causes underfitting of the data [10, 139, 140]. This is a result of there being both "good" and "bad" eigenvalues (Fourier coefficients) transmitted in any real data set, which correspond to strongly and weakly transmitted eigenvectors (spatial frequencies). The use of the expectation M , therefore, which assumes an equal degree of contamination for all measured eigenvalues, results in an overly conservative estimate of image structure. As neither of these problems is particularly serious in relation to other anticipated difficulties (especially in relation to approximations in the scattering model), the discrepancy principle has been used extensively throughout this thesis as part of MaxEnt and other regularised image reconstruction algorithms.

Over surfaces defining the maxima of $Q(\hat{\mathbf{f}})$, χ^2 increases monotonically with α , and therefore any simple approach to achieving $\chi^2 \cong M$, such as a binary search on α , is effective [50]. In the MNLS and CLS implementations used later in this chapter, the *Numerical Recipes in C* bracketing function **zbrent** [70], which incorporates inverse quadratic interpolation, was used with good results. The method used to set α in MaxEnt image reconstruction, although also essentially a binary search, requires additional explanation, which is deferred until Chapter 4.

3.3 MaxEnt Regularisation

The previous sections of this chapter introduced deterministic and probabilistic interpretations for regularised image reconstruction in terms of the forward scattering model derived in Chapter 2. This section will summarise the properties of conventional MaxEnt as it applies to image reconstruction in general, leaving the details of its adaptation to coherent ultrasonic imaging to Chapter 4.

As noted in the introduction, MaxEnt arose from the reinterpretation of the work of Shannon [12] by Jaynes [13, 131], who saw the Shannon entropy, defined for a set of discrete, independent and exhaustive probabilities p_i ($i=1,2,\dots,N$) by

$$S = -\sum_{i=1}^N p_i \ln p_i \quad (3.34)$$

to be applicable to modelling the uncertainty associated with problems of inference. In image processing, Frieden [141] first described the basic characteristics of maximum entropy images through the super-resolution of one-dimensional (1-D) incoherent objects with a diffraction-limited imaging system. Frieden's work showed that by using entropy as a regulariser and equating pixel intensities to probabilities, it was possible to *extrapolate* frequency information meaningfully outside a given pass-band, thereby providing a useful super-resolution effect. The main properties of the images $\hat{\mathbf{f}}$ reconstructed via Frieden's method are summarised as follows:

- $\hat{\mathbf{f}}$ is restricted to be positive, as required for an image of an incoherent object (or where a complex-valued image consists of real and imaginary parts; both must be positive).
- Each measurement in the data vector $\mathbf{d}$ contributes a degree of freedom (DOF) to $\hat{\mathbf{f}}$. As the number of measurements decreases, $\hat{\mathbf{f}}$ becomes increasingly conservative (smooth), reflecting only the minimum *configurational information* consistent with the measurements.
- $\hat{\mathbf{f}}$ is not intrinsically band limited. The resolution obtained may significantly exceed the Rayleigh limit.

In subsequent work in Fourier synthesis radio-astronomy, Wernecke and D’Addario [1] employed a form of entropy proposed by Burg [16] for autoregressive spectral estimation, given by

$$S_{Burg} = \sum_{i=1}^N \ln f_i \quad (3.35)$$

where the f_i are pixel intensities as opposed to normalised probabilities. The results obtained were complementary to those of Frieden in that they showed that maximum entropy (albeit with a different entropy) allowed sparse frequency data to be *extrapolated* to produce images with low distortion and superior resolution. The formal theory of MaxEnt, based on the Shannon entropy, was later popularised by Skilling *et al* [2, 45, 46] through a complicated but robust algorithm that made it possible for the first time to apply MaxEnt to large-scale problems. That algorithm, known as the “Cambridge Algorithm”, is the one that has been extended to ultrasonic image reconstruction in Chapter 4 and subsequently applied to synthetic aperture and array based imaging experiments in Chapters 5 and 6 of this thesis.

In Section 3.2, it was noted that regularising operators used in image reconstruction should be selected to enforce desirable properties in reconstructed images. Although the framework of Tikhonov regularisation is well justified by both probabilistic and deterministic considerations, the selection of priors to enforce these desirable characteristics is usually *ad hoc*. In contrast, the Shannon entropy, which is used to assign prior probabilities in MaxEnt, has strong justifications from both probabilistic and deterministic points of view, as will be illustrated shortly.

3.3.1 The Shannon Entropy

The question of which mathematical function is *the* uniquely appropriate measure of probabilistic uncertainty in image reconstruction, and in fact, whether there is any such function, has been the subject of much controversy. Although Wernecke and D’Addario [1] achieved good results with the Burg entropy, this was later shown to be technically inappropriate as a measure of general image entropy [50].

Many do not agree that there is any "magical function" for image regularisation. Landl and Anderssen [101], for example, have suggested that the popularity of maximum entropy relates more to the way it enforces positivity, than to its philosophical credentials. Ponman [142], despite acknowledging the effectiveness of

the Shannon entropy as a regulariser, has argued that the specific expression used is not critical, citing alternative expressions purported to have similar properties. In the opinion of Press *et al* [70], the definition of entropy is an open one. They state that various simple functions other than the strict Shannon entropy may be employed with little impact on the resulting images. Similar views have been expressed by Nityananda and Narayan [143] and Högbom [144].

That the Shannon entropy S is at least fundamentally different from the edge-preserving regulariser given by Equation (3.25) is apparent from the fact that it does not difference neighbouring pixels or rely in any way on the image gradient ∇f . Because it does not penalise sharp transitions in reconstructed images, S is an edge preserving regulariser, although by being convex, it does not suffer from the convergence problems associated with many other edge-preserving regularisers.

Despite its global as opposed to local characteristics, the Shannon entropy does achieve a useful smoothing effect, however. This can be seen by considering the entropy of a binary ensemble assigned the complementary probabilities p and $(1 - p)$ as illustrated in Figure 3.3. In this simple case, as in any case, the maximum Shannon entropy is attained for equal probabilities (normalised pixel intensities) $p = 1 - p = 0.5$. Thus, maximising S favours smooth images.




Figure 3.3: The binary entropy function $S(p, 1 - p)$ as a function of p .

Invoking isomorphisms between pixel intensities and probabilities is an intuitive way of justifying the use of entropy as a regulariser, but because it does not always make sense to equate pixel intensities to probabilities, it is useful to dispense with normalisation [145]. Introducing an initial estimate or *model* $\mathbf{m}$ for $\mathbf{f}$, Equation (3.34) becomes the cross-entropy [124, 146] or Kullback-Leibler distance [147, 148] given by

$$S = - \sum_{i=1}^N f_i \ln (f_i / m_i) . \quad (3.36)$$

In the absence of detailed preconceptions as to the structure of $\mathbf{f}$, the m_i can be replaced by a uniform initial estimate m , whereby Equation (3.36) becomes

$$S = - \sum_{i=1}^N f_i \ln (f_i / m) . \quad (3.37)$$

In most MaxEnt implementations, m is replaced by eA so that instead of the f_i converging to $e^{-1}m$ in the absence of relevant information, they converge to the user-defined value A [3]. In any case, the *shape* of the image $\hat{\mathbf{f}}$ is virtually independent of A provided that it is sufficiently small [2, 149].

While other functions, such as the Burg entropy possess similar mathematical properties, it can still be argued that Equation (3.37) is the unique best choice for regularising images without more detailed prior information [145, 146, 150-153]. To justify this elevated status, both probabilistic and deterministic arguments can be invoked, among which are the monkey and kangaroo arguments.

3.3.2 The Monkey Argument

According to Jaynes [131], prior probability assignments for problems of inference should have maximum entropy while agreeing with prior knowledge. In a frequentist context, the probability distribution that maximises the entropy is numerically identical to the frequency distribution that can be realised in the greatest number of ways. The monkey argument, or less colourfully, the multiplicity or combinatorial argument, was the first by which MaxEnt was formally justified [44, 70, 154, 155]. According to this argument, any image $\hat{\mathbf{f}}$ with N pixels can be viewed as a discrete distribution of indivisible quanta (balls) distributed by some random agent (traditionally a troupe of monkeys). If the quanta are denoted δ such that

$$f_i = u_i \delta, \quad (3.38)$$

then the *proportion* of quanta in each pixel f_i is given by

$$f_i = \frac{u_i}{U}, \quad (3.39)$$

where U is the total number of quanta in the image and the u_i are the integral numbers of quanta per pixel. The normalisation conditions are therefore

$$\sum_{i=1}^N u_i = U \quad \text{and} \quad \sum_{i=1}^N f_i = 1 \quad (3.40)$$

Under these conditions, the probability of obtaining any particular image $\hat{\mathbf{f}}$ by chance is proportional to the number of ways it can be produced, as given by [149]

$$p(\hat{\mathbf{f}}) \propto \frac{U!}{\prod_{i=1}^N u_i!}. \quad (3.41)$$

For large u_i , Sterling's approximation to the factorial

$$\ln(x!) \cong x \ln x - x \quad (3.42)$$

allows equation (3.41) to be expressed in the form [70]

$$p(\hat{\mathbf{f}}) \propto \exp \left[- \sum_{i=1}^N u_i \ln \left(\frac{u_i}{U} \right) + \frac{1}{2} \left(\ln U - \sum_{i=1}^N \ln u_i \right) \right], \quad (3.43)$$

which, by ignoring minor terms reduces to

$$p(\hat{\mathbf{f}}) \propto \exp \left[- \sum_{i=1}^N u_i \ln \left(\frac{u_i}{U} \right) \right], \quad (3.44)$$

giving

$$\frac{1}{U} \ln p(\hat{\mathbf{f}}) \cong - \sum_{i=1}^N f_i \ln f_i = S(\hat{\mathbf{f}}). \quad (3.45)$$

The prior probability is therefore given by

$$p(\hat{\mathbf{f}}) \propto \exp[\alpha S(\hat{\mathbf{f}})], \quad (3.46)$$

where the constant α is essentially the total number of quanta U . The interpretation of the entropy here is as the logarithm of the prior probability required by Bayes'

theorem. Unfortunately, when noise is included, this model cannot be normalised, for if U is allowed to become arbitrarily large, then the entropy swamps the likelihood and the image converges to a featureless scene [153, 155]. The usual way of overcoming this problem in MaxEnt is to partially destroy the Bayesian structure of the problem and replace the likelihood function with a hard constraint on χ^2 according (for example) to the discrepancy principle discussed in Section 3.2.5.

As Equation (3.46) is based on very general assumptions, it is known as a minimally informative, or *free-form* prior. This prior is not completely Bayesian, firstly because of the frequentist arguments associated with Equation (3.41), and secondly because χ^2 has been treated as a hard constraint rather than a “nuisance parameter” to be marginalised away according to standard Bayesian practice [10, 134, 156]. In [129], Skilling suggests that MaxEnt image reconstruction is to be preferred to the direct Bayesian approach, which leads to “impossibly large problems of computation and display in the space of all images.”

The monkey argument, although offering intuitive justifications for MaxEnt, is not universally convincing, as the association between images and probability distributions is an abstract one [49]. In addition, the physical interpretation of the quantisation parameter U is unclear without detailed knowledge of the physical *microstates* of specific problems [153]. To address these defects, there has been a slight paradigm shift in recent years towards justifying MaxEnt on the grounds of decision-theoretic self-consistency. The corresponding argument, originally known as the kangaroo argument, is discussed next.

3.3.3 The Kangaroo Argument

To avoid the association of images with probability distributions, and the dubious requirement for quantisation, Shore and Johnson [146, 157] gave an axiomatic derivation of MaxEnt that was independent of probabilistic considerations. A related derivation from a more physical perspective was later given by Tikochinsky *et al.* [158], and this was subsequently simplified by Gull and Skilling [49] to what is now known as the kangaroo, or logical consistency argument.

The kangaroo argument is concerned with the consistency of inferences made in the absence of certainty. The case originally cited concerned the eye colour and handedness of kangaroos, which in being (notionally) independent events, could be

represented as marginal data for a 2×2 contingency table. In answering the question of how best to infer the proportion of kangaroos that are both blue-eyed and left-handed, Gull and Skilling [49] demonstrated that MaxEnt, and specifically the Shannon entropy, gave correct answers; that is, they did not introduce unwarranted correlations between handedness and eye colour.

Paralleling the slightly more formal version of the kangaroo argument given by Livesey and Skilling [153], it is easily demonstrated that the requirement for “system independence” [146] (no unwarranted correlations) is at the heart of the MaxEnt phenomenon. In the two-dimensional $a \times b$ contingency table illustrated in Figure 3.4, the p_{ij} are *proportions* related to the marginal data u_i and v_j by

$$\left. \begin{aligned} u_i &= \sum_{j=1}^b p_{ij} \quad (i=1,2,\dots,a) \\ v_j &= \sum_{i=1}^a p_{ij} \quad (j=1,2,\dots,b) \end{aligned} \right\} \quad (3.47)$$

		$\rightarrow j$					
	p_{11}	p_{12}	$\dots$	$\dots$	p_{1b}	$\rightarrow$	u_1
	p_{21}	p_{22}	$\dots$	$\dots$	p_{2b}	$\rightarrow$	u_2
	$\vdots$	$\vdots$			$\vdots$	$\rightarrow$	$\vdots$
	p_{a1}	p_{a2}	$\dots$	$\dots$	p_{ab}	$\rightarrow$	u_a
$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$		$\downarrow$
	v_1	v_2	$\dots$	$\dots$	v_b	$\rightarrow$	1

Figure 3.4: Two-dimensional contingency table and normalised marginal data.

Given the geometry of the problem, there is no reason to expect correlations between the row and column structure of Figure 3.4, and the only rule for combining the marginals such that this is true is $p_{ij} = u_i v_j$ [153]. In attempting to infer the p_{ij} from the scalar or vector functions $U(u)$ and $V(v)$, however, there are two legitimate approaches. Firstly, given their independence, it is acceptable to solve for

the marginals u_i or v_j from $U(u)$ and $V(v)$ separately, and then compute the $p_{ij} = u_i v_j$ indirectly. Secondly, as the result of combining two disjoint (independent) subsets of information in a *consistent* mode of inference should not alter the outcome, it should also be acceptable to solve for the p_{ij} directly using $U(u)$ and $V(v)$ simultaneously. Proceeding via the first method, and solving for the row and column structures independently, the MaxEnt formulation for each sub-problem consists of an entropy term, a constraint equation and a normalisation equation, viz.

$$\left. \begin{aligned} \max_i S(u) &\equiv \max_i \left\{ - \sum_{i=1}^a u_i \ln u_i \right\} & U(u) & \sum_{i=1}^a u_i = 1 \\ \max_j S(v) &\equiv \max_j \left\{ - \sum_{j=1}^b v_j \ln v_j \right\} & V(v) & \sum_{j=1}^b v_j = 1 \end{aligned} \right\}. \quad (3.48)$$

By augmenting the constraint equations with Lagrange multipliers and equating the total differentials to zero, the u_i and v_j can be obtained by separately maximising

$$\left. \begin{aligned} \ln u_i &= -\lambda \frac{\partial U}{\partial u_i} - \alpha - 1 \\ \ln v_j &= -\mu \frac{\partial V}{\partial v_j} - \beta - 1 \end{aligned} \right\}, \quad (3.49)$$

whence the p_{ij} can then be obtained indirectly by maximising

$$\ln p_{ij} = \ln u_i + \ln v_j = -\lambda \frac{\partial U}{\partial u_i} - \mu \frac{\partial V}{\partial v_j} - \alpha - \beta - 2, \quad (3.50)$$

which is also precisely the result of applying MaxEnt to infer the p_{ij} directly by using $U(u)$ and $V(v)$ together. In particular, it can be noted that the constraint differentials have added by virtue of the functional form of the Shannon entropy. Of fundamental importance is the fact that the entropy gradient is logarithmic, as the logarithm is the only function that converts products into sums [145]. The rationale for image entropy being a simple sum over pixels also follows, as this is in keeping with minimising unnecessary correlations between neighbouring pixels [145, 159].

This simple example illustrates that the Shannon entropy separates the influence of data in different dimensions by rejecting unnecessary correlations. In non-trivial

problems, the "dimensions" are no longer completely orthogonal, and the marginals are replaced by complex integral transformations. In these cases also, if consistent inference is to be applied, then MaxEnt is the method of choice.

As noted by Skilling [51], "Programming maximum entropy is not trivial." Due to the nonlinearity of the Shannon entropy, a constrained non-linear optimisation is required. In relation to ultrasonic image reconstruction, there are additional problems, as the images are complex-valued, and entropy is not even strictly defined. For this reason, detailed discussion of the optimisation algorithm used in MaxEnt and its adaptation to complex-valued images is left until Chapter 4.

3.4 The CLEAN Algorithm

To conclude this survey of image reconstruction concepts, the non-linear CLEAN algorithm will now be briefly reviewed. The CLEAN algorithm is important in any critical evaluation of MaxEnt, as the two approaches have been in competition for many years in astronomical circles. Originally invented by Högbom [48] as a means of overcoming sparse Fourier-space sampling in synthesis radio-astronomy, CLEAN was subsequently applied to coherent microwave imaging by Tsao and Steinberg [35]. In this thesis, CLEAN has been extended to reconstruct images from broadband ultrasonic data for comparison with those generated by MaxEnt.

The essential feature of CLEAN is that it iteratively finds peaks in so-called "dirty images" that have been corrupted by side-lobe artefacts associated with an imaging system. With knowledge of the "dirty beam" responsible for these artefacts, it is possible to subtract discrete targets from the image, recording their locations and intensities, and at the same time remove their associated side-lobe contributions. Following this, the next-strongest target is found and the process is repeated until legitimate targets can no longer be distinguished from noise impulses. The final phase of the reconstruction is to convolve the set of recorded targets with a "clean beam" that is devoid of any side-lobes. This results in a CLEANed image.

The main assumption that CLEAN makes in relation to the image distribution is that it is dominated by high-contrast impulsive targets. If the images are of a substantially different character, with extended homogeneous regions, the algorithm is known to perform poorly in terms of both reconstruction speed and quality.

Unlike regularised images, the results of CLEAN cannot be expressed in any simple mathematical form, and consequently it is difficult to make statements as to the effects of noise [160]. The SNR of the data generally influences the number of iterations that can be completed before noise peaks are selected as legitimate targets, and knowledge of the SNR permits this condition to be identified.

Details of CLEAN algorithms can be found in [35, 48, 144, 161]. The implementation discussed here is broadband, combining several different frequencies to enable side-lobe rejection in both lateral and range dimensions. The process of calculating the dirty beam is precisely the same as that discussed in Section 2.3.3 for calculating the incident field generated by an arbitrary aperture. In the case of monostatic data acquisition, the two-way beam pattern is used. Initially, the dirty image is calculated by matched filtering, as discussed in Section 3.1.2, giving

$$\hat{\mathbf{f}}_{dirty} = \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{d}_k, \quad (3.51)$$

Following detection of the peak amplitude and storage of its real and imaginary parts in both an accumulating vector $\mathbf{s}$ and a vector $\mathbf{t}$, which is zeroed at the commencement of each iteration, a new dirty beam

$$\mathbf{b}_{dirty} = \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{t}}, \quad (3.52)$$

is generated specifically for this target and is then subtracted from the previous dirty image according to

$$\hat{\mathbf{f}}_{dirty} \Leftarrow \hat{\mathbf{f}}_{dirty} - \mathbf{b}_{dirty}. \quad (3.53)$$

During this process, the algorithm must be protected from spurious target detection. Tsao and Steinberg [35] proposed a scheme involving both absolute and relative intensity thresholds that was successfully adapted here to broadband image reconstruction. In this work, only the absolute threshold was employed, which modelled the peak of the background intensity level of the image below which new targets could not be reliably detected. The absolute threshold increases with each target detected, and at each iteration is given by [35]

$$T_{abs} = \frac{B_p}{M} \left(\|\mathbf{s}\|^2 [1 - \exp(-\sigma_\phi^2)] + \frac{1}{2\text{SNR}} \right), \quad (3.54)$$

where B_p is the ratio of the peak intensity of the (broadband) beam pattern to its average intensity, and M is the number of measurements, being the product of the number of active array elements D and the number of frequencies K . By ignoring the phase errors arising from element misplacement and medium anomalies, the phase error variance σ_ϕ^2 can be related solely to the receiver noise by [35]

$$\sigma_\phi^2 = \frac{1}{2\text{SNR}}. \quad (3.55)$$

In the final stage of the CLEAN algorithm, the target vector $\mathbf{s}$ is convolved with an m -dimensional Gaussian kernel with a width comparable to that of the main beam of the aperture under consideration. The benefits of CLEAN, therefore, are primarily in side-lobe cancellation as opposed to resolution enhancement [162].

3.5 Comparative Simulations

The comparative simulations in this section are intended to illustrate the superior performance of MaxEnt relative to the other approaches to image reconstruction discussed in this Chapter. Simulations are believed to be necessary in this regard, as they allow fine control over experimental parameters such as the amplitude and statistical distribution of the noise, and the scattering behaviour of the targets.

3.5.1 Simulation Parameters

Figure 3.5 illustrates the combination of scatterers simulated, which included a 13.5 mm diameter disk with a relative reflectivity of 20 %, and four point scatterers with relative reflectivities of 100 %. All scatterers were modelled as being impenetrable with Neumann boundary conditions. The ROI was discretised over a $64 \times 64 \times 64$ grid with a uniform spacing of 0.5 mm. The range offset R from the aperture plane to the ROI (see Figure 2.2) was 50.0 mm, and the wave velocity in the medium was set to that of H_2O at 25°C (approximately 1490 ms^{-1}). Monostatic data acquisition was simulated for a transducer with a 6dB beam-width of 90° .




Figure 3.5: Plan and elevation of the objects used in reconstruction simulations.

Numerically, the procedure used for simulating the scattered fields was identical to that previously described in Section 2.5 for the flat-bottom hole simulation. Again, a Gaussian weighted pulse was simulated, with time and frequency-domain characteristics as illustrated in Figure 3.6. Gaussian noise was simulated via the *Numerical Recipes in C* function **gasdev** [70]. In Figure 3.6 (b), the range of frequencies actually used in image reconstruction is indicated. In order to pose a more severe test for the image reconstruction algorithms considered, the number of discrete frequencies was restricted to ten in the range from 0.92 MHz to 1.09 MHz.


Figure 3.6: The synthetic pulse used in image reconstruction simulations. (a) Time-domain; (b) Frequency domain, $f_c = 1$ MHz $BW_{6dB} = 0.8$ MHz (10 frequencies used).

Table 3.1 summarises the experimental conditions that were simulated for each of the six image reconstruction algorithms compared, which included inverse filtering (unconstrained least-squares), matched filtering, minimum-norm least-squares, constrained least-squares, broadband CLEAN and complex-valued MaxEnt.

	Average SNR	Aperture Sampling
TEST 1	20 dB	100 %
TEST 2	20 dB	10 %
TEST 3	0 dB	10 %

Table 3.1: Simulated experimental conditions for algorithm comparison.

In Figure 3.7, the real parts of the aperture data computed at 1.0 MHz are shown for SNR values of 20 dB and 0dB, along with the random pattern used to simulate sparse sampling. Noteworthy is the strong specular reflection from the planar disc, which renders the waves from the point scatterers almost invisible.




Figure 3.7: Simulated data at 1.0 MHz and sparse sampling pattern for comparison of reconstruction methods; (a) SNR = 20dB; (b) SNR = 0dB; (c) 10 % sampling pattern.

3.5.2 Simulation Results

The next six pages contain images reconstructed via the six algorithms considered. Each set of images is arranged vertically according to the test conditions in Table 3.1, and shown in plan view and elevation according to the colour scale in Figure 3.5.




Figure 3.8: Inverse filter (unconstrained least-squares) image reconstruction results. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB. (Plan views at left – elevations at right).




Figure 3.9: Results of matched filter image reconstruction. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB. (Plan views at left – elevations at right).




Figure 3.10: Results of minimum-norm least-squares (MNLS) image reconstruction. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB. (Plan views at left – elevations at right).




Figure 3.11: Results of constrained least squares (CLS) image reconstruction. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB. (Plan views at left – elevations at right).




Figure 3.12: Results of broadband CLEAN image reconstruction. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB. (Plan views at left – elevations at right).




Figure 3.13: Results of complex-valued MaxEnt image reconstruction. (a) Full aperture data, SNR = 20 dB; (b) 10 % random sampling, SNR = 20 dB; (c) 10 % random sampling, SNR = 0 dB (Plan views at left – elevations at right).

3.5.3 Discussion

Based on the images presented in Figures 3.8 to 3.13, the following are apparent:

- Inverse filtering is indicative of the ill-posedness of inverse scattering problems. While Figures 3.8 (a) and (b) have a slightly higher resolution than that obtained by matched filtering (Figure 3.9), image stability decreases with the number of aperture samples and the SNR such that Figure 3.8 (c) is useless.
- Both lateral and range side-lobes are prominent features in images reconstructed by inverse filtering, matched filtering, MNLS and CLS algorithms. In particular, the small number of discrete frequencies has seriously degraded the range resolution achieved by these algorithms.
- The ability to detect impulsive scatterers has been seriously degraded by background clutter in Figures 3.8 to 3.11. The apparent amplitudes of point scatterers that can be discerned are much lower than their actual amplitudes.
- Background clutter and image break-up are marginally improved by MNLS and CLS regularisation, but the additional complexity of these algorithms could not be justified on the grounds of these improvements. Linear regularisation appears not to be effective when a significant proportion of the data are missing.
- Figures 3.12 (a) and (b) produced by broadband CLEAN show marked improvements over the previous images in terms of side-lobe and clutter rejection, but not in terms of resolution. Under conditions of low SNR and sparse sampling, the CLEAN algorithm reconstructs only that part of the target distribution that can be distinguished by the absolute threshold.
- The MaxEnt images in Figure 3.13 are fundamental improvements over all other images. In particular, it can be noted that the range resolution achieved by MaxEnt with only ten discrete frequencies is almost single pixel, as is the resolution in the lateral dimensions. At high SNR, sparse sampling does not significantly affect the ability of MaxEnt to accurately reconstruct both smooth

and impulsive scatterers. In all cases, the relative reflectivities of the targets appear to be correctly reconstructed. Under conditions of low SNR and sparse sampling, MaxEnt retains its high-resolution properties, detecting two out of four point scatterers while maintaining a virtually zero-amplitude background.

In addition to these observations regarding the quality of the reconstructed images, aspects such as computational complexity are also important in comparing algorithm performance. Table 3.2 summarises the execution time, in CPU seconds, required by each of the six algorithms tested. Clearly, on the basis of performance vs. complexity, the MNLS and CLS algorithms can be immediately dismissed from further attention, as their execution times were comparable to, and even exceeded those of MaxEnt, yet the resulting image quality was vastly inferior. Similarly, iterative inverse filtering can also be dismissed, although its inclusion here was mainly as a demonstration of ill-posedness in any case.

	20 dB - 100%	20 dB - 10%	0 dB - 10%
Inverse Filter	1844	3653	2746
Matched Filter	244	244	244
MNLS	4690	3534	9082
CLS	6234	4694	12714
CLEAN	1954	1433	426
MaxEnt	6573	6566	6668

Table 3.2: Algorithm execution times on a 200 MHz R10000 CPU (system seconds).

In relation to MNLS and CLS in particular, some explanation is required for the fact that their execution times exceeded those of MaxEnt, despite the former algorithms being linear, and MaxEnt being highly non-linear. Both the MNLS and CLS algorithms tested here utilised a simple outer loop search for α , which required α to be initially bracketed between zero and some upper bound. In the simulations in this chapter, a conservative upper bound was used; meaning that each algorithm

spent additional time optimising α as opposed to minimising the cost function. By lowering the upper bound for α , or improving the sophistication of its selection, both the MNLS and CLS algorithms could be made to execute in a comparable time to MaxEnt. In [2], Burch *et al.* concur that for problems involving sparse data, unlike those in conventional image restoration and/or reconstruction, the convergence time for CLS is expected to be comparable to that of MaxEnt. This is because there is no simple Fourier-domain mapping between data and image space for sparse data problems; that is, the overall matrix transformation is no longer block-circulant. In any case, efficiency alone was not the reason that these algorithms were dismissed from further consideration. Rather, they were demonstrated to be fundamentally limited in their ability to handle sparse data.

In consideration of Table 3.2 in conjunction with the reconstructed images, it can be seen that matched filtering offers a reasonable ratio of overall performance to complexity while not suffering severe stability problems. The next best algorithm in this regard is CLEAN, which speeds up considerably for sparse data and low SNR as a result of detecting fewer legitimate targets. The execution times of MaxEnt remained relatively static across the range of experimental conditions considered.

3.6 Summary

In common with other forms of image reconstruction, ultrasonic imaging possesses simultaneous and contradictory requirements for resolution preservation and smoothing. The objective of this chapter has been to explain MaxEnt as a logical extension of regularisation theory that appears to be capable of fulfilling these requirements. In particular, the relationships between MaxEnt and the powerful Bayesian framework have been noted, however MaxEnt offers tractable solutions, whereas the former appears to be constrained by the "tyranny of dimensionality".

Although it has been claimed that the positivity constraint inherent in MaxEnt is largely responsible for its capabilities, especially super-resolution [101, 163], the results of the simulations in this chapter suggest that this is not the case. In addition, the successful demonstration of MaxEnt in reconstructing both extended and impulsive objects suggests that the "near-blackness" restriction imposed by some authors [163] is an unnecessary one. Details of how to reconstruct complex-valued maximum entropy images will be given in the next Chapter.

Chapter 4

Extensions to

The Cambridge Algorithm

This chapter describes the algorithmic requirements for reconstructing complex-valued maximum entropy images from coherent broadband ultrasonic data. Due to the extreme non-linearity of the Shannon entropy, calculating maximum entropy images in any case requires a highly specialised optimisation strategy. Unlike general-purpose algorithms, the Cambridge algorithm was specifically designed to solve large maximum entropy problems efficiently and robustly, and so it is well suited to multi-dimensional image reconstruction from sparse aperture data.

As standard implementations of the Cambridge algorithm are unable to handle complex-valued images, two new methods for overcoming this limitation are discussed and the necessary modifications to the basic algorithm are detailed in each case. The chapter concludes with an evaluation of the computational complexity of the Cambridge algorithm as applied to broadband ultrasonic image reconstruction.

4.1 The Cambridge Algorithm

As noted in Chapter 3, regularised image reconstruction generally reduces to optimisation, and the specific optimisation algorithm depends on the mathematical properties of the regularising operator used. According to the notation in Chapters 2 and 3, and the definition of the number of broadband measurements M in Section 3.2.5, the MaxEnt problem can be stated in its variational form as

$$\hat{\mathbf{f}}_{MaxEnt} = \arg \min_{\mathbf{f}} [C - \alpha S] \quad \text{for } C \cong M, \quad (4.1)$$

where the Chi-squared error is given by

$$C = \chi^2 \equiv \frac{1}{2\sigma^2} \sum_{k=1}^K \left\| \left(\mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} - \mathbf{d}_k \right) \right\|^2, \quad (4.2)$$

and the entropy term is given by

$$S = - \sum_{i=1}^N \hat{f}_i \ln (\hat{f}_i / eA). \quad (4.3)$$

In view of the non-linearity of S , minimising $J = C - \alpha S$ requires considerable subtlety. First-order algorithms based purely on the gradients of C and S meet with limited success, as the gradient of S tends to dominate at low pixel amplitudes, leading to the slow development of image structure [46, 160]. True second-order methods, such as those based on Newton and quasi-Newton algorithms, require the computation and storage of $N \times N$ Hessian (second partial derivative) matrices, and are thus unsuitable for problems of realistic size [46, 70, 141, 164, 165]. Conjugate gradient algorithms, as approximations to second-order algorithms, build up information about Hessian matrices by carrying second partial derivative (curvature) information forward to successive iterations. Unfortunately, again, the entropy function is sufficiently far from quadratic to lead to poor performance [1, 46].

In [160], Cornwell proposed a simplified Newton algorithm that used a diagonal approximation to the Hessian matrix. As the effect of such an approximation equates to neglecting the side-lobes of the point-spread function (PSF), it is only acceptable in cases where the underlying PSF is dominated by a sharp central maximum. As noted in Section 3.2.4 in relation to Bayesian inverse scattering solutions, time-harmonic Green's functions have decidedly non-local characteristics, and hence a more general approach is required in ultrasonic image reconstruction.

The Cambridge algorithm, originally developed by Skilling and Bryan [51] for applications in astronomy, was selected for further development in this thesis because it achieves similar performance to second-order optimisation algorithms without the associated storage requirements. Software based on this algorithm is now marketed commercially [166], however the implementations used here were coded in C from information available in the literature [2, 46, 50, 51, 166]. The salient aspects of the Cambridge algorithm will now be reviewed to facilitate discussion of the extensions required for complex-valued image reconstruction.

4.1.1 Derivative Information

In common with second order algorithms, the Cambridge algorithm uses first and second partial derivatives to guide it to the minimum of the objective function, and a large proportion of its execution time is spent in their calculation. In broadband ultrasonic imaging (assuming real-positive pixels $\hat{f}_i$), these derivatives are given by

$$[\nabla C]_i = \frac{\partial C}{\partial \hat{f}_i} = \frac{1}{\sigma^2} \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} (\mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} - \mathbf{d}_k) \quad (4.4)$$

$$[\nabla \nabla C]_{ij} = \frac{\partial^2 C}{\partial \hat{f}_i \partial \hat{f}_j} = \frac{1}{\sigma^2} \sum_{k=1}^K [\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k]_{ij} \quad (4.5)$$

$$[\nabla S]_i = \frac{\partial S}{\partial \hat{f}_i} = \ln A - \ln \hat{f}_i \quad (i = 1, 2, \dots, N) \quad (4.6)$$

$$[\nabla \nabla S]_{ij} = \frac{\partial^2 S}{\partial \hat{f}_i \partial \hat{f}_j} = -\frac{\delta_{ij}}{\hat{f}_i} \quad (i = 1, 2, \dots, N) \quad (4.7)$$

where the first partial derivatives with respect to the image, ∇C and ∇S , are N -element vectors and the second partial derivatives, $\nabla \nabla C$ and $\nabla \nabla S$, are (formally) $N \times N$ matrices. As $\nabla \nabla S$ is diagonal, it does not present any difficulties as far as its calculation or storage; however, for reasons noted in Section 3.1.1, $\nabla \nabla C$ will not generally be semi-block circulant in the sense defined in Section 2.3.4, and thus it cannot be calculated in diagonal form via DFTs. $\nabla \nabla C$ can, however, be *multiplied* by arbitrary vectors, by sequentially applying the operators in equation (4.5), and this is how it is used in the Cambridge algorithm, as discussed shortly.

From a computational perspective, it is particularly important that Equations (4.2), (4.4) and (4.5) are handled efficiently. As noted in Section 3.1.1, the projection operator $\mathbf{P}$ should be applied in the Fourier domain prior to applying the sampling operator $\mathbf{S}$, as it is faster to calculate the DFT of a sum than it is to sum separate DFTs. As the main interest here is in processing sparse aperture data such that $\mathbf{S} \neq \mathbf{I}$, the subtraction of $\mathbf{d}_k$ and sampling by $\mathbf{S}$ require additional inverse DFTs that would not be necessary if $\mathbf{S}$ were an identity matrix.

Similarly, the form of $\mathbf{E}_k$ has a bearing on the efficiency of computing ∇C and $\nabla \nabla C$, as in particular, when the incident field is not invariant with respect to the

receiving aperture, $\mathbf{E}_k$ must be applied in the spatial domain. Assuming either perpendicular plane-wave insonification or monostatic data acquisition, the summations over k in Equations (4.4) and (4.5) can be carried out entirely in the Fourier domain and followed by a single inverse DFT.

4.1.2 Search Directions

In addition to requiring only order N storage, the Cambridge algorithm is efficient because it uses a sub-space of search directions during each iteration rather than a single search direction. The rationale for this is that by projecting quantities such as derivatives from image space into a compact subspace, more sophisticated control procedures can be employed than if the image space quantities were manipulated directly. Calculating projections between vectors via dot products is also much faster than calculating the vectors themselves, and so it is advantageous to capitalise on any vectors that have been calculated by formulating such a subspace.

As originally derived by Skilling and Bryan [46], the recommended vectors $\mathbf{e}_\mu$ ($\mu=1,2,3$) comprising the subspace of search directions can be expressed in terms of the intermediate (real-positive) image estimate $\hat{\mathbf{f}}$ as

$$\left. \begin{aligned} \mathbf{e}_1 &= \hat{\mathbf{f}}(\nabla S) \\ \mathbf{e}_2 &= \hat{\mathbf{f}}(\nabla C) \\ \mathbf{e}_3 &= \hat{\mathbf{f}}(\nabla \nabla C) \hat{\mathbf{f}}(\nabla J) \end{aligned} \right\}, \quad (4.8)$$

where the notation $\hat{\mathbf{f}}(\dots)$ implies pointwise multiplication. According to the condition $\nabla J = 0$, the trial value of α used in calculating $\mathbf{e}_3$ is given by

$$\alpha' = \sqrt{\sum [\hat{\mathbf{f}}(\nabla S)^2] / \sum [\hat{\mathbf{f}}(\nabla C)^2]}. \quad (4.9)$$

The rationale behind the various factors of $\hat{\mathbf{f}}$ in Equations (4.8) and (4.9) follows from the nature of the objective function to be minimised. As mentioned previously in relation to the failure of first-order optimisation strategies, ∇S tends to dominate ∇J at low pixel amplitudes, and consequently, the formation of high-amplitude image structure must be accelerated in some way.

To achieve this acceleration, Skilling and Bryan applied an *entropy metric* to effectively distort the computational space in which images are reconstructed. By virtue of the entropy metric, the effective distances between images become dependent on the curvature of S . In the terminology of differential geometry, the relevant *contravariant metric tensor* g^{ij} in image space is given by¹

$$g^{ij} = -[\nabla\nabla S]_{ij}^{-1}, \quad (4.10)$$

which in consideration of Equation (4.7), can be written as

$$g^{ij} = \delta_{ij} f^i, \quad (4.11)$$

where the f^i are elements of a (contravariant) image vector.

In Equations (4.8) and (4.9), therefore, pre-multiplication by $\hat{\mathbf{f}}$ can be viewed simply as a computational device that favourably distorts the *geometry* of the space to be negotiated by the optimisation algorithm, without altering the underlying problem to be solved [167]. In addition to fitting the Kullback-Leibler interpretation of entropy as a distance measure [13-15], using the entropy metric in this way also provides a natural means of establishing a "trust region" or distance limit within which the overall objective function can be approximated quadratically. This is used in controlling the convergence of the algorithm, as discussed later.

As for the particular combination of gradients and Hessian matrices that make up the three search directions in Equation (4.8), $\mathbf{e}_1$ and $\mathbf{e}_2$ are readily identified as components of the steepest descent direction ∇J , suitably pre-multiplied by $\hat{\mathbf{f}}$. The third direction $\mathbf{e}_3$ is a "look ahead" term that corresponds to the search direction that would be generated by a conjugate gradient algorithm following its first iteration [70]. In this case, it is pointless operating on ∇J with the full Hessian $\nabla\nabla J$, as $\hat{\mathbf{f}}\nabla\nabla S$ is simply an identity matrix, which does not correspond to a useful search direction [46]. There is, however, the possibility of recursively operating on $\mathbf{e}_1$, $\mathbf{e}_2$ and $\mathbf{e}_3$ with $\hat{\mathbf{f}}\nabla\nabla C$ to generate still further search directions, but only the standard three search directions in Equation (4.8) were used in this thesis.

¹ In the original literature, the Cambridge algorithm was framed in a differential geometry formalism that included the concepts of covariant and contravariant vectors. While providing useful insights into the justification for the algorithm, these concepts are not essential to this discussion.

4.1.3 Subspace Projection

As noted previously, attempting to directly optimise multi-dimensional functions of high dimensionality is complicated by the sizes of the Hessian matrices involved, even in cases where the functions are well approximated by quadratic forms. In the present case, the regularisation parameter α must also be chosen such that C attains its expected value M , according to the discrepancy principle in Section 3.2.5, and this would ordinarily require an additional "outer loop" optimisation.

To solve both the optimisation and the parameter selection problems simultaneously, each iteration of the Cambridge algorithm begins by projecting $\nabla C, \nabla \nabla C, \nabla S$ and $\nabla \nabla S$ into the subspace of search directions given by Equation (4.8). These projections are carried out via dot products according to

$$\left. \begin{aligned} S_\mu &= \mathbf{e}_\mu \cdot \nabla S & g_{\mu\nu} &= \mathbf{e}_\mu \cdot \nabla \nabla S \cdot \mathbf{e}_\nu \\ C_\mu &= \mathbf{e}_\mu \cdot \nabla C & M_{\mu\nu} &= \mathbf{e}_\mu \cdot \nabla \nabla C \cdot \mathbf{e}_\nu \end{aligned} \right\}, \quad (4.12)$$

where the S_μ and C_μ are three component gradient vectors and the $g_{\mu\nu}$ and $M_{\mu\nu}$ are 3×3 (symmetric) Hessian matrices in the subspace. Having computed the actual values of chi-squared C_0 and entropy S_0 for an intermediate image $\hat{\mathbf{f}}$, quadratic models for these quantities within a certain trust region can be formulated as a pair of uncoupled Taylor series approximations given by²

$$\left. \begin{aligned} \tilde{S}(x) &= S_0 + S_\mu x^\mu + \frac{1}{2} g_{\mu\nu} x^\mu x^\nu \\ \tilde{C}(x) &= C_0 + C_\mu x^\mu + \frac{1}{2} M_{\mu\nu} x^\mu x^\nu \end{aligned} \right\}, \quad (4.13)$$

where the x_μ ($\mu = 1, 2, 3$) are increments by which the search directions are multiplied in updating the current image estimate according to

$$\hat{\mathbf{f}}_{n+1} = \hat{\mathbf{f}}_n + \delta \hat{\mathbf{f}} \quad \text{where} \quad \delta \hat{\mathbf{f}} = \sum_{\mu=1}^3 x^\mu \mathbf{e}_\mu. \quad (4.14)$$

Maximising $\tilde{S}(x)$ while simultaneously controlling α to achieve the condition $\tilde{C}(x) \cong M$ constitutes most of the real sophistication of the Cambridge algorithm,

² To retain contact with the original literature, the summation convention [168] is also used here.

although computationally, this task is trivial in relation to calculating and projecting the search directions themselves.

4.1.4 Algorithm Control

As implemented in this work, the procedure for controlling the convergence of the Cambridge algorithm was essentially that recommended by Skilling and Bryan [46]. As there appear to be very few complete accounts of this control procedure that would allow its coding, the following discussion is included for completeness.

In the first stage of the control procedure, the nine component Hessian matrices $g_{\mu\nu}$ and $M_{\mu\nu}$ are subjected to simultaneous diagonalisation and normalisation to simplify the subsequent algebra. As part of the diagonalisation, eigenvalue filtering via a standard singular value decomposition (SVD) algorithm is used to protect against linear dependencies (rank deficiencies) in the subspace³. In the case of $g_{\mu\nu}$, the relevant decomposition in terms of the 3×3 orthonormal matrix $\mathbf{R}$ and the diagonal matrix of singular values a_μ is given by

$$g_{\mu\nu} = \mathbf{R} \text{diag}[a_1, a_2, a_3] \mathbf{R}^T. \quad (4.15)$$

By defining $\mathbf{y} = \mathbf{R}\mathbf{x}$, or conversely $\mathbf{x} = \mathbf{R}^T\mathbf{y}$, the model equations become

$$\left. \begin{aligned} \tilde{S}(y) &= S_0 + \sqrt{a_\mu^{-1}} R_{\mu\nu}^T S_\mu y_\mu + \frac{1}{2} y_\mu y_\mu \\ \tilde{C}(y) &= C_0 + \sqrt{a_\mu^{-1}} R_{\mu\nu}^T C_\mu y_\mu + \frac{1}{2} \sqrt{a_\mu^{-1}} R_{\mu\nu}^T M_{\mu\nu} R_{\mu\nu} \sqrt{a_\nu^{-1}} y_\mu y_\nu \end{aligned} \right\}, \quad (4.16)$$

where the summation convention for repeated indices [168] has again been used. Employing a second orthogonal decomposition such that

$$\sqrt{a_\mu^{-1}} R_{\mu\nu}^T M_{\mu\nu} R_{\mu\nu} \sqrt{a_\nu^{-1}} = \mathbf{V} \text{diag}[b_1, b_2, b_3] \mathbf{V}^T, \quad (4.17)$$

and defining $\mathbf{z} = \mathbf{V}\mathbf{y}$, or conversely, $\mathbf{y} = \mathbf{V}^T\mathbf{z}$, the model equations become

³ The *Numerical Recipes in C* function **dsvdcmp** [70] was used to perform these decompositions.

$$\left. \begin{aligned} \tilde{S}(z) &= S_0 + \sqrt{b_\mu^{-1}} V_{\mu\nu}^\top \sqrt{a_\mu^{-1}} R_{\mu\nu}^\top S_\mu z_\mu + \frac{1}{2} z_\mu z_\mu \\ \tilde{C}(z) &= C_0 + \sqrt{b_\mu^{-1}} V_{\mu\nu}^\top \sqrt{a_\mu^{-1}} R_{\mu\nu}^\top C_\mu z_\mu + \frac{1}{2} b_\mu z_\mu z_\mu \end{aligned} \right\}. \quad (4.18)$$

These equations are now in a form that can be manipulated to solve the optimisation and parameter selection problems simultaneously in the subspace, and hence in image space. The key to achieving this is in controlling the convergence of the algorithm on C rather than explicitly on α . Given the quadratic model for $\tilde{C}(z)$, the minimum value for Chi-squared in the subspace is given by

$$\tilde{C}_{\min} = C_0 - \frac{1}{2} b_\mu^{-1} (\sqrt{b_\mu^{-1}} V_{\mu\nu}^\top \sqrt{a_\mu^{-1}} R_{\mu\nu}^\top C_\mu)^2, \quad (4.19)$$

which is clearly not the desired result, as the regularising influence of the entropy is then not utilised. Skilling and Bryan [46] suggest instead aiming for

$$\tilde{C}_{\text{aim}} = \max\left(\frac{2}{3} \tilde{C}_{\min} + \frac{1}{3} C_0, M\right). \quad (4.20)$$

Returning to the global objective, which is that of minimising $J = C - \alpha S$ such that $C \cong M$, it is clear that this can be accomplished in the subspace by combining the two models in Equation (4.18), setting $\tilde{J}(z)$ to zero, and computing the increments

$$z_\mu = \frac{(\alpha S_\mu - C_\mu)}{(b_\mu - \alpha)} \quad (4.21)$$

in conjunction with a binary search on α . More sophisticated approaches to optimising α , such as that mentioned in Section 3.2.5 would be difficult to justify in view of the small proportion of time spent on subspace calculations.

Unfortunately, some additional complications enter the above scheme when the computed increments z_μ yield an image that is too distant from the image $\hat{\mathbf{f}}$ for the quadratic approximation to be valid. Utilising the subspace counterpart to the aforementioned entropy metric, which is the second order term in the quadratic model for $\tilde{S}$ in Equation (4.18), the recommended distance limit is [46, 51]

$$l^2 = z_\mu z_\mu \geq l_0^2 \cong 0.1 \sum_{i=1}^N f_i. \quad (4.22)$$

In the event that this distance limit is exceeded, Skilling and Bryan [46] suggest that it is better to attempt to maximise $\tilde{S}$ without further reducing $\tilde{C}$ in the current iteration, whereby $\tilde{C}_{\text{aim}}$ in equation (4.20) reverts to C_0 . If there are *no* computed increments that correspond to an image within the distance limit, an explicit distance penalty must be introduced whereby a factor P is added to the eigenvalues b_μ in Equation (4.21). The corresponding value of chi-squared in this augmented subspace is denoted C_P . By implementing a binary search on P outside the binary search on α and thus minimising the distance to the trust region, an image can always be found that satisfies the condition $\tilde{C}_{\text{aim}} \leq \tilde{C} < \tilde{C}_P \leq C_0$ [46].

Having found satisfactory increments z_μ , all that remains is to map these to the increments x_μ used in formulating the original quadratic models in Equation set (4.13). By retracing the sequence of transformations used to arrive at Equation set (4.18), the overall inverse transformation is found to be

$$x_\mu = R_{\mu\nu} \sqrt{a_\mu^{-1}} V_{\mu\nu} z_\mu. \quad (4.23)$$

Following this inverse transformation, the next image estimate can be calculated according to Equation (4.14), and the cycle of search direction calculation, subspace projection and optimisation can proceed once more. The stopping criteria for the Cambridge algorithm are based on either the approach to zero of the gradient ∇J , or a check for parallelism between its two components ∇C and ∇S . The convergence properties of the algorithm will be discussed in further detail following the introduction of extensions that allow complex-valued image reconstruction.

4.2 Complex-Valued Reconstructions

MaxEnt applications involving complex-valued, or *coherent* imagery have been far less common than those involving intensity-based optical imagery. Certainly, many distinctions separate acoustic and microwave images from optical images. Most notably, images reconstructed from coherent wave-fields generally display magnitude and phase, and hence their status as probability distributions or even as positive-additive distributions is doubtful [169].

Philosophy aside, however, the main difficulty in computing complex-valued MaxEnt solutions is that the logarithms of negative numbers are not usually meaningful in a signal processing context. As noted in the introduction, such considerations did not slow the application of MaxEnt in nuclear magnetic resonance (NMR) spectroscopy [8], where it was first realised that MaxEnt could be useful in reconstructing complex-valued signals in addition to real-positive signals

The first contributions in this area were due to Laue *et al.* [17], who demonstrated that MaxEnt could be adapted to automatically reconstruct NMR spectra containing both absorption and dispersion peaks. This was achieved by extending the definition of the Shannon entropy from the real-positive axis, first to the negative axis, and then to the positive and negative imaginary axes. In terms of the spectral amplitudes f_ω , the resulting entropy expression was

$$S = -\sum \hat{f}_\omega^+ \ln \hat{f}_\omega^+ - \sum \hat{f}_\omega^- \ln \hat{f}_\omega^- - \sum \hat{f}_\omega^i \ln \hat{f}_\omega^i - \sum \hat{f}_\omega^{-i} \ln \hat{f}_\omega^{-i}. \quad (4.24)$$

An alternative definition of the entropy of a complex-valued signal was suggested by Wright and Belton [170], who argued that the above approach posed unnecessary and unrealistic constraints on reconstructions. In particular, they pointed out that the phase of reconstructions should be *a priori* evenly distributed between 0° and 360° rather than restricted to angles of $0^\circ, 90^\circ, 180^\circ$ or 270° . To overcome the requirement for splitting complex signals into multiple channels, it was proposed to regularise complex-valued spectra on the entropy of their squared moduli. Unfortunately, no suggestions were made as to how this optimisation could be implemented in practice, and no results were presented.

Daniell and Hore [18] examined various other avenues to extending the definition of entropy into the complex plane, and suggested, on quantum-mechanical grounds, that the entropy of NMR spectra should be calculated from their complex amplitudes. Denoting these amplitudes by $\hat{z}_i$ such that

$$\hat{z}_i = \|\hat{f}_i\| = \sqrt{(x_i^2 + y_i^2)}, \quad (4.25)$$

they derived an expression for the so-called spin- $1/2$ entropy given by [8, 171]

$$S = -\sum_{i=1}^N \left[\frac{|\hat{z}_i|}{m} \ln \left(\frac{|\hat{z}_i|/m + \sqrt{4 + |\hat{z}_i|^2/m^2}}{2} \right) - \sqrt{4 + |\hat{z}_i|^2/m^2} \right], \quad (4.26)$$

where, as in Equation (3.37), m has the status of an initial estimate or model amplitude. Although the derivation of the spin- $1/2$ entropy is specialised to a specific physical interpretation of NMR, the general idea of basing entropy calculations on complex amplitudes is conceptually applicable to any complex-valued signal. A similar approach, without the quantum mechanical overtones, will be discussed later in this chapter in relation to ultrasonic image reconstruction.

Beyond NMR, the concept of one-dimensional (1-D) complex-valued signal reconstruction via MaxEnt has been applied to ultrasonic A-scan deconvolution by Smith *et al.* [172-174] with the aim of enhancing the resolution of surface profile measurements. In this work, it was concluded that MaxEnt is able to reconstruct reflectivity profiles with a useful degree of super-resolution, even when a substantial portion of the original frequency content of the data is missing. As in any form of coherent signal processing, the ability to handle complex-valued signals in this application is important given the range of acoustic boundary conditions that are possible, and in particular, the likelihood of phase inverted returns.

Aside from 1-D signal reconstruction, there have been several attempts to apply complex-valued MaxEnt to *image* reconstruction from coherent wavefield data, most notably in diffraction tomography [98, 175, 176] and inverse synthetic aperture radar (ISAR) [100, 177, 178]. The most common approach to computing the entropy of complex-valued images has been to substitute the amplitudes of complex pixels for the real-positive quantities used in conventional MaxEnt. This is essentially the prescription of Equation (4.25), however in image processing, the expression for S that is used is simply the Shannon entropy of Equation (3.37).

Unfortunately, although the motivation for adapting MaxEnt to coherent wavefield image reconstruction has been strong, there have been few examples of its anticipated superior performance [179], particularly under realistic experimental conditions. The main reason for this has been the lack of simple, yet powerful MaxEnt algorithms that are easily adapted to complex-valued image reconstruction. In most cases, authors have based their approaches on either second-order optimisation algorithms such as Newton-Raphson [176-178], or conjugate-gradient algorithms [175]; neither of which is well suited to large-scale MaxEnt calculations. In this thesis, two separate methods were used to adapt the Cambridge algorithm to complex-valued image reconstruction. Referred to as the “composite image” method and the “complex metric” method, each of these will now be discussed.

4.2.1 The Composite Image Approach

In [180] (see also Appendix D for complete details) Battle *et al.* adapted the Cambridge algorithm to ultrasonic image reconstruction via an image model similar in motivation to that proposed by Bajkova [177], and Frieden and Bajkova [178, 181]. In the composite image approach, the complex image $\hat{\mathbf{f}}$ is replaced by an ensemble of four sub-images $\hat{\mathbf{f}}_q$, ($q = 1,2,3,4$); each phase shifted to occupy a prescribed phase quadrant while maintaining positivity of its real and imaginary parts. The forward data model of Equation (3.1) thus becomes

$$\mathbf{d}_k = \mathbf{S} \left[\mathbf{P} \mathbf{H}_k \mathbf{E}_k \sum_{q=1}^4 \hat{\mathbf{f}}_q \exp\left(\frac{jq\pi}{2}\right) + \mathbf{n} \right], \quad (4.27)$$

which leads to a revised form for the chi-squared error given by

$$C = \frac{1}{2\sigma^2} \sum_{k=1}^K \left\| \left[\mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \sum_{q=1}^4 \hat{\mathbf{f}}_q \exp\left(\frac{jq\pi}{2}\right) - \mathbf{d}_k \right] \right\|^2 \quad (4.28)$$

Use of the composite image model requires that the search directions in Equation (4.8) be split into $\mathbf{e}_{\mu q}$ where $\mu = 1,2,3$ and $q = 1,2,3,4$. The calculations for $\mathbf{e}_1$ and $\mathbf{e}_2$ proceed as before for each of the (real-positive) sub-images, however the multiplications by $\nabla \nabla C$ in the calculation of $\mathbf{e}_3$ are complicated by cross terms in the second partial derivatives of Equation (4.28) such that...

$$\begin{bmatrix} \nabla \nabla C_{\mu 1} \\ \nabla \nabla C_{\mu 2} \\ \nabla \nabla C_{\mu 3} \\ \nabla \nabla C_{\mu 4} \end{bmatrix} = \frac{1}{\sigma^2} \begin{bmatrix} 1 & \exp(\frac{j\pi}{2}) & -1 & \exp(-\frac{j\pi}{2}) \\ \exp(-\frac{j\pi}{2}) & 1 & \exp(\frac{j\pi}{2}) & -1 \\ -1 & \exp(-\frac{j\pi}{2}) & 1 & \exp(\frac{j\pi}{2}) \\ \exp(\frac{j\pi}{2}) & -1 & \exp(-\frac{j\pi}{2}) & 1 \end{bmatrix} \begin{bmatrix} \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \mathbf{e}_{\mu 1} \\ \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \mathbf{e}_{\mu 2} \\ \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \mathbf{e}_{\mu 3} \\ \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \mathbf{S} \mathbf{P} \mathbf{H}_k \mathbf{E}_k \mathbf{e}_{\mu 4} \end{bmatrix} \quad \dots (4.29)$$

When calculating the entropy, the summation in Equation (4.3) is extended over the four real-positive sub-images to act as the regularising functional for the complex-valued image as a whole. Using the composite image model does increase the complexity of calculating the subspace basis vectors, however in most respects,

the Cambridge algorithm behaves as reported by its original authors [46]. On satisfying the stopping criteria, the reconstructed image is obtained through addition of the sub-images and calculating its magnitude from its real and imaginary parts.

In the early stages of the research undertaken for this thesis, the composite image algorithm was used extensively, and even extended to the reconstruction of three-dimensional (3-D) images from broadband data. Not only did the resulting images display super-resolution, but it was demonstrated that their quality did not seriously degrade in the absence of large proportions of the aperture data.

Following the publication of [180], the basic scattering model was generalised to that described in Chapter 2 to permit greater flexibility. Particularly because of the use of multiple frequencies, the composite image algorithm then became overly complicated, and an alternative approach was developed, based on a similar treatment to that used in [8]. Referred to here as the "complex metric" approach, this approach is more closely allied to the idea of replacing the real-positive pixels in conventional MaxEnt with the magnitudes of complex-valued pixels.

4.2.2 The Complex Metric Approach

The complex metric approach is so-named because the only substantial alteration required to the standard Cambridge algorithm relates to the entropy metric discussed in Section 4.1.2. In this method, the image is treated as a distribution of magnitudes and the entropy is computed irrespective of phase. Of course, in computing the chi-squared error term according to Equation (4.2), the phase retains its importance.

For the purposes of explanation, complex-valued matrices and vectors will be considered in terms of their equivalent real representations; that is, matrix elements will be viewed as having the submatrix form

$$c = a + j b = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}, \quad (4.30)$$

and image vector elements will be expressed as

$$f_i = x + j y = \begin{pmatrix} x \\ y \end{pmatrix}. \quad (4.31)$$

By Equation (4.25), the corresponding generalisation of the Shannon entropy is

$$S = -\sum_{i=1}^N \hat{z}_i \ln(\hat{z}_i / eA). \quad (4.32)$$

By Equation (4.10), the entropy metric was seen to equate to the negative inverse of the second partial derivative of the entropy. For purely real images, this results in the entropy metric (interpreted as a matrix) being strictly diagonal, and hence the factors of $\hat{f}_i$ in Equations (4.8) and (4.9). As a starting point to deriving an equivalent metric for complex-valued images, each element of the first partial derivative of S can be calculated as

$$\nabla S(z_i) = \frac{\partial S(z_i)}{\partial \hat{x}_i} + j \frac{\partial S(z_i)}{\partial \hat{y}_i} = \frac{x_i \ln\left(\frac{z_i}{A}\right)}{z_i} + j \frac{y_i \ln\left(\frac{z_i}{A}\right)}{z_i}. \quad (4.33)$$

By calculating the mixed partial derivatives of $\nabla S(z_i)$ with respect to both x_i and y_i , the Hessian for S assumes a block diagonal form, with each block given by

$$\nabla \nabla S(z_i) = \begin{bmatrix} \frac{\partial^2 S(z_i)}{\partial x_i^2} & \frac{\partial^2 S(z_i)}{\partial x_i \partial y_i} \\ \frac{\partial^2 S(z_i)}{\partial y_i \partial x_i} & \frac{\partial^2 S(z_i)}{\partial y_i^2} \end{bmatrix} = \begin{bmatrix} \frac{-x_i^2 - y_i^2 \ln\left(\frac{z_i}{A}\right)}{z_i^3} & \frac{x_i y_i \left(1 - \ln\left(\frac{z_i}{A}\right)\right)}{z_i^3} \\ \frac{x_i y_i \left(1 - \ln\left(\frac{z_i}{A}\right)\right)}{z_i^3} & \frac{-y_i^2 - x_i^2 \ln\left(\frac{z_i}{A}\right)}{z_i^3} \end{bmatrix}. \quad (4.34)$$

As before, the entropy metric is found by taking the negative inverse of $\nabla \nabla S$. As $\nabla \nabla S$ is block-diagonal, $-\left[\nabla \nabla S\right]^{-1}$ is also block-diagonal, with each block given by⁴

$$-\left[\nabla \nabla S(z_i)\right]^{-1} = \begin{bmatrix} \frac{\left(x_i^2 + \frac{y_i^2}{\ln\left(\frac{z_i}{A}\right)}\right)}{z_i} & \frac{x_i y_i \left(1 - \frac{1}{\ln\left(\frac{z_i}{A}\right)}\right)}{z_i} \\ \frac{x_i y_i \left(1 - \frac{1}{\ln\left(\frac{z_i}{A}\right)}\right)}{z_i} & \frac{\left(y_i^2 + \frac{x_i^2}{\ln\left(\frac{z_i}{A}\right)}\right)}{z_i} \end{bmatrix} \quad (4.35)$$

⁴ This matrix inverse was computed via the *Mathematica* 3.0 system [182].

Unfortunately, Equation (4.35) is not immediately useful as a metric, as the inverse logarithmic terms diverge as z_i approaches A . This is unacceptable, as it is the specific aim of MaxEnt to drive image amplitudes to this value in the absence of information to the contrary. As this pathology is only associated with pixel amplitudes z_i near A , which in the present application can be assumed to be a small background amplitude, it is acceptable to augment the logarithmic terms in Equation (4.35) with a stabilising term such that the metric matrix $\mathbf{M}$ becomes

$$\mathbf{M} = \text{Ident}[N] \otimes \begin{bmatrix} \frac{x_i^2 + \frac{y_i^2}{\ln\left(\frac{z_i + A}{A}\right)}}{z_i} & \frac{x_i y_i \left(1 - \frac{1}{\ln\left(\frac{z_i + A}{A}\right)}\right)}{z_i} \\ \frac{x_i y_i \left(1 - \frac{1}{\ln\left(\frac{z_i + A}{A}\right)}\right)}{z_i} & \frac{y_i^2 + \frac{x_i^2}{\ln\left(\frac{z_i + A}{A}\right)}}{z_i} \end{bmatrix} \quad (4.36)$$

where the $\otimes$ symbol denotes the Kronecker product [58]. Using this metric, the set of search directions previously established for real-positive images becomes

$$\left. \begin{aligned} \mathbf{e}_1 &= \mathbf{M} \nabla S \\ \mathbf{e}_2 &= \mathbf{M} \nabla C \\ \mathbf{e}_3 &= \mathbf{M} \nabla \nabla C \mathbf{M} \nabla J \end{aligned} \right\}, \quad (4.37)$$

where the trial value of α used in calculating $\mathbf{e}_3$ is now given by


$$\alpha' = \sqrt{\sum [\mathbf{M}(\nabla S)^2] / \sum [\mathbf{M}(\nabla C)^2]}. \quad (4.38)$$

These modifications to the Cambridge algorithm subsume much of the complexity of the composite image approach discussed in Section 4.2.1 and are all that is required to apply MaxEnt to complex-valued image reconstruction without further restriction.


4.3 Algorithm Convergence

As noted in Section 4.1.4, the stopping criteria for the Cambridge algorithm are based on either the approach to zero of the composite gradient ∇J , or a check for


parallelism between its two vector components. Figure 4.1 illustrates the typical convergence behaviour of the complex-metric variant of the Cambridge algorithm in reconstructing a 3-D ultrasonic image from ten frequencies at an SNR of 20 dB.



(a)



(b)



(c)

Figure 4.1: Typical convergence of the complex-metric variant of the complex-valued MaxEnt algorithm (a) Chi-squared $\div M$; (b) Entropy; (c) Angular error (SNR = 20 dB).

In particular, it can be seen that the initial rates at which chi-squared and the entropy decrease are quite slow, and this has been observed to be related to the value of the background parameter A . With large values of A (perhaps an order of magnitude less than the average image amplitude), convergence is much faster, however in such cases the super-resolution and side-lobe reduction effects of MaxEnt are less apparent. As the images presented in subsequent chapters are of objects in homogeneous media with very low background scattering amplitudes, small values of A were appropriate. A value of $A = 1 \times 10^{-6}$ relative to the average image amplitude, as used in Figure (4.1), was found to give a good compromise between image quality and rate of convergence.

Following the initial decrease in entropy that is always associated with the departure of the image from its initial flat estimate, the mode of convergence of the Cambridge algorithm changes. In the second phase of convergence, chi-squared is maintained at its expected value of M while the entropy increases to its global maximum, as evidenced by the reduction in the angle between ∇C and ∇S . Unless otherwise stated, all of the images presented in subsequent chapters achieved an angular error of less than 1° . In reality, images do not change perceptibly below approximately 10° in the final stages of convergence.

Generally, it has been observed that the failure of the Cambridge algorithm to achieve its chi-squared objective within fifty iterations is indicative of either modelling errors or an overly optimistic SNR estimate. This can be seen as a useful feature in that the algorithm is less likely to reconstruct meaningless images if provided with incorrect parameters or data. In ultrasonic inverse scattering problems, there is also the possibility of multiple scattering effects, and as these cannot be easily modelled, SNR estimates should be especially conservative.

4.4 Computational Complexity

To gauge the computational complexity of MaxEnt as applied to ultrasonic image reconstruction, the number of floating point operations (FLOPs) typically required for algorithm convergence was estimated as a function of image size⁵. To permit comparisons with conventional Fourier-domain beamforming, the number of FLOPs needed to implement matched filtering was also estimated.

⁵ Calculations associated with incident field and Green's function operators were not included in these estimates.

In Section 4.1.1, it was noted that fewer inverse DFTs are required under certain conditions of insonification, and for this reason, two versions of each algorithm were evaluated. In version A of each algorithm, the incident field was assumed to be invariant with respect to the receiving aperture. This corresponds to plane-wave insonification in the plane of the aperture, or bistatic data acquisition, whereby $\mathbf{E}_k$ and $\mathbf{H}_k$ can be combined. In either case, only a single inverse DFT is required in Equations (3.14), (4.2), (4.4) and (4.5). In version B, arbitrary insonification was assumed such that each discrete frequency required a separate inverse DFT.

Figure 4.2 illustrates the $n \times n \times d$ image for which the computational complexities were estimated, with n being the lateral image dimension and d being the number of depth levels. Assuming a complexity of $5n \log_2 n$ for the basic 1-D Cooley-Tukey FFT algorithm [183], the number of FLOPs *per iteration* of MaxEnt algorithms A and B respectively were estimated to be approximately

$$O_A \cong 363dn^2 + 272Kdn^2 + 200dn^2 \log_2[n] + 200Kn^2 \log_2[n] \quad (4.39)$$

and

$$O_B \cong 399dn^2 + 284Kdn^2 + 200Kdn^2 \log_2[n] + 200Kn^2 \log_2[n] \quad (4.40)$$

Figure 4.3 indicates the total number of FLOPs estimated for algorithms A and B to reconstruct a cubical image with $d = n$ pixels on a side from ten frequencies with typical convergence in 50 iterations. The number of FLOPs required to implement matched filtering under the same conditions is also indicated.




Figure 4.2: 3-D image consisting of $n \times n \times d$ volume elements (voxels).




Figure 4.3: The total number of FLOPs needed to implement 3-D MaxEnt via the Cambridge algorithm compared with matched filtering. Algorithm A in each case corresponds to an arbitrary incident field, whereas algorithm B corresponds to either plane-wave insonification parallel to the aperture or monostatic data acquisition.

In view of Figure 4.3, it is apparent that MaxEnt requires approximately three orders of magnitude more computation than matched filtering to reconstruct a three-dimensional image of the same size from the same number of discrete frequencies. In interpreting this, however, it must be remembered that matched filtering requires a wider temporal bandwidth, and hence a larger number of frequencies than MaxEnt to obtain acceptable range resolution, as shown by the simulations in Chapter 3. In relation to the overall complexity of either algorithm, the additional operations required to account for arbitrary incident fields are significant, and should be avoided if possible by employing plane wave insonification or monostatic data acquisition.

4.5 Summary

In this chapter, the general theory behind the Cambridge MaxEnt algorithm was presented, and two specific adaptations of the algorithm to complex-valued imagery were discussed. Although some of the advantages of MaxEnt have already been illustrated through comparative simulations, these simulations in themselves have not proved the suitability of MaxEnt for reconstructing images from actual ultrasonic data, and nor have they adequately demonstrated its full potential. The next two chapters of this thesis are concerned with experimental validation of MaxEnt under conditions that more closely parallel those of practical imaging applications.

Chapter 5

MaxEnt Evaluation in Synthetic Aperture Imaging

In previous chapters, ultrasonic image reconstruction was formulated as a linearised inverse-scattering problem, to which various linear and non-linear approaches to image reconstruction were shown to be inadequate solutions, particularly in cases involving incomplete data. Although the simulations in Chapter 3 suggested that MaxEnt is capable of producing images of superior resolution and clarity, simulations alone do not constitute proof of these capabilities.

The aim of this chapter is to report the results of synthetic aperture imaging experiments in which images of two and three-dimensional scatterers were reconstructed from real ultrasonic data. Two monostatic data-acquisition geometries commonly applied in NDE were investigated, particularly in relation to the effects of sparse sampling. In addition to MaxEnt reconstructions obtained via the complex-metric variant of the Cambridge algorithm described in Chapter 4, matched filter and CLEAN reconstructions are also presented for comparison.

5.1 Synthetic Aperture Imaging

As noted in Chapter 1, synthetic aperture imaging offers a means of obtaining focussed images that would otherwise require the use of inordinately large focussed transducers, or complex arrays of fixed transducers. By transmitting and receiving pulsed ultrasonic fields via a single omni-directional element, or an approximation thereof, sufficient information can be obtained to reconstruct images of scatterers within the region of interest (ROI). In [180], (see also Appendix D) Battle *et al.*

presented 2-D MaxEnt images reconstructed from time-harmonic monostatic data. The results in this chapter constitute an extension of this work to broadband data.

5.1.1 Equipment and Infrastructure

Figure 5.1 illustrates the arrangement of the major items of equipment used in the experiments described in this chapter, including the computers used for data acquisition, image reconstruction and visualisation. According to the usual practice in monostatic synthetic aperture imaging [80, 180, 184], focussed transducers were used to approximate line and point sources, which were then scanned in one and two dimensions to acquire broadband data. A standard immersion-scanning tank was used, with all controlling software and device drivers either written from scratch by the author or ported from other operating systems. The Linux™ operating system facilitated easy communication with high-performance computing resources in addition to enabling code compatibility across the three systems shown in Figure 5.1.




Figure 5.1: Arrangement of major equipment for synthetic aperture experiments.

A Krautkramer USIP-12 flaw detector was used as the pulser/receiver. The USIP-12 is a laboratory class instrument, which due to its extensive range of front panel controls and rear panel connections is ideally suited to experimental work. In this series of experiments, the pulse amplitude was set to 150 volts and the damping impedance was set to $75\ \Omega$. The centre frequency of the broadband filter was set according to the frequency of the transducer used in each experiment.

The single radio frequency (RF) output from the rear of the USIP-12 was connected via 3 m of $50\ \Omega$ coaxial cable to a Gage CS225 analogue-to-digital converter (ADC) card, installed in an Intel 486 personal computer. A connection was also made from the trigger output at the rear of the USIP-12 to the trigger input of the CS225 in order to synchronise A-scan capture with pulse transmission.

In its single channel configuration, the CS225 card supported a wide range of sampling rates up to 50 megasamples per second (MSPS). In none of the experiments reported here was it strictly necessary to sample data at this rate. However, it was observed that the trigger uncertainty (jitter) of the CS225 (specified at 10 ps rms) was higher than expected at low sample rates, which contributed noise to the sampled wavefields. In order to minimise the effect of this jitter, the maximum sampling rate of 50 MSPS was always used.

The scanning tank had a minimum stepping increment of 0.12 mm and stepper-motor pulses were supplied by a Parker-Hannifin PC-23 controller card. Although this arrangement provided adequate indexing control, the PC-23 had no facility for generating CPU interrupts during continuous scanning. To trigger A-scan capture, therefore, a separate counter/timer card was used to generate interrupts by dividing the stepper-motor pulse train by a predetermined value.

5.1.2 Pre-processing

The most important aspect of pre-processing following data acquisition was the registration and Fourier transformation of the acquired A-scans. Following data acquisition, the required number of Fourier coefficients from each A-scan was automatically selected near the peak of the transducer pulse-echo frequency response. In this work, the Fourier coefficients were always selected consecutively, although other selections (random, for example) could also have been made.

As the signals originating from within the ROI were offset from the aperture plane, the Fourier coefficients of the acquired A-scans had to be referred *back* to the aperture plane in order to allow image reconstruction at the correct range. For a range offset of R (m) between the aperture and ROI (see Section 2.1.1), the *medium delay* T_m is given by

$$T_m = \frac{2R}{c_m} (\text{s}), \quad (5.1)$$

where c_m is the appropriate (longitudinal or transverse) velocity in the medium, and the factor of two accounts for two-way propagation.

In the opposite sense to T_m , there is generally also a delay between triggering the transmitter and launching an incident wave into the medium. The effect of this delay is to move the time datum *forward* from the instant of transducer excitation, to the instant of *physical* transmission. In monostatic imaging, where focussed transducers are commonly used, this *coupling delay* T_c is given by

$$T_c = \frac{2F}{c_c} (\text{s}), \quad (5.2)$$

where F is the focal length of the transducer, and c_c is the propagation velocity in the coupling medium. Subtracting the coupling delay from the medium delay, the overall phase correction to be applied to the FFT coefficients is given by

$$\phi_k = 2\pi f_k (T_m - T_c) \text{ (radians)}. \quad (5.3)$$

Also important in relation to phase compensation is the temporal spectrum of the transmitted waveform, as one of the aims of image reconstruction is to synthesise as close an approximation to a delta function response as possible. In this series of experiments, the spectral characteristics of the transmitted pulse were incorporated directly into the Green's function operators, and for this reason, the phase parameters of the transmitted waveform were not required during pre-processing. In any case, it should be noted that the image reconstruction algorithms tested in this chapter were not sensitive to the *absolute phases* of the Fourier coefficients of the data. The *relative* phases (and amplitudes) of these Fourier coefficients, however, were important, and had to be incorporated into the observation model.

5.2 Two-Dimensional Imaging

Due to the NDE emphasis of this thesis and the current importance of 2-D ultrasonic imaging in industry (TOFD imaging for example [185]), 2-D image reconstruction was investigated as a prelude to 3-D image reconstruction. MaxEnt was previously investigated in this context by Burch [186], but as no allowance for the complex-valued nature of the images was made, the 2-D problem is worth revisiting.

5.2.1 Data Acquisition

In the 2-D synthetic aperture imaging experiments, a Krautkramer-Branson immersion transducer with a nominal peak frequency of 4.0 MHz was utilised. This transducer was fitted with a plastic cylindrical lens giving a focal length of approximately 47-mm in H₂O, and a 6-dB beam angle of approximately 30° in aluminium. The targets to be imaged consisted of six side-drilled holes¹ in a 20-mm thick aluminium plate with a longitudinal velocity of 6367 ms⁻¹ (see Figure 5.2). Data were acquired by scanning the focal point of the transducer along the top of the plate, according to the 2-D imaging geometry described in Section 2.3.1.




Figure 5.2: Experimental configuration for 2-D synthetic aperture imaging.

Within the Al plate, the region of interest (ROI) to be reconstructed covered an area of 61.44 mm $\times$ 30.72 mm. Each of the 256 $\times$ 128 pixels measured 0.24 mm² so that each scan increment corresponded to two stepper-motor increments. To gauge

¹ With respect to the scattering theory discussed in Chapter 2, the holes were modelled as Dirichlet boundaries.

the effects of uniform under-sampling on the quality of the reconstructed images, various multiples of this increment were also tested, as will be discussed in Section 5.2.2. Figure 5.3 illustrates the raw B-scan data acquired from the block, which consisted of 256 A-scans arranged vertically. Each A-scan was $9.65 \mu\text{s}$ in duration.






Figure 5.3: B-scan acquired during 2-D synthetic aperture imaging.

In addition to the main echoes from the drilled holes, considerable secondary structure is evident in Figure 5.3. Part of the reason for this is undoubtedly the non-ideal response of the transducer-lens combination, which is illustrated in Figure 5.4 (a). In particular, the echo approximately $1.5 \mu\text{s}$ after the main pulse was believed to be due to lens reverberation. Clearly, not all of the secondary structure in Figure 5.3 can be attributed to the characteristics of the pulse-echo response, however, and further probable causes for this phenomenon will be discussed in Section 5.2.2.


Figure 5.4 (b) is an amplitude plot of the FFT spectrum of (a). The vertical bars indicate the frequency range selected for image reconstruction. As in the simulations in Chapter 3, ten consecutive FFT coefficients were utilised, in this case ranging from 3.22 MHz to 4.20 MHz. Figure 5.3 (c) illustrates the phase spectrum of the pulse-echo response from which the linear component was removed by translating the peak of the analytic envelope back to $t = 0$ prior to FFT processing [187].



(a)



(b)




(c)


Figure 5.4: Measured time and frequency domain responses for the cylindrically focussed transducer. (a) Pulse-echo waveform; (b) Amplitude spectrum; (c) Phase.

5.2.2 Results and Discussion


Figure 5.5 illustrates the results of image reconstruction from the data in Figure 5.3 via matched filtering (effectively conventional SAFT), broadband CLEAN and complex-valued MaxEnt respectively. Based on Figure 5.3, the peak SNR was initially estimated at 12 dB. The colour scale in all of the following images is the same as that used in Figure 3.5 (red = maximum reflectivity, black = background).



(a)



(b)



(c)


Figure 5.5: Images reconstructed from the B-scan data in Figure 5.3 for an estimated SNR of 12 dB. (a) Matched filtered image; (b) CLEAN image; (c) MaxEnt image.

All of the images in Figure 5.5 display non-ideal properties to varying degrees. In all cases, six holes are evident, although additional artefacts, characteristic of the respective algorithms used are also present. In Figure 5.5 (a), the limited temporal bandwidth has blurred the holes to approximately three times their actual diameter, and the amplitudes of three of the holes appear disproportionately low. Coupled with this, the relatively high-amplitude side-lobe clutter adjacent to each hole is characteristic of images reconstructed via matched filtering.


In Figure 5.5 (b), broadband CLEAN is seen to have reconstructed five false targets before reaching its absolute threshold corresponding to the estimated SNR of 12 dB. The elliptical distribution of these targets suggests multiple scattering, whereby the sum of distances from points on the circumference of the ellipse to its two foci is constant, as are the delays incurred by multiple reflections between the holes. Other second-order effects such as creeping waves [22] would also go some way to explaining the secondary structure apparent in Figure 5.3.

In Figure 5.5 (c), MaxEnt also appears to have reconstructed false targets, and in precisely the same locations as those reconstructed by CLEAN. It can be noted, however, that the amplitudes of these targets are somewhat lower than those in Figure 5.5 (b), and they are accompanied by other low-level background clutter. In fact, Figure 5.5 (c) is *not* a maximum entropy image in the sense defined in Chapter 4, as on this occasion, the Cambridge algorithm failed to achieve its chi-squared objective, much less maximise the entropy. This is one occasion where MaxEnt has warned of its failure in view of irreconcilable modelling errors. To achieve the chi-squared objective, additional image structure would have had to be introduced at the expense of decreasing the image entropy.

Thankfully, the presence of modelling errors does not preclude the possibility of obtaining better images via CLEAN or MaxEnt if it is considered that such errors can be accounted for to some degree by simply reducing the SNR estimate. Unfortunately, however, the appropriate effective SNR is then difficult to estimate from the data alone, and it is likely that several images would have to be reconstructed for different SNR estimates in order to select an optimal value. To evaluate this approach, the SNR estimate for the data in Figure 5.3 was decreased from 12 dB to 6 dB, and the CLEAN and MaxEnt algorithms were run again. Figure 5.6 (a) and (b) are the results of this second attempt at image reconstruction (the matched filtered image was independent of the SNR and hence remained the same).



(a)



(b)

Figure 5.6: Images reconstructed according to the same parameters in Figure 5.5, except for a reduced SNR estimate of 6 dB. (a) CLEAN image; (b) MaxEnt image.

In Figure 5.6, many of the problems previously believed to be associated with modelling errors have disappeared. The CLEAN algorithm detected six targets correctly before being stopped by its absolute target threshold. MaxEnt also converged successfully within fifty iterations, achieving no less resolution than previously, but significantly fewer artefacts (a single, faint artefact is visible between the two uppermost holes). The execution times for matched filtering, CLEAN and MaxEnt were 16, 24 and 383 seconds respectively.

Visually, it could be argued that the CLEAN image in Figure 5.6 is superior to the MaxEnt image, perhaps because the holes in the former look "round", whereas in the latter, they seem over-resolved laterally and stretched vertically. In fact, these impressions are superficial, as Figure 5.6 (a) contains nothing but six impulses, with no possibility of any


deeper interpretation. In Figure 5.6 (b), the fact that the "holes" appear narrow has more to do with their scattering behaviour and the limited temporal bandwidth than any deficiencies on the part of MaxEnt (recall that Kirchhoff theory assumes normally orientated facets, and only the top of each hole would fit this description). For this reason, it should be kept in mind that CLEAN is limited to reconstructing images as collections of impulses, which then become blurred together to form continuous structures. MaxEnt, on the other hand, is inherently a "free-form" solution, and as such is capable of superior performance.

To evaluate the effects of under-sampling on images reconstructed by the above three algorithms, two additional sets of images were reconstructed, again for an estimated SNR of 6 dB. At 4.20 MHz, the longitudinal wavelength in the aluminium was 1.52 mm, thus the original sample spacing of 0.24 mm represented a significant over-sampling relative to the 0.38 mm ($\lambda/4$) monostatic sampling interval (see Section 2.3.3). In the first under-sampling experiment, the results of which are illustrated in Figure 5.7, only every eighth A-scan was retained. This corresponded to an effective sampling interval of 1.92 mm and an under-sampling factor of 500 %.


In view of the substantial grating lobes that have manifested, Figure 5.7 (a) has obviously suffered severe aliasing. While all six holes are still visible, it would be difficult to discriminate between the aliased images and the true images, as both have similar amplitudes. In contrast, the CLEAN and MaxEnt images show no evidence of aliasing. CLEAN detected six holes correctly and stopped. MaxEnt also reconstructed six holes, with the only artefact being a slight smudge on the second hole from the right (possibly due to creeping wave generation). In this case, the execution times for matched filtering, CLEAN and MaxEnt were 16, 24 and 355 seconds respectively on a 200 MHz R10000 CPU.

Noteworthy in Figure 5.7 (b) is the amplitude uniformity of the reconstructed holes. In particular, in comparison with Figure 5.6 (a), it appears that the centre-left hole was better reconstructed from the sparsely sampled data than it was from the full aperture data. In other trials (reported here only in summary) it was concluded that the amplitudes of distinct targets reconstructed by CLEAN were subject to substantial variations, depending on the degree of undersampling and the SNR. In the present case, where the assumed SNR was 6dB, this variation was ± 10 % for sampling intervals up to 1000 % of the critical (Nyquist) interval.


In the second under-sampling experiment, the sample interval was increased to 3.84 mm, with only 16 A-scans retained. The margin of under-sampling was thus 1000 %. In this case, the degree of aliasing in the matched filtered image was worse, and CLEAN completely misplaced one of the holes. Remarkably, however, MaxEnt still reconstructed the six holes in their correct positions, with minimal artefacts. In addition, the resolution of MaxEnt appears not to have suffered substantially from sparse sampling. The execution times for the three algorithms in this case were 16, 22 and 301 seconds respectively.



(a)




(b)




(c)


Figure 5.7: Images reconstructed from 32 A-scan samples at a spacing of 1.92 mm for an estimated SNR of 6 dB (a) Matched filtering; (b) CLEAN; (c) MaxEnt.



(a)



(b)



(c)

Figure 5.8: Images reconstructed from 16 A-scan samples at a spacing of 3.84 mm for an estimated SNR of 6 dB (a) Matched filtering; (b) CLEAN; (c) MaxEnt.

5.3 Three-Dimensional Imaging

The 2-D experiments reported in the previous section showed that MaxEnt is substantially more robust against under-sampling than either matched filtering or CLEAN. In this section, the more challenging problem of 3-D image reconstruction is investigated to further exemplify the advantages of MaxEnt over simpler approaches, particularly in relation to sparse aperture performance and resolution.

5.3.1 Data Acquisition

In most respects, the data acquisition procedure for 3-D imaging was similar to that described for 2-D imaging, except that the transducer was scanned in a 2-D raster as opposed to a 1-D line. Due to the difficulty of arranging for "interesting" 3-D scatterers in a solid medium, data acquisition was performed entirely in immersion with a test target assembled from 2 mm diameter copper wire as illustrated in Figure 5.9². The ROI was a $64 \times 64 \times 64$ cubical grid with a uniform spacing of 1.2 mm and a range offset R from the aperture plane of 180 mm (see Figure 2.2).





Figure 5.9: Test target used for 3-D synthetic aperture imaging experiments.


The transducer used was a KB-Aerotech 1 MHz (nominal) immersion unit, spherically focussed at 25.4 mm. Figure 5.10 (a) illustrates the pulse-echo response of this transducer in the time domain, as measured from the top surface of a 100-mm thick Al block positioned at the focal distance and normal to the beam axis. Figure

² With respect to the scattering theory discussed in Chapter 2, the wires were modelled as Neumann boundaries.


5.10 (b) is the FFT spectrum of the time-domain response. The vertical bars indicate the range of frequencies actually used in image reconstruction. As in all previous experiments, ten discrete frequencies near the spectral peak were selected, in this case spanning the range from 0.88 MHz to 0.93 MHz.



(a)



(b)



(c)

Figure 5.10: Measured time and frequency domain responses for the spherically focussed transducer. (a) Pulse-echo response; (b) Amplitude spectrum; (c) Phase.

At 0.93 MHz, the wavelength of sound in H_2O is approximately 1.6 mm, giving a monostatic sampling increment ($\lambda/4$) of 0.4 mm; substantially smaller than the 1.2 mm increment used here. With reference to Figure 1.4 and Equation (1.2), however, it can be shown that aliasing will not occur given the above sampling parameters, provided the target is further than 100 mm from the aperture. Using as coarse a sampling grid as possible is advantageous, as it maximises the size of the ROI while lessening the demand for CPU time and memory. In addition, the above combination of sampling parameters was previously used successfully in [180], and it was desired to retain continuity with this work for comparative reasons.

Figure 5.11, illustrates the real part of the time-harmonic wavefield sampled with the above parameters at the highest frequency considered. This corresponds to a 100 % sampled or "filled" aperture, which is one of the sampling conditions investigated in the present series of experiments. As the main interest here is in reconstructing images from sparse data, however, various other sampling patterns, ranging from 25 % sparse, down to 3 % sparse were also investigated, and these patterns have been illustrated in Figures 5.12 (a) to (f). Based on the visual appearance of the data in Figure 5.11, the peak SNR was estimated at 20 dB.




Figure 5.11: The real part of the full aperture data at 0.93 MHz (area = 76.8 mm^2).




Figure 5.12: Aperture sampling patterns, (a) Real part of full aperture data at 0.93 MHz. (b) 25 % Square; (c) 6 % Random; (d) 6 % Uniform; (e) 6 % Star; (f) 3 % Cross.

5.3.2 Results and Discussion

Synthetic aperture images of the test target in Figure 5.9 were computed via matched filtering, coherent CLEAN and MaxEnt after re-sampling the original aperture data with the sampling patterns in Figure 5.12. These images are presented in Figures 5.13 to 5.18 respectively. As was done previously with the simulation results in Chapter 3, each 3-D image is given in plan and elevation.

The images in Figure 5.13 were computed from the full aperture data. In particular, it can be noted that the range resolution of the matched filtered image is poor, with each wire segment blurred to many times its actual vertical extent. This is a direct result of restricting the temporal bandwidth of the original data by utilising only ten Fourier coefficients of the scattered field. Both CLEAN and MaxEnt have fared considerably better than matched filtering in terms of resolution, although the CLEAN image shows signs of its underlying nature as a collection of blurred impulses. There also appear to be several low amplitude targets between the two uppermost wire segments in the CLEAN image, possibly associated with multiple scattering. Thus, it is possible that improvements could have been obtained by lowering the SNR estimate, as was done previously in Section 5.2.2. Of the three algorithms tested with full aperture data, MaxEnt yielded the best image, with Figure 5.13 (c) being very close to the visual appearance of the test target.

The points made above in relation to the full aperture reconstructions have been amplified in Figure 5.14 by restricting the aperture to 25 % of its original size. As expected, the matched filtered image has suffered a considerable loss of resolution, with the bottom wire segment barely resolved. The CLEAN image too, has been degraded by the aperture restriction, as in the full aperture reconstruction the wire segments appeared continuous (despite being dominated by impulses), whereas now only the impulses remain. In this regard, the MaxEnt image is superior to the CLEAN image, as the wires have retained their visual continuity.

The images in Figure 5.15 were reconstructed after randomly re-sampling 6 % of the original aperture data according to the pattern in Figure 5.12 (c). As noted in Chapter 1, random sampling achieves most of the benefits of large apertures without the grating lobes typically associated with uniform under-sampling. Certainly, the matched filtered image in Figure 5.15 (a) has retained the original resolution of the

full aperture image, however the level of background clutter due to the irregular side-lobes of the sampling pattern has become unacceptably high.

In comparing the CLEAN and MaxEnt images in Figure 5.15, it is apparent that CLEAN has managed, in this case, to preserve some of the visual continuity of the wire segments, although the image is still dominated by high-amplitude impulses. The MaxEnt image, in contrast, has suffered only minor degradation, with the test target remaining as sharply resolved as it was in the full-aperture case.

The images in Figure 5.16 were reconstructed after re-sampling the original aperture data according to the 6 % uniform grid of Figure 5.12 (d). These images are possibly the most startling in this chapter, as they demonstrate the dramatic improvements that can be obtained via non-linear as opposed to linear image reconstruction. As noted in Chapter 1, and in the previous paragraph, uniform under-sampling would usually be associated with grating lobes, which serve to replicate the true image at regular intervals throughout the field of view. This is precisely what has happened in the matched filtered image in Figure 5.16 (a). Remarkably though, neither the CLEAN, nor MaxEnt images contain such replicas. In fact, Figures 5.16 (b) and (c) are comparable to their respective full aperture images, although the MaxEnt image is again slightly superior on account of its smoothness.

The above comments also apply to the images in Figure 5.17, which were reconstructed from aperture data sampled according to the star pattern in Figure 5.12 (e). As with uniform under-sampling, the side-lobes of this sampling pattern are very strong due to the concentration of the samples in four principal directions. Predictably, the matched filtered solution is seriously distorted, whereas CLEAN has reconstructed a collection of impulses that barely qualifies as an image. In keeping with the character of maximum entropy images generally, Figure 5.17 (c) displays an absolute minimum of structure, and in doing so, remains visually smooth whilst retaining most of the resolution of the full aperture image.

In the final set of images in Figure 5.18, only 3 % of the aperture data were retained according to the cross pattern in 5.12 (f). At this level of sparsity, which probably represents the smallest amount of information from which any reasonably intelligible image can be reconstructed, neither matched filtering nor CLEAN was able to produce a meaningful image. The MaxEnt image in Figure 5.18 (c), on the other hand, still indicates the presence of four distinct, although not necessarily identifiable structures.







Figure 5.13: Reconstructed images of the 2 mm diameter wire test target in immersion from 100 % of the aperture data taken at 10 discrete frequencies. (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).



(a)




(b)




(c)


Figure 5.14: Reconstructed images of the 2 mm diameter wire target in immersion from 25 % of the aperture data sampled uniformly according to Figure 5.12 (b). (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).



(a)



(b)



(c)

Figure 5.15: Reconstructed images of the 2 mm diameter wire target in immersion from 6 % of the aperture data sampled randomly according to Figure 5.12 (c). (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).




Figure 5.16: Reconstructed images of the 2 mm diameter wire target in immersion from 6 % of the aperture data sampled according to the uniform grid in Figure 5.12 (d). (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).







Figure 5.17: Reconstructed images of the 2 mm diameter wire target in immersion from 6 % of the aperture data sampled according to the star pattern Figure 5.12 (e). (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).



(a)



(b)



(c)

Figure 5.18: Reconstructed images of the 2 mm diameter wire target in immersion from 3 % of the aperture data sampled according to the cross pattern Figure 5.12 (f). (a) Matched filtering; (b) CLEAN; (c) MaxEnt; (Plan views at left – elevations at right, 1.2 mm per pixel).

To conclude this comparative evaluation of the results of 3-D synthetic aperture image reconstruction, Table 5.1 summarises the execution times of the three algorithms tested for each of the aperture sampling patterns in Figure 5.12. From this table, it can be seen that MaxEnt was approximately twenty-five times slower than matched filtering, and approximately seven times slower than CLEAN. As previously noted in Chapter 3, the execution times of CLEAN were strongly dependent on the specific image to be reconstructed and the quality of the data. In this instance, MaxEnt was also seen to display some dependence on these factors.

	(a) Full Aperture	(b) 25 % Full	(c) 6 % Random	(d) 6 % Uniform	(e) 6 % Star	(f) 3 % Cross
Matched Filter	245	245	245	245	245	245
CLEAN	1063	1088	1032	937	1031	804
MaxEnt	7800	7600	6785	7753	5704	3855

Table 5.1: 3-D reconstruction times in seconds with a 200 MHz R10000 CPU.

5.4 A Note on Aliasing

The following comments are made in relation to the phenomenon of aliasing as it pertains to ultrasonic image reconstruction. Conventional wisdom has it that a band-limited function can be reconstructed without loss of information provided it has been sampled at a sufficiently high rate, as given by the Whittaker-Shannon (or Nyquist) sampling theorem [29, 57], and not otherwise. Certainly, in relation to the classical direction-of-arrival (DOA) estimation problem discussed in Chapter 1, which equates image reconstruction to spectral estimation, this wisdom holds true. In view of the limitations imposed by the sampling theorem, however, some of the images presented in this chapter may appear to contradict important axioms. How could images such as 5.16 (c) possibly have been reconstructed from aliased data?

In fact, from this work, and earlier work [180], there appear to be at least two avenues by which the ambiguity of aliasing can be resolved. Firstly, important distinctions must be drawn between the types of *imaging* problems that have been considered in this thesis, and the classical DOA problem. In the far field, time-harmonic wavefields radiated or scattered by an object are almost plane, so that image reconstruction from far field data can be interpreted as spectral estimation. In the near field, however, or in any region in which the wavefield intercepted by the aperture displays *curvature*, point-spread functions possess spatially localised structure, and the underlying problem passes from being one of spectral estimation to being one of image reconstruction. As discussed in Chapter 3, the Bayesian framework provides powerful techniques for handling uncertainty, and aliasing is simply a manifestation of uncertainty. Treating image reconstruction as spectral estimation, therefore, ignores cogent prior information, such as the PSF, and this is one way of explaining the apparent contradiction of the sampling theorem.

Secondly, it is well known that images reconstructed from sparsely sampled broadband data are less susceptible to side-lobe clutter than narrow-band images [25, 28, 40]. This is a result of side-lobes at neighbouring frequencies cancelling, and is an entirely linear phenomenon. With non-linear algorithms, such as MaxEnt, however, the cancellation effect can be exploited to an even greater degree, as broadband images are no longer restricted to being simply summations of narrow-band images. Moreover, after taking into account the specific frequency dependence of the scattering mechanism involved, there is only *one* image of the scatterer, and this is independent of frequency.

Whereas the action of CLEAN is readily understood in terms of picking bright targets and removing their associated side-lobes, the action of MaxEnt is more subtle, which explains its special abilities in solving ill-posed inverse problems. Specifically, the “kangaroo” or self-consistency property of MaxEnt works to reconstruct an image that fits the data at all frequencies, but rejects image structure that is not correlated across the frequency range. That is, by invoking a measure of configurational information, such as entropy, image features that are peculiar to a single frequency, including side-lobes and grating lobes, will tend to be suppressed. This is another reason why the sampling theorem does not necessarily apply to broadband image reconstruction. Forming a linear sum of narrow-band images is

entirely different to processing multiple frequencies simultaneously via a non-linear algorithm.

Lastly, it should be noted that although image reconstruction from aliased data is possible via MaxEnt and other algorithms, it is essential that aliasing is avoided during intermediate calculations involving Green's functions and the scattering model. This requires that the spacing of the m -dimensional computational grid be suitably chosen so that the sampling theorem applies in at least the first $m-1$ dimensions (this aspect of the grid specification was discussed in Section 2.3.2).

5.5 Summary

In this chapter, MaxEnt was critically evaluated under realistic experimental conditions by comparing its image reconstruction performance with that of competing algorithms. In addition to being highly dependent on temporal bandwidth to achieve acceptable range resolution, matched filtering was seen to pose major limitations to the quality of images reconstructed from sparse data. Although CLEAN is simpler than MaxEnt and displays similar properties of robustness under some conditions, the images produced are of a fundamentally different character to that generally desired, particularly in the case of extended scatterers.

For every combination of scan geometry and sampling pattern considered in this chapter, MaxEnt resulted in superior images. In cases where the scattering model did not adequately reflect the observed physics, the Cambridge algorithm stalled, thus providing unambiguous warning of its failure. In most cases, however, given a suitably realistic SNR estimate, MaxEnt reconstructed high-resolution images while properly preserving the smooth structure of extended objects. Aside from its added complexity, therefore, MaxEnt was superior to the other algorithms tested.

Chapter 6

MaxEnt Application to

Orthogonal Array Imaging

In Chapter 5, synthetic aperture imaging experiments confirmed that MaxEnt was superior to both matched filtering and coherent CLEAN as an algorithm for ultrasonic image reconstruction, particularly in its handling of sparse data. For a variety of reasons, as noted in Chapter 1, however, synthetic aperture imaging using mechanically scanned transducers is not always practical, and transducer arrays remain the favoured mode of data acquisition in many applications.

In this chapter, the previously developed principles of complex-valued MaxEnt are applied to array-based 3-D imaging. The approach taken here differs from conventional approaches to array imaging in many respects. In particular, it was desired to overcome the difficulties of fabricating and controlling the large numbers of small transducer elements required, even in sparse 2-D arrays.

The first steps to achieving this were in approximating plane-wave insonification and focussing entirely in reception. In this regard, the use of MaxEnt is highly advantageous, as its super-resolution properties compensate for the resolution loss that normally accompanies unfocussed transmission. Secondly, to avoid making connections to millimetre-sized array elements with high-impedances, a novel 2-D array structure known as an *orthogonal* or *edge-connected* array was used. As opposed to the *spatially* sparse data encountered in previous chapters, the data acquired by an orthogonal array is sparse in the *spatial-frequency* domain. Again, MaxEnt is eminently suitable as an image reconstruction algorithm, as it is able to extrapolate the measured spatial-frequencies into unmeasured regions.

6.1 2-D Array Technology

As noted by Plummer *et al.* [188], "Against the background of present efforts to secure high-resolution real-time ultrasonic imaging with extended viewing fields and scanning ability, the development of capable area arrays emerges as a basic requirement." In view of this statement, the current level of research interest in 2-D ultrasonic arrays is not surprising. What is, perhaps, more surprising is that there have been relatively few innovations in the basic philosophy of 2-D data acquisition, which is central to the problem of 3-D ultrasonic imaging.

The reasons for this are clearly related to the technological precedents set by linear transducer arrays, which have been enormously successful in medical imaging. 3-D ultrasonic imaging requires data acquisition over 2-D apertures, and the most natural way of achieving this is by extending existing linear array technology to a second dimension. As discussed in Chapter 1, however, there are many difficulties associated with this approach. Chief amongst these are the prodigious numbers of array elements required [188-190], the small size of each element [191-193] and the side-lobe clutter associated with sparse sampling patterns [40, 194, 195].

While several approaches to fabricating and controlling 2-D array structures have been developed [188-190, 192, 193, 196, 197], most of these share the same heritage of classical beamforming applied to point-sampled array data. This is because free-space Green's functions, which constitute the basic solutions to the inhomogeneous Helmholtz equation, are associated with geometric point sources in m dimensions, as discussed in Section 2.1.4. Mathematically of course, there are alternatives to sampling 2-D functions via delta functions, and in an image-processing context, any linear transfer function is acceptable, provided it can be included in the underlying observation model. In the next section, the particular transfer function displayed by orthogonal array structures will be investigated.

6.2 The Orthogonal Array

Whereas an $n \times n$ array, in its fully populated state, requires n^2 connections to address each element separately, an orthogonal array addresses the same number of elements via only $2n$ electrical connections along its edges. The ultrasonic

orthogonal array illustrated in Figure 6.1 consists of a single sheet of pressure sensitive (piezoelectric) material between sets of row and column electrodes.




Figure 6.1: Electrode configuration of a 2-D orthogonal array.

6.2.1 Array Transfer Function

Schau [198] investigated the properties of $n \times n$ orthogonal or "edge-connected" arrays in relation to acoustic beamforming. The problem of approximating n^2 independent 2-D potential differences m_{ij} , by $2n$ row and column voltages, r_i and c_j respectively, was stated in the form of the matrix equation

$$\mathbf{W} = \mathbf{A}\mathbf{V}, \quad (6.1)$$

in which the vector

$$\mathbf{V} = [r_1, r_2, \dots, r_n, c_1, c_2, \dots, c_n]^T, \quad (6.2)$$

was related to the vector

$$\mathbf{W} = [m_{11}, m_{12}, \dots, m_{nn}]^T, \quad (6.3)$$

by the $n^2 \times 2n$ connection matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \dots & 0 & -1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \dots & 0 & 0 & 0 & \dots & -1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -1 & 0 & \dots & 0 \\ 0 & 0 & \dots & 1 & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & 0 & 0 & \dots & -1 \end{bmatrix}. \quad (6.4)$$

To solve Equation (6.1) in a least-squares sense, a pseudo-inverse matrix

$$\mathbf{A}^\# = [\mathbf{A}^\top \mathbf{A}]^{-1} \mathbf{A}^\top \quad (6.5)$$

was introduced such that

$$\hat{\mathbf{V}} = \mathbf{A}^\# \mathbf{W}. \quad (6.6)$$

Implicit in Equation (6.6) and in the definition of $\mathbf{A}^\#$ is the requirement for an overall normalisation for the row and column voltages r_i and c_j , due to the fact that the elements of $\mathbf{W}$ are differential voltages. This normalisation adds an additional constraint to Equation (6.1), which is necessary to make $[\mathbf{A}^\top \mathbf{A}]$ non-singular. Henceforth, the normalisation assumed here is such that

$$\sum_{i=1}^n r_i + \sum_{j=1}^n c_j = 0, \quad (6.7)$$

which leads to the pseudo-inverse operator [198]

$$\mathbf{A}^\# = \left[\frac{1}{n} \mathbf{A}^\top + \frac{1}{2n^2} \mathbf{G} \right], \quad (6.8)$$

where $\mathbf{G}$ is given by the $2n \times n^2$ matrix

$$\mathbf{G} = \begin{bmatrix} -1 & -1 & \cdots & -1 \\ -1 & -1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & -1 \\ \hline 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{bmatrix}. \quad (6.9)$$

By analysing the eigenvectors of the matrix $\mathbf{A}\mathbf{A}^\#$, the cyclic impulse response of the general orthogonal array can be examined [198]. Moreover, by applying $\mathbf{A}^\#$ to any image vector consisting of a lexicographically ordered unit impulse, the shift-invariant 2-D impulse response can be found directly. One such typical impulse response for a 4×4 orthogonal array is shown in Figure 6.2. Clearly, isolated impulses cannot be reproduced perfectly due to the pattern of electrode connections, which results in electrical coupling between elements throughout the array.




Figure 6.2: The 2-D impulse response for a 4×4 orthogonal array, with corresponding numerical values for the differential voltages $W_{ij} = r_i - c_j$.

In the Fourier domain, the above impulse response takes a particularly simple form. In Figure 6.3, it can be seen that only the DFT coefficients on the $f_x = 0$ and $f_y = 0$ axes have been retained in the orthogonal array data. These coefficients correspond to spatial frequencies propagating in the column and row directions respectively. Setting the off-axis DFT coefficients of an arbitrary 2-D wavefield to zero, therefore, is a more direct means of modelling the transfer function of an orthogonal array than employing Equation (6.6) and subsequently calculating $W_{ij} = r_i - c_j$. By virtue of the same Fourier-slice theorem discussed in Section 2.2.4, the DFT coefficients of the orthogonal array data can also be calculated by taking 1-D DFTs of the original field values, averaged in the row and column directions¹. This latter approach will be discussed further in relation to the observation model for orthogonal array imaging in Section 6.2.3.




Figure 6.3: The DFT amplitudes corresponding to the impulse response in Figure 6.2.

¹ Averaging a discrete set of values is equivalent to calculating the 0th or DC term of its associated Fourier series.

Regardless of how they are modelled, orthogonal array transfer functions obviously introduce severe distortion into 2-D wavefields. In practice, therefore, orthogonal arrays must present some substantial advantages over conventional point-sampled arrays if their use is to be justified. Aside from the fact that only $2n$ connections are required, the major advantage of orthogonal arrays in relation to ultrasonic imaging is their low electrical impedance, as discussed next.

6.2.2 Array Impedance

One of the major problems in designing 2-D arrays is the small size of the individual elements, which have correspondingly small capacitances and thus high impedances [189, 193, 199]. Even when special precautions are taken in the design of pre-amplifiers and cabling, the SNR of the signals received by high-impedance array elements can be seriously degraded [200]. For an $n \times n$ orthogonal array, therefore, it is of interest to calculate the source impedance seen by a differential amplifier placed across an arbitrary pair of row and column electrodes.

The following analysis considers only row-column electrode capacitances and neglects the effects of stray inter-electrode and electrode-screen coupling capacitances, although these could also have been included if a more accurate analysis had been desired. Figure 6.4 depicts the capacitance model for an $n \times n$ orthogonal array, where row n has been arbitrarily assigned as a reference node and a hypothetical current I_{cn} has been injected into column n .




Figure 6.4: Inter-electrode capacitance model for an $n \times n$ orthogonal array.

If C is the capacitance of each of the n^2 2-D array elements, and $Y = 1/Z = j\omega C$ is the associated admittance, then the node-voltage matrix equation for the orthogonal array connection takes the form

$$\begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ I_{Cn} \end{bmatrix} = \begin{bmatrix} nY & 0 & \cdots & 0 & -Y & -Y & -Y & \cdots & -Y \\ 0 & nY & \ddots & \vdots & -Y & -Y & -Y & \cdots & -Y \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & nY & -Y & -Y & -Y & \cdots & -Y \\ \hline -Y & -Y & \cdots & -Y & nY & 0 & 0 & \cdots & 0 \\ -Y & -Y & \cdots & -Y & 0 & nY & 0 & \cdots & 0 \\ -Y & -Y & \cdots & -Y & 0 & 0 & nY & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \ddots & 0 \\ -Y & -Y & \cdots & -Y & 0 & 0 & \cdots & 0 & nY \end{bmatrix} \begin{bmatrix} V_{R1} \\ V_{R2} \\ \vdots \\ V_{Rn-1} \\ \hline V_{C1} \\ V_{C2} \\ V_{C3} \\ \vdots \\ V_{Cn} \end{bmatrix} \quad (6.10)$$

By solving Equation (6.10) for V_{Cn} , the expression for the differential (row-column) impedance Z_D as a function of the single element impedance Z is found to be²

$$Z_D = \frac{(2n-1)Z}{n^2} \cong \frac{2Z}{n} \quad (n \gg 1). \quad (6.11)$$

The exact variation of Z_D/Z as a function of n has been plotted in Figure 6.5.




Figure 6.5: Variation in row-column impedance with orthogonal array size (n).

² This expression was computed via the *Mathematica* 3.0 system [182].

From Figure 6.5, it is evident that the differential impedances of an orthogonal array decrease monotonically with n . This provides a natural advantage over point-sampled arrays in respect to driving connecting cables and pre-amplifiers, which load the output of the array elements. Pre-amplifiers are a major source of noise in ultrasonic imaging systems, and signal attenuation through capacitive loading results in a reduced input signal amplitude and hence a lower SNR. In addition to the noise voltage contributed directly by pre-amplifiers, noise voltage is generated by the passage of amplifier noise current through the array element and cable impedances, and by the resistances of the array elements themselves [193, 201]. Both of these latter noise components decrease with decreasing element impedance, and thus vary inversely with the number of interconnected elements in an orthogonal array.

6.2.3 Observation Model

Now that the salient aspects of orthogonal array theory have been introduced, the aim is to incorporate the transfer function of the orthogonal array into the matrix observation model developed in Section 2.3.4. The only aspect of this model that needs to be altered is the sampling operator $\mathbf{S}$, which was previously defined in the spatial domain via Equation (2.48). As seen in Section 6.2.1, the wavefield data acquired via an orthogonal array is sparse, however in this case, the sparseness is in the Fourier domain rather than the spatial domain. To represent this, the Fourier domain sampling operator $\mathbf{T}$ is introduced such that

$$\mathbf{T} = \text{diag}[t_i] \quad (i = 1, 2, \dots, A) \quad (6.12)$$

where $t_i = 1 + j0$ for each of the 2-D Fourier coefficients on the $f_x = 0$ and $f_y = 0$ axes, and $t_i = 0 + j0$ otherwise. A is the total number of elements (sampled and unsampled) in the array. Rather than to completely replace the spatial sampling operator $\mathbf{S}$, the intention here is to augment the matrix observation model to allow both spatial *and* Fourier domain sampling concurrently via multiple orthogonal array cells. Such “hybrid sampling” is essential, as the data obtained from a single orthogonal array cell would otherwise only permit the reconstruction of separate orthogonal views of a scattering object, rather than a true 3-D image as desired³.

³ This situation is similar to that in tomography when reconstructing an image from two orthogonal projections.

By using orthogonal array unit cells, and having L such cells distributed throughout a 2-D aperture, the combined spatial and Fourier domain diversity in the acquired broadband data permits 3-D images to be reconstructed. Figure 6.6 illustrates how this scheme was implemented in the prototype array that was constructed to test the concept of MaxEnt image reconstruction from orthogonal array data. In this instance, four orthogonal cells were used ($L = 4$), which allowed electrical connections to be made entirely around the edges of the array.




Figure 6.6: Implementation of hybrid sampling in the prototype orthogonal array.

In practice, the extension of hybrid aperture sampling to larger numbers of orthogonal array cells would not significantly increase the computational complexity of image reconstruction, as the pre-processing associated with each cell involves only 1-D Fourier transformations of its row and column data, suitably zero-padded to the overall dimensions of the aperture⁴. Zero-padding is necessary in order to express the total aperture as a linear sum of L sub-apertures. As with purely spatial sampling, there is no requirement for the aperture to be completely sampled.

⁴ If one of the $2n$ connections from each orthogonal array cell was used as a reference, then the row and column voltages could be Fourier transformed directly. In this work, differential signals were acquired from each array cell, and then averaged or “projected” as illustrated in Figure 6.6 to improve the overall SNR of the data.

Considering the l^{th} orthogonal array cell at the k^{th} discrete frequency, the Fourier-domain data, in terms of the matrix operators defined in Chapter 2, and the operator $\mathbf{T}$ defined above, is given by

$$D_{kl} = \mathbf{T} \mathbf{F} \mathbf{S}_l (\mathbf{P} \mathbf{H}_k \mathbf{E}_k \mathbf{f} + \mathbf{n}_k), \quad (6.13)$$

where $\mathbf{F}$ is a 2-D DFT matrix, and $\mathbf{S}_l$ is the sampling operator associated with the l^{th} cell. Although Equation (6.13) is concise, $\mathbf{F}$ is only required to facilitate Fourier domain sampling. As discussed in relation to Figure 6.6, efficient implementation of the operators $\mathbf{T}$ and $\mathbf{F}$ involves only 1-D DFTs of the projected cell data.

To adapt the Cambridge algorithm to orthogonal array imaging, all that is required is to find an expression for the Chi-squared error and differentiate this with respect to the image to find its first and second partial derivatives. Making the same simplifications as in Section 3.1.1 in relation to the binary-diagonal sampling operators; that is, $\mathbf{S}_l^{*T} \mathbf{S}_l = \mathbf{S}_l \mathbf{S}_l = \mathbf{S}_l$ and $\mathbf{T}_l^{*T} \mathbf{T}_l = \mathbf{T}_l \mathbf{T}_l = \mathbf{T}_l$, the chi-squared error and its first and second partial derivatives are given by

$$C = \chi^2 = \frac{1}{2\sigma^2} \sum_{k=1}^K \left\| \sum_{l=1}^L \mathbf{T} (\mathbf{F} \mathbf{S}_l \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} - D_{kl}) \right\|^2 \quad (6.14)$$

$$[\nabla C]_i = \frac{\partial C}{\partial \hat{f}_i} = \frac{1}{\sigma^2} \sum_{k=1}^K \mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \sum_{l=1}^L \mathbf{S}_l \mathbf{F}^{-1} \mathbf{T} (\mathbf{F} \mathbf{S}_l \mathbf{P} \mathbf{H}_k \mathbf{E}_k \hat{\mathbf{f}} - D_{kl}) \quad (6.15)$$

$$[\nabla \nabla C]_{ij} = \frac{\partial^2 C}{\partial \hat{f}_i \partial \hat{f}_j} = \frac{1}{\sigma^2} \sum_{k=1}^K [\mathbf{E}_k^* \mathbf{H}_k^{*T} \mathbf{P}^T \sum_{l=1}^L \mathbf{S}_l \mathbf{F}^{-1} \mathbf{T} \mathbf{F} \mathbf{S}_l \mathbf{P} \mathbf{H}_k \mathbf{E}_k]_{ij} \quad (6.16)$$

When implementing the discrepancy principle to control the convergence of the Cambridge algorithm, the expected value of chi-squared must be properly defined relative to the noise variance according to the particular DFT normalisation in force. In addition, the effective number of measurements M must reflect the fact that the row and column data from individual array cells have been zero-padded to the overall aperture size prior to Fourier transformation. For a general $n_1 \times n_2$ aperture consisting of L orthogonally connected cells, M is given by

$$M = (n_1 + n_2 - 1) L K, \quad (6.17)$$

which is independent of the number of rows and columns in each cell.

6.3 The Prototype Imaging System

This section is dedicated to describing the design and construction of a prototype 3-D ultrasonic imaging system based on the orthogonal array principles previously described in this chapter, and the MaxEnt signal processing concepts described in previous chapters. The construction and testing of this system represented the culmination of the research undertaken for this thesis, which first sought improved methods for image reconstruction from sparse ultrasonic data, and then identified MaxEnt as exceptionally well suited to the task. In constructing the following hardware, it was not expected that a working imaging system would result, but rather it was hoped that some of the important parameters of a future bistatic imaging system suited to NDE applications might be identified. Nevertheless, the imaging system to be described possessed capabilities that far exceeded expectations.

6.3.1 Array Construction

To simplify construction, the receiving array was based on standard 110 μm thick polyvinylidene fluoride (PVDF) sheet with silver ink metallisation, manufactured by AMP Piezo-Sensors Inc. (Harrisburg USA) [202]. Due to the low cost of PVDF sheet, several researchers have experimented with simplified array construction techniques, utilising adhesive tapes as bonding agents [203] and printed circuit boards as substrates [204, 205]. The same techniques were exploited in this work.

Due to the manual construction techniques used, the choice of PVDF film thickness for the receiver array was influenced by a requirement to minimise the operating frequency in order to achieve manageable array element dimensions. 110 μm PVDF film was the thickest (lowest frequency) flexible film available from AMP, and although its $\lambda/2$ resonant frequency was higher than desired, its frequency response has been found to extend well below 1 MHz in practice [203, 206]. Following the same arguments made in Section 5.3.1 in relation to the permissible sample spacing, the major parameters of the receiving array were specified as follows:

Array element pitch	-	1.5 mm
Number of elements	-	32×32
Nominal frequency	-	1.0 MHz

For the reasons noted at the beginning of this chapter, the mode of data acquisition was bistatic. In H_2O , the above parameters would not result in under-sampling provided the ROI was offset from the aperture by at least 80 mm. To reduce the level of electrical interference between the transmitter and the receiver circuitry, a separate 50.8 mm² thick-film PVDF tile (also manufactured by AMP [202]) was used as a transmitting element. This tile was bonded to the opposite side of the 1.75-mm printed circuit board (PCB) substrate with Loctite 480 epoxy, so that it effectively transmitted *through* the receiver array. In an effort to minimise transmission losses, a recess was milled out of the rear of the PCB to reduce its thickness to 0.5 mm. This recess was then sputter coated with a 0.1- μ m layer of gold to form a ground connection for the transmitter tile, which was clad on both sides with 250 μ m of aluminium. At 500 μ m in thickness, the $\lambda/4$ (rigidly backed) resonant frequency of the transmitter tile was approximately 1 MHz, matching the specified frequency of the system. Figure 6.7 is a cross-section of the assembled array transducer illustrating its various adhesive, insulating and shielding layers.






Figure 6.7: Cross-section of the prototype bistatic array transducer.

To define the orthogonal electrode pattern on the receiver array, the silver ink metallisation on the 110 μ m thick PVDF film was deeply scored with a scalpel at precise 1.5 mm increments, as illustrated in Figure 6.8 (a). Following bonding of the array film to the PCB via 3M 75 μ m double sided acrylic tape, 50 μ m tinned wires

were attached to the electrodes via Tracon Tra-duct 2902 room temperature curing silver-loaded epoxy, as illustrated in Figure 6.8 (b).



(a)




(b)

Figure 6.8: Magnified views of the orthogonal array film and connections. (a) Front electrode pattern; (b) Conductive epoxy connections to the array film. Note the lines associated with scoring the orthogonal electrode pattern on the reverse side of the film.


As mentioned earlier, the receiving array consisted of four separate 16×16 orthogonal cells, which formed quadrants of the 48 mm^2 array aperture, as seen in Figure 6.9 (a). This required 128 connections to the array film (32 per edge). In order to access the rows and columns of the four cells, half of the connections were made to the front of the film, and half were made to the rear. Two 3-mm slots were milled through the PCB either side of the transmitting tile to gain access to the latter, as shown in Figure 6.9 (b). With trial and error, the most successful means of making the array connections was found to be pre-soldering the wires to the PCB, bonding the array film, and then applying the conductive epoxy with a pinhead. After complete curing of the epoxy, the resistances of the array connections were all of the order of a few hundred milliohms. This was comparable to the resistances of the silver ink electrodes themselves, measured along their length.

As seen in Figure 6.9 (a), the PCB included four copper pads, each of which was connected to the ground plane behind the array film in addition to one of the four 34 pin headers that served as cable connectors. After applying two coats of water-resistant polyurethane lacquer (Radio-Spares number 569-290) over the entire face of the array, each ground pad was soldered to remove a small area of lacquer. To assist in protecting the array connections, the soldered ends of the header pins were covered with masking tape prior to applying the lacquer. Having sealed the array and its connections, a $0.1 \mu\text{m}$ layer of gold was sputtered over its front face, as shown in Figure 6.10 (a). This layer made contact with the ground pads and thus shielded the array film electrostatically. Following sputtering, the array was subjected to an additional two coats of polyurethane lacquer as waterproofing.

From the perspective of obtaining a broadband ultrasonic response from the PVDF transmitter and receiver, providing a backing for the transducer was not strictly necessary [203, 207]. In this instance, however, early prototype testing revealed severe reverberation of both transmitted and received acoustic waves from the front and back faces of the transducer. To redirect and attenuate these reverberations, an araldite backing-block was cast over the rear of the transmitting tile, angled at approximately 30° to the aperture plane. This had precisely the desired effect and resulted in a clean pulse-echo response with little evidence of reverberation (see Section 6.4.1). Figure 6.10 (b) is a view of the completed array transducer in which the epoxy backing, 34-pin headers and coaxial transmitter cable are all visible.




(a)




(b)

Figure 6.9: (a) A front view of the orthogonal array showing inter-quadrant isolation (horizontal line) and the ground plane pads; (b) Rear view showing the transmitter tile.



(a)



(b)

Figure 6.10: The completed array transducer. (a) A view of the front face of the array showing the sputtered gold screen; (b) The epoxy backing layer and transmitter cable.

Prior to sputtering the front screen onto the array, the row-column capacitances, measured with a Hewlett-Packard 4192 impedance analyser, averaged approximately 16 pF. Assuming an inter-electrode gap of 0.1 mm and a permittivity of $11.5 \epsilon_0$ for PVDF [207], this capacitance agreed reasonably well with the value of 15 pF calculated according to Equation (6.11). Upon completion of the array, the average row-column capacitance was approximately 28 pF due to coupling with the screen.

6.3.2 Interface Electronics

The array interface electronics, which consisted of a 128-channel multiplexer and high-impedance differential pre-amplifier, were constructed on a single PCB. The multiplexer, shown schematically in Figure 6.11, was based on low-voltage analogue CMOS technology. Each of the four quadrants of the array was sampled via two MAX396 16-channel switch arrays arranged in row-column format, with their common outputs fed to a MAX384 4-channel differential multiplexer. Using this arrangement, a 10-bit binary address supplied to the interface board was sufficient to select the active quadrant, row and column of the array. Given the shunt capacitances of the MAX396 and MAX384 (specified by the manufacturer to be 90 pF and 34 pF respectively at 1.0 MHz), the capacitance seen by the differential preamplifier, in addition to that of the array, was approximately 248 pF.




Figure 6.11: Schematic diagram of the 128 channel orthogonal array multiplexer.

The differential pre-amplifier was based on the standard instrumentation configuration [201], with some additional bootstrapping of the input bias resistances to increase its input impedance. OPA604 FET-input operational amplifiers were selected for the inputs due to their low noise ($e_n = 10 \text{ nV}/\sqrt{\text{Hz}}$, $i_n = 7 \text{ fA}/\sqrt{\text{Hz}}$ @ 1 MHz) and large gain-bandwidth product (20MHz). The output stage consisted of a MAX437 low noise bipolar operational amplifier. Given the finite open loop gains of both devices, the overall differential gain of the pre-amplifier was approximately 43 (33 dB) at 1 MHz, and the 3dB bandwidth was approximately 4 MHz.






Figure 6.12: Differential preamplifier used with the prototype orthogonal array.

In view of the many high-impedance signal paths routed between the multiplexer and pre-amplifier, one entire side of the interface board was used as a ground plane, as shown in Figure 6.13 (a). The component layout on the PCB was largely determined by the geometry of the array connections, which took the form of four 34-pin headers reverse soldered onto the PCB, as shown in Figure 6.13 (b). The ± 12 volt power supply for the interface electronics (regulated to ± 5 volts on the board) was derived from the same PC30-B I/O card used to generate interrupts in the synthetic aperture experiments in Chapter 5. In this case, the primary function of the PC30-B was to generate array addresses during data acquisition.



(a)



(b)

Figure 6.13: (a) The component side of the interface board showing the multiplexer chips surrounding the pre-amplifier; (b) A view of the solder side and array connectors.

6.3.3 System Overview

Figure 6.14 shows the assembled "front end" of the prototype orthogonal array imaging system. The array transducer and the interface board were connected via four 100-mm lengths of 34 way ribbon cable terminated in female IDC connectors. For the same reasons noted in Chapter 5 in relation to the difficulty of procuring interesting scatterers in solid materials, only immersion imaging experiments were conducted with the array. To facilitate this, a 60-mm² window was cut in the bottom of a die-cast metal box, in front of which the array was secured by four screws and sealed around its periphery with vacuum grease. In addition to waterproofing the array, the metal box served as electromagnetic shielding for the ribbon cables.




Figure 6.14: A view of the assembled imaging system showing the 4×34 way ribbon cables connecting the PVDF array to the multiplexer/pre-amplifier board and the separate coaxial transmitter cable in the foreground (waterproof enclosure not shown).

Figure 6.15 shows the general arrangement of equipment that comprised the prototype orthogonal array imaging system. In the synthetic aperture experiments reported in Chapter 5, transmitted pulses were generated by the USIP-12 flaw detector. In the case of the array transducer, however, it was found that significantly higher pulse-echo amplitudes could be obtained by using an ENI LA-325 RF power amplifier (Gain = 50 dB, 250 kHz – 150 MHz) in conjunction with a Tektronix AWG 2020 arbitrary waveform generator. In addition to enabling higher pulse-mode SNR, this arrangement allowed the consideration of alternative transmitter signals such as chirps and pseudo-random sequences [208-210], which although likely to benefit future work, were not used in the experiments described in this chapter.




Figure 6.15: The equipment configuration for the array imaging experiments.

As in the earlier synthetic aperture experiments, the USIP-12 flaw detector was used as the main amplifier. In this case, the transducer mode switch was set to "dual" and the output of the array interface board was connected directly to the receiver input. The USIP-12 display time-base and the Gage CS225 ADC were synchronised by triggering them in parallel from the SYNC output of the AWG2020.

6.4 Array Imaging Experiments

6.4.1 Array Response

Figure 6.16 shows the pulse-echo response of the bistatic array measured at the intersection of the last row and last column of the first array quadrant (corresponding approximately to the centre of the array). The test reflector used was a 100 mm thick, 200 mm² aluminium slab in immersion, approximately 120 mm distant from the array and perpendicular to its beam axis. The driving signal programmed into the AWG2020 was a Gaussian weighted pulse (see Figure 2.12 or 3.6) with a centre frequency of 1.0 MHz and a 6 dB bandwidth of 1.0 MHz. The peak pulse amplitude produced by the ENI LA-325 power amplifier was approximately 250 volts.





Figure 6.16: The pulse-echo response of the array transducer.


Although indicative of the overall gain of the transmitter, array and preamplifier, Figure 6.16 is not representative of the transmitted (or incident) field needed to calibrate the amplitudes and phases of the Green's function operators prior to image reconstruction⁵. This is because the array film is a distributed structure, which is likely to distort the received waveform. To determine the time and frequency-domain characteristics of the transmitted field alone, a hydrophone was fabricated by bonding 5 mm² of 110 μ m thick PVDF film to a 50 mm Perspex rod with 75 μ m double-sided tape. This was intended to mimic the reception characteristics of the array itself, but over a much smaller area. The characteristics of the signal received

⁵ In most of the following experiments, approximate plane wave transmission was assumed, which enabled the Green's function and incident field operators to be combined, as discussed previously in Section 2.3.3.


by the hydrophone when located on the axis of the array and at a range of 120 mm are illustrated in Figure 6.17. The measured centre frequency of the transmitted signal was 0.98 MHz and its bandwidth was 740 kHz.



(a)



(b)



(c)

Figure 6.17: The signal transmitted by the PVDF tile measured at a range of 120 mm. (a) Time domain waveform; (b) Frequency spectrum; (c) Normalised phase.

6.4.2 Data Acquisition

As in previous chapters, ten discrete temporal frequencies were selected for the initial series of image reconstruction experiments. These frequencies were centred on 0.98 MHz and ranged from 0.94 to 1.05 MHz, as indicated by the vertical bars in Figure 6.17 (b). The effect of using a wider reconstruction bandwidth, including up to forty frequencies, as well as using incident field approximations other than a simple plane wave will also be discussed later in this chapter.

Figure 6.18 illustrates the four immersion test targets that were used in the array imaging experiments. Targets (a), (b) and (d) were fabricated from 2 mm diameter enamelled copper wire as in Chapter 5, and target (c) was the cut end of an aluminium extrusion identical to that used in [180]. All targets were modelled as impenetrable, rigid (Neumann) boundaries in the forward modelling phase of image reconstruction. The region of interest (ROI) to be reconstructed in each case was a 48 mm^3 cube, divided into $32 \times 32 \times 32$ voxels, with a range offset R from the array aperture of approximately 100 mm (See Figure 2.2).




Figure 6.18: The test targets used in the array imaging experiments.

Because orthogonal array structures are only sensitive to spatial frequencies in their row and column directions, the imaging performance of the prototype array was evaluated with the test targets in various non-preferred orientations. As the diffracted components of acoustic fields tend to originate from edges [93], each test target was imaged with its straight edges at intermediate angles between the row and column directions of the orthogonal array. In all cases, it was observed that the array data contained significant spatial frequencies, irrespective of the target orientation. Figure 6.19 illustrates the real part of the data acquired from target (d) with its wire segments oriented at 45° to the x -axis. The discontinuities in the data, indicated by dashed lines, are due to the use of four electrically isolated orthogonal array cells.




Figure 6.19: The real part of the array data at 1.05 MHz corresponding to the bistatic field scattered from target (d) (area = 48 mm^2 , blue - min negative, red - max positive).

To compare the transfer functions of the array cells manufactured from PVDF film with that derived theoretically in Section 6.2.1, 2-D FFTs were computed for each quadrant of the array data in Figure 6.19 (including its imaginary part). The resulting Fourier coefficients, plotted in normalised amplitude in Figure 6.20, clearly

display the expected orthogonal array response, whereby only the coefficients on the $f_x = 0$ and $f_y = 0$ axes have been retained. These Fourier coefficients do not correspond directly to the data used in image reconstruction, which must be calculated from zero-padded cell data, as discussed in Section 6.2.3. Using a Fourier domain sampling operator such as that in Figure 6.3, however, is equally valid in the case of zero-padded data, as the Fourier coefficients on the $f_x = 0$ and $f_y = 0$ axes still uniquely describe the unpadded cell data. This is demonstrated in Appendix C.




Figure 6.20: Normalised FFT amplitudes $|D_i(f_x/\Delta f_x, f_y/\Delta f_y)|$ for the data in Figure 6.19.

In relation to the total time required to scan the orthogonal array, it had been intended that the rate of A-scan capture would be limited only by the pulse repetition frequency (PRF) of the transmitter. At the PRF of 500 Hz that was used, this would have meant that A-scans could have been acquired from all 1024 array elements in approximately 2 seconds. Unfortunately, following address transitions, it was found that the preamplifier needed approximately 100 ms to recover from saturation, which was probably caused by multiplexer charge-injection transients [201]. This extended the array scanning time to over 200 seconds, but did not impact on the quality of data, which in all cases displayed a peak differential-mode SNR of approximately 12 dB. This SNR value was estimated from field plots such as that in Figure 6.19.

6.4.3 Results and Discussion

In view of previous results, only MaxEnt was used to reconstruct 3-D images from orthogonal array data. Figures 6.21 to 6.24 are images of the targets in Figures 6.18 (a) to (d) respectively, reconstructed from ten discrete frequencies. The average time required to reconstruct these images (with algorithm A in Section 4.4) was 1120 seconds on a 200 MHz R10000 CPU. In addition to being shown in plan view and elevation according to the colour scale in Figure 3.5, each 3-D image has been rendered with translucent iso-surfaces at 20, 40, 60 and 80 % of its peak amplitude.

The first image in Figure 6.21 clearly shows the 25-mm long wire target, oriented at approximately 30° to the x -axis and parallel to the array aperture. Although the ends of the wire appear to have been reconstructed with higher amplitudes than the middle section, the image is still obviously that of a continuous object. As noted in Section 6.4.2, diffracted fields tend to be generated by discontinuities, and in this case, the wide range of spatial frequencies generated by the ends of the wire has made them more visible to the orthogonal array.

In addition to being non-preferentially oriented relative to the rows and columns of the array, the 30-mm wire target in Figure 6.22 was tilted at approximately 30° to the aperture plane. This resulted in the specularly reflected part of the scattered field not being intercepted by the array aperture, and hence only the ends of the wire were reconstructed. This behaviour is not a shortcoming of either the orthogonal array or MaxEnt, but is a well-known aspect of acoustic imaging whenever reflections from specular targets are involved (see for example [93] or page 273 of [22]).




Figure 6.21: An image of test target (a) reconstructed from orthogonal array data at ten frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).




Figure 6.22: An image of test target (b) reconstructed from orthogonal array data at ten frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).




Figure 6.23: An image of test target (c) reconstructed from orthogonal array data at ten frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).




Figure 6.24: An image of test target (d) reconstructed from orthogonal array data at ten frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).

The next target imaged was the end of the aluminium extrusion depicted in Figure 6.18 (c). With reference to [180], which is included as Appendix D, it can be seen that the orthogonal array image in Figure 6.23 is comparable to that obtained by using 6 % of the monostatic data arranged in a star pattern. In this case, however, the lower corner of the extrusion is smudged, and the border is not as continuous as in the monostatic image. Clearly, if images of more complicated targets were to be reconstructed with the array, then more information would be required, either in the form of additional temporal frequencies or a larger number of array cells.

The 3-D wire target depicted in Figure 6.18 (d) was the most complex target imaged in this series of experiments, and occupied the largest proportion of the ROI. As seen in Figure 6.24, MaxEnt reconstructed all four of the wire segments reasonably well, however the top and bottom segments in particular appear slightly distorted. This distortion was thought to be associated with inadequate modelling of the transducer incident field, which was approximated by a simple plane wave. In most respects, the image in Figure 6.24 is comparable to that reconstructed in the synthetic aperture experiments in Chapter 5 by using 6 % of the monostatic data arranged in a star pattern. As in the case of target (c), improving the orthogonal array image would require additional frequencies or array cells.

The effect on image quality of increasing the reconstruction bandwidth was tested by first using twenty, and then forty discrete frequencies. Surprisingly, as seen in Figures 6.25 and 6.26, the improvements obtained were very modest. In fact, the image reconstructed from forty frequencies shows some deterioration relative to that reconstructed from twenty frequencies. Again, this was attributed to the inadequacies of the plane wave approximation used during image reconstruction.

To test the accuracy of the plane wave assumption itself, incident fields were calculated numerically at ten frequencies according to Equation (2.46). Figure 6.27, which is an amplitude plot of the incident field calculated at a frequency of 1.05 MHz and a range of 100 mm, indicates the inaccuracy of the plane wave assumption in the near field of the PVDF tile⁶. In particular, the rapid variation in pressure in the plane of the tile, and the tapering-off at its edges would go some way to explaining the distortion observed, especially at the extremities of Figures 6.24 to 6.26.

⁶ According to diffraction theory [22], the transition region between the near field and the far field for a square transducer of width w is at $z = 1.2w^2 / 4\lambda$, or 516 mm for the transmitter tile used in the array transducer.




Figure 6.25: An image of test target (c) reconstructed from orthogonal array data at twenty frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).




Figure 6.26: An image of test target (c) reconstructed from orthogonal array data at forty frequencies via complex-valued MaxEnt. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).




Figure 6.27: An amplitude plot of the time-harmonic incident field pressure generated by the PVDF transmitter tile at a frequency of 1.05 MHz and a range of 100 mm, calculated according to Equation (2.46) (arbitrary units, area = 48 mm²).

Figure 6.28 shows the result of using algorithm B (discussed in section 4.4) to reconstruct an image of test target (d) from the same data used in Figure 6.24. Unfortunately, this image is very poor in relation to those in Figures 6.24-6.26, indicating that the original plane wave assumption more accurately represented the actual incident field than the fields calculated specifically for the PVDF transmitter tile. The full reason for this is unknown, but is probably related to the construction of the array transducer. If the incident field could be made more homogeneous, or at least be more accurately modelled, then it is likely that the quality of the reconstructed images would be higher than that achieved here.

In contrast to the 1120 seconds required by algorithm A to reconstruct images from ten discrete frequencies, image reconstruction using twenty and forty frequencies required 2030 and 3609 seconds respectively on an R10000 CPU. Using algorithm B with ten frequencies, the reconstruction time was 4287 seconds. Clearly, the plane wave approximation allows significant acceleration of MaxEnt image reconstruction, and is therefore desirable in practical applications.




Figure 6.28: An image of test target (d) reconstructed via complex-valued MaxEnt and based on numerically calculated incident fields. (a) Plan; (b) Elevation; (c) iso-surfaces at 20,40,60 and 80 % amplitude ($32 \times 32 \times 32$ voxels, 1.5 mm^3 per voxel).

Aside from the difficulties of accurately modelling the incident fields, the MaxEnt images obtained in this chapter display two impressive characteristics. These are the lateral resolution achieved from bistatic data acquisition, and the range resolution achieved with only ten discrete frequencies. Given the combination of target range, wavelength and aperture-size, the lateral resolution predicted by Equation 1.4 is almost 4 mm. As many of the features reconstructed were of the order of a single pixel (1.5 mm) wide, it is apparent that a fair degree of super-resolution was achieved. In relation to the range resolution achieved, this was single pixel in most cases, confirming that MaxEnt yields excellent solutions to broadband inverse scattering problems, even when temporal bandwidth is severely restricted.

As a final point of interest, the real parts of the orthogonal array data acquired when imaging target (d) were compared with those reconstructed by the Cambridge algorithm in its final forward modelling phase prior to convergence. In Figure 6.29, it is clear by visual comparison that MaxEnt has faithfully reproduced the row and column structure in the orthogonal array data at each of the ten frequencies. As the reconstructed data also contain features that could not have been recorded by the array, it is clear that unmeasured spatial frequencies have been extrapolated from the measured data. In consideration of Figure 6.24, it is actually surprising that such an image could have been reconstructed from so little information.

6.5 Summary

In this chapter, an alternative scheme for 2-D bistatic array sampling was demonstrated that holds important advantages over conventional point-sampling in terms of simplicity of array construction and low element impedance. Although the level of performance demonstrated by the prototype array would be useful in some existing applications, it is likely that an imaging system tailored to the specific requirements of NDE would require more sophisticated manufacturing techniques. In particular, a real NDE imaging system would use a smaller aperture and piezoceramic materials better matched to the acoustic impedances of metallic structures. Given that an NDE imaging system would operate in the near field, as was the case with the prototype system, it would be important to accurately model the incident field, or better still, adaptively estimate it during image reconstruction. These issues will be discussed further in the conclusion.




Figure 6.29: Comparisons of the real parts of the orthogonal array data for target (d) with the data reconstructed via MaxEnt (left – array data, right – MaxEnt data).

Chapter 7

Conclusion

In the introduction to this thesis, the question of which image reconstruction algorithm is capable of producing the best image from a given set of ultrasonic data was posed. Although it is unlikely that there will ever be a definitive answer to this question, information theory suggests that a strong candidate for the “best” image is that which is *minimally informative* with respect to the available information. In addition to prior knowledge, therefore, reconstructed images should reflect only the information contained in the data, without spurious features associated with characteristics of the imaging system. Jaynes' maximum entropy principle, and MaxEnt in particular, are useful approximations that allow *configurational information* to be measured in a sense, and minimised to yield solutions that are consistent with the data while being minimally informative in the same sense.

This thesis has taken a largely experimental and hence comparative path to exploring the benefits of MaxEnt regularisation in ultrasonic image reconstruction. The major inadequacies of conventional approaches to solving such problems were identified, not in relation to the linearisation of the underlying inverse problem, but rather the inadequate treatment of incomplete data. It was demonstrated that resolution and side-lobe clutter could be substantially improved by extending the definition of the Shannon Entropy into the complex domain and using this as a regularising functional in much the same way as for real-positive images.

In comparison with the range of alternative approaches investigated, complex-valued MaxEnt emerged as the clear leader in terms of resolution and overall image quality. Matched filtering was seen to produce reasonable images in circumstances where the aperture data were substantially complete, yet even then, the physical characteristics of propagation lead to severe resolution limitations. With incomplete

aperture data, the performance of matched filtering rapidly deteriorated, and the addition of linear regularisation was seen not to improve the situation.

Certainly, the coherent CLEAN algorithm, which was specially adapted to broadband image reconstruction for this thesis, displays some of the desirable properties of MaxEnt, such as side-lobe clutter rejection, and can be faster when only simple scatterers are involved. CLEAN lacks any real super-resolution capability, however, and in the limit of sparse data and low SNR, was shown to display poor performance due to its treatment of images as collections of discrete sources.

Apart from its obvious successes, there is currently a scarcity of *philosophical* justification for applying MaxEnt to complex-valued image reconstruction, which reflects earlier confusion regarding the meaning of entropy itself. Given this unresolved controversy, various authors have seen fit to distinguish between "conventional" MaxEnt, and *entropy-like* regularisation, which is conceptually applicable to a wider class of problems, including those involving complex values [101, 211]. From this thesis, it is apparent that MaxEnt is no less *experimentally* justified for complex-valued images than it is for real-positive images.

One interpretation for complex-valued MaxEnt is as a form of marginalisation for image phases, which can be considered as "nuisance parameters" in a Bayesian context [134, 156]. This interpretation emphasises the association between MaxEnt and Bayesian methods generally, whereby the Shannon entropy can be viewed as a means of assigning prior probabilities to images which are compatible with a given data set. As discussed in Chapter 3, the Bayesian framework is the most powerful and general for dealing with problems of inference, but as a special case, MaxEnt holds the advantage of being computationally tractable for large-scale problems.

The major conclusions reached in this research were as follows:

- The Shannon entropy retains powerful regularising characteristics in the absence of a positivity constraint, meaning that super-resolution and clutter rejection properties carry over to image reconstruction problems involving coherent data.
- In view of the low amplitude of the background scattering in homogeneous media, a small value for the parameter ' A ' in Equation (4.3) is equally applicable in ultrasonic imaging as it is in astronomical imaging. This low amplitude "default model" provides a powerful regularising effect.

- Using MaxEnt image reconstruction, high-quality images of scatterers can be reconstructed from very small numbers of aperture samples, and the specific sampling pattern used has little bearing on the resulting image.
- MaxEnt regularisation is an effective anti-aliasing measure in cases involving PSFs with large regions of support, such as free-space Green's functions. This is because MaxEnt is biased against image features that are not essential in achieving agreement between the forward scattering model and the data.
- The benefits of MaxEnt in ultrasonic imaging are further amplified through using multiple frequencies in broadband image reconstruction. This is because MaxEnt explicitly seeks to reconstruct a single image that agrees with the data, rather than simply summing time-harmonic images that correspond to different frequencies.
- In adapting the Cambridge MaxEnt algorithm to handle complex-valued images, both the composite image and complex metric approaches lead to similar convergence behaviour to the unmodified Cambridge algorithm [2, 46, 51].
- By using conservative estimates of the SNR in MaxEnt image reconstruction, the impact of non-linear and second order phenomena, such as multiple scattering, on image reconstruction can be minimised.

7.1 MaxEnt Implementation

As this thesis was mainly concerned with the properties of complex-valued maximum entropy images, rather than devising new approaches to optimisation, many of the developments here were based on the Cambridge algorithm, which was already highly refined. Despite its success in reconstructing the images in this thesis, the computational complexity of the Cambridge algorithm guarantees that the development of new MaxEnt algorithms will be ongoing.

As is customary in digital signal processing (DSP) research, the basic properties of the complex-valued MaxEnt algorithm have been investigated using general-purpose hardware, however this is by no means the preferred platform on which to implement the algorithm. The Cambridge algorithm is FFT intensive and amenable

to parallel computation. In this regard, it should be noted that the FFT is one of the most highly optimised signal processing algorithms in existence, and as such, it is well suited to direct hardware acceleration or efficient DSP implementation. Employing a large number of inexpensive commodity DSP devices to perform the bulk of the computation would appear to be the most economical path to applying MaxEnt algorithms to practical problems in ultrasonic imaging.

7.2 Prototype Imaging System

In relation to the prototype imaging system discussed in Chapter 6, it was demonstrated that MaxEnt enabled considerable flexibility in the way data was acquired from 2-D apertures, provided the resulting transfer function was linear. The number of orthogonal array cells used in this instance was probably the minimum useful number, below which, the data would become too dependent on the orientation and scattering behaviour of the target(s).

Many of the characteristics of the prototype array discussed in Chapter 6 would be directly applicable to ultrasonic NDE, provided different materials, a finer element pitch, and a higher operating frequency were used. Problems noted in relation to the speed of data acquisition could be solved with relatively minor modifications to the multiplexer. In any case, the time-critical activity would be that of image reconstruction rather than data acquisition, especially as the targets to be imaged in NDE are stationary. It should be noted that the difficulties in modelling the incident fields in Chapter 6 were mainly due to the ROI being in the near field of the transmitter. In this region, the amplitudes and phases of the incident fields vary rapidly and are sensitive to the boundary conditions at the aperture. Notwithstanding adaptive estimation of the incident field, which is discussed further in the next section, many of the problems experienced would disappear when operating in the far field of the transmitter, as might be case in a large-scale application, such as underwater imaging, or when using a point source.

In underwater imaging, a scaled-up and optimised version of the prototype imaging system presented here would be superior in terms of image quality to systems based on point-sampled arrays and conventional linear beamforming. The price of simplifying the "front-end" array structure would be paid for by a substantial increase in the complexity of the "back-end" signal processing. Nevertheless, the use of orthogonal array sampling could offer advantages over point-sampling in many array-based imaging applications, regardless of whether MaxEnt was employed for image reconstruction or not.

7.3 Suggestions for Further Research

Aside from investigating improvements to the basic complex-metric implementation of the Cambridge algorithm, or devising entirely new algorithms, the most important area for future research appears to be in adaptively modelling both the incident field and the frequency dependence of scattering during image reconstruction. As discussed in Chapter 2, the scattering behaviour of simple objects can be approximated reasonably well provided a dominant scattering mechanism; that is, penetrable or impenetrable scattering can be identified. When this is not the case, however, the discrepancies between the reconstructed and measured field data might adversely affect image reconstruction, particularly when large numbers of frequencies are employed. Similarly, errors in the incident field modelling can also affect the quality of reconstructed images, as was the case in Chapter 6.

One promising path to achieving such adaptive modelling involves interleaving the maximisation steps of an expectation-maximisation (EM) algorithm between successive iterations of the Cambridge algorithm to progressively refine estimates of the incident field and scattering parameters. Such an approach would be inconceivable with the majority of optimisation algorithms, however as the Cambridge algorithm does not carry curvature information between iterations, its convergence behaviour should not be affected by updating the scattering model.

Interestingly, it is well known that combinatorial approaches to calculating conditional expectations in the EM algorithm (cf. the monkey argument in Section 3.3.2) have a very similar effect to that of entropy [212]. Replacing the expectation steps in the EM algorithm with iterations of the Cambridge algorithm, therefore, would seem to be a natural way of embedding adaptive parameter estimation within the MaxEnt framework. Overall, the effect of this adaptive estimation would be to make MaxEnt more robust against unknown parameters in the scattering model. This additional robustness would be required if wide bandwidths were to be employed, as would be the case when using spread-spectrum signals such as chirps or pseudo-random sequences to improve the SNR during data acquisition. Decoupling the scattering model from the parameters of the scattering medium would also allow the methods developed here to be applied to more general and complex problems, such as those encountered in medical imaging.

In concluding this thesis, it must be emphasised that the MaxEnt algorithms developed here represent only the first generation of such algorithms applicable to ultrasonics. For now, therefore, the potential user should be encouraged by the possibilities, rather than discouraged by the computational complexity.

Appendix A

Derivation of Formula for Penetrable Scattering

To arrive at a simple, linear model for penetrable scattering, several approximations and assumptions are necessary. These have been made explicit in this derivation, which starts with the inhomogeneous form of the time-harmonic Helmholtz equation

$$[\nabla^2 + k^2] p(\mathbf{r}, \omega) = -o(\mathbf{r}, \omega), \quad (\text{A.1})$$

where the secondary source term is given by [20]

$$o(\mathbf{r}, \omega) = k^2 \gamma_\kappa(\mathbf{r}) p(\mathbf{r}, \omega) - \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega)]. \quad (\text{A.2})$$

To solve this equation for the scattered pressure field $p_s(\xi', \omega)$, the scalar Green's function $g_m(|\xi' - \mathbf{r}|, \omega)$ is required, which is the solution to the canonical problem

$$[\nabla^2 + k^2] g_m(|\xi' - \mathbf{r}|, \omega) = -\delta(|\xi' - \mathbf{r}|) \quad (\text{A.3})$$

involving a delta function in m dimensions (see section 2.4). Multiplying Equation (A.3) by $p(\mathbf{r}, \omega)$ and subtracting from it $g_m(|\xi' - \mathbf{r}|, \omega)$ times Equation (A.1) gives

$$\begin{aligned} g_m(|\xi' - \mathbf{r}|, \omega) \nabla^2 p(\mathbf{r}, \omega) - p(\mathbf{r}, \omega) \nabla^2 g_m(|\xi' - \mathbf{r}|, \omega) \\ = p(\mathbf{r}, \omega) \delta(|\xi' - \mathbf{r}|) - g_m(|\xi' - \mathbf{r}|, \omega) o(\mathbf{r}, \omega). \end{aligned} \quad (\text{A.4})$$

Now integrating each side over V_0 , which is selected to encompass the secondary sources of interest, and using the sifting property of the delta function; that is,

$$\int_{V_0} p(\mathbf{r}, \omega) \delta(|\xi' - \mathbf{r}|) d^m \mathbf{r} = p(\xi', \omega), \quad (\text{A.5})$$

the sound pressure at any point ξ' in the receiving aperture can be written as the following sum of volume integrals

$$p(\xi', \omega) = \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) o(\mathbf{r}, \omega) d^m \mathbf{r} + \int_{V_0} [g_m(|\xi' - \mathbf{r}|, \omega) \nabla^2 p(\mathbf{r}, \omega) - p(\mathbf{r}, \omega) \nabla^2 g_m(|\xi' - \mathbf{r}|, \omega)] d^m \mathbf{r} \quad (\text{A.6})$$

The second volume integral can be transformed into a surface integral by invoking Green's Theorem [168]; that is,

$$\int_V (\phi \nabla^2 \psi - \psi \nabla^2 \phi) d^m \mathbf{r} = \oint_S (\phi \nabla \psi - \psi \nabla \phi) \cdot \hat{\mathbf{n}} da, \quad (\text{A.7})$$

where ϕ and ψ are any scalar fields. This yields the sound pressure in the receiving aperture ξ' in terms of a convolution between $g_m(|\xi' - \mathbf{r}|, \omega)$ and the secondary source function $o(\mathbf{r}, \omega)$ plus a surface integral term over $S_1 + S_2$ (see Figure 2.1)

$$p(\xi', \omega) = \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) o(\mathbf{r}, \omega) d^m \mathbf{r} + \oint_{S_1 + S_2} [g_m(|\xi' - \mathbf{r}|, \omega) \nabla p(\mathbf{r}, \omega) - p(\mathbf{r}, \omega) \nabla g_m(|\xi' - \mathbf{r}|, \omega)] \cdot \hat{\mathbf{n}} da \quad (\text{A.8})$$

The surface integral is required to satisfy boundary conditions associated with $S_1 + S_2$ in addition to the incident field generated by the transmitting aperture. All outward travelling waves from V_0 will either contribute nothing to this integral by virtue of the Sommerfeld radiation condition [29], or generate a reflected contribution at boundaries of the medium. In practice, unwanted reflections from boundaries can usually be gated out of the measurements in a similar manner to unwanted wave modes, as discussed in Section 2.2. Having dispensed with unwanted boundary reflections, Equation (A.9) is seen to represent the pressure field in the aperture as the sum of incident and scattered field components according to

$$p(\xi', \omega) = p_i(\xi', \omega) + p_s(\xi', \omega), \quad (\text{A.9})$$

where the scattered pressure field is given by the volume integral

$$p_s(\xi', \omega) = \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) o(\mathbf{r}, \omega) d^m \mathbf{r}. \quad (\text{A.10})$$

At this point it is worth reflecting on the form of Equation (A.10), as it is approaching the form of a convolution that is desirable for efficient numerical calculation. However, the secondary source $o(\mathbf{r}, \omega)$ remains field dependent, and the total (incident plus scattered) field has not been ascertained. Substituting Equation (A.2) into Equation (A.10) gives

$$p_s(\xi', \omega) = \int_{V_0} k^2 \gamma_\kappa(\mathbf{r}) p(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} - \int_{V_0} \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega)] g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \quad (\text{A.11})$$

Equation (A.11) can be simplified by applying the general vector identity [168]

$$\phi \nabla \cdot \mathbf{a} = \nabla \cdot (\phi \mathbf{a}) - \mathbf{a} \cdot \nabla \phi, \quad (\text{A.12})$$

whereby

$$\begin{aligned} \int_{V_0} \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega)] g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} &= \\ \int_{V_0} \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega)] d^m \mathbf{r} & \\ - \int_{V_0} \gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega) \cdot \nabla g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} & \end{aligned} \quad (\text{A.13})$$

With the aid of the divergence theorem [168], it can be shown that

$$\begin{aligned} \int_{V_0} \nabla \cdot [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega)] d^m \mathbf{r} &= \\ \oint_S [\gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega)] \cdot \hat{\mathbf{n}} da & \end{aligned} \quad (\text{A.14})$$

and this surface integral can be eliminated entirely by choosing any suitable boundary S such that $\gamma_\rho = 0$ (i.e. any closed surface outside the inhomogeneous region). The scattered field $p_s(\xi', \omega)$ is then given by

$$\begin{aligned} p_s(\xi', \omega) &= \int_{V_0} k^2 \gamma_\kappa(\mathbf{r}) p(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \\ &+ \int_{V_0} \gamma_\rho(\mathbf{r}) \nabla p(\mathbf{r}, \omega) \cdot \nabla g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \end{aligned} \quad (\text{A.15})$$

and it remains to evaluate the gradients of both the Green's function and total pressure field $p(\mathbf{r}, \omega)$ with respect to the co-ordinate system depicted in Figure 2.1.

For any of the Green's functions listed in Table 2.1, it can be shown that the gradient $\nabla g_m(|\xi' - \mathbf{r}|, \omega)$ can be approximated in the far field by [62, 64]

$$\nabla g_m(|\xi' - \mathbf{r}|, \omega) = \nabla g_m(|\mathbf{r} - \xi'|, \omega) \cong jk g_m(|\mathbf{r} - \xi'|, \omega) \frac{(\mathbf{r} - \xi')}{|\mathbf{r} - \xi'|}. \quad (\text{A.16})$$

It is now required to evaluate the gradient of the total pressure field $p(\mathbf{r}, \omega)$, which by Equation (A.9), is the sum of the incident and scattered fields, the latter of which is an unknown quantity. By invoking the Born approximation [20], which holds for weak scattering, $p(\mathbf{r}, \omega)$ may be approximated by the incident field; that is,

$$p(\mathbf{r}, \omega) \cong p_i(\mathbf{r}, \omega). \quad (\text{A.17})$$

By further assuming that $p_i(\mathbf{r}, \omega)$ can be approximated by a plane wave propagating in the direction $\mathbf{r} - \xi$, the gradient $\nabla p_i(\mathbf{r}, \omega)$ is given by

$$\nabla p_i(\mathbf{r}, \omega) \cong jk p_i(\mathbf{r}, \omega) \frac{(\mathbf{r} - \xi)}{|\mathbf{r} - \xi|}. \quad (\text{A.18})$$

Substituting Equations (A.16) and (A.17) back into (A.15) and neglecting the directivity of the dipole secondary sources due to the density variation, or equivalently, assuming that the angle between the incident and scattered fields is approximately zero; that is

$$\frac{(\mathbf{r} - \xi)}{|\mathbf{r} - \xi|} \cdot \frac{(\mathbf{r} - \xi')}{|\mathbf{r} - \xi'|} \cong 1, \quad (\text{A.19})$$

Equation (A.15) can be rewritten as

$$\begin{aligned} p_s(\xi', \omega) = & \int_{V_0} k^2 \gamma_\kappa(\mathbf{r}) p_i(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \\ & - \int_{V_0} k^2 \gamma_\rho(\mathbf{r}) p_i(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r} \end{aligned} \quad (\text{A.20})$$

By substituting $f(\mathbf{r}) = \gamma_\kappa(\mathbf{r}) - \gamma_\rho(\mathbf{r})$ for the reflectivity function, this reduces to

$$p_s(\xi', \omega) = k^2 \int_{V_0} f(\mathbf{r}) p_i(\mathbf{r}, \omega) g_m(|\xi' - \mathbf{r}|, \omega) d^m \mathbf{r}, \quad (\text{A.21})$$

which is a formula for penetrable scattering with the desired convolutional form.

Appendix B

Derivation of Formula for Impenetrable Scattering

Appendix A gave the derivation of the penetrable scattering formula. The corresponding derivation for impenetrable scattering commences with the integral theorem of Helmholtz and Kirchhoff [29], which has the same form as the surface integral of Equation (A.8), except in this case it is over the characteristic surface of the scatterer S_c , such that

$$p_s(\xi', \omega) = \oint_{S_c} [p(\mathbf{r}, \omega) \nabla g_m(|\xi' - \mathbf{r}|, \omega) - g_m(|\xi' - \mathbf{r}|, \omega) \nabla p(\mathbf{r}, \omega)] \cdot \hat{\mathbf{n}}' da, \quad (\text{B.1})$$

where the normal $\hat{\mathbf{n}}'$ is defined in Figure 2.1. At the two extremes, impenetrable scatterers may possess boundary conditions which are either perfectly soft (stress-free), or perfectly rigid. Cracks and voids in solid materials are well modelled as stress-free boundaries and will be considered first. Provided the time-domain displacement $\mathbf{u}(\mathbf{r}, t)$ associated with the pressure field $p(\mathbf{r}, t)$ is always less than the width of the scatterer [67], the characteristic surface S_c may be assumed to represent a Dirichlet boundary condition at all frequencies, i.e.

$$p(\mathbf{r}, \omega) = 0 \quad \mathbf{r} \in S_c \quad (\text{B.2})$$

Substituting this boundary condition into Equation (B.1) gives

$$p(\xi', \omega) = - \oint_{S_c} g_m(|\xi' - \mathbf{r}|, \omega) \nabla p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' da. \quad (\text{B.3})$$

Due to the fact that wave energy does not penetrate into the scatterer, the entire scattering function can be represented by the *singular function* [52, 68]

$$\gamma(\mathbf{r}) = \nabla \Gamma(\mathbf{r}) \cdot \hat{\mathbf{n}}', \quad (\text{B.4})$$

where $\Gamma(\mathbf{r})$ is the associated *characteristic function* given by

$$\Gamma(\mathbf{r}) = \begin{cases} 1 & \mathbf{r} \in V_c \\ 0 & \mathbf{r} \notin V_c \end{cases}. \quad (\text{B.5})$$

By incorporating Equation (B.4) into (B.3), and invoking the integral properties of delta functions, Equation (B.3) can be transformed into the volume integral

$$p_s(\xi', \omega) = - \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) \gamma(\mathbf{r}) \nabla p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}}' d^m \mathbf{r}. \quad (\text{B.6})$$

To proceed with the analysis, it is necessary to linearise Equation (B.6), as the total field $p(\mathbf{r}, \omega)$ is still dependent on the scattered field $p_s(\xi', \omega)$. Weak scattering arguments are not applicable here as they were in Appendix A, as impenetrable scatterers are hardly "weak" in the sense of not disrupting the incident field. Kirchhoff theory, however, offers a simple approximation for $p(\mathbf{r}, \omega)$ whereby each point on the illuminated portion of the singular function, $\gamma'(\mathbf{r})$, is assumed to behave as if it were part of an infinite plane with a prescribed reflection coefficient [53, 69]. This is a high-frequency approximation, as the degree to which an object looks specular is determined by its size relative to the wavelength at which it is observed. Using Kirchhoff theory, the total field derivative at the boundary V_c of a perfectly soft scatterer is taken to be simply twice the incident field derivative that would exist in the absence of the boundary; that is,

$$\nabla p(\mathbf{r}, \omega) \cdot \mathbf{n}' \cong 2 \nabla p_i(\mathbf{r}, \omega) \cdot \mathbf{n}' \quad \mathbf{r} \in S_c \quad (\text{B.7})$$

$$\cong 2 j k p_i(\mathbf{r}, \omega) \frac{(\mathbf{r} - \xi)}{|\mathbf{r} - \xi|}. \quad (\text{B.8})$$

As in Appendix A, the form of the incident field $p_i(\mathbf{r}, \omega)$ remains undefined, however a plane wave approximation is assumed to be adequate when computing its gradient. Substituting (B.8) into (B.6) gives

$$p(\xi', \omega) = -2jk \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) \gamma'(\mathbf{r}) p_i(\mathbf{r}, \omega) \frac{(\mathbf{r} - \xi)}{|\mathbf{r} - \xi|} \cdot \hat{\mathbf{n}}' d^m \mathbf{r}. \quad (\text{B.9})$$

The final simplifying assumption will be that the surfaces of the impenetrable scatterer are effectively normal to the incident beam, i.e.

$$\frac{(\mathbf{r} - \xi)}{|\mathbf{r} - \xi|} \cdot \hat{\mathbf{n}}' \cong -1. \quad (\text{B.10})$$

This is a far field assumption, the accuracy of which in the near field is about 5% if the angle remains less than 18° [29]. The final form of the forward scattering equation for impenetrable scatterers is given by

$$p_s(\xi', \omega) \cong 2jk \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) f(\mathbf{r}) p_i(\mathbf{r}, \omega) d^m \mathbf{r}, \quad (\text{B.11})$$

where the image function $f(\mathbf{r})$ has been substituted for the illuminated portion of the singular function. The derivation of the scattering equation for Neumann boundary conditions, representing perfectly rigid scatterers, i.e.

$$\nabla p(\mathbf{r}, \omega) \cdot \hat{\mathbf{n}} = 0 \quad \mathbf{r} \in S_c \quad (\text{B.12})$$

proceeds in a similar manner to that above, giving

$$p_s(\xi', \omega) \cong -2jk \int_{V_0} g_m(|\xi' - \mathbf{r}|, \omega) f(\mathbf{r}) p_i(\mathbf{r}, \omega) d^m \mathbf{r}. \quad (\text{B.13})$$

Thus Equation (B.11) serves as a general formula for impenetrable scattering under both Dirichlet and Neumann boundary conditions, with the sign of $f(\mathbf{r})$ inverted in the latter case to reflect the change in sign of the reflectivity.

Appendix C

The Orthogonal Array Transfer Function and Zero-Padded Data

The purpose of this appendix is to illustrate that zero-padding has no effect on the orthogonal array transfer function discussed in Section 6.2.1. This is important, because in a multi-cell array, the sub-apertures represented by individual cells must be zero-padded to the size of the total aperture if they are to be considered additive.

Starting with quadrant 1 (cell 1) of the data acquired from target (d) in Figure 6.18, zeroing extraneous off-axis frequencies yields a set of Fourier coefficients that describes the orthogonal field, as shown in Figure C.1. This corresponds to multiplying the Fourier data D_1 by the transfer function T defined in Section 6.2.3.




Figure C.1: Quadrant 1 of the data acquired from target (d) in Figure 6.18 with off-axis frequencies removed. (a) Filtered array data (real part); (b) FFT amplitude spectrum.

To show that this transfer function is invariant under zero-padding, the original array data d_1 from quadrant 1 is zero-padded to twice its original dimensions, as would be the case in reconstructing images from multiple orthogonal array data. The amplitude of the corresponding spectrum D'_1 shows considerable off-axis energy, as would be expected of any 2-D function with finite support.




Figure C.2: (a) The original data from quadrant 1, zero-padded in the x and y directions (real part shown); (b) FFT amplitude spectrum showing off-axis Fourier coefficients.

By filtering the zero-padded orthogonal array data d'_1 in Figure C.2 (a) in a similar manner to Figure C.1, the data and FFT amplitude spectrum in Figure C.3 result.




Figure C.3: (a) The result of filtering the zero-padded array data in a similar manner to Figure C.1; (b) FFT amplitude spectrum (The operator $\mathbf{T}'$ in this case operates on D'_1).

After sampling the field in the lower-left of Figure C.3 (a) with an appropriate spatial domain operator $\mathbf{S}_1$, it is apparent that the result is identical to Figure C.1 (a). This shows that the Fourier coefficients on the $f_x = 0$ and $f_y = 0$ axes of the FFT spectrum in Figure C.3 (b) uniquely define the original unpadded array data. The orthogonal array transfer function is therefore invariant under zero-padding, which allows the data from many orthogonal array cells to be combined during image reconstruction, according to Equations (6.13) to (6.16).

Maximum Entropy Image Reconstruction from Sparsely Sampled Coherent Field Data

David J. Battle, *Member, IEEE*, Robert P. Harrison, and Mark Hedley, *Member, IEEE*

Abstract—There are many practical problems in which it is required to detect and characterize hidden structures or remote objects by virtue of the scattered acoustic or electromagnetic fields they generate. It remains an open question, however, as to which reconstruction algorithms offer the most informative images for a given set of field measurements. Commonly used time-domain beamforming techniques, and their equivalent frequency-domain implementations, are conceptually simple and stable in the presence of noise, however, large proportions of missing measurements can quickly degrade image quality. In this paper, we apply a new algorithm based on the maximum entropy method (MEM) to the reconstruction of images from sparsely sampled coherent field data. The general principles and limitations of the new method are discussed in the framework of regularization theory, and the results of monostatic imaging experiments confirm that superior resolution and artifact suppression are obtained relative to a commonly used linear inverse filtering approach.

I. INTRODUCTION

SINCE lensless image reconstruction from scattered fields was first reported by Gabor in 1948 [1], the theory of inverse scattering has been extensively developed and applied to the fields of optics, radar, sonar, medical ultrasonics, geophysical exploration, and the nondestructive evaluation of materials (NDE).

Medical ultrasonics has stimulated many developments in acoustic imaging, and in particular, has led to a high state of refinement in linear array technology. Commercial linear array scanners are able to provide near ideal cross-sectional images by dynamically steering and focusing a sound beam in two dimensions. Beyond this, the natural progression toward more sophisticated imaging capabilities has led to the consideration of two-dimensional (2-D) arrays as a means of achieving true three-dimensional (3-D) imaging. When coupled with advanced visualization technology, 3-D acoustic imaging is a potentially powerful tool, particularly in the field of NDE, where knowledge of the geometry of material defects is critical in many situations.

Although several workers have reported new developments in 2-D array technology [2]–[5], progress toward a practical device has been hampered by the complexities of fabricating

and controlling the large numbers of elements required. In this paper, we address the problem of imaging with so-called “sparse arrays” from which a large proportion of the elements have been removed for reasons of economy. While it is well known that sparse arrays allow substantial improvements in resolution over filled arrays for a given number of active elements [6], the resulting images usually suffer from unwanted side-lobe artifacts, which are a function of the array element pattern. Of interest here, are improved methods for processing array data that minimize image distortion due to missing information.

The maximum entropy method (MEM) represents one attempt to incorporate *information* or *complexity* measures into image reconstruction, which has been found useful due to its robust treatment of both noisy and incomplete data [7]. Wernecke and D’Addario [8] demonstrated that superior images of astronomical radio sources were resolved via MEM reconstruction relative to direct Fourier transformation. In that case, the Van Cittert–Zernike theorem established a direct Fourier transform relation between the spatial autocorrelation function of the *incoherent* electric fields measured by interferometers, and the image distribution, which took the form of a nonnegative power spectrum. This paper aims to demonstrate a new application of the MEM principle, wherein high quality images of acoustic scatterers are reconstructed from sparsely sampled *coherent* fields. For simplicity, the discussion and experimental validation are limited to monostatic imaging, although the method is not restricted to this arrangement. A brief account of the relevant scattering model is given in Section II. Section III describes a commonly used inverse filtering approach to image reconstruction from wavefield data, and the limitations imposed by linear methods. The concept of regularization is introduced in Section IV, and the particular optimization algorithm developed by Skilling and Bryan [9] is adapted to coherent wavefields in Sections V–VIII. The remainder of the paper constitutes a comparison between the new results and those obtained by inverse filtering. The aim of this work is to show that significant redundancy exists in wavefield data, and that the imaging performance of sparse sensor arrays may approach that of filled arrays, when the appropriate signal processing is employed.

II. SCATTERING MODEL

In this paper, we concentrate on monostatic data acquisition due to its simplicity and suitability for both mechanically scanned and array based imaging systems. We consider either a single transducer scanned through a set of aperture coordi-

Manuscript received August 30, 1995; revised July 12, 1996. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Hua Lee.

D. J. Battle and R. P. Harrison are with the Materials Division, Australian Nuclear Science and Technology Organization, Lucas Heights, Australia 2234 (email: davidb@ansto.gov.au).

M. Hedley is with the Department of Electrical Engineering, Sydney University, Sydney, Australia 2006 (email: hedley@ee.su.oz.au).

Publisher Item Identifier S 1057-7149(97)05477-8.

nates ξ , or equivalently, a physical array of fixed transducers deployed over the same set of coordinates. The scattered continuous wave (CW) scalar field Φ_s is assumed to originate from the interaction between an incident field Φ_i and a scattering potential $f(\mathbf{r})$, which is a function of the position vector $\mathbf{r}$. All fields are assumed to be propagating as longitudinal waves in a lossless medium, and mode conversion effects have been neglected. As our immediate motivation is toward imaging defects in metallic components, we specialize to the case of impenetrable scatterers. For such scatterers, $f(\mathbf{r})$ takes on the character of a singular function, or visible surface with respect to the measurement aperture, and this is what we seek to reconstruct. Following the customary physical optics assumptions which hold adequately for convex scatterers at high frequencies [10], Φ_s for the monostatic arrangement is found to be related to $f(\mathbf{r})$ through the convolution integral

$$\Phi_s(\xi) \approx 2jk \iiint G^2(|\xi - \mathbf{r}|) f(\mathbf{r}) d^3\mathbf{r} \quad (1)$$

where

$$G(|\xi - \mathbf{r}|) = \frac{\exp(jk|\xi - \mathbf{r}|)}{4\pi|\xi - \mathbf{r}|} \quad (2)$$

and k is the wave number of the surrounding medium at the given frequency. The squared Green's function $G^2(|\xi - \mathbf{r}|)$ accounts for two-way propagation between points in the aperture plane and points within the interrogated volume, and in an image processing context, can be interpreted as a shift invariant point spread function (PSF). Given the bandlimited nature of scattered wavefields which results from the loss of evanescent components within a few wavelengths of a scatterer [11], $\Phi_s(\xi)$ may be sampled without loss of information, allowing the continuous model in (1) to be replaced by a discrete matrix model. Our hypothetical imaging system can thus be represented by the standard matrix equation

$$\mathbf{d} = \mathbf{H}\mathbf{f} + \mathbf{n} \quad (3)$$

where $\mathbf{H}$ is a block circulant matrix operator, $\mathbf{d}$ and $\mathbf{f}$ are row or column ordered vectors containing samples of $\Phi_s(\xi)$ and $f(\mathbf{r})$, respectively, and $\mathbf{n}$ is a vector of complex noise terms, (assumed here to be Gaussian with uniform variance $2\sigma^2$). To permit the use of fast Fourier transforms (FFT's) in subsequent block-circulant matrix multiplications, wavefield and image vectors are nominally sampled over a rectangular grid of N points. For reasons noted in the introduction, however, it may not be possible to sample the incident wavefield at all grid points within the aperture, so a sampling operator

$$\mathbf{S} = \text{diag}[s_i] \quad (i = 1, \dots, N) \quad (4)$$

is introduced where $s_i = 1 + j0$ for each of the m measured complex samples, and $s_i = 0 + j0$ otherwise. Equation (3) then becomes

$$\mathbf{d} = \mathbf{S}(\mathbf{H}\mathbf{f} + \mathbf{n}). \quad (5)$$

Equation (5) constitutes the discrete *forward* scattering model used throughout this paper. The process of forming an image estimate $\hat{\mathbf{f}}$ from data $\mathbf{d}$ is termed the *inverse* scattering problem.

III. LINEAR INVERSE FILTERING

A popular direct approach to solving the inverse scattering problem relies upon finding the pseudoinverse operators $\mathbf{H}^\#$ and $\mathbf{S}^\#$ such that the image estimate $\hat{\mathbf{f}}$ is obtained (symbolically) as

$$\hat{\mathbf{f}} = \mathbf{H}^\# \mathbf{S}^\# \mathbf{d}. \quad (6)$$

In practice, allowance for the noise term is incorporated through additional linear filtering, which we omit here for brevity. The properties of images obtained through this type of reconstruction are well documented [12]–[14], and generalizations to 3-D imaging with broadband (impulse or coded) signals are easily derived [10], [15]. Direct application of the pseudoinverse $\mathbf{H}^\#$ implies the minimum norm solution to (3), which implicitly assumes zero values for all unmeasured wavefield samples external to the aperture. This results in a broad main lobe in the beam pattern, and restricts the image resolution to that determined by the Rayleigh criterion [11].

While the lateral resolution of array systems employing linear beamforming can be improved by using larger apertures, this means that either additional elements are required if the array is to remain filled at the Nyquist spacing of $\lambda/2$, or the same number of elements must be redeployed in some thinned configuration. In (5), missing elements are reflected as zeros in the diagonal sampling operator $\mathbf{S}$. Ideally, an inverse sampling operator $\mathbf{S}^\#$ is required to counter the effects of these missing elements. Unlike $\mathbf{H}^\#$ however, the operator $\mathbf{S}^\#$ has no simple interpretation, and usually entails some type of interpolation operation in either the spatial or frequency domains [16]. In general, such interpolation schemes need to be coupled to the appearance of the reconstructed image through iterative means, as in the regularized algorithm presented here, or the image may suffer severe loss of dynamic range due to side-lobe clutter. Following Wernecke and D'Addario [8], we set $\mathbf{S}^\#$ to the identity matrix in computing inverse filtered reconstructions for comparison with regularized results, as there is often no better approximation.

IV. REGULARIZED INVERSION

In view of the inherent deficiencies of linear filtering as a solution to the inverse scattering problem, we based our solution on iterative application of the *forward* scattering model, which is always well posed, at least as far as the approximations inherent in the scattering model. The regularized approach to the inversion of ill-posed problems differs from simple inverse filtering by allowing a compromise between fidelity to the data, and fidelity to some *a priori* information about the nature of the solution [17]. In general, the available information in the form of the data $\mathbf{d}$, will be insufficient to determine a unique image estimate $\hat{\mathbf{f}}$, but can be expected to define a “null” space or set of feasible solutions $\hat{\mathbf{F}}$ such that $\hat{\mathbf{f}} \in \hat{\mathbf{F}}$ [18].

A. The Feasible Solution Set

The feasibility of a solution $\hat{\mathbf{f}}$ is conveniently measured through a single Chi-squared statistic, which is defined here

in matrix form as

$$\chi^2 = \frac{1}{2\sigma^2} \|\mathbf{S}(\mathbf{H}\hat{\mathbf{f}} - \mathbf{d})\|^2. \quad (7)$$

From the definition adopted in (7) and the assumed Gaussian distribution of $\mathbf{n}$, the expected value of χ^2 for all $\hat{\mathbf{f}} \in \hat{\mathbf{F}}$ is m [19].

B. The Entropy Function

For detailed accounts of maximum entropy principles, the reader is referred to [8], [9], [20], and [21]. Frieden first described the basic characteristics of maximum entropy reconstructions through the super-resolution of one-dimensional (1-D) incoherent objects with a diffraction limited imaging system. Some of the main properties of reconstructions $\hat{\mathbf{f}}$ via Frieden's method are summarized as follows.

- 1) $\hat{\mathbf{f}}$ is restricted to be positive, as required for an image of an incoherent object (or where a complex valued image consists of real and imaginary parts; both are positive).
- 2) Each measurement in $\mathbf{d}$ contributes a degree of freedom (DOF) to $\hat{\mathbf{f}}$. As the number of measurements decreases, $\hat{\mathbf{f}}$ becomes increasingly conservative (smooth), reflecting only the minimum *configurational information* consistent with the measurements.
- 3) $\hat{\mathbf{f}}$ is not intrinsically bandlimited. The resolution obtained may significantly exceed the Rayleigh limit.

Numerous practical and philosophical rationale exist for using entropy to regularize image reconstructions [20]. The fact remains that entropy is only one of a class of nonlinear regularizing functionals that may be used to resolve the degeneracy associated with ill-posed inverse problems. Although it is often claimed that the positivity constraint inherent in MEM formulations is largely responsible for their super-resolution capabilities [19], it has been noted that the nonlinearity also plays an important role in sharpening high intensity image features while suppressing spurious background structure [18], [22]. Similarities between the source distributions imaged in radioastronomy, and those often encountered in NDE, make the method potentially attractive for imaging cracks and other defects, which are characterized as small, high contrast scatterers, or "nearly black objects" for which super-resolution has been shown to be especially effective [22]. Given these similarities, the entropy function adopted by Skilling and Bryan [9], is also adopted here. Hence

$$E = - \sum_{j=1}^N \hat{\mathbf{f}}_j \log_e \left(\frac{\hat{\mathbf{f}}_j}{eA} \right) \quad (8)$$

where N is the number of pixels in the image, and A is a default intensity, whose exact value is only weakly related to the actual appearance of the reconstructed image [19]. As noted by Press *et al.* [18], the definition of entropy appears to be an open one. Various simple functions other than the strict Shannon entropy may be employed with little impact on the resulting reconstructions.

V. OPTIMIZATION ALGORITHM

The quantity to be minimized in finding the maximum entropy image is given by

$$J = \chi^2 - \lambda E \quad (9)$$

with the constraint that χ^2 achieves its expected value of m (the number of measurements). Maximum entropy algorithms are necessarily iterative due to the nonlinearity of (8). Nonlinear optimization problems are notoriously difficult, and in this particular case, we have the added complication that the image must at all times remain positive. Various algorithms based on steepest ascent and conjugate gradients have been used with some degree of success [8], [21], however these algorithms suffer from the common problem that even small numbers of negative pixels appearing at intermediate stages of the optimization severely impede progress toward the solution. In this work, we based our approach on a general optimization algorithm, which avoids problems associated with intermediate nonpositive values by calculating subspace search directions which favor image positivity [9]. The salient aspects of the optimization algorithm are presented here so as to allow discussion of our extension to the basic model in (3). For complete details of this standard algorithm, the reader is directed to [9] and [19].

A. Derivative Formulae

In common with other image reconstruction methods based on optimizations, that of Skilling and Bryan requires the calculation of first and second partial derivatives of the objective function(s) with respect to the image pixel values. The gradients and second partial derivative (Hessian) matrices for χ^2 and E are given by

$$\nabla \chi^2 = \frac{\partial(\chi^2)}{\partial \hat{\mathbf{f}}_i} = \frac{1}{\sigma^2} \mathbf{H}^* \mathbf{T} \mathbf{S} (\mathbf{H}\hat{\mathbf{f}} - \mathbf{d}) \quad (10)$$

$$\nabla \nabla \chi^2 = \frac{\partial^2(\chi^2)}{\partial \hat{\mathbf{f}}_i \partial \hat{\mathbf{f}}_j} = \frac{1}{\sigma^2} [\mathbf{H}^* \mathbf{T} \mathbf{S} \mathbf{H}]_{ij} \quad (11)$$

$$\nabla E = \frac{\partial E}{\partial \hat{\mathbf{f}}_i} = \log_e A - \log_e \hat{\mathbf{f}}_i \quad (12)$$

$$\nabla \nabla E = \frac{\partial^2 E}{\partial \hat{\mathbf{f}}_i \partial \hat{\mathbf{f}}_j} = -\frac{\delta_{ij}}{\hat{\mathbf{f}}_j}. \quad (13)$$

The block circulant structure of $\mathbf{H}$ means that the matrix multiplications in (10) and (11) may be carried out efficiently via FFT algorithms [18], [23]. Explicit calculation and storage of the Hessian $\nabla \nabla \chi^2$ is never required, as matrix-vector multiplications involving $\nabla \nabla \chi^2$ are carried out via sequential application of $\mathbf{H}$, $\mathbf{S}$, and $\mathbf{H}^* \mathbf{T}$. It must be noted that all operations involving the sampling operator $\mathbf{S}$ must be performed in the spatial domain, as in general, $\mathbf{S}$ will not be block circulant and, therefore, will not be diagonalized by the FFT. In addition to the two FFT operations required when $\mathbf{S}$ is an identity matrix (corresponding to a filled aperture), there is an overhead of two FFT operations for each multiplication involving $\nabla \nabla \chi^2$ when sparse data are considered.

B. Subspace Basis Vectors

The optimization algorithm proceeds by calculating three basis vectors given by $\mathbf{e}_\mu$ ($\mu = 1, 2, 3$) [9], [18] where

$$\left. \begin{aligned} \mathbf{e}_1 &= \hat{\mathbf{f}}(\nabla\chi^2) \\ \mathbf{e}_2 &= \hat{\mathbf{f}}(\nabla E) \\ \mathbf{e}_3 &= \frac{\hat{\mathbf{f}}(\nabla\nabla\chi^2\mathbf{e}_2)}{\sqrt{\hat{\mathbf{f}} \cdot (\nabla E)^2}} - \frac{\hat{\mathbf{f}}(\nabla\nabla\chi^2\mathbf{e}_1)}{\sqrt{\hat{\mathbf{f}} \cdot (\nabla\chi^2)^2}} \end{aligned} \right\} \quad (14)$$

and the notation $\hat{\mathbf{f}}(\dots)$ implies pointwise multiplication.¹ These three vectors form a subspace into which the quantities (10)–(13) are projected via dot products to form two quadratic models given by

$$\tilde{E}(x) = E + E_\mu x^\mu - \frac{1}{2} g_{\mu\nu} x^\mu x^\nu \quad (15)$$

$$\tilde{\chi}^2(x) = \chi^2 + C_\mu x^\mu + \frac{1}{2} M_{\mu\nu} x^\mu x^\nu \quad (16)$$

where

$$\left. \begin{aligned} E_\mu &= \mathbf{e}_\mu \cdot \nabla E \\ g_{\mu\nu} &= \mathbf{e}_\mu \cdot \mathbf{e}_\nu \\ C_\mu &= \mathbf{e}_\mu \cdot \nabla\chi^2 \\ M_{\mu\nu} &= \mathbf{e}_\mu \cdot \nabla\nabla\chi^2 \cdot \mathbf{e}_\nu \end{aligned} \right\} \quad (17)$$

and the x^μ ($\mu = 1, 2, 3$) correspond to increments by which the search directions are multiplied in updating the current image estimate according to

$$\hat{\mathbf{f}}_{n+1} = \hat{\mathbf{f}}_n + \delta\hat{\mathbf{f}} \quad \text{where } \delta\hat{\mathbf{f}} = x^\mu \mathbf{e}_\mu. \quad (18)$$

In projecting the first and second partial derivatives into the subspace, the quantity J in (9) is remodeled as locally quadratic at each iteration. As curvature information is not propagated forward to successive iterations, the stagnation problems associated with conjugate gradient algorithms are avoided [18].

C. Control Steps

As part of the convergence control algorithm, the nine component Hessian matrices $g_{\mu\nu}$ and $M_{\mu\nu}$ are subjected to simultaneous diagonalization and eigenvalue filtering via a standard singular value decomposition (SVD) algorithm [18]. The use of SVD in this regard is important, due to the fact that the search directions (14) may occasionally develop linear dependencies. Progress toward the ME solution is

¹For a discussion of the significance of these three search directions, the reader is directed to [9], [18], [19], and [24].

made by calculating the increments x^μ ($\mu = 1, 2, 3$) in the subspace such that the model entropy $\tilde{E}(x)$ is maximized as $\tilde{\chi}^2$ converges to m . The precise algorithm to achieve this is detailed at length in [9] and [24], and although embodying most of the sophistication of the overall algorithm, requires negligible computation relative to calculating the search directions and their projections. Stopping criteria for the optimization algorithm are based on either the approach to zero of the composite gradient ∇J , or a check for parallelism between its two vector components.

VI. EXTENSION TO COHERENT FIELDS

The positivity constraint inherent in MEM formulations places potential restrictions on the type of images that can be reconstructed from coherent field data. Burch [25] reported a series of experiments in which multiple drilled holes in a metal block were well resolved via MEM reconstruction, provided a corresponding PSF was *measured* under precisely the same conditions. In general however, range, frequency, and velocity errors serve to corrupt the phases of the echo signals, and the real or imaginary parts of pixel values in the reconstruction may be required to be zero or negative in order for χ^2 to achieve its expected value of m . Nonpositive pixel values are incompatible with (8) and are therefore inadmissible even as intermediate results in MEM algorithms. To overcome this problem, we used a revised image model in which the complex image $\mathbf{f}$ was replaced by an ensemble of four sub-images $\mathbf{f}_\alpha$, ($\alpha = 1, 2, 3, 4$); each phase shifted to occupy a prescribed phase quadrant while maintaining positivity of its real and imaginary parts. Equation (3) thus becomes

$$\mathbf{d} = \mathbf{H} \sum_{\alpha=1}^4 \mathbf{f}_\alpha \exp\left(\frac{j\alpha\pi}{2}\right) + \mathbf{n} \quad (19)$$

leading to the following revised form for (7):

$$\chi^2 = \frac{1}{2\sigma^2} \left\| \mathbf{S} \left(\mathbf{H} \sum_{\alpha=1}^4 \hat{\mathbf{f}}_\alpha \exp\left(\frac{j\alpha\pi}{2}\right) - \mathbf{d} \right) \right\|^2. \quad (20)$$

Use of the composite image model requires that the basis vectors in (14) be split into $\mathbf{e}_{\mu\alpha}$ where $\mu = 1, 2, 3$ and $\alpha = 1, 2, 3, 4$. The calculations for $\nabla\chi^2$, ∇E , $\nabla\nabla E$, $\mathbf{e}_1$, and $\mathbf{e}_2$ proceed as before for each of the sub-images, however the multiplications by $\nabla\nabla\chi^2$ in the calculation of $\mathbf{e}_3$ are complicated by cross terms in the second partial derivatives of (20) as shown in (21), at the bottom of the page.

$$\begin{bmatrix} \nabla\nabla\chi^2\mathbf{e}_{\mu 1} \\ \nabla\nabla\chi^2\mathbf{e}_{\mu 2} \\ \nabla\nabla\chi^2\mathbf{e}_{\mu 3} \\ \nabla\nabla\chi^2\mathbf{e}_{\mu 4} \end{bmatrix} = \frac{1}{\sigma^2} \begin{bmatrix} 1 & \exp\left(\frac{j\pi}{2}\right) & -1 & \exp\left(\frac{-j\pi}{2}\right) \\ \exp\left(\frac{-j\pi}{2}\right) & 1 & \exp\left(\frac{j\pi}{2}\right) & -1 \\ -1 & \exp\left(\frac{-j\pi}{2}\right) & 1 & \exp\left(\frac{j\pi}{2}\right) \\ \exp\left(\frac{j\pi}{2}\right) & -1 & \exp\left(\frac{-j\pi}{2}\right) & 1 \end{bmatrix} \begin{bmatrix} \mathbf{H}^{*T} \mathbf{S} \mathbf{H} \mathbf{e}_{\mu 1} \\ \mathbf{H}^{*T} \mathbf{S} \mathbf{H} \mathbf{e}_{\mu 2} \\ \mathbf{H}^{*T} \mathbf{S} \mathbf{H} \mathbf{e}_{\mu 3} \\ \mathbf{H}^{*T} \mathbf{S} \mathbf{H} \mathbf{e}_{\mu 4} \end{bmatrix} \cdots \quad (21)$$

When calculating the entropy, the summation in (8) is extended over the four positive sub-images to act as the regularizing functional for the complex image as a whole. Using the composite image model does increase the complexity of calculating the subspace basis vectors, however in most respects, the optimization algorithm behaves as reported by the original authors [9]. On satisfying the stopping criteria, the reconstructed image is obtained through addition of the sub-images and calculating the magnitude from the real and imaginary parts.

VII. COMPUTATIONAL COMPLEXITY

As coded in ANSI C, the composite-image MEM algorithm takes approximately 5 min to reconstruct a 64×64 pixel CW image on one CPU of a SPARCstation model 502 workstation. The algorithm requires 45 multidimensional FFT's at each iteration, and approximately 50 iterations are usually necessary before convergence for signal to noise ratios of 20 dB. As noted by Skilling and Bryan [9], reconstructions at lower noise levels require larger numbers of iterations. FFT's dominate the execution time such that significant performance improvements could be realized by using high-speed digital signal processor (DSP) hardware specially designed for FFT operations. Parallel processing is an attractive option given the number of image transformations in each iteration.

VIII. EXPERIMENTAL PROCEDURE

A. General Arrangement

Fig. 1 depicts the arrangement of key equipment for an imaging experiment conducted to investigate the performance of the above algorithm. The main parameters for the experiment were as follows. ...

- *Aperture:* 76.8×76.8 mm with 64×64 samples on a 1.2 mm grid
- *Transducer:* 1.0 MHz focused at 25.4 mm, scanned monostatically
- *Target:* Aluminum extrusion in H_2O , 200 mm from aperture plane

Transducer positioning and echo data acquisition were controlled by a personal computer. The combination of a USIP-12 flaw detector and TDS-520 digital oscilloscope performed the functions of a pulser/receiver. At each of the 4096 sample locations, a 500-point broadband echo signal was acquired by the oscilloscope and immediately Fourier transformed into its CW components. For the purposes of testing the current algorithm, only CW reconstructions at a temporal frequency of 1.0 MHz were attempted, although work is well advanced on a multifrequency algorithm.

B. PSF Model Validation

The Greens function expression in (2) models a spherically spreading wave, and therefore the response of an omnidirectional transmitter or receiver, but does not take into account the directivity pattern associated with a practical array element. The monostatic directivity pattern of the focused transducer used to simulate individual array elements in our experiment




Fig. 1. (a) Monostatic scanning arrangement for the acoustic imaging experiment. (b) Test target dimensions.

was measured in order to more correctly model the PSF of the imaging system. Fig. 2 shows a central slice through the real part of the 2-D PSF obtained experimentally from a 1 mm diameter spherical target at a range of 200 mm and a frequency of 1.0 MHz. The directivity pattern of the transducer was modeled by a Gaussian function approximating the envelope of this signal. It may be observed that the noise amplitude has increased between samples 20 and 40 in Fig. 2. This was attributed to nonlinearity in the analog electronics, which introduced errors into the FFT coefficients. Under the prevailing experimental conditions, it was estimated that the average SNR of the data was approximately 20 dB.

Fig. 3 shows good agreement between the experimentally and theoretically determined PSF's for the given imaging experiment. The theoretical PSF in Fig. 3(b) was calculated by sampling $G^2(|\xi - \mathbf{r}|)$ at the target range and applying the same amplitude taper as that measured for the experimental PSF. An additional phase shift was then applied in order to allow better comparison with the experimental data. Due to the fact that the composite image model allows the reconstruction algorithm to respond to the *phase curvature* of the data rather than its principle value phase, this comparison remains valid.




Fig. 2. Cross section of the experimental PSF. A Gaussian curve with a standard deviation of 14 samples was used to model the transducer directivity profile (real part shown; arbitrary amplitude units).




Fig. 3. Comparison between (a) experimental and (b) theoretical PSF's for the imaging experiment described above (real parts shown).

C. Experimental Results

The experimental arrangement in Fig. 1(a) was used to model four types of monostatic array. Fig. 4 depicts a comparison between reconstructions from the full 1.0 MHz array data via inverse filtering and MEM algorithms, respectively. Sparse




Fig. 4. Comparison of reconstruction results from the full array data consisting of 4096 complex samples. (a) Real part of array data. (b) Inverse filtered reconstruction. (c) MEM reconstruction. Both images are displayed with ten linear magnitude contours. Image coordinates correspond to 1.2 mm sample increments.

array configurations were simulated by selectively zeroing 94% to 97% of the diagonal elements of the sampling operator

$\mathbf{S}$, thereby decoupling the image from the corresponding field measurements. Figs. 5–7 compare the experimental results for random, star, and cross array configurations, respectively.

IX. DISCUSSION

The full array reconstruction results in Fig. 4 reveal that the test target has been more highly resolved by the MEM algorithm than by inverse filtering. This indicates a similar super-resolution property to that observed in other applications where maximum entropy has been successfully applied [7]–[9], [19], [21], [26]. Super-resolution can be partly ascribed to the aperture extension effect whereby the data vector $\hat{\mathbf{d}} = \mathbf{H}\hat{\mathbf{f}}$ is estimated over the support for $\mathbf{H}$ rather than that of $\mathbf{d}$. This contrasts with the inverse filtering approach, wherein $\mathbf{d}$ is regarded as completely specified, and zero for all missing samples. An additional tendency toward super-resolution enters through the nonlinear entropy function E , which discriminates against small values, and sharpens high intensity point sources [18], [22].

Fig. 5 depicts the results of image reconstruction from a 6% filled random array. Conventional random array techniques aim to lessen side-lobe levels by achieving beam patterns that are optimally uncorrelated outside the main beam [27]. The effect of random element placement is therefore to spread the side-lobe energy as incoherently as possible, so as to minimize clutter concentrations in the resulting images. Notwithstanding adaptive beamforming techniques, the end result is that both complicated and simple objects are imaged uniformly poorly, whereas from the perspective of information theory, it is desirable to make full use of the available channel capacity under all circumstances. Ideally, simple objects could be reconstructed as perfectly as permitted by the degrees of freedom (DOF's) in the measurements, with the image quality expected to degrade as the complexity of the object exceeded the DOF of the imaging system. The MEM reconstruction in Fig. 5(c) displays this behavior, with almost no degradation relative to the full array result in Fig. 4(c). Figs. 5(b) and 6(b) display marked differences in their background structures, which are functions of the respective array configurations. It is apparent that regularized image reconstructions show far less dependence on the array configuration than do inverse filtered reconstructions.

Fig. 7 shows the result of further reducing the number of array elements so that only 3% of the samples remain in a cross configuration. While neither image would allow identification of the target shape, Fig. 7(c) is localized to the true dimensions of the target, and displays none of the extended artifacts that characterize Fig. 7(b).

X. CONCLUSIONS

The application of maximum entropy principles to ultrasonic imaging has shown that considerable improvements are achievable relative to conventional inverse filtering. These improvements manifest as enhanced resolution with full array data, and substantial side-lobe artifact reductions for sparse array configurations. Given the general nature of the scattering model employed, similar results should be obtainable with




Fig. 5. Reconstruction results for a 6% filled random array. (a) Sampling mask for full array data. (b) Inverse filtered reconstruction versus (c) MEM reconstruction (ten linear magnitude contours shown).

other types of wave energy (microwaves, seismic waves, etc.). While there is a significant computational burden associated






Fig. 6. Reconstruction results for a 6% filled star array. (a) Sampling mask for full array data. (b) Inverse filtered reconstruction versus (c) MEM reconstruction (ten linear magnitude contours shown).






Fig. 7. Reconstruction results for a 3% filled cross array. (a) Sampling mask for full array data. (b) Inverse filtered reconstruction versus (c) MEM reconstruction (ten linear magnitude contours shown).

with maximum entropy algorithms, there are many situations where it would be preferable to employ enhanced “back-end”

signal processing as an alternative to investing in a more complex “front-end” array structure. As part of the future

research, it is intended to improve the performance of the MEM algorithm by utilizing DSP hardware, and to develop an approach to multifrequency imaging, whereby the quality of images may again be substantially improved.

ACKNOWLEDGMENT

The authors thank the anonymous reviewers for their helpful comments in relation to the manuscript.

REFERENCES

- [1] D. Gabor, "A new microscopic principle," *Nature*, vol. 161, pp. 777–778, 1948.
- [2] D. N. Hall, G. Hayward, and Y. Gorfu, "Theoretical and experimental evaluation of a two-dimensional composite matrix array," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 40, pp. 704–709, 1993.
- [3] D. H. Turnbull and S. F. Foster, "Fabrication and characterization of transducer elements in two-dimensional arrays for medical ultrasound imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 39, pp. 464–475, 1992.
- [4] S. W. Smith, H. G. Pavy, and O. T. von Ramm, "High-speed ultrasound volumetric imaging system—Part I: Transducer design and beam steering," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 38, pp. 100–108, 1991.
- [5] J. D. Plummer *et al.*, "Two-dimensional transmit/receive ceramic piezoelectric arrays: Construction and performance," *IEEE Trans. Son. Ultrason.*, vol. 25, pp. 273–280, 1978.
- [6] R. M. Leahy and B. D. Jeffs, "On the design of maximally sparse beamforming arrays," *IEEE Trans. Antennas Propagat.*, vol. 39, pp. 1178–1186, 1991.
- [7] S. F. Gull and J. Skilling, "Maximum entropy method in image processing," *Proc. Inst. Electr. Eng.*, vol. 131, pp. 646–659, 1984.
- [8] S. J. Wernecke and L. R. D'Addario, "Maximum entropy image reconstruction," *IEEE Trans. Comput.*, vol. C-26, pp. 351–364, 1977.
- [9] J. Skilling and R. K. Bryan, "Maximum entropy image reconstruction: General algorithm," *Mon. Not. R. Astronom. Soc.*, vol. 211, pp. 111–124, 1984.
- [10] K. Mayer *et al.*, "Three-dimensional imaging system based on Fourier transform synthetic aperture focusing technique," *Ultrasonics*, vol. 28, pp. 241–255, 1990.
- [11] J. W. Goodman, *Introduction to Fourier Optics*, Electrical and Electronic Engineering Series. New York: McGraw-Hill, 1968, pp. 36–37.
- [12] A. J. Devaney, "A fast filtered backpropagation algorithm for ultrasound tomography," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 34, pp. 330–339, 1987.
- [13] V. I. Neeley, *Development of a Production Prototype Pressure Vessel Imaging System*, EPRI NP-3253, Electr. Power Res. Inst., Washington, DC, 1983.
- [14] R. Kiefer, V. Schmitz, and H. Wendel, "Fast recording and numerical reconstruction of holograms," in *Proc. Third Int. Conf. Non-Destructive Evaluation in the Nuclear Industry*, ASME, 1981.
- [15] L. J. Busse, "Three dimensional imaging using a frequency domain synthetic aperture focussing technique," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 39, pp. 174–179, 1992.
- [16] A. Klug and R. A. Crowther, "Three dimensional image reconstruction from the viewpoint of information theory," *Nature*, vol. 238, pp. 435–440, 1972.
- [17] G. Demoment, "Image reconstruction and restoration," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. 37, pp. 2024–2036, 1989.
- [18] H. P. Press *et al.*, *Numerical Recipes in C*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 1992, p. 994.
- [19] S. F. Burch, S. F. Gull, and J. Skilling, "Image restoration by a powerful maximum entropy method," *Comput. Vis. Graph. Image Processing*, vol. 23, pp. 113–128, 1983.
- [20] E. T. Jaynes, "On the rationale of maximum entropy methods," *Proc. Inst. Electr. Eng.*, vol. 70, pp. 939–952, 1982.
- [21] B. R. Frieden, "Restoring with maximum likelihood and maximum entropy," *J. Opt. Soc. Amer.*, vol. 62, pp. 511–518, 1972.
- [22] D. L. Donoho *et al.*, "Maximum entropy and the nearly black object," *J. Roy. Stat. Soc. B.*, vol. 54, pp. 41–81, 1992.
- [23] B. R. Hunt, "A matrix theory proof of the discrete convolution theorem," *IEEE Trans. Audio Electron.*, vol. 19, pp. 285–288, 1971.
- [24] H. Stark, *Image Recovery Theory and Application*, 4th ed. New York: Academic, 1987, p. 543.
- [25] S. F. Burch, "Comparison of SAFT deconvolution methods for the improvement of ultrasonic B-Scan images," *Ultrasonics*, vol. 25, pp. 259–265, 1987.
- [26] G. J. Daniell and S. F. Gull, "Maximum entropy algorithm applied to image enhancement," *Proc. Inst. Electr. Eng.*, vol. 127, pp. 170–172, 1980.
- [27] M. I. Skolnik, Ed., *Radar Handbook*, 2nd ed. New York: McGraw-Hill, 1991.



David J. Battle (S'86–M'87) was born in Brisbane, Australia. He received the B.Eng. degree in electrical engineering in 1988 from the Queensland University of Technology, Brisbane. He is a candidate for the Ph.D. degree at the Electrical Engineering Department, University of Sydney, Australia.

He is currently employed as Nondestructive Evaluation Engineer at the Australian Nuclear Science and Technology Organization, Sydney. His research interests include inverse problems in

ultrasonic imaging, array transducers, and novel approaches to high-resolution array signal processing.



Robert P. Harrison was educated in England and Australia, and received the B.Sc. degree in physics and pure mathematics from the University of Sydney, Australia, in 1974, and the M.Sc. in solid state physics from the University of New South Wales, in 1978.

He is currently manager of the Materials Assessment Section, Materials Division, Australian Nuclear Science and Technology Organization. His research interests include the development of NDE techniques, particularly ultrasound-based, the study of aging of research reactor components, and the determination of material properties using miniature samples.



Mark Hedley (S'89–M'95) was born in Sydney, Australia. He received the B.Sc. degree in pure mathematics and computer science in 1987, and the B.E. and Ph.D. degrees in electrical engineering in 1989 and 1993, respectively, all from the University of Sydney, Australia.

He has been with Telecom Australia and is currently a lecturer in the Electrical Engineering department, University of Sydney. His research interests include real-time image processing, stereo vision, image reconstruction, and engineering education.

Bibliography

- [1] S. J. Wernecke and L. R. D'Addario, "Maximum entropy image reconstruction," *IEEE Transactions on Computers*, vol. 26, pp. 351-364, 1977.
- [2] S. F. Burch, S. F. Gull, and J. Skilling, "Image restoration by a powerful maximum entropy method," *Computer Vision Graphics and Image Processing*, vol. 23, pp. 113-128, 1983.
- [3] S. F. Gull and G. J. Daniell, "Image restoration from incomplete and noisy data," *Nature*, vol. 272, pp. 686-690, 1978.
- [4] C. T. Mottershead, "Maximum entropy tomography," in *Maximum Entropy and Bayesian Methods*, K. M. Hanson and R. N. Silver, Eds. Dordrecht: Kluwer, 1996, pp. 425-436.
- [5] M. C. Lawrence, M. A. Jaffer, and B. T. Sewell, "The application of the maximum entropy method to electron microscopic tomography," *Ultramicroscopy*, vol. 31, pp. 285-301, 1989.
- [6] R. T. Smith, C. K. Zoltani, G. J. Klem, and M. W. Coleman, "Reconstruction of tomographic images from sparse data sets by a new finite element maximum entropy approach," *Applied Optics*, vol. 30, pp. 573-582, 1991.
- [7] K. M. Hanson, "Bayesian and related methods in image reconstruction from incomplete data," in *Image Recovery: Theory and Application*, H. Stark, Ed. New York: Academic Press, 1987, pp. 79-127.
- [8] J. C. Hoch and A. S. Stern, *NMR Data Processing*, 1st ed. New York: Wiley-Liss, 1996.
- [9] J. B. Abbiss, C. De Mol, and H. S. Dhadwal, "Regularized iterative and non-iterative procedures for object restoration from experimental data," *Optica Acta*, vol. 30, pp. 107-124, 1983.
- [10] A. R. Davies and R. S. Anderssen, "Optimization in the regularization of ill-posed problems," *Journal of the Australian Mathematical Society B.*, vol. 28, pp. 114-133, 1985.

- [11] A. N. Tikhonov and V. Y. Arsenin, *Solutions of Ill-Posed Problems*. New York: Wiley, 1977.
- [12] C. E. Shannon, "A mathematical theory of communication," *Bell Systems Technical Journal*, vol. 27, pp. 379-423, 1948.
- [13] E. T. Jaynes, "Information theory and statistical mechanics II," *Physical Review*, vol. 108, pp. 171-190, 1957.
- [14] E. T. Jaynes, "Bayesian methods: general background," in *Maximum Entropy and Bayesian Methods in Applied Statistics*, J. H. Justice, Ed. Cambridge: Cambridge University Press, 1986, pp. 1-25.
- [15] T. J. Loredo, "From Laplace to Supernova SN 1987A: Bayesian inference in astrophysics," in *Maximum Entropy and Bayesian Methods*, P. F. Fougère, Ed. Dordrecht: Kluwer, 1990, pp. 81-142.
- [16] J. P. Burg, "Maximum entropy spectral analysis," presented at the 37th Meeting of the Society of Exploration Geophysicists, 1967.
- [17] E. D. Laue, M. R. Mayger, J. Skilling, and J. Staunton, "Reconstruction of phase-sensitive two-dimensional NMR spectra by maximum entropy," *Journal of Magnetic Resonance*, vol. 68, pp. 14-29, 1986.
- [18] G. J. Daniell and P. J. Hore, "Maximum entropy and NMR - A new approach," *Journal of Magnetic Resonance*, vol. 84, pp. 515-536, 1989.
- [19] P. J. Hore, "Maximum entropy and nuclear magnetic resonance," in *Maximum Entropy in Action*, B. Buck and V. A. Macaulay, Eds. Oxford: Clarendon Press, 1991, pp. 41-72.
- [20] P. M. Morse and K. U. Ingard, *Theoretical Acoustics*. New York: McGraw-Hill, 1968.
- [21] K. J. Langenberg, "Applied inverse problems for acoustic, electromagnetic and wave scattering," in *Basic Methods of Tomography and Inverse Problems*, *Malvern Physics Series*, P. C. Sabatier, Ed., 1st ed. Bristol: Adam Hilger, 1987, pp. 127-476.
- [22] G. S. Kino, *Acoustic Waves: Devices, Imaging and analog signal processing*, 1st ed. New Jersey: Prentice hall, 1987.
- [23] R. B. Thompson and D. O. Thompson, "Ultrasonics in nondestructive evaluation," *Proceedings of the IEEE*, vol. 73, pp. 1716-1755, 1985.
- [24] M. Fink, "Ultrasound puts materials to the test," *Physics World*, vol. 11. No. 2, pp. 41-45, 1998.

- [25] S. Haykin, "Array Signal Processing," New Jersey: Printice Hall, 1985, 433.
- [26] S. R. Doctor, T. E. Hall, and L. D. Reid, "SAFT - the evolution of a signal processing technology for ultrasonic testing," *NDT International*, vol. 19, pp. 163-167, 1986.
- [27] K. J. Langenberg, M. Berger, T. Kreutter, K. Mayer, and V. Schmitz, "Synthetic aperture focussing technique signal processing," *NDT International*, vol. 19, pp. 177-189, 1986.
- [28] H. Krim and M. Viberg, "Two decades of array signal processing," *IEEE Signal Processing Magazine*, vol. 13 No. 4, pp. 67-94, 1996.
- [29] J. W. Goodman, *Introduction to Fourier Optics*. San Francisco: McGraw-Hill, 1968.
- [30] A. W. Rihaczek and S. J. Hershkowitz, *Radar Resolution and Complex Image Analysis*, 1st ed. Boston: Artech House, 1996.
- [31] H. Lee, "Resolution enhancement by wavefield extrapolation," *IEEE Transactions on Sonics and Ultrasonics*, vol. 31, pp. 642-645, 1984.
- [32] M. Lach, M. Platte, and A. Ries, "Piezoelectric materials for ultrasonic probes," . Online NDT Journal: Internet Transducer Workshop, 1996.
- [33] H. Lee and G. Wade, "High-resolution imaging for systems with small apertures," *Journal of The Acoustical Society of America*, vol. 72, pp. 2033-2035, 1982.
- [34] P. N. Keating, T. Sawatari, and G. Zilinskas, "Signal processing in acoustic imaging," *Proceedings of the IEEE*, vol. 67, pp. 496-509, 1979.
- [35] J. Tsao and B. D. Steinberg, "Reduction of sidelobe and speckle artifacts in microwave imaging: The CLEAN techniqe," *IEEE Transactions on Antennas and Propagation*, vol. 36, pp. 543-552, 1988.
- [36] J. Krautkramer and H. Krautkramer, *Ultrasonic Testing of Materials*, 1st ed. Berlin: Springer-Verlag, 1969.
- [37] H. Lee, "Inverse filter design for holographic imaging systems with small apertures," in *Acoustical Imaging*, vol. 14, A. J. Berkhout, J. Ridder, and L. F. van der Wal, Eds. New York: Plenum Press, 1985, pp. 715-718.
- [38] E. T. Jaynes, "Detection of extra-solar system planets," in *Maximum-Entropy and Bayesian Methods in Science and Engineering*, vol. 1: Foundations, *Fundamental Theories of Physics*, G. J. Erickson and C. R. Smith, Eds., 1st ed. Dordrecht: Kluwer, 1988, pp. 147-160.

- [39] M. I. Skolnik, "Radar Handbook," New York: McGraw-Hill, 1991,
- [40] B. Steinberg, *Principles of Aperture and Array System Design*. New York: Wiley, 1976.
- [41] R. M. Leahy and B. D. Jeffs, "On the design of maximally sparse beamforming arrays," *IEEE Transactions on Antennas and Propagation*, vol. 39, pp. 1178-1186, 1991.
- [42] H. W. Jones, "Superresolution in ultrasonic imaging," in *Acoustical Imaging*.
- [43] A. Schatzberg and A. J. Devaney, "Super-resolution in diffraction tomography," in *Inverse Problems*, vol. 8, 1992, pp. 149-164.
- [44] G. J. Daniell and S. F. Gull, "Maximum entropy algorithm applied to image enhancement," *IEE Proceedings*, vol. 127, pp. 170-172, 1980.
- [45] S. F. Gull and J. Skilling, "Maximum entropy method in image processing," *IEE Proceedings*, vol. 131, pp. 646-659, 1984.
- [46] J. Skilling and R. K. Bryan, "Maximum entropy image reconstruction: General algorithm," *Monthly Notices of the Royal Astronomical Society*, vol. 211, pp. 111-124, 1984.
- [47] S. Sibisi, J. Skilling, R. G. Brereton, E. D. Laue, and J. Staunton, "Maximum entropy signal processing in practical NMR spectroscopy," *Nature*, vol. 311, pp. 446-447, 1984.
- [48] J. A. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astronomy and Astrophysics (Supplement)*, vol. 15, pp. 417-426, 1974.
- [49] S. F. Gull and J. Skilling, "The maximum entropy method," in *Indirect Imaging*, J. A. Roberts, Ed. Cambridge: Cambridge University Press, 1984, pp. 267-279.
- [50] J. Skilling and S. F. Gull, "Algorithms and applications," in *Maximum Entropy and Bayesian Methods in Inverse Problems - Wyoming 1982*, C. R. Smith and W. T. Grandy, Eds. Dordrecht: Kluwer, 1985, pp. 83-132.
- [51] J. Skilling, "The Cambridge maximum entropy algorithm," in *Maximum Entropy and Bayesian Methods in Applied Statistics*, J. H. Justice, Ed. Cambridge: Cambridge University Press, 1986, pp. 179-193.
- [52] N. Bleistein, *Mathematical Methods for Wave Propagation*, 1st ed. Orlando: Academic Press, 1984.

- [53] J. A. Ogilvy and I. D. Culverwell, "Elastic model for simulating ultrasonic inspection of smooth and rough defects," *Ultrasonics*, vol. 29, pp. 490-496, 1991.
- [54] T. Kreutter and K. J. Langenberg, "Multimode diffraction tomography with elastic waves," in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 9, D. O. Thompson and D. E. Chimenti, Eds. New York: Plenum Press, 1990, pp. 431-446.
- [55] S. D. Silverstein and L. J. Thomas, "Analytical comparison of sensor signal processing enhancements for NDT synthetic aperture ultrasonic imaging," *IEEE Transactions on Image Processing*, vol. 2, pp. 60-67, 1993.
- [56] A. C. Kak and M. Slaney, *Principles of Computerised Tomographic Imaging*, 1st ed. New York: IEEE, 1988.
- [57] R. C. Gonzalez and P. Wintz, *Digital Image Processing*, 2nd ed. Reading, Massachusetts: Addison-Wesley, 1987.
- [58] A. K. Jain, *Fundamentals of Digital Image Processing*, 1st ed. New Jersey: Prentice-Hall, 1989.
- [59] H. Lee and G. Wade, "Alternative viewpoint on wave propagation," *IEEE Transactions on Sonics and Ultrasonics*, vol. 30, pp. 331-332, 1983.
- [60] G. Prokoph and H. Ermert, "A comparison of broadband holographic and tomographic imaging concepts," in *Acoustical Imaging*, vol. 18, H. Lee and G. Wade, Eds. New York: Plenum Press, 1991, pp. 380-390.
- [61] A. J. Devaney, "A fast filtered backpropagation algorithm for ultrasound tomography," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 34, pp. 330-339, 1987.
- [62] R. Y. Chiao and L. J. Thomas, "Analytic evaluation of sampled aperture ultrasonic imaging techniques for NDE," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 41, pp. 484-493, 1994.
- [63] A. J. Devaney, "Inverse source and scattering problems in ultrasonics," *IEEE Transactions on Sonics and Ultrasonics*, vol. 30, pp. 355-364, 1983.
- [64] S. J. Norton and M. Linzer, "Ultrasonic reflectivity imaging in three dimensions: exact inverse scattering solutions for plane, cylindrical, and spherical apertures," *IEEE Transactions on Biomedical Engineering*, vol. 28, pp. 202-220, 1981.

- [65] A. Ishimaru, *Wave Propagation and Scattering in Random Media*, vol. 1: Single scattering and transport theory, 1st ed. New York: Academic Press, 1978.
- [66] J. F. Greenleaf, "Computerised tomography with ultrasound," *Proceedings of the IEEE*, vol. 71, pp. 330-337, 1983.
- [67] J. A. Ogilvy and J. A. G. Temple, "Diffraction of elastic waves by cracks: application to time-of-flight inspection," *Ultrasonics*, pp. 259-269, 1983.
- [68] K. Mayer, R. Marklein, K. J. Langenberg, and T. Kreutter, "Three-dimensional imaging system based on Fourier transform synthetic aperture focusing technique," *Ultrasonics*, vol. 28, pp. 241-255, 1990.
- [69] J. A. Ogilvy, "Computer simulation of acoustic wave scattering from rough surfaces," *Journal of Physics D: Applied Physics*, vol. 21, pp. 260-277, 1988.
- [70] H. P. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C*, 2nd ed. New York: Cambridge University Press, 1992.
- [71] S. Twomey, "On the numerical solution of Fredholm integral equations of the first kind by inversion of the linear systems produced by quadrature," *Journal of the Association for Computing Machinery*, vol. 10, pp. 97-101, 1963.
- [72] S. X. Pan and C. K. Avinash, "A computational study of reconstruction algorithms for diffraction tomography: Interpolation versus filtered backpropagation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 31, pp. 1262-1274, 1983.
- [73] A. J. Devaney, "Reconstructive tomography with diffracting wavefields," in *Inverse Problems*, vol. 2. Great Britain: IOP, 1986, pp. 161-183.
- [74] A. J. Witten, J. E. Molyneux, and J. E. Nyquist, "Ground penetrating radar tomography: algorithms and case studies," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, pp. 461-467, 1994.
- [75] K. T. Ladas and A. J. Devaney, "Generalized ART algorithm for diffraction tomography," in *Inverse Problems*, vol. 7: IOP Publishing Ltd., 1991, pp. 109-125.
- [76] W. C. Chew, *Waves and Fields in Inhomogeneous Media*. New York: van Nostrand Reinhold, 1990.
- [77] W. A. Veronesi and J. D. Maynard, "Nearfield acoustic holography (NAH) II. Holographic reconstruction algorithms and computer implementation," *Journal of The Acoustical Society of America*, vol. 81, pp. 1307-1322, 1987.

- [78] J. E. Molyneux and A. Whitten, "Diffraction tomographic imaging in a monostatic measurement geometry," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 31, pp. 507-511, 1993.
- [79] M. Soumekh, "Echo imaging with physical and synthesised arrays," *Optical Engineering*, vol. 29, pp. 545-554, 1990.
- [80] J. T. Ylitalo and H. Ermert, "Ultrasound synthetic aperture imaging: monostatic approach," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 41, pp. 333, 1994.
- [81] M. G. Silk, "The transfer of ultrasonic energy in the diffraction technique for crack sizing," *Ultrasonics*, pp. 113-121, 1979.
- [82] L. W. Schmerr and A. Sedov, "The flat-bottom hole: an ultrasonic scattering model," in *Research in Nondestructive Evaluation*, vol. 1. New York: Springer-Verlag, 1889, pp. 181-196.
- [83] D. P. Orofino and P. C. Pedersen, "Efficient angular spectrum decomposition of acoustic sources - Part I: Theory," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 40, pp. 238-249, 1993.
- [84] J. Hadamard, *Lectures on the Cauchy Problem in Linear Partial Differential Equations*. New Haven: Yale University Press, 1923.
- [85] J. Biemond, R. L. Lagendijk, and R. M. Mersereau, "Iterative methods for image deblurring," *Proceedings of the IEEE*, vol. 78, pp. 856-883, 1990.
- [86] A. J. Devaney and G. C. Sherman, "Nonuniqueness in inverse source and scattering problems," *IEEE Transactions on Antennas and Propagation*, vol. 30, pp. 1034-1041, 1982.
- [87] M. Fischer and K. Langenberg, "Limitations and defects of certain inverse scattering theories," *IEEE Transactions on Antennas and Propagation*, vol. 32, pp. 1080-1088, 1984.
- [88] N. N. Bojarski, "Comments on 'Nonuniqueness in Inverse Source and Scattering Problems'," *IEEE Transactions on Antennas and Propagation*, vol. 30, pp. 1037-1038, 1982.
- [89] W. R. Stone, "Comments on 'Nonuniqueness in Inverse Source and Scattering Problems'," *IEEE Transactions on Antennas and Propagation*, vol. 30, pp. 1039-1041, 1982.
- [90] C. W. Helstrom, *Statistical Theory of Signal Detection*, vol. 9, 1st ed. London: Pergamon, 1960.

- [91] Y. Aoki, "Optical and numerical reconstruction of images from sound-wave holograms," *IEEE Transactions on Audio and Electroacoustics*, vol. 18, pp. 258-267, 1970.
- [92] B. P. Hilderbrand, T. J. Davis, A. J. Boland, and R. L. Silta, "A portable digital ultrasonic holography system for imaging flaws in heavy section materials," *IEEE Transactions on Sonics and Ultrasonics*, vol. 31, pp. 287-294, 1984.
- [93] K. J. Langenberg, M. Fischer, M. Berger, and G. Weinfurter, "Imaging performance of generalised holography," *Journal of The Optical Society of America*, vol. 3, pp. 329-339, 1986.
- [94] R. K. Mueller, "Acoustic holography," *Proceedings of the IEEE*, vol. 59, pp. 1319-1971, 1971.
- [95] J. Ylitalo, E. Alasaarela, A. Tauriainen, K. Tervola, and J. Koivukangas, "Three-dimensional ultrasound C-scan imaging using holographic reconstruction," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 33, pp. 731-739, 1986.
- [96] L. J. Busse, "Three dimensional imaging using a frequency domain synthetic aperture focussing technique," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, pp. 174-179, 1992.
- [97] R. M. Leahy and C. E. Goutis, "An optimal technique for constraint-based image restoration and reconstruction," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 34, pp. 1629-1642, 1986.
- [98] M. K. Nguyen and A. Mohammad-Djafari, "Bayesian approach with the maximum entropy principle in image reconstruction from microwave scattered field data," *IEEE Transactions on Medical Imaging*, vol. 13, pp. 254-262, 1994.
- [99] A. Mohammad-Djafari and G. Demoment, "Image restoration and reconstruction using entropy as a regularization functional," in *Maximum Entropy and Bayesian Methods in Science and Engineering*, vol. 2, *Fundamental Theories of Physics*, G. J. Erickson and C. R. Smith, Eds. Dordrecht: Kluwer, 1988, pp. 341-355.
- [100] B. Borden, "Maximum entropy regularization in inverse synthetic aperture radar imagery," *IEEE Transactions on Signal Processing*, vol. 40, pp. 969-973, 1992.

- [101] G. Landl and R. S. Anderssen, "Non-negative differentially constrained entropy-like regularization," *Inverse Problems*, vol. 12, pp. 35-53, 1996.
- [102] D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 413-424, 1986.
- [103] A. H. Delaney and Y. Bresler, "Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography," *IEEE Transactions on Image Processing*, vol. 7, pp. 204-221, 1998.
- [104] C. Bouman and K. Sauer, "A generalised Gaussian image model for edge-preserving MAP estimation," *IEEE Transactions on Image Processing*, vol. 2, pp. 296-310, 1993.
- [105] L. Bedini and A. Tonazzini, "Image restoration preserving discontinuities: the Bayesian approach and neural networks," *Image and Vision Computing*, vol. 10, pp. 108-118, 1992.
- [106] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 721-741, 1984.
- [107] R. L. Lagendijk and J. Biemond, *Iterative Identification and Restoration of Images*, 1st ed. Boston: Kluwer, 1991.
- [108] M. I. Sezan and A. M. Tekalp, "Survey of recent developments in digital image restoration," *Optical Engineering*, vol. 29, pp. 393-404, 1990.
- [109] Y. M. Wang and W. C. Chew, "Reconstruction of two-dimensional refractive index distribution using the Born iterative and distorted Born iterative method," in *Acoustical Imaging*, vol. 18, H. Lee and G. Wade, Eds. New York: Plenum Press, 1991, pp. 105-114.
- [110] J. Lin and W. C. Chew, "Ultrasonic imaging by local shape function method with CGFFT," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, pp. 956-969, 1996.
- [111] J. L. Sutton, "Underwater acoustic imaging," *Proceedings of the IEEE*, vol. 67, pp. 554-566, 1979.
- [112] B. R. Hunt, "The application of constrained least squares estimation to image restoration by digital computer," *IEEE Transactions on Computers*, vol. 22, pp. 805-812, 1973.

- [113] D. L. Phillips, "A technique for the numerical solution of certain integral equations of the first kind," *Journal of the Association for Computing Machinery*, vol. 9, pp. 84-97, 1962.
- [114] M. R. Banham and A. K. Katsaggelos, "Digital Image Restoration," in *IEEE Signal Processing Magazine*, vol. 14, 1997, pp. 24-41.
- [115] R. Cristi, "Markov and recursive least squares methods for the estimation of data with discontinuities," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, pp. 1972-1980, 1990.
- [116] M. E. Zervakis, "Edge-preserving image restoration based on robust stabilizing functionals," *SPIE Applications of digital image processing*, vol. 1657, pp. 170-180, 1992.
- [117] P. Charbonnier, L. Blanc-Feraud, and M. Barlaud, "An adaptive reconstruction method involving discontinuities," presented at ICASSP, Minneapolis, MN, 1993.
- [118] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, pp. 629-639, 1990.
- [119] S. Teboul, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Variational approach for edge-preserving regularization using coupled PDE's," *IEEE Transactions on Image Processing*, vol. 7, pp. 387-397, 1998.
- [120] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 367-383, 1992.
- [121] A. Blake and A. Zisserman, *Visual Reconstruction*. Cambridge, MA: MIT Press, 1987.
- [122] C. R. Vogel and M. E. Oman, "Fast numerical methods for total variation minimization in image reconstruction," presented at SPIE Advanced Signal Processing Algorithms, San Diego, CA, 1995.
- [123] T. F. Chan and C. K. Wong, "Total variation blind deconvolution," *IEEE Transactions on Image Processing*, vol. 7, pp. 370-375, 1998.
- [124] C. L. Byrne, "Iterative image reconstruction algorithms based on cross-entropy minimization," *IEEE Transactions on Image Processing*, vol. 2, pp. 96-103, 1993.

- [125] B. R. Hunt, "Bayesian methods in non-linear digital image restoration," *IEEE Transactions on Computers.*, vol. 26, pp. 219-229, 1977.
- [126] W. H. Richardson, "Bayesian based iterative method of image restoration," *Journal of The Optical Society of America*, vol. 62, pp. 55-59, 1972.
- [127] G. Demoment, "Image reconstruction and restoration," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, pp. 2024-2036, 1989.
- [128] A. Klug and R. A. Crowther, "Three-dimensional image reconstruction from the viewpoint of information theory," *Nature*, vol. 238, pp. 435-440, 1972.
- [129] J. Skilling, "Theory of maximum entropy image reconstruction," in *Maximum Entropy and Bayesian Methods in Applied Statistics*, J. H. Justice, Ed. Cambridge: Cambridge University Press, 1984, pp. 156-178.
- [130] J. Bernoulli, *The Art of Conjecture*, 1713.
- [131] E. T. Jaynes, "Prior probabilities," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, pp. 227-241, 1968.
- [132] T. Bayes, "An essay towards solving a problem in the doctrine of chances," *Philosophical Transactions of the Royal Society of London*, vol. 53, 1763.
- [133] P. S. Laplace, *Theorie Analytique des Probabilités*. Paris: Courcier, 1812.
- [134] J. K. O Ruanaidh and W. J. Fitzgerald, *Numerical Bayesian Methods Applied to Signal Processing*, 1st ed. New York: Springer-Verlag, 1996.
- [135] T. Hebert and R. Leahy, "A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors," *IEEE Transactions on Medical Imaging*, vol. 8, pp. 194-202, 1989.
- [136] D. S. Lalush and B. M. W. Tsui, "Simulation evaluation of Gibbs prior distributions for use in maximum *a posteriori* SPECT reconstructions," *IEEE Transactions on Medical Imaging*, vol. 11, pp. 267-275, 1992.
- [137] Z. Zhou, R. M. Leahy, and J. Qi, "Approximate maximum likelihood hyperparameter estimation for Gibbs priors," *IEEE Transactions on Image Processing*, vol. 6, pp. 844-861, 1997.
- [138] M. Nikolova, J. Idier, and A. Mohammad-Djafari, "Inversion of large-support ill-posed linear operators using a piecewise Gaussian MRF," *IEEE Transactions on Image Processing*, vol. 7, pp. 571-585, 1998.
- [139] D. M. Titterton, "General structure of regularisation procedures in image reconstruction," *Astronomy and Astrophysics*, vol. 144, pp. 381-387, 1985.

- [140] S. F. Gull, "Developments in maximum entropy data analysis," in *Maximum Entropy and Bayesian Methods, Fundamental Theories of Physics*, J. Skilling, Ed. Dordrecht: Kluwer, 1989, pp. 525.
- [141] B. R. Frieden, "Restoring with maximum likelihood and maximum entropy," *Journal of The Optical Society of America*, vol. 62, pp. 511-518, 1972.
- [142] T. J. Ponman, "Maximum entropy methods," *Nuclear Instruments and Methods in Physics Research*, vol. 221, pp. 72-76, 1984.
- [143] R. Nityananda and R. Narayan, "Maximum entropy image reconstruction - a practical non-information theoretic approach," *Journal of Astrophysics and Astronomy*, vol. 3, pp. 419-450, 1982.
- [144] J. A. Högbom, "The introduction of a priori knowledge in certain processing problems," in *Image Formation from Coherence Functions in Astronomy*. Groningen: D. Reidel Publishing Company, 1978, pp. 417-426.
- [145] J. Skilling, "The axioms of maximum entropy," in *Maximum-entropy and Bayesian Methods in Science and Engineering*, vol. 1: Foundations, G. J. Erickson and C. R. Smith, Eds. Dordrecht: Kluwer, 1988, pp. 173-186.
- [146] J. E. Shore and R. W. Johnson, "Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy," *IEEE Transactions on Information Theory*, vol. 26, pp. 26-37, 1980.
- [147] S. Kullback and R. Leibler, "On information and sufficiency," *The Annals of Mathematics and Statistics*, vol. 22, pp. 79-86, 1951.
- [148] S. Kullback, *Information Theory and Statistics*. New York: Wiley, 1959.
- [149] G. J. Daniell, "Of maps and monkeys: an introduction to the maximum entropy method," in *Maximum entropy in action*, B. Buck and V. A. Macaulay, Eds. Oxford: Clarendon Press, 1991, pp. 1-18.
- [150] I. Csiszár, "Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems," *The Annals of Statistics*, vol. 19, pp. 2032-2066, 1991.
- [151] S. F. Gull and J. Skilling, "The entropy of an image," in *Maximum Entropy and Bayesian Methods in Inverse Problems - Wyoming 1982*, C. R. Smith and W. T. Grandy, Eds. Dordrecht: Kluwer, 1985, pp. 287-301.
- [152] E. T. Jaynes, "On the rationale of maximum-entropy methods," *IEE Proceedings*, vol. 70, pp. 939-952, 1982.

- [153] A. K. Livesey and J. Skilling, "Maximum entropy theory," *Acta Crystallographica*, vol. A41, pp. 113-122, 1985.
- [154] E. T. Jaynes, "Where do we go from here?," in *Maximum-Entropy and Bayesian Methods in Inverse Problems*, C. R. Smith and W. T. Grandy, Eds. Dordrecht: D. Reidel Publishing Company, 1985, pp. 21-58.
- [155] E. T. Jaynes, "Monkeys, kangaroos, and N," in *Maximum Entropy and Bayesian Methods in Applied Statistics*, J. H. Justice, Ed. Cambridge: Cambridge University Press, 1986, pp. 26-58.
- [156] D. S. Sivia, *Data Analysis - A Bayesian Tutorial*, 1st ed. Oxford: Clarendon Press, 1996.
- [157] R. W. Johnson and J. E. Shore, "Comments on and correction to "axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy," *IEEE Transactions on Information Theory*, vol. 29, pp. 942-943, 1983.
- [158] Y. Tikochinsky, N. Z. Tishby, and R. D. Levine, "Consistent inference of probabilities for reproducible experiments," *Physical Review Letters*, vol. 52, pp. 1357-1360, 1984.
- [159] R. K. Bryan, "Maximum entropy and the phase problem in protein crystallography," in *Maximum Entropy and Bayesian Methods in Science and Engineering. Volume 2: Applications*, G. J. Erickson and C. R. Smith, Eds. Dordrecht: Kluwer, 1988, pp. 155-169.
- [160] T. J. Cornwell and K. F. Evans, "A simple maximum entropy deconvolution algorithm," *Astronomy and Astrophysics*, vol. 143, pp. 77-83, 1985.
- [161] J. A. Högbom, "Clean as a pattern recognition procedure," in *Indirect Imaging*, J. A. Roberts, Ed. Cambridge: Cambridge University Press, 1984, pp. 247-254.
- [162] D. S. Briggs, "Superresolution of SN 1987A: A comparative deconvolution study," presented at the second conference on restoration of HST images and spectra, Space Telescope Science Institute, 1994.
- [163] D. L. Donoho, I. M. Johnstone, J. C. Hoch, and A. S. Stern, "Maximum entropy and the nearly black object," *Journal of the Royal Statistical Society B*, vol. 54, pp. 41-81, 1992.

- [164] N. Agmon, Y. Alhassid, and R. D. Levine, "An algorithm for finding the distribution of maximal entropy," *Computational Physics*, vol. 30, pp. 250-258, 1979.
- [165] R. Wilczek and S. Drapatz, "A high accuracy algorithm for maximum entropy image restoration in the case of small data sets," *Astronomy and Astrophysics*, vol. 142, pp. 9-12, 1985.
- [166] S. F. Gull and J. Skilling, "MEMSYS5 User's Manual," Maximum Entropy Data Consultants 1991.
- [167] V. A. Macaulay, "Bayesian inversion with applications to physics," PhD Thesis, Oxford University, 1992.
- [168] J. D. Jackson, *Classical Electrodynamics*, 2nd ed. New York: John Wiley & Sons, 1975.
- [169] B. Buck and V. A. Macaulay, "Maximum Entropy in Action," Oxford: Clarendon Press, 1991, 220.
- [170] K. M. Wright and P. S. Belton, "A new definition of the information content of NMR spectra suitable for use in maximum entropy signal processing," *Molecular Physics*, vol. 58, pp. 485-495, 1986.
- [171] J. C. Hoch, A. S. Stern, D. L. Donoho, and I. M. Johnstone, "Maximum entropy reconstruction of complex (phase-sensitive) spectra," *Journal of Magnetic Resonance*, vol. 86, pp. 236-246, 1990.
- [172] P. F. Smith, M. A. Player, and D. A. L. Collie, "The performance of the maximum entropy method: deconvolution and the frequency content of data," *Journal of Physics D: Applied Physics*, vol. 22, pp. 906-914, 1989.
- [173] P. F. Smith, M. A. Player, and D. A. L. Collie, "The determination of surface topology by the signal processing of ultrasonic pulses," *Journal of Physics. E: Scientific Instrumentation*, vol. 21, pp. 397-402, 1988.
- [174] P. F. Smith and M. A. Player, "The application of maximum entropy signal processing to ultrasonic surface parameterisation," in *Maximum Entropy and Bayesian Methods - Wyoming 1990, Fundamental Theories of Physics*, W. T. Grandy and L. H. Schick, Eds. Dordrecht: Kluwer, 1991, pp. 137-146.
- [175] A. Mohammad-Djafari and G. Demoment, "Maximum entropy Fourier synthesis with application to diffraction tomography," *Applied Optics*, vol. 26, pp. 1745-1754, 1987.

- [176] M. Baribaud, "Microwave imagery: analytical method and maximum entropy method," *Journal of Physics D: Applied Physics*, vol. 23, pp. 269-288, 1990.
- [177] A. T. Bajkova, "The generalisation of the maximum entropy method for the reconstruction of complex functions," *Astronomical and Astrophysical Transactions*, vol. 1, pp. 313-320, 1992.
- [178] B. R. Frieden and A. T. Bajkova, "Bayesian cross-entropy reconstruction of complex images," *Applied Optics*, vol. 33, pp. 219-226, 1994.
- [179] Heidbreder, "Maximum entropy applications in radar," in *Maximum Entropy and Bayesian Methods, Fundamental Theories of Physics*, W. T. Grandy and L. H. Schick, Eds. Dordrecht: Kluwer, 1991, pp. 127-136.
- [180] D. J. Battle, R. P. Harrison, and M. Hedley, "Maximum entropy image reconstruction from sparsely sampled coherent field data," *IEEE Transactions on Image Processing*, vol. 6, pp. 1139-1147, 1997 (See also Appendix D).
- [181] B. R. Frieden and A. T. Bajkova, "Reconstruction of complex signals using minimum Rényi information," *Applied Optics*, vol. 34, pp. 4086-4093, 1995.
- [182] S. Wolfram, "Mathematica Version 3.0," . Champaign Il.: Wolfram Research, 1996.
- [183] R. N. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed. New York: McGraw-Hill, 1978.
- [184] S. Ganapathy, W. S. Wu, and B. Schmolt, "Analysis and design considerations for a real-time system for non-destructive evaluation in the nuclear industry," *Ultrasonics*, pp. 249-256, 1982.
- [185] M. G. Silk, "The interpretation of TOFD data in the light of ASME XI and similar rules," *British Journal of NDT*, vol. 31, pp. 242-251, 1989.
- [186] S. F. Burch, "Comparison of SAFT deconvolution methods for the improvement of ultrasonic B-Scan images," *Ultrasonics*, vol. 25, pp. 259-265, 1987.
- [187] S. F. Burch and A. T. Ramsey, "Digital techniques to derive and display the phases of ultrasonic pulses," *British Journal of NDT*, vol. 31, pp. 185-188, 1989.
- [188] J. D. Plummer, R. G. Swartz, M. G. Maginness, J. R. Beaudouin, and J. D. Meindl, "Two-dimensional transmit/receive ceramic piezoelectric arrays: Construction and performance," *IEEE Transactions on Sonics and Ultrasonics*, vol. 25, pp. 273-280, 1978.

- [189] M. Greenstein, P. Lum, H. Yoshida, and M. S. Seyed-Bolorforosh, "A 2.5 MHz 2D array with z-axis electrically conductive backing," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, pp. 970-977, 1997.
- [190] R. E. Davidsen and S. W. Smith, "Two-dimensional arrays for medical ultrasound using multilayer flexible circuit interconnection," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 45, pp. 338-348, 1998.
- [191] S. W. Smith and E. D. Light, "Two-dimensional transducers using hybrid connection technology," presented at the IEEE Ultrasonics Symposium, 1992.
- [192] D. H. Turnbull and S. F. Foster, "Fabrication and characterisation of transducer elements in two-dimensional arrays for medical ultrasound imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, pp. 464-475, 1992.
- [193] R. L. Goldberg and S. W. Smith, "Multilayer piezoelectric ceramics for two-dimensional array transducers," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 41, pp. 761-771, 1994.
- [194] S. W. Smith, H. G. Pavy, and O. T. von Ramm, "High-speed ultrasound volumetric imaging system- Part I: Transducer design and beam steering," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 38, pp. 100-108, 1991.
- [195] S. Holm and G. Dahl, "Properties of the beampattern of weight- and layout-optimized sparse arrays," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, pp. 983-991, 1997.
- [196] D. N. Hall, G. Hayward, and Y. Gorfu, "Theoretical and experimental evaluation of a two-dimensional composite matrix array," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 40, pp. 704-709, 1993.
- [197] O. T. von Ramm, S. W. Smith, and H. G. Pavy, "High-speed ultrasound volumetric imaging system- Part II: Parallel Processing and Image Display," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 38, pp. 109-114, 1991.

- [198] H. C. Schau, "Edge-connected, crossed-electrode array for two-dimensional projection and beamforming," *IEEE Transactions on Signal Processing*, vol. 39, pp. 289-297, 1991.
- [199] L. Daane and M. Greenstein, "A demountable interconnect system for a 50×50 ultrasonic imaging transducer array," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, pp. 978-982, 1997.
- [200] C. G. Oakley, "Calculation of ultrasonic transducer signal-to-noise ratios using the KLM model," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, pp. 1018-1026, 1997.
- [201] P. Horowitz and W. Hill, *The Art of Electronics*, 2nd ed. Cambridge: Cambridge University Press, 1991.
- [202] AMP, "Piezo Film Sensors Product Guide," Harrisburg, Pensilvania. Catalog 65711, 1995.
- [203] L. F. Brown and J. L. Mason, "Disposable PVDF ultrasonic transducers for nondestructive testing applications," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, pp. 560-568, 1996.
- [204] D. J. Powell and G. Hayward, "Flexible ultrasonic transducer arrays for nondestructive evaluation applications - Part II: Performance assessment of different array configurations," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, pp. 393-402, 1996.
- [205] D. J. Powell and G. Hayward, "Flexible ultrasonic transducer arrays for nondestructive evaluation applications - Part I: The theoretical modelling approach," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, pp. 385-391, 1996.
- [206] P. D. Wilcox, R. S. C. Monkhouse, P. Cawley, M. J. S. Lowe, and B. A. Auld, "Development of a computer model for an ultrasonic polymer film transducer system," *NDT International*, vol. 31, pp. 51-64, 1998.
- [207] R. G. Swartz and J. D. Plummer, "On the generation of high-frequency acoustic energy with polyvinylidene flouride," *IEEE Transactions on Sonics and Ultrasonics*, vol. SU-27, pp. 295-303, 1980.
- [208] F. K. Lam, "Microcomputer-based digital pulse compression system for ultrasonic NDT," *Ultrasonics*, vol. 25, pp. 166-171, 1987.
- [209] G. Hayward and Y. Gorfu, "Low power ultrasonic imaging using digital correlation," *Ultrasonics*, vol. 27, pp. 288-296, 1989.

- [210] M. O'Donnell, "Coded excitation system for improving the penetration of real-time phased array imaging systems," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, pp. 341-351, 1992.
- [211] G. Landl and R. S. Anderssen, "On the Utility of Non-Negative Differentially Constrained Entropy-Like Regularizers," Australian National University, Mathematics Research Report MRR 045-95, July 1995.
- [212] T. K. Moon, "The expectation - maximization algorithm," in *IEEE Signal Processing Magazine*, vol. 13, 1996, pp. 47-60.