

IMPERIAL COLLEGE LONDON

**TOWARDS NEXT GENERATION
ULTRASONIC IMAGING**

by

Tim Hutt

A thesis submitted to Imperial College London for the degree of
Doctor of Philosophy

Department of Mechanical Engineering
Imperial College London
London SW7 2AZ

November 2011

Abstract

Recently the use of ultrasonic arrays for imaging defects in metal components has become economically attractive in Non-Destructive Testing. Given a certain array, the image quality strongly depends on how the measurements are processed into an image. The current state-of-the-art imaging algorithm in actual use is delay-and-sum beamforming, which has a resolution capability that is fundamentally limited by the physical approximation used to describe how waves interact with matter.

This thesis explores the practical use of alternative non-linear “super-resolution” imaging algorithms that use more accurate physical models, and can theoretically achieve unlimited resolution. This is made possible by utilising additional sources of information contained within the measurements, in particular the small amplitude multiply scattered signals.

The distribution of information contained in the measurements, and utilised by the imaging algorithms is studied in the context of information capacity of signals. We discover some insights into the limits of imaging which depend on the signal-to-noise ratio.

The accuracy of non-linear imaging algorithms can be strongly dependent on the accuracy of the measurements. Therefore several experiments are performed to assess their performance in practice. The experimental implementation of these methods poses a number of challenges, including removal of the incident field, and compensating for array element directivity.

Super-resolution capability is demonstrated in a highly attenuative medium for the first time. To further improve the image quality we explore the possibility of using mirror reflections. This gives an increase in the effective aperture. We perform simulated and experimental reconstructions of a complex scatterer and find that the completeness of the image is improved.

The mirror interface also allows quantitative speed-of-sound imaging of penetrable scatterers using the HARBUT algorithm. This is tested experimentally for the first

time.

Acknowledgements

I would like to thank Dr Simonetti, and Professor Cawley for their expert guidance and infinite patience! This thesis would have been impossible without their support and encouragement.

I would also like to thank the other members of the NDT lab for making it such a friendly and sociable working environment. Especially to Remo, Rosalba and Andrea for introducing me to Italian customs, culture and time keeping; to Vatche for his irreverent musings, and to Robin, Peter H. and Andrew for keeping the lab sane.

Finally, thanks to the RCNDE and EPSRC for funding this project.

Contents

1	Introduction	28
1.1	Background	28
1.2	Thesis Plan	37
2	Formulation of the Imaging Problem	39
2.1	Background	39
2.2	Forward acoustic scattering	40
2.2.1	Helmholtz Decomposition	41
2.2.2	Solution of the Acoustic Wave Equation	44
2.2.3	The Object Function	45
2.2.4	Far-field scattering	48
2.2.5	Spaces and Operators	51
2.3	Inverse scattering	54
2.3.1	Classical Imaging Methods	55
2.3.2	Diffraction Tomography	56
2.3.3	Beamforming	60

2.4	Resolution Limits	63
2.5	Super Resolution	68
2.6	Conclusion	71
3	The Discrete Inversion	73
3.1	Background	73
3.2	Degrees of freedom of an image	74
3.2.1	Discretisation of the far-field operator	76
3.3	Spatial harmonics	78
3.3.1	Dynamic range	81
3.4	Limited apertures	84
3.5	Temporal diversity	85
3.6	Broadband array measurements and conclusions	88
4	Broadband Super-Resolution	92
4.1	Background	92
4.2	Monochromatic sampling methods	93
4.2.1	The Linear Sampling Method	93
4.2.2	Factorisation of the Far-field Operator	96
4.3	Broadband Linear Sampling Method	100
4.3.1	Pre-beamforming	105
4.4	Experimental validation of Time-Domain LSM	108

4.5	Results and Discussion	109
4.6	Conclusion	111
5	Broadband super resolution in elastic solids	114
5.1	Background	115
5.2	Simulating the Effect of Gating	118
5.2.1	Simulations and Results	120
5.2.2	Directionality	122
5.2.3	Super-Resolution	123
5.3	Experiments	124
5.3.1	Experimental Setup	124
5.3.2	Experimental procedure	129
5.3.3	Coupling compensation	129
5.3.4	Incident field removal	132
5.3.5	Images and discussion	133
5.4	Conclusion	137
6	Using mirror reflections for imaging—Theory	138
6.1	Background	138
6.2	Information content of the reflection	141
6.2.1	Singular Value Decomposition of the Far-Field Operator	146
6.3	Imaging impenetrable scatterers	148

6.3.1	Beamforming	148
6.3.2	Factorisation Method	150
6.3.3	Numerical Simulations	151
6.3.4	Noise Analysis	156
6.4	Imaging penetrable scatterers	160
6.5	Conclusions	162
7	Using mirror reflections for imaging—Experiments	164
7.1	Background	164
7.2	Practical considerations	165
7.2.1	Incident field removal	165
7.2.1.1	Gating	166
7.2.1.2	Field-fitting	166
7.2.2	Beam directionality	169
7.3	Impenetrable scatterer experiments	172
7.4	Penetrable Scatterer experiments	174
7.5	Conclusions	175
8	Conclusions	178
8.1	Theoretical conclusions	179
8.2	Practical progress	180
8.3	Future direction	181

Nomenclature

$\mathbf{x}, \mathbf{r}, \mathbf{p}, \mathbf{x}'$ etc.	Position vectors. Vectors are indicated with a bold typeface.
$\hat{\mathbf{x}}$	A normalised vector.
$\boldsymbol{\sigma}, \boldsymbol{\epsilon}$	Stress and strain tensors.
p	Pressure.
t	Time.
$H(\cdot)$	Heaviside step function.
$H_n^{(1)}(\cdot)$	Hankel function of the first kind and order n .
ω	Angular frequency (rad/s).
$u(\mathbf{x})$ or $u(\mathbf{x}, t)$	Total wave field (often only the spatial varying component of a single frequency coefficient).
$u^i(\mathbf{x})$ or $u^i(\mathbf{x}, t)$	Incident wave field.
$u^s(\mathbf{x})$ or $u^s(\mathbf{x}, t)$	Scattered wave field.
$u_\infty(\mathbf{x})$ etc.	A far-field wave field. (the angular component of a field measured at an infinite distance from the origin).
$\Gamma(\mathbf{x})$	Fixed radial component of a far-field wave field.
$o(\mathbf{x})$	Object function.
k	Angular wavenumber (radians per distance).
$G(\mathbf{x}, \mathbf{x}')$ or $ g_z\rangle$	Green's function.
\mathcal{D}	Scatterer support.

$\partial\mathcal{D}$	Boundary of the scatterer.
F	The far-field operator.
$v(\hat{\mathbf{r}}, \hat{\mathbf{x}}),$ $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k),$ $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t)$ etc.	The scattering amplitude.
M	Imaging operator.
H	Herglotz operator.
S	Single layer operator.
\mathcal{F}	Fourier transform.
$\mathcal{O}(\mathbf{X})$	Fourier transform of the object function, at spatial frequency \mathbf{X} .
σ_n, ξ_n, ψ_n	Singular values and vectors of the far-field operator.
$\frac{s}{\varsigma}$	Signals (s) to noise (ς) ratio.
ρ	Radius of circle (or sphere) enclosing scatterer.
I	Information capacity.
$\mathcal{I}(\mathbf{x})$	Image value at \mathbf{x} .

Acronyms

BF	Beamforming
DT	Diffraction Tomography
FM	Factorisation Method
HARBUT	Hybrid Algorithm for Robust Breast Ultrasound Tomography
LSM	Linear Sampling Method
TDLSM	Time Domain Linear Sampling Method
ToFD	Time of Flight Diffraction
ToFT	Time of Flight Tomography

List of Figures

- 1.1 Images of a dragonfly that are a) ideal, b) low resolution, c) noisy, and d) low contrast. Certain features, such as its legs and the pattern on its wings, are rendered invisible by the degradation. Photo by Billy Lindblom. 31
- 1.2 Typical construction of 4-element linear array. A composite of PZT pillars in a polymer matrix is sandwiched between element electrodes. Each array element is made from several pillars vibrating in concert. . 32
- 1.3 A schematic showing the process of generating and receiving ultrasound for one channel of an array. The input generator can be a pulse system, or an arbitrary function generator and power amplifier. This signal is fed to the array element which transmits the wave. The resulting weak reflected signal is amplified by a high gain amplifier and digitised by an ADC. 33
- 1.4 In confocal microscopy (a) a lens is used to focus light at a point, and also used to detect reflections from that point. By scanning the focal point an image can be built. Ultrasonic beamforming (b) uses a similar approach, but an array of elements is used to transmit and receive the focused waves, and usually the same array is used for transmission and reception, although this need not be the case 34

1.5 Common array geometries. a) 1D or Linear; b) “1.5D” or Hybrid; c) 2D or Matrix; d) Annular. Although there are some common configurations, essentially any geometry can be manufactured. Arrays also do not necessarily have flat faces, they can be curved or even circular. 34

1.6 A medical ultrasound image of the jugular (a), compared with an NDE image of three side-drilled holes in a stainless steel block (b). The structure of the tissue and blood vessels in (a) can be seen due to the speckle contrast, which cannot be seen in (b). Additionally, the impenetrable nature of the holes results in shadowing of the backwall. 35

2.1 a) In the far-field scattering problem an incident wave $u_{\infty}^i(\hat{\mathbf{x}})$ illuminates the scatterer. b) The incident wave interacts with the scatterer, which is labelled \mathcal{D} with boundary $\partial\mathcal{D}$, and produces a scattered wave $u_{\infty}^s(\hat{\mathbf{x}})$. The far-field scattering problem is to determine $u_{\infty}^s(\hat{\mathbf{x}})$ given $u_{\infty}^i(\hat{\mathbf{x}})$ 48

2.2 The paraxial approximation. When $|\mathbf{x}|$ is very large, the component of \mathbf{x}' perpendicular to \mathbf{x} can be ignored when calculating the distance from \mathbf{x}' to \mathbf{x} , therefore $|\mathbf{x} - \mathbf{x}'| \approx |\mathbf{x}| - \frac{\mathbf{x}' \cdot \mathbf{x}}{|\mathbf{x}|}$ 50

2.3 a) The far-field operator, F maps any incident far-field to the corresponding scattered far-field for a particular scatterer. b) The M operator maps the scatterer (represented by o or in the frequency domain by \mathcal{O}) to the *complete set* of far-field measurements v , and therefore also F which depends only on v 52

2.4 Factorisation of the far-field operator. 54

2.5 The Ewald circle and the Ewald limiting circle. Outside the Ewald limiting circle, spatial frequencies of the scatterer cannot be measure using DT. But within it, every point can be measured by illuminating and observing the scatterer from many directions $\hat{\mathbf{r}}$ and $\hat{\mathbf{x}}$ 58

2.6	Scattering measurements and interpolation in the k-space for diffraction tomography. See the main text for a description.	59
2.7	An example of time-domain beamforming. The images show successive points in time as a wave is transmitted from a line of point sources and converges on a focal point. The transmission by each point on the line is delayed and apodized in order to produce a wave that converges on the focal point in the forth image. The delays are directly calculated from the distance from the focal point to each source point.	62
2.8	Scattering from independent scatterers to the far-field depends only on the distances a_i and b_i (and the scattering coefficients of the scatterers). See the main text for a complete description.	64
2.9	The first 100 singular values corresponding to scattering from the sound-soft “kite” and circle scatterer shown. It can be seen that the singular values quickly become extremely small. Theoretically the downward trend would continue, however the addition of noise raises the singular values to an approximately constant level. In this case the noise is actually caused by the finite numerical precision.	66
2.10	The magnitude of the incident fields produced by some right singular vectors ($\langle \psi_n $) for $n = 1, 2, 20, 30, 40, 60$. The higher magnitude singular values produce fields with higher incident values on the scatterer surface which therefore scatter strongly. The magnitude of reflection is determined by the singular value, so zero singular values produce non-scattering fields. Note that the singular vectors are normal, so the power of the transmission is constant.	66
2.11	The real component of the incident field on the surface of the scatterer for fields shown in figure 2.10. Note that the scale of each sub-plot is different, and that the fields are more rapidly varying for small singular values. They also tend to have more energy concentrated within the occluded area between the two scatterers.	67

2.12 A demonstration of the formation of evanescent waves. a) A single harmonic point source generates an expanding circular wave with the labelled wavelength. b) The single point is replaced with a continuous line of point sources which have phases that vary faster than the natural wavelength in the medium. This generates a near-field of evanescent waves which cannot propagate. The rapid phasing of the sources mean that the waves they produce sum destructively in the far-field. 69

2.13 The mechanism of NSOM. An evanescent field is produced in the background medium. If another medium with a low enough sound speed c_1 is introduced into the evanescent field, it will produce radiating waves that can propagate towards a sensor. In NSOM the probe is much narrower than shown here—less than a wavelength—so that high wavenumber evanescent fields can be measured. 70

2.14 An illustration of multiple scattering between two circular scatterers. A harmonic wave is incident from above the scatterers. In (a) the scattered field from each circle is calculated separately, and then the fields are summed. In (b) the correct scattered wave is found by considering the two circles together. (c) shows the difference between (a) and (b) which corresponds to the waves that scatter multiple times between the circles. It can be seen that the multiply scattered waves originate between the scatterers, as would be expected, and that they have a significant effect on the value of the scattered wave field. . . . 71

3.1 A plot of $|\alpha_n|^2$ for several values of ρk . The points $n = \rho k$ where α_n starts to become exponentially small are marked. The negative harmonics (not shown) are have symmetric values (that is, $|\alpha_n|^2 = |\alpha_{-n}|^2$). 81

3.2 Total information capacity of a measurement for various $k\rho$ at $\frac{s}{\zeta} = 10$ for different numbers of transducers, N_x 83

3.3 In most practical situations the full circle (or sphere) cannot be measured. Instead a subset of angles $[0, \delta)$ are accessible. 83

3.4 Singular values as a function of aperture for $\rho k = 81$ (compare the $\delta = 2\pi$ plot with $\rho k = 81$ in figure 3.1). The number of accessible singular values varies linearly with aperture. 85

3.5 Continuous Wavelet Transform (CWT) of the plotted sample signal into different bases on the time-frequency continuum. In (a) the basis wavelets are very short and broadband, so the signal is localised well in the time domain. The basis functions are increased in length and narrow in bandwidth through (a)–(f) so that (f) is almost equivalent to the Fourier transform of the signal. The energy is most concentrated in the fewest coefficients somewhere between these extremes, at approximately (c). The plotted Morlet wavelets A-D are the wavelets used to compute (a)–(d) at the centre frequency. Note that the CWT is used rather than the DWT merely to aid visualisation (see [1]). . . 87

3.6 Additional sources of data in a linear array measurement. Existing single-frequency super-resolution methods only use data from the region marked with a white line before the backwall reflection. However, useful signal energy is received after this, and also in other frequencies. 91

4.1 The Linear Sampling Method relies on the fact that it is possible to create field values on the boundary of a scatterer, such that a field that appears to radiate only from a single point within the scatterer. However it is not possible to do the same for a point outside the scatterer. 94

4.2 A demonstration of the frequency-domain sampling methods. a) The far-field operator F is measured for a kite-shaped sound-soft scatterer (i.e. a hole). The axes are in units of λ (or equivalently $\lambda = 1$). b) The resulting LSM image, clipped at an ad-hoc threshold to make the scatterer visible. c) The FM image. No clipping was applied in this case. F was calculated using a circular array of 256 transducers with a radius of 10, using the Boundary Element Method (BEM). The formula for the kite-shaped scatterer is given in [2]. 97

4.3 The values $|\langle \xi_n | G_{\mathbf{z}} \rangle|^2$, for the first, last, and some intermediate n 's, and the cumulative sum for LSM and FM. 99

4.4 Non-causality of Fourier solutions to deconvolution. See the text for a description. 102

4.5 Structure of the discrete time-domain far-field operator $F_{t,d}$ in block Toeplitz matrix form for a) a single transducer and b) three transducers. The exact clipping of the time signals does not necessarily have to allow for the complete convolution of the input and output, nor do the input and output have to be the lengths given here. For example we could use only the top half of each block if desired. In all the TDLSM images in this thesis, the structure was as shown here, with each block having a dimension of $N_t \times (2N_t - 1)$ 104

4.6 a) Time-domain beamforming, b) FM, and c) TDLSM imaging of two squares beneath 8 transducers. The wavelength is $\frac{\pi}{2}$. It can be seen that both sampling methods resolve the scatterers, whereas BF does not. 107

4.7 Measurement of the attenuation of ultrasound in glycerol was achieved by transmitting a broadband pulse over various distances. A sub-set of the reflections are shown overlaid in (a). As the distance d increases the pulses are delayed and attenuated. The rate of decay of each frequency component of the pulse was calculated, and used to find the frequency-dependent attenuation, this is shown with a quadratic fit in (b). Above 6 MHz the signals decay to quickly too measure accurately. 109

4.8 Experimental set-up. The array is positioned above the crossed hairs, such that they appear as two dots in the image plane. The hairs are at a distance of approximately 10 mm from the array, giving an angular aperture of 46° . The array is scanned along the hairs, which vary in separation from 0.8 mm to -0.8 mm. a) Perspective schematic, b) Cross section through image plane, c) Bird's-eye view. 110

4.9 The critical images where resolution of the hairs is lost. The top row of images are produced with traditional beamforming, the bottom with time-domain LSM. As separation of the hairs decreases from left to right the hairs remain resolved longer with TDLSM than BF. The TDLSM images are also freer of artefacts. 111

4.10 The single-frequency and broadband sampling images for one particular hair position (highlighted in red in figure 4.9). The broadband image is much more stable than the individual frequency images but retains almost the same resolution. See the text for a complete description. 112

5.1 The displacement fields corresponding to A_0 , A_1 , C_0 and C_1 . The central region is not shown so the smaller magnitude vectors further out can be seen. 117

5.2 The process of gating signals to remove shear waves. a) The original signals showing different mode paths. b) The expected arrival time of the wave modes. c) The signals with shear waves gated out. This particular example shows a pulse-echo signal, where the PS and SP modes always coincide. 119

5.3 Layout of the simulations. A 32-element array is positioned above a series of horizontal and vertical cracks. The background material was steel and a 2 MHz 5 cycle hamming tone burst was excited vertically at each transducer location. This 14 mm square denotes the imaging region. 120

5.4 The effect of gating on imaging the 3 mm horizontal and vertical cracks. a-d are without gating, e-h have unwanted shear modes removed. Horizontal and vertical cracks imaged with BF and FM are shown. The true crack length is shown as a white line, offset from the true location which is at the centre of the images. It can be seen that gating has little to no effect, and that FM performs better than BF. . 121

5.5 FM imaging of the 3 mm vertical and horizontal cracks with P (a, c) and S (b, d) wave speeds. 122

5.6 By considering the direction nature of shear wave generation, the magnitude of image artefacts can be reduced. Shown here are FM images of 1.5 mm vertical crack with (a) and without (b) consideration of the directionality. 123

5.7 Demonstration of super-resolution imaging using FM (a-f) and BF (g-l) in elastic media. Various lengths of horizontal (a-c, g-h) and vertical (d-f, j-l) cracks are imaged. The FM image sizes follow the crack sizes with much better accuracy than the BF images. 123

5.8 The experimental set-up. a) A 0.5, 1.0 or 2.0 deep EDM notch is cut into the side of a $600 \times 50 \times 50$ mm aluminium block. The notch is approximately 0.3 mm thick. b) Photograph of the array and block with a 2 mm notch. 125

5.9 a) The waves received by each transducer when a simulated tone-burst is transmitted on element 36 in a defect-free block. b) The various modes corresponding to the reflections in (a). Red and blue denote P and S waves respectively. Where mode conversion occurs the order is not important, i.e. P-S and S-P waves arrive simultaneously. 126

5.10 The simulated pulse-echo signals for a block containing a 1 mm crack. The three simultaneous reflections are the P-P, P-P-P-P and P-P-P-P-P-P reverberations between the front and back walls. The largest reflection from the crack is the S-S mode at $45 \mu\text{s}$. The P-P and P-S/S-P reflections from the crack can also be seen around 20 and $28 \mu\text{s}$ respectively, but they are both very weak. 127

5.11 Simulated wave fields resulting from transmission from element 36 for the 1 mm crack. P and S waves are coloured red and blue respectively. It can be seen that the “S-S” reflection is actually due to a more complex interaction of the mode converted S-P wave in addition to the S wave. The strength and high degree of directionality of the reflection shows the advantage of using a large aperture array and imaging with shear waves. 128

5.12 a) Experimental pulse echo signals for the 1 mm crack. The S-S reflections are clearly visible and P-P is just visible. b) The received signals for all array elements when transmitting on element 36. Both (a) and (b) closely match the simulations in figures 5.9 and 5.10. (There are actually slight timing differences due to the actual block being 51.04 mm thick, rather than 50 mm as in the simulation.) . . . 129

5.13 Average spectrum of the experimental signals. This was calculated by averaging the spectrum over all $5\ \mu\text{s}$ Hamming windows in all recorded signals. The centre frequency is approximately 2.4 MHz and the half-power bandwidth is 0.5 MHz. 130

5.14 (The square root of) the channel energies before (a) and after (b) coupling compensation. The slight anomalies seen every 32 transducers are due to the fact that the array is constructed from four 32 element sub-arrays. 131

5.15 Beamforming images of the cracks. a) 0.5 mm b) 1.0 mm c) 2.0 mm. Note that the backwall is nominally at $y = 0$, but there is a small error in the position of the cracks so the region below the backwall is shown. 133

5.16 FM images of the cracks at 2.5 MHz. 134

5.17 The first 500 measured singular values of the time domain operator $F_{d,t}$ for the 1 mm crack. Note that the values come in nearly equal pairs. The symmetry of the scatterer means that there are pairs of odd and even symmetric functions with the same singular value. . . . 135

5.18 Pre-beamformed signals. a) Received for transmitter 36, b) Pulse-echo, c) The form of the target signal, its exact location and amplitude depends on the receiving transducer and image point. But the key point is due to the pre-beamforming all possibilities occur later than the one shown, but before $10\ \mu\text{s}$. The signals are 220 samples long, so the $F_{t,d}$ matrix operator has dimensions 56192×28160 135

5.19 The final TDLSM images (left) for the a) 0.5, b) 1.0 and c) 2.0 mm cracks. The 1 mm and 2 mm cracks are both undersized, however, the 2 mm crack is noticeably larger than the others, which was not the case for beamforming (reproduced on the right). The wavelength is 1.3 mm. 136

6.1 The effect of array aperture on the completeness of the resulting image. A kite-shaped void (a) is imaged using the beamforming algorithm with b) a full-view ring array surrounding the void; and c) the top 90° of the array. The aperture restriction leads to an incomplete image of the scatterer. 139

6.2 A schematic of the solution to the limited-view problem proposed by Mora. A strongly scattering reflector is used as a mirror to provide additional views of the object to be imaged. 140

6.3 Two related scattering models. a) A limited-view configuration with a semi-circular array and an impenetrable planar interface along the line $y = 0$. b) A related full-view configuration with the scatterer and transducers mirrored and the interface removed. 142

6.4 A diagram of the matrix representation of the far-field operator. Only a square subset is accessible in limited-view situations. The main diagonal of the matrix corresponds to ‘pulse-echo’ measurements in which the same transducer is used for transmission and reception of the waves. The ‘through-transmission’ diagonals are those in which opposite transducers are used for transmission and reception. 143

6.5 The block structure of the bisymmetric far-field operator F . The superscript T and F refer to the transpose and flip-transpose. A is a symmetric matrix and B is persymmetric. 145

6.6 Schematic of the numerical examples with a circular array. The dotted red squares indicate the region that is imaged in the following figures. \mathbb{A} , \mathbb{B} and \mathbb{C} correspond to the situations described in the text (\mathbb{A} and \mathbb{C} are shown at a different scale to \mathbb{B}). The measurements for \mathbb{A} and \mathbb{B} correspond to the F and T matrices respectively. 151

- 6.7 The phase and normalised magnitude of the multistatic matrices for the three simulated situations show in figure 6.6 and described in the text. θ and ϕ refer to the illumination and detection angles respectively according to the convention given in figure 6.3a. It can be seen that the structure of the F matrix (\mathbb{A}) matches the block structure shown in figure 6.5, and that all matrices are symmetric as expected. 152
- 6.8 An example of the singular values for full-view, limited-view with a sound-soft mirror and limited-view with a sound-hard mirror configurations. The singular values for \mathbb{A} (solid black line) are equal to the combination of the singular values of the limited-view case with a sound-hard and sound-soft mirror (\mathbb{B}), which give rise singular values corresponding to the symmetric (red crosses) and anti-symmetric (blue circles) singular functions. The dashed black line is the singular values for \mathbb{C} , but with each element repeated to allow comparison with \mathbb{B} (the blue circles). It can be seen that the addition of the mirror increases the magnitude of the singular values. This corresponds with the increase in the available scattering information. 154
- 6.9 The results of imaging the kite void using the multistatic matrices in figure 6.7. The first row shows these results of imaging the region marked in figure 6.6 with the Beamforming algorithm, and the second row with the Factorisation Method. For imaging of the T matrix (b and e) the modified Green's function [eq. (6.22)] was used. Each image is normalised with respect to its maximum, and a linear scale is used. The size of the regions shown is $10 \times 10\lambda$ 155
- 6.10 Schematic of the numerical example with a linear array. The array has 90 elements, a length of 100λ and is positioned 100λ from the backwall. The array has an aperture of 90° 156

6.11 Results of imaging the kite void with the linear array for the configurations \mathbb{A} , \mathbb{B} , and \mathbb{C} shown in figure 6.10. With the decreased aperture afforded by the linear array, the BF images are now unusable. The FM images are mostly unchanged from the results for the circular array. 157

6.12 The results of imaging situation \mathbb{B} with various noise levels and apertures. The beamforming results are only shown for two noise levels since the noise has so little effect. 158

6.13 The FM results of imaging the kite at a variety of distances from a sound-soft backwall with a linear array. The angle subtended is the dotted angle marked on the left hand diagram. As the subtended angle increases, a larger and larger aperture is required for successful imaging. The noise level is 2.7%. 159

6.14 HARBUT Algorithm. See the main text for a description. 161

7.1 A simple method of gating out the incident field. a) The arrival time of the backwall reflection is calculated from geometry. b) This is turned into a smooth mask. c) The mask is multiplied by the total field to obtain d) an estimate of the incident field. 167

7.2 Very small phase errors make baseline subtraction impractical. a) The total field, and incident field measured with a small timing error, where the scattered field should be zero. When subtracted there is significant residual error even for a small timing error. b) The magnitude of the residual error vs the timing/phase error. 167

7.3 The incident field fitting algorithm. a) The parameters to be fit (r and θ); b) The algorithm (see main text for a complete description). . 169

7.4 Experimental incident field fitting results showing the phase of a) the measured T matrix, b) the simulated R matrix, and c) the experimentally measured R matrix found by removing the scatterer. The matrices in (b) and (c) should match. 170

7.5 a) Schematic of beam directionality measurement. b) Measured directivity for a 0.3 mm element in water at 4.5 MHz ($\lambda = 0.33$ mm). . . . 171

7.6 The results for backwall imaging of a 4 mm diameter air-filled straw in water using a) beamforming and b) FM. The size and position of the straw are marked. 173

7.7 The results for backwall imaging of a 5 mm razor blade suspended in two different locations in water using FM. 173

7.8 a) Layout of the tubes of alcohol in the experiment showing the array (crosses) and scatterers. The water surface is at $y = 0$ (the figure is upside-down compared to the physical experiment, for consistency with previous figures). b) Percentages of alcohol used in the experiment, and the resulting speeds of sound. 174

7.9 ToFT, DT and HARBUT images for a) 1%, b) 2%, and c) 5% changes in the speed of sound. 176

Chapter 1

Introduction

1.1 Background

The remarkable advances in electronics and computing over the past half-century have enabled a rapid succession of new ultrasound technologies. The ability to form images of the interior of an object, merely by vibrating its surface has proven to be an immensely useful tool for a wide variety of medical and industrial applications. In medicine it allows us to see inside the body without using invasive surgery or ionising radiation and in industry, ultrasound is used for Non-Destructive Evaluation (NDE), where it enables damage to components caused by fatigue and corrosion to be detected before it becomes critical.

The use of NDE to prevent component failure enables the implementation of *damage-tolerant* or *fail-safe design*, in which systems can be designed with expected component lifetimes shorter than the system lifetime [3]. The ability to detect and monitor damage allows the remaining lifetime of a component to be measured after its manufacture so safety factors can be less conservative, and parts can be replaced at the end of their actual lifetime rather than before their minimum expected lifetime, which was the previous approach (known as *safe life design*). This can be a more cost effective approach provided that the cost of each inspection is inexpensive compared to the cost of replacing components, and the time between inspections is

sufficiently long that the number of inspections can be minimized throughout the system lifetime. Additionally inspection is only worthwhile when the cost of failure is high, so ultrasound NDE is often used in safety-critical systems such as rail and air transport, and nuclear power. Similar cost considerations apply to medicine, where ultrasound can be used to diagnose medical conditions. As with NDE there are costs associated with failure to detect disease, but also with erroneous detections which may result in unnecessary follow-up examinations, surgery and stress for the patient.

Ultrasound is widely used in both NDE and medical diagnostics [4]. This can be attributed in part to the relatively good performance of ultrasound as a defect detection method. The performance of ultrasound—and any screening or detection method—is characterised by two metrics: the Probability of Detection (PoD) and the Probability of False Detection (PoFD). In medicine two related parameters are used: the sensitivity and specificity. If a method has high PoD (and therefore sensitivity) it is good at detecting the target when it *is* present. If it has a low PoFD (high specificity) it is good at *not* detecting the target when it *is not* present. The objective is to achieve a high PoD and low PoFD simultaneously. Ideally the PoD will be 100%, and the PoFD will be 0% but in reality less perfect numbers such as 90%, 10% will be found. The acceptable values of PoD and PoFD depend on the application and the costs associated with incorrect detections.

With ultrasonic NDE, the PoD and PoFD also depend on the size of the defect we wish to detect. Detecting small defects is more difficult than large defects because the ultrasound reflects less strongly from them, and the reflections are therefore hard to distinguish from other signals such as noise and reflections from component features. The ability to detect small defects is the key to reducing costs, as it allows earlier diagnosis in medicine and longer periods between inspections in NDE since it takes longer for a small defect to grow to a critical size than a large one [5]. The pursuit of higher PoDs and lower PoFDs is a key area of modern NDE research.

One way to raise the PoD and lower the PoFD is by *imaging*. In contrast with medical ultrasound, imaging is not widely used in NDE. Instead, ad-hoc methods are used that use the recorded signals (“A-scans”) directly rather than processing

them into an image. The PoD and PoFD of image-based NDE depends on the quality of the images, which can be characterised by a number of metrics including contrast, resolution, noise level, artefacts, completeness and accuracy. The concepts of contrast, resolution and noise are familiar from digital photography. Examples of these factors are shown in figure 1.1 where they adversely affect the detection of small features like the legs of a dragonfly. These metrics are affected by the array and array controller performance and also the sophistication of the imaging algorithm which is used to map the measurements to an image.

The maximum resolution of an image is subject to a trade-off with imaging depth. An image's resolution can be characterised by the Rayleigh criterion which is the minimum separation at which two point scatterers are just resolvable. This separation is proportional to the wavelength of the waves used for imaging, so higher frequencies result in a higher image resolution. However, higher frequencies generally result in higher attenuation and a shorter penetration distance. Therefore there is a trade-off where low frequencies waves can achieve high penetration, but have low resolution, and high frequency wave give high resolution but do not penetrate very far. The frequency of waves that can be used is constrained by the penetration depth required so arbitrarily high resolutions cannot be achieved. If imaging is constrained to low resolution, it is difficult to detect and characterise small defects.

To see how ultrasound images can be improved we must discuss how they are produced in the first place. The key to producing images is the use of ultrasonic arrays. Ultrasound is produced using Lead Zirconate Titanate (PZT), a piezoelectric material that expands and contracts when a voltage is applied across it. Single PZT transducers can be used for the electronic generation and detection of ultrasound but in order to produce an image without scanning the transducer it is necessary to use an array of PZT elements so that ultrasonic waves can be transmitted from and received at many locations simultaneously. This also provides more data than scanning a single transducer since the transmit and receive locations can be different. In their final form, arrays, and array elements are not formed of a solid block of PZT, but from PZT composite—PZT rods or sheets embedded in a polymer matrix [6].




Figure 1.1: Images of a dragonfly that are a) ideal, b) low resolution, c) noisy, and d) low contrast. Certain features, such as its legs and the pattern on its wings, are rendered invisible by the degradation. Photo by Billy Lindblom.

The PZT rods/sheets are diced mechanically from a solid block and the gaps are filled with polymer. PZT composite has several advantages over bulk PZT. It expands only in the desired direction, it can be deformed to make curved arrays for example, and it reduces cross-talk between elements to allow them to operate independently. The construction of a 1D four element array is shown in figure 1.2. The elements are defined by their electrodes which span multiple PZT pillars.

The array is driven electronically by an array controller, which contains the electronics necessary to generate and receive ultrasound. An array controller generates the input wave for each element, and amplifies and digitises the received waves. Figure 1.3 shows a high level schematic of one channel. An input generator produces a high voltage pulse or tone burst that causes the PZT to vibrate and produce waves. These are reflected by component features and cause a weak electrical signal to be generated by the PZT. This is amplified and digitised to be used for input to the imaging algorithm. In recent years array controller technology has greatly advanced, with increasing numbers of channels and a greater degree of parallel recording. For example, a typical array controller from a few years ago may have 128 channels, and




Figure 1.2: Typical construction of 4-element linear array. A composite of PZT pillars in a polymer matrix is sandwiched between element electrodes. Each array element is made from several pillars vibrating in concert.

be able to produce pulse trains and record 32 channels simultaneously, but systems now are available with arbitrary wave generators, and that are capable of recording up to 2048 channels in parallel (or more using multiplexing), a feat which produces a vast amount of data and is only possible with the availability of large cheap RAM. The ability to transmit and receive waves from many points in space is the key to imaging as it allows the transmission and reception of focused waves.

Once an array has transmitted and received ultrasound waves, a particular imaging algorithm is applied to compute an image from the measurements. Many imaging algorithms exist, but one of the simplest and most intuitive is *beamforming* [7, 8]. Beamforming is equivalent to confocal microscopy [9], which works by focusing a beam of light to a point inside the component and then focussing in reception on the light scattered from that point. The magnitude of the scattered light depends predominantly on the transmittance of the material at the focal point, so it is used to colour the image at that point. In confocal microscopy, a lens is used to focus the transmitted and scattered light (figure 1.4a), whereas in ultrasound




Figure 1.3: A schematic showing the process of generating and receiving ultrasound for one channel of an array. The input generator can be a pulse system, or an arbitrary function generator and power amplifier. This signal is fed to the array element which transmits the wave. The resulting weak reflected signal is amplified by a high gain amplifier and digitised by an ADC.

an array of transducers is used instead (figure 1.4b). In both cases focussing is achieved by delaying the waves that are received or transmitted at different points along the aperture. In a lens the varying thickness produces greater delays at the centre, whereas the delays in an array are produced electronically. Because these electronic delays are not fixed they can be dynamically changed to produce focal laws equivalent to any lens. Arrays can be made with elements laid out in a wide variety of configurations, some of which are shown in figure 1.5. The geometry of the array determines the range of locations that be focused on, and the size of the focal spot. Currently the most common are 1D linear arrays, which only allow focusing—and therefore imaging—in a plane. 2D arrays allow focusing anywhere, but need many more elements to achieve this, and hence are only recently becoming more common as high channel count array controllers become available. A review of the use of arrays and focusing in NDE is given by Drinkwater [10].

The beamforming imaging method can be used identically in NDE and medical ultrasound, however there are significant differences between tissue and metal that make imaging in tissue easier, and have lead to its widespread use in medicine. The most significant of these differences is the strength and distribution of scatterers. In tissue, features such as blood vessels are generally weakly scattering and distributed throughout the medium. In metal, scatterers such as cracks and voids are infrequent and strongly scattering. This allows medical ultrasound to take advantage of speckle,

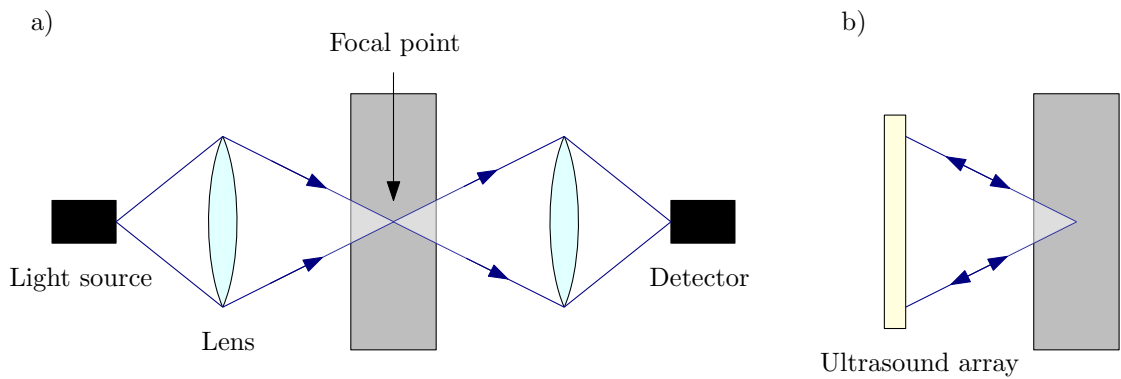


Figure 1.4: *In confocal microscopy (a) a lens is used to focus light at a point, and also used to detect reflections from that point. By scanning the focal point an image can be built. Ultrasonic beamforming (b) uses a similar approach, but an array of elements is used to transmit and receive the focused waves, and usually the same array is used for transmission and reception, although this need not be the case*




Figure 1.5: *Common array geometries. a) 1D or Linear; b) "1.5D" or Hybrid; c) 2D or Matrix; d) Annular. Although there are some common configurations, essentially any geometry can be manufactured. Arrays also do not necessarily have flat faces, they can be curved or even circular.*

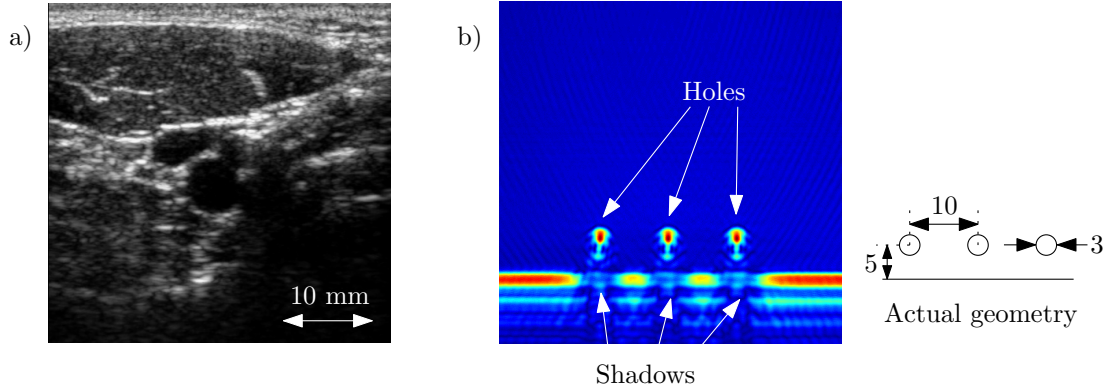


Figure 1.6: A medical ultrasound image of the jugular (a), compared with an NDE image of three side-drilled holes in a stainless steel block (b). The structure of the tissue and blood vessels in (a) can be seen due to the speckle contrast, which cannot be seen in (b). Additionally, the impenetrable nature of the holes results in shadowing of the backwall.

which is a random but deterministic pattern caused by the addition of the reflected waves from randomly distributed point-like scatterers. A medical ultrasound image showing speckle is shown in figure 1.6a. Although the speckle does not correspond directly to physical features, its contrast does depend on the type of tissue. The tissue is also penetrable which means we can see speckle throughout its volume. This makes it possible to see the complete size and shape of features such as blood vessels and even some individual cells. In contrast, speckle in metal is very weak, and the impenetrable nature of the scatterers means that we can only see their top surfaces. An example of this is shown in figure 1.6b, which shows three holes in metal above an air interface. Only the tops of the holes can be seen, and they produce shadows which obscure parts of the interface.

It can be seen from figure 1.6 that conventional beamforming imaging reveals significantly less structure and detail in NDE than medicine. It is highly desirable to improve the NDE image so that the complete shape of the scatterers can be seen and higher resolution can be achieved, so that the PoD and PoFD may be maximised and minimised respectively.

It can be shown that measuring the scattered waves from transmission of each element of an array sequentially will result in “complete” scattering data. Any additional measurements will be redundant. The optimal physical measurement procedure is

therefore fixed. To improve image quality we must focus on how information is decoded from the measured data to form an image. The central theme of this thesis is that *all information encoded in the scattered waves should be used in order to achieve as high an image quality as possible*. The volume of information can be quantified and we look at how it is distributed in the temporal and spatial domains of the measured signals. The sources of information were identified that are usually neglected:

1. The multiply scattered signals that interact with multiple parts of the damage (this includes shadows). These waves are small in tissue, but can be significant in metal. Multiple scattering is the mechanism by which the high-resolution spatial information is conveyed from non-propagating evanescent waves to the vicinity of the array. This information is not considered by beamforming or other linear imaging algorithms.
2. Signals received after the reflections from structural features of a component (e.g. its outer surfaces). These are normally discarded but they allow us to send and receive probing waves to and from the far side of the scatterers. In essence the outside surfaces of the component can be used as mirrors to see the far side of the scatterers, in the same way the far side of an object can be seen if a mirror is placed behind it. The reflection of the waves from the component boundaries allow a more complete reconstruction where all sides of the scatterer can be seen, not just the side nearest to the array.

The information encoded in multiply scattered waves can be decoded using recent “super-resolution” methods. These are highly sensitive to noise due to being single-frequency methods. This is solved by using a time-domain super-resolution algorithm which is demonstrated experimentally to achieve a resolution beyond that of beamforming in an attenuative medium where maximum wave frequency is highly constrained. The backwall reflection data can be incorporated into the reconstruction, but it is necessary to know the position of the backwall so the path of the reflected waves can be calculated. A method to measure this is developed and

experimental results are given. Both techniques are further complicated by elastic mode conversion.

These ideas have been introduced in other fields previously (and references will be given later). However the results so far have been mostly theoretical and use simulated data; it is not clear if they are feasible in practice. In this thesis we aim to combine the two sources of information given above and to solve the practical difficulties that have kept their use theoretical to date. The aim of this thesis is therefore to improve the quality of ultrasound images in NDE (although the results are applicable to other imaging fields including medical ultrasound and geophysics) by utilising information that is usually discarded, focusing on multiply scattered signals, and signals occurring after the back-wall echoes. These are studied from an entropic information-encoding perspective. Utilising these two sources of information should allow improvements to image resolution, accuracy and completeness, and should ultimately lead to a higher PoD and lower PoFD, particularly when imaging small defects.

1.2 Thesis Plan

The thesis is organised as follows:

In chapter 2, the forward and inverse scattering problems are formulated. The forward scattering problem consists of finding the scattered wave field from an object, given the object's structure and the incident wave field. This is well-posed and easily solvable. The inverse scattering problem is to find the object's structure given the incident and scattered wave fields. This problem is ill-posed and much harder to solve. We will explain the classical imaging methods and how they lead to resolution limits, and then introduce more advanced 'super-resolution' methods which break those limits.

In chapter 3, we will look at how the information about the object structure is actually encoded into the scattered wave fields. Shannon's sampling theorem is extended to

provide an upper bound on the amount of information embedded in the ultrasound signals. This information is not distributed uniformly, but is concentrated in finite regions of time, and temporal and spatial frequencies. This allows us to see where classical algorithms are discarding information, and therefore how improvements to the images might be obtained.

In chapter 4, we describe the theory, implementation and experimental results for a broadband super-resolution algorithm. Super-resolution algorithms theoretically allow exact reconstructions, however they are usually single-frequency methods which makes them particularly sensitive to noise. By using a broadband method, the stability of linear imaging methods can be retained, while still achieving super-resolution. This is demonstrated by imaging two point scatterers in a highly attenuative medium where super-resolution is most useful.

In chapter 5, the effect of elastic waves on the super-resolution algorithm is investigated. These provide additional complications to the imaging algorithms which must be overcome. The broadband super-resolution method is tested experimentally on EDM notches in an aluminium block.

In chapter 6, we look at the theory of using information contained in signals after their reflection from the component edges. In other words the possibility of using the backwall or side walls of a component as mirrors in order to gain a view behind a scatterer and build a more complete image. The scattering operation is decomposed into terms involving the mirror, and the effect of various parameters such as aperture and noise on the quality of the reconstructions are considered.

In the penultimate chapter, the theory of chapter 6 is tested experimentally, and in chapter 8 the conclusions and main contributions of the thesis are summarised.

Chapter 2

Formulation of the Imaging Problem

2.1 Background

This chapter introduces the physical and mathematical basis of the scattering and imaging problems. The scattering and imaging steps effectively encode and decode the scatterer's structural information in the radiated wave fields. The imaging operation is the inverse of the *forward* or *direct* scattering operation, so it is also known as *inverse scattering*. While forward scattering transforms or maps the scatterer's physical structure onto disturbances to a wave field, inverse scattering attempts to invert the process and recover the physical structure from the scattered wave field.

The forward scattering problem is well posed [11], and a number of numerical methods exist that provide nominally identical solutions. However, the inverse problem is ill-posed [12] and many different imaging algorithms exist which can produce vastly different results. The inverse scattering problem is also non-linear so the most common algorithms provide approximate solutions by assuming the problem is linear.

After formulating the forward and inverse scattering problems, the linear beamforming

imaging algorithm and its resolution limit will be described, and non-linear algorithms will be introduced that can break this limit. This formulation will provide the foundation for the following chapters where these advanced algorithms will be described and investigated. It will also provide the tools needed to identify the areas in which ultrasonic imaging can be improved.

2.2 Forward acoustic scattering

The forward scattering problem is that of calculating the response of an object subjected to a known vibration. The input vibrations propagate through the object as waves, according to laws of physics and the physical properties of the object. The problem is to calculate these wave fields, given the input vibrations, the physical model, and the material properties of the object. We will first derive a wave equation that models the system, and describe its solution.

Consider a solid 3D medium, with material properties (density, Young's modulus, shear modulus, etc.) that vary with position \mathbf{x} . Suppose a time-varying input body force $\mathbf{F}(\mathbf{x}, t)$ is applied, with $\mathbf{F}(\mathbf{x}, t < 0) = 0$, and the resulting displacement of the material at each point is $\mathbf{d}(\mathbf{x}, t)$, also with $\mathbf{d}(\mathbf{x}, t < 0) = 0$. The forward scattering problem is then to find $\mathbf{d}(\mathbf{x}, t)$ given the material properties and $\mathbf{F}(\mathbf{x}, t)$ and appropriate boundary conditions. The accuracy with which we can find the solution to this problem depends on the accuracy of our physical model of wave propagation. It is possible to consider a large number of physical phenomena, such as thermal and viscous effects, but for most problems linear elasticity alone is sufficient. Therefore we can model the forces in the material according to the Hooke's law:

$$\boldsymbol{\sigma} = \mathbf{C} : \boldsymbol{\epsilon} \quad (2.1)$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\epsilon}$ are 2nd order stress and strain tensors, and \mathbf{C} is the 4th order stiffness tensors.¹ By applying Newton's laws of motion to an infinitesimal volume we find

¹The stress, strain and stiffness vary with position, but we have dropped the explicit reference

that the displacement must obey the following partial differential equation [13].

$$\rho \ddot{\mathbf{d}} = (\lambda + 2\mu)\nabla(\nabla \cdot \mathbf{d}) - \mu\nabla \times (\nabla \times \mathbf{d}) + \mathbf{F} \quad (2.2)$$

where ρ is the density. This is *Navier's equation*, the *vector* or *elastic wave equation* [14]. It can also be written using the vector Laplacian $\nabla^2 = (\nabla \cdot \nabla)$ as

$$\rho \ddot{\mathbf{d}} = (\lambda + \mu)\nabla(\nabla \cdot \mathbf{d}) - \mu\nabla^2 \mathbf{d} + \mathbf{F}, \quad (2.3)$$

where λ and μ are the Lamé constants. The forward scattering problem is then to find $\mathbf{d}(\mathbf{x}, t)$ given $\lambda(\mathbf{x})$, $\mu(\mathbf{x})$ and $\mathbf{F}(\mathbf{x}, t)$, and boundary and initial conditions. Although numerical solutions to this equation are readily available, inverse scattering with elastic waves is significantly more difficult than with acoustic waves. Therefore, in medicine, NDE and geophysics, elastic waves are usually analysed as if they were acoustic. This is not without motivation because elastic waves can be decomposed into the superposition of acoustic and a shear waves. This is known as the Helmholtz decomposition.

2.2.1 Helmholtz Decomposition

The Helmholtz theorem states that any vector field \mathbf{d} can be decomposed into a solenoidal field \mathbf{d}_t with $\nabla \cdot \mathbf{d}_t = 0$ and an irrotational field \mathbf{d}_l with $\nabla \times \mathbf{d}_l = 0$, provided the divergence and curl of \mathbf{d} vanish at infinity. Then

$$\mathbf{d} = \mathbf{d}_t + \mathbf{d}_l \quad (2.4)$$

$$\mathbf{d}_t = \nabla \times \mathbf{s} \quad \mathbf{d}_l = \nabla p \quad (2.5)$$

to \mathbf{x} for brevity.

Where \mathbf{s} is a *vector potential*, and p is a scalar potential. Note that $\nabla \cdot \nabla \times \mathbf{s} = 0$ and $\nabla \times \nabla p = 0$. Because we have decomposed a vector field into a vector and scalar field, we have added an additional degree of freedom. This degree of freedom arises from the fact that multiple vector potentials \mathbf{s} correspond to identical values of \mathbf{d}_t . In general, $\nabla \times \mathbf{s} = \nabla \times (\mathbf{s} + \nabla m)$ where m is a continuously differentiable scalar field. In order to fix this we must choose a *guage condition* so that \mathbf{s} only has two degrees of freedom.

Due to the linearity of equation (2.2), any solution can be split up into the two solutions \mathbf{d}_t and \mathbf{d}_l . Substituting \mathbf{d}_t and \mathbf{d}_l into equation (2.2) gives the following equations.

$$\rho \ddot{\mathbf{d}}_t = -\mu \nabla \times \nabla \times \mathbf{d}_t + \mathbf{F}_t \quad (2.6)$$

$$\rho \ddot{\mathbf{d}}_l = (\lambda + 2\mu) \nabla (\nabla \cdot \mathbf{d}_l) + \mathbf{F}_l \quad (2.7)$$

The identity $\nabla \times \nabla \times \mathbf{d} = \nabla (\nabla \cdot \mathbf{d}) - \nabla^2 \mathbf{d}$ can be used with equation (2.6) to eventually obtain

$$\rho \frac{\partial^2 \mathbf{d}_t}{\partial t^2} = \mu \nabla^2 \mathbf{d}_t + \mathbf{F}'_t \quad (2.8)$$

(recall $\nabla \cdot \mathbf{d}_t = 0$). Again ∇^2 is the vector Laplacian and is defined as $\frac{\partial^2}{\partial x_i^2}$ (the same as the scalar Laplacian). Since μ and ρ are both positive we can substitute $c_s^2 = \frac{\mu}{\rho}$ to give the transverse (shear) wave equation

$$\frac{\partial^2 \mathbf{d}_t}{\partial t^2} - c_s^2 \nabla^2 \mathbf{d}_t = \mathbf{F}'_t \quad (2.9)$$

where $\mathbf{F}'_t = \frac{\mathbf{F}_t}{\rho}$ and c_s is the bulk velocity of shear waves in the material. For 2D waves, \mathbf{s} is a scalar (or equivalently, its x and y components are zero, and there is only one shear potential component. In 3D there are three shear potential components

(corresponding to shear in the X, Y and Z axes), but only two possible shear waves that can travel in a certain direction—horizontally and vertically polarised waves.

Now we take the divergence of equation (2.7) to obtain

$$\rho \nabla \cdot \frac{\partial^2 \mathbf{d}_l}{\partial t^2} = (\lambda + 2\mu) \nabla^2 (\nabla \cdot \mathbf{d}_l) + \nabla \cdot \mathbf{F}_l \quad (2.10)$$

Using the following substitutions

$$p = \nabla \cdot \mathbf{d}_l \quad (2.11)$$

$$F'_l = \nabla \cdot \mathbf{F}_l \quad (2.12)$$

and the fact that $\frac{\partial^2 p}{\partial t^2} = \nabla \cdot \frac{\partial^2 \mathbf{d}_l}{\partial t^2}$ we find

$$\rho \frac{\partial^2 p}{\partial t^2} - (\lambda + 2\mu) \nabla^2 p = F'_l \quad (2.13)$$

Now ρ and $\lambda + 2\mu$ are positive, so we can replace them with $c_l^2 = \frac{\lambda+2\mu}{\rho}$ to give the longitudinal wave equation:

$$\frac{\partial^2 p}{\partial t^2} - c_l^2 \nabla^2 p = F'_l \quad (2.14)$$

where c_l is the bulk velocity of longitudinal waves in the material. Equations (2.9) and (2.14) describe two independent wave fields of dilatatory and deviatoric stresses. These fields are only independent where there is no restriction on the total wave field \mathbf{d} . At scatterers they are dependent and mode conversion can occur. However, in the bulk material the equation for longitudinal waves is identical to the acoustic wave equation. In other words, P-waves in solids behave identically to acoustic waves when they are not scattering. The acoustic wave equation is

$$\frac{\partial^2 p}{\partial t^2} - c^2 \nabla^2 p = F'_l \quad (2.15)$$

where c is the wave speed and $p(\mathbf{r}, t)$ is the pressure field.

2.2.2 Solution of the Acoustic Wave Equation

In order to solve equation (2.15), we first solve the problem for a point source $F'_l = \delta(\mathbf{r}, t)$ (where δ is the Dirac delta function).

The Green's function is the solution to

$$\frac{\partial^2 p}{\partial t^2} - c^2 \nabla^2 p = \delta(\mathbf{x}, t) \quad (2.16)$$

which in 3D is given by

$$p(\mathbf{x}, t) = \frac{\delta(t - \frac{|\mathbf{x}|}{c})}{4\pi|\mathbf{x}|} \quad (2.17)$$

and in 2D by [15]

$$p(\mathbf{x}, t) = \frac{H(t - \frac{|\mathbf{x}|}{c})}{\sqrt{t^2 - \frac{|\mathbf{x}|^2}{c^2}}} \quad (2.18)$$

where $H(\cdot)$ is the Heaviside step function. Convolution of a time-varying point source with the 3D Green's function results in a travel-time delay, and attenuation due to geometrical beam-spreading, and is simple to calculate. On the other hand, convolution with the 2D Green's function results in a complex frequency-dependent distortion of the waveform that is difficult to calculate due to the singularity at $t = 0$. However, due to linearity the solution can be simplified by Fourier transforming the problem into the frequency domain, thereby removing the time dependence of $p(\mathbf{x}, t)$. The spatial variation of $p(\mathbf{x}, t)$ at a particular frequency is denoted $u(\mathbf{x}, \omega)^2$, such

²In this thesis u is pressure in the frequency domain, not displacement.

that

$$p(\mathbf{x}, t) = \frac{1}{2\pi} \int u(\mathbf{x}, \omega) e^{-i\omega t} d\omega \quad (2.19)$$

where ω is the angular frequency. u is complex in general. For a single frequency ω , this transforms the scalar wave equation into the *Helmholtz equation*

$$\nabla^2 u + k^2 u = F \quad (2.20)$$

where $k = \omega/c$ is the angular wavenumber and we have re-purposed F to be the Fourier component of F'_t . This equation can be solved analytically in free space, however this implies that the material properties are constant, in which case no scattering can occur and the solution will be trivial. There are two ways of modelling scatterers, depending on whether they are penetrable or impenetrable. For impenetrable scatterers, a boundary condition such as Neumann or Dirichlet is enforced on the scatterers surface. For sound-soft (empty) scatterers the Dirichlet condition applies ($u = 0$); for sound-hard (rigid) scatterers the Neumann condition applies ($\frac{\partial u}{\partial \mathbf{n}} = 0$, where \mathbf{n} is the surface normal). Penetrable scatterers can be modelled using a variable k within the scatterer. This is made more convenient by using the *object function*.

2.2.3 The Object Function

Consider a penetrable scatterer with support \mathcal{D} . The object function is defined as

$$o(\mathbf{x}) = \frac{k_d^2(\mathbf{x})}{k^2} - 1 \quad \mathbf{x} \in \mathcal{D} \quad (2.21)$$

where $k_d(\mathbf{x})$ is the wavenumber in \mathcal{D} and k is the background wavenumber.³ This alteration to the wavenumber can be considered to be an external forcing function

³In this thesis o stands for 'object', and is not related to little-o notation.

located at the scatterer such that $F = -o(\mathbf{x})k^2u$. Then equation (2.27) becomes

$$\nabla^2 u + k^2(1 + o(\mathbf{x}))u = 0 \quad (2.22)$$

$$\nabla^2 u + k^2u = -k^2o(\mathbf{x})u \quad (2.23)$$

It is helpful to decompose the physical field (the *total field*) u into the incident field u^i and the scattered field u^s so that

$$u = u^i + u^s. \quad (2.24)$$

Therefore we have the equation

$$(\nabla^2 + k^2)(u^i + u^s) = -k^2o(\mathbf{x})u \quad (2.25)$$

and we define the incident field as the solution where there are no scatterers

$$(\nabla^2 + k^2)u^i = 0. \quad (2.26)$$

These equations can be solved using the Green's function of the Helmholtz equation which is the solution to:

$$\nabla^2 u(\mathbf{x}) + k^2u(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}') \quad (2.27)$$

subject to the Sommerfeld radiation condition which precludes non-physical solutions containing incoming waves [12]. In 3D this gives the Green's function

$$G_{3D}(\mathbf{x}, \mathbf{x}') = -\frac{e^{ik|\mathbf{x}-\mathbf{x}'|}}{4\pi|\mathbf{x} - \mathbf{x}'|} \quad (2.28)$$

and in 2D

$$G_{2D}(\mathbf{x}, \mathbf{x}') = -\frac{i}{4}H_0^{(1)}(k|\mathbf{x} - \mathbf{x}'|). \quad (2.29)$$

where $H_0^{(1)}$ is the Hankel function of the first kind and order 0. This form can be very accurately approximated where $|\mathbf{x} - \mathbf{x}'| > k^{-1}$ as

$$G_{2D}(\mathbf{x}, \mathbf{x}') = -\frac{i+1}{4\sqrt{\pi k|\mathbf{x} - \mathbf{x}'|}}e^{ik|\mathbf{x} - \mathbf{x}'|}. \quad (2.30)$$

Using the standard Green's function method it can be shown that

$$u^s(\mathbf{x}) = -k^2 \int_{\mathcal{D}} o(\mathbf{x}')u(\mathbf{x}')G(\mathbf{x}, \mathbf{x}')d\mathbf{x}'. \quad (2.31)$$

If we set the incident wave to a plane wave travelling in the $\hat{\mathbf{r}}$ direction, $e^{ik\hat{\mathbf{r}} \cdot \mathbf{x}}$ we have

$$u(\mathbf{x}) = e^{ik\hat{\mathbf{r}} \cdot \mathbf{x}} - k^2 \int_{\mathcal{D}} o(\mathbf{x}')u(\mathbf{x}')G(\mathbf{x}, \mathbf{x}')d\mathbf{x}'. \quad (2.32)$$

This is called the *Lippmann-Schwinger* equation [16]. It describes how the total wave field depends on the incident wave and recursively on its own values within \mathcal{D} . This reflects the way the waves bounce repeatedly between scatterers and is the source of the forward scattering problem's non-linearity in $o(\mathbf{x})$. In other words, suppose we have two different object functions, $o_1(\mathbf{x})$ and $o_2(\mathbf{x})$ subject to the same incident field that produce total fields $u_1(\mathbf{x})$ and $u_2(\mathbf{x})$. Then the field for an object function $o_1(\mathbf{x}) + o_2(\mathbf{x})$ is *not* equal to $u_1(\mathbf{x}) + u_2(\mathbf{x})$. It can also be seen that this equation is linear with respect to $u^i(\mathbf{x})$.

Many methods have been created to solve the forward scattering problem. They can be divided in to two categories, volumetric simulations and boundary value simulations. The two main volumetric simulation methods are the Finite Difference




Figure 2.1: *a) In the far-field scattering problem an incident wave $u_{\infty}^i(\hat{\mathbf{x}})$ illuminates the scatterer. b) The incident wave interacts with the scatterer, which is labelled \mathcal{D} with boundary $\partial\mathcal{D}$, and produces a scattered wave $u_{\infty}^s(\hat{\mathbf{x}})$. The far-field scattering problem is to determine $u_{\infty}^s(\hat{\mathbf{x}})$ given $u_{\infty}^i(\hat{\mathbf{x}})$.*

Method (FDM) [17] and Finite Element Method (FEM). They calculate the wave fields throughout the entire volume, based on a direct discretisation of the wave equation and can be formulated in the frequency domain or the time domain. Boundary value methods are based on the Green's function of the wave equation, so they can skip direct calculation of field values at points other than the boundaries. They can generally only be formulated in the frequency domain, due to the complexity of Green's function in the time domain (in 2D at least). Because they only calculate the fields at boundaries, they can be very fast, but scale badly with scatterer size, and are also not suitable for penetrable scatterers with varying k_d . Examples of these methods include the Boundary Element Method (BEM) [18] and the Distributed Point Source Method (DPSM) [19]. In this thesis we have used both FDM, and BEM.

2.2.4 Far-field scattering

In the physical world we cannot excite and measure the wave fields at every point, but rather at a set of points in the far-field of the scatterer. Standard far-field formulations consider measurements at an infinite distance, so that only propagating waves can be

measured, and measurements are performed outside the scatterer, \mathcal{D} . In the far-field, wave fields can be separated in polar coordinates, so that $u(\mathbf{x}) = \Gamma(|\mathbf{x}|)u_\infty(\hat{\mathbf{x}})$, where $\hat{\mathbf{x}} = \frac{\mathbf{x}}{|\mathbf{x}|}$. The radial component, $\Gamma(|\mathbf{x}|)$ only depends on geometric beam-spreading and the direction of the waves (incoming or outgoing), and is independent of the scatterer. Therefore only the angular component $u_\infty(\hat{\mathbf{x}})$ contains information about the scatterer. The *far-field* forward-scattering problem is then to find $u_\infty^s(\hat{\mathbf{x}})$ for all directions $\hat{\mathbf{x}}$ given $u_\infty^i(\hat{\mathbf{x}})$, as shown in figure 2.1. The subscript ∞ refers to the fact that these are measurements of only the angular component of the fields and are taken at an infinite distance from the scatterer. The far-field is given by

$$\Gamma u_\infty(\hat{\mathbf{x}}) = \lim_{x \rightarrow \infty} u(\mathbf{x}) \quad (2.33)$$

where Γ given by

$$\Gamma_{3D} = -\frac{e^{ikx}}{4\pi x} \quad (2.34)$$

$$\Gamma_{2D} = -\frac{(i+1)e^{ikx}}{4\sqrt{\pi kx}}. \quad (2.35)$$

and we have used the notation $x = |\mathbf{x}|$. This is valid for out-going wave fields; for incoming waves such as $u^i(\mathbf{x})$, Γ^* is used instead.

We can calculate the result of propagation of near-field waves to and from the far-field using the paraxial approximation of the Green's function. The paraxial approximation is $\lim_{x \rightarrow \infty} |\mathbf{x} - \mathbf{x}'| = x - \frac{\mathbf{x} \cdot \mathbf{x}'}{x}$ (see figure 2.2) and it allows us to consider only the angular component of G as with u_∞ . We have (Γ has the same definition as before)

$$\lim_{x \rightarrow \infty} G(\mathbf{x}, \mathbf{x}') = \Gamma G_\infty(\hat{\mathbf{x}}, \hat{\mathbf{x}}') \quad (2.36)$$

where G_∞ is a plane wave given by




Figure 2.2: *The paraxial approximation. When $|\mathbf{x}|$ is very large, the component of \mathbf{x}' perpendicular to \mathbf{x} can be ignored when calculating the distance from \mathbf{x}' to \mathbf{x} , therefore $|\mathbf{x} - \mathbf{x}'| \approx |\mathbf{x}| - \frac{\mathbf{x}' \cdot \mathbf{x}}{|\mathbf{x}|}$*

$$G_\infty(\hat{\mathbf{x}}, \mathbf{x}') = e^{-ik\hat{\mathbf{x}} \cdot \mathbf{x}'}. \quad (2.37)$$

We can now find the incident near-field due to an incident far-field wave with density $u_\infty^i(\hat{\mathbf{r}})$.

$$u^i(\mathbf{x}) = \int_S u_\infty^i(\hat{\mathbf{r}}) G_\infty^*(\hat{\mathbf{r}}, \mathbf{x}) d\hat{\mathbf{r}} \quad (2.38)$$

where integration is performed over all angles (or solid angles) on the unit circle (or sphere) S . From equation (2.31) we obtain

$$u_\infty^s(\hat{\mathbf{x}}) = -k^2 \int_{\mathcal{D}} o(\mathbf{x}') u(\mathbf{x}') G_\infty(\hat{\mathbf{x}}, \mathbf{x}') d\mathbf{x}'. \quad (2.39)$$

These two equations map the incident far-field to the near field, and then the total (or scattered) field in the object $u(\mathbf{x})$ to the scattered far-field $u_\infty^s(\hat{\mathbf{x}})$. As was mentioned previously, scattering is linear in $u^i(\hat{\mathbf{x}})$ and therefore o determines a linear operator which maps from $u_\infty^i(\hat{\mathbf{x}})$ to $u_\infty^s(\hat{\mathbf{x}})$. That operator is called the *far-field scattering operator* F , and

$$u_\infty^s(\hat{\mathbf{x}}) = F u_\infty^i(\hat{\mathbf{x}}). \quad (2.40)$$

Because this operator is linear the scattered field resulting from any incident field

transmitted from the far-field can be calculated from a linear combination of the scattering of incident plane waves. We can now define the *scattering amplitude* $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ as the solution to equation (2.40) for $u_\infty^i(\hat{\mathbf{x}}) = \delta(\hat{\mathbf{x}} - \hat{\mathbf{r}})$, which is a plane wave incident in the $\hat{\mathbf{r}}$ direction. Therefore each value of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ corresponds to the complex amplitude of the scattered wave in direction $\hat{\mathbf{x}}$ due to an incident wave in direction $\hat{\mathbf{r}}$. Using $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ we can calculate the scattered wave resulting from any incident wave giving the explicit form of F as [20, 21]

$$Fu_\infty^i(\hat{\mathbf{x}}) = \int_S u_\infty^i(\hat{\mathbf{r}})v(\hat{\mathbf{r}}, \hat{\mathbf{x}})d\hat{\mathbf{r}}. \quad (2.41)$$

$v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ can be measured experimentally or calculated (assuming $o(\mathbf{x})$ is known) by transmitting all possible incident plane waves, and recording the resulting values of $u_\infty^s(\hat{\mathbf{x}})$ in each case, thus allowing F to be found.

In far-field imaging, the far-field waves are the only data that is accessible, and all knowledge about the scatterer is obtained by transmitting waves from the far-field and measuring the perturbation to the incident field. If we have complete knowledge of F we can calculate the result of any scattering experiment, and can therefore learn all information that can be learnt by performing actual scattering experiments. Since $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ uniquely defines F , it must also contain all information about the scatterer accessible from the far-field. We can conclude that any imaging algorithm with access to $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ or F cannot be improved by additional scattering measurements. Note that this is assuming that scattering is linear, and neglecting random noise—images can of course be improved by averaging repeated measurements of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ to reduce the noise level [7].

2.2.5 Spaces and Operators

There are a number of spaces and operators that map between them that are involved in the scattering problems. We have already seen that the far-field operator F maps incident far-field waves to scattered far-field waves. These waves correspond to a

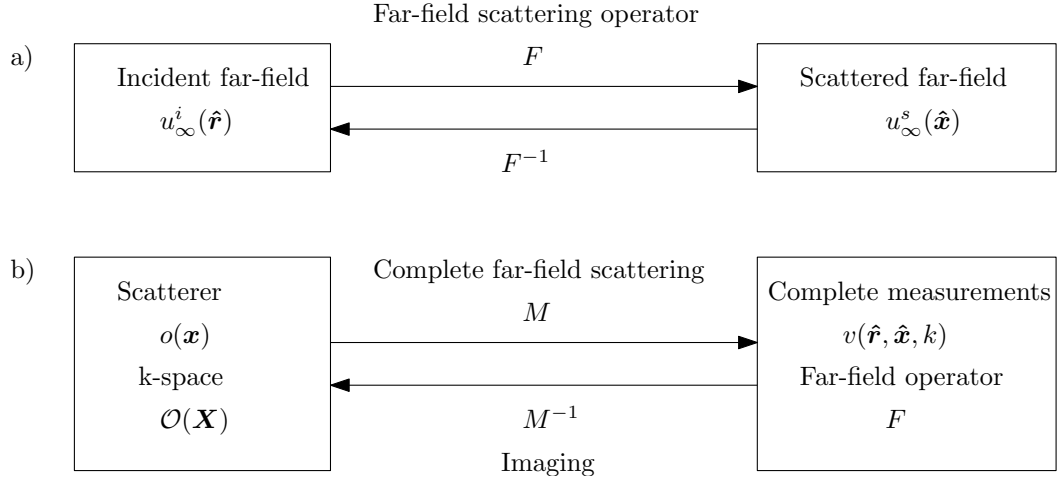


Figure 2.3: a) The far-field operator, F maps any incident far-field to the corresponding scattered far-field for a particular scatterer. b) The M operator maps the scatterer (represented by o or in the frequency domain by \mathcal{O}) to the complete set of far-field measurements v , and therefore also F which depends only on v .

single scattering experiment, but F itself is defined by the result of *all* possible scattering experiments which depend only on $o(\mathbf{x})$. Therefore there exists a mapping from $o(\mathbf{x})$ to $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ ⁴ (and therefore F) which we shall call the *complete* far-field scattering operator M such that

$$v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k) = Mo(\mathbf{x}). \quad (2.42)$$

Its inverse is the imaging operator M^{-1} . Unlike F , these operators are non-linear and M^{-1} is ill-posed and generally hard to calculate. These mappings are shown in figure 2.3. The problem of calculating M^{-1} is *inverse scattering*, or imaging, and its formulation is the subject of the rest of this thesis. In practice, a complete measurement of v is impossible and only a sub-set is known. This causes M^{-1} to be increasingly ill-posed [21].

Because F operates from the far-field, it can be decomposed into three distinct sub-operators. These sub-operators are formed by segregating space into two regions along the boundary of the scatterer, $\partial\mathcal{D}$. Because the field in a region is determined by the value and normal derivative of the field on its boundary, the fields inside

⁴Henceforth we have made the dependence of v on k explicit.

and outside \mathcal{D} are coupled by the value and normal derivative of the field on $\partial\mathcal{D}$. Therefore we can decompose F into an operator that maps the far-field incident waves to $\partial\mathcal{D}$, then to the scattered wave field on $\partial\mathcal{D}$, and finally to the far-field scattered waves. The first sub-operator is the Herglotz operator, H^* and was already given in equation (2.38). It maps an incident far-field wave with density $u_\infty^s(\hat{\mathbf{r}})$ to the incident field on the boundary of the scatterer.

$$H^*u_\infty^s(\hat{\mathbf{r}}) = u^i(\mathbf{x}) \quad \mathbf{x} \in \partial\mathcal{D} \quad (2.43)$$

The second operator, S , maps $u^i(\mathbf{x})$ on $\partial\mathcal{D}$ to a continuous distribution of monopole and dipole sources on $\partial\mathcal{D}$ that produce the scattered field. Its form depends on the type of scatterer and the scattering model. For impenetrable scatterers it is exactly the operator implemented by the Boundary Element Method (BEM). This results in a set of point sources on the boundary of scatterer that result in a scattered field that satisfies the boundary conditions on \mathcal{D} :

$$Su^i(\mathbf{x}) = g(\mathbf{x}) \quad \mathbf{x} \in \partial\mathcal{D} \quad (2.44)$$

The final stage is propagation of the scattered field on $\partial\mathcal{D}$ back to the far-field. This is simply given by

$$Hg(\mathbf{x}) = u_\infty^s(\hat{\mathbf{r}}) \quad \mathbf{x} \in \partial\mathcal{D} \quad (2.45)$$

where H is the adjoint of H^* and is calculated as follows

$$u_\infty^s(\hat{\mathbf{r}}) = \int_{\partial\mathcal{D}} g(\mathbf{x})G_\infty(\hat{\mathbf{r}}, \mathbf{x})d\mathbf{x} \quad \mathbf{x} \in \partial\mathcal{D} \quad (2.46)$$

Therefore the far-field scatterer F is decomposed as




Figure 2.4: Factorisation of the far-field operator.

$$F = HSH^*. \quad (2.47)$$

This factorisation is shown in figure 2.4, and is used for an inverse scattering method called the *Factorisation Method*. It is useful because H only depends on $\partial\mathcal{D}$, so if we can find H (and it turns out we can), then we can determine $\partial\mathcal{D}$ and produce an image of the support of the scatterer. For impenetrable scatterers this provides a complete image of the scatterer. This method will be described in a later chapter. The next section discusses inverse scattering using linear methods and penetrable scatterers, for which we attempt to recover $o(\mathbf{x})$ rather than just $\partial\mathcal{D}$.

2.3 Inverse scattering

The inverse scattering—or imaging—problem is to construct an explicit form of M^{-1} that allows calculation of $o(\mathbf{x})$ given $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$. The mapping between $o(\mathbf{x})$ and $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ is unique [12, 22]. This uniqueness means that given exact knowledge of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ we can exactly reconstruct $o(\mathbf{x})$ with unlimited resolution and accuracy. However in reality our measurement of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ is limited by finite sampling and corruption of the measurement by noise and sensor inaccuracies, therefore we do not have exact knowledge of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ and cannot exactly recover $o(\mathbf{x})$. A small

amount of noise can have a large effect on our estimation of $o(\mathbf{x})$ because M^{-1} is ill-posed [12]. This means that the value of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ does not vary continuously with $o(\mathbf{x})$ and a small uncertainty in the value of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ (due to noise) can lead to an arbitrarily large uncertainty in the value of $o(\mathbf{x})$ resulting in significant ambiguity as to which value of $o(\mathbf{x})$ caused the measured $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$.

In addition to the problems of ill-posedness, the non-linearity of M^{-1} makes an explicit formulation difficult, and to date none have been found. A simple implicit method for calculating $M^{-1}(v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k))$ is to treating it as an iterative optimisation problem. $M(o(\mathbf{x}))$ is simple to calculate, so we can attempt to find an estimate of $o(\mathbf{x})$, $\hat{o}(\mathbf{x})$ that minimizes

$$\hat{o}(\mathbf{x}) = \underset{\hat{o}(\mathbf{x})}{\operatorname{argmin}} \|M(\hat{o}(\mathbf{x})) - v(\hat{\mathbf{r}}, \hat{\mathbf{x}})\| \quad (2.48)$$

under a suitable norm. However the $o(\mathbf{x})$ search space is generally vast and contains many local minima. Despite this, and the time-consuming nature of repeatedly calculating $M(\hat{o}(\mathbf{x}))$ there are successful implementations of iterative methods. These are often made feasible by using a simplified version of M and hierarchical search techniques in which progressively high resolution estimates of o are sought [23, 24]. Because of the problems of using iteration, classical imaging methods provide closed form solutions to linearised approximations of M .

2.3.1 Classical Imaging Methods

In order to make non-iterative imaging tractable, a linear scattering approximation is often used. The most common are the Born, and Rytov approximations [25]. The Born approximation linearises M by modelling the scattered field as an addition to the total field whose magnitude is small. It provides a linear relationship between the values of the scattered field and $o(\mathbf{x})$. The Rytov approximation assumes that the scattered field may have a large magnitude, but causes a small phase change in the total field, and provides a linear relationship between the phase perturbation

and $o(\mathbf{x})$. Of the two, the Born approximation is more commonly used.

The Born approximation of M , M_{Born} allows a direct non-iterative form of M^{-1} , M_{Born}^{-1} to be explicitly formulated that is well-posed and exact. The Born approximation is based on the observation that for weak scatterers, we can approximate

$$u^s \ll u^i \tag{2.49}$$

in equation (2.39) to give

$$u_{\infty}^s(\hat{\mathbf{x}}) = -k^2 \int_{\mathcal{D}} o(\mathbf{x}') u^i(\mathbf{x}') e^{-ik\hat{\mathbf{x}} \cdot \mathbf{x}'} d\mathbf{x}'. \tag{2.50}$$

This is now linear in $o(\mathbf{x})$, and hence M_{Born} and M_{Born}^{-1} are linear and well-posed [26]. Both operators have simple explicit forms which allow fast but approximate imaging. The simplest exact derivation of M_{Born}^{-1} leads to a classical imaging algorithm known as *Diffraction Tomography* (DT), which at the high frequency limit is equivalent to Filtered Back-Projection from X-ray CT [27] (an alternative way to calculate the DT image is using Filtered Back-*Propagation* [28]). Two linear methods based on the Born approximation will now be discussed in detail: diffraction tomography, and beamforming.

2.3.2 Diffraction Tomography

The Born approximation, specifically the *first-order* Born approximation (because it assumes *single* scattering) allows a direct mapping from v to the spatial frequencies of $o(\mathbf{x})$ which is called the k-space.

$$\mathcal{O}(\mathbf{X}) = \mathcal{F}[o](\mathbf{X}) \tag{2.51}$$

where \mathcal{F} denotes the Fourier transform. If we consider an incident plane wave

$u^i(\mathbf{x}') = e^{ik\hat{\mathbf{r}}\cdot\mathbf{x}'}$ then from equation (2.50) the scattering amplitude under the Born approximation is

$$v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k) = -k^2 \int_{\mathcal{D}} o(\mathbf{x}') e^{ik(\hat{\mathbf{r}}-\hat{\mathbf{x}})\cdot\mathbf{x}'} d\mathbf{x}' \quad (2.52)$$

which is proportional to $\mathcal{O}(k(\hat{\mathbf{r}} - \hat{\mathbf{x}}))$. Therefore we have the direct mapping

$$v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k) \propto \mathcal{O}(k(\hat{\mathbf{r}} - \hat{\mathbf{x}})). \quad (2.53)$$

This is the basis of DT. It shows that for a particular illumination direction $\hat{\mathbf{r}}$ and wavenumber k , values of $\mathcal{O}(\mathbf{X})$ on a sphere (or circle in 2D) known as the Ewald sphere (or circle) can be measured (shown in 2D in figure 2.5). By sampling $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ at a number of points we can determine the values of $\mathcal{O}(\mathbf{X})$ for many values of \mathbf{X} and inverse Fourier transform this data to obtain $o(\mathbf{x})$. The maximum spatial frequencies of $o(\mathbf{x})$ are obtained when $\hat{\mathbf{x}} = -\hat{\mathbf{r}}$ and lie on the *Ewald limiting sphere* (or circle).

Implementation of DT requires finding the values of $\mathcal{O}(\mathbf{X})$ on a regular grid of points. This is not trivial because a regular sampling of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ maps to an irregular sampling of $\mathcal{O}(\mathbf{X})$ in the spatial frequency domain, so the Fast Fourier Transform cannot be used directly and values must be interpolated, either in the $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ or $\mathcal{O}(\mathbf{X})$ domains. Therefore the complete DT algorithm is (shown in figure 2.6)

1. Transmit and receive many plane waves.
2. Each transmission, reception and wavenumber combination directly gives the value of a point in the k-space.
3. By transmitting and receiving many waves, the k-space is populated on an irregular grid.
4. The samples are interpolated onto a regular grid, which can be inverse Fourier transformed to obtain o .

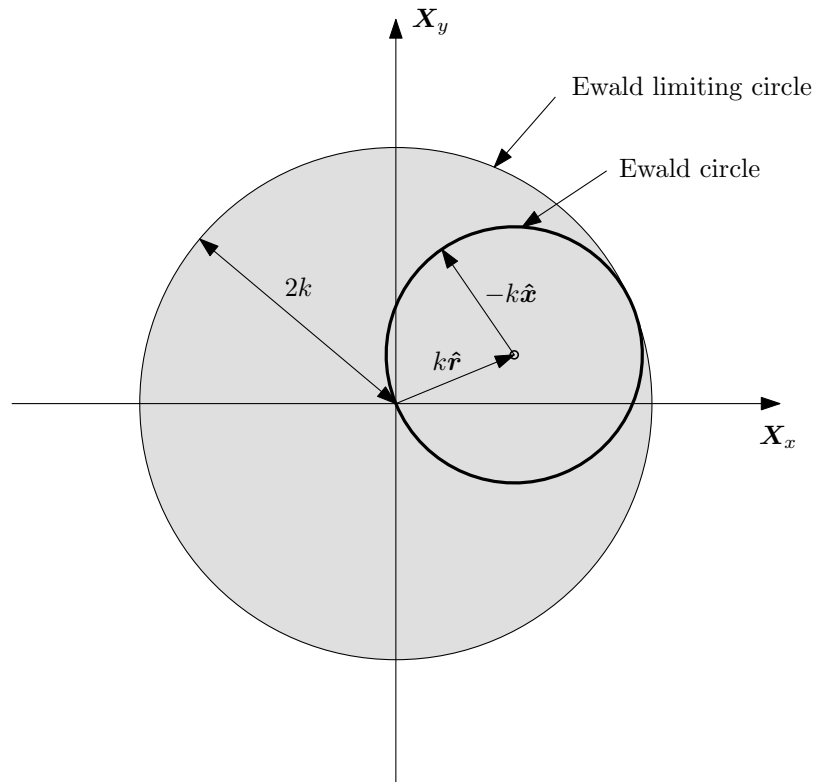


Figure 2.5: *The Ewald circle and the Ewald limiting circle. Outside the Ewald limiting circle, spatial frequencies of the scatterer cannot be measured using DT. But within it, every point can be measured by illuminating and observing the scatterer from many directions \hat{r} and \hat{x} .*

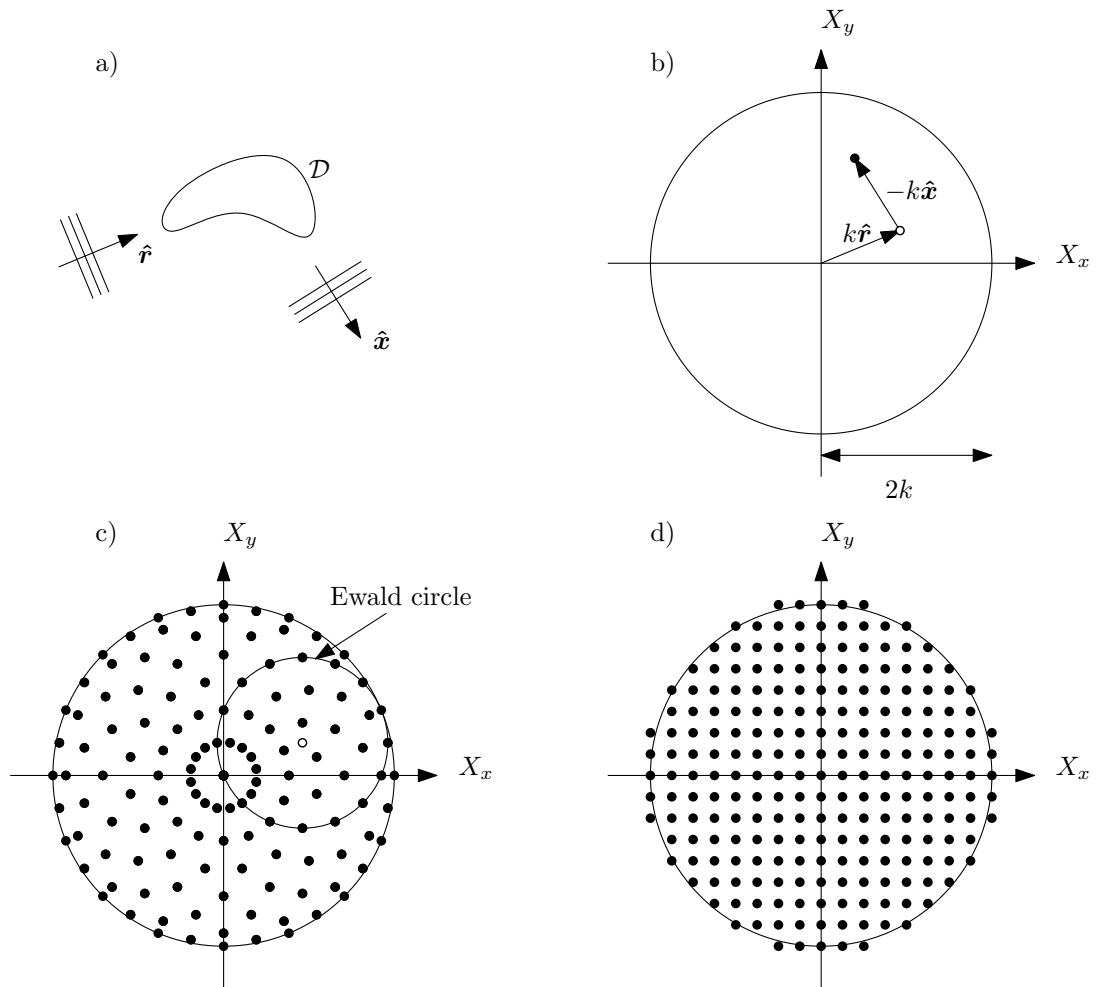


Figure 2.6: Scattering measurements and interpolation in the k -space for diffraction tomography. See the main text for a description.

An implementation that is very fast for linear arrays, where only a subset of v is sampled, is the wavenumber algorithm [29]. This also takes into account the production of plane waves from a linear array of point sources.

DT provides an exact method for calculating M_{Born}^{-1} , and therefore an optimal reconstruction of $o(\mathbf{x})$ where the Born approximation applies. In order for the Born approximation to be valid, we require $u^s(\mathbf{x}) \ll u^i(\mathbf{x})$. Several conditions for the satisfaction of this criterion have been proposed [30]. The simplest is [31]

$$o(\mathbf{x})kL \ll 1 \tag{2.54}$$

where L is the characteristic size of the scatterer. An extension of this is [32]

$$o(\mathbf{x})\{1, kL, (kL)^2\} \ll 1 \tag{2.55}$$

which means that the scatterer not only has to be weakly scattering ($o(\mathbf{x}) \ll 1$), but also small compared to the wavelength. In practice, DT is rarely used. Instead imaging based on *Beamforming* (BF) is used, which does not produce accurate images even if the Born approximation is satisfied, but is much easier to implement and understand.

2.3.3 Beamforming

Beamforming is a principle that is useful in all wave-related fields. It is the process of producing a focused wave (“forming a beam”) through the use of an array of phased transmitters (or receivers for focusing in reception) [7, 8, 33]. This is used in radar [8], Synthetic Aperture Radar (SAR) and Synthetic Aperture Sonar (SAS) imaging [34], microwave telecoms [35], acoustic imaging [36], and essentially any situation involving an array of wave transmitters or sensors. There are also a number of similar algorithms based on the same principle, such as the Synthetic Aperture Focusing Technique (SAFT) [37].

Physically, beamforming⁵ works by sending a focused wave to each point in the object, and focusing on reflections from that point. The corresponding pixel in the image is coloured according to the magnitude of reflections. A beam is formed that focuses on a point by transmitting a wave packet from each transmitter, delayed such that they arrive at the focal point simultaneously. An example of this is shown in figure 2.7 for a line of point sources. To focus on reflections from a point, delays are applied to the recorded signals so that the reflections arrive at each receiver simultaneously. This process is often performed electronically using a delay-and-sum beamformer, but its results can also be calculated numerically using F . The value of the BF image at a point is given in the frequency domain by⁶

$$\hat{o}_{BF}(\mathbf{x}') = - \int_S \int_S v(\hat{\mathbf{r}}, \hat{\mathbf{x}}) \left[k^2 G_\infty^*(\hat{\mathbf{r}}, \mathbf{x}') G_\infty(\hat{\mathbf{x}}, \mathbf{x}') \right]^{-1} d\hat{\mathbf{r}} d\hat{\mathbf{x}} \quad (2.56)$$

$$= -k^{-2} \int_S \int_S v(\hat{\mathbf{r}}, \hat{\mathbf{x}}) G_\infty(\hat{\mathbf{r}}, \mathbf{x}') G_\infty^*(\hat{\mathbf{x}}, \mathbf{x}') d\hat{\mathbf{r}} d\hat{\mathbf{x}}. \quad (2.57)$$

In this context the Green's functions are called the steering functions or focal laws, as they produce the phased waves that steer the wave to the focal point at \mathbf{x}' . It has been shown in [38] that for a full-view configuration, the estimate of $o(\mathbf{x})$ produced by BF, $\hat{o}_{BF}(\mathbf{x})$ is a filtered version of $\hat{o}_{DT}(\mathbf{x})$ and therefore the DT image may be obtained by deconvolving the BF image. This is the approach used by HARBUT [39], because it allows calculation of the DT image in a medium with slowly varying background sound speed. However, since $\mathcal{O}_{DT}(\mathbf{X}) = 0 \forall X > 2k$, both methods have the same fundamental resolution limit.

⁵A short note on nomenclature: when we use the word “beamforming” in this thesis it refers to *imaging using beamforming*. Strictly speaking this usage is not correct as beamforming can be used for other purposes.

⁶Note that the simplification of the inverse only applies because we are using the far-field measurements, and G_∞ (equation 2.37) does not include the geometric beam-spreading term. If imaging using near-field transducers we must use the G_∞^{-1} form.




Figure 2.7: *An example of time-domain beamforming. The images show successive points in time as a wave is transmitted from a line of point sources and converges on a focal point. The transmission by each point on the line is delayed and apodized in order to produce a wave that converges on the focal point in the forth image. The delays are directly calculated from the distance from the focal point to each source point.*

2.4 Resolution Limits

From our previous discussion of DT it was found that the maximum spatial frequencies of $o(\mathbf{x})$ are obtained when $\hat{\mathbf{x}} = -\hat{\mathbf{r}}$ and lie on the *Ewald limiting sphere* (or circle). This limit means that spatial frequencies with $|\mathbf{X}| > 2k$ cannot be retrieved. One definition for the resolution of an image is the maximum spatial frequency content of the image. Using this definition any image formed assuming the Born approximation will have a resolution limit of $2k$.

The reason for this resolution limit can be seen intuitively by considering the scattering of a single plane wave from a collection of point scatterers, as shown in figure 2.8. Under the Born approximation, each scattered wave is independent and the phase of the scattering from each point depends only on the total distance $a_i + b_i$ (assuming the scatterers have the same scattering coefficients). The scattered waves with similar $(a_i + b_i) \bmod \lambda$ will sum coherently. These points correspond to the interference pattern between the planar wave fronts. Any spatial variations that occur at a different wavenumber to the interference pattern will not contribute to the scattered wave in the $\hat{\mathbf{x}}$ direction. The interference pattern with maximum wavenumber occurs for pulse echo measurements where the wavenumber is $2k$.

The Born approximation that led to the resolution limit results from the neglect of multiply scattered waves. We can see how multiply scattered waves affect the BF image by considering the Singular Value Decomposition (SVD) of F , which is given by

$$F = \sum_{n=1}^{\infty} \sigma_n |\xi_n\rangle \langle \psi_n| \quad (2.58)$$

$$F|\psi_n\rangle = \sigma_n |\xi_n\rangle \quad (2.59)$$

$$F^*|\xi_n\rangle = \sigma_n |\psi_n\rangle \quad (2.60)$$




Figure 2.8: Scattering from independent scatterers to the far-field depends only on the distances a_i and b_i (and the scattering coefficients of the scatterers). See the main text for a complete description.

where F^* is the conjugate transpose of F and the ket symbol is used according to Dirac's bra-ket notation [40]. Note that F is symmetric due to reciprocity, and therefore $\bar{F} = F^*$ where \bar{F} is the complex conjugate of F . Now it can be seen why the left singular vectors (ξ_n) with non-zero σ_n correspond to the range of F . It can be seen that equation (2.59) corresponds to a normalised incident far-field wave $u_\infty^i = |\psi_n\rangle$ which results in a scattered wave $u_\infty^s = |\xi_n\rangle$ with magnitude σ_n . Therefore we can immediately see that when $\sigma_n = 0$, $|\psi_n\rangle$ is a non-scattering wave field. The other singular vectors correspond to independent wave fields that capture different aspects of the scattering. These are sorted by the magnitude of the reflection, σ_n which corresponds to the scattering strength of the features each incident wave field interacts with.

The singular values of an example scatterer are shown in figure 2.9. The F operator was measured by simulated scattering (using BEM) of the kite and circle scatterer shown in the figure and the first 100 singular values are plotted on a log scale. The scatterer is a void, with the boundary condition $u(\mathbf{x}) = 0; \mathbf{x} \in \partial\mathcal{D}$, and is approximately 3λ across. Therefore the first few singular values will produce fields

that strongly scatter from the large features of scatterer (i.e. its overall shape), and will be a low order field. These are shown in figure 2.10. Two things that can be seen from this figure. First, the lowest singular values (the first two are shown) produce incident fields with high magnitude on the boundary of the scatterer and low spatial frequency which scatter strongly. The small singular values produce fields that are very weak on the scatterer. As this makes it hard to see the actual field values on the scatterer, these have been plotted separately in figure 2.11. The figure shows the incident field values produced by the right singular vectors ($|\psi_n\rangle$) on $\partial\mathcal{D}$. From this figure we can also see that as n increases and the magnitude of σ_n decreases, the incident field on the scatterer not only increases in spatial frequency, but has more energy concentrated in the occluded surfaces between the kite and the circle.

An alternative interpretation of the SVD is found by considering a random incident field. This random field can be decomposed into the sum of the right singular vectors ($|\psi_n\rangle$) with random phase, and hence the resulting field is the sum of $\sigma_n|\xi_n\rangle$ with random phase. Therefore, the singular values correspond to how much of the reflection of a random incident field is made up of that left singular vector. Since a random incident field will produce little energy between the circle and kite (because it is occluded), it will not affect the scattering much. The small details that do not affect the scattered field much are encoded in small singular values (and their associated vectors).

We can now see how the components of F are weighted in the BF:

$$\begin{aligned}\hat{o}_{BF}(\mathbf{x}) &= \langle g_{\mathbf{z}}|F|g_{\mathbf{z}}\rangle \\ &= \sum_{n=1}^{\infty} \sigma_n \langle g_{\mathbf{z}}|\xi_n\rangle \langle \psi_n|g_{\mathbf{z}}\rangle\end{aligned}$$

where $|g_{\mathbf{z}}\rangle = G_{\infty}(\hat{\mathbf{x}}, \mathbf{z})$. Therefore the BF image is dominated by the fields with large singular values, and hence the small singular values which contain high resolution and occluded detail will be lost. The retrieval of this high resolution data is called *super-resolution*.

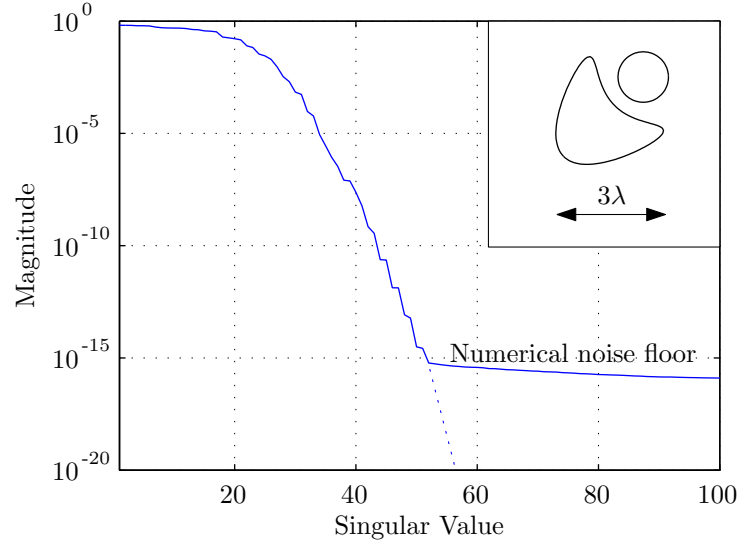


Figure 2.9: The first 100 singular values corresponding to scattering from the sound-soft “kite” and circle scatterer shown. It can be seen that the singular values quickly become extremely small. Theoretically the downward trend would continue, however the addition of noise raises the singular values to an approximately constant level. In this case the noise is actually caused by the finite numerical precision.

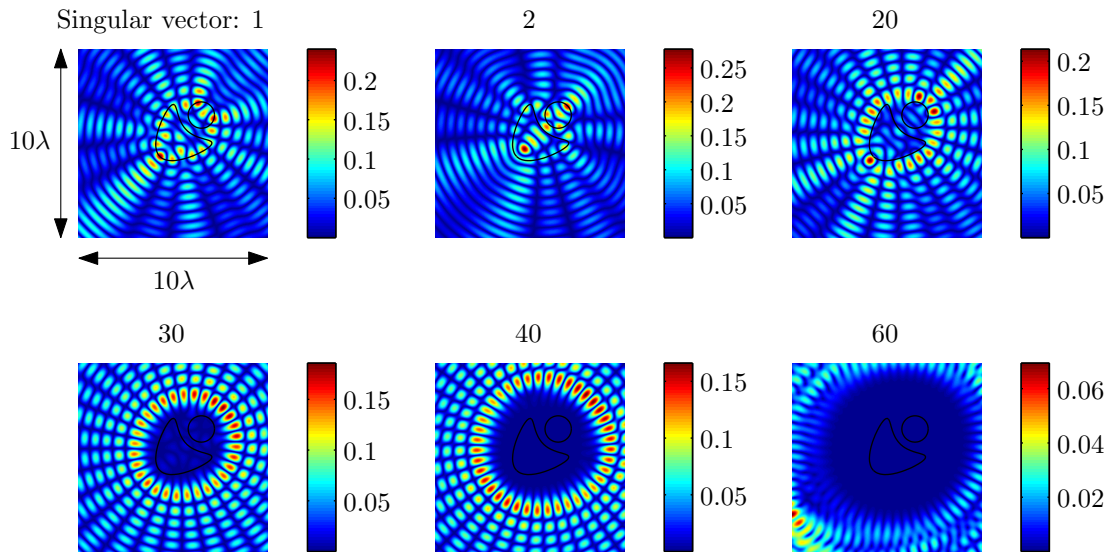


Figure 2.10: The magnitude of the incident fields produced by some right singular vectors ($|\psi_n|$) for $n = 1, 2, 20, 30, 40, 60$. The higher magnitude singular values produce fields with higher incident values on the scatterer surface which therefore scatter strongly. The magnitude of reflection is determined by the singular value, so zero singular values produce non-scattering fields. Note that the singular vectors are normal, so the power of the transmission is constant.




Figure 2.11: *The real component of the incident field on the surface of the scatterer for fields shown in figure 2.10. Note that the scale of each sub-plot is different, and that the fields are more rapidly varying for small singular values. They also tend to have more energy concentrated within the occluded area between the two scatterers.*

2.5 Super Resolution

One way to achieve super-resolution (the measurement of $\mathcal{O}(\mathbf{X} > 2k)$) is by measuring the near-field of the scattered waves. The near-field can contain evanescent waves which have temporal frequency ω , but spatial wavenumbers greater than $\frac{\omega}{c}$. These properties mean that they cannot propagate losslessly through free-space. Evanescent waves are generated when the field values on the boundary of the scatterer vary rapidly enough that no plane waves can exist with compatible field values on the boundary.

Although they may sound exotic, evanescent waves are not strictly a distinct phenomenon from normal scattering, and can only truly occur with infinite boundaries. The simplest case to consider is a line of point sources at $y = 0$ in 2D. If the monochromatic source at a point $\mathbf{x} = (x, 0)$ is given by e^{iax} (where a is a constant) with a frequency ω and sound-speed in free space of c then we can calculate the field at a point $\mathbf{r} = (0, h)$ by integrating along the line as follows

$$u(\mathbf{r}) = \int_{-\infty}^{\infty} G_{2D}(\mathbf{x}, \mathbf{r}) e^{iax} dx \quad (2.61)$$

The result of this integration depends on a . For $a > \frac{\omega}{c}$ the magnitude of $u(\mathbf{r})$ rapidly decays with increasing h and the resulting wave-field is insignificant more than a few wavelengths from $y = 0$. As with the far-field DT resolution limit, the evanescent waves can be seen as the superposition of phased wave sources that sum destructively in the far-field. The field resulting from $a = 5$ and $\frac{\omega}{c} = 3$ is shown in figure 2.12.

A physical interpretation of this phenomenon can be seen by considering the wave field generated for $a < \frac{\omega}{c}$. In this case, a plane wave will result with an angle that is geometrically compatible with the e^{iax} phasing. Such a plane wave is only possible when $a \leq \frac{\omega}{c}$, therefore no plane wave can be a solution for $a > \frac{\omega}{c}$. Because these high resolution waves decay exponentially with distance, they must be measured within a wavelength or so of their source. Direct measurement of the evanescent waves is the approach taken by Near-field Scanning Optical Microscopy (NSOM) and other optical

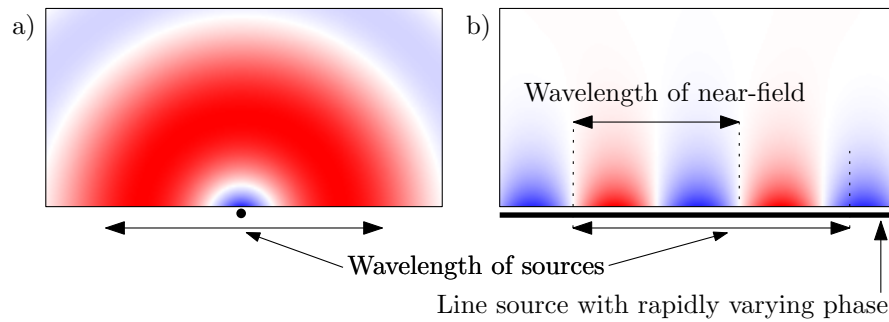


Figure 2.12: A demonstration of the formation of evanescent waves. a) A single harmonic point source generates an expanding circular wave with the labelled wavelength. b) The single point is replaced with a continuous line of point sources which have phases that vary faster than the natural wavelength in the medium. This generates a near-field of evanescent waves which cannot propagate. The rapid phasing of the sources mean that the waves they produce sum destructively in the far-field.

microscopy super-resolution methods. A material with a low speed of sound, c_1 is introduced into the evanescent field such that $\frac{\omega}{c} < a < \frac{\omega}{c_1}$. This material has an $\frac{\omega}{c}$ which is higher than the wavenumber of the evanescent waves, so they can propagate through it and be measured. This principle is illustrated in figure 2.13, where a plane wave travels from a low sound-speed material to a high sound-speed material. Because no plane waves with compatible boundary conditions can propagate, an evanescent field of waves travelling along the surface is produced. These wave decay rapidly away from the boundary, but before they fully decay they reach another material (the NSOM probe) which has a low sound speed. The evanescent waves produce radiating plane waves in this material that can be transmitted to a sensor and recorded.

The result is that although plane waves cannot propagate in the background medium, they are able to traverse small gaps via evanescent waves, as long as the gap is less than a wavelength or so. This traversal is called “tunnelling” because the same evanescent coupling in the Schrödinger wave equation leads to quantum tunnelling where particles can tunnel through a thin barrier in which there are no propagating solutions to the wave equation. In NSOM a sharp glass optical fibre is scanned very close to the surface of the object in order to allow the evanescent waves to radiate

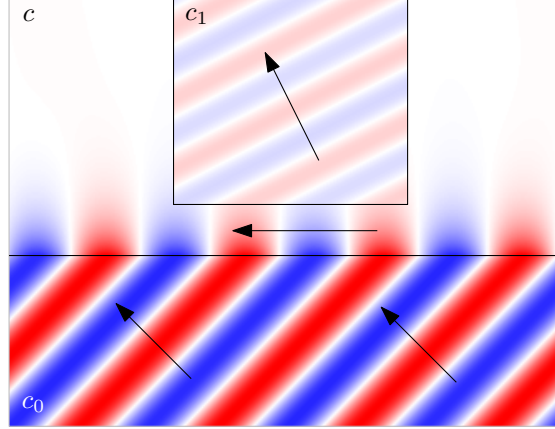


Figure 2.13: *The mechanism of NSOM. An evanescent field is produced in the background medium. If another medium with a low enough sound speed c_1 is introduced into the evanescent field, it will produce radiating waves that can propagate towards a sensor. In NSOM the probe is much narrower than shown here—less than a wavelength—so that high wavenumber evanescent fields can be measured.*

along the fibre to an optical sensor. It is the multiple scattering of the scattered wave from the surface of the object and the NSOM probe that allows the evanescent waves to be measured in the far-field; if the evanescent waves only scattered once, by definition they could not interact with the probe and would decay exponentially before reaching the far-field where they can be measured.

NSOM techniques can achieve resolutions of $\lambda/20$ [41], or up to $\lambda/14$ using the recently reported white-light nanosphere technique [42]. Resolutions of up to $\lambda/1000$ have been achieved with near-field microwave imaging [43]. However for sub-surface imaging, we can only measure the far-field because we can only excite and measure the surface of the component, which in general is far from the scatterers. Therefore we cannot use this near-field approach to achieve super-resolution.

Neither of the far-field techniques discussed consider multiple scattering, which can be a significant effect (see figure 2.14 for a numerical example). It has been shown explicitly by Simonetti [20] that when multiple scattering is considered, the value of $v(\hat{\mathbf{r}}, \hat{\mathbf{x}}, k)$ depends not on a single value of \mathcal{O} , but on *all* values of \mathcal{O} . This corroborates the various uniqueness results for forward scattering, which imply that all of \mathcal{O} must be encoded into v , and therefore there is no fundamental resolution limit. Since a




Figure 2.14: An illustration of multiple scattering between two circular scatterers. A harmonic wave is incident from above the scatterers. In (a) the scattered field from each circle is calculated separately, and then the fields are summed. In (b) the correct scattered wave is found by considering the two circles together. (c) shows the difference between (a) and (b) which corresponds to the waves that scatter multiple times between the circles. It can be seen that the multiply scattered waves originate between the scatterers, as would be expected, and that they have a significant effect on the value of the scattered wave field.

single-scattering model leads to a resolution limit, and a multiple-scattering model does not we can deduce two things: first, it must be multiply scattered waves that contain high resolution information, and second, super-resolution will be harder to achieve with weak scatterers.

2.6 Conclusion

By formulating the forward scattering and imaging problems as a mapping between the measurements u^s and the object o we can see that a simplified linear forward model gives rise to a resolution limit when imaging. The linear model excludes multiply scattered waves which contain information about the high frequencies of o . We have also defined a number of useful entities for scattering:

- $o(\mathbf{x})$ - The object function, which describes the physical properties of the scatterer.
- $\mathcal{O}(\mathbf{X})$ - The Fourier transform of the object function (the k-space).

- $u^i(\mathbf{x}), u_\infty^i(\hat{\mathbf{x}})$ - The incident field, and incident far-field.
- $u^s(\mathbf{x}), u_\infty^s(\hat{\mathbf{x}})$ - The scattered field, and scattered far-field.
- $u(\mathbf{x})$ - The total field.
- $v(\hat{\mathbf{r}}, \hat{\mathbf{x}})$ - The scattered wave in direction $\hat{\mathbf{x}}$ for an incident wave from $\hat{\mathbf{r}}$.
- F - The far-field operator, which maps $u_\infty^i(\hat{\mathbf{r}})$ to $u_\infty^s(\hat{\mathbf{x}})$.
- HSH^* - The factorisation of the far-field operator.
- $\sigma_n, |\psi_n\rangle, |\xi_n\rangle$ - The SVD of the far-field operator.
- M - The operator which maps o to v (and F).
- M^{-1} - The (perfect) imaging operator.
- M_{Born}^{-1} - The Born approximation imaging operator (as calculated by DT and equivalent techniques).

In the next chapter we will look in detail at the information content of v in real situations where it is sampled on a discrete grid.

Chapter 3

The Discrete Inversion

3.1 Background

In the previous chapter all parameters of the imaging process were assumed to be continuous whereas in real digital systems they are discrete. This discretisation results in fundamental limits on the capability of an imaging system. The discretisation occurs in the domains of space, time, and magnitude. First, it was assumed that the incident and scattered fields could be excited and measured at arbitrary locations in the far-field. In practice, waves cannot be excited or sensed in a continuum of independent points, but instead with a discrete set of transmitters and receivers. Typically ultrasound arrays have 32, 64 or 128 transducers, although much higher transducer counts are possible [44, 45]. In addition to this spatial discretisation there is a temporal discretisation. The fields cannot be measured continuously in time, but again at a discrete set of times whose number depend on the recording length and the sampling rate. Sampling rates up to 100 MHz are typical [10, 46].

Finally, it was assumed the field values could be measured exactly, whereas real sensors have a finite dynamic range and are subject to noise. The digital signals recorded by a typical ultrasound system are 8-16 bits, which causes quantisation noise in addition to the Gaussian process noise. Quantisation noise can be minimised by varying the amplification of the analogue signal such that it spans the full range

of the ADC (Analogue to Digital Converter), and for high bit depths is usually less significant than random noise. If the object is static, then it is also possible to repeat the measurement with several gains and combine them to form a single measurement with high dynamic range (this is the same technique used in High Dynamic Range photography, where several exposures of a single scene are taken).

This discretisation of measurement in time, space and amplitude, puts a fundamental limit on the amount of information that can be stored in the recorded signals. In turn, this determines a limit on the accuracy and resolution of the image of the scatterer. The information capacity of the measurements, and the number of degrees of freedom of an image have been studied in the past [47, 48, 49]. The aim of this chapter is to further this work by looking at this distribution of information in the measurements in terms of the measurement parameters, to see how they affect the limits on image quality. We begin in section 2 by discussing the concepts of information capacity and the degrees of freedom of an image, and how they are related.

There is no single equation that can tell us the information capacity of a complete set of measurements, however we can derive upper bounds on the information capacity of the signals when looked at from two different perspectives—the temporal diversity of a single time domain signal, and the spatial diversity of a monochromatic array measurement. These can give an insight into the true bounds on the whole measurement. This is done in sections 3 and 5. Similarly to the Nyquist criterion, these bounds can give us insights into measurement criterion for broadband array measurements and tell us what to measure in order to record all the available wave field information.

3.2 Degrees of freedom of an image

The end result of inverse scattering is an image of the scatterer. In computing, the image is represented by a rectangular grid of independent samples of the underlying image function. However, this is not the only possible representation: images are often encoded in the frequency domain (e.g. JPEG) or the wavelet domain (JPEG2000).

If each image sample is independent, then the image has that many degrees of freedom. For example, a 10×10 pixel image has 100 degrees of freedom. If we wish to exactly reconstruct a 10×10 pixel image we must have measurements with also have at least 100 degrees of freedom. Conversely if our measurements have 100 degrees of freedom, we will only be able to reconstruct an image with at most 100 pixels (or frequency or wavelet coefficients). Therefore the degrees of freedom of the measurement imposes a limit on the resolution of the image.

The information capacity of the measurements—i.e. the amount of information, in bits, that it can contain—is a related but distinct concept. An exact measurement of a single value can contain infinite information, but has only one degree of freedom. In fact, arithmetic coding, a type of data compression, explicitly works by encoding all of the data into a single number between 0 and 1 [50]. A larger amount of data requires the number to be specified to more significant figures, and with infinite significant figures an infinite amount of data can be encoded.

When each degree of freedom is known to only finite precision, for example due to noise or quantisation, it has a finite information capacity. If we can distinguish between two different values (e.g. 0 or 1) then the information capacity is one bit. Shannon showed that this information capacity of a single noisy measurement is [51]

$$I = \log_2 \left(1 + \frac{s}{\zeta} \right) \quad (3.1)$$

where $\frac{s}{\zeta}$ is the Signal-to-Noise Ratio (SNR; SNR is a power or energy ratio).

In addition to the degrees of freedom limitation, the information capacity of the output image cannot be greater than the information capacity of the measurements.

The information capacity of the measurements is important because it determines how much detail about the scatterer we can extract. Or more precisely, how many different subsets we can divide all of the possible scatterers into. As an example, suppose we measure a single variable—the total scattered energy from an incident plane wave. Suppose further that noise, or inferior sensor technology means we

can only detect two different levels of energy—high or low. This means that we measure 1 bit of information, and can only distinguish between weak scatterers, and strong scatterers. We have divided scatterer-space into two. With more bits we can distinguish more scatterer classes and provide more detailed information and images.

The fact that the noise level of the measurements can impact the level of achievable image resolution can be appreciated from the information capacity of noisy systems [52, 53]. Since there is an upper bound on the amount of information that our discrete measurements contain, there must be a similar limit on the amount of information that the resulting image contains. It is not easy to draw conclusions about the nature of this limit without knowing how the measurement coefficients affect the image. In other words, we do not know how scatterer-space is divided into distinguishable subsets by the measurements. All we know is that with I bits of information in the measurements, we can distinguish 2^I different subsets of the scatterer space. If we make the rather large assumption that the measurement coefficients correspond to the pixels of the image, then scatterer-space is divided into $X \times Y \times D$ subsets, where X, Y are the image size, and D is the image bit depth. If D is fixed, then it follows that the image's resolution will be limited by the information content of the measurements.

Before proceeding with the analysis we must formulate the discretisation of the far-field operator.

3.2.1 Discretisation of the far-field operator

For discrete scattering measurements, the discrete time-domain far-field operator $F_{d,t}$ maps an incident wave-field composed of time-varying waves incident from a discrete set of directions $\hat{\mathbf{r}}_i = \hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2 \cdots \hat{\mathbf{r}}_{N_r}$ to the scattered waves in another set of directions, $\hat{\mathbf{x}}_i = \hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2 \cdots \hat{\mathbf{x}}_{N_x}$. For simplicity we shall set $\hat{\mathbf{r}}_i = \hat{\mathbf{x}}_i$. In the time domain, we assume only a finite number of samples are recorded at $t_i = t_1, t_2, \cdots, t_{N_t}$, and as sampling is digital, each sample can take one of only N_v values. Therefore $F_{d,t}$ maps

$$F_{d,t}u_{\infty}^i(\hat{\mathbf{x}}, t) = u_{\infty}^s(\hat{\mathbf{x}}, t) \quad (3.2)$$

where we mean that $F_{d,t}$ maps the entire vector $u_{\infty}^i(\hat{\mathbf{x}}, t)$, not just its value at $(\hat{\mathbf{x}}, t)$. $F_{d,t}$ is a convolution operator determined by the impulse response measurements from $\hat{\mathbf{r}}$ to $\hat{\mathbf{x}}$

$$v_{d,t}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t) \quad (3.3)$$

$$F_{d,t}(u_{\infty}^s(\hat{\mathbf{x}}_i, t)) = \sum_{r=1}^{N_x} u_{\infty}^i(\hat{\mathbf{r}}_r, t) * v_{d,t}(\hat{\mathbf{r}}_r, \hat{\mathbf{x}}_i, t) \quad (3.4)$$

where $*$ is the convolution operator. We can now consider the discrete time-domain operator $M_{d,t}$ which maps the object function $o(\mathbf{x})$ to $v_{d,t}$ and $F_{d,t}$. The number of different values that $v_{d,t}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t)$ can have is $N_v^{N_x^2 N_t}$, and by definition the information capacity of such a measurement of $v_{d,t}$ is $I = \log_2(N_v^{N_x^2 N_t}) = N_x^2 N_t \log_2(N_v)$ bits. This is the information capacity if all possible values are equally likely. However, the prior probability of each $v_{d,t}$ is not equal. The prior probability is the probability of each measurement occurring in the real world. It depends on the probability of each scatterer occurring. For example we are unlikely to see a measurement where signals arrive an hour after the transmission, so that value of $v_{d,t}$ has a low prior probability. Suppose the prior probability of encountering an object function $o(\mathbf{x}) = o_0(\mathbf{x})$ is $P(o(\mathbf{x}) = o_0(\mathbf{x}))$, then the measurement probability is

$$P(v_{d,t}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t) = Mo_0(\mathbf{x})) = P(o(\mathbf{x}) = o_0(\mathbf{x})) \quad (3.5)$$

which is not constant. Shannon showed that the true information capacity (the entropy) of a measurement is [51]

$$I = \sum_{v_0} P(v_{d,t} = v_0) \log_2(P(v_{d,t} = v_0)) \quad (3.6)$$

where the sum is performed over all $N_v^{N_x N_t}$ possible values of $v_{d,t}$. The set of possible values of the $v_{d,t}$ function is called the *ensemble* [54]. Most of these possible signals never occur or are very unlikely, therefore the true information capacity is much less than $N_x^2 N_t \log_2(N_v)$, but finding the prior probabilities of $o(\mathbf{x})$ or $v_{d,t}$ is not possible. In order to do so we could in theory perform a survey of a very large number of scatterers found in the real world, and state that “measurements / scatterers similar to these are likely” under some similarity metric. In fact this is exactly what we already do mentally through experience, so we know that cracks are more likely to be encountered than triangular voids for example.

However the measurement space is so large that it would be impossible to obtain a sufficient sampling of the probability function. Instead we can draw on the range of I by transforming $v_{d,t}$ into other domains, such as the frequency domain, in which we have more knowledge about the expected values of $v_{d,t}$. In the following sections we shall look at two decompositions that provide insight into how the information capacity is distributed in $v_{d,t}$. This can be used to determine resolution bounds, and also to determine an optimal sampling strategy. We begin by looking at the information capacity of the measurement of a single monochromatic field in the spatial harmonic domain, measured by an array of spatially diverse sensors.

3.3 Spatial harmonics

The first perspective we can consider is to Fourier transform an outgoing wave field into the frequency domain, and then consider the spatial frequencies of the monochromatic wave fields. The assumption we make here to provide a prior probability on $v_{d,t}$ is that the scatterer is contained within a sphere or circle of radius ρ . Under this assumption we will find that the higher order harmonics can be expected to have low magnitude and contain little energy. Since the information capacity depends on the Signal-to-Noise Ratio (SNR), this will also be low for high order harmonics.

The wave fields outside $r \geq \rho$ can be decomposed into circular/spherical harmonics.

The circular harmonic decomposition of a monochromatic 2D wave field is given in [55] as

$$u(r, \theta) = \sum_{n=-\infty}^{\infty} \left[A_n H_n^{(1)}(kr) + B_n H_n^{(2)}(kr) \right] e^{in\theta} \quad (3.7)$$

where $H_n^{(1)}$ and $H_n^{(2)}$ represent outgoing and incoming waves respectively, so fields made of only outgoing (incoming) waves have the coefficients $B_n = 0$ ($A_n = 0$). The angular component $e^{in\theta}$ is simply the Fourier transform of the field around the circle of radius r , therefore a circular transducer array consisting of N equidistant transducers can measure harmonics up to order N without aliasing. For a spherical harmonics we need $(N + 1)^2$ transducers to measure up to order N [56]. On a sphere the transducers can only be equidistant if they are positioned on the face of a platonic solid (tetrahedron, cube, octahedron, dodecahedron or icosahedron) so this is approximate for other values of N .

The outgoing field on the circle of radius ρ is

$$u^s(\rho, \theta) = \sum_{n=-\infty}^{\infty} A_n H_n^{(1)}(k\rho) e^{in\theta}. \quad (3.8)$$

The corresponding angular component of the far-field is

$$u_{\infty}^s(\theta) = \lim_{r \rightarrow \infty} \frac{1}{\Gamma_{2D}(r)} \sum_{n=-\infty}^{\infty} A_n H_n^{(1)}(kr) e^{in\theta} \quad (3.9)$$

where $\Gamma_{2D}(x) = -\frac{(i+1)e^{ikx}}{4\sqrt{\pi kx}}$ as in equation (35) in chapter 2. The ratio of the magnitude of each harmonic order from the near-field ($u^s(\rho, \theta)$) to the far-field ($u_{\infty}^s(\theta)$) is defined as

$$\alpha_n = \lim_{r \rightarrow \infty} \frac{1}{\Gamma_{2D}(r)} \frac{H_n^{(1)}(kr)}{H_n^{(1)}(k\rho)}. \quad (3.10)$$

If α_n is small then it means that the n^{th} harmonic order does not propagate strongly

to the far-field and is difficult to measure. The asymptotic behaviour of $H_n^{(1)}(x)$ for $x \gg |n^2 - \frac{1}{4}|$ is

$$\lim_{r \rightarrow \infty} H_n^{(1)}(kr) = \sqrt{\frac{2}{\pi kr}} e^{i(kr - \frac{n\pi}{2} - \frac{\pi}{4})} \quad (3.11)$$

therefore

$$\alpha_n = \frac{4e^{-i\pi\frac{n}{2}}}{H_n^{(1)}(k\rho)} \quad (3.12)$$

and the energy (or power) ratio is

$$|\alpha_n|^2 = \frac{16}{|H_n^{(1)}(k\rho)|^2}. \quad (3.13)$$

If the magnitude of each harmonic order on $r = \rho$ is fixed, so $A_n H_n^{(1)}(k\rho) = 1$ then the far-field magnitude of that order is α_n and the power is $|\alpha_n|^2$. The power ratio is plotted in figure 3.1 for various values of ρk . It becomes exponentially small with spatial order where $n > k\rho$. Therefore with sensors with a limited dynamic range, harmonics above order $k\rho$ will be difficult to measure. A standard sampling criterion for circular arrays is based on this limit, and only measures up to order $k\rho$, for which we need the following number of transducers [55, 57]

$$N_x \approx k\rho \quad (3.14)$$

in 2D, and in 3D the criterion is

$$N_x \approx \frac{1}{2}(\pi k\rho)^2. \quad (3.15)$$

It can be seen from equation (3.12) that α_n never becomes zero, so we can conclude that all wave modes actually propagate, and none of the waves are truly evanescent.




Figure 3.1: A plot of $|\alpha_n|^2$ for several values of ρk . The points $n = \rho k$ where α_n starts to become exponentially small are marked. The negative harmonics (not shown) have symmetric values (that is, $|\alpha_n|^2 = |\alpha_{-n}|^2$).

Therefore with infinite transducers we can recover all orders of the wave field. Although all wave modes can be measured in principle, in practice these transducers will have limited dynamic range due to quantisation and noise, and very small values cannot be measured. The optimum sampling strategy depends on the noise level and for sensors with high dynamic range it will be beneficial to sample slightly above $N_x = k\rho$. The effect of dynamic range on the total information capacity of all modes is discussed next.

3.3.1 Dynamic range

If the wave field values can be measured with unlimited precision, then the small expected amplitudes of the high order harmonics do not affect their information capacity (which is infinite with no noise). However in practice measurement of field values cannot be arbitrarily accurate. The ability to measure small signals will be constrained by quantisation, noise, and sensor inaccuracies. Although small signals

can be sometimes measured accurately by amplification and averaging to reduce noise, this is not possible in every case, in particular when small and large signals are present simultaneously.

We can model a single measurement, \hat{x} as a random variable composed of the true value x , and noise n , both of which we assume to be zero mean Gaussians.

$$\hat{x} = x + n \tag{3.16}$$

with

$$x \sim \mathcal{N}(0, s) \tag{3.17}$$

$$n \sim \mathcal{N}(0, \varsigma) \tag{3.18}$$

Therefore from equation 3.1 the total information capacity for a measurement composed of a number of such coefficients is

$$I = \sum_{i=1}^N \log_2 \left(1 + \frac{s_i}{\varsigma_i} \right) \tag{3.19}$$

where s_i , ς_i are the expected signal and noise energies for each coefficient. Therefore using the circular harmonic decomposition of an outgoing wave field, its information capacity is

$$I = \sum_{n=-N_x}^{N_x} \log_2 \left(1 + \frac{|\alpha_n^2|s_n}{\varsigma_n} \right) \tag{3.20}$$

where s_n is the energy of the near-field modes, and ς_n is the energy of the far-field noise modes (which can be assumed to be constant). This function is plotted in figure 3.2, which shows that the maximum information capacity can be significantly higher if sampling is performed at a greater rate than $N_x = k\rho$, and also that




Figure 3.2: Total information capacity of a measurement for various $k\rho$ at $\frac{\sigma}{\zeta} = 10$ for different numbers of transducers, N_x .

maximum total information capacity is linearly proportional to $k\rho$. In the next section we look at how the total information capacity of a monochromatic field changes if the full 360° aperture is not available.

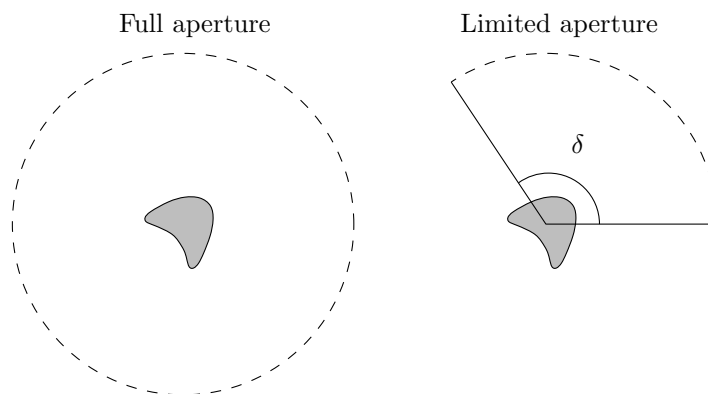


Figure 3.3: In most practical situations the full circle (or sphere) cannot be measured. Instead a subset of angles $[0, \delta)$ are accessible.

3.4 Limited apertures

Practical ultrasound systems often (but not always) are constrained to transmit and receive waves from an aperture that is smaller than the full circle/sphere. If only the far-field subtended by angle δ is available (see figure 3.3), then the limited aperture results in only part of the far-field operator being knowable. Consider (in 2D) the operator Q which maps the near-field to the far-field

$$Qu(\rho, \theta) = u_\infty(\theta) \quad (3.21)$$

given by

$$Q = \mathcal{F}^* \Lambda \mathcal{F} \quad (3.22)$$

where \mathcal{F} is the discrete Fourier transform matrix, \mathcal{F}^* is its inverse, and Λ is a matrix whose diagonal elements are α_n . If the columns of \mathcal{F} are rearranged so that α_n are sorted in descending order, then this is the SVD of Q and its singular values are α_n and its singular vectors are $e^{in\theta}$. The α_n determine the information capacity of full-view measurements (equation 3.20), and for the full view these are equivalent to the singular values of Q . So for limited apertures, instead of α_n we can plot the singular values of Q .

Therefore we can extend the analysis to consider the information capacity of a limited aperture measurement of a wave field by taking the SVD of a limited aperture subset of Q by removing some of its rows. The resulting singular values are plotted in figure 3.4, with every other singular value removed (since half correspond to the negative harmonic orders). It can be seen that the number of high-magnitude singular values depends linearly on the aperture, so a large aperture is very important in maximising the information content of the signals. For more background on using the SVD for information capacity see [48].




Figure 3.4: *Singular values as a function of aperture for $\rho k = 81$ (compare the $\delta = 2\pi$ plot with $\rho k = 81$ in figure 3.1). The number of accessible singular values varies linearly with aperture.*

3.5 Temporal diversity

The second perspective which we consider in this chapter is the frequency domain representation of the individual signals. So far we have been looking at only a single frequency component of the signals. By summing equation 3.20 over frequencies we can obtain an upper bound on the information capacity of a broadband field. However it is simpler to first consider the information capacity of a single signal. In the time domain, each coefficient of equation 3.19 corresponds to a single sample at a point in time, so we have (the same as equation 3.19)

$$I = \sum_{t=1}^{N_t} \log_2 \left(1 + \frac{s_t}{\varsigma_t} \right) \tag{3.23}$$

where N_t is the number of independently measured samples, and s_t and ς_t are the signal and noise energies of each sample. This equation was derived by Shannon in the context of telecoms, where the sample values are chosen in order to maximise

the channel capacity. In that case, the resulting signal will be stationary, in that its statistical properties, including the SNR $\frac{s_t}{\zeta_t}$, are independent of time. However, this is not the case for ultrasound signals. Recall that the ensemble of signals is the set of all possible values the signal can take. The information content of each sample is then determined by the SNR of that sample, across the *ensemble* rather than across time. For stationary processes these are the same. So s_t is the ensemble signal power at time t , and ζ_t is the ensemble noise power, which is likely to be constant, but may not be if Time-Varying Gain (TVG) is used.

Equation 3.23 provides an upper bound on the information content of a signal in terms of the time domain decomposition. But it is only one possible upper bound. We can use any orthogonal decomposition of the signals to try to find the lowest upper bound possible. The best decomposition (resulting in the lowest I) will be one with high SNR concentrated in a few coefficients, rather than distributed evenly over all coefficients.

An important feature of ultrasound signals is that their energy is localised in both time *and* frequency. Therefore a purely time-domain or frequency-domain representation will not result in an optimal energy distribution. Instead a hybrid-domain representation such as the Short Time Fourier Transform (STFT), Modified Discrete Cosine Transform (MDCT) or Discrete Wavelet Transform (DWT) [58] allows us to concentrate the signal energy into as few coefficients as possible and therefore obtain a low upper bound on the information capacity of the signals. In order to demonstrate this, the Continuous Wavelet Transform (CWT) of an example signal, as the wavelet basis is varied from a temporally to harmonically localised form is shown in figure 3.5. Somewhere between the time and frequency domains the signal energy is concentrated in a small area of the domain (figure 3.5c) resulting in a tight upper bound on the information capacity. The figure shows the energy for a single ensemble member (a single signal). In order to calculate $\frac{s}{\zeta}$ for the ensemble we need to determine the probabilistically weighted power across all possible signals in the ensemble.

The frequency component of the ensemble energy is easy to estimate, as it will

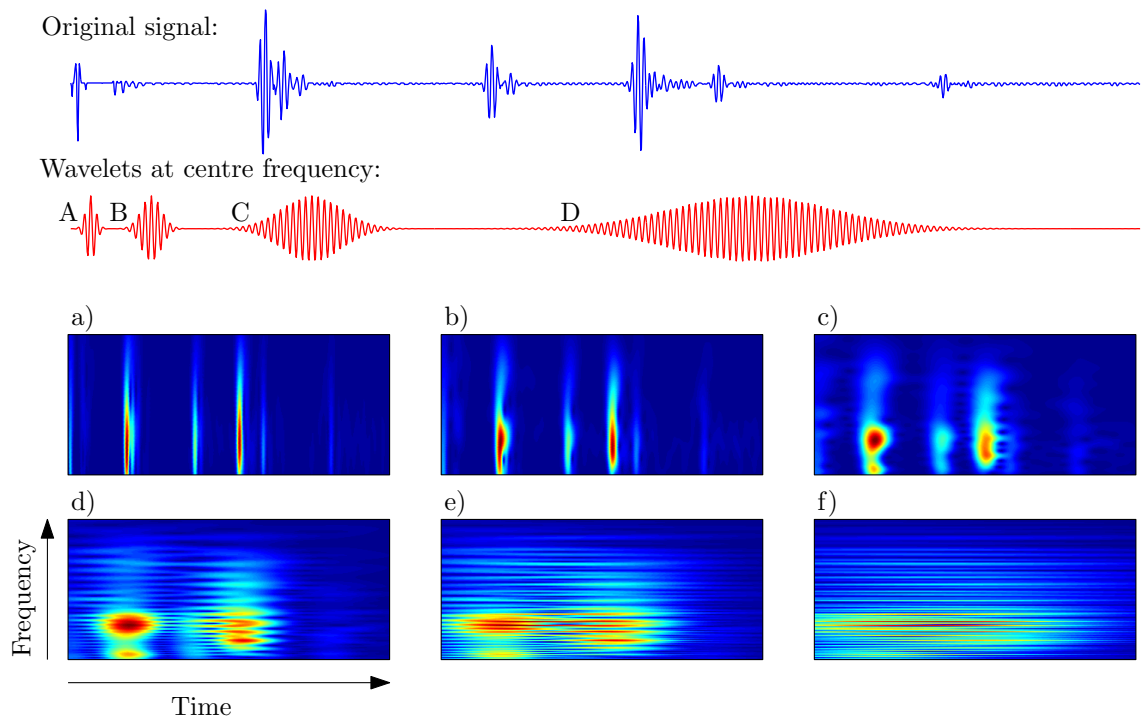


Figure 3.5: *Continuous Wavelet Transform (CWT) of the plotted sample signal into different bases on the time-frequency continuum. In (a) the basis wavelets are very short and broadband, so the signal is localised well in the time domain. The basis functions are increased in length and narrow in bandwidth through (a)–(f) so that (f) is almost equivalent to the Fourier transform of the signal. The energy is most concentrated in the fewest coefficients somewhere between these extremes, at approximately (c). The plotted Morlet wavelets A–D are the wavelets used to compute (a)–(d) at the centre frequency. Note that the CWT is used rather than the DWT merely to aid visualisation (see [1]).*

predominantly depend on the frequency response of the transducers. The time component roughly depends on ρ . With single scattering we can exactly determine the bounds t_{\min} and t_{\max} within which all signal energy is contained, however with multiple scattering t_{\max} will be greater because multiply scattered waves can take longer paths through the scatterer, and hence take longer to arrive. Therefore there is potential for multiply scattered signals to contain more information than singly scattered signals.

We can therefore define an approximate box defined by f_{\min} , f_{\max} , t_{\min} , t_{\max} which contains the most of the energy. So that the lower bound on the information capacity of the signal in the wavelet domain is given by

$$I = \sum_{t=t_{\min}}^{t_{\max}} \sum_{f=f_{\min}}^{f_{\max}} \log_2 \left(1 + \frac{s(t, f)}{\varsigma(t, f)} \right).$$

3.6 Broadband array measurements and conclusions

As stated in the introduction, there is an upper bound on the information capacity of a scattering measurement of

$$I = \log_2 \left(N_v N_x^2 N_t \right) \quad (3.24)$$

This is a conservative upper bound, but formulating tighter bounds is not easy. Therefore we looked at different aspects of the broadband array measurements separately. By studying effects of spatial diversity and temporal diversity on their own we could derive tighter bounds on monochromatic array measurements, and single broadband measurements.

The circular or spherical harmonic expansion of wave fields was used to show that with low sensitivity measurements a standard sampling criterion can be derived that

depends on the frequency and the size of the scatterer. But with high sensitivity measurements, higher sampling rates than this criterion still yield additional information by allowing measurement of higher order modes. The effect of aperture was analysed in the same way by calculating the singular values of subsets of the operator that maps near-field waves to the far-field. This showed that the information capacity of limited aperture measurements of a field depends linearly on the aperture.

We also discussed the sampling of a single broadband signal in both the time and frequency domains, which are not independent. Since measurements in time and frequency must be finite, the signal can be thought about more intuitively in the wavelet domain. The information capacity bound on the signal is also lowest in this domain since the signal energy is concentrated in fewer coefficients. We can draw bounding box around the signal energy, which allows us to calculate sampling rate and duration criteria.

Broadband array measurements will have information capacity between single-frequency or single-transducer limits, and the loose upper bound of equation 3.24. It is not possible to actually find the information capacity of the signals, however we can say that compared to broadband array measurements, current methods neglect information from two sources:

1. Broadband signals. The super-resolution algorithms described so far operate only on a single frequency of the data. This ignores the huge amount of information contained in the other frequencies.
2. Backwall reflections. Normally signals are only recorded up to the first reflection from the boundaries of a component. By extending t_{\max} to include signals beyond this point we can add additional information.

Figure 3.6 illustrates the limited use of signal information with single-frequency imaging methods. Only the region occupied by the white line is used. This corresponds to only a single frequency, and only extends to the backwall reflection. Most of the energy reflected from the defect arrives later than this and is broadband. The spatial

sampling and temporal sampling is usually already sufficient, and although apertures are often less than is desired, this is mainly due to physical practicality. Super-resolution methods rely on larger numbers of independent measurements (i.e. a larger amount of information), so single-frequency methods will be more sensitive to noise. However, most ultrasonic measurements are broadband, and neighbouring frequencies should contain complementary information about the scatterer—independent measurements of the same underlying information. The complementary nature of different frequencies is also suggested by information capacity theory, where thanks to the invariance theorem, the spatial bandwidth of a microscope can be extended at the expense of the temporal bandwidth, leading to the time domain multiplexing methods [59]. By utilising the broadband information we should be able to increase the stability of super-resolution methods.

The Linear Sampling Method has recently been extended to the broadband case, which allows us to take advantage of the greater number of independent measurements, resulting in much greater image stability than single-frequency super-resolution methods while retaining their super-resolving capabilities. In the next chapter the single-frequency and broadband super-resolution methods will be described, and we perform an experiment that demonstrates the high resolution and stability of the broadband LSM in an attenuative medium. In chapter 6 we look at using the backwall reflections, which carry information about the far side of a scatterer which is positioned in front of a planar reflector.

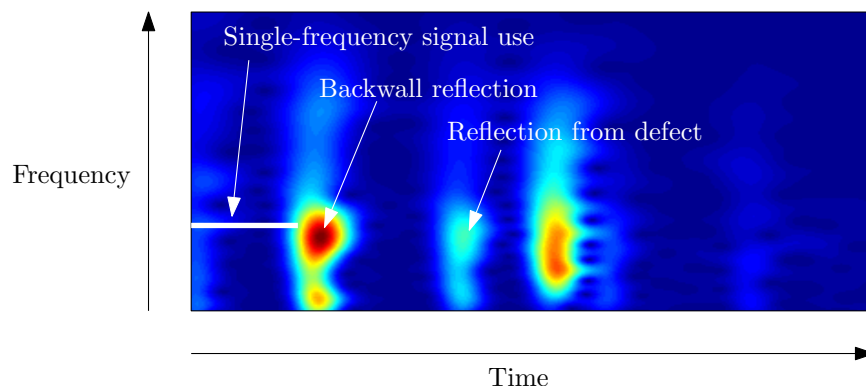


Figure 3.6: *Additional sources of data in a linear array measurement. Existing single-frequency super-resolution methods only use data from the region marked with a white line before the backwall reflection. However, useful signal energy is received after this, and also in other frequencies.*

Chapter 4

Broadband Super-Resolution

4.1 Background

Recently, new imaging algorithms have been developed that allow “super-resolution”, that is, reproducing spatial frequencies of the scatterer above $2k$ or resolving two point scatterers spaced more closely than the Rayleigh limit dictates [60]. Among these algorithms are the sampling methods, which were initially developed for monochromatic (single frequency) measurements [61, 62]. The first practical method—the Linear Sampling Method (LSM) [2, 63], was numerically and experimentally demonstrated to achieve super-resolution [20]. It is attractive because it is simple to implement, fast to compute and achieves potentially unlimited resolution. However it was not clear exactly why LSM works [64, 65]. A more advanced monochromatic sampling method called the Factorisation Method (FM) was then developed [66, 67] which was proven to produce an exact reconstruction of impenetrable scatterers with no noise. However due to their inverse nature and use of only single frequencies, both methods are highly sensitive to noise—more so than linear methods. By utilising only a single frequency, many independent measurements from the broadband signals are needlessly discarded. Broadband imaging methods allow the use of data from a range of frequencies which contain complementary information, and allow increased robustness in the presence of noise. Therefore it is desirable to use sampling methods

that combine data from different frequencies to give improved robustness while maintaining their super-resolving capabilities.

Broadband super-resolution can be achieved in an ad-hoc way by summing single frequency images, however this lacks satisfying justification and does not work well if the constituent images are poor quality due to noise. This is particularly common in highly attenuative materials where the attenuation of the signal leads to a low SNR. In attenuative materials it is also usually the case that the attenuation increases rapidly with frequency, so resolution cannot be improved by using shorter wavelengths and super-resolution would be useful. By using a time-domain variant of LSM, we can incorporate all frequencies of the signals into an integrated reconstruction [68]. The Time Domain Linear Sampling Method (TDLSM) does not improve resolution, but it does improve robustness to noise. The mathematical framework behind the sampling methods is extremely involved, therefore we shall provide a simple intuitive explanation of the monochromatic sampling methods and refer the reader to the literature for details. The monochromatic sampling methods LSM and FM, and TDLSM are described first, together with some improvements we have made to TDLSM. Then we give an experimental demonstration of the improved stability TDLSM provides in an attenuative medium. This is the first experimental use of a broadband sampling method. This work has been submitted for publication to IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control [69].

4.2 Monochromatic sampling methods

4.2.1 The Linear Sampling Method

In recent years sampling methods have been created that offer super-resolution imaging of a scatterer (i.e. images with spatial frequency content above $2k$) [70]. These work by computing a spatially varying function based on the far-field operator, F , that tends to infinity outside the scatterer, and so an image is formed by plotting its inverse which tends to zero outside the scatterer. This only allows reconstruction



Figure 4.1: *The Linear Sampling Method relies on the fact that it is possible to create field values on the boundary of a scatterer, such that a field that appears to radiate only from a single point within the scatterer. However it is not possible to do the same for a point outside the scatterer.*

of the scatterer support \mathcal{D} rather than $o(\mathbf{x})$, but in many cases this is sufficient, and for impenetrable scatterers it is the only relevant quantity. For LSM,

$$\lim_{\mathbf{x} \rightarrow \overline{\mathcal{D}}} \mathcal{I}_{\text{LSM}}(\mathbf{z}) = 0. \quad (4.1)$$

where $\mathcal{I}_{\text{LSM}}(\mathbf{z})$ is the image value at the point \mathbf{z} . The linear sampling method was the first non-iterative super-resolution algorithm that did not rely on prior knowledge of the scatterer's geometry, connectedness or boundary conditions [71]. The idea of LSM is to use the *range* of F to determine whether \mathbf{z} is inside or outside the scatterer. The range of F corresponds to the set of all possible scattered fields that can be received after transmitting an arbitrary incident wave for this particular scatterer. LSM produces an image from the observation that a scattered wave cannot be produced that emanates from \mathbf{z} for $\mathbf{z} \notin \mathcal{D}$. This is illustrated graphically in figure 4.1.

Recall that the far-field operator F maps any incident field to the corresponding scattered field. Its inverse, F^{-1} which is computable (assuming F is non-singular), and allows us to calculate the incident field required to produce any scattered field, and whether or not such a field exists. The key insight in LSM is that scattered waves cannot originate from outside the scatterer, and therefore the following equation will have no solution for $u_{\infty}^i(\hat{\mathbf{x}})$.

$$Fu_\infty^i(\hat{\mathbf{x}}) = G_\infty(\mathbf{z}, \hat{\mathbf{x}}) \quad \mathbf{z} \notin \mathcal{D} \quad (4.2)$$

In other words, $u_\infty^i(\hat{\mathbf{x}})$ is the incident field that results in a scattered field equal to that of a point source at \mathbf{z} . It can be shown mathematically that no such incident field exists when \mathbf{z} is outside \mathcal{D} , and therefore that scattered field is outside the range of F . If we attempt to solve

$$u_\infty^i(\hat{\mathbf{x}}) = F^{-1}G_\infty(\mathbf{z}, \hat{\mathbf{x}}) \quad \mathbf{z} \notin \mathcal{D} \quad (4.3)$$

we will find that $|u_\infty^i(\hat{\mathbf{x}})|$ becomes infinitely large. By plotting $\mathcal{I}(\mathbf{z}) = 1/|u_\infty^i(\hat{\mathbf{x}})|^2$ an image of \mathcal{D} is formed. Note that although it is always true that there is no solution to equation 4.2 for $\mathbf{z} \notin \mathcal{D}$, it is not guaranteed that there *is* a solution for $\mathbf{z} \in \mathcal{D}$. However the method does work in practice.

Rather than explicitly calculating the the inverse of F we can instead use the SVD of F to find its range. Recall that the SVD of F is

$$F = \sum_{n=1}^{\infty} \sigma_n |\xi_n\rangle \langle \psi_n| \quad (4.4)$$

$$F|\psi_n\rangle = \sigma_n |\xi_n\rangle \quad (4.5)$$

$$F^*|\xi_n\rangle = \sigma_n |\psi_n\rangle. \quad (4.6)$$

or we can also use the matrix notation

$$F = U\Lambda V^*$$

$$FV = \Lambda U$$

$$F^*U = \Lambda V$$

where $|\xi_n\rangle$ and $|\psi_n\rangle$ are the columns of U and V respectively and Λ is a diagonal matrix with elements σ_n . Picard's theorem says that $|G_\infty(\mathbf{z})\rangle$ is in the range of F iff [72]

$$\left[|u_\infty^i|^2 = \sum_{n=1}^{\infty} \frac{|\langle \xi_n | G_\infty(\mathbf{z}) \rangle|^2}{\sigma_n^2} \right] < \infty \quad (4.7)$$

therefore the LSM image is given by

$$\mathcal{I}_{\text{LSM}}(\mathbf{z}) = \left(\sum_{n=1}^{\infty} \frac{|\langle \xi_n | G_\infty(\mathbf{z}) \rangle|^2}{\sigma_n^2} \right)^{-1} \quad (4.8)$$

An example of the LSM image of a scatterer is shown in figure 4.2b. An intuitive interpretation of LSM is that we are plotting the magnitude of the largest scattered wave we can produce from \mathbf{z} given a normalised input wave. If there is no scatterer at \mathbf{z} then we cannot produce a large reflection from that point. In a sense, LSM is the same as BF where the optimal focal laws in transmission are calculated automatically and can take advantage of multiple scattering. LSM works, but often produces images with hugely varying values within \mathcal{D} , and significant artefacts. Various methods can be used to improve the result, such as using Tikhonov regularisation (essentially a small constant is added to each singular value), but it can be seen from figure 4.2c that FM gives much better results with no regularisation. This method is explained next.

4.2.2 Factorisation of the Far-field Operator

An improvement to LSM was made by Kirsch [66] which is better justified and produces higher quality images in practice. It is based on the factorisation of F given previously

$$F = HSH^*. \quad (4.9)$$




Figure 4.2: A demonstration of the frequency-domain sampling methods. *a)* The far-field operator F is measured for a kite-shaped sound-soft scatterer (i.e. a hole). The axes are in units of λ (or equivalently $\lambda = 1$). *b)* The resulting LSM image, clipped at an ad-hoc threshold to make the scatterer visible. *c)* The FM image. No clipping was applied in this case. F was calculated using a circular array of 256 transducers with a radius of 10, using the Boundary Element Method (BEM). The formula for the kite-shaped scatterer is given in [2].

Recall that H is the operator that maps

$$Hu^s(\mathbf{x}) = u_\infty^s(\hat{\mathbf{r}}) \quad \mathbf{x} \in \partial\mathcal{D}. \quad (4.10)$$

It follows that if

$$Hu(\mathbf{r}) = G_\infty(\mathbf{z}, \hat{\mathbf{x}}) \quad \mathbf{r} \in \partial\mathcal{D}, \mathbf{z} \notin \mathcal{D} \quad (4.11)$$

has no solution, then neither does

$$Fu_\infty^i(\hat{\mathbf{x}}) = G_\infty(\mathbf{z}, \hat{\mathbf{x}}) \quad \mathbf{z} \notin \mathcal{D} \quad (4.12)$$

and an image can be formed by interrogating the range of H instead of F . It was shown in [66] that the ranges of $(F^*F)^{\frac{1}{4}}$ and H coincide, and therefore we can calculate the range of H using $(F^*F)^{\frac{1}{4}}$ which is directly computable. Alternatively, with the singular value decomposition $F = U\Lambda V^*$, we have

$$(F^*F)^{\frac{1}{4}} = (V\Lambda^*U^*U\Lambda V^*)^{\frac{1}{4}} \quad (4.13)$$

$$= (V\Lambda^2V^*)^{\frac{1}{4}} \quad (4.14)$$

$$= V\Lambda^{\frac{1}{2}}V^*. \quad (4.15)$$

Therefore if the singular system for $(F^*F)^{\frac{1}{4}}$ is $\sigma'_n, |\xi'_n\rangle, |\psi'_n\rangle$ we have

$$\sigma'_n = \sqrt{\sigma_n} \quad (4.16)$$

$$|\xi'_n\rangle = |\bar{\psi}_n\rangle \quad (4.17)$$

$$|\psi'_n\rangle = |\bar{\psi}_n\rangle \quad (4.18)$$

Note that if F is symmetric, we have

$$F = F^T = \bar{F}^* \quad (4.19)$$

$$= V^*\Lambda U \quad (4.20)$$

therefore an equivalent singular system of $(F^*F)^{\frac{1}{4}}$ for $F = F^T$ is

$$\sigma'_n = \sqrt{\sigma_n} \quad (4.21)$$

$$|\xi'_n\rangle = |\xi_n\rangle \quad (4.22)$$

$$|\psi'_n\rangle = |\xi_n\rangle \quad (4.23)$$

Using the former of these transformations, we find that \mathbf{z} is in \mathcal{D} if $|G_{\mathbf{z}}\rangle$ is in the range of H which is true iff

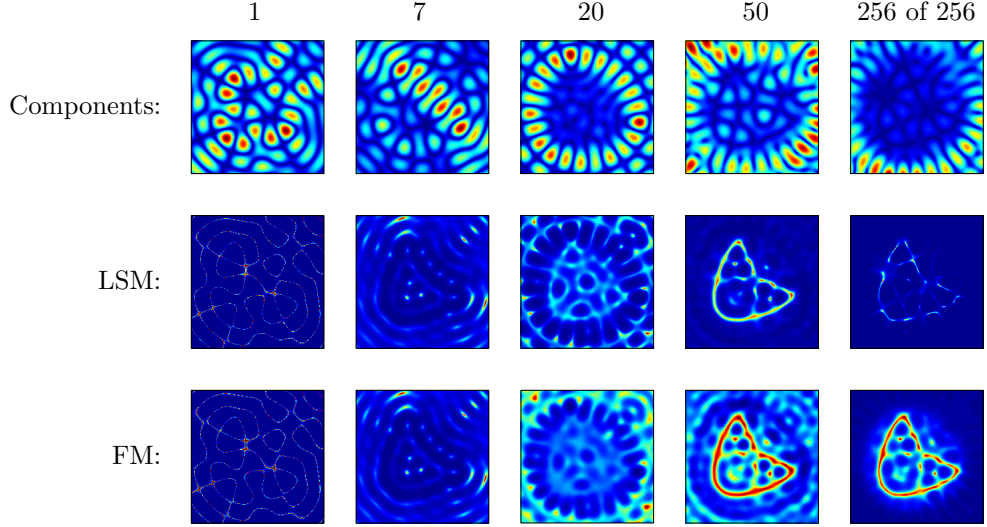


Figure 4.3: The values $|\langle \xi_n | G_{\mathbf{z}} \rangle|^2$, for the first, last, and some intermediate n 's, and the cumulative sum for LSM and FM.

$$\sum_{n=1}^{\infty} \frac{|\langle \xi'_n | G_{\infty}(\mathbf{z}) \rangle|^2}{\sigma_n'^2} < \infty \quad (4.24)$$

which is equivalent to

$$\sum_{n=1}^{\infty} \frac{|\langle \bar{\psi}_n | G_{\infty}(\mathbf{z}) \rangle|^2}{\sigma_n} < \infty \quad (4.25)$$

and for symmetric F ,

$$\sum_{n=1}^{\infty} \frac{|\langle \xi_n | G_{\infty}(\mathbf{z}) \rangle|^2}{\sigma_n} < \infty. \quad (4.26)$$

Therefore the FM image is given, for symmetric F (which is expected due to reciprocity) by

$$\mathcal{I}_{\text{FM}}(\mathbf{z}) = \left(\sum_{n=1}^{\infty} \frac{|\langle \xi_n | G_{\infty}(\mathbf{z}) \rangle|^2}{\sigma_n} \right)^{-1} \quad (4.27)$$

Unlike LSM, this function is known to be finite on the boundary $\partial\mathcal{D}$ and zero outside it. A numerical example of FM imaging is shown in figure 4.2c.

Contrasting LSM and FM with the BF image which is given by

$$\mathcal{I}_{\text{BF}}(\mathbf{z}) = \sum_{n=1}^{\infty} \sigma_n \langle G_{\infty}(\mathbf{z}) | \xi_n \rangle \langle \psi_n | G_{\infty}(\mathbf{z}) \rangle \quad (4.28)$$

we can see that the difference between BF, LSM and FM is primarily how σ_n affects the weighting of each of the terms involving $\langle G_{\infty}(\mathbf{z}) | \xi_n \rangle$ in the sum. The information about the shape of the scatterer is contained in these terms, and with BF the terms associated with the highest σ_n contribute most to the image. On the other hand, with LSM and FM the terms associated with the *lowest* σ_n contribute most to the image. The $|\langle \xi_n | G_{\mathbf{z}} \rangle|^2$ terms, and cumulative value of the LSM and FM image sums are plotted for the kite scatterer in figure 4.3.

The difference in term weighting affects how noise degrades the image. Gaussian noise will be equally distributed across all singular vectors, but it is amplified in the LSM and FM case for the terms with small σ_n . One way to try to reduce the effects is by using regularisation to reduce the impact of the small singular values. Tikhonov regularisation changes the sum to [66]

$$\mathcal{I}_{\text{FM}}(\mathbf{z}) = \left(\sum_{n=1}^{\infty} \frac{|\langle \xi_n | G_{\infty}(\mathbf{z}) \rangle|^2}{\sigma_n + \alpha \sigma_n^{-1}} \right)^{-1}$$

where α is a small regularisation parameter. However, we have not found this to be very effective. A more powerful way to improve robustness is to combine the data from multiple frequencies.

4.3 Broadband Linear Sampling Method

A significant problem with LSM and FM when applied to time-domain signals, such as those recorded by an ultrasound array, is that they operate upon only a single frequency. This means that we discard all information contained in other frequencies. One solution is to sum the images generated at each frequency (or wavenumber)

$$\mathcal{I}_{\text{LSM}}(\mathbf{z}) = \left(\sum_{k=-\infty}^{\infty} \sum_{n=1}^{\infty} \frac{|\langle \xi_{n,k} | G_{\infty}(\mathbf{z}, k) \rangle|^2}{\sigma_{n,k}^2} \right)^{-1}. \quad (4.29)$$

where now we have a singular system $\sigma_{n,k}$, $|\xi_{n,k}\rangle$, $|\psi_{n,k}\rangle$ and a different $G_{\infty}(\mathbf{z})$ for each k . The interpretation of this function is as follows. Suppose we try to solve the time-domain equivalent of equation 4.2

$$F_t u_{\infty}^i(\hat{\mathbf{r}}, t) = G_{\infty}(\mathbf{z}, \hat{\mathbf{x}}, t) \quad (4.30)$$

where F_t is the time-domain far-field operator, and $G_{\infty}(\mathbf{z}, \hat{\mathbf{x}}, t)$ is the time-domain far-field Green's function (i.e. the field resulting from an impulse at \mathbf{z}). The sum in equation 4.29 is then the magnitude of $u_{\infty}^i(\hat{\mathbf{r}}, t)$ in equation 4.30, calculated in the frequency domain. By Parseval's theorem we obtain

$$|u_{\infty}(\hat{\mathbf{r}}, t)|^2 = \sum_{k=-\infty}^{\infty} \sum_{n=1}^{\infty} \frac{|\langle \xi_{n,k} | G_{\infty}(\mathbf{z}, k) \rangle|^2}{\sigma_{n,k}^2} \quad (4.31)$$

Therefore the (inverse) quantity we are plotting is the energy of the time-domain input required to produce a delta-function point source output at \mathbf{z} . However, there is one practical detail that means this is not quite correct: the finite Fourier transform assumes the signal is periodic. Since all our measurements are necessarily finite, equation 4.29 allows *non-causal* solutions.

Suppose we are trying to solve the following convolution system for $u(t)$

$$u(t) * v(t) = s(t)$$

where $v(t)$ is the impulse response and $s(t)$ is the desired output. An example of this is shown in figure 4.4. $v(t)$ and $s(t)$ are plotted as figure 4.4a and b respectively. It can be seen that there should not be a solution in this case, since the wave packet in $v(t)$ occurs after the wave in $s(t)$, so even if $u(t)$ contains a transmission at $t = 0$,




Figure 4.4: *Non-causality of Fourier solutions to deconvolution. See the text for a description.*

the response will be too late to produce $s(t)$. However, if we solve for $u(t)$ using Fourier deconvolution, which assumes that all signals are periodic, then we obtain the solution shown in figure 4.4c, given by

$$u_{\mathcal{F}}(t) = \mathcal{F}^{-1} \left[\frac{\mathcal{F}[s(t)]}{\mathcal{F}[v(t)]} \right].$$

This is clearly not the correct solution; if we calculate the time-domain convolution $u_{\mathcal{F}}(t) * v(t)$ we obtain the incorrect solution shown in figure 4.4d (it should be equal to figure 4.4b). Therefore in order to maintain causality we must perform the deconvolution in the time domain. This example also demonstrates that we need to delay the target signal $G_{\infty}(\mathbf{z}, \hat{\mathbf{x}}, t)$ in equation 4.30 so that it “arrives” after the reflections in $v_{\infty}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t)$.

In order to allow only causal solutions we can instead formulate the time-domain far-field operator F_t directly in the time domain. F_t is a convolution operator, which is given in the continuous case by

$$u_{\infty}^s(\hat{\mathbf{x}}, t) = F_t u_{\infty}^i(\hat{\mathbf{x}}, t) = \int_S \int_{-\infty}^{\infty} v_{\infty}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t - \tau) u_{\infty}^i(\hat{\mathbf{r}}, t) d\tau d\hat{\mathbf{r}} \quad (4.32)$$

where $v_{\infty}(\hat{\mathbf{r}}, \hat{\mathbf{x}}, t)$ is the impulse response when transmitting from the direction $\hat{\mathbf{r}}$ to

$\hat{\mathbf{x}}$. Due to the form of the 2D time-domain Green's function this is not well defined in 2D. Therefore we shall now assume that instead of transmitting and receiving waves from infinity, we have a discrete set of transducers at \mathbf{x}_i recording a time-limited discrete set of samples, as in the previous chapter. Therefore in the discrete case $F_{t,d}$ is

$$u^s(\mathbf{x}_i, t) = F_{t,d} u^i(\mathbf{x}_i, t) = \sum_{j=1}^{N_x} \sum_{\tau=-\infty}^{\infty} v(\mathbf{x}_j, \mathbf{x}_i, \tau) u^i(\mathbf{x}_j, t - \tau) \quad (4.33)$$

or alternatively,

$$\begin{bmatrix} v(\mathbf{x}_1, \mathbf{x}_1, t) & v(\mathbf{x}_2, \mathbf{x}_1, t) & \cdots & v(\mathbf{x}_{N_x}, \mathbf{x}_1, t) \\ v(\mathbf{x}_1, \mathbf{x}_2, t) & & & \\ \vdots & & \ddots & \vdots \\ v(\mathbf{x}_1, \mathbf{x}_{N_x}, t) & \cdots & v(\mathbf{x}_{N_x}, \mathbf{x}_{N_x}, t) \end{bmatrix} * \begin{bmatrix} u^i(\mathbf{x}_1, t) \\ u^i(\mathbf{x}_2, t) \\ \vdots \\ u^i(\mathbf{x}_{N_x}, t) \end{bmatrix} = \begin{bmatrix} u^s(\mathbf{x}_1, t) \\ u^s(\mathbf{x}_2, t) \\ \vdots \\ u^s(\mathbf{x}_{N_x}, t) \end{bmatrix}.$$

$F_{t,d}$ can be written in explicit matrix form as a block Toeplitz matrix with the structure shown in figure 4.5. Since this operator is now a simple matrix it can be inverted or SVD'd using standard methods. However in practical scenarios the dimension of the matrix is huge (on the order of $N_x N_t$ rows and columns) and using standard SVD methods is impossible. Instead, as in [68] we use the Arnoldi iteration method to calculate the largest few thousand singular values and vectors, in order to evaluate the range of $F_{t,d}$.

There are two practical issues that must be considered in a real application. The first is that in general we cannot measure the true impulse response of the scatterer, but a band-limited impulse response that depends on the frequency response of the transducers, the scatterer and the background material. Equation 4.30 will have no solution if $F_{t,d}$ does not contain all frequencies. Therefore instead of attempting to solve for an impulse output, we try to solve for a band-limited output

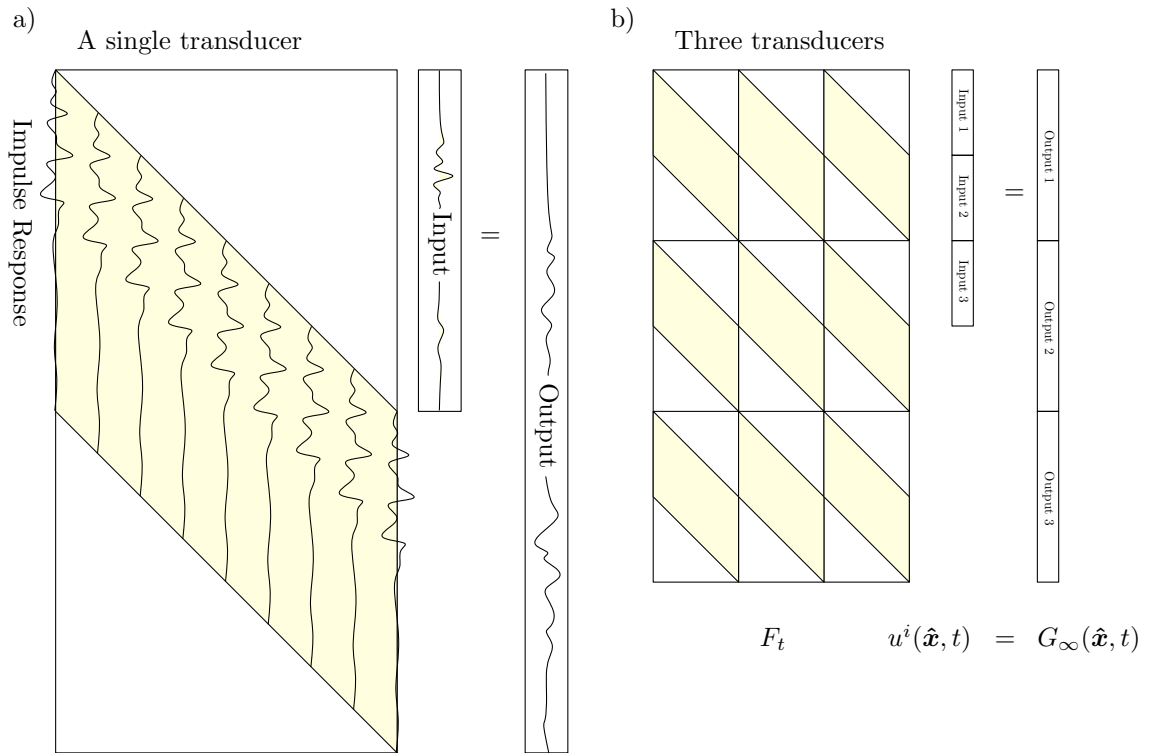


Figure 4.5: Structure of the discrete time-domain far-field operator $F_{t,d}$ in block Toeplitz matrix form for a) a single transducer and b) three transducers. The exact clipping of the time signals does not necessarily have to allow for the complete convolution of the input and output, nor do the input and output have to be the lengths given here. For example we could use only the top half of each block if desired. In all the TDLSM images in this thesis, the structure was as shown here, with each block having a dimension of $N_t \times (2N_t - 1)$.

$$F_{t,d}u^i(\mathbf{x}, t) = G(\mathbf{z}, \mathbf{x}, t) * s(t) \quad (4.34)$$

where $s(t)$ contains approximately the same frequencies as $F_{t,d}$. In [68], an apparently arbitrary exponentially windowed function was chosen. This, and several other waveforms were tested for $s(t)$, and it was found that setting $s(t)$ to equal the transmitted wave-packet used to measure $F_{t,d}$ (and $v(\mathbf{r}, \mathbf{x}, t)$) gave the best results empirically.

The second practical consideration is the length of time to consider for $u^i(\hat{\mathbf{r}}, t)$ and $G(\mathbf{z}, \mathbf{x}, t) * s(t)$. We cannot actually compute the convolution over all time, so arbitrary limits must be placed on the convolution. We must allow enough time from the input for the waves to propagate to the scatterer and back to the transducers. A new optimisation we have made is to pre-beamform the time-signals relative to the centre of the scatterer (or the image) so that a shorter time for $G_\infty(\mathbf{z}, \hat{\mathbf{x}}, t) * s(t)$ can be considered. This can significantly speed up the computation—especially for large linear apertures—and is now described in detail.

4.3.1 Pre-beamforming

If we assume the scatterer is mainly within the image region around the point \mathbf{b} , we can assume that the impulse responses will all be zero up to a certain point, which depends on the distances of each transducer from \mathbf{b} (in practice the transducers are not an infinite distance away). Denote the time taken for a wave to travel from transducer i to \mathbf{b} (or vice versa) as T_i . Then usually $v(\mathbf{x}_i, \mathbf{x}_j, t) = 0 \quad \forall t < T_i + T_j$. Since $u^i(\mathbf{x}, t)$ is zero before $t = 0$, we require $u^s(\mathbf{x}, t) = 0 \quad \forall t < \max(2T_i)$, where $u^s(\mathbf{x}, t) = G(\mathbf{z}, \mathbf{x}, t) * s(t)$, otherwise a correct solution is impossible. In order to make this true we add the τ constant to delay y . However the problem is if T_i varies a lot, then all the signals need to be very long, which makes the SVD calculation impossibly slow.

The solution is to delay the transmissions and receptions by p_i and q_i respectively,

so that the important portions of $v(\mathbf{x}_i, \mathbf{x}_j, t)$ around $t = T_i + T_j$ are aligned, then we can consider a shorter period for all signals. In other words, we calculate

$$\begin{bmatrix} v(\mathbf{x}_1, \mathbf{x}_1, t + p_1 + q_1) & v(\mathbf{x}_2, \mathbf{x}_1, t + p_1 + q_2) & \cdots & v(\mathbf{x}_{N_x}, \mathbf{x}_1, t + p_1 + q_{N_x}) \\ v(\mathbf{x}_1, \mathbf{x}_2, t + p_2 + q_1) & & & \\ \vdots & & \ddots & \vdots \\ v(\mathbf{x}_1, \mathbf{x}_{N_x}, t + p_{N_x} + q_1) & \cdots & v(\mathbf{x}_{N_x}, \mathbf{x}_{N_x}, t + p_{N_x} + q_{N_x}) \end{bmatrix} * \begin{bmatrix} u^i(\mathbf{x}_1, t - q_1) \\ u^i(\mathbf{x}_2, t - q_2) \\ \vdots \\ u^i(\mathbf{x}_{N_x}, t - q_{N_x}) \end{bmatrix} = \begin{bmatrix} u^s(\mathbf{x}_1, t + p_1) \\ u^s(\mathbf{x}_2, t + p_2) \\ \vdots \\ u^s(\mathbf{x}_{N_x}, t + p_{N_x}) \end{bmatrix}.$$

with $p_i = q_i = T_i$. The norm of $u^i(\mathbf{x}, t - q_i)$ is the same as $u^i(\mathbf{x}, t)$ so this will give the substantially the same image but in a much shorter time.

Therefore the TDLSM image is then formed by considering a grid of points in space and calculating the same functional as for LSM, but using the time-domain operator F_t instead of F , with time-shifted signals as described above.

$$\mathcal{I}_{\text{TDLSM}}(\mathbf{z}) = \left(\sum_{n=1}^{\infty} \frac{|\langle \xi_n | G(\mathbf{z}) * s(t) \rangle|^2}{\sigma_n^2} \right)^{-1} \quad (4.35)$$

Where $|\xi_n\rangle$ and σ_n are the singular vectors and values of F_t as calculated by Arnoldi iteration. Due to computational limitations, only the largest N_σ terms of the sum are calculated. We used $N_\sigma = 5\,000$ or $N_\sigma = 10\,000$.

Arnoldi iteration does not require an explicit representation of $F_{t,d}$ as a matrix; it only requires us to calculate the value $F_{t,d}u$ or $F_{t,d}^*u$ when presented with a test vector u , thus we can implement the operator however we wish and using no more




Figure 4.6: a) *Time-domain beamforming*, b) *FM*, and c) *TDLSM* imaging of two squares beneath 8 transducers. The wavelength is $\frac{\pi}{2}$. It can be seen that both sampling methods resolve the scatterers, whereas BF does not.

memory than the source data occupies. A fast implementation of $F_{t,d}$ was written in C++ and parallelised using MPI [73]. Arnoldi iteration was performed using the P_ARPACK Fortran library, which is a version of the ARPACK Arnoldi iteration library¹ also parallelised using MPI. In the experiments in this chapter, signals consisting of 700 samples were recorded with 32 transducers giving an operator dimension of 44799×22400 . We calculated the 5000 largest singular values and their associated singular vectors, which was sufficient for convergence of equation 4.35. The value of $I(\mathbf{z})$ was calculated on a 300×300 pixel grid.

An examples of TDLSM compared with FM and time-domain BF is shown in figure 4.6 for two squares that are closer than the Rayleigh limit of linear imaging. The centre wavelength is $\frac{\pi}{2}$, so the squares are 1, or 0.64λ apart. The array aperture taken at the point between the squares is 84° therefore the Rayleigh limit is $\frac{0.61\lambda}{\sin(42^\circ)} = 1.43 = 0.91\lambda$. The first 1000 singular values were used to calculate the TDLSM image, and the centre wavelength is $\frac{\pi}{2}$. It can be seen that TDLSM retains the super-resolution capabilities of LSM. In the next section we demonstrate experimentally that it is much more robust to noise than LSM.

¹<http://www.caam.rice.edu/software/ARPACK/>

4.4 Experimental validation of Time-Domain LSM

We measured the horizontal resolution capabilities of BF, FM and TDLSM using two human hairs suspended in glycerol, both materials that are readily available. Glycerol was used because it is highly attenuative and liquid, allowing the scatterers to be easily positioned. The speed of sound in glycerol is 1900 m/s. The frequency-dependent attenuation of glycerol was calculated by measuring the spectrum of the reflection of a broadband pulse from a planar reflector at various distances. The pulses for different reflector distances are shown in figure 4.7a. At each frequency the attenuation was found by calculating an exponential fit to the magnitude-vs-distance curve for that frequency. This allowed calculation of the attenuation up to 6 MHz reliably. Above this frequency the attenuation was too great for the signals to be measured beyond the dead-zone of the transducer. A best-fit curve was then calculated for the attenuation, using a quadratic function as shown in figure 4.7b. This gave an attenuation of 0.27 dB/cm/MHz² (at 26.4°C) which agrees with the data in the literature: 0.22 dB/cm/MHz² in [74] and 0.19 dB/cm/MHz² in [75].

Data was captured at 80 positions along the hair (separation increments of 0.02 mm) using an Ultrasonix RP medical scanner². 32 elements of the L14-5/38 array were used with an element size of 0.3 mm and a pitch of 0.6 mm, placed a distance of approximately 10 mm from the hairs. In order to finely control the separation of the hairs, we chose to fix them in position in a cross shape as shown in figure 4.8. The hairs were almost—but not quite parallel—so that scanning the array in the x direction gradually changes the measured hair separation. Because the angle between the hairs was so small their separation could be changed in very small steps. The true separation of the hairs was estimated by fitting lines to their 3D reconstruction. The scattering data was measured by transmitting a pulse from each element and recording the reflections on all other elements. A single negative pulse was transmitted which was filtered by the frequency response of the array and lead to a centre frequency of 5.5 MHz with a bandwidth of 5.5 MHz. 1500 samples were recorded at a rate of 20 MHz. This gave 32×32 signals, which are the discretisation

²<http://www.ultrasonix.com/>

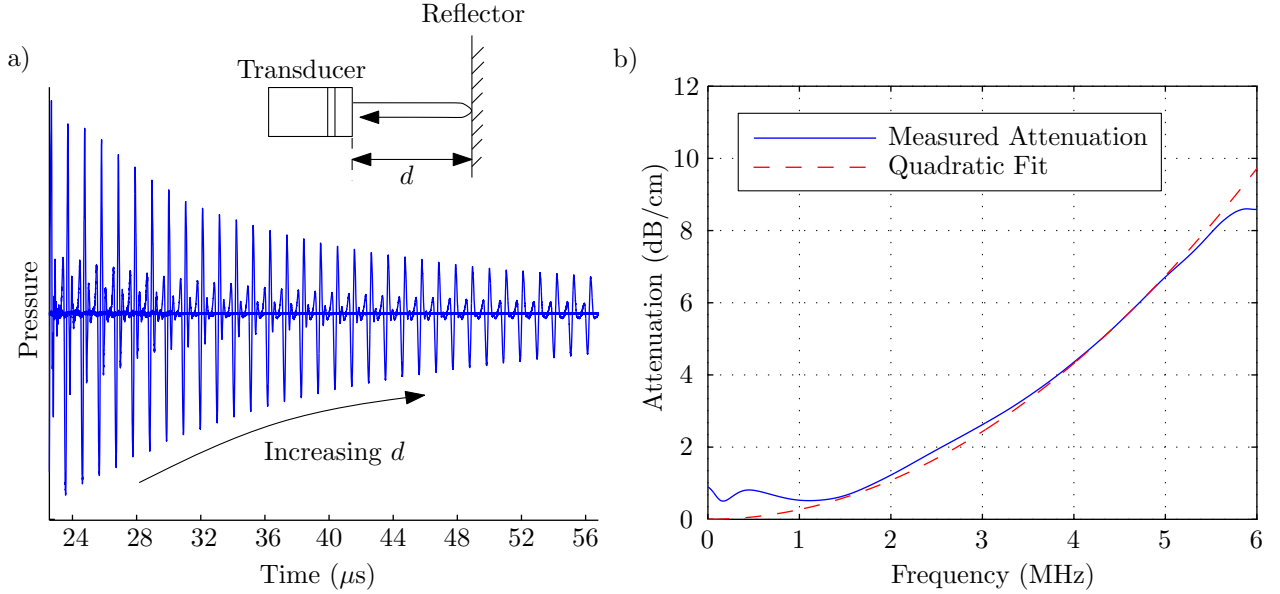


Figure 4.7: Measurement of the attenuation of ultrasound in glycerol was achieved by transmitting a broadband pulse over various distances. A sub-set of the reflections are shown overlaid in (a). As the distance d increases the pulses are delayed and attenuated. The rate of decay of each frequency component of the pulse was calculated, and used to find the frequency-dependent attenuation, this is shown with a quadratic fit in (b). Above 6 MHz the signals decay to quickly too measure accurately.

of $u^\infty(\boldsymbol{\theta}, \boldsymbol{\phi}, t)$ for a limited range of $\boldsymbol{\theta}$, $\boldsymbol{\phi}$ and t .

The resulting data was then processed using time-domain beamforming and the time domain Linear Sampling Method. Beamforming was performed using the standard ellipses method. The details of the time-domain LSM calculation are given in the next sub-section.

4.5 Results and Discussion

All image slices were processed with time-domain BF, the time-domain LSM, and single-frequency FM. Several slices around the critical separation at which the hairs become unresolved are shown in figure 4.9. The time-domain LSM algorithm is able to resolve the hairs at closer separation than with BF. It is also free from ringing artefacts in the BF images, which are caused by the band-limited nature of the signals,




Figure 4.8: *Experimental set-up. The array is positioned above the crossed hairs, such that they appear as two dots in the image plane. The hairs are at a distance of approximately 10 mm from the array, giving an angular aperture of 46° . The array is scanned along the hairs, which vary in separation from 0.8 mm to -0.8 mm. a) Perspective schematic, b) Cross section through image plane, c) Bird's-eye view.*

however it does suffer from a lower vertical resolution. This is caused by the relaxed assumptions made by the algorithm, combined with the limited aperture. Intuitively BF works by using the time-of-flight of reflections to calculate their distance, whereas LSM uses a process similar to triangulation. This means that a small aperture can lead to large uncertainties in the range direction. In the horizontal direction however, it leads to the increased resolution shown here.

Images for the closest separation at which the hairs are still *just* resolved using LSM are shown in figure 4.10. The first seven images show the single-frequency FM reconstructions using frequencies from 3 to 9 MHz with the attenuation and wavelength in each case. At 3 MHz, the frequency is too low to resolve the hairs, and a single merged object can be seen. The position and shape of the scatterer image is not stable, and varies wildly between the images. The background of these images also has significant noise. These are both a result of the limited amount of information contained in the single-frequency measurement. At 4, 5 and 6 MHz the hairs are resolved, however the images remain unstable, and by 8 MHz the hair is indistinguishable from the background artefacts. Super-resolution is achieved with the single-frequency method, but the reconstructions are very sensitive to noise.

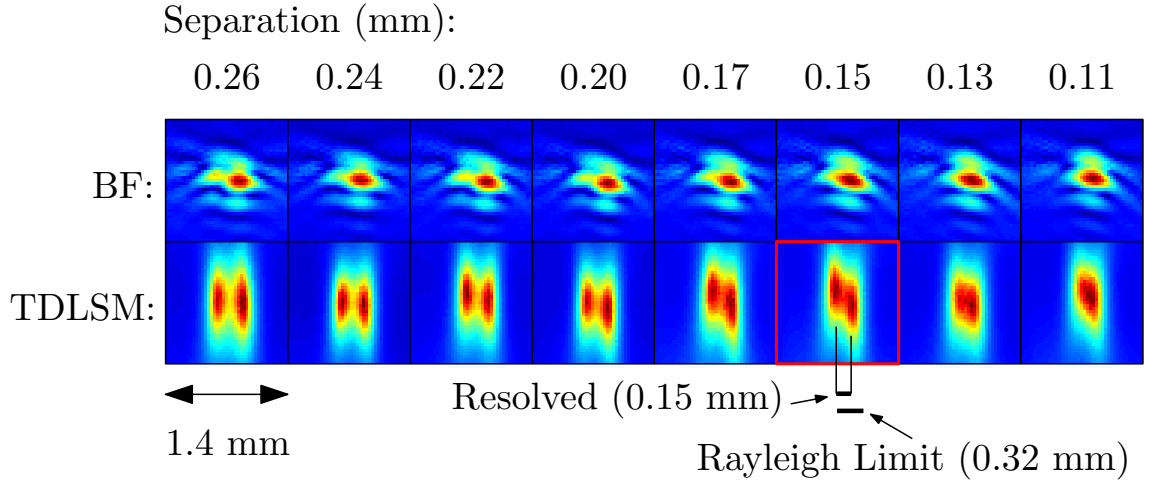


Figure 4.9: *The critical images where resolution of the hairs is lost. The top row of images are produced with traditional beamforming, the bottom with time-domain LSM. As separation of the hairs decreases from left to right the hairs remain resolved longer with TDLSM than BF. The TDLSM images are also freer of artefacts.*

The eighth image shows the time-domain LSM reconstruction. Although the resolution is slightly reduced compared to the resolved FM images, the stability is much better, and the images is completely free of artefacts. Note that this image is not merely the sum of single-frequency images, but is formed by combining the information from all frequencies during the imaging process (rather than in a post-processing step). This confirms our initial hypothesis that the greater information content in the broadband images can improve the effective SNR of the image.

The high attenuation of the glycerol prevents the propagation of broadband high frequency waves, which are necessary for high resolution beamforming. The separation of the hairs in figure 4.10 is 0.15 mm. The aperture angle is 46° which leads to a Rayleigh resolution limit of 0.32 mm which is approximately double the resolved distance with LSM.

4.6 Conclusion

In this chapter, the main sampling methods were described, and the possibility of super-resolution imaging in attenuative media using a broadband sampling method

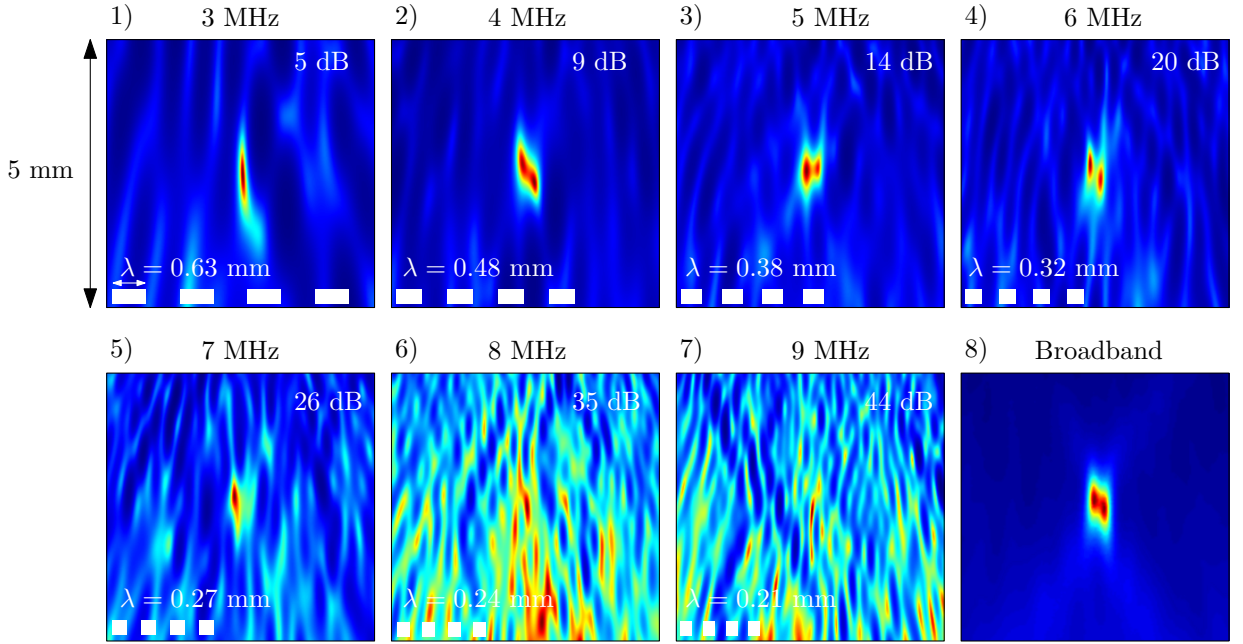


Figure 4.10: *The single-frequency and broadband sampling images for one particular hair position (highlighted in red in figure 4.9). The broadband image is much more stable than the individual frequency images but retains almost the same resolution. See the text for a complete description.*

was demonstrated. This allows small details in the image to be seen, beyond the Rayleigh criterion which limits the resolution of images produced using classical imaging algorithms. Super-resolution is especially important in attenuative media where high frequency waves do not propagate very far, precluding the possibility of increasing the resolution by using shorter wavelengths. High resolution with low frequencies has additional technological advantages, such as allowing the use of fewer elements.

Super-resolution was demonstrated by attempting to resolve two hairs at a number of different separations. These appear as two blobs in the image, and below a certain separation they merge into one. It was found that the time domain Linear Sampling Method was able to resolve the hairs at closer separations than BF (approximately half the Rayleigh limit).

The recently developed time domain LSM algorithm was used for the first time experimentally. The use of multiple frequencies provides a greater number of independent

measurements of the scatterer, which greatly increases the stability of the image, with no loss of resolution, as can happen with BF methods.

This opens the door to resolution increases in sub-surface imaging in attenuative media, such as in medical and seismic imaging. In the next chapter we look at the second source of improved sub-surface imaging by exploiting multiple reflections involving smooth boundaries of a component.

Chapter 5

Broadband super resolution in elastic solids

Although scalar scattering captures most of the physics involved in inverse scattering, several important imaging applications occur in multi-mode media. In particular, Ultrasonic Testing (UT) and geophysical imaging are performed in solid media. In solid materials both shear (S) and longitudinal (P) wave modes are generated, and conversion between these modes at inhomogeneities is generally unavoidable. The additional wave modes contain extra information about the scatterers they illuminate, so it would be beneficial to use both modes for imaging.

However, due to the complexity of elastic scattering a common approach to the problem of imaging with elastic waves is to assume scattering is actually acoustic, with acoustic waves travelling at the speed of P- or S-waves, and to ignore the other mode(s). Since beamforming is linear, acoustic imaging in an elastic medium should result in an image of the scatterer plus some additional artefacts caused by the additional modes, whose strength will depend on the angular diversity of the waves. Beamforming with S or P waves in elastic materials is not new, for example Long [76] uses mode conversion from the backwall of a component to generate shear waves to detect weld defects. Shear waves are also commonly used in other imaging modes, such as for Time of Flight Diffraction (ToFD) [77].

Simonetti [78] demonstrated that super-resolution could be achieved in steel, however in that case, due to the limited aperture, mode conversion was minimal and it was possible to separate the P and S waves in the time domain due to the difference in their speeds. However it is not known how the sampling methods are affected by elastic waves in general. It is the aim of this chapter to investigate this, and to show that time domain LSM can work in a practically significant experiment.

In the first section we give some background on elastic scattering, in the second section simulations are performed for cracks in an elastic medium to investigate the effect of elastic modes. Finally in the last section an experimental demonstration of using TDLSM on an aluminium sample to image small cracks is provided.

5.1 Background

Elastic waves travel in bulk material as either a longitudinal P-wave, or as one of two S-wave modes (or a single S-wave mode in 2D). Each of these wave modes behaves according to the acoustic wave equation. From chapter 2, we have

$$\frac{\partial^2 p}{\partial t^2} - c_l^2 \nabla^2 p = F'_l \quad (5.1)$$

$$\frac{\partial^2 \mathbf{s}}{\partial t^2} - c_s^2 \nabla^2 \mathbf{s} = \mathbf{F}'_s \quad (5.2)$$

$$\mathbf{d} = \nabla p + \nabla \times \mathbf{s} \quad (5.3)$$

where in 2D, \mathbf{s} is a scalar field (or a vector field with zero x and y components). All wave modes behave as acoustic waves in the bulk material where there are no constraints on the total displacement. We cannot determine the value of these modes by measuring at a single point because all modes can cause displacements in all directions, but we can distinguish the wave modes if we can measure all displacement

components on a sufficiently well sampled linear or circular aperture.

As before, we can separate the radial and circumferential components dependence of \mathbf{d} by decomposition into spherical (or circular) harmonics. This can be used to formulate an elastic far-field operator \mathbf{F} . The circular harmonic expansion of an elastic wave field in 2D is

$$p(r, \theta) = \sum_{n=-\infty}^{\infty} \left[A_n H_n^{(1)}(k_p r) + B_n H_n^{(2)}(k_p r) \right] e^{in\theta} \quad (5.4)$$

$$s(r, \theta) = \sum_{n=-\infty}^{\infty} \left[C_n H_n^{(1)}(k_s r) + D_n H_n^{(2)}(k_s r) \right] e^{in\theta} \quad (5.5)$$

$$\mathbf{d} = \nabla p + \nabla \times (s \mathbf{k}) \quad (5.6)$$

where \mathbf{k} is the out-of-plane unit vector, and k_p and k_s are the wavenumbers for P and S waves respectively. The fields corresponding to the coefficients A_0 , A_1 , C_0 and C_1 are shown in figure 5.1. These coefficients represent out-going waves, whereas B_n and D_n represent incoming waves. We can write the far-field out-going waves as before

$$p_{\infty}^s(\theta) = \lim_{r \rightarrow \infty} \frac{1}{\Gamma_{2D}(r)} \sum_{n=-\infty}^{\infty} A_n H_n^{(1)}(k_p r) e^{in\theta} \quad (5.7)$$

$$s_{\infty}^s(\theta) = \lim_{r \rightarrow \infty} \frac{1}{\Gamma_{2D}(r)} \sum_{n=-\infty}^{\infty} C_n H_n^{(1)}(k_s r) e^{in\theta} \quad (5.8)$$

and we can write the elastic far-field operator, \mathbf{F} as

$$\mathbf{F} \begin{bmatrix} p_{\infty}^i(\hat{\mathbf{x}}) \\ s_{\infty}^i(\hat{\mathbf{x}}) \end{bmatrix} = \begin{bmatrix} p_{\infty}^s(\hat{\mathbf{x}}) \\ s_{\infty}^s(\hat{\mathbf{x}}) \end{bmatrix}.$$




Figure 5.1: *The displacement fields corresponding to A_0 , A_1 , C_0 and C_1 . The central region is not shown so the smaller magnitude vectors further out can be seen.*

\mathbf{F} can be decomposed as $\mathbf{F} = \mathbf{H}\mathbf{S}\mathbf{H}^*$ as before. This can be written more explicitly as

$$\begin{bmatrix} F_{pp} & F_{sp} \\ F_{ps} & F_{ss} \end{bmatrix} = \begin{bmatrix} H_{pp} & H_{sp} \\ H_{ps} & H_{ss} \end{bmatrix} [\mathbf{S}] \begin{bmatrix} H_{pp}^* & H_{sp}^* \\ H_{ps}^* & H_{ss}^* \end{bmatrix}$$

We can conclude, firstly that all the results of chapter 3 also apply to elastic scattering, and secondly it should in principle be possible to formulate FM and LSM for elastic scattering. However in practice ultrasound arrays are incapable of measuring the displacement of the medium in all directions, which is necessary to produce and detect the P- and S-waves independently. Although arrays can be constructed that excite the medium in the vertical or horizontal directions, they cannot do both, therefore we must consider an alternative method of separating the waves.

In chapter 2 it was shown that the speeds of P- and S-waves are $c_l^2 = \frac{\lambda+2\mu}{\rho}$ and $c_s^2 = \frac{\mu}{\rho}$ respectively. Since μ must be positive, it can be seen that as long as λ is positive—which it is in non-exotic materials—then S-waves will always be slower than P-waves and can be distinguished based on their arrival times. In the next section we consider imaging using only P or S waves, and gating out the unwanted wave modes based on their arrival times.

5.2 Simulating the Effect of Gating

In all real materials, bulk shear waves have a slower velocity than longitudinal waves, therefore we can attempt to separate the modes by calculating the expected arrival times of the various wave modes. We will use the notation “PS” to refer to waves that are transmitted as P waves and mode converted to S waves. The PP, PS/SP and SS modes should have different arrival times, so we can suppress the unwanted parts of the signals. This will not work in general, for instance if there are several scatterers, because the modes will sometimes overlap. The PS and SP signals will always overlap in pulse-echo signals since the distance travelled by each mode is the




Figure 5.2: *The process of gating signals to remove shear waves. a) The original signals showing different mode paths. b) The expected arrival time of the wave modes. c) The signals with shear waves gated out. This particular example shows a pulse-echo signal, where the PS and SP modes always coincide.*

same. This gating method was considered by Simonetti [78] in order to isolate the PP signals, however in that case the aperture of the array was small compared to its distance from the scatterers and no significant S wave reflection was measured. When this is the case gating is fairly trivial, however, a large aperture also results in better image quality [79], so we have used a large aperture array close to the scatterer for our investigation.

The gating algorithm we used is illustrated in figure 5.2. For each transmit/receive combination we calculate the expected arrival times of each wave mode from the scatterer. We then use these times and the width of the transmitted signal to calculate a gating point, finally we use a smooth sigmoid function to gate the signals. In our simulations we know exactly where the scatterer is, however with no prior knowledge of its location one could assume there is a single scatterer located at the current image point, and gate accordingly.

We tested the effect of this gating on simulated signals, the generation of which will now be described.




Figure 5.3: *Layout of the simulations. A 32-element array is positioned above a series of horizontal and vertical cracks. The background material was steel and a 2 MHz 5 cycle hamming tone burst was excited vertically at each transducer location. This 14 mm square denotes the imaging region.*

5.2.1 Simulations and Results

In order to evaluate the effectiveness of gating different wave modes, we simulated 2D scattering from a series of horizontal and vertical cracks embedded in steel. As the main strength of FM is its ability to image small details, we used relatively short cracks of lengths 3, 1.5 and 1 mm. A typical close range ultrasound inspection frequency is 2 MHz, and in steel this gives a longitudinal wavelength of 3 mm and a shear wavelength of 1.6 mm. The input signal was a 5-cycle hamming windowed tone burst. A 32 element array with 1 mm pitch was positioned 10 mm above the scatterer, which was close enough that separation of the shear and longitudinal waves is not trivial. This layout is shown in figure 5.3.

Each dataset was gated in several ways in order to remove or preserve the different modes prior to imaging. Imaging was then performed on these datasets, and on the ungated signals using BF and FM at 10 frequencies from 1.5 to 2.5 MHz. The single-frequency images were summed to produce one image. For BF this improves the vertical resolution, as the (complex) images coherently sum over the scatterer. For FM, it acts as an averaging process to reduce random variations and artefacts in the images.

The first result of our analysis is that gating has little effect on image quality.




Figure 5.4: *The effect of gating on imaging the 3 mm horizontal and vertical cracks. a-d are without gating, e-h have unwanted shear modes removed. Horizontal and vertical cracks imaged with BF and FM are shown. The true crack length is shown as a white line, offset from the true location which is at the centre of the images. It can be seen that gating has little to no effect, and that FM performs better than BF.*

Figure 5.4 shows images of the 3 mm vertical and horizontal cracks generated using BF and FM with the P wave speed. That is, the imaging algorithms were applied assuming the waves behaved acoustically with a speed equal to the P wave speed. The top row (a-d) shows the results with no gating, the bottom row shows the results with the PS, SP and SS waves removed. It can be seen that in this case (and the others we looked at), gating makes no appreciable difference to the results, and furthermore imaging with FM works well in both cases. The shape and position of the image of the cracks are not as accurate with BF as they are for FM. The fact that FM imaging works well even with the additional shear modes present is fortuitous because it means that we can perform imaging successfully even in the cases where gating would be infeasible, such as when imaging scatterers with a large vertical extent. For the rest of the chapter, no gating is performed.

The second result is that, different modes produce different results for the images. Since we perform no gating we can perform imaging using the P or S wave speeds on the same dataset. This allows us to generate two different images of the defect, and depending on the defect one or the other of the images may be better quality.




Figure 5.5: FM imaging of the 3 mm vertical and horizontal cracks with P (a, c) and S (b, d) wave speeds.

Figure 5.5 shows the results of imaging vertical and horizontal cracks with FM using both the P and S wave speeds. It can be seen that the vertical crack is imaged more accurately with the S wave speed, and the horizontal crack is imaged more accurately with the P wave speed. This effect is due to the directionality of the incident wave-field as explained in the next section.

5.2.2 Directionality

A separate issue that becomes apparent when imaging with shear waves, is the directional nature of the source. The vertical vibration associated with an array element produces shear waves that vary in strength according to the cosine of the angle with the horizontal. This means that little to no shear waves are transmitted or received (by reciprocity) in a vertical direction from the array. Most of the shear wave energy travels in the direction parallel to the array making shear waves more sensitive to vertical cracks as shown in figure 5.6.

Normally it is assumed that the array elements are omnidirectional, which is acceptable in the region of interest with longitudinal waves and a small aperture. However since we are imaging directly below the array, the directionality of the shear waves becomes important. By considering the directionality we can reduce the magnitude of the artefacts in the shear wave images, and improve the accuracy of the true image. The method used was detailed in [80]; we alter the directionality of the Green's function in the BF and FM equations to match that of the transducers. The effect on imaging is shown in figure 5.6, where a 1.5 mm vertical crack is imaged with and without consideration of the directionality. The magnitude of the artefacts to the




Figure 5.6: By considering the direction nature of shear wave generation, the magnitude of image artefacts can be reduced. Shown here are FM images of 1.5 mm vertical crack with (a) and without (b) consideration of the directionality.




Figure 5.7: Demonstration of super-resolution imaging using FM (a-f) and BF (g-l) in elastic media. Various lengths of horizontal (a-c, g-h) and vertical (d-f, j-l) cracks are imaged. The FM image sizes follow the crack sizes with much better accuracy than the BF images.

side of the crack are greatly reduced using this method.

5.2.3 Super-Resolution

At the 2 MHz centre frequency we used for simulation and imaging, the three crack lengths correspond to 1 , $\frac{1}{2}$ and $\frac{1}{3}$ of the longitudinal wavelength (3 mm). The results in figure 5.7 show that the FM images closely follows the length of the crack even down to the third-wavelength scale, whereas the BF image remains unchanged as the crack length varies. Imaging was performed without gating the signals, and using the P wave speed for imaging.

5.3 Experiments

5.3.1 Experimental Setup

In order to test the elastic imaging of cracks experimentally, we obtained several $50 \times 50 \times 600$ mm extruded aluminium blocks and cut 2D cracks in them using EDM (Electrical Discharge Machining). In order to match the simulation, it would be necessary to use a very large block of metal, and cut a crack in the middle of it, which is difficult to fabricate. Instead we used a more realistic thickness 50 mm, and cut cracks from the outside. The 600 mm length was chosen so that reflections from the ends of the block did not interfere with the measurements. Three cracks were cut with depths of 0.5, 1.0 and 2.0 mm, all with widths of approximately 0.3 mm. Additionally, the face of a aluminium spare block was machined flat (rather than left with an extruded finish) to see if the smoother surface made a difference to the coupling, which was initially quite poor. The machining made no difference, and it was found that the poor coupling was due to the use of water rather than gel as a couplant. With gel, very consistent coupling was achieved. The layout for the experiments is shown in figure 5.8.

Due to the results of simulated imaging with various apertures (presented later in chapter 6) and from chapter 3 we know that a large aperture is very important for imaging, therefore a very wide array was used to image the cracks. This array has 128 elements, and a pitch of 1.5 mm. Additionally the elements have a relatively small width (0.45 mm) so that waves can be transmitted in every direction, and large length (15 mm) to give as high an SNR as possible. The large pitch can result in under-sampling, as the P and S wavelengths in aluminium at the 2.4 MHz centre frequency of the array are 2.7 and 1.3 mm respectively.

Before performing the experiment, a time-domain simulation of the layout was performed using the Finite Difference Method (FDM) discretisation of the elastic wave equation. The code uses Convolutional-Perfectly Matched Layers (C-PMLs) to absorb waves that reach the edge of the simulation domain [17]. A displacement is




Figure 5.8: *The experimental set-up. a) A 0.5, 1.0 or 2.0 deep EDM notch is cut into the side of a $600 \times 50 \times 50$ mm aluminium block. The notch is approximately 0.3 mm thick. b) Photograph of the array and block with a 2 mm notch.*

excited at the location of one transducer and the displacement field is recorded at the locations of all the transducers. Additionally the entire wave field is outputted as a video, with P-waves displayed blue and S-waves displayed red. This colouring is achieved by setting the red channel equal to the isostatic stress and the blue channel equal to the minimum principal stress. Due to the proximity of the source to the front wall, an unstable (and unrealistic; it isn't observed in the experiments) surface wave is generated that was removed by subtracting the results of simulation with only the front wall present. This is only mentioned because it is visible in the wave fields shown later.

First we may examine the simulation results for the case with no scatterer present. The labelled signals received for a single transmission are shown in figure 5.9. The labelled wave modes are: 1: P-P, 2: P-S, 3: P-P-P-P, 4: P-P-P-S, 5: P-P-P-P-P-P, 6: P-P-S-S, 7: S-S. The modes can be in any order; P-P-P-S arrives at exactly the same time as S-P-P-P, P-S-P-P and P-P-S-P.

One can see the reflected P-P wave arrives first, followed by the mode converted




Figure 5.9: a) The waves received by each transducer when a simulated tone-burst is transmitted on element 36 in a defect-free block. b) The various modes corresponding to the reflections in (a). Red and blue denote P and S waves respectively. Where mode conversion occurs the order is not important, i.e. P-S and S-P waves arrive simultaneously.

P-S and S-P waves. The S-S waves are the slowest, and are followed by waves which are reflected multiple times between the front and back walls (e.g. the P-P-P-P waves). The directionality of the transmission can also be seen—for example no shear waves are received in pulse-echo because the vertical excitation of the transducer only results in P waves in the vertical direction.

The pulse-echo signals from each transducer with the 1 mm crack introduced are shown in figure 5.10. Although P-P and P-S/S-P reflections do occur, by far the strongest reflection from the crack is the S-S mode, therefore we have used the




Figure 5.10: *The simulated pulse-echo signals for a block containing a 1 mm crack. The three simultaneous reflections are the P-P, P-P-P-P and P-P-P-P-P-P reverberations between the front and back walls. The largest reflection from the crack is the S-S mode at 45 μs . The P-P and P-S/S-P reflections from the crack can also be seen around 20 and 28 μs respectively, but they are both very weak.*

S-speed for imaging. Frames from the video of scattering from transducer 36 are shown in figure 5.11. They show the highly directional nature of the reflections from the crack, and the complexity of the scattering, where several reverberations of the P mode occur before the S-S reflection from the crack is received. Fortunately in this geometry the S-S reflection from the crack does not overlap with any others, so it can be extracted by gating out the other signals.

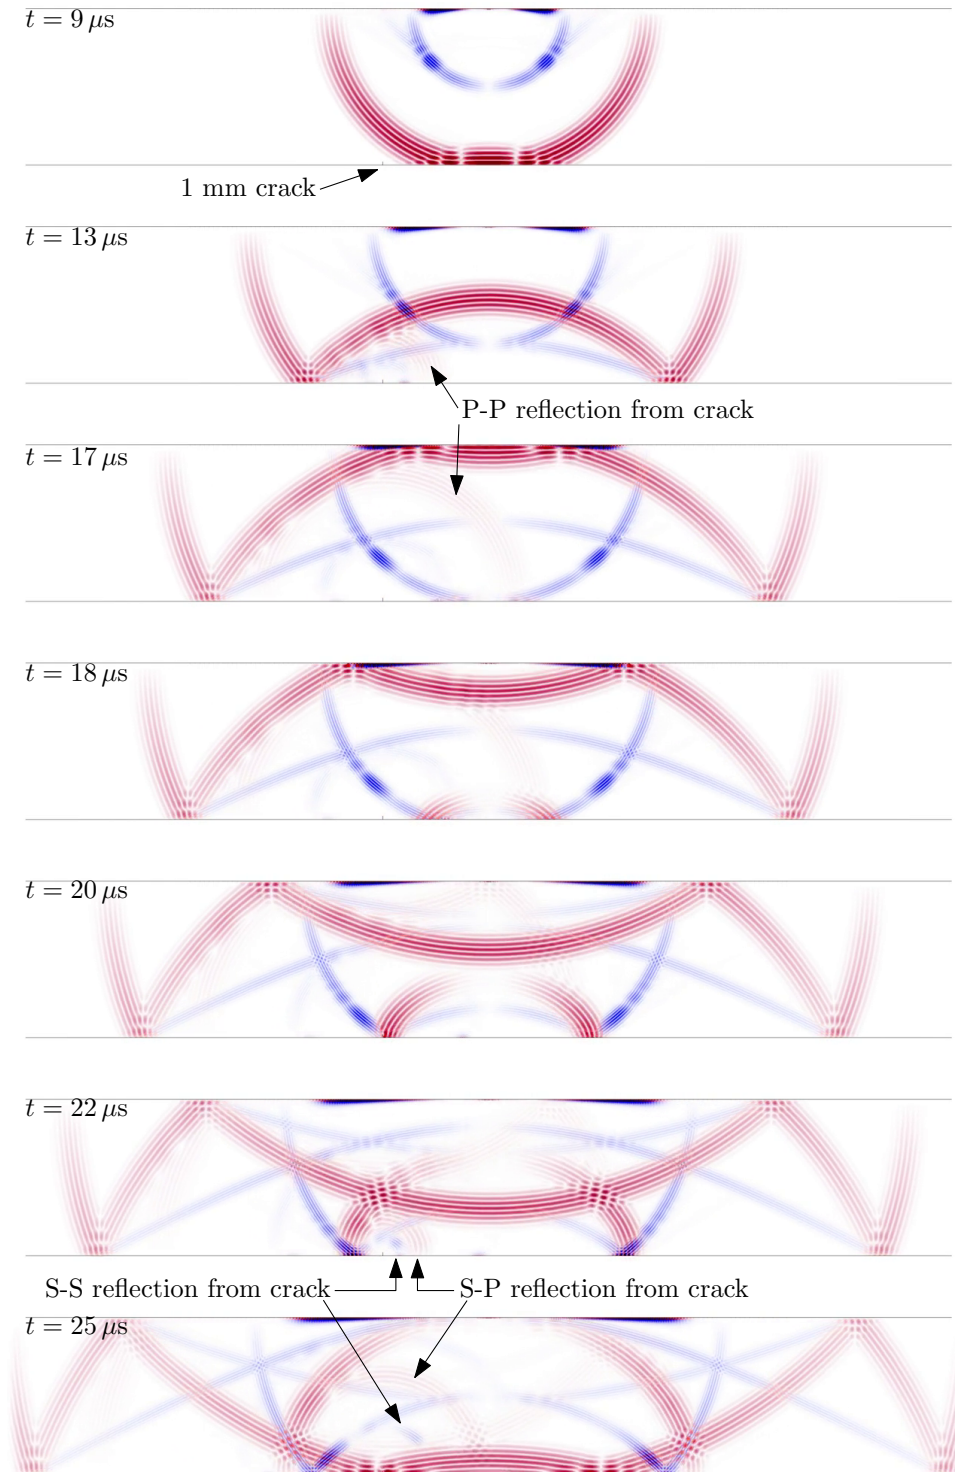


Figure 5.11: Simulated wave fields resulting from transmission from element 36 for the 1 mm crack. P and S waves are coloured red and blue respectively. It can be seen that the “S-S” reflection is actually due to a more complex interaction of the mode converted S-P wave in addition to the S wave. The strength and high degree of directionality of the reflection shows the advantage of using a large aperture array and imaging with shear waves.

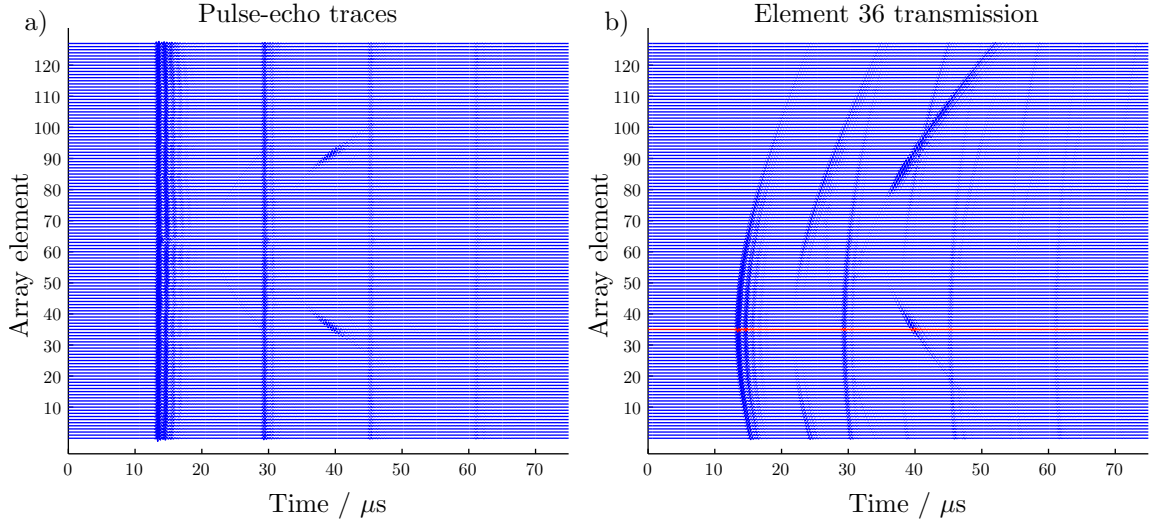


Figure 5.12: a) *Experimental pulse echo signals for the 1 mm crack. The S-S reflections are clearly visible and P-P is just visible.* b) *The received signals for all array elements when transmitting on element 36. Both (a) and (b) closely match the simulations in figures 5.9 and 5.10. (There are actually slight timing differences due to the actual block being 51.04 mm thick, rather than 50 mm as in the simulation.)*

5.3.2 Experimental procedure

The experiment was performed by transmitting a single pulse on each array element, and recording the signals at 20 MHz, for 1500 samples. The gain was constant, and 31 times averaging was used to improve the SNR. The signals corresponding to figures 5.9 and 5.10 are shown in figure 5.12. The average spectrum for all the signals is shown in figure 5.13. $11 \mu\text{s}$ worth of each signal was used for the imaging, shifted using the pre-beamforming method described in the previous chapter. This gave an overall dimension of $F_{d,t}$ of 56192×28160 . Before imaging we first had to perform coupling compensation and incident field removal.

5.3.3 Coupling compensation

Although coupling using gel was very consistent, it could still be improved by numerically correcting for the transmission and reception coupling coefficients which were assumed to be independent. A simple but effective algorithm was developed to




Figure 5.13: Average spectrum of the experimental signals. This was calculated by averaging the spectrum over all $5 \mu\text{s}$ Hamming windows in all recorded signals. The centre frequency is approximately 2.4 MHz and the half-power bandwidth is 0.5 MHz.

achieve this, based on reciprocity. Most of the “coupling” errors turned out to be due to electrical or physical differences between the actual array elements, rather than actual differences in the coupling between the array and the aluminium, however these have the same effect and can be lumped in with the actual coupling effects. The “before” and “after” channel energies ($\int u(t)^2 dt$) for each transmit–receive combination and each crack length are shown in figure 5.14. The increase in consistency is quite dramatic; the only remaining inconsistencies, which occur at 32 element intervals are due to the four-part construction of the array. This is also the reason the first 32 elements are stronger than the others when transmitting and receiving within that sub-array.

The algorithm proceeds as follows:

1. Calculate the channel energy for each transmit–receive combination (see figure 5.14a).

$$E_0(\mathbf{x}_i, \mathbf{x}_j) = \int v(\mathbf{x}_i, \mathbf{x}_j, t)^2 dt$$




Figure 5.14: (The square root of) the channel energies before (a) and after (b) coupling compensation. The slight anomalies seen every 32 transducers are due to the fact that the array is constructed from four 32 element sub-arrays.

- Calculate the mean logarithmic (i.e. the geometric mean) difference in energy of each row or column from the previous row or column.

$$d_1^{\text{tx}} = d_1^{\text{rx}} = 0$$

$$d_i^{\text{tx}} = \frac{1}{N_x - 1} \sum_{k=2}^{N_x} (\log E_0(\mathbf{x}_i, \mathbf{x}_k) - \log E_0(\mathbf{x}_{i-1}, \mathbf{x}_{k-1})) \quad i = 2 \dots N_x$$

$$d_i^{\text{rx}} = \frac{1}{N_x - 1} \sum_{k=2}^{N_x} (\log E_0(\mathbf{x}_k, \mathbf{x}_i) - \log E_0(\mathbf{x}_{k-1}, \mathbf{x}_{i-1})) \quad i = 2 \dots N_x$$

- Calculate the cumulative difference (so now each value represents roughly the ratio of energy transmitted/received by a transducer compared to the first one).

$$m_i^{\text{tx}} = \exp \left[\sum_{k=1}^i d_k^{\text{tx}} \right]$$

- Compute the updated energy matrix:

$$E_1(\mathbf{x}_i, \mathbf{x}_j) = \frac{E_0(\mathbf{x}_i, \mathbf{x}_j)}{m_i^{\text{tx}} m_j^{\text{rx}}}$$

5. Repeat this process, using E_1 in place of E_0 to obtain E_2 . Finally scale each channel so that its energy is equal the geometric mean of E_1 and E_2 :

$$v'(\mathbf{x}_i, \mathbf{x}_j, t) = v(\mathbf{x}_i, \mathbf{x}_j, t) \frac{E_1(\mathbf{x}_i, \mathbf{x}_j) E_2(\mathbf{x}_i, \mathbf{x}_j)}{E_0(\mathbf{x}_i, \mathbf{x}_j)^2}$$

The second iteration is necessary to remove errors that arise due to the fact that transmission and reception coupling are conflated within a single channel.

5.3.4 Incident field removal

A rather significant detail that has been overlooked until now is the fact that the sampling methods require the scattered field measurements ($u^s(\mathbf{r}, \mathbf{x}, t)$) as input, not the total field that we actually measure ($u(\mathbf{r}, \mathbf{x}, t)$). In order to obtain u^s we must measure u , and also u^i by performing the measurement on a defect-free but otherwise identical specimen. We can then perform the baseline subtraction $u^s = u - u^i$. However in practice, very slight differences in the measurements of u and u^i where they should be equal cause significant errors in u^s . This is especially problematic because u^i contains much more energy than u^s , so any errors in u^i can easily drown out u^s . In the next chapter, we present a method for performing this baseline subtraction in acoustic media that minimises errors, however due to the complexity of the incident field in the elastic case, we gated out the incident waves in the time domain.

The incident fields $u^i(\mathbf{r}, \mathbf{x}, t)$ were measured for each \mathbf{r}_i , then enveloped using the Hilbert transform, and thresholded. The resulting mask was manipulated using dilation and erosion operators to remove anomalies, and finally blurred to allow a smooth transition from gated to non-gated regions. The total signals, $u(\mathbf{r}, \mathbf{x}, t)$, were then multiplied by this mask to remove the incident field.

The downside to this approach is that it removes the shadows of the scatterer, which




Figure 5.15: *Beamforming images of the cracks. a) 0.5 mm b) 1.0 mm c) 2.0 mm. Note that the backwall is nominally at $y = 0$, but there is a small error in the position of the cracks so the region below the backwall is shown.*

are usually the most energetic part of the scattered field. One can see that the differences in channel energy between the crack sizes in figure 5.14 are primarily due to shadowing of the pulse-echo, and near-pulse-echo signals, and there are clear differences in the magnitude of the shadows. Sizing can be performed based on the magnitude of the scattered field [81, 82], however this method has downsides, primarily the need for accurate calibration, and potentially prior knowledge of the shape of the defect.

5.3.5 Images and discussion

We computed the beamforming images in the time domain, and the results are shown in figure 5.15. Although the difference in size of the crack is only on the edge of the Rayleigh limit, the images are essentially useless for sizing the cracks; the size of the main “blob” does not correlate with crack size at all. The magnitude does (note the scales), however as noted before, there are significant downsides to magnitude-based sizing, and we wish to achieve sizing based on the image geometry instead.

Next, FM at 2.5 MHz, and TDLSM were used to produce images of the 0.5, 1.0 and 2.0 mm datasets. The FM images are shown in figure 5.16. Due to the limited information content of single frequency measurements the image does not reconstruct the crack.

Next, TDLSM was used to perform imaging. The following inversion parameters




Figure 5.16: *FM images of the cracks at 2.5 MHz.*

were used:

Singular Values computed: 8000

Target signal ($s(t)$): 4-cycle Hamming windowed 2.55 MHz tone-burst. Although the centre frequency of the reflections was actually 2.4 MHz, a slightly higher frequency target was chosen, as deconvolution of a few sample traces gave smoother results at this frequency. In other words, the incident field calculated by TDLSM was smoother. This intuition was apparently correct; the TDLSM images using 2.55 MHz are somewhat better than those using a target at 2.4 MHz.

Signal window length (τ): 11 μ s.

Sound speed (shear): 3170 m/s.

The first few singular values for the 1 mm crack are shown in figure 5.17, and some examples of the pre-beamformed signals, and the corresponding target are shown in figure 5.18.

The resulting images are shown on the left in figure 5.19, together with the BF images for comparison. A good image was obtained for the 0.5 mm crack, however the 1.0 and 2.0 mm cracks are both undersized. The results are not as good as had been hoped, however the difference between then 2.0 mm crack and the others mean that it could be distinguished from smaller cracks, which was not the case with beamforming.




Figure 5.17: The first 500 measured singular values of the time domain operator $F_{d,t}$ for the 1 mm crack. Note that the values come in nearly equal pairs. The symmetry of the scatterer means that there are pairs of odd and even symmetric functions with the same singular value.

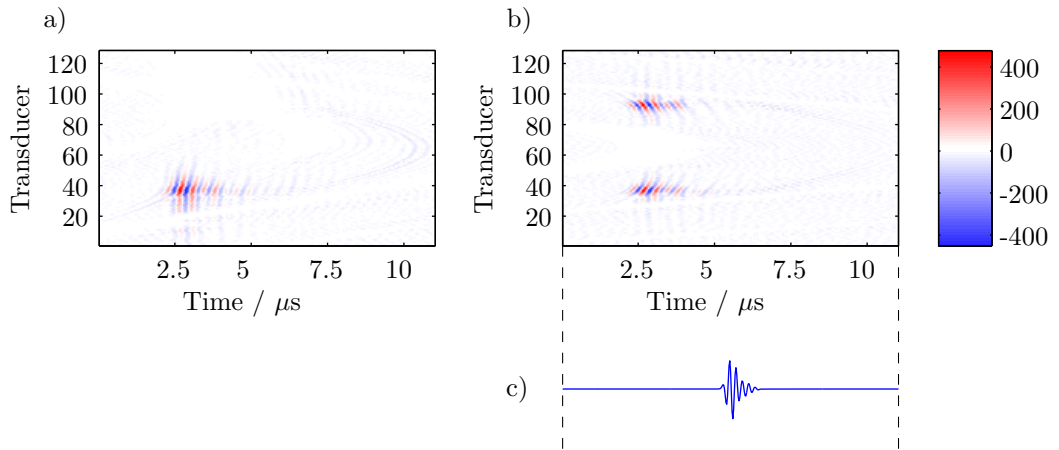


Figure 5.18: Pre-beamformed signals. a) Received for transmitter 36, b) Pulse-echo, c) The form of the target signal, its exact location and amplitude depends on the receiving transducer and image point. But the key point is due to the pre-beamforming all possibilities occur later than the one shown, but before $10 \mu\text{s}$. The signals are 220 samples long, so the $F_{t,d}$ matrix operator has dimensions 56192×28160 .

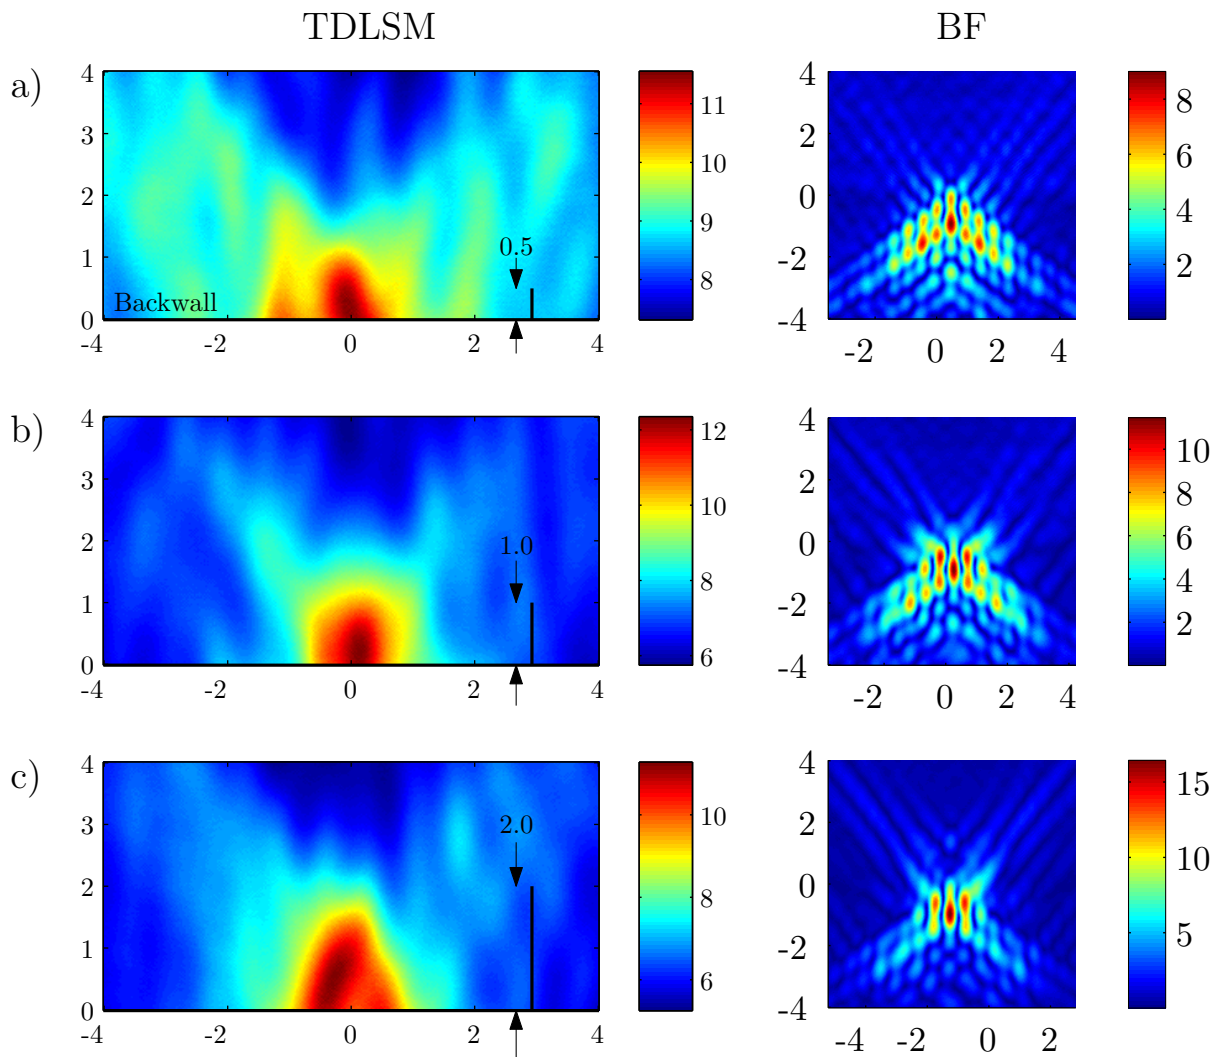


Figure 5.19: The final TDLSM images (left) for the a) 0.5, b) 1.0 and c) 2.0 mm cracks. The 1 mm and 2 mm cracks are both undersized, however, the 2 mm crack is noticeably larger than the others, which was not the case for beamforming (reproduced on the right). The wavelength is 1.3 mm.

5.4 Conclusion

From the simulated images produced we can conclude that the Factorisation Method is robust to the presence of additional wave modes with unexpected speeds. Gating out these additional modes is not necessary to produce high quality reconstructions. The simulation also shows that sampling methods can theoretically achieve super-resolution imaging of a crack in free space in elastic media where both shear and longitudinal waves are present, and cannot be measured independently. Imaging can be performed with both shear and longitudinal wave speeds, the choice of the best mode to use depends on the scatterer being imaged.

In experimental situations however, the noise is too great for the single-frequency sampling methods, so TDLSM is necessary to achieve a remotely coherent image. Both TDLSM and broadband BF are vast improvements over the single frequency methods. TDLSM's sizing capability is slightly better than BF's and it is also much better at accurately locating the cracks as it is less sensitive to global timing errors in the measurements than BF.

The TDLSM image is not perfect however, and this is probably due to the low SNR of the tip reflections in comparison to the corner reflection from the root of the crack, and the fact that the free space Green's function was used, rather than one that considers backwall reflections. Backwall reflections for the simpler acoustic imaging case are considered in the next chapter.

Chapter 6

Using mirror reflections for imaging—Theory

6.1 Background

Commonly when imaging industrial components using ultrasound, only one side of the component is accessible, or both sides are accessible but it is difficult to align two arrays on opposite sides, examples include pipes, storage tanks and pressure vessels. This means that it is only possible to insonify the component from a small subset of directions, i.e. from a limited aperture. In order to form an accurate image of an object, one would ideally like to illuminate the object from every direction, and record the scattered waves in all directions. This would require the use of a circular or spherical array of sensors that surround the object—a configuration that is rarely practical.

The limited view aperture can dramatically degrade the reconstruction, as demonstrated in figure 6.1. In this numerical example, sound waves are transmitted from an array of transducer elements through a 2D fluid, and scattered from a kite-shaped hole in the fluid. The kite scatterer is shown in Fig 6.1a. In Fig 6.1b a circular array of transducers (not shown) is used to excite waves and record the reflections. The image is then produced from this data using the beamforming approach described




Figure 6.1: *The effect of array aperture on the completeness of the resulting image. A kite-shaped void (a) is imaged using the beamforming algorithm with b) a full-view ring array surrounding the void; and c) the top 90° of the array. The aperture restriction leads to an incomplete image of the scatterer.*

in chapter 2. This image is far superior to that of Fig 6.1c, which only uses data from the top quarter of the array. In this limited-view case there are many scattered waves that cannot be detected which leads to severe image degradation. Therefore, the aperture used for imaging has a dramatic effect on the completeness of the reconstructed image.

The problem of the limited aperture is a pervasive one, and has been widely studied in the literature. In the context of Diffraction Tomography, Mora [24] and Pan [83] show that the effects of a limited view can be described by a linear image filter. However, this is only true for linear scattering models such as the Born approximation in which it is assumed that the scatterer is small and mostly transparent. These models are not very realistic in many situations, so this result cannot be generalised to situations involving strong scatterers such as holes. Zinn [21] has attempted non-linear imaging of an impenetrable scatterer using an aperture of 90° . As expected, the numerical results are far from complete reconstructions. Various studies of the effects of small apertures on imaging have been made [84, 85, 86]. Ochs [87] has shown that an aperture of 180° can be sufficient, numerically at least, to completely image a scatterer. Unfortunately even a 180° aperture is rarely practical. Alternative solutions have been proposed by Mora [24], Nolan [88] and more recently Natterer [89]. Nolan uses a vertical reflector to the side of the horizontal transceiver array in order to provide additional views of the scatterer. Beamforming is used to separate the waves

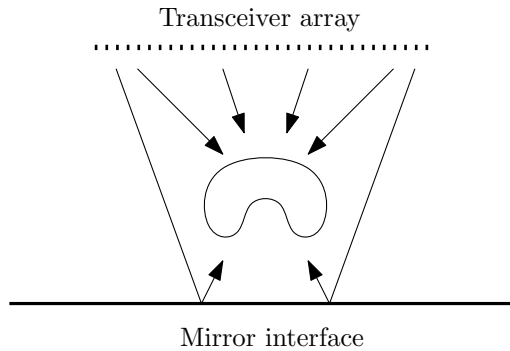


Figure 6.2: *A schematic of the solution to the limited-view problem proposed by Mora. A strongly scattering reflector is used as a mirror to provide additional views of the object to be imaged.*

that travel directly to the scatterer from those that are reflected from the known mirror. Mora uses a horizontal reflector beneath the scatterer to provide additional views. A schematic of this solution is shown in figure 6.2. In his paper, Mora uses an iterative inversion method to reconstruct a weak scatterer above the strongly scattering planar interface. The slowness field is iteratively refined until simulations of the wave propagation match the measurements. Although this provides a good reconstruction of the scatterer, there are some problems with iterative inversion algorithms. They tend to be very computationally expensive, and the presence of local minima in the objective function can cause convergence problems [90, 91].

For this reason, we have attempted to combine the use of a non-linear non-iterative inversion algorithm with Mora’s mirror-like reflector. The use of a non-iterative inversion technique allows us to explore the effects of various aspects of the imaging problem while avoiding the computational burden of an iterative inversion scheme. In addition, the geometric configuration of figure 6.2 is of practical interest as there are physical imaging problems that closely match this situation. For example in Non-Destructive Testing (NDT), it is common to want to image defects in thick metal plates or pipe walls where only one side is accessible. The back wall of the component (or the inside of the pipe) is a very strong reflector that can be used as a mirror.

The purpose of this chapter is to discover how information about the geometrical

properties of the object is conveyed to the scattering measurements, and to find a link between imaging parameters such as aperture size and noise level, and the level of detail we can hope to extract from the measurements. In section 6.2 we define a mathematical formalism that links the information contained in the limited aperture measurements with a mirror, to a related ideal full-view case. This relationship is independent of the scattering model used, and applies to both penetrable and impenetrable scatterers. In section 6.3 we then simulate scattering from an impenetrable scatterer, and invert the measurements using Beamforming, and the Factorisation Method. This avoids the problems of global optimisation and scattering approximations. We find that combining this advanced imaging method with the additional data contained in reflections from the mirror results in a reconstruction that is superior to that of standard beamforming.

The presence of the backwall also enables through-transmission measurements, which are required for Time of Flight Tomography (ToFT). The ToFT results provide an accurate but low resolution map of the speed of sound, which is used to augment the BF/DT image using an algorithm developed by Huthwaite called Hybrid Algorithm for Robust Breast Ultrasound Tomography (HARBUT) [92]. This algorithm is briefly described in section 6.4. The results of section 6.2 also apply to this algorithm.

The conclusions of this chapter are drawn from mathematical and numerical analyses. The material in this chapter has been previously published in [79]. Experimental considerations and results are presented in the next chapter.

6.2 Information content of the reflection

The scattering scenario that we are attempting to image is shown in figure 6.3a. An array of transducers is present in the far-field above a strongly scattering mirror interface. A scattering object is embedded in background medium above the mirror interface. Each transducer can transmit and receive acoustic waves through the 2D background medium. Acoustic waves are pure pressure waves as found in liquids and gases; there are no shear waves.

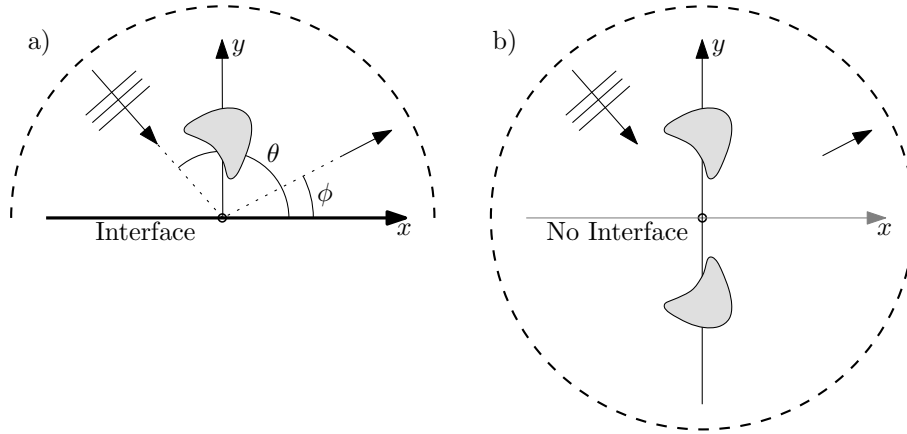


Figure 6.3: *Two related scattering models. a) A limited-view configuration with a semi-circular array and an impenetrable planar interface along the line $y = 0$. b) A related full-view configuration with the scatterer and transducers mirrored and the interface removed.*

In figure 6.3b a related full-view scenario is shown in which the scatterer and transducer array have been reflected about the mirror interface, which has been removed. This configuration provides the ideal imaging scenario that has led to the image in figure 6.1b. These two cases are closely related. To show how, consider the scattering of a wave transmitted from a transducer in the limited-view case. The scattering of this wave is equivalent to the scattering of an identical wave in the full-view case, plus the scattering of a wave transmitted from the mirrored transmitter with a suitable phase shift to account for the boundary condition at the mirror. Because these two waves—the direct and mirrored waves—are transmitted simultaneously, the reflections due to each of them are mixed and cannot be separated. Therefore some information is lost relative to the ideal full-view configuration where each element of the array can be excited separately. It is our aim to characterize this information loss and understand the effect it has on image reconstruction.

In previous chapters we have defined the far-field operator, F in terms of unit direction vectors $\hat{\mathbf{r}}$ and $\hat{\mathbf{x}}$. For this chapter we use a slightly different notation, using angles θ , and ϕ instead, so that F is defined by

$$Fu_{\infty}^s(\phi) = \int v(\theta, \phi) u_{\infty}^i(\theta) d\theta, \quad (6.1)$$




Figure 6.4: A diagram of the matrix representation of the far-field operator. Only a square subset is accessible in limited-view situations. The main diagonal of the matrix corresponds to ‘pulse-echo’ measurements in which the same transducer is used for transmission and reception of the waves. The ‘through-transmission’ diagonals are those in which opposite transducers are used for transmission and reception.

where $u_{\infty}^i(\theta)$ is the complex amplitude of the incident plane waves in the θ direction, and $v(\theta, \phi)$ is the scattered field in the scattering direction ϕ when the object is illuminated with a plane wave from θ . The subscript ∞ refers to the fact that the measurement is performed in the far-field as before. The values of $v(\theta, \phi)$ can be measured experimentally by exciting a plane wave for each angle θ and measuring the scattered field at each angle ϕ . If a sufficient number of measurements are performed such that the data is not under-sampled, the continuous F operator can then be recovered [93]. For numerical computation, the discrete version of F is found by measuring the scattering matrix for a circular array of evenly spaced transducers in the far-field.

The complete far-field operator F can only be obtained in full-view configurations; i.e. when the transducer array completely surrounds the scatterer. For a 180° limited-view in free space only a limited part of F can be measured. This is illustrated in figure 6.4 which shows a diagram of the matrix associated with F . Each entry of the matrix corresponds to one pair of illumination and detection angles. The main diagonal corresponds to the so-called pulse-echo measurements in which the same transducer excites an incident field and then receives the scattered field. The diagonals labelled ‘through-transmission’ are the measurements in which one transducer excites a wave and then the *opposite* transducer records the scattered field. The through-transmission measurements cannot be obtained in the limited-view case.

We now turn our attention to the limited-view problem when a mirror interface is added. Consider the two configurations in figure 6.3 and let T be the subset of the far-field operator measured with the limited view shown in figure 6.3a and F the operator corresponding to figure 6.3b. The two operators are related by

$$T(\theta, \phi) = F(\theta, \phi) + \alpha F(-\theta, \phi) \quad \theta, \phi \in [0, \pi], \quad (6.2)$$

where $\alpha = 1$ or -1 for a sound-hard or sound-soft mirror interface respectively. Note that T is defined for $\theta, \phi \in [0, \pi]$, while F is defined over the wider interval $[-\pi, \pi]$. The first term on the right hand side of eq. (6.2) contains the direct waves that do not interact with the mirror interface. The second contains waves that reflect from the mirror interface. These appear to have come from mirrored sources, therefore the sign of θ is changed. The correct choice of α allows the boundary conditions on the interface to be satisfied. By considering the symmetry of the model in figure 6.3b about the $y = 0$ line we can see that F has the following property

$$F(\theta, \phi) = F(-\theta, -\phi), \quad (6.3)$$

and by considering the reciprocity of non-attenuative scattering we find that

$$F(\theta, \phi) = F(\phi, \theta). \quad (6.4)$$

In matrix terms, eq. (6.3) is equivalent to F being centrosymmetric (symmetric about the centre of the matrix), and eq. (6.4) is equivalent to F being symmetric. From these properties it follows that F must be bisymmetric, i.e. symmetric with respect to both its diagonal and its antidiagonal. It has the block structure shown in figure 6.5 where A is a symmetric matrix and B is persymmetric (symmetric with respect to its antidiagonal) [94]. The superscript T and F refer to the transpose and the flip transpose. For a finite number of illumination and observation angles eq. (6.2) can be given in matrix form as




Figure 6.5: The block structure of the bisymmetric far-field operator F . The superscript T and F refer to the transpose and flip-transpose. A is a symmetric matrix and B is persymmetric.

$$T = A + \alpha BR, \quad (6.5)$$

where A and B are the quadrants of F as shown in figure 6.5. R is the ‘exchange matrix’—the identity matrix rotated by 90° . It has the same effect as the negation of θ in eq. (6.2). T is therefore the sum or difference of the two independent quadrants of F . Due to reciprocity, T is also symmetric.

If we could measure both A and B then it would be possible to determine the full view matrix F using the block structure of figure 6.5. However we can only measure T , which is a linear combination of A and B . This shows that the mirror does not allow the direct retrieval of the complete full-view scattering information, thus resulting in a partial loss of information.

Our goal is to determine how the loss of information contained in the measurement of T affects the reconstruction. The nature of the image degradation will depend on the algorithm used to produce the image. As stated before, we are using the Beamforming and Factorisation Methods and will now investigate how they behave when T is used for imaging instead of F . For this it is convenient to formulate these methods in terms of the singular value decomposition of the far-field operator.

6.2.1 Singular Value Decomposition of the Far-Field Operator

The full-view operator F can be decomposed using Singular Value Decomposition (SVD) as before:

$$F = \sum_{n=-\infty}^{\infty} \sigma_n |\xi_n\rangle \langle \psi_n| \quad (6.6)$$

so that

$$F|\psi_n\rangle = \sigma_n |\xi_n\rangle \quad (6.7)$$

$$F^*|\xi_n\rangle = \sigma_n |\psi_n\rangle, \quad (6.8)$$

where F^* is the conjugate transpose of F , and the ket symbol is used according to Dirac's bra-ket notation [40]. Similarly for T we have

$$T_\alpha = \sum_{n=-\infty}^{\infty} \lambda_{n;\alpha} |v_n\rangle_\alpha \langle u_n|_\alpha \quad (6.9)$$

$$T_\alpha |u_n\rangle_\alpha = \lambda_{n;\alpha} |v_n\rangle_\alpha \quad (6.10)$$

$$T_\alpha^* |v_n\rangle_\alpha = \lambda_{n;\alpha} |u_n\rangle_\alpha, \quad (6.11)$$

where the subscript α refers to the boundary condition. We can now show that the singular functions of F ($|\xi_n\rangle$ and $|\psi_n\rangle$) can be given in terms of the singular functions of $T_{\alpha=1}$ and $T_{\alpha=-1}$. By using the block structure of F in figure 6.5 and eq. (6.5) it can be shown that

$$F \begin{bmatrix} \alpha R|u_n\rangle_\alpha \\ |u_n\rangle_\alpha \end{bmatrix} = \lambda_{n;\alpha} \begin{bmatrix} \alpha R|v_n\rangle_\alpha \\ |v_n\rangle_\alpha \end{bmatrix} \quad (6.12)$$

$$F^* \begin{bmatrix} \alpha R|v_n\rangle_\alpha \\ |v_n\rangle_\alpha \end{bmatrix} = \lambda_{n;\alpha} \begin{bmatrix} \alpha R|u_n\rangle_\alpha \\ |u_n\rangle_\alpha \end{bmatrix}. \quad (6.13)$$

Therefore the complete SVD of F can be obtained from the SVDs of two T operators measured for a sound-soft ($\alpha = -1$) and sound-hard ($\alpha = 1$) interface. In particular the sound-hard configuration provides the symmetric component of the SVD of F

$$\psi_n^{\text{sym}} = \frac{1}{\sqrt{2}} \begin{bmatrix} R|u_n\rangle_{\alpha=1} \\ |u_n\rangle_{\alpha=1} \end{bmatrix} \quad (6.14)$$

$$\xi_n^{\text{sym}} = \frac{1}{\sqrt{2}} \begin{bmatrix} R|v_n\rangle_{\alpha=1} \\ |v_n\rangle_{\alpha=1} \end{bmatrix} \quad (6.15)$$

$$\sigma_n^{\text{sym}} = \lambda_{n;\alpha=1}, \quad (6.16)$$

and the sound-soft interface gives the antisymmetric components

$$\psi_n^{\text{ant}} = \frac{1}{\sqrt{2}} \begin{bmatrix} -R|u_n\rangle_{\alpha=-1} \\ |u_n\rangle_{\alpha=-1} \end{bmatrix} \quad (6.17)$$

$$\xi_n^{\text{ant}} = \frac{1}{\sqrt{2}} \begin{bmatrix} -R|v_n\rangle_{\alpha=-1} \\ |v_n\rangle_{\alpha=-1} \end{bmatrix} \quad (6.18)$$

$$\sigma_n^{\text{ant}} = \lambda_{n;\alpha=-1}. \quad (6.19)$$

The $1/\sqrt{2}$ factors are necessary to normalise the singular functions. Therefore, depending on the boundary condition at the mirror interface we will be able to recover either the symmetric or antisymmetric singular vectors and singular values of

F . This characterises the information loss of the limited-view measurements with a mirror, relative to the ideal full-view case and is independent of the scattering model and material properties of the scatterer. The same result holds if the eigenvector decomposition is performed instead of SVD.

6.3 Imaging impenetrable scatterers

6.3.1 Beamforming

Under the full-view configuration of figure 6.3b the BF image at \mathbf{z} is given by [38]

$$I_{bf}(\mathbf{z}) = |\langle g_{\mathbf{z}} | F | g_{\mathbf{z}} \rangle|, \quad (6.20)$$

or in terms of the SVD of F

$$\begin{aligned} I_{BF}^F(\mathbf{z}) &= \sum_{n=-\infty}^{\infty} \sigma_n \langle g_{\mathbf{z}} | \xi_n \rangle \langle \psi_n | g_{\mathbf{z}} \rangle \\ &= \sum_{n=-\infty}^{\infty} \sigma_n^{\text{sym}} \langle g_{\mathbf{z}} | \xi_n^{\text{sym}} \rangle \langle \psi_n^{\text{sym}} | g_{\mathbf{z}} \rangle \\ &\quad + \sum_{n=-\infty}^{\infty} \sigma_n^{\text{ant}} \langle g_{\mathbf{z}} | \xi_n^{\text{ant}} \rangle \langle \psi_n^{\text{ant}} | g_{\mathbf{z}} \rangle, \end{aligned} \quad (6.21)$$

where the functional has been split into symmetric and anti-symmetric components.

The same functional in eq. (6.20) can be applied to T to produce a limited view image. However, instead of considering the free space Green's function (equations 34 and 35 in chapter 2) we use a Green's function that accounts for the interface according to the standard approach used in Kirchhoff migration [95]. In this case, optimum focussing with the mirror present can be achieved using the Green's function of the half space bounded by the mirror

$$|g_{\mathbf{z};\alpha}^\dagger\rangle = |g_{\mathbf{z}^+}\rangle + \alpha R|g_{\mathbf{z}^-}\rangle \quad \theta \in [0, \pi], \quad (6.22)$$

where $|g_{\mathbf{z}^+}\rangle$ is the free space Green's function and $|g_{\mathbf{z}^-}\rangle$ is the free space Green's function corresponding to a source that mirrors the actual source. Therefore the BF image of T_α is

$$I_{bf}^{T_\alpha}(\mathbf{z}) = \left| \langle g_{\mathbf{z};\alpha}^\dagger | T_\alpha | g_{\mathbf{z};\alpha}^\dagger \rangle \right| \quad (6.23)$$

$$= \sum_{n=-\infty}^{\infty} \lambda_{n;\alpha} \langle g_{\mathbf{z};\alpha}^\dagger | v_{n;\alpha} \rangle \langle u_{n;\alpha} | g_{\mathbf{z};\alpha}^\dagger \rangle. \quad (6.24)$$

Now consider the following equality:

$$\langle \psi_n | g_{\mathbf{z}} \rangle = \left\langle \frac{\alpha R u_{n;\alpha}}{\sqrt{2}} \middle| g_{\mathbf{z}^-} \right\rangle + \left\langle \frac{u_{n;\alpha}}{\sqrt{2}} \middle| g_{\mathbf{z}^+} \right\rangle \quad (6.25)$$

$$= \frac{\langle \alpha R g_{\mathbf{z}^-} + g_{\mathbf{z}^+} | u_{n;\alpha} \rangle}{\sqrt{2}} \quad (6.26)$$

$$= \frac{\langle g_{\mathbf{z}}^\dagger | u_{n;\alpha} \rangle}{\sqrt{2}} \quad (6.27)$$

$$\langle g_{\mathbf{z}} | \xi_n \rangle = \frac{\langle v_{n;\alpha} | g_{\mathbf{z}}^\dagger \rangle}{\sqrt{2}}. \quad (6.28)$$

Substituting these into eq. (6.21) and using eq. (6.23) one obtains

$$2I_{BF}^F(\mathbf{z}) = I_{BF}^{T_{\alpha=1}}(\mathbf{z}) + I_{BF}^{T_{\alpha=-1}}(\mathbf{z}). \quad (6.29)$$

Therefore, the doubled full-view image $2I_{BF}^F$ is the sum of the limited-view images with sound-soft and sound-hard interfaces that would be reconstructed via eq. (6.23). Moreover $I_{BF}^{T_{\alpha=1}}$ and $I_{BF}^{T_{\alpha=-1}}$ in eq. (6.29) correspond to the symmetric and antisymmetric terms in eq. (6.21), i.e.

$$I_{BF;\alpha=1}^T = 2 \sum_{n=-\infty}^{\infty} \sigma_n^{\text{sym}} \langle g_{\mathbf{z}} | \xi_n^{\text{sym}} \rangle \langle \psi_n^{\text{sym}} | g_{\mathbf{z}} \rangle$$

and

$$I_{BF;\alpha=-1}^T = 2 \sum_{n=-\infty}^{\infty} \sigma_n^{\text{ant}} \langle g_{\mathbf{z}} | \xi_n^{\text{ant}} \rangle \langle \psi_n^{\text{ant}} | g_{\mathbf{z}} \rangle.$$

6.3.2 Factorisation Method

The full view FM image is obtained by calculating the functional

$$I_{FM}^F(\mathbf{z}) = \left(\sum_{n=-\infty}^{\infty} \frac{|\langle \psi_n | g_{\mathbf{z}} \rangle|^2}{\sigma_n} \right)^{-1} \quad (6.30)$$

$$= \left(\sum_{n=-\infty}^{\infty} \frac{|\langle \psi_n^{\text{sym}} | g_{\mathbf{z}} \rangle|^2}{\sigma_n^{\text{sym}}} + \sum_{n=-\infty}^{\infty} \frac{|\langle \psi_n^{\text{ant}} | g_{\mathbf{z}} \rangle|^2}{\sigma_n^{\text{ant}}} \right)^{-1}. \quad (6.31)$$

For the limited view image scenario we propose to use the modified steering function $g_{\mathbf{z}}^{\dagger}$ as in the beamforming case:

$$I_{FM}^T(\mathbf{z}) = \left(\sum_{n=-\infty}^{\infty} \frac{|\langle u_n | g_{\mathbf{z}}^{\dagger} \rangle|^2}{\lambda_n} \right)^{-1}. \quad (6.32)$$

Using the property $\langle \psi_n | g_{\mathbf{z}} \rangle = \frac{\langle u_n | g_{\mathbf{z}}^{\dagger} \rangle}{\sqrt{2}}$ as before, and $\sigma_n = \lambda_n$ we find that

$$I_{FM\alpha=1}^T(\mathbf{z}) = \left(\sum_{n=-\infty}^{\infty} \frac{2 |\langle \psi_n^{\text{sym}} | g_{\mathbf{z}} \rangle|^2}{\sigma_n} \right)^{-1} \quad (6.33)$$

$$I_{FM\alpha=-1}^T(\mathbf{z}) = \left(\sum_{n=-\infty}^{\infty} \frac{2 |\langle \psi_n^{\text{ant}} | g_{\mathbf{z}} \rangle|^2}{\sigma_n} \right)^{-1} \quad (6.34)$$

which implies that




Figure 6.6: Schematic of the numerical examples with a circular array. The dotted red squares indicate the region that is imaged in the following figures. \mathbb{A} , \mathbb{B} and \mathbb{C} correspond to the situations described in the text (\mathbb{A} and \mathbb{C} are shown at a different scale to \mathbb{B}). The measurements for \mathbb{A} and \mathbb{B} correspond to the F and T matrices respectively.

$$\frac{1}{I_{FM}^F(\mathbf{z})} = \frac{1}{2I_{FM_{\alpha=1}}^T(\mathbf{z})} + \frac{1}{2I_{FM_{\alpha=-1}}^T(\mathbf{z})}. \quad (6.35)$$

This result is very similar to eq. (6.29) for beamforming.

6.3.3 Numerical Simulations

The SVD representations of BF and FM allow us to study the effect of the mirror on the quality of the reconstruction. For this purpose we have generated forward scattering data for the three scattering situations in figure 6.6:

- \mathbb{A} . The full-view situation from figure 6.3b with the scatterer mirrored and a 360° array with 360 point sources at a distance of 100λ . This is the ideal imaging scenario.
- \mathbb{B} . The limited-view situation from figure 6.3a with a 180° array with 180 point sources at a distance of 100λ , and a sound-soft interface at $y = 0$.
- \mathbb{C} . The same as \mathbb{B} , but with the mirror interface removed. This situation is representative of current practice in NDT where signals received after the backwall reflection are ignored.

In each case the scatterer used was the same sound-soft kite shape we have used before. It was introduced by Potthast[96] and is parametrized by




Figure 6.7: The phase and normalised magnitude of the multistatic matrices for the three simulated situations show in figure 6.6 and described in the text. θ and ϕ refer to the illumination and detection angles respectively according to the convention given in figure 6.3a. It can be seen that the structure of the F matrix (A) matches the block structure shown in figure 6.5, and that all matrices are symmetric as expected.

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} -\frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 6 \end{bmatrix} \begin{bmatrix} 1.5 \sin(t) \\ \cos(t) + 0.65 \cos(2t) - 0.65 \\ 1 \end{bmatrix}, \quad (6.36)$$

where $t = [0, 2\pi]$ and the dimensions of x and y are in units of λ . We solved the Helmholtz equation in the frequency domain using the Boundary Element Method [97]. This method is particularly suited to our purposes since it can calculate multistatic matrices for a single frequency very quickly. Additionally it does not suffer from problems associated with meshing methods (FEM and FDM) such as mesh noise, interpolation between nodes, and the problems of representing curved surfaces on a regular grid.

The magnitude and phase of the resulting multistatic matrices for each simulation are shown in figure 6.7. It can be confirmed that the structure of the F matrix (the matrix for situation A) matches the block structures given in figure 6.5. It is

also interesting to note that the regions with the highest scattering magnitude are the pitch-catch measurements. This is because the receiver in those measurements lies in the shadow of the scatterer; the scattered field is approximately equal to the negative of the incident field, producing a small total field in the region of the shadow. Moreover \mathbb{B} shows a richer structure than \mathbb{C} that varies smoothly with θ and ϕ . This suggests that \mathbb{B} contains more information than \mathbb{C} .

The SVD of the three multistatic matrices was performed. As expected it was found that the singular values of F (case \mathbb{A}) are the combination of those for $T_{\alpha=-1}$ (case \mathbb{B}) and $T_{\alpha=1}$. These singular values, and the singular values for case \mathbb{C} are shown in figure 6.8. From the figure it can also be seen that the singular values of \mathbb{B} are larger than those of \mathbb{C} thanks to the presence of the mirror. The number of large singular values correlates with the amount of information we can retrieve (see chapter 3), and is affected by the size of the aperture, as explained later in this section, and by the noise level as discussed in section 6.3.4.

Imaging was performed using both the FM and BF methods and the three simulated multistatic matrices as input. For \mathbb{B} , the modified Green's function of eq. (6.22) was used. The results of these imaging processes are shown in figure 6.9. The images in the figure correspond to the marked regions in figure 6.6.

For the full-view configuration (\mathbb{A}), the BF and FM images are both good reconstructions, with FM providing higher resolution and fewer artefacts. The results for the limited-view case with no mirror (\mathbb{C}) are much worse; both BF and FM provide incomplete images, with only the top edge of the scatter visible. When the mirror is added to the limited-view case (\mathbb{B}) the FM image is restored so that it is nearly identical to the full-view FM image. The BF image is similarly restored, although the level of ringing artefacts becomes significantly higher. In all cases there are severe ringing artefacts in the BF images. These may be improved by the use of time-domain beamforming, but will not achieve the clarity of the FM images. From these results it is clear that in this situation the FM produces much better reconstructions than BF.




Figure 6.8: An example of the singular values for full-view, limited-view with a sound-soft mirror and limited-view with a sound-hard mirror configurations. The singular values for \mathbb{A} (solid black line) are equal to the combination of the singular values of the limited-view case with a sound-hard and sound-soft mirror (\mathbb{B}), which give rise singular values corresponding to the symmetric (red crosses) and anti-symmetric (blue circles) singular functions. The dashed black line is the singular values for \mathbb{C} , but with each element repeated to allow comparison with \mathbb{B} (the blue circles). It can be seen that the addition of the mirror increases the magnitude of the singular values. This corresponds with the increase in the available scattering information.




Figure 6.9: The results of imaging the kite void using the multistatic matrices in figure 6.7. The first row shows these results of imaging the region marked in figure 6.6 with the Beamforming algorithm, and the second row with the Factorisation Method. For imaging of the T matrix (b and e) the modified Green's function [eq. (6.22)] was used. Each image is normalised with respect to its maximum, and a linear scale is used. The size of the regions shown is $10 \times 10\lambda$.

This improvement is even more evident when apertures smaller than 180° are considered. Indeed, in reality it is often impossible to achieve a 180° aperture. The transducers must usually be placed in a line, making it impossible to excite and receive waves that travel horizontally. Even if a very long array is used, the transducers may be directional, which reduces the effective aperture. A more realistic set-up with a linear array is shown in figure 6.10. The array contains 90 point sources evenly distributed over a distance of 100λ , at a distance of 100λ from the mirror, i.e. from $(-100\lambda, 100\lambda)$ to $(100\lambda, 100\lambda)$. This provides an aperture of approximately 90° . The same three situations as before were simulated but with the linear array. The results are shown in figure 6.11.

One can immediately see that the BF results are unacceptably degraded. It is impossible to determine the shape of the scatterer from even configuration **A**. In contrast, the FM results are largely unchanged from the previous results with the




Figure 6.10: Schematic of the numerical example with a linear array. The array has 90 elements, a length of 100λ and is positioned 100λ from the backwall. The array has an aperture of 90° .

circular array. The kite is fully reconstructed with no artefacts after both full-view imaging and limited-view imaging with the mirror. The main reason why FM outperforms BF is that it uses all of the available data. When waves interact with fine details of the scatterer or areas that lie in shadow the reflections they produce are weak. These weak reflections cause correspondingly small variations in the measured multistatic matrix. When the matrix is decomposed by the SVD, small variations become encoded in the singular vectors associated with singular values that are small in magnitude. The functional for beamforming [eq. (6.21)] shows that this data is largely ignored. Each term in the functional is proportional to the singular value, hence the terms associated with the small singular values will have little effect on the total series and hence the image. On the other hand, the terms in the FM functional [eq. (6.30)] are inversely proportional to the magnitude of the singular value. Therefore the small variations in the multistatic matrix can be used to produce a more detailed and complete image than that of Beamforming. In the limited-view case half of the singular values are lost so it is especially important to use as much of the data as possible.

6.3.4 Noise Analysis

The previous simulations did not consider noise besides numerical errors. This is an important factor because, as was discussed before, FM makes use of all the singular values including those with very small magnitude. These singular values will be




Figure 6.11: Results of imaging the kite void with the linear array for the configurations \mathbb{A} , \mathbb{B} , and \mathbb{C} shown in figure 6.10. With the decreased aperture afforded by the linear array, the BF images are now unusable. The FM images are mostly unchanged from the results for the circular array.

corrupted by the addition of noise and some will be completely lost. The loss of the small singular values and their associated singular functions will cause more problems for FM than BF because BF does not use them—this is the source of its noise robustness.

In order to investigate the combined effect of noise and apertures smaller than 180° on the imaging process, various levels of additive Gaussian noise were added to the T multistatic matrix measured in the presence of a mirror and with various angular apertures. The 180° T matrix is shown in figure 6.7 \mathbb{B} . The noise was given a random phase uniformly distributed in $[-\pi, \pi]$, and normally distributed magnitude. The noise matrix was scaled by an amount relative to the RMS magnitude of the 180° multistatic matrix, so for example 1% noise corresponds to noise with an RMS magnitude of 1% of the RMS magnitude of the T matrix. After adding noise, the different apertures were simulated by cropping the multistatic matrix for the 180° case. Although this cropping operation leaves the absolute level of noise unchanged,

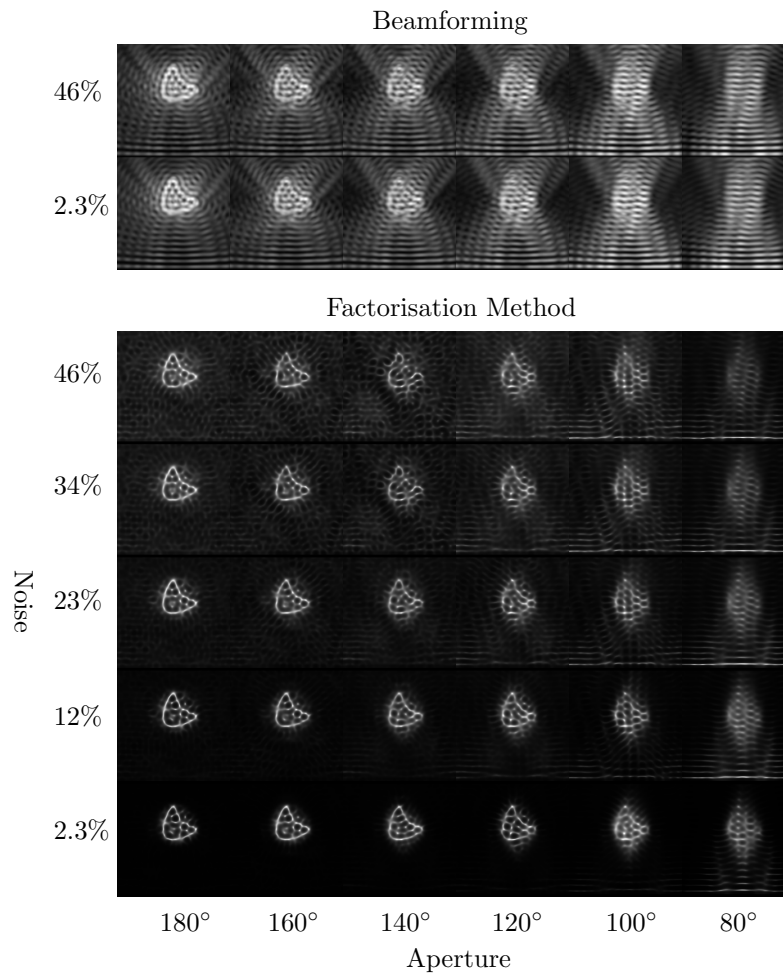


Figure 6.12: *The results of imaging situation \mathbb{B} with various noise levels and apertures. The beamforming results are only shown for two noise levels since the noise has so little effect.*

the signal-to-noise ratio (SNR) increases slightly as a result of the different RMS magnitude of the cropped matrix. The reconstructions using BF and FM for each of these matrices are shown in figure 6.12.

With a full 180° aperture both BF and FM were found to be very robust; even under very large noise levels. However, as the aperture decreases two main trends can be observed. The BF image rapidly degrades as the aperture decreases whilst for a given aperture the reconstruction is not affected by noise. On the other hand, the FM reconstructions are less affected by a reduction in the aperture but are more sensitive to the noise level than BF. These trends can be explained by considering the singular value decomposition of T . Without noise as the aperture decreases, the




Figure 6.13: *The FM results of imaging the kite at a variety of distances from a sound-soft backwall with a linear array. The angle subtended is the dotted angle marked on the left hand diagram. As the subtended angle increases, a larger and larger aperture is required for successful imaging. The noise level is 2.7%.*

number of larger singular values shown in figure 6.8 decreases, thus implying that more information is carried by the smaller singular values. Since BF mainly uses the information carried by the largest singular values [eq. (6.21)] the reconstruction must degrade rapidly as the aperture decreases. Moreover since the larger singular values are not significantly affected by noise this results in the robustness of BF to noise. On the other hand, FM attempts to extract the information associated with the smaller singular values and therefore it leads to better reconstructions than BF even for small apertures. However, as noise increases the small singular values are lost and with them the information they carry leading to increasing image degradation.

We posit that the aperture required for successfully imaging a scatterer using the FM depends not only on the noise, but also on the angle subtended by the scatterer, relative to its mirror image (see diagram in figure 6.13). The logic behind this hypothesis is that as the angle subtended increases, a larger and larger aperture is needed in order to produce waves that interact with sufficient energy with the back of the scatterer. Or in other words, if the aperture remains the same and

the scatterer moves towards the mirror, more and more information is encoded in the small singular values at the expense of the large ones. To test this, the kite scatterer was imaged at a variety of distances from the backwall, for a number of different apertures. Each distance corresponds with a certain subtended angle. The reconstruction with 2.3% noise are shown in figure 6.13. Although there is not a simple relationship between the angle subtended by the scatterer and the aperture required for a good image, the trend is clear: as the scatterer approaches the mirror (i.e. as the subtended angle increase) it is imaged less and less accurately if the aperture is kept constant. From this we can conclude that a large aperture is required to correctly image scatterers very close to the backwall. If they are separated slightly then it is possible to use apertures as small as 100° with 2.3% random noise. In both tests performed in this section aperture size has a large effect on the quality of both FM and BF reconstructions. Any method of increasing the available aperture would certainly improve the results of real-life imaging problems.

6.4 Imaging penetrable scatterers

When imaging impenetrable scatterers, the goal is not to merely find their support, \mathcal{D} , but to find the object function (or sound speed) throughout the scatterer $o(\mathbf{x})$. The image functional used for the sampling methods has no direct correlation with the sound speed, so another method must be used. The canonical method, described in chapter 2 is Diffraction Tomography (DT). However, DT is difficult to formulate for arbitrary incident waves and non-constant background speeds. An alternative method derives from the observation that for a circular aperture, the BF image is a linear function of the DT image. Therefore the DT image can be formed by calculating the BF image and deconvolving the result. This is the method used by the HARBUT algorithm developed by Huthwaite and Simonetti [92].

The motivation for HARBUT is that for many imaging problems, in particular imaging tumours in breasts the background sound speed outside the scatterer we are interested in is not constant. By removing the effects of the varying background sound

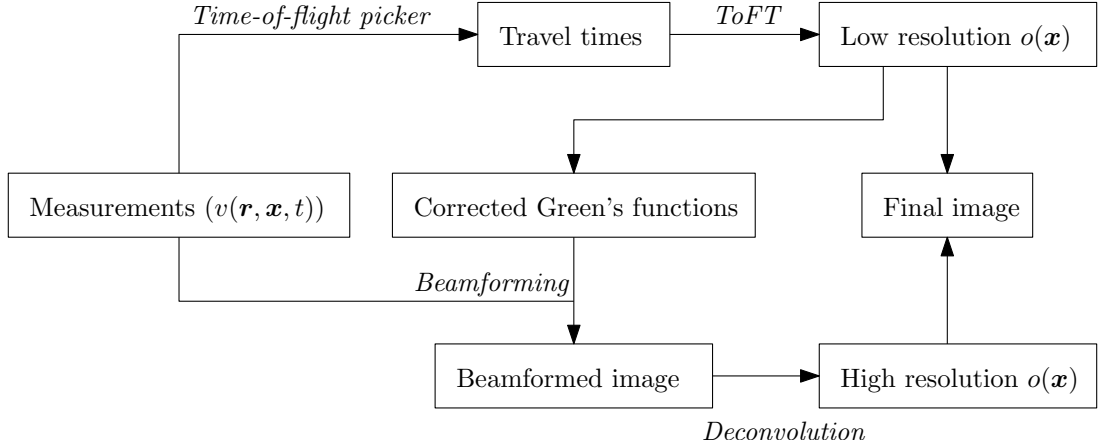


Figure 6.14: *HARBUT* Algorithm. See the main text for a description.

speed we can obtain a clearer image of the scatterer. In HARBUT the background sound speed is measured with low resolution using Time of Flight Tomography (ToFT). ToFT is good at accurately reconstructing sound speeds, but can only do this with very low resolution. In HARBUT a bent-ray ToFT method is used. The resulting sound speed map is then used to calculate the correct, spatially varying Green’s function which takes into account the different travel times in the medium. Finally the modified Green’s function is used to calculate the BF image, which is deconvolved to obtain the DT image which is a high resolution measure of the sound speed in the scatterer.

In the breast application of HARBUT, a full circular array is used, however HARBUT can also be used with a linear array and a mirror surface. When used with a linear array, the mirror surface is necessary to obtain through-transmission travel times (since ToFT requires waves that more or less travel straight through the scatterer). This mode of imaging is of interest in NDT as small changes in sound speed can be used to detect damage precursors, measure residual stresses [98] and material strength, as in the case of concrete [99]. The ability to detect and size these regions depends on both their size, and the magnitude of sound-speed change. The full algorithm is then as follows (also depicted graphically in figure 6.14):

1. Perform Full Matrix Capture (FMC) with the array, i.e. measure $v(\mathbf{r}, \mathbf{x}, t)$.
2. Calculate the travel time for each transmit-receive combination using the

method in [100].

3. Apply the bent-ray ToFT algorithm to produce a low resolution sound-speed map.
4. Calculate the travel time from each transducer to each point in space, this can be done using a simple algorithm based on the application of Fermat's principle. This is used to calculate the frequency domain Green's function for each point.
5. Use the modified Green's function to calculate the single frequency BF image, as described in chapter 2.
6. Deconvolve the resulting image using the frequency response given in [38].
7. Combine the sound speed maps calculated in steps 3 and 6.

In the next chapter we test the capability of ToFT to detect sound-speed changes experimentally.

6.5 Conclusions

We have studied a method of improving the image obtained of an object when we only have a limited view of it. If a planar reflector is situated behind the object then it becomes possible to measure waves scattered from the back of the object and a greater amount of information can be retrieved. However, compared to a related full-view configuration there is still some loss of information. In order to assess this loss, we studied the singular value decomposition of the multistatic matrix and demonstrated that half of the singular values/vectors of the full-view case are retained in the limited-view case; the other half are lost.

Scattering information about the fine details of an object, and parts of the object that lie in shadows tend to be encoded in the small singular values, so this information is lost due to noise, especially when using a small aperture. This was confirmed by

performing imaging on numerical examples using the BF and FM algorithms. BF does not use the small singular values to reconstruct the image, so it is very robust to noise, however this also means that it cannot produce satisfactory images when the aperture is small. On the other hand, FM always uses all of the singular values. This enables it to produce detailed reconstructions with small apertures at low noise levels, which gradually degrade as the noise level is increased. The position of the object relative to the mirror also affects the quality of the reconstruction, particularly at small apertures. If the object is close to the mirror then the reflections from its back are small in magnitude and become lost in the noise.

Limited-view imaging with a mirror interface beneath the object provides dramatically better images than would be obtained without the mirror. This is true of imaging with both BF and FM. Our numerical simulations suggest that in order to obtain an accurate reconstruction under realistic noise levels an aperture of at least 100° is necessary. Experimental results using this technique to image impenetrable scatterers, and also using the briefly described HARBUT algorithm to image penetrable scatterers, are presented in the next chapter.

Chapter 7

Using mirror reflections for imaging—Experiments

7.1 Background

In the previous chapter, we studied the theory of imaging penetrable and impenetrable scatterers using a mirror interface to provide additional scattering information. In this chapter we describe several experimental tests of this method, and discuss practical factors that are not included in the simulations.

Experiments where reflective interfaces are used to augment ultrasonic imaging are rare-to-non-existent. The only previous work that we are aware of is [76], in which mode converted reflections from the backwall are used to detect defects in welds. However, this imaging is performed from a very small aperture, and provides “indications” rather than a true image of defects.

The measurement procedure is the same as in the simulations, and as in standard imaging—Full Matrix Capture is used to record $v(\mathbf{r}, \mathbf{x}, t)$. But there is a significant difference between the processing of simulated and experimental data.

As with the imaging of the aluminium crack in chapter 5 we must remove the incident field from the total field to obtain the scattered field. Some of the experiments in

this chapter are simple enough that this can be done by baseline subtraction of a measured or a simulated incident field. Accurate measurement of the incident field is only possible in experiments in fluids. Otherwise the incident field must be removed by gating, which results in an inaccurate measurement of the scattered field.

A second difference from the simulations is the directionality of the array elements. In the simulations they were assumed to be point sources, however real transducers excite wave-fields with phase and magnitude that varies with angle. We therefore need a method for measuring and compensating for this variation.

These practical backwall imaging aspects are discussed in the next section. We then present three experiments that exploit the backwall. These are: two experiments with impenetrable scatterers in water, and one with penetrable scatterers in water. The “backwall” is simulated using the water surface or a metal block. This work has been published in [80]. The final experiment images penetrable scatterers—specifically tubes of alcohol in water. The results of these experiments are presented and discussed in section 7.4.

7.2 Practical considerations

7.2.1 Incident field removal

As discussed previously, imaging a scatterer ideally requires knowledge of the scattered, rather than total field. Since the total field u is the one that is actually measured we must recover u^s from u . One possibility is to identify the temporal regions in which the incident field is expected to be non-zero, and gate out those regions by setting u to be zero there.

7.2.1.1 Gating

Gating out the incident field in the time domain was used with the elastic scatterer in chapter 5. There, due to the complexity of the multiple elastic reflections and mode conversions, the gated locations were found by measuring the incident field of a real defect-free component experimentally. In water we can simply calculate the expected arrival time of the reflected waves by geometry, as shown in figure 7.1.

The main disadvantage of this method is that any components of the scattered field that overlap the incident field will be lost. Since there will always be such components—e.g. the shadows—we cannot recover the scattered field in its entirety. In particular the scattered field in the frequency domain will never be correct because it depends on the values of u^s at all times, and times where the shadow is expected will be incorrectly zeroed.

Since $u^s = u - u^i$ we can see that in order to know u^s everywhere we must know the actual values of u^i , irrespective of their domain of representation. This can be achieved in experiments in fluids by carefully removing the scatterer and repeating the FMC measurement. However in most practical applications we cannot remove the scatterer and we must use another approach, such as the one we describe next.

7.2.1.2 Field-fitting

To obtain the incident field when the scatterer cannot be removed, we can measure it on a defect-free copy of the component. However, subtracting the resulting waves from the measured waves of another component does not work in practice.

The reason for this is illustrated in figure 7.2a. Due to the large maximum gradient of a sine wave relative to its maximum magnitude (1:1), a very small error in the relative phase of two nominally equal waves can lead to a large residual error when they are subtracted. Changes in magnitude have a similar but smaller effect. Phase errors can arise from small differences in coupling, dimensions, sound speeds, and electronic timing. We found that even taking two measurements on different positions

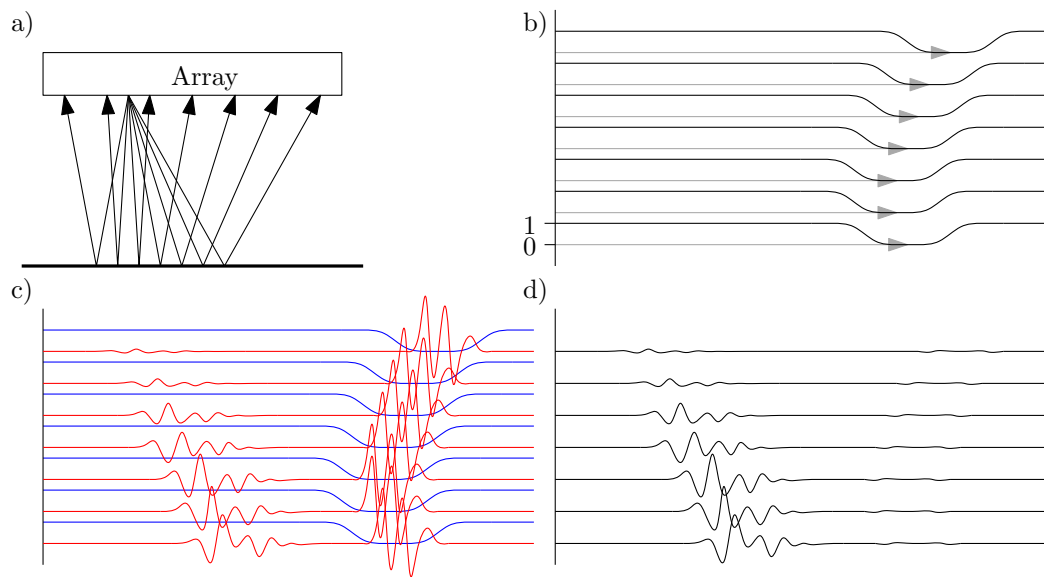


Figure 7.1: A simple method of gating out the incident field. a) The arrival time of the backwall reflection is calculated from geometry. b) This is turned into a smooth mask. c) The mask is multiplied by the total field to obtain d) an estimate of the incident field.

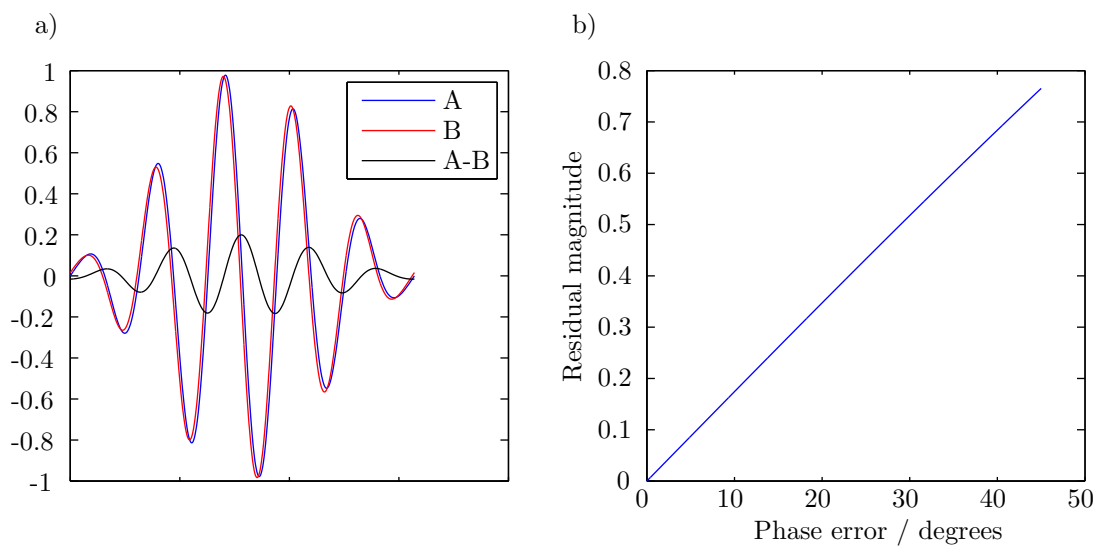


Figure 7.2: Very small phase errors make baseline subtraction impractical. a) The total field, and incident field measured with a small timing error, where the scattered field should be zero. When subtracted there is significant residual error even for a small timing error. b) The magnitude of the residual error vs the timing/phase error.

along an extruded aluminium block leads to waves that do not fully cancel.

However, for the simple case of single-frequency imaging of an object above a metal block in water, we found that it is possible to avoid the need for measuring the incident field by fitting a simulated incident field to the total field. In this process it is important to accurately locate the position of the backwall. To simplify matters we assume the backwall is straight, then we only have two parameters to find: the distance to the array and its angle relative to the array, as shown in figure 7.3a. We must also assume that the scatterer is reasonably small compared to the array. If this is not the case then there will be few reflections from the backwall and although we will not be able to locate it, there will be little benefit in using this imaging method anyway.

In our method, an initial estimate of the backwall location is made using the pulse-echo responses at the ends of the array. It is assumed that there are no scatterers to obstruct these signals, and that the strongest signals will be the reflections from the backwall. By timing these reflections the distance to the backwall at each end of the array can be found and used to find an initial estimate of the backwall location. This initial estimate is then refined using the following iterative method.

First we produce a monochromatic simulation of the FMC without the scatterer for the current backwall position. This results in a matrix R . The magnitude of each transmit–receive pair is determined by the directionality of the elements (see the next section), and the transmission distance. The magnitude is not sensitive to the backwall position. The phase is only determined by the transmission distance and is highly sensitive to the backwall position.

The mismatch between the measured total field matrix T and the simulated reference field matrix R is calculated using the following error metric (R and T are both complex matrices)

$$e = \sum_i \sum_j H \left(|R_{ij} - T_{ij}| - \frac{1}{2} \|R - T\|_{\max} \right) \quad (7.1)$$

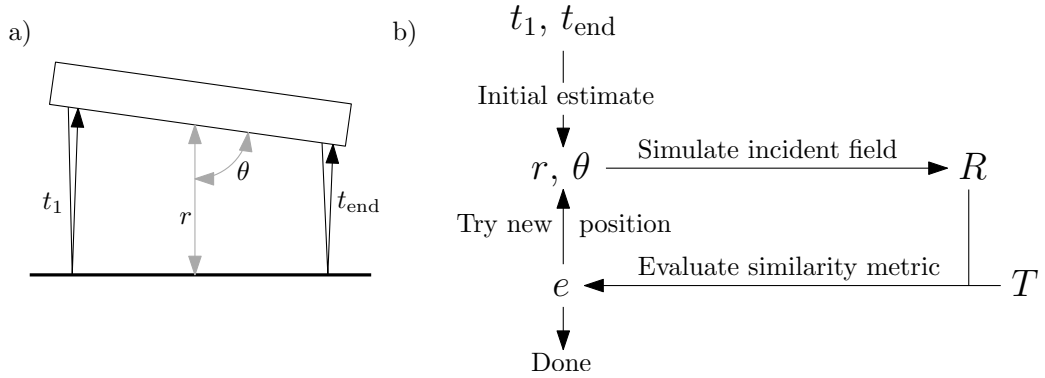


Figure 7.3: *The incident field fitting algorithm. a) The parameters to be fit (r and θ); b) The algorithm (see main text for a complete description).*

where $H(\cdot)$ is Heaviside step function, i and j are the matrix indices (which correspond to the array elements) and $\|R - T\|_{\max}$ is the maximum element of $|R - T|$. In other words the metric counts the number of matrix coefficients that differ by more than half the maximum difference. The two position parameters of the backwall (distance and angle) are iteratively adjusted to minimize this metric using the downhill simplex method built into Matlab (but any optimization method could be used).

The results for one experiment in which a cylindrical scatterer was used are shown in figure 7.4. The measured T matrix is shown (figure 7.4a) along with the predicted R matrix (figure 7.4b). For comparison the measured matrix obtained by removing the scatterer is shown in figure 7.4c. They match very closely with a discrepancy of a fraction of a wavelength.

7.2.2 Beam directionality

In our simulations and theoretical studies, point source (equivalent to plane wave sources in the far field) are used to illuminate the scatterer. In reality, the PZT elements used in ultrasonic arrays are approximately strip sources. The extended nature of the source results in a directional field whose strength depends on the observation angle relative to the normal to the face of the element. This dependence approximates a sinc function [10]. However, the exact shape is distorted significantly by mechanical and electrical coupling of adjacent array elements, and the effects of




Figure 7.4: *Experimental incident field fitting results showing the phase of a) the measured T matrix, b) the simulated R matrix, and c) the experimentally measured R matrix found by removing the scatterer. The matrices in (b) and (c) should match.*

matching / coupling layers. Therefore it must be measured. We developed a rapid and simple procedure to do this which proceeds as follows.

FMC is performed near a planar reflector with no scatterers and the amplitude of each echo from the backwall is calculated. The field excited by a single transmitter and received by all array elements can be thought of as the field produced by a virtual transmitter at the mirror position relative to the interface plus an appropriate phase shift as shown in figure 7.5a. The received signal magnitudes can therefore be related to the directivity function.

The amplitude of each signal is a function of the distance travelled by the wave and the directivity function for transmission and reception (these are the same by reciprocity). This amplitude is given by

$$a = \frac{D(\theta)^2}{\sqrt{r}} \quad (7.2)$$

where $D(\theta)$ is the directivity function, θ is the angle of reception and transmission and r is the distance between the transmitter and receiver (see figure 7.5a). Since r and θ are known geometric parameters, $D(\theta)$ can be found.

As the transmitter element is changed, the same directivity function should be

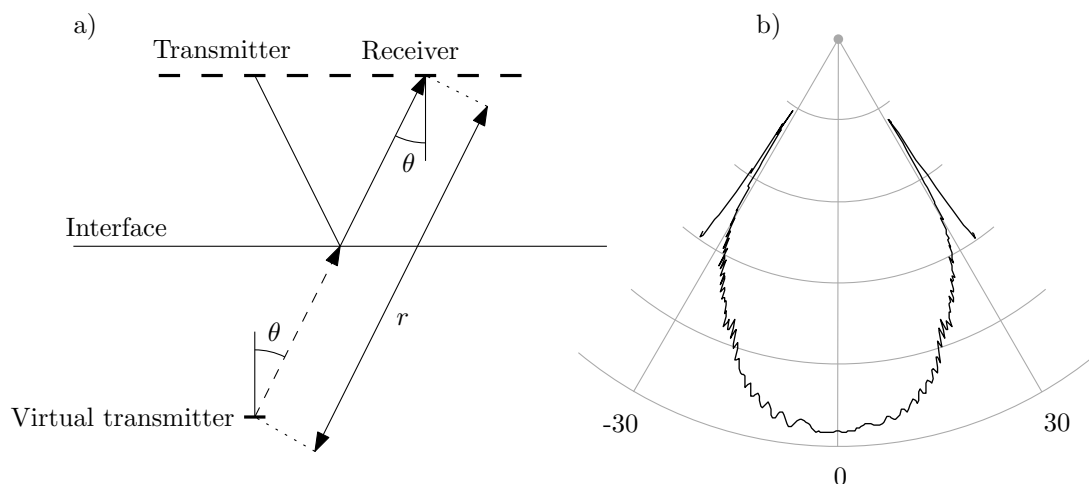


Figure 7.5: a) Schematic of beam directionality measurement. b) Measured directivity for a 0.3 mm element in water at 4.5 MHz ($\lambda = 0.33$ mm).

obtained. This fact can be used to reduce noise by averaging the measurements for all transmit-receive pairs at each angle. The directivity function was measured at a frequency of 4.5 MHz for a medical ultrasound array and is plotted in figure 7.5b. The element pitch was 0.3 mm and the wavelength was 0.33 mm. Ideally the elements of our array would be omnidirectional as this gives the greatest field of view. As the element width increases and the field becomes more directional fewer reflections are received from the scatterer until only the transducer elements that are directly above the scatterer receive echoes.

If we consider the form of LSM and FM it is clear that if the array elements are highly directional, the equations will be impossible to solve because some of the elements will receive no reflections because they are at an angle to the scatterer which has a very small directionality coefficient. To solve this we multiply the Green's function by appropriate directionality coefficients, corresponding to the measured directivity function, when evaluating it at each transducer element. This provides a small but noticeable improvement to imaging with directional elements. However it does not solve the underlying problem—that highly directional elements effectively cause a smaller aperture which is very detrimental to the high resolution imaging methods.

7.3 Impenetrable scatterer experiments

Our initial experiment was performed in water using a metal block as the “backwall” and a small aperture high frequency array. This array has 128 elements, a pitch of 0.3 mm and a centre frequency of approximately 5 MHz. Two scatterers were used: a razor blade, and an air-filled drinking straw.

Because in this case the scatterer is an object placed in water, it is actually possible to simply remove the object without disturbing the array in order to measure the incident field. This would not be possible in an industrial application. Imaging for this experiment was performed using the single frequency factorisation method, and the time domain beamforming method.

FMC data was recorded for scattering from the straw and the razor blade. The backwall location was found and baseline subtraction performed. The resulting data was imaged using traditional beamforming and the modified factorization method. Imaging was performed at ten frequencies with a bandwidth of 2 MHz and the resulting images were summed to give the final results shown in figures 7.6 and 7.7. Figures 7.6a and 7.6b show the data from the straw imaged using the unmodified beamforming algorithm and with the modified mirror factorization method respectively. The actual location of the straw is marked with a circle. Although the factorization method does not reconstruct the shape of the straw entirely correctly, its extent is closer to the desired result than beamforming achieves.

The factorization method images for the razor blade in two different positions are shown in figure 7.7a and 7.7b. In these cases the imaging entirely depends on data received after the backwall echoes because no reflections are recorded before them. Even if tip diffraction echoes were detectable they could only be used to see the tip of the blade. This could not be distinguished from a point scatterer. However, using the backwall reflections the full extent of the blade is visible. The length and shape of the reconstruction accurately match the scatterer.

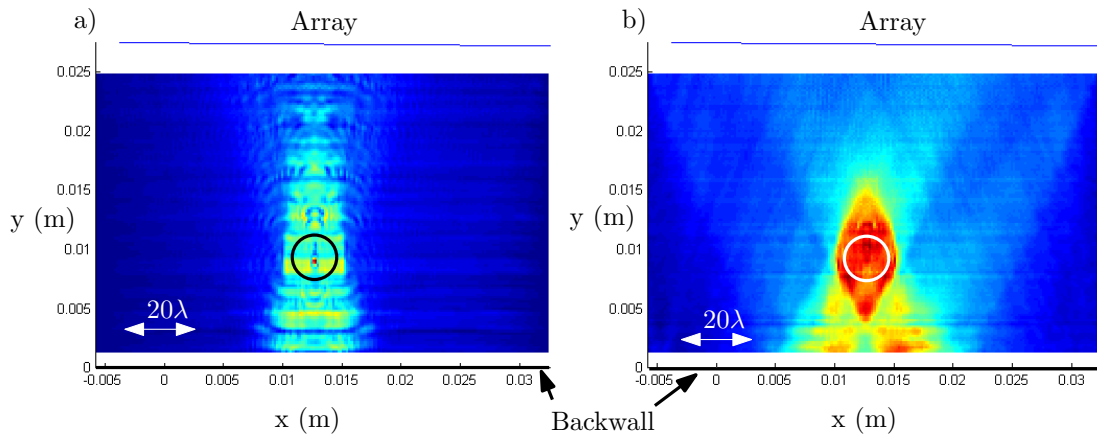


Figure 7.6: The results for backwall imaging of a 4 mm diameter air-filled straw in water using a) beamforming and b) FM. The size and position of the straw are marked.

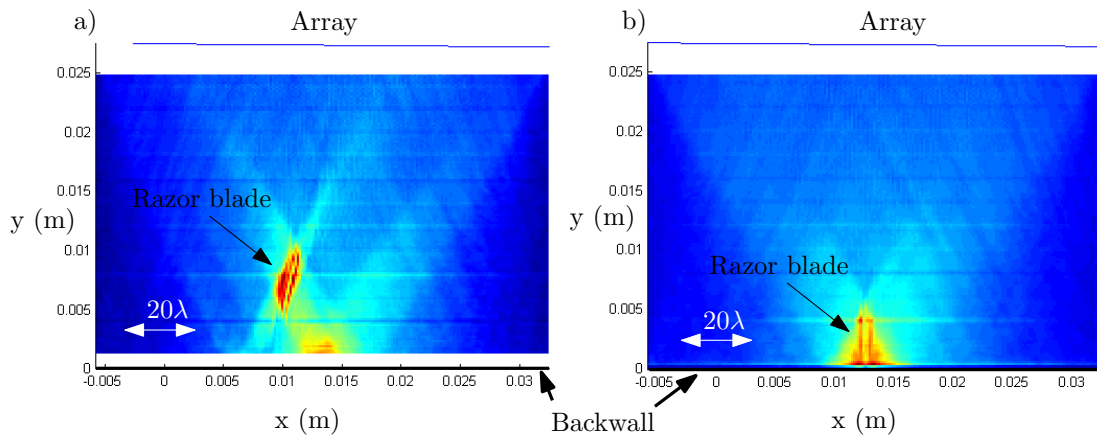


Figure 7.7: The results for backwall imaging of a 5 mm razor blade suspended in two different locations in water using FM.

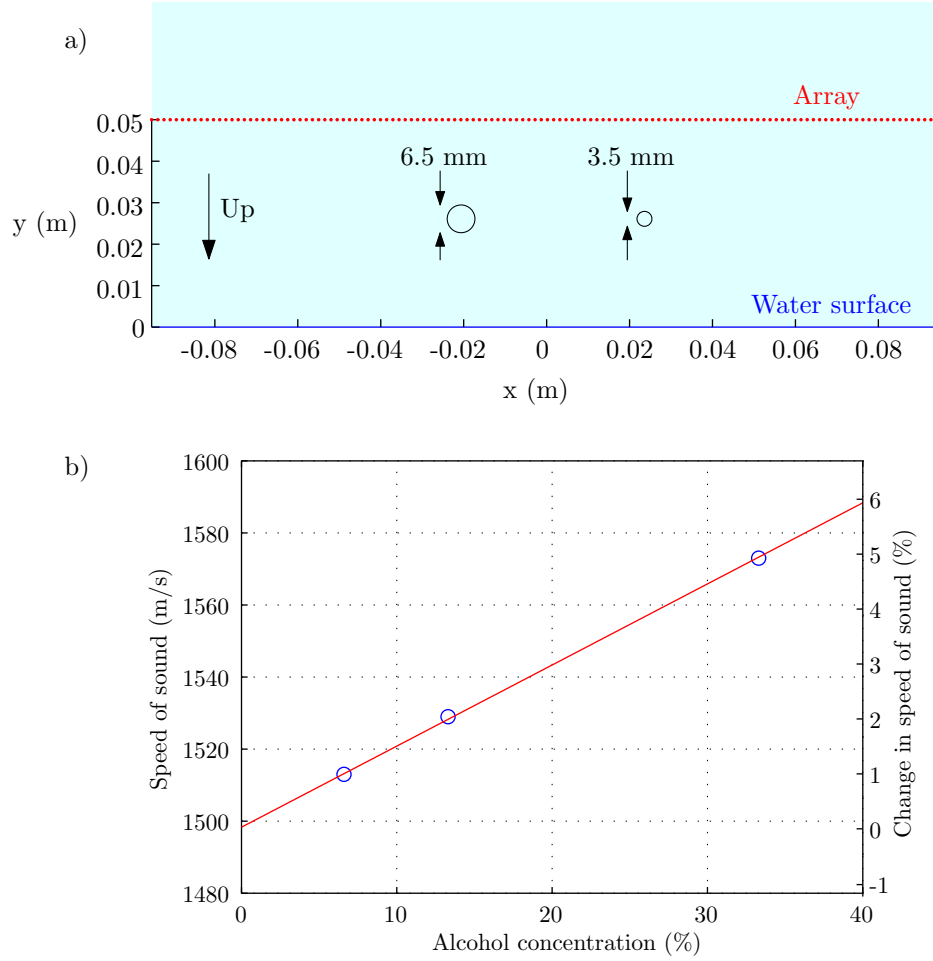


Figure 7.8: a) Layout of the tubes of alcohol in the experiment showing the array (crosses) and scatterers. The water surface is at $y = 0$ (the figure is upside-down compared to the physical experiment, for consistency with previous figures). b) Percentages of alcohol used in the experiment, and the resulting speeds of sound.

7.4 Penetrable Scatterer experiments

To determine the effectiveness of the HARBUT algorithm for imaging penetrable scatterers above a backwall, we performed experiments using thin plastic tubes (heat shrink) containing diluted alcohol (specifically Tesco own-brand vodka, which is much easier to obtain than denatured ethanol). A diagram of the experimental configuration is shown in figure 7.8a. Three percentages of alcohol were tested, corresponding to approximately 1%, 2% and 5% changes in sound speed. The percentages of alcohol and measured speed of sounds are shown in figure 7.8b. They agree with the speed–concentration relationship given in [101] at 25°C.

The times of flight of each dataset, and also the incident field dataset with no scatterers present, were extracted using a threshold-based time-of-flight picker. These times, and the full waves were used to generate ToFT images, DT images, and the combined HARBUT images for each dataset as described in the previous chapter. The datasets were processed as described in the previous chapter, resulting in the images shown in figure 7.9, for 1%, 2% and 5% speed-of-sound changes.

From the figure one can see that DT provides much higher resolution information than ToFT, but mainly reconstructs the edges of the scatterer rather than the speed-of-sound within it (e.g. compare the ToFT and DT images of the larger scatterer in figure 7.9c). When combined, the HARBUT images retain the higher resolution of the DT image, but also reconstruct the sound-speed inside the tubes correctly.

7.5 Conclusions

We have experimentally investigated the use of backwall reflections to improve the reconstructions of strongly scattering defects in thick components. Several issues encountered in the experimental procedure have been discussed; these are element directionality, localization of the backwall, and baseline subtraction. Element directionality was addressed by modifying the Green's function used in the imaging algorithms. The position of the backwall was accurately found by iteratively adjusting its position until the measured wave fields matched simulated ones. Model experiments were performed in water and have shown that this imaging approach could lead to a substantial improvement over standard beamforming by improving the reconstruction of crack-like and extended defects.

Additionally we demonstrated experimentally the HARBUT algorithm for imaging penetrable scatterers above a mirror interface. It was able to produce quantitative measures of the speed of sound in the scatterer and detected a 1% change in speed of sound over 6.5 mm (11 wavelengths). Both techniques are promising for imaging defects in practical elastic configurations. However, in order to do this the algorithms will need to be further developed in order to take into account the effects of mode

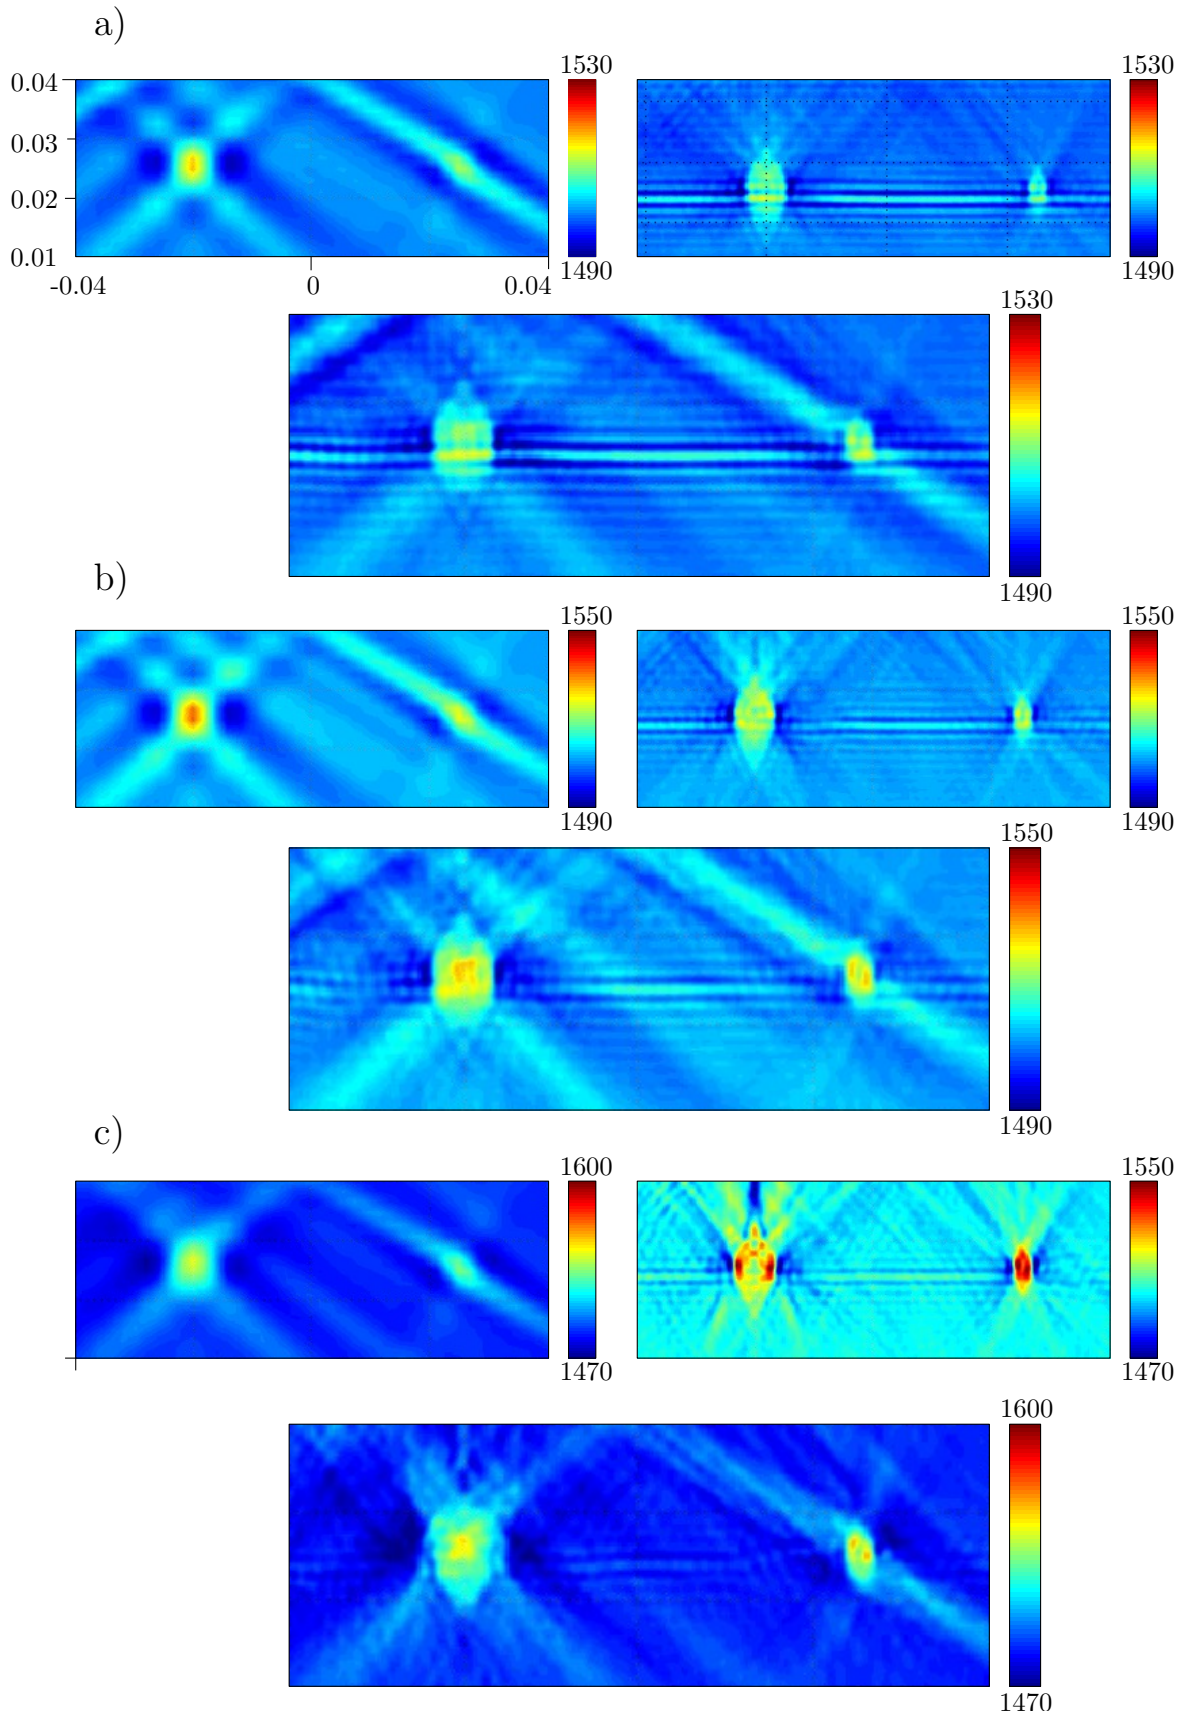


Figure 7.9: *ToFT*, *DT* and *HARBUT* images for a) 1%, b) 2%, and c) 5% changes in the speed of sound.

conversion at the backwall.

Chapter 8

Conclusions

The aim of this research was to provide higher quality images of scatterers when imaged using ultrasound arrays. Although we have concentrated on ultrasound, many of the results are applicable to other wave-imaging fields such as seismic imaging used in geophysics. The current state-of-the-art imaging algorithms are linear algorithms based on beamforming or diffraction tomography. We looked at two sources of usually-discarded data to improve the images: multiply scattered signals, and reflections from known planar interfaces such as component boundaries.

Non-linear algorithms were used to decode the multiply scattered signals, leading to super-resolution, and the reflections from the component backwalls were used to provide more complete images of impenetrable scatterers, and quantitative measurements of the sound speed in penetrable scatterers. Both methods have been proposed in the literature, but neither have been made experimentally practical before.

Our results can be divided into theoretical conclusions on the informational limits of signals, and the conditions necessary for high resolution imaging to work; and the practical progress we have made in terms of implementing the methods experimentally. These are briefly described first and then we provide some considerations on future research directions, guided by our experience gained during experimental testing.

8.1 Theoretical conclusions

Prior to this research we knew little about the necessary conditions for the sampling methods to work, or the limits to their performance. Simulations, and the mathematical results of others, showed that infinite resolution was possible, but in practice noise prevents this. In chapter 3 we investigated the distribution of information in the scattering measurements. Although it is not possible to give tight bounds on the information capacity of a set of measurements we found some upper bounds on the information content which provide insight into what is worth measuring, and the achievable image quality. This was achieved by decomposing the data into different domains (temporal, spatial and wavelet).

Importantly, we have shown that broadband signals and the backwall reflections can contain data that can be exploited to improve imaging.

In chapter 5 we developed a mathematical framework for imaging a scatterer using the reflections from a mirror interface behind it. We found that half of the singular values of the full view multistatic matrix (far-field operator) are accessible in the limited view case with the mirror. Simulations showed that the mirror imaging method is able to fully reconstruct all sides of an impenetrable scatterer with only a limited view. Crucially we found that in the presence of noise, a large aperture array is critical to achieving a good reconstruction, especially when imaging with the sampling methods. These results have led to a paper published in the *Journal of Applied Physics* [79].

We implemented the time domain LSM on a parallel computing cluster, and developed a pre-processing beamforming step in order to allow its use with wide-aperture arrays with many elements.

We also showed through numerical simulation that the super-resolution sampling methods originally developed for acoustic media can achieve super-resolution in elastic materials without needing to gate out signals resulting from mode conversion at the scatterer. Where both shear and longitudinal waves are present, imaging can

be performed with either shear and longitudinal wave speeds even if they cannot be measured independently. The choice of the best mode to use depends on the scatterer being imaged. In our simulations, shear waves worked best for vertical cracks, and longitudinal waves worked best for horizontal cracks.

In addition to the theoretical work, we performed many experiments and made several advancements relating to the the practical implementation of these techniques.

8.2 Practical progress

In terms of practical progress we have developed some useful techniques and demonstrated some successful results. We experimentally demonstrated the super-resolving capabilities of time domain LSM method for the first time by imaging two point-like scatterers spaced more closely than the Rayleigh resolution limit (approximately half the separation). These scatterers—hairs—were suspended in glycerol, which is highly attenuative and prevents the use of higher frequencies. This work has been submitted for publication to IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control [69].

These experiments have shown that the use of multiple frequencies provides a greater number of independent measurements of the scatterer, which greatly increases the stability of the image, with no loss of resolution compared to single frequency methods.

Time domain LSM was also used experimentally to image cracks in aluminium. Although it is a vast improvement over other methods, like BF it is unable to accurately size the cracks. This is probably due to the low SNR of the tip reflections in comparison to the corner reflection from the root of the crack, and the fact that the free space Green's function was used, rather than one that considers backwall reflections.

For mirror imaging, several acoustic experiments were performed in water using a metal block or the water surface as the mirror interface. While BF can only be used

to detect the location of a scatterer, we have demonstrated that by using the backwall it is possible to extract information about the shape of a scatterer which is important for sizing. Penetrable scatterers consisting of thin plastic tubes containing diluted alcohol were imaged using the HARBUT algorithm which was able to reconstruct a 1% change in sound speed over 6.5 mm. Some of these results have been published in [80].

We developed a reliable technique for automatically compensating for differences in coupling between array elements, and a method for rapid measurements of the directivity of elements. We also developed a simple method for gating out the incident fields, and a more advanced method for subtracting it from the total field.

These results and techniques show potential for industrial implementations that can approach the limit of what is possible, by exploiting all sources of scattering information.

8.3 Future direction

In terms of future research, the first step is to perform more extensive testing in metals with a wider variety of defects. Ideally this would include complex defects in free space, where the effects of the backwall reflections can be separated. Before commercialisation it would be necessary to investigate the effects of 3D defects, and whether a full 2D array is required to image them.

In our experiments in metal, we did not consider the reflections from the backwall, which are more complex than in acoustic media. This meant that our physical scattering model did not match the experiment. Fixing this requires adding reflected waves (including mode converted waves) to the Green's function used for imaging and account for them properly. This should lead to improved image quality with fewer artefacts when imaging near a backwall.

Finally we should continue to explore accurate algorithms that make the best use of broadband scattering information to achieve high resolution and robustness to

noise. The Time Domain Linear Sampling Method is a first attempt at such an algorithm, but is too slow (on current hardware) to be commercially useful in many circumstances. Future algorithms may be able to achieve comparable images but run at the speed of single-frequency methods.

Bibliography

- [1] W. C. Lang and K. Forinash, “Time-frequency analysis with the continuous wavelet transform,” *American Journal of Physics*, vol. 66, no. 9, pp. 794–797, 1998.
- [2] D. Colton and A. Kirsch, “A simple method for solving inverse scattering problems in the resonance region,” *Inverse Problems*, vol. 12, pp. 383–393, 1996.
- [3] H. K. Reddick, “Safe-life and damage-tolerant design approaches for helicopter structures,” tech. rep., 1983.
- [4] P. N. T. Wells, “Ultrasonic imaging of the human body,” *Reports on Progress in Physics*, vol. 62, pp. 671–722, 1999.
- [5] D. L. McDowell, “Basic issues in the mechanics of high cycle metal fatigue,” *International Journal of Fracture*, vol. 80, pp. 103–145, Apr. 1989.
- [6] C. H. Sherman and J. L. Butler, *Transducers and arrays for underwater sound*. Springer, 2007.
- [7] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, “Post-processing of the full matrix of ultrasonic transmit-receive array data for non-destructive evaluation,” *NDT & E International*, vol. 38, pp. 701–711, 2005.
- [8] B. D. Van Veen and K. M. Buckley, “Beamforming: A Versatile Approach to Spatial Filtering,” *IEEE ASSP Magazine*, no. April, 1988.

- [9] C. J. R. Sheppard and A. Choudhury, "Image formation in the scanning microscope," *Optica Acta*, vol. 24, no. 10, pp. 1051–1073, 1977.
- [10] B. Drinkwater and P. Wilcox, "Ultrasonic arrays for non-Destructive evaluation: a review," *NDT & E International*, vol. 39, pp. 525–541, Oct. 2006.
- [11] A. Kirsch, R. Kress, P. Monk, and A. Zinn, "Two methods for solving the inverse acoustic scattering problem," *Inverse Problems*, vol. 4, pp. 749–770, 1988.
- [12] D. Colton, "Inverse Acoustic and Electromagnetic Scattering Theory," *Inverse Problems*, vol. 47, pp. 67–110, 2003.
- [13] K. Hashiguchi, *Elastoplasticity Theory*. New York: Springer, 1st ed., 2009.
- [14] Wolfram, "Navier's Equation."
- [15] W. Jianguo and F. Ruyu, "Time-domain volume integral equation for transient scattering from inhomogeneous objects—2D TE case," *Journal of Electronics*, vol. 18, no. 4, pp. 366–369, 2001.
- [16] P. Hähner, "Scattering by Media," in *Scattering* (E. R. Pike and P. C. Sabatier, eds.), pp. 74–94, Academic Press, 2002.
- [17] D. Komatitsch and R. Martin, "An unsplit convolutional perfectly matched layer improved at grazing incidence for the seismic wave equation," *Geophysics*, vol. 72, no. 5, p. SM155, 2007.
- [18] W. S. Hall, *The Boundary Element Method*. Kluwer, 1994.
- [19] D. Placko and T. Kundu, *DPSM for modeling engineering problems*. Wiley, 2007.
- [20] F. Simonetti, "Multiple scattering: the key to unravel the subwavelength world from the far-field pattern of a scattered wave," *Physical Review E*, vol. 73, pp. 1–13, Mar. 2006.

- [21] A. Zinn, "On an optimisation method for the full- and the limited-aperture problem in inverse acoustic scattering for," *Inverse Problems*, vol. 5, pp. 239–253, 1989.
- [22] D. Colton and B. D. Sleeman, "Uniqueness Theorems for the Inverse Problem of Acoustic Scattering," *IMA Journal of Applied Mathematics*, vol. 31, pp. 253–259, 1983.
- [23] a. Abubakar, P. VandenBerg, and S. Semenov, "A Robust Iterative Method for Born Inversion," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, pp. 342–354, Feb. 2004.
- [24] P. Mora, "Inversion = migration + tomography," *Geophysics*, vol. 54, no. 12, pp. 1575–1586, 1989.
- [25] J. F. Greenleaf, "Computerized Tomography with Ultrasound," *Proceedings of the IEEE*, vol. 71, no. 3, pp. 330–337, 1983.
- [26] M. Brandfass, A. D. Lanterman, and K. F. Warnick, "A comparison of the Colton-Kirsch inverse scattering methods with linearized tomographic inverse scattering," *Computer Engineering*, vol. 17, pp. 1797–1816, 2001.
- [27] K. Iniewski, *Medical imaging: principles, detectors, and electronics*. Wiley-Interscience, 2009.
- [28] A. J. Devaney, "A Fast Filtered Backpropagation Algorithm for Ultrasound Tomography," *Control*, vol. 34, no. 3, 1987.
- [29] A. J. Hunter, B. W. Drinkwater, and P. D. Wilcox, "The Wavenumber Algorithm for Full-Matrix Imaging Using an Ultrasonic Array," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 55, no. 11, pp. 2450–2462, 2008.
- [30] J. Li, X. Wang, and T. Wang, "On the validity of Born approximation," *Progress in Electromagnetics Research*, vol. 107, pp. 219–237, 2010.
- [31] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. SIAM, 2001.

- [32] K. Shariff and A. Wray, "Analysis of the radar reflectivity of aircraft vortex wakes," *Journal of Fluid Mechanics*, vol. 463, pp. 121–161, July 2002.
- [33] K. E. Thomenius, "Evolution of Ultrasonic Beamforming," in *1996 IEEE Ultrasonics Symposium*, pp. 1615–1622, 1996.
- [34] C. J. Nolan and M. Cheney, "Synthetic aperture inversion," *Inverse Problems*, vol. 18, pp. 221–235, 2002.
- [35] R. Arnott, S. Ponnekanti, C. Taylor, and H. Chaloupka, "Advanced base station technology," *IEEE Communications Magazine*, vol. 36, no. 2, pp. 96–102, 1998.
- [36] A. Cigada, F. Ripamonti, and M. Vanali, "The delay & sum algorithm applied to microphone array measurements: numerical analysis and experimental validation," *Mechanical Systems and Signal Processing*, vol. 21, pp. 2645–2664, Aug. 2007.
- [37] 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, no. 3, pp. 163–167, 1986.
- [38] F. Simonetti and L. Huang, "From beamforming to diffraction tomography," *Journal of Applied Physics*, vol. 103, no. 103110, 2008.
- [39] P. Huthwaite, F. Simonetti, and L. Huang, "The Different Structural Scales of the Breast and Their Impact on Time-of-Flight and Diffraction Tomography," in *Medical Imaging 2010: Ultrasonic Imaging, Tomography, and Therapy* (S. D'hooge, J and McAleavey, ed.), vol. 7629 of *Proceedings of SPIE-The International Society for Optical Engineering*, SPIE, 2010.
- [40] G. Afken and H. Weber, *Mathematical Methods for Physicists*. Academic Press, 6th ed., 2005.
- [41] D. W. Pohl, W. Denk, and M. Lanz, "Optical stethoscopy: Image recording with resolution $\lambda/20$," *Applied Physics Letters*, vol. 44, no. 7, pp. 651–653, 1984.

- [42] Z. Wang, W. Guo, L. Li, B. Luk'yanchuk, A. Khan, Z. Liu, Z. Chen, and M. Hong, "Optical virtual imaging at 50 nm lateral resolution with a white-light nanoscope," *Nature Communications*, vol. 2, pp. 1–6, Mar. 2011.
- [43] M. Tabib-Azar, N. S. Shoemaker, and S. Harris, "Non-destructive characterization of materials by evanescent microwaves," *Measurement Science and Technology*, vol. 4, pp. 583–590, Dec. 1993.
- [44] A. T. Fernandez, K. L. Gammelmark, J. J. Dahl, C. G. Keen, R. C. Gauss, and G. E. Trahey, "Synthetic elevation beamforming and image acquisition capabilities using an 8 x 128 1.75D array.," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 50, pp. 40–57, Jan. 2003.
- [45] R. Waag, F. Lin, T. Varslot, and J. Astheimer, "An Eigenfunction Method for Reconstruction of Large-Scale and High-Contrast Objects," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 54, pp. 1316–1332, July 2007.
- [46] J. H. Lee and S. W. Choi, "A parametric study of ultrasonic beam profiles for a linear phased array transducer.," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 47, pp. 644–50, Jan. 2000.
- [47] I. J. Cox and C. J. R. Sheppard, "Information capacity and resolution in an optical system," *Journal of the Optical Society of America A*, vol. 3, pp. 1152–1158, Aug. 1986.
- [48] E. Pike, J. McWhirter, M. Bertero, and C. de Mol, "Generalised information theory for inverse problems in signal processing," *IEE Proceedings F Communications, Radar and Signal Processing*, vol. 131, no. 6, pp. 660–667, 1984.
- [49] A. Brancaccio, G. Leone, and R. Pierri, "Information content of Born scattered fields: results in the circular cylindrical case," *Journal of the Optical Society of America A*, vol. 15, no. 7, pp. 1909–1917, 1998.

- [50] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic coding for data compression," *Communications of the ACM*, vol. 30, pp. 520–540, June 1987.
- [51] C. Shannon, "Communication In The Presence Of Noise," *Proceedings of the IEEE*, vol. 86, pp. 447–457, Feb. 1998.
- [52] E. a. Marengo, F. K. Gruber, and F. Simonetti, "Time-reversal MUSIC imaging of extended targets.," *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 16, pp. 1967–84, Aug. 2007.
- [53] M. E. Testorf and M. A. Fiddy, "Superresolution Imaging—Revisited," in *Advances in Imaging and Electron Physics*, ch. 5, pp. 165–218, Elsevier, 2010.
- [54] C. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.
- [55] F. Simonetti, L. Huang, and N. Duric, "On the spatial sampling of wave fields with circular ring apertures," *Journal of Applied Physics*, vol. 101, no. 083103, pp. 1–6, 2007.
- [56] S. Yan, H. Sun, U. P. Svensson, X. Ma, and J. M. Hovem, "Optimal Modal Beamforming for Spherical Microphone Arrays," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, pp. 361–371, Feb. 2011.
- [57] O. M. Bucci and T. Isemia, "Electromagnetic inverse scattering: Retrievable information and measurement strategies," *Radio Science*, vol. 32, no. 6, pp. 2123–2137, 1997.
- [58] S. Qian, *Introduction to time-frequency and wavelet transforms*. 2002.
- [59] A. Shemer, D. Mendlovic, Z. Zalevsky, J. Garcia, and P. G. Martinez, "Superresolving optical system with time multiplexing and computer decoding," *Applied Optics*, vol. 38, no. 35, pp. 7245–7251, 1999.
- [60] J. W. Goodman, *Introduction to Fourier Optics*. McGraw-Hill, 2nd ed., 1996.

- [61] D. Colton and P. Monk, “A novel method for solving the inverse scattering problem for time-harmonic acoustic waves in the resonance region,” *Society for Industrial and Applied Mathematics*, vol. 45, no. 6, pp. 1039–1053, 1985.
- [62] D. Colton and P. Monk, “A novel method for solving the inverse scattering problem for time-harmonic acoustic waves in the resonance region II,” *SIAM Journal on Applied Mathematics*, vol. 46, no. 3, pp. 506–523, 1986.
- [63] F. Cakoni and D. Colton, “The linear sampling method for cracks,” *Inverse Problems*, vol. 19, pp. 279–295, 2003.
- [64] T. Arens, “Why linear sampling works,” *Inverse Problems*, vol. 20, pp. 163–173, Feb. 2004.
- [65] M. Hanke, “Why linear sampling really seems to work,” *Inverse Problems and Imaging*, vol. 2, no. 3, pp. 373–395, 2008.
- [66] A. Kirsch, “Characterization of the shape of a scattering obstacle using the spectral data of the far field operator,” *Inverse Problems*, vol. 14, pp. 1489–1512, 1998.
- [67] A. Kirsch, “The MUSIC algorithm and the factorization method in inverse scattering theory for inhomogeneous media,” *Inverse Problems*, vol. 18, pp. 1025–1040, 2002.
- [68] Q. Chen, H. Haddar, A. Lechleiter, and P. Monk, “A sampling method for inverse scattering in the time domain,” *Inverse Problems*, vol. 26, p. 085001, Aug. 2010.
- [69] T. Hutt and F. Simonetti, “Experimental observation of super-resolution in highly attenuative media,” *Submitted to Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2011.
- [70] F. Simonetti, N. Duric, and O. Rama, “Imaging beyond the born approximation: an experimental investigation with an ultrasonic ring array,” *Physical Review E*, vol. 76, pp. 1–10, Sept. 2007.

- [71] F. Cakoni, “Recent Developments in the Qualitative Approach to Inverse Scattering Theory,” *Journal of Computational and Applied Mathematics*, vol. 204, pp. 242–255, July 2007.
- [72] D. Colton and P. Monk, “Mathematical and Numerical Methods in Inverse Acoustic Scattering Theory,” *Zamm*, vol. 81, pp. 723–731, Nov. 2001.
- [73] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra, *MPI : The Complete Reference*. MIT Press, 1996.
- [74] L. Mordfin, ed., *Handbook of Reference Data for Nondestructive Testing*. ASTM, 2002.
- [75] D. E. Gray, ed., *American Institute of Physics Handbook*. McGraw-Hill, 3rd ed., 1957.
- [76] R. Long, J. Russel, P. Cawley, and N. Habgood, “Ultrasonic phased array inspection of flaws on weld fusion faces using Full Matrix Capture,” *Review of Quantitative Nondestructive Evaluation*, vol. 28, pp. 848–855, 2009.
- [77] G. Baskaran, K. Balasubramaniam, and C. Lakshmanarao, “Shear-wave time of flight diffraction (S-TOFD) technique,” *NDT & E International*, vol. 39, pp. 458–467, Sept. 2006.
- [78] F. Simonetti, “Localization of pointlike scatterers in solids with subwavelength resolution,” *Applied Physics Letters*, vol. 89, no. 094105, pp. 1–3, 2006.
- [79] T. Hutt and F. Simonetti, “Reconstructing the shape of an object from its mirror image,” *Journal of Applied Physics*, vol. 108, no. 064909, pp. 1–11, 2010.
- [80] T. Hutt and F. Simonetti, “Reconstructing the back of a defect from its mirror image,” *Insight*, vol. 52, no. 2, pp. 82–86, 2010.
- [81] I. Komura, “Crack detection and sizing technique by ultrasonic and electromagnetic methods,” *Nuclear Engineering and Design*, vol. 206, pp. 351–362, June 2001.

- [82] J. Zhang, B. W. Drinkwater, and P. D. Wilcox, “The Use of Ultrasonic Arrays to Characterize Crack-Like Defects,” *Journal of Nondestructive Evaluation*, vol. 29, pp. 222–232, Sept. 2010.
- [83] X. Pan and M. A. Anastasio, “On a Limited-View Reconstruction Problem in Diffraction Tomography,” *IEEE Transactions on Medical Imaging*, vol. 21, no. 4, pp. 413–416, 2002.
- [84] N. Bleistein, “Large wave number aperture-limited Fourier inversion and inverse scattering,” *Wave Motion*, vol. 11, pp. 113–136, 1989.
- [85] G. Kogan and J. H. Rose, “Limited aperture effects on ultrasonic image reconstruction,” *Journal of the Acoustic Society of America*, vol. 77, no. April, pp. 1342–1351, 1984.
- [86] R. Mager and N. Bleistein, “An examination of the limited aperture problem of physical optics inverse scattering,” *IEEE Transactions on Antennas and Propagation*, vol. 26, pp. 695–699, Sept. 1978.
- [87] R. L. Ochs, “The limited aperture problem of inverse acoustic scattering: Dirichlet boundary conditions,” *Journal of Applied Mathematics*, vol. 47, no. 6, pp. 1320–1341, 1987.
- [88] C. J. Nolan, M. Cheney, T. Dowling, and R. Gaburro, “Enhanced angular resolution from multiply scattered waves,” *Inverse Problems*, vol. 22, pp. 1817–1834, Oct. 2006.
- [89] F. Natterer, “Reflectors in wave equation imaging,” *Wave Motion*, vol. 45, pp. 776–784, June 2008.
- [90] W. A. Mulder and R.-E. Plessix, “Exploring some issues in acoustic full waveform inversion,” *Geophysical Prospecting*, vol. 56, no. 6, pp. 827–841, 2008.
- [91] C. Bunks, F. M. Saleck, S. Zaleski, and G. Chavent, “Multiscale seismic waveform inversion,” *Geophysics*, vol. 60, no. 5, pp. 1457–1473, 1995.

- [92] P. Huthwaite and F. Simonetti, “High-resolution imaging without iteration: a fast and robust method for breast ultrasound tomography,” *The Journal of the Acoustical Society of America*, vol. 130, pp. 1721–1734, Sept. 2011.
- [93] M. A. Doron and E. Doron, “Wavefield modeling and array processing. I. Spatial sampling,” *IEEE Transactions on Signal Processing*, vol. 42, no. 10, pp. 2549–2559, 1994.
- [94] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore: John Hopkins, 3rd ed., 1996.
- [95] J. C. Bancroft, “A visual relationship between Kirchhoff migration and seismic inversion,” *CREWES Research Report*, vol. 14, pp. 1–24, 2002.
- [96] R. Potthast, “A fast new method to solve inverse scattering problems,” *Inverse Problems*, vol. 12, pp. 731–742, 1996.
- [97] A. J. Burton and G. F. Miller, “The application of integral equation methods to the numerical solution of some exterior boundary-value problems,” *Proceedings of the Royal Society of London A*, vol. 323, pp. 201–210, 1971.
- [98] C. Ruud, “A review of selected non-destructive methods for residual stress measurement,” *NDT International*, vol. 15, pp. 15–23, Feb. 1982.
- [99] P. Purnell, T. Gan, D. Hutchins, and J. Berriman, “Noncontact ultrasonic diagnostics in concrete: A preliminary investigation,” *Cement and Concrete Research*, vol. 34, pp. 1185–1188, July 2004.
- [100] C. Li, L. Huang, N. Duric, H. Zhang, and C. Rowe, “An improved automatic time-of-flight picker for medical ultrasound tomography,” *Ultrasonics*, vol. 49, pp. 61–72, Jan. 2009.
- [101] A. Giacomini, “Ultrasonic Velocity in Ethanol–Water Mixtures,” *Journal of the Acoustical Society of America*, vol. 19, no. 4, pp. 701–702, 1947.