

IMPERIAL COLLEGE LONDON

QUANTITATIVE IMAGING WITH
MECHANICAL WAVES

by

Peter Edward Huthwaite

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

Department of Mechanical Engineering
Imperial College London
London SW7 2AZ

March 2012

Declaration of Originality

The content of this thesis is my own work, with supervision from Dr. Francesco Simonetti and Professor Michael Lowe. Where I have made use of the work of others, I have made this clear and provided appropriate references.

Peter Huthwaite

12/3/2012

Abstract

Quantitative imaging complements structural imaging by providing quantitative estimations of subsurface material properties as opposed to the sizes, shapes and positions of scatterers available from structural methods. The ability to reconstruct material properties from a series of wave measurements is extremely valuable in a range of applications as it potentially allows diagnostic technology with superior sensitivity and selectivity. Breast cancer, for example, is stiffer and hence of higher sound velocity than the surrounding tissue, so reconstructing velocity from ultrasonic measurements could allow cancer detection. Using this concept, breast ultrasound tomography has the potential to significantly improve the cost, safety and reliability of breast cancer screening and diagnosis over mammography, the gold-standard. Key to unlocking this potential is the availability of an accurate, fast, robust and high-resolution algorithm to reconstruct wave velocity. This thesis introduces HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography, a new imaging approach combining the complementary strengths of low resolution bent-ray tomography and high resolution diffraction tomography. HARBUT's theoretical foundation is explained and applied to simulated and experimental, in-vivo, breast ultrasound tomography data, confirming that it generates a step change in image quality over existing techniques, revealing lesions that would not be visible on a mammogram. This thesis also shows how, by combining data from many slices, the out-of-plane resolution can be significantly improved compared to treating each slice independently. HARBUT is applied to alternative problems including guided wave tomography, which aims to quantify the remaining wall thickness of a potentially corroded, inaccessible plate-like structure. Thickness estimates within 1mm for a 10mm nominal thickness plate were demonstrated for both simulated and experimental data. The thesis finally investigates HARBUT's performance with limited view configurations, and introduces VISCIT, the Virtual Image Space Component Iterative Technique, which accounts for the missing data, significantly improving the reconstructed image.

Acknowledgements

Firstly, I must thank my supervisors Dr. Francesco Simonetti and Professor Mike Lowe. Without their technical input, guidance and ability to put up with my frequently inappropriate sense of humour, both my time in the lab and this project would have been much poorer.

From the international breast imaging community, I would like to thank Dr. Nicole Ruiter and her group at the Karlsruhe Institute of Technology for the interesting discussions, and I would also like to thank Professor Neb Duric and his group at the Karmanos Cancer Institute, Wayne State University for making the in vivo breast ultrasound data used in Chapter 3 available.

I would like to thank Dr. Tim Hutt and Dr. Remo Ribichini for collaborating and sharing data. Acknowledgement should also be given to Dr. Pierre Belanger, who contributed the guided wave tomography data of Chapter 5.

Imperial's non-destructive testing group is a great place to work, and a large amount of the credit for this must go to Professor Peter Cawley. Along similar lines, my thanks also go out each one of my friends and colleagues in the lab who have made my time here hugely enjoyable.

Contents

1	Introduction	27
1.1	Motivation	27
1.2	Outline of thesis	33
2	Forward and inverse scattering theory	35
2.1	Introduction	35
2.2	Forward Scattering Theory	36
2.3	The HARBUT Method	40
2.3.1	An implementation of DT	40
2.3.2	Combining DT and BRT	43
2.4	Summary	45
3	Breast imaging with HARBUT	46
3.1	Introduction	46
3.2	Forward Model	51
3.2.1	Physical model	51
3.2.2	Array model	53

3.2.3	Numerical simulation	54
3.2.4	Data Processing	55
3.3	Results	57
3.4	Practical Considerations	60
3.4.1	3D effects and sampling conditions	60
3.4.2	The subtraction problem	65
3.4.3	Density	66
3.4.4	Amplitude correction	68
3.5	In vivo results	69
3.6	Summary	71
4	Three dimensional breast ultrasound tomography	73
4.1	Introduction	73
4.2	Theory	75
4.3	Model	80
4.4	Results	82
4.5	Summary	84
5	Guided Wave Tomography	86
5.1	Introduction	86
5.2	Method	88
5.2.1	Iterative HARBUT	90

5.2.2	Convergence criterion	91
5.3	Numerical results	92
5.3.1	Numerical model	92
5.3.2	Data processing	92
5.3.3	Reconstructions	94
5.4	Experimental models and results	100
5.4.1	Experimental setup	100
5.4.2	Data processing	102
5.4.3	Experimental results	103
5.5	Summary	108
6	Limited view quantitative imaging	110
6.1	Introduction	110
6.2	The limited view problem and regularisation	113
6.2.1	Consequences of the limited view	113
6.2.2	Regularisation to obtain additional information	116
6.3	A new approach to the limited view problem	118
6.3.1	Virtual transducers	118
6.3.2	Storing virtual components in the image space	120
6.3.3	Regularisation methods	122
6.4	Algorithm	123
6.5	Numerical model parameters	127

6.6	Numerical results	128
6.6.1	Five circular inclusions model	129
6.6.2	Gaussian profile example	131
6.6.3	Complex scatterer	133
6.7	Experimental results	135
6.8	Summary	138
7	Conclusions	140
7.1	Thesis Review	140
7.2	Main findings	141
7.3	Areas for future work	143
A	Ultrasound transmission through a penetrable cone	146
A.1	Introduction	146
A.2	The cone transmission model	148
A.3	Numerical study	150
A.4	Amplitude drop within the transducer plane	154
A.4.1	Plane wave illumination with finite-sized receiver	155
A.4.2	Finite transmitter and receiver	158
A.4.3	Parametric study	163
A.5	Deflection outside the transducer plane	166
A.5.1	The ‘smile’ effect	166

A.5.2	Amplitude drop outside the transducer plane	167
A.6	Conclusions	170
B	Beamforming to diffraction tomography for a non-circular array	172
	References	178
	List of Publications	191

List of Figures

1.1	Breast ultrasound tomography. In diagram (a), the patient lies prone with the breast suspended in a water tank. A transducer array begins at the chest wall and gathers sets of data at many slices through the breast. A cross section through the transducer array is shown in (b), where the scattered field from a single illumination is measured by the transducer array.	30
2.1	DT implementation achieved by filtering the beamforming image . . .	40
2.2	Illustration of the coordinates used in the scattering integral. The integral is performed across all points in the scatterer, Ω , to produce a value for the signal at the receiver \mathbf{x} for the source \mathbf{y}	41
2.3	Flowchart of the stages which make up HARBUT, combining DT and BRT	43
3.1	Schematic of the system for breast ultrasound tomography. (a) The patient lies prone with the breast suspended in a water tank. A transducer array begins at the chest wall and gathers sets of data at many slices through the breast. Shown in (b) is a single illumination and the scattered field produced which is measured by the transducer array.	48
3.2	3D numerical model of a breast with non-uniform sound speed	52

3.3	Density field for the 3D numerical breast model	53
3.4	Through transmission time traces from the 3D FDTD simulation. The source is at coordinates (+60mm, 0) and the receiver is diametrically opposite, at (-60mm, 0) in the x-y plane. (a) is the incident time trace, i.e. the time trace without any scatterer present to distort the wave field. The wave packet has become elongated by the numerical dispersion from the coarse FDTD grid. (b) is the time trace with the phantom model of Sec. 3.2.1 present. The signal arrives slightly earlier because of the higher average sound speed along the path between the source and receiver when the phantom is present. There is also an amplitude loss due to the wave's energy being scattered by the phantom.	55
3.5	Matrices of processed data, plotted for all send-receive pairs. (a) gives the arrival times of the modelled signals relative to the incident signal, estimated by a frequency domain deconvolution. (b) presents the amplitude, taken at 750kHz. In both plots, the regions marked 1 correspond to the transmit-receive pairs with a line-of-sight passing through the subcutaneous fat layer and those marked 2 to the transducer pairs that 'see through' the volume of the phantom.	56
3.6	(a) The original central slice of the sound speed map as in Fig. 3.2, with the locations of the transducer array marked. (b) is the BRT sound speed reconstruction. This is used as the background for the corrected beamforming at 750kHz in (c), which is then filtered to get the object function perturbation component O_δ given in (d). (e) is the full object function O generated by combining (d) and the background object function O_b calculated from (b). (f) is the hybrid sound speed reconstruction from the object function (e).	58

3.7	Standard DT reconstruction. The size and the contrast of the original phantom (a) are large enough that the standard Born approximation is invalid, causing the reconstruction (b) to have extensive artefacts that obscure the inclusions.	59
3.8	Comparison of the HARBUT reconstruction from 3D data (a) and 2D data (b). The boundaries of the glandular region and the phantom itself in the 3D reconstruction are blurred in comparison to the 2D reconstruction. The random medium which makes up the glandular region is also more homogeneous in the 3D reconstruction. These effects are a result of averaging in the out-of-plane direction.	61
3.9	Normalised Point Spread Function at the centre of the array for the system modelled in this chapter – 12mm tall transducers with an array diameter of 120mm. The PSF is thin – about 1mm ($\lambda/2$) wide – within the plane due to the resolution of the Born approximation used in the reconstruction. As shown in (d), taking a threshold at -6dB relative to the maximum, the PSF extends in the region $-4.5\text{mm} < z < 4.5\text{mm}$, making its height around 9mm. This is significantly wider than the in-plane PSF dimensions.	62
3.10	Schematic diagram of the effect of PSF averaging. The averaging in the z direction has little effect where the 3D boundaries are well aligned with the z axis, as with the inclusion, forming the sharper boundary marked. When this is not the case, as with the oblique boundaries of the phantom, significant blurring is visible in the 2D reconstruction. This effect is seen in the reconstruction of Fig. 3.6(f).	63

3.11	(a) is a schematic 3D diagram of the arrangement of scatterers in the spiral staircase model. (b) gives the 2D reconstruction of data from such a model, simulated with the 3D FDTD method, for heights 0-10mm at 0.5mm gaps and radii of 15-50mm at 5mm gaps. The artefacts surrounding each scatterer are a result of the relatively coarse FDTD mesh used, rather than the imaging process, and are ignored. The transducers are modelled as the 12mm tall line sources used in all the simulations in this chapter and are at a radius of 60mm. There is a clear drop in response as the height of the scatterer is increased due to the transducer beam height. Following Fig. 3.9(d), the 4.5mm z offset points lie around the -6dB threshold, indicating this is the boundary of the slice captured by the transducer array.	63
3.12	Image generated by imaging $\psi_\delta = \psi_b$ in the corresponding background velocity field. This is present in the final image if the subtraction is not performed. The error is within $\pm 1\text{m/s}$ through the majority of the imaging domain.	66
3.13	Image of the density component present in a reconstruction from a 2D simulation of the central slice	67
3.14	Amplitude of signal around the array relative to the incident signal at 750kHz. Averaging has been performed across all illuminations by matching up the measurements in the same positions relative to the source. The average amplitude drop of the signal, caused by passing through the phantom, was taken to be 0.5.	69
3.15	Sound speed reconstructions through the breast using in vivo data. (a) is the BRT method, the dominant method in the field of breast ultrasound tomography. This reconstruction was performed with one of the most up-to-date methods. (b) is HARBUT, showing a significant resolution improvement over BRT.	70

4.1	Signals from scatterers outside the plane of the array. Since waves scattered from A will travel further than from B, by a distance of $(a_1 + a_2) - (b_1 + b_2)$, they will encode an extra phase shift which contains information about the out-of-plane offset.	74
4.2	Comparison of 2D and 3D k-space vectors. In 2D, as shown in (a), there are only two source-receiver vector combinations which correspond to a particular point in the k-space. However, in 3D, as shown in (b), there is an infinite number.	76
4.3	Mapping of measured components into the k-space Ewald sphere with the array at a fixed axial location. (a) shows a schematic diagram of scattering with the array at a particular axial location, offset from the scatterer. (b) shows how the measurements at each axial location map into a single disc of the Ewald sphere.	77
4.4	Transducer beam variation out of the plane of the array. Shown is the amplitude of beams travelling through a distance of 200mm (the diameter of a typical array) for (a) a 4mm transducer, (c) a 12mm transducer and (e) a 20mm transducer. Normalisation is performed against the peak value within each horizontal line. (b), (c) and (f) plot the amplitudes at a distance of 100mm (i.e. the centre of the array) and 200mm (the opposite side of the array) for the 4mm, 12mm and 20mm transducers respectively.	79
4.5	Algorithm for generating a tree-like model. Branches are drawn by stepping along, randomly varying direction and reducing thickness and step length as the algorithm progresses. After a certain number of iterations, each branch splits into several smaller branches and the process repeats for each of these.	81

4.6	Reconstructions of a tree-like model. (a) gives the original model, generated by the method in Fig. 4.5. (b) presents the results obtained by treating each set of data separately under the 2D framework then stacking the results. (c) presents the image reconstructed using all the data sets together using the 3D framework discussed in this chapter.	83
5.1	(a) shows the dispersion curves for the fundamental Lamb wave modes in aluminium, for both phase and group velocity. The relationship between thickness and velocity at a given frequency provides a means to determine thickness. (b) shows the guided wave tomography configuration considered. Waves are excited by a single source, hit the defect and are scattered in all directions. The scattered field is measured from all directions by the transducer array. The process is repeated with each transducer in turn providing the illumination, until data for all possible send and receive combinations has been obtained.	88
5.2	Flowchart showing the stages of iterative HARBUT.	90
5.3	The HARBUT algorithm applied to a single central defect. The defect is circular, 60mm in diameter, and has a thickness of 50% of the plate itself (i.e. 5mm remaining thickness) - see the original map in (a). The 50kHz frequency component was used in the beamforming and DT-based reconstructions. The λ and $\sqrt{\lambda L}$ scales shown are 37.7mm and 174mm respectively. (b) shows the bent-ray tomography reconstruction, which is low resolution so cannot capture much of the detail. This is used to correct the beamforming algorithm in (c), which, when filtered produces the O_δ of (d). This is combined with the background object function to give the total object function O in (e). Converting this to thickness via the dispersion curves is shown in (f).	95
5.4	Cross sections of the thickness reconstructions from simulated data for a single circular defect.	96

5.5	Reconstructions of a two defect model. The central defect is 60mm in diameter and 5mm thick, and the offset defect is 100mm in diameter and 7mm thick, as shown in the original map of (a). (b) shows the bent-ray tomography image and (c) gives the HARBUT image at 50kHz. (d) applies 5 iterations of HARBUT to the reconstruction.	97
5.6	Cross sections of the reconstructions from simulated data for a plate with two circular defects.	98
5.7	Reconstructions of a complex corrosion patch. (a) is the original thickness map. (b) is the standard DT reconstruction, and (c) is the bent-ray tomography reconstruction. (d) gives the standard HARBUT reconstruction, showing an improvement over both the bent-ray and DT. (e), (f), (g) give the results after 2, 4 and 8 iterations. (h) gives the reconstruction after 8 iterations with artificial noise added to the data. (i) presents the convergence parameter Q .	99
5.8	Cross sections of reconstructions of the complex defect along the line marked in Fig. 5.7(a).	101
5.9	Experimental time trace from an axisymmetric central defect. Taking the defect as the origin, the source was at (400mm, 0) and the measurement was taken at (-324mm, 235mm) – corresponding to transducer 24 of 64. Noise was removed from the time trace by multiplying by the gating function. The gating function was defined as 1 between t_{arr} (the arrival time, $248\mu s$ in this case), and $t_{arr} + 120\mu s$. Outside this region the function smoothly returned to zero via a sine curve over a time of $40\mu s$.	102
5.10	Imaging a single central defect, diameter 60mm, depth 50%, as in Fig. 5.3(a). (a) gives bent-ray tomography, (b) shows diffraction tomography. (c) is a HARBUT reconstruction, and in (d) four HARBUT iterations have been applied.	103

5.11	Cross sections of the reconstructions from experimental data for a single circular defect.	104
5.12	Reconstructions from experimental data for a two defect model. The central defect is 60mm diameter and 50% (5mm) deep, and the offset defect is 100mm diameter and is 30% deep, as shown in Fig. 5.5(a). (a) is the bent-ray tomography method, (b) is diffraction tomography, (c) is single iteration HARBUT, and (d) is HARBUT with 6 iterations. In (e), I have taken (d) as a background, and averaged together HARBUT images from all the frequency components between 30 and 60kHz. This makes better use of the bandwidth to minimise noise in the data. (f) plots the convergence parameter Q	106
5.13	Cross sections of the reconstructions from experimental data for two circular defects in an aluminium plate.	107
6.1	Scattering from an object with a backwall reflector. (a) shows two scattering components. (b) shows how the backwall reflection component, expected to be much larger than any other component due to the low contrast, can be treated as though it was obtained from pair of parallel arrays.	113
6.2	The mapping of far field measurements to the 2D Fourier transform of the scatterer under the Born approximation. (a), (c) and (e) show schematic diagrams of the transducer layouts and (b), (d) and (f) show the corresponding Fourier components which can be obtained from each layout. (a) and (b) show a full view array as considered in earlier chapters; (c) and (d) are for a linear array obtaining data in reflection and (e) and (f) show a pair of parallel linear arrays measuring transmission data.	115

6.3	Schematic example of the approach used by Sung et al. to obtain values for the unknown Fourier components via a regularisation method. Fourier components from a limited view configuration with an array of sources above and receivers below are assembled in the k-space in (a). By Fourier transforming, the reconstruction can be obtained, as in (b), then regularisation is applied in (d), by setting any value less than zero to zero. (c) demonstrates that this regularisation process gives values to the unknown Fourier components. The estimates for these unknown values are combined with the known values in (e). By replacing (a) by (e), the iterative process is repeated until the unknown values converge such that they conform with both the regularisation method and the known values.	117
6.4	Schematic outline of the thresholding regularisation method. The initial reconstruction is shown in (a), with a threshold applied in (b) and the subsequent VISCIT reconstruction is shown in (c).	124
6.5	HARBUT for a limited view configuration using VISCIT. The stages in blue are needed to incorporate a background into the algorithm, which is necessary to extend the range of the algorithm beyond the Born approximation.	125

- 6.6 Algorithm performance with different sizes and contrasts. The wavelength is 0.8mm. The transducer array lies at the top surface of the image and the backwall coincides with the bottom surface. (a) gives the original sound speed map, consisting of three 2% contrast (1530m/s) inclusions sized 5mm, 10mm and 15mm in diameter, and two 10mm inclusions of 1% (1515m/s) and 5% (1575m/s). Reflection imaging in (b) shows no quantitative information; (c) shows that BRT is poor resolution and exhibits diagonal, ‘X’ shaped artefacts. The Born approximation is violated for most of the inclusions, as is clear from (d), the DT reconstruction, which is also poor because of the limited view. It is clear that HARBUT does provide an improvement, as shown in (e), but the high contrast inclusion in particular is poorly reconstructed with many artefacts. (f) presents HARBUT with VISCIT, which has few artefacts and is able to accurately capture the sound speeds and shapes of all five inclusions. 130
- 6.7 Reconstructions from a model representing stress damage within a metal component. The original model (a) is formed by summing a 5m/s Gaussian of radius 5mm to a random background. This background had a correlation length of 1.5mm and varied between 1499 and 1501m/s. (b) shows that without the presence of sharp boundaries, reflection imaging is completely unsuitable for detecting this early stage damage. (c) gives the BRT image, which, as in Fig. 6.6(c) features diagonal artefacts. (d) shows diffraction tomography, again with the diagonal artefacts. Because of the low contrast of the scatterer, the Born approximation works well so there is no sign of the Born artefacts seen before. HARBUT, in (e), shows a slight improvement over (d) although the diagonal artefacts are still present and the reconstructed values are underestimates of the true values. (f) shows HARBUT with VISCIT, confirming that the method is suitable for more realistic sound speed maps. 132

6.8	Reconstructions of a complex feature. (a) shows the original, complex scatterer. Little detail is visible in standard reflection imaging (b). The resolution of BRT, as shown in (c) is too poor to be able to identify any of the features. (d), showing diffraction tomography, does capture some of the features, but there are significant artefacts and the sound speeds are not accurate. (e) is the HARBUT reconstruction, which obtains a more accurate reconstruction, but again shows significant limited view artefacts. (f) shows how the VISCIT method obtains better sound speed estimates, although a few artefacts are visible.	134
6.9	Limited view GWT results; these correspond to the full view images of Fig. 5.12. All images shown here are monochromatic and use the 50kHz component from a set of experimental data. (a) shows the true model, a 10mm plate with two circular defects; the central one has thickness 5mm and diameter 60mm and the offset one has thickness 7mm and diameter 100mm, and is offset by 200 mm from the array centre. (b) shows the standard full view result for comparison. For the remaining images, the first column, i.e. (c), (e), (g) and (i) shows standard limited view imaging with HARBUT and the second, i.e. (d), (f), (h) and (j) shows VISCIT. In (c) and (d) the top half of the array is used as sources and the bottom half as receivers. In (e) and (f) the transducer arrays have been reduced; only 75% (i.e. 24 sources, 24 receivers) are used. (g) and (h) reduce that to 50% and (i) and (j) only use 25%, i.e. eight sources and eight receivers. All the VISCIT reconstructions used six HARBUT iterations and a total of 63 VISCIT iterations in the geometric progression discussed in Sec. 6.4.137	
6.10	Cross sections through the VISCIT images in Fig. 6.9, taken along a horizontal line passing through the centre of the array. As the amount of available information is reduced, the reconstructed contrast is reduced.	138

A.1	A practical example where understanding scattering from a penetrable cone is useful: breast ultrasound tomography. The bulk effect of the waves passing through the breast is approximated by acoustic transmission through a cone.	147
A.2	The cone distorts and scatters the field produced by the source transducer. The resulting field is measured by the receiver. The output plane is used to plot the output of the FDTD simulations in Sec. A.3.	148
A.3	3D transmission through the cone, as calculated by a FDTD simulation. Each shows the amplitude of the beam across the output plane of Fig. A.2 for different cone velocity values at 750kHz.	152
A.4	(a) gives the amplitudes along the $y = 0$ lines for each of the simulations in Fig. A.3, showing how the deflection increases with contrast. (b) estimates the receiver transducer response through a Huygens approach of summing across the surface of the transducer. Values in (b) are normalised against the value that would be generated if no cone was present.	153
A.5	Comparison between the 2D and 3D amplitude on the central plane. Each is normalised such that the maximum value of the incident beam is 1 in the corresponding 2D or 3D model.	155
A.6	Plane wave interaction with the cone, where $\beta = 30$ degrees (a) The refraction causes the wave field to be rotated. The effect of the transducer is to integrate the wavefield across its surface. (b) The estimated measurements from the transducer relative to the equivalent incident wavefield, for both the FDTD (dashed lines) and the analytical result (solid lines) of eq. (A.4). As the field gets rotated more, the integration causes the measurements to cancel, ultimately resulting in the zero for the 15mm tall transducer with the 1350m/s wedge. This is shown in (c) and (d).	157

A.7	Analytical model for estimating transmission response. The effects of the refraction of the cone (a) can be estimated from the straight beam as shown in (b).	159
A.8	Comparison of the analytical model of Fig. A.7 and the FDTD data for a range of different receiver transducer sizes. Normalisation is calculated relative to the same measurement without the cone present. The source transducer height is 12mm and $\beta = 30$ degrees.	161
A.9	The effect of transducer size on the response with different δ_2 values. The transmitter and receiver have the same size.	164
A.10	3D incident angles. The angle θ is a combination of θ_{xy} and θ_{rz} such that $\sin \theta = \sin \theta_{xy} \sin \theta_{rz}$	167
A.11	Paths through the cone plotted for three different rays. The deflection is clearly increased as two incident angles are combined. A ‘smile’ pattern can be seen in a similar way to the result of the 3D FDTD model in Fig. A.3.	168
A.12	Estimates of the transducer response as the transducer is moved away from the transducer plane. The solid lines are the values derived by summing sections of the FDTD responses of Fig. A.8. The dashed lines are the analytical estimates.	169
B.1	Coordinates for calculating BF weightings for a linear transducer array	174

B.2	Near field position dependency of filters. Note that just the angular spacing component is present; the weightings of eq. (B.7) are not included for clarity. (a) and (d) show the circular and square array configurations respectively, and mark the two points at which the filter is calculated. (b), the circular array filter at the centre, and (c), the circular filter at the offset location, show that the filter remains relatively constant when the position changes for the circular array. (e) shows that the square array filter is far more complex, and (f) shows that as the position varies, discontinuities appear in the filter which move as a function of position; as such the filter for the square array cannot be approximated as constant in space.	176
B.3	The angular ‘contribution’ of a transducer, used in the BF integral discretisation. Note that the extent of the transducer here is considered to be from halfway from the previous transducer to halfway to the next transducer; this definition is such that the sum of all angular contributions should sum to 2π if a full view configuration is used, or to the angular extent of a limited view array.	177

List of Tables

3.1	Material properties of the structures in the breast phantom	53
3.2	Dimensions of the inclusions in the breast phantom. Inclusion numbering is performed clockwise from the top as shown in Figs. 3.2 and 3.3.	54
A.1	Actual responses for the cone simulations of Sec. A.3 compared to the estimated responses from eq. (A.11). Responses are given as percentages of the signals measured without the cone present. Both the transmitter and receiver have width 4mm and height 20mm.	165

Nomenclature

c	Sound speed
δ	Dirac delta
f	Frequency
G	Green's function
k	Wavenumber
λ	Wavelength
n	Iteration count
Ω	Domain of the scatterer
O	Object function
ψ	Wavefield scalar potential
p	Pressure
Q	Convergence parameter: average thickness change between iterations over the damage area, relative to the background thickness
q	Scattering potential
ρ	Density
S	Defect area; 1 within the defect, 0 outside
T	Thickness
\mathbf{x}	Receiver transducer location
\mathbf{y}	Source transducer location

Subscripts

b	inhomogeneous background
δ	remainder between true model and inhomogeneous background
s	scattered, i.e. remainder between true model and homogeneous background
u	uniform background

Abbreviations

BF	beamforming
BRT	bent ray tomography
DT	diffraction tomography
FDFD	finite difference, frequency domain
FDTD	finite difference, time domain
FWI	full wave inversion
GWT	guided wave tomography
HARBUT	hybrid algorithm for robust breast ultrasound tomography
NDT	non-destructive testing
PSF	point spread function
RMS	root mean square
VISCIT	virtual image space component iterative technique

Chapter 1

Introduction

1.1 Motivation

The power of subsurface imaging has given humanity the ability to see beyond what is visible with our own eyes, providing us with information that can be used to make vital decisions such as whether it is necessary to perform an – often traumatic – biopsy on a patient with suspected cancer, where to spend huge resources digging for oil, or whether a pipe with a crack is likely to fail causing an expensive period of inactivity. The initial discovery of X-rays over a century ago [1] has been followed by the development of ultrasound [2], X-ray CT (Computed Tomography) [3] and MRI (Magnetic Resonance Imaging) [4] imaging techniques in medical imaging, with similar advances in geophysics using mechanical waves to provide subsurface information. Non-destructive testing (NDT) is a field where imaging is increasingly being used in the characterisation of defects [5]. These imaging techniques have made a previously unachievable level of information available to help make more informed decisions.

There is huge diversity in the types of waves that are used for imaging. An excellent example of the diversity of electromagnetic waves used comes from astronomy, where optical, infra-red, ultraviolet, gamma, radio and X-ray sections of the spectrum are used to make discoveries about the universe. X-rays are also used in the high-profile

backscatter security scanners being introduced at airports around the world [6], and radar - using radio waves - is used for applications such as detecting aircraft and producing maps of weather patterns [7]. Imaging with mechanical waves likewise has huge variety, with hand-held ultrasonic probes for medical imaging using megahertz frequencies to produce acoustic waves [8], or geophysical imaging using seismic waves of tens of hertz [9].

All of these wave-based imaging methods can be classified by whether they provide structural or quantitative information. Often, a situation will exist where we wish to detect a large discontinuity in material properties. An example in non-destructive testing is crack characterisation, where the size and shape (i.e. the structure) of the crack provides vital information about the life of the component, but the material properties of the substance (typically air) within the crack itself is of no interest. Similarly when using radar to scan for aircraft, the useful information obtained is the position of each aircraft – obtained from the sharp impedance discontinuity between the air and the aircraft reflecting measurable waves back to the receiver – rather than any quantifiable data about each aircraft’s material properties. These methods are classified as structural imaging, providing information about the locations, shapes and sizes of structures. Structural imaging is well developed, with a wide range of applications in areas such as ultrasonic medical diagnosis [8], radar [7], geophysics [10] and non-destructive testing [5].

More subtle changes in local properties can often provide useful indications of certain physical phenomena, however. An example comes from the detection of cancer masses within the breast, which are known to be stiffer than the surrounding tissue. If the stiffness at all points throughout the breast can be determined in some way, this would provide a mechanism to detect breast cancer. Methods reconstructing material properties like this are classified as quantitative imaging methods, in contrast to the structural imaging methods discussed above.

Breast cancer diagnosis therefore has huge potential to benefit from advances in quantitative imaging. Mammograms – which image the breast by physically compressing it between plates then X-raying it – have become the gold standard in the

diagnosis of breast cancer and are used extensively in screening programmes. The method relies on identifying cancer masses by their high density and hence increased absorption of photons over the surrounding medium. A significant weakness of the method is the problem of dense breast, where the density of the structures within the breast is high, often masking the presence of cancer. Coupled with the greater cancer risk associated with dense breasts, this is a major drawback of the method. Additionally, the technique is often painful and relies on ionising radiation, which paradoxically increases the risk of cancer.

One alternative solution is to use quantitative imaging to exploit the increased stiffness of breast cancer, which is done in the field of elastography [11]. Standard sonography hardware is used to reconstruct a pair of images – one undeformed and one when a known external load is applied – and from these it is possible to determine the strain field by correlating the images. The Young’s and shear moduli, which define the stiffness, can be determined by solving the inverse problem using this strain field. This quantitative imaging method has the potential to aid cancer detection by improving the sensitivity of sonography, but is challenging to implement because of the ill-posed nature of the inverse problem and its sensitivity to noise [12]. These problems are in addition to the operator dependency problem which is inherent in sonography.

Another example of quantitative imaging with ultrasonic waves is supersonic shear imaging [13]. A moving, high amplitude source is synthesised by illuminating sequentially along an array. By moving this effective source faster than the shear wave velocity, a shock wave can be produced, exciting an intense shear wave within the medium. The effects of this shear wave are measured through repeated high speed (5000 frames/s) imaging, using standard sonography methods. As with elastography, a cross correlation method is used to establish the effects of the waves on the medium and from this quantify tissue elasticity. This suffers from similar drawbacks to standard elastography.

Common to all wave imaging methods – regardless of the wave type or frequency – is the concept of an array of transducers; a number of measurements need to be taken




Figure 1.1: Breast ultrasound tomography. In diagram (a), the patient lies prone with the breast suspended in a water tank. A transducer array begins at the chest wall and gathers sets of data at many slices through the breast. A cross section through the transducer array is shown in (b), where the scattered field from a single illumination is measured by the transducer array. (a) is from [14].

at different locations or in different directions in order to obtain enough information to reconstruct an image. Modern advances in electronic hardware have allowed for rapid improvements in both the numbers of channels available and the quality of the measured data, and coupled with the availability of cheap computing power these advances unlock significant potential for improved imaging technology.

These advances in array technology have allowed breast ultrasound tomography to become an attractive alternative to mammography, while avoiding the inversion problems of elastography. This uses the relationship that high stiffness corresponds to high sound speed; therefore a quantitative image of sound speed would greatly aid the diagnosis and detection of cancer. In breast ultrasound tomography, as shown in Fig. 1.1, a circular transducer array surrounds the breast and illuminates it with ultrasound waves from different directions. A matrix of scattering data can be produced by collecting a column of scattered measurements for each illumination and repeating for illuminations from all directions. The challenge is to utilise these measurements in order to obtain a quantitative image of sound speed through the breast, given the complex inhomogeneous nature of the breast.

The algorithm must deal with a variety of scales, from the sub-millimetre scales of the milk ducts up to the hundreds of millimetres for the breast itself. In the middle of these scales are the cancer masses themselves, which the algorithm must reliably and accurately reconstruct. This multi-scale medium provides a significant challenge

to imaging algorithms.

An additional challenge is that the complexity of the breast extends into three dimensions, when the circular array is more suitable for two. Obtaining complete data from a 3D problem is much more challenging, since the number of transducers must – in effect – be squared in order to capture information from the third dimension. Devices have been built which attempt to capture 3D data, for example the system developed by Karlsruhe Institute of Technology [15], but such systems are generally highly undersampled due to the limited number of transducers, and hence struggle to produce satisfactory images. A common alternative is to sweep or rotate a linear array to obtain data sets at several out-of-plane locations (see e.g. [16–18]). The data available from this method is still less than would be available from a full 3D imaging array since the receivers are all in the same plane as the sources, and the acquisition time is also significantly increased which causes problems with moving subjects.

Non-destructive testing (NDT) is another area with the potential to benefit from quantitative imaging. Early stage stress damage (i.e. prior to the appearance of cracks) in a component is often associated with a localised change in ultrasonic sound speed. While there are no sharp material discontinuities associated with the damage which would appear with structural imaging, quantitative imaging methods can instead be used to generate a map of sound speed and hence allow the severity and extent of the damage to be determined.

Another NDT application for quantitative imaging is guided wave tomography. Often, large diameter pipes or similar plate-like structures can suffer from corrosion which can locally reduce the thickness of the wall. By reconstructing wave velocity, the thickness can then be estimated using the known dispersion relationship. Guided wave tomography therefore aims to use quantitative imaging methods to reconstruct a map of the remaining wall thickness, enabling judgements to be made about whether or not the wall has thinned enough to make failure likely. By contrast, structural imaging would only provide an indication that there is a defect present and its location, without any quantitative information about the wall thickness.

The aim of this thesis is to develop quantitative imaging methods for ultrasound tomography. In contrast to structural imaging, despite the clear benefits of being able to determine subsurface properties, quantitative imaging is relatively underdeveloped and hence holds huge potential for advances, with simple ray assumptions dominating the majority of quantitative imaging applications, including X-ray CT [19], breast ultrasound tomography [20–22], and guided wave tomography [23, 24]. These ray approximations are highly limiting for ultrasound tomography because the resolution is severely reduced by ignoring diffraction [25]. Higher resolution methods are available, such as diffraction tomography, which instead relies on the Born or Rytov approximations [19]. However, the former is limited to low contrast, small scatterers and the latter relies on phase unwrapping so is only suitable for simple, low contrast scatterers with very little noise. In both cases these restrictions are impractical for most ultrasound tomography applications.

One widely proposed alternative is non-linear, full wave inversion methods [26, 27]. These use numerical approximations of wave propagation (such as the finite difference method) to estimate the measurements which would be obtained from a particular initial model (commonly the low-resolution BRT – bent-ray tomography – image is used for this purpose). This model is then iteratively updated, usually using a gradient descent method, in order to improve the model to better match the measured data. While in principle such an approach can address the limitations of the current methods, these are difficult to implement in practice because of the presence of local minima and the significant computational expense of calculating many forward models. Experimental uncertainties are difficult to account for in the forward model, which means such algorithms are rarely robust and require very high signal-to-noise ratios.

A new ultrasound tomography algorithm was developed for breast imaging, to address the need for a robust, high-resolution method. Accurate sound speed reconstructions through the breast have the potential to significantly improve breast cancer diagnosis, leading to earlier detection and ultimately saving lives. The algorithm is built on general principles which are widely applicable; the algorithm was

therefore subsequently applied to other applications in NDT.

1.2 Outline of thesis

This thesis investigates quantitative imaging using mechanical waves. A new quantitative method, HARBUT (Hybrid Algorithm for Robust Breast Ultrasound Tomography), has been developed to obtain high-resolution sound speed maps through the breast by combining the complementary strengths of bent-ray tomography and diffraction tomography. In this thesis I outline the theory behind HARBUT and demonstrate its performance when applied to breast ultrasound tomography. The method is then extended to guided wave tomography and stress damage detection in the field of non-destructive testing, and modifications to the algorithm to improve its performance for these different applications are investigated and recorded.

Chapter 2 introduces the forward scattering problem, and provides the mathematical framework used throughout the thesis. The forward problem is first simplified using the Born approximation. However, for many applications, notably breast ultrasound tomography, the size and contrast of the scatterer is too great for the Born approximation to be valid. I make modifications to increase the range of applicability of the model using a known inhomogeneous background. HARBUT is then introduced, solving the inverse problem using this modified model and allowing a reconstruction of wave velocity to be generated from a series of ultrasonic measurements.

HARBUT is applied to breast ultrasound tomography in Chapter 3. The method is applied to both numerical and simulated data and is shown to provide robust, fast resolution improvements over the widely used bent-ray tomography method. I perform my reconstructions in 2D with data from a ring array surrounding the breast, and I discuss the problems the 2D approximation causes in the resulting image.

In Chapter 4 I introduce a 3D framework to solve the breast imaging problem more accurately than in 2D. By using data obtained with the array at several axial

locations – a synthetic aperture approach – I demonstrate how it is possible to improve the resolution in the out-of-plane direction.

In Chapter 5 I apply HARBUT to a true 2D problem, guided wave tomography. The small but high contrast corrosion defects of interest in guided wave tomography are challenging to HARBUT, and I explore the solution of repeatedly applying HARBUT to iteratively improve the solution, better capturing these high contrast defects.

All the problems so far have considered a full view configuration, where the transducer array surrounds the scatterer allowing illumination from all angles. In many cases it is not possible to completely surround the scatterer, rendering such a configuration impractical. In Chapter 6 I apply HARBUT to a limited view configuration, with a linear array placed on the surface of a component, parallel to a reflecting backwall surface. Even taking advantage of the backwall reflections to provide additional illumination angles, there is insufficient information from this data alone to reconstruct the image. This chapter discusses how HARBUT performs with this limitation. I also demonstrate VISCIT (the Virtual Image Space Component Iterative Technique), showing how I can iteratively apply regularisation and store the resulting information in image space components, which are then used to significantly improve the limited view reconstructions.

The conclusions of the thesis are summarised in Chapter 7, along with potential areas for future work.

Chapter 2

Forward and inverse scattering theory

2.1 Introduction

The interaction of waves with a scatterer encodes information about the scatterer's properties. Extraction of this information from measurements of the field is an ill-posed inverse problem [28] which is challenging to solve because of the complex, non-linear nature of wave scattering. Therefore, any imaging algorithm typically utilises a set of simplifying assumptions about how the waves propagate through the object. Suitable assumptions depend hugely on the parameters of the problem being considered, for example the scatterer size relative to wavelength or the impedance contrast of the scatterer; the goal for any imaging problem is to find a suitable set of assumptions which allow inversion to be performed in a fast and robust manner for all anticipated scatterers.

An excellent example of this comes from X-ray CT. One might attempt to perform a reconstruction which accounts for diffraction, refraction and reflection of the X-rays as they propagate. However, when X-rays propagate through the objects typically imaged, these phenomena have such a small effect that they can be completely ignored. Instead, the waves can accurately be described by a straight ray model,

where the amplitude of each measured signal is obtained via a line integral of energy absorption along the ray path. This straight ray assumption allows for fast, robust CT reconstructions [19].

This chapter introduces the underlying theory of ultrasonic scattering and the simplifying assumptions made to the model which will be used throughout this thesis. This background theory is based on the acoustic model, which is the accepted model of the behaviour of soft tissue, such as the breast, but can also be easily generalised to other applications, as will be discussed later in the thesis. The second half of the chapter then explains how, having made these simplifying assumptions, they can be incorporated into the new imaging algorithm, HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography.

2.2 Forward Scattering Theory

The standard acoustic wave equation [29] is the accepted model used to describe ultrasound propagation in many materials

$$\rho(\mathbf{r})\nabla\cdot\left[\frac{1}{\rho(\mathbf{r})}\nabla p(\mathbf{r})\right]-\frac{1}{c(\mathbf{r})^2}\frac{\partial^2 p(\mathbf{r})}{\partial t^2}=0 \quad (2.1)$$

where $p(\mathbf{r})$ is the pressure at point \mathbf{r} , $\rho(\mathbf{r})$ is the density and $c(\mathbf{r})$ is the sound speed. By converting this to the temporal frequency domain it can be rewritten as

$$(\nabla^2 + k_u^2)\psi = -O\psi \quad (2.2)$$

where ψ is the scalar potential of the field (equal to the Fourier transform of the pressure) and $k_u = 2\pi f/c_u$ is the wavenumber of the background where f is the frequency and c_u is the sound speed in the background. Throughout this thesis the subscript u is used to refer to this homogeneous background, which represents the uniform material properties of the domain if no scatterer was present. The object function, $O(\mathbf{r})$, is the mathematical representation of the scatterer and is defined as

$$O(\mathbf{r}) = k_u^2 \left[\left(\frac{c_u}{c(\mathbf{r})} \right)^2 - 1 \right] - \rho^{1/2}(\mathbf{r}) \nabla^2 \rho^{-1/2}(\mathbf{r}). \quad (2.3)$$

The aim of tomography is to reconstruct this object function. The second term in (2.3) accounts for variations in the local density field, ρ [30]. The dependence of the density term on the Laplacian means that it is only significant at an interface where rapid changes in density occur. This term will therefore be low in many applications and is considered negligible for the majority of this thesis.

In order to solve (2.2), the background wave field ψ_u is defined as the solution to the case where there is no scatterer present,

$$(\nabla^2 + k_u^2)\psi_u = 0 \quad (2.4)$$

and the Green's function, G_u , as the solution to

$$(\nabla^2 + k_u^2)G_u = \delta \quad (2.5)$$

where δ is the Dirac delta. Equation (2.5) can be solved analytically in 2D and 3D to give the solutions

$$G_u^{2D}(\mathbf{x}, \mathbf{x}_0) = -\frac{i}{4} H_0^{(1)}(k_u R) \quad (2.6)$$

$$G_u^{3D}(\mathbf{x}, \mathbf{x}_0) = \frac{\exp(ik_u R)}{4\pi R} \quad (2.7)$$

where $R = |\mathbf{x} - \mathbf{x}_0|$ is the distance from the source to the measurement point.

Equations (2.2), (2.4) and (2.5) are combined [19,31] with Green's theorem to obtain an integral equation over the domain of the scatterer, Ω , for the total field relative to the uniform background field

$$\psi = \psi_u - \int_{\Omega} G_u O \psi d\mathbf{x}' \quad (2.8)$$

which is the classic Lippmann-Schwinger equation. Physically, this can be interpreted as the total field equalling the background field plus a superposition of sources. These sources are responsible for the scattering effect of the medium, and each has a strength $-O\psi$, which is the total field multiplied by the object function at the location of interest.

The goal of tomography is to solve (2.8) given knowledge of ψ_u and ψ around the aperture of the array. This problem is non-linear due to the presence of the unknown total field ψ under the integral, and requires the use of iterative methods for its solution. A direct solution to (2.8) can be obtained with the Born approximation, i.e. $\psi \approx \psi_u$ within the integral, so that

$$\psi \approx \psi_u - \int_{\Omega} G_u O \psi_u d\mathbf{x}'. \quad (2.9)$$

This corresponds to assuming that the incident field is not perturbed as it travels within the object. However, due to the sound speed variation in the object relative to the homogeneous background, the phase of ψ will differ from that of ψ_u . The maximum phase error will increase as the sound speed contrast increases and as the size of the object increases. For large objects of high contrast, this phase error can approach π as the wavefield travels inside the object, so that $\psi \approx -\psi_u$, which renders the Born approximation completely invalid. A criterion can therefore be defined for the use of the Born approximation by requiring that the maximum phase distortion must be less than π [19]. It is this that limits the use of the standard Born approximation imaging methods for many applications. In the case of breast ultrasound tomography, for example, the breast is typically around 100mm (50λ at $\lambda = 2\text{mm}$) in diameter meaning that its sound speed contrast relative to the water bath background must be less than 1% to meet this criterion, which is an unrealistic assumption (see for example [21, 32]).

In this thesis I address the phase problem by dividing the object function into a sum of an ‘artificial’ inhomogeneous background object function O_b , and a small

perturbation O_δ so that

$$O(\mathbf{r}) = O_b(\mathbf{r}) + O_\delta(\mathbf{r}). \quad (2.10)$$

If only the inhomogeneous background were present, the background wave field, ψ_b , and its corresponding object function O_b would satisfy (2.2)

$$(\nabla^2 + k_u^2)\psi_b = -O_b\psi_b. \quad (2.11)$$

On the other hand, by keeping (2.2) in terms of the total field ψ but substituting in (2.10), propagation within the actual object is described by

$$(\nabla^2 + k_u^2 + O_b)\psi = -O_\delta\psi. \quad (2.12)$$

As with the homogeneous case, a Green's function G_b can now be defined, which gives the response due to a point source in the inhomogeneous background

$$(\nabla^2 + k_u^2 + O_b)G_b = \delta. \quad (2.13)$$

Using the same procedure as for the homogeneous background case the following inhomogeneous equivalent of (2.8) is derived

$$\psi = \psi_b - \int_{\Omega} G_b O_\delta \psi d\mathbf{x}'. \quad (2.14)$$

Provided O_δ is sufficiently small, the ψ term under the integral can be approximated with the background term ψ_b so that

$$\psi \approx \psi_b - \int_{\Omega} G_b O_\delta \psi_b d\mathbf{x}' \quad (2.15)$$

which is a more accurate version of (2.9). Equation (2.15) is central to the DBIM (distorted Born iterative method) that aims to solve (2.8) through an iterative scheme [33]. For (2.15) to be sufficiently accurate the background medium has




Figure 2.1: DT implementation achieved by filtering the beamforming image

to be selected so as to ensure that the phase difference between ψ and ψ_b is much less than π . This is the inhomogeneous equivalent of the standard Born criterion as given in [19].

In many cases, the sound speed contrast within the scatterer itself is generally low; it is instead the overall contrast of the scatterer relative to the homogeneous background that breaks the standard Born approximation. This is particularly the case for breast ultrasound tomography. Therefore, provided a background is available that accounts for the average speed through the scatterer then the approximation should be valid. Bent Ray Tomography (BRT), which ignores diffraction as its simplifying assumption, provides a low-resolution estimate of sound speed through the scatterer and hence provides a suitable background for many applications, as will be demonstrated in later chapters.

2.3 The HARBUT Method

This section uses the formulation from the previous section as a basis for HARBUT, the new imaging algorithm. The starting point is a particular implementation of diffraction tomography (DT) that has been introduced in [34].

2.3.1 An implementation of DT

This method is outlined in Fig. 2.1 and consists of two main steps: beamforming (BF) and the application of a filter. The beamforming algorithm is discussed first.

The coordinates \mathbf{x} and \mathbf{y} are introduced as the coordinates of the receiver transducer and source transducer as shown in Fig. 2.2. Using these, and defining the scattered




Figure 2.2: Illustration of the coordinates used in the scattering integral. The integral is performed across all points in the scatterer, Ω , to produce a value for the signal at the receiver \mathbf{x} for the source \mathbf{y} .

field as $\psi_s = \psi - \psi_u$, eq. (2.9) can be rewritten as

$$\psi_s(\mathbf{x}, \mathbf{y}) \approx - \int_{\Omega} G_u(\mathbf{x}, \mathbf{x}') O(\mathbf{x}') \psi_u(\mathbf{y}, \mathbf{x}') d\mathbf{x}'. \quad (2.16)$$

The symmetry of this equation can be demonstrated by recognising that $\psi_u(\mathbf{y}, \mathbf{x}') = G_u(\mathbf{y}, \mathbf{x}')$ i.e. the field produced by an illumination at \mathbf{y} can simply be replaced by the equivalent Green's function. This leads to

$$\psi_s(\mathbf{x}, \mathbf{y}) \approx - \int_{\Omega} G_u(\mathbf{x}, \mathbf{x}') O(\mathbf{x}') G_u(\mathbf{y}, \mathbf{x}') d\mathbf{x}'. \quad (2.17)$$

Now consider a single point scatterer in a homogeneous background. In this case, the equation is simplified from an integral to

$$\psi_s(\mathbf{x}, \mathbf{y}) = G_u(\mathbf{x}, \mathbf{x}') q G_u(\mathbf{y}, \mathbf{x}') \quad (2.18)$$

where the scattering potential of the point is q , located at \mathbf{x}' . Knowing the location of the point scatterer, (2.18) can be rearranged to determine the exact scattering potential from a single scattering measurement, i.e.

$$q = \frac{\psi_s(\mathbf{x}, \mathbf{y})}{G_u(\mathbf{x}, \mathbf{x}') G_u(\mathbf{y}, \mathbf{x}')}. \quad (2.19)$$

In this equation, it can be considered that the $1/G_u$ factors provide corrections to the scattered field ψ_s so as to account for the phase shift and amplitude change as the wave propagates through the medium. In general, however, the location of each scatterer is unknown, and there are multiple scatterers present so that the signals interfere and make the use of (2.19) unfeasible in its current form.

The solution considered here is to take advantage of the multiple send-receive pairs in the data, rather than just the single pair as above. By integrating (2.19) for all the receivers \mathbf{x} and all the sources \mathbf{y} in the array, the image at a prescribed point \mathbf{z} can be defined as

$$I^{BF}(\mathbf{z}) = \int_{\mathbb{S}} \int_{\mathbb{S}} \frac{\psi_s(\mathbf{x}, \mathbf{y})}{G_u(\mathbf{x}, \mathbf{z}) G_u(\mathbf{z}, \mathbf{y})} d\mathbf{x} d\mathbf{y} \quad (2.20)$$

where \mathbb{S} is the aperture of the transducer array. For real, sampled data the continuous integrals are replaced by discrete sums. Due to this integral/summing process, if there is a scatterer present at the imaging point then the integrand will sum coherently leading to a large value, but if there is no scatterer at the point then the summing will be incoherent and the result will be much weaker, with the values cancelling themselves out.

Equation (2.20) is a slightly modified version of the beamforming algorithm. Conventionally, beamforming is performed in the time domain by applying a backwards time shift to account for the shift that occurs as the wave propagates through the medium. It does not make any correction for amplitude. As such, the true frequency domain BF equation would be an altered version of (2.20),

$$I^{BF*}(\mathbf{z}) = \int_{\mathbb{S}} \int_{\mathbb{S}} \frac{\psi_s(\mathbf{x}, \mathbf{y})}{\text{sgn}[G_u(\mathbf{x}, \mathbf{z})] \text{sgn}[G_u(\mathbf{z}, \mathbf{y})]} d\mathbf{x} d\mathbf{y} \quad (2.21)$$

where the sign function is defined as $\text{sgn}(x) = x/|x|$ so that only the phase component of G_u is used. In this thesis the full value of the Green's functions is included as in (2.20) since this approach is more general. This would allow us, for example, to account for background media which cause significant amplitude changes, such as through attenuation.

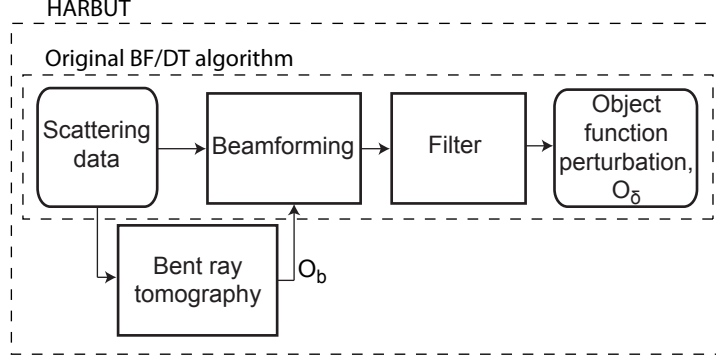


Figure 2.3: Flowchart of the stages which make up HARBUT, combining DT and BRT

As discussed in [34], the BF image is a distorted version of the DT reconstruction with different weights applied to different spatial frequencies. The DT reconstruction is obtained by removing this weighting from the BF image, which is done by Fourier transforming the BF image, applying the weighting corrections, and transforming the result back to the geometrical space. This approach is more flexible than directly generating the DT image, and is essential for HARBUT, as presented in the next section.

2.3.2 Combining DT and BRT

The case of imaging in an inhomogeneous background is now considered. This process is similar to the homogeneous case, with the addition of the BRT algorithm to provide a suitable background. Fig. 2.3 illustrates the process.

The first step is to perform the beamforming algorithm, this time accounting for the background. Starting from (2.15), $\psi_b = G_b$ is substituted and the source and receiver coordinates \mathbf{x} and \mathbf{y} are included

$$\psi(\mathbf{x}, \mathbf{y}) \approx \psi_b(\mathbf{x}, \mathbf{y}) - \int_{\Omega} G_b(\mathbf{x}, \mathbf{x}') O_{\delta}(\mathbf{x}') G_b(\mathbf{y}, \mathbf{x}') d\mathbf{x}'. \quad (2.22)$$

A BRT image is used to provide the background sound speed field. In this thesis BRT image is generated by the method in [35], although the actual method used is not critical to the success of the algorithm, provided it gives a reasonable low resolution

reconstruction of the sound-speed. The Green's functions for this background field need to be calculated via a numerical simulation, where each illumination is calculated separately to provide wavefield values at all points in the imaging domain. An eikonal equation solver [36] provides sufficiently accurate results throughout the domain, with significant speed advantages over alternatives such as frequency domain finite difference; a secondary advantage of using ray-approximation-based methods is that such solvers (and the solutions they provide) are often already available as part of the BRT algorithm. The numerical method is only used to provide a phase correction relative to the Green's function in the homogeneous background, so that

$$G_b(\mathbf{x}', \mathbf{y}) = G_u(\mathbf{x}', \mathbf{y}) e^{i\omega\Delta t} \quad (2.23)$$

where Δt is the difference in arrival time between the propagation in the background model and in the uniform medium, calculated via the eikonal solver.

If a point scatterer in an inhomogeneous medium is now considered, it is possible to follow through exactly the same logic as in Sec. 2.3.1 except starting with (2.22) instead of (2.17) to give the BF image of the perturbation relative to the background,

$$I_\delta^{BF}(\mathbf{z}) = \int_{\mathbb{S}} \int_{\mathbb{S}} \frac{\psi_\delta(\mathbf{x}, \mathbf{y})}{G_b(\mathbf{x}, \mathbf{z}) G_b(\mathbf{z}, \mathbf{y})} d\mathbf{x} d\mathbf{y} \quad (2.24)$$

where $\psi_\delta = \psi - \psi_b$ is the perturbation of the measured field relative to the background. Given that all of these quantities can be established, either from measurements or numerical models, the modified beamforming image can be formed from this equation. This is similar in approach to the Kirchhoff migration method used in geophysics [37]. Kirchhoff migration is only used on reflected data, however, and reconstructs the interfaces of impedance variations in the subsurface rather than material properties.

The BF image generated from (2.24) is then converted to the DT image using the filter introduced in [34]. Having obtained O_δ by filtering I_δ^{BF} , it is combined with the background velocity field according to (2.3) and (2.10), forming the final

HARBUT image. It should be observed that HARBUT solves (2.15); this process corresponds to the first iteration of the DBIM [33]. However, the HARBUT approach uses a different scheme to solve (2.15) based on the BRT starting model and the combination of BF and DT. This combination represents the main novelty of the method and is key to its robustness and speed.

2.4 Summary

This chapter has introduced the scattering framework used throughout this thesis and from it developed HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography. Linearising the scattering problem using the Born approximation allows for a relatively simple reconstruction algorithm, but since the Born approximation is only suitable for small, low contrast scatterers, the range of problems to which it can be applied is extremely narrow. HARBUT uses an existing estimate of the velocity through the scatterer, from the Bent Ray Tomography (BRT) algorithm, to greatly increase the range of applicability of the algorithm; the scatterer need only be of low contrast relative to the BRT velocity estimate for this method to be valid. The next chapter investigates the suitability of HARBUT for breast ultrasound tomography with the aid of a numerical model and a set of experimental in vivo data.

Chapter 3

Breast imaging with HARBUT

3.1 Introduction

Breast cancer is by far the most common cancer among women throughout the world, with 411 000 deaths per year [38]. Breast cancer screening programmes, where women over the age of 40 to 50 years have their breasts regularly checked for cancer, have been shown to reduce death rates [39, 40]. The current screening gold-standard is mammography, which projects X-rays through the breast so that the absorption of the photons by the structures within produces shadows which form an image. The detection of cancer is based on the assumption that the cancer mass is denser – and hence absorbs more X-ray radiation – than the surrounding tissue. The sensitivity (true positive rate) of the method is estimated at 68% to 88% [41], but this drops to around 30% to 40% in radiographically dense breasts where structures within the breast – parenchyma and stroma – mask the presence of cancer masses. Dense breast is a common occurrence, affecting approximately 50% of women under 50 years and a third over [42], and it is also in the latter group that the risk of developing cancer is the highest.

There are two main diagnostic tools complementary to mammography that are now routinely used in the clinic. The objective of these is to increase sensitivity and specificity (true negative rate). Magnetic Resonance Imaging (MRI) is one method,

producing accurate images and achieving high sensitivity [43]. However, its use leads to a large number of unneeded biopsies due to its low specificity [44]. Additionally the cost of examination is typically an order of magnitude more expensive than mammography, and it relies on the injection of contrast agents, making it unsuitable for widespread screening.

A second option used to aid diagnosis is sonography, which produces an image via a handheld ultrasound array. Sonography is more sensitive in dense breasts than mammography because it can distinguish between structures with similar density but different acoustic impedance. However, being handheld, it is highly operator dependent and its use is limited to situations where the areas of interest – an ambiguous mass for example – have already been identified. This makes the technique in its current form unsuitable for screening.

The possibility of improving the sensitivity and specificity of sonography has been investigated since the 1970s within the framework of breast ultrasound tomography [20, 21, 45–47]. Instead of the handheld probe, the breast is mechanically scanned with a toroidal ultrasound array that encircles the breast as shown in Fig. 3.1. The patient lies prone with the breast suspended in the water bath and the array is repeatedly moved down to image slices through it. At each slice, a single transducer provides an illumination, with the total field being recorded around the breast. This process is repeated with the next transducer providing the illumination and so on to obtain a full matrix of scattering data for each vertical location of the array. This matrix is then used to reconstruct the mechanical properties of the materials within the breast with the goal of using these to distinguish cancer from healthy tissue. Identifying the characteristic mechanical properties of cancer within the breast is key to achieve high sensitivity and specificity. Greenleaf first proposed that cancer masses are characterised by higher sound speed and attenuation than the surrounding medium [48]; recent work is suggesting a similar pattern [21].

Breast ultrasound tomography has the potential to address the main limitations of mammography and sonography. By using ultrasound, there is no need to rely on ionising radiation. Also, the design of the system means that the often painful




Figure 3.1: Schematic of the system for breast ultrasound tomography. (a) The patient lies prone with the breast suspended in a water tank. A transducer array begins at the chest wall and gathers sets of data at many slices through the breast. Shown in (b) is a single illumination and the scattered field produced which is measured by the transducer array.

breast compression needed by mammography is unnecessary, and the results are operator independent. Breast ultrasound tomography, as sonography, is unaffected by radiographically dense breasts that cause the high false negative rates with mammograms, and is relatively inexpensive which would allow widespread use as a screening technology.

The dominant approach in breast ultrasound tomography is Bent-Ray Tomography (BRT) [21, 22, 45, 46, 49–53] which applies a ray-based approach to arrival times for transmitted signals in order to produce – either directly in the case of straight rays or with iterations for bent rays – a sound-speed map. This follows the approach in X-ray CT based reconstruction systems which can rely on simple straight ray approximations [54, 55]. The diffraction that occurs around small scale objects is ignored under this approximation, leading to a resolution limit of $\sqrt{L\lambda}$ where L is the maximum distance between the transmitter and receiver and λ is the wavelength [25, 56]. In order to penetrate the breast, the minimum practical wavelength is 1mm, and with a typical 200mm array diameter the resolution is around 10mm – too low to detect small structures such as cancer at an early stage.

Higher resolution imaging methods based on the Born or Rytov approximations, such as Diffraction Tomography (DT) [19, 31, 57, 58], are suitable for imaging fine details of the breast architecture. However, DT is of little use in breast imaging

because the object to be imaged must be small relative to λ and of low contrast relative to the background for the approximations to be valid. The breast is a large object (around 50λ across at $\lambda = 2\text{mm}$) and the contrast is high enough that the criterion for the validity of the approximation, as given in [19], cannot generally be satisfied for breast ultrasound tomography.

An alternative solution being considered is to improve the resolution of the BRT image with an iterative full-wave inversion technique [26, 27, 59]. The method uses the low resolution BRT algorithm to reconstruct a starting model of the breast mechanical properties. The algorithm then runs a numerical computer simulation to predict the signals that would be measured with the system in Fig. 3.1 for this model. The target of the algorithm is to minimise the residual between the resulting signal and the measured signal by updating the material properties of the breast; the breast model that minimises the residual provides the final image. The model refinement is generally achieved by a gradient-based stepping method. The first issue with this technique is that the algorithm will only converge to the nearest local minimum rather than the global minimum of the problem. Because of this, the starting model – i.e. the image produced by the BRT algorithm – must already be close to the global minimum. Also, the technique is susceptible to uncertainties not considered in the forward model – for example transducer characteristics, 3D effects, and noise which can lead to the the algorithm converging to an incorrect solution. Speed is another significant drawback: a full set of illuminations needs to be simulated at each step (and more to calculate the gradient) with many iterations needed to generate the final image.

The previous chapter introduced HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography, which aims to address the need for a fast, robust, high-resolution breast ultrasound tomography method, by adjusting DT to use a background sound speed estimate such as from the BRT method. This provides a mechanism to combine the complementary strengths of the BRT and DT algorithms to reconstruct a high-resolution sound speed map.

This chapter applies the algorithm to a realistic numerical simulation and a set of

experimental data of a human breast to investigate its performance. It should be noted that the sound speed has the largest effect on the wave propagation and is therefore of fundamental importance in reconstructing other parameters; the focus of this chapter is therefore on sound speed. It is recognised, however, that density could subsequently be determined from the reconstructed object function by using multiple frequencies according to the approach outlined in [60].

This chapter evaluates the effects of the main experimental constraints which will affect the nature and accuracy of the measurements. The factors investigated are:

1. 3D effects and sampling conditions. The array architecture shown in Fig. 3.1 is suitable for imaging 2D (i.e. uniform in the out-of-plane direction) objects. The anatomy of the breast is fully 3D which would require a spherical array to perform all the measurements required to satisfy the Nyquist sampling criterion. Therefore it is important to understand whether the reconstruction can be treated as 2D and what type of artefacts one could expect as a result.
2. Knowledge of ψ_b . In order to form the beamforming (BF) image with (2.24), the scattered field $\psi_\delta = \psi - \psi_b$ needs to be known. Although ψ is directly available from the measurements, ψ_b cannot be measured. It would therefore be necessary to use a forward solver that predicts the outcome of the measurements that would be taken if only the background medium was present. However, for the calculations to be sufficiently accurate, one would need a very accurate model of transducer response and to be able to account for 3D effects. This challenge is magnified because $|\psi_\delta| \ll \{|\psi|, |\psi_b|\}$, meaning that even small errors in the estimate of ψ_b would lead to large errors in ψ_δ . Therefore, it is critical to understand whether ψ could be used in (2.24) instead of ψ_δ .
3. Density. Once the object function has been reconstructed, the sound speed and density have to be obtained by inverting (2.3). However, $O(\mathbf{r})$ at a single frequency does not contain sufficient information to extract both material properties. As a result, it is important to understand if the density term in (2.24) can be neglected to obtain sound speed only.

4. Attenuative effects. As the incident wave travels through the breast, its amplitude will be reduced due to the attenuative properties of human tissue. It is important to understand the effect this has on the resulting image and whether it can be accounted for by the algorithm.

Sections 3.3 and 3.4 show how the proposed method addresses these points, but first I introduce the forward model used to provide numerical data with which to test the algorithm.

3.2 Forward Model

This section details a numerical model used to test HARBUT, as introduced in Chapter 2. The aim is to demonstrate the robustness of the new algorithm by generating data which reproduces realistic experimental conditions.

Such data is provided by a 3D model solved numerically with the Finite Difference Time Domain (FDTD) method. This models sound speed and density in 3D, as well as accounting for the size of the transducers in the z direction.

3.2.1 Physical model

Selecting a suitable, realistic breast model is challenging since the necessary material properties of the breast are difficult to measure and tend to vary from person to person. The wide range in quoted values (see for example [61–63]) reflects this variation. Therefore, I use the more extreme, higher contrast values (which are more challenging to breast ultrasound tomography algorithms) in order provide a thorough test of the new algorithm.

The model is fully 3D to represent the actual shape of the human breast, as shown in Figs. 3.2 and 3.3 for the sound speed and density respectively. A glandular region forms the bulk of the breast model, with an irregular subcutaneous fat layer around

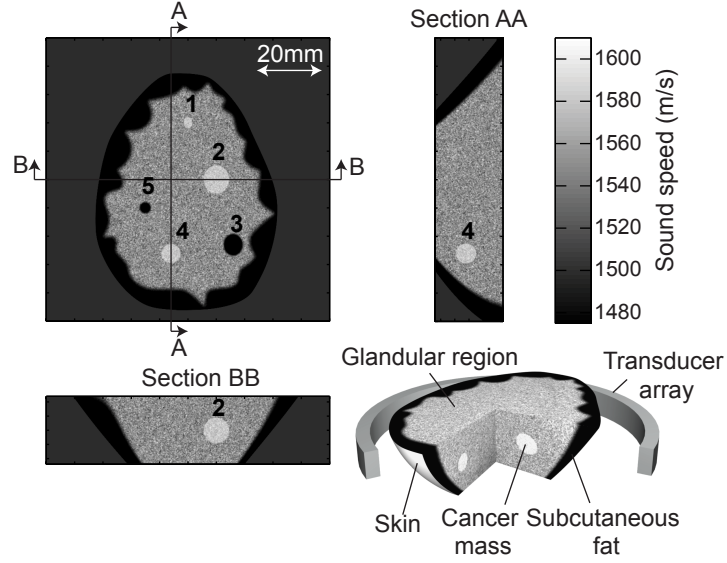


Figure 3.2: 3D numerical model of a breast with non-uniform sound speed

the boundary. Material properties for the glandular region and fat are given in Table 3.1. Existing numerical studies have so far used data generated by 2D models used to test breast ultrasound tomography algorithms; a novel aspect of this work is the use of more realistic 3D data for this purpose.

The glandular region is represented by a random medium. Importantly at the low frequencies used in breast ultrasound tomography, the presence of random variation is the main cause of signal attenuation due to scattering. This attenuation is an additional challenge for the imaging algorithm to overcome. I use a random medium model based on that outlined in [63], although with longer correlation lengths due to the relatively large spacing of the FDTD grid. Following the empirical relationship that density tends to vary linearly with sound speed [32], I use the same random field pattern for both density and velocity.

Five inclusions are placed in the model: three representing cancer masses and two representing fat spheres. The material properties of these are given in Table 3.1 and the dimensions and locations in Table 3.2. The goal of the imaging algorithm is to be able to detect these and identify whether each inclusion is cancer from the reconstructed mechanical properties.

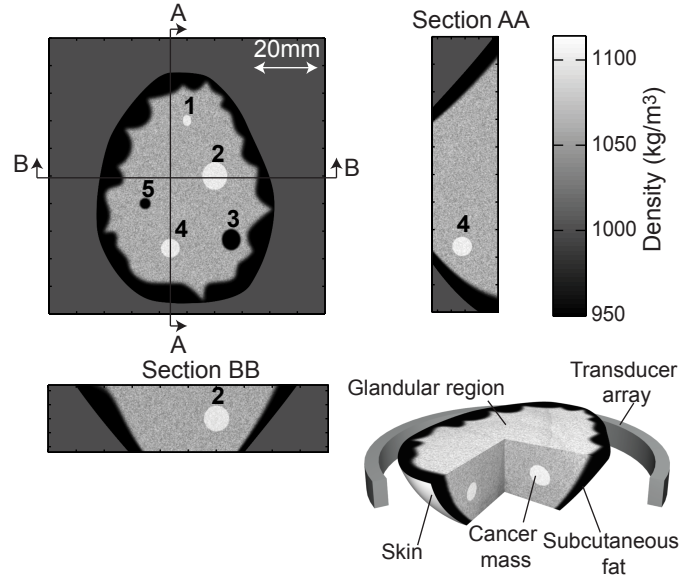


Figure 3.3: Density field for the 3D numerical breast model

Table 3.1: Material properties of the structures in the breast phantom

Structure	Sound speed (m/s)	Density (kg/m ³)	Standard deviation (%)	Correlation Length (mm)
Water	1500	1000	-	-
Glandular region	1550	1060	2	1.5
Fat	1470	950	-	-
Cancer masses	1580	1100	1	1.5

3.2.2 Array model

A 450 transducer array (sufficient elements according to the sampling criterion in [64]) is considered. The transducers are 12mm tall (in the out-of-plane, z , direction) but thin within the plane of the array so as to act as line sources. The array has a diameter of 120mm to limit the size of the propagation domain and hence the computational burden of the 3D model. Each array element is excited with a 3-cycle Hann windowed toneburst at a centre frequency of 750kHz.

Table 3.2: Dimensions of the inclusions in the breast phantom. Inclusion numbering is performed clockwise from the top as shown in Figs. 3.2 and 3.3.

Inclusion Number	Location (mm)			Diameter (mm)			Type
	x	y	z	x	y	z	
1	0	-20	0	3	4	4	Cancer
2	10	0	0	9	10	9	Cancer
3	16	23	-1	7	8	8	Fat sphere
4	-6	26	1	7	7	7	Cancer
5	-15	10	0	4	4	4	Fat sphere

3.2.3 Numerical simulation

A 3D FDTD modelling method is used with a standard Yee grid [65], with the mesh terminated with efficient convolutional perfectly matched layers [66] to minimise reflections from the boundary of the domain. 14 nodes were used per 2mm wavelength, so that the grid spacing is $1/7$ mm. A domain of 861 by 861 by 189 nodes was used. This is large enough for the 120mm array in the x and y directions and is 24mm tall in the z direction to allow the beam to diffract as it travels into the domain. A Courant number of 0.95 is chosen which gives a time step of 2.96×10^{-8} s (based on a maximum velocity of 1620m/s), therefore needing 4056 time steps for the 0.12ms simulation (long enough for a wave in water to travel 1.5 times the array diameter). The array locations are rounded to the nearest node so the recorded coordinates for each transducer are adjusted accordingly. 450 separate simulations have to be performed – one for each illumination. To model the effect of the transducer out of the plane, each transducer is modelled as a set of in-phase point sources at all nodes along a 12mm tall line.

A single illumination for this configuration takes around 4 hours to run on a single core of an AMD Opteron 8384 2.7GHz processor. Parallelism was achieved in a coarse-grained manner by running many – typically 64 – illuminations simultaneously. Using a cluster of two quad-core Intel Xeon E5462 2.8GHz processors per

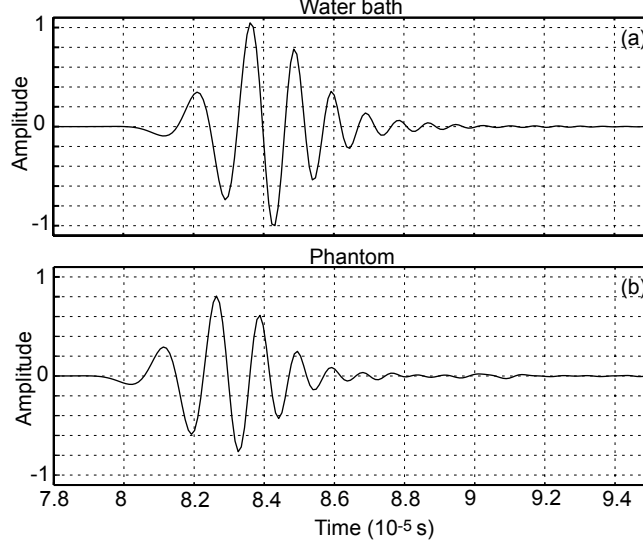


Figure 3.4: Through transmission time traces from the 3D FDTD simulation. The source is at coordinates $(+60\text{mm}, 0)$ and the receiver is diametrically opposite, at $(-60\text{mm}, 0)$ in the x-y plane. (a) is the incident time trace, i.e. the time trace without any scatterer present to distort the wave field. The wave packet has become elongated by the numerical dispersion from the coarse FDTD grid. (b) is the time trace with the phantom model of Sec. 3.2.1 present. The signal arrives slightly earlier because of the higher average sound speed along the path between the source and receiver when the phantom is present. There is also an amplitude loss due to the wave’s energy being scattered by the phantom.

node and 40 nodes, the running of all 450 simulations was completed within around 48 hours, with quite a large dependency on the queueing system in place on the cluster. These times are for the forward simulation only; the actual imaging algorithm, running on a single computer, is several orders of magnitude faster.

3.2.4 Data Processing

Two sets of data are generated: an incident data set where the material properties throughout the model are the same as that of water and the set for the case where the breast phantom described in Sec. 3.2.1 is present. When performing experimental measurements, the incident data set can be generated by taking measurements when there are no objects present in the water bath. These data need to be processed prior to passing them through the imaging algorithms.

Fig. 3.4 is a comparison of the unprocessed time traces from through transmission

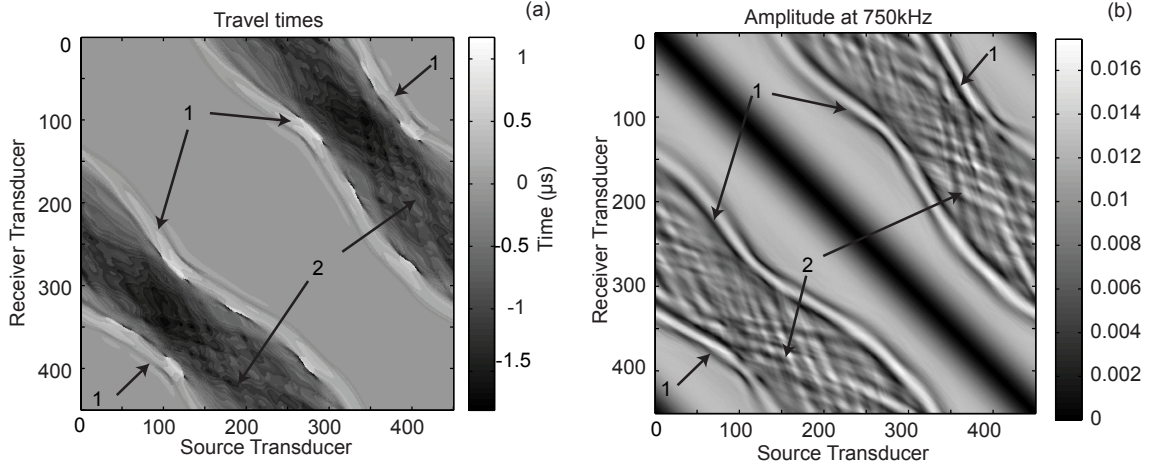


Figure 3.5: Matrices of processed data, plotted for all send-receive pairs. (a) gives the arrival times of the modelled signals relative to the incident signal, estimated by a frequency domain deconvolution. (b) presents the amplitude, taken at 750kHz. In both plots, the regions marked 1 correspond to the transmit-receive pairs with a line-of-sight passing through the subcutaneous fat layer and those marked 2 to the transducer pairs that ‘see through’ the volume of the phantom.

from the transducer with coordinates $(+60\text{mm}, 0)$ to the receiver at point $(-60\text{mm}, 0)$. Fig. 3.4(a) shows the time trace for transmission through the water bath, i.e. the incident field. The coda present is a result of numerical dispersion occurring in the FDTD method due to the relatively coarse mesh. Fig. 3.4(b) is the same measurement, taken when the phantom is present, showing the wave packet arriving earlier than the incident wave. This is because the average sound speed along the path between the source and receiver is higher when the phantom is present.

Fig. 3.5(a) presents this difference in arrival times, estimated for all send-receive pairs relative to the incident field. The arrival time is defined as the time at which the first disturbance is measured at the receiver. These have to be established quantitatively in order to be able to perform the bent-ray tomography (BRT) algorithm; here I use a deconvolution via the frequency domain. The regions marked as 2 correspond to the send-receive pairs with a line of sight through the glandular region of the model; the main effect is to make the waves arrive earlier due to the higher sound speed of the glandular material. To either side of these diagonal regions, in the boundaries marked 1, the waves do not pass into the glandular region but only pass through the subcutaneous fat, which is slower than the water bath and leads to later arrival

of the waves. These therefore appear lighter in the figure.

Fig. 3.5(b) gives the amplitude of each send-receive pair, produced at 750kHz by taking a discrete Fourier transform of the measured data. The data have been normalised such that they represent what would be seen if a unit point source was used to provide the illumination. If there was no scatterer present then the field would be the same as the Green's function sampled around the array. Under this condition, a singularity would be present along the principal diagonal due to the measurement being taken at the source location. However, this singularity is removed by gating off the incident signal for the source and surrounding measurements, which leaves the dark stripe down the diagonal.

The bright boundaries (marked 1) are the waves which have passed solely through the subcutaneous fat. The relatively high amplitude indicates that some form of focusing is occurring. Waves passing through the bulk of the glandular region, 2, have lower amplitude due to the effects of the random glandular material that scatters sound in all directions in space. The diagonal crossing patterns throughout this region are due to the presence of the inclusions.

3.3 Results

Fig. 3.6 shows the separate stages of HARBUT compared to the original sound speed model of the central slice of the phantom in Fig. 3.6(a). The BRT reconstruction in Fig. 3.6(b) shows the expected low resolution characteristics since diffraction is not accounted for. The reconstruction allows the detection of the subcutaneous fat layer and the glandular region; however, the inclusions are not reliably detectable. Enhanced BRT reconstructions are likely to be possible – for example through improved arrival time estimates – but the improvements will still be fundamentally limited by the algorithm's inability to deal with diffraction.

Fig. 3.6(c) is the modified BF image using the background from Fig. 3.6(b). In this reconstruction, the total field ψ is used rather than the perturbation ψ_δ . The




Figure 3.6: (a) The original central slice of the sound speed map as in Fig. 3.2, with the locations of the transducer array marked. (b) is the BRT sound speed reconstruction. This is used as the background for the corrected beamforming at 750kHz in (c), which is then filtered to get the object function perturbation component O_δ given in (d). (e) is the full object function O generated by combining (d) and the background object function O_b calculated from (b). (f) is the hybrid sound speed reconstruction from the object function (e).




Figure 3.7: Standard DT reconstruction. The size and the contrast of the original phantom (a) are large enough that the standard Born approximation is invalid, causing the reconstruction (b) to have extensive artefacts that obscure the inclusions.

features are not clear at this stage. Converting to the DT image in Fig. 3.6(d), using the filter from [34] based on the free space k_u , allows all five inclusions to be identified. This perturbation component is sufficiently small that the approximation $\psi = \psi_b$ necessary for the algorithm to work is valid. The total object function in Fig. 3.6(e) is calculated by combining O_b from Fig. 3.6(b) and the O_δ in Fig. 3.6(d) according to (2.10).

Fig. 3.6(f) is the final sound speed reconstruction obtained from Fig. 3.6(e) by inverting (2.3) and ignoring density effects. All five inclusions can be very clearly seen and all the irregular features of the subcutaneous fat layer are reconstructed. The sharp boundaries at the edge of the model and at the edge of the glandular region are blurred in the reconstruction; also the random medium representing the glandular region appears more homogeneous. This blurring is due to an averaging effect in the z direction that determines the so called slice thickness as discussed in Sec. 3.4.1. However, the sound speed is correctly reconstructed and there are no artefacts. The image is a dramatic improvement over the BRT reconstruction of Fig. 3.6(b).

To demonstrate the effectiveness of HARBUT, Fig. 3.7(b) shows the reconstruction obtained with the standard BF/DT algorithm [34] with a homogeneous water

background. Due to the size and contrast of the phantom relative to the water background, the condition for validity of the Born approximation is violated. As the illuminating field travels inside the phantom, it accumulates a phase delay larger than π . As a result, the total field that appears under the integral in (2.8) is in opposition of phase with the free space incident field (which replaces it under the Born approximation) leading to the artefacts in Fig. 3.7(b).

The current implementation of the BF/DT stage uses Matlab. To generate the 481 by 481 pixel image given the background correction data currently takes around 60 seconds on an HP z600 dual quad-core workstation without significant optimisation. The background correction data, required to calculate G_b , are taken from the final iteration of the BRT algorithm and therefore have no associated overhead. A C++ version of the algorithm has been written, using the excellent Eigen matrix library [67] which generates the image in around 5 seconds.

3.4 Practical Considerations

In this section I discuss the practical aspects that were introduced at the end of Sec. 3.1, points 1) to 4).

3.4.1 3D effects and sampling conditions

Fig. 3.8 compares the hybrid image from the 3D data as in Fig. 3.6(f) with a reconstruction using data from a similar simulation, except performed in 2D using the central slice of Fig. 3.2. The 2D reconstruction has sharp boundaries at the edge of the glandular region and the breast itself, which are both blurred in the 3D version. The granular appearance of the random medium is also better defined in the 2D reconstruction than the 3D. These differences are caused by blurring in the z (out-of-plane) direction in the 3D reconstruction due to the finite height of the transducer beam, as will be explained in this section.




Figure 3.8: Comparison of the HARBUT reconstruction from 3D data (a) and 2D data (b). The boundaries of the glandular region and the phantom itself in the 3D reconstruction are blurred in comparison to the 2D reconstruction. The random medium which makes up the glandular region is also more homogeneous in the 3D reconstruction. These effects are a result of averaging in the out-of-plane direction.

For this purpose it is necessary to consider the 3D Point Spread Function (PSF) which gives the response of the imaging system to a point scatterer. If the PSF is space invariant, i.e. does not depend on the position of the point scatterer, then the image is a convolution of the PSF with the original object function. The PSF is space invariant for standard DT [34], assuming plane wave illuminations and measurements taken in the far field, and here I assume that it can also be considered space invariant for HARBUT due to the relatively low contrast of the background sound speed map.

To generate the 3D PSF, a point scatterer at the origin – the centre of the array at $z = 0$ – is imaged with the transducer array at several axial locations; these images are then stacked to form the point scatterer response. Fig. 3.9 gives the PSF for the system considered in this chapter. In the x and y directions it is 1mm ($\lambda/2$) thick because of the Born approximation resolution limit, but in the z direction – as shown in Fig. 3.9(d) – the response stretches out to around $\pm 4.5\text{mm}$, using a threshold of -6dB relative to the maximum.

Due to the convolution, each point in the final image will be a weighted average of the object function in the z direction with the weights defined by the PSF projec-

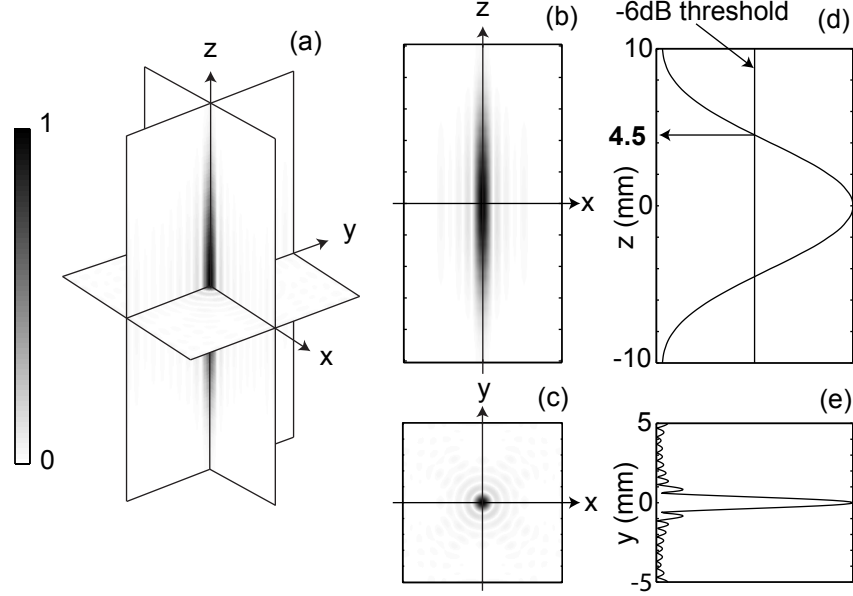


Figure 3.9: Normalised Point Spread Function at the centre of the array for the system modelled in this chapter – 12mm tall transducers with an array diameter of 120mm. The PSF is thin – about 1mm ($\lambda/2$) wide – within the plane due to the resolution of the Born approximation used in the reconstruction. As shown in (d), taking a threshold at -6dB relative to the maximum, the PSF extends in the region $-4.5\text{mm} < z < 4.5\text{mm}$, making its height around 9mm. This is significantly wider than the in-plane PSF dimensions.

tion in Fig. 3.9. Fig. 3.10 illustrates this, showing how the tall nature of the PSF blurs the reconstruction in the z direction and the effect this has on the 2D projection. The boundaries of the reconstruction of the elliptical inclusion are relatively sharp because of the 3D boundaries being primarily in the z direction, however, the oblique boundaries of the phantom itself are significantly more blurred. The phantom reconstruction in Fig. 3.6(f) shows this blurring at the boundary.

The extent of the interval in the z direction over which material properties are averaged defines the slice thickness. To estimate the slice thickness, a simple ‘spiral staircase’ model is used. This consists of point scatterers at a series of heights and a series of radii, with the points at each particular height being arranged in one radial direction, forming the steps in the staircase. Fig. 3.11(a) shows this schematically.

The scatterers are placed at heights of 0-10mm with 0.5mm gaps and radii of 15-50mm with 5mm gaps. Fig. 3.11(b) is the image obtained with the hybrid method at one position of the array. For an ideal imaging system with 0mm slice thick-




Figure 3.10: Schematic diagram of the effect of PSF averaging. The averaging in the z direction has little effect where the 3D boundaries are well aligned with the z axis, as with the inclusion, forming the sharper boundary marked. When this is not the case, as with the oblique boundaries of the phantom, significant blurring is visible in the 2D reconstruction. This effect is seen in the reconstruction of Fig. 3.6(f).




Figure 3.11: (a) is a schematic 3D diagram of the arrangement of scatterers in the spiral staircase model. (b) gives the 2D reconstruction of data from such a model, simulated with the 3D FDTD method, for heights 0-10mm at 0.5mm gaps and radii of 15-50mm at 5mm gaps. The artefacts surrounding each scatterer are a result of the relatively coarse FDTD mesh used, rather than the imaging process, and are ignored. The transducers are modelled as the 12mm tall line sources used in all the simulations in this chapter and are at a radius of 60mm. There is a clear drop in response as the height of the scatterer is increased due to the transducer beam height. Following Fig. 3.9(d), the 4.5mm z offset points lie around the -6dB threshold, indicating this is the boundary of the slice captured by the transducer array.

ness the image should contain only a single set of eight scatterers along the radial direction in the plane at $z = 0$. Instead, due to the spreading of the PSF in the z direction, weaker reconstructions of the scatterers from different heights can be observed. However, the amplitude of the reconstructed scatterers decays as the corresponding distance from the plane of the array increases. In particular a -6dB drop in amplitude can be observed for the scatterers at $z = 4.5\text{mm}$, thus verifying that the slice thickness is around 9mm .

The finite slice thickness is a result of the reduction in sensitivity of the transducer array to scatterers at greater distances from $z = 0$. The sensitivity of the array drops because there is a reduction in both 1) the amplitude of the illuminating beam incident on the scatterer and 2) the sensitivity of the receiving transducer to waves from the scatterer. Here, I consider only a point scatterer along the axis of the array so that the distances to all transducers are the same, and therefore, by the principle of reciprocity, 1) and 2) will both cause the same amplitude drop. Therefore, the -6dB drop in array sensitivity, which defines the boundary of the slice thickness, will be achieved when the illuminating beam and the receiver sensitivity each drop by -3dB . Considering the Fraunhofer zone of a line transducer [68], the thickness of the beam, B , at the centre of the array with a -3dB threshold is

$$B = \frac{0.884\lambda r_{arr}}{h} \quad (3.1)$$

where r_{arr} is the radius of the array and h is the transducer height. For the case considered here, the -6dB slice thickness becomes 8.8mm at the centre of the array, which is close enough to the 9mm slice thickness to verify the validity of (3.1). Equation (3.1) shows that it is possible that the slice thickness could be reduced by increasing the height of the transducers, although the benefit of this is limited because the Fraunhofer approximation becomes invalid with large transducers. Alternatively, a synthetic aperture approach could be used to reduce the slice thickness as investigated in Chapter 4.

3.4.2 The subtraction problem

The beamforming algorithm in (2.24) uses the field ψ_δ , which is calculated as the difference between the measured field and the background field

$$\psi_\delta = \psi - \psi_b. \quad (3.2)$$

However, to obtain the reconstructions in Fig. 3.6 I have used the total field ψ , which is directly available from the measurements, in place of ψ_δ . Here, I justify why this is possible.

In principle, ψ_b could be calculated by solving (2.11) with the FDTD method. However, this would not be reliable in practice because the perturbation field ψ_δ is small compared to ψ and ψ_b , so any small errors in the estimation of ψ_b will cause large errors in ψ_δ . Errors in the ψ_b estimation are unavoidable, mainly because of uncertainty in the transducer response.

Here, I show that it is not necessary to perform the subtraction and it is sufficient to form the BF image from the measured total field directly, i.e.

$$I_\delta(\mathbf{z}) = \int \frac{\psi}{G_b G_b} d\mathbf{x}'. \quad (3.3)$$

For this purpose it is observed that the BF and filtering steps in HARBUT are linear (for a fixed background) with respect to the measurements. This means that I can define a linear operator, \mathfrak{J}_{DT} , that maps the measurements or data, d , onto an image, i , i.e.

$$i = \mathfrak{J}_{DT}(d). \quad (3.4)$$

When the data correspond to ψ_δ then $O_\delta = \mathfrak{J}_{DT}(\psi_\delta)$. If instead, the data correspond to the total field then

$$\mathfrak{J}_{DT}(\psi) = \mathfrak{J}_{DT}(\psi_\delta) + \mathfrak{J}_{DT}(\psi_b) = O_\delta + \mathfrak{J}_{DT}(\psi_b). \quad (3.5)$$

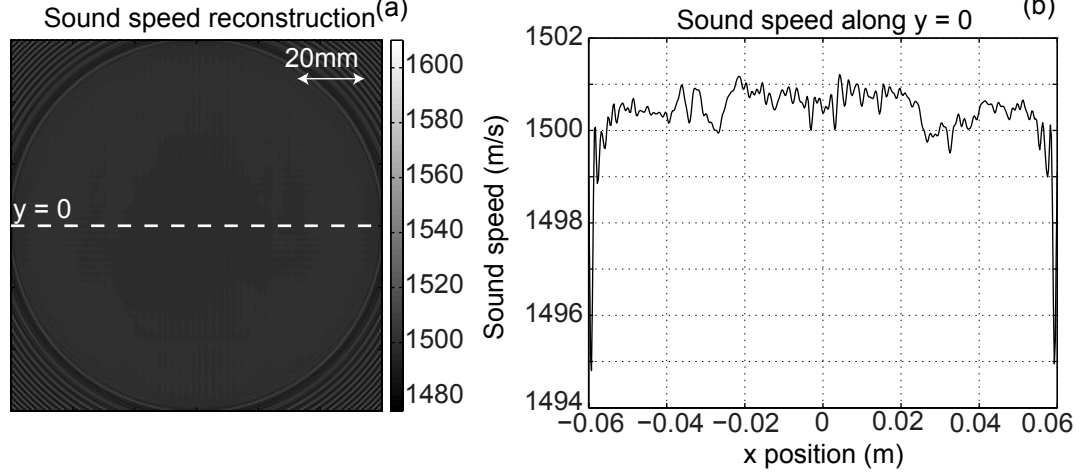


Figure 3.12: Image generated by imaging $\psi_\delta = \psi_b$ in the corresponding background velocity field. This is present in the final image if the subtraction is not performed. The error is within $\pm 1\text{m/s}$ through the majority of the imaging domain.

The term $\mathfrak{J}_{DT}(\psi_b)$ therefore represents the error caused by making the assumption that the background field does not need to be subtracted. Fig. 3.12 shows $\mathfrak{J}_{DT}(\psi_b)$ converted to a velocity image for the ψ_b calculated for the BRT background of Fig. 3.6(b). The velocity map is within $\pm 1\text{m/s}$ of the background except the region outside the array where ring artefacts appear, which is an acceptably small error given that the structures of interest have a sound speed contrast relative to the background in the order of 5%.

This convenient property makes the proposed approach very robust because it means that it is not necessary to estimate ψ_b , thus avoiding significant sources of error.

3.4.3 Density

Here I show that density variations within the breast can be neglected.

Equation (2.3) defines the object function, which includes a term dependent on the density field. The form of this term means that in order for the density to contribute a significant amount to $O(\mathbf{r})$, a large density gradient must be present. Within soft tissue, density varies continuously; therefore the object function, O , is mainly defined by the sound speed. Even if the density were to vary suddenly at




Figure 3.13: Image of the density component present in a reconstruction from a 2D simulation of the central slice

the interface of a cancer mass, the limited density contrast would still make the density term in (2.3) negligible. This is shown in Fig. 3.13, which is the image from a 2D model using uniform sound speed and the central slice of the density model of Fig. 3.3. The sound speed map is obtained assuming that the density term in (2.3) is negligible. If this were true the reconstructed velocity should be 1500m/s (the same as the water background) across the image plane, but instead different values of sound speed are seen where density discontinuities occur, showing that density affects the sound speed reconstruction to some extent. However, the values of sound speed significantly differ from the background velocity only at the boundaries of sudden density variations. This leads to two main conclusions: 1) the absolute value of density does not affect the velocity estimate; 2) the errors in velocity at sudden density variations help the visualisation of these boundaries and therefore aid the definition of complex morphologies.

3.4.4 Amplitude correction

The amplitudes of the waves drop as they pass through the 3D breast phantom. One cause of this is the random medium of the glandular region, which scatters energy from the waves in all directions. Some of the energy is scattered out of the plane of the array and is lost; 2D reconstructions ignore this so it becomes a form of attenuation. A second amplitude loss occurs due to the oblique angle at which the wave hits the phantom boundary. The wave is refracted slightly upwards (or downwards depending on the relative sound speeds of the breast and the water) which causes an amplitude drop due to the misalignment of the wave with the receiver [69]. Appendix A contains a detailed analysis of this amplitude drop. In experimental data both these effects will be present, along with the attenuation caused by the material itself. This loss in amplitude, if it is not accounted for, will cause the reconstructed object function perturbation O_δ to be too small. One solution to this problem would be to correct the amplitude loss in the same way phase is corrected. A ray-based attenuation image could be formed in the manner of the BRT image, which would include material attenuation, as well as 3D scattering and deflection since the effects are inseparable. From this, the wave amplitudes at all points in the domain for all illuminations could be calculated using a forward model, which could then correct the Green's functions used in the BF algorithm.

The simulations show that it is sufficient to assume that all received signals are attenuated by a constant factor. This allows the correction to be applied by simply multiplying all the measured signals by this factor. By taking the measurements of Fig. 3.5(b), dividing by the equivalent incident field (i.e. the field corresponding to a unit point source as defined by $G_u(\mathbf{x}, \mathbf{y})$), and averaging for the positions around the array relative to the source, Fig. 3.14 can be produced. From this, the amplitude drop factor was taken to be 0.5. Accordingly, prior to reconstructing the images from this model, all the signals were multiplied by a correction factor of 2.




Figure 3.14: Amplitude of signal around the array relative to the incident signal at 750kHz. Averaging has been performed across all illuminations by matching up the measurements in the same positions relative to the source. The average amplitude drop of the signal, caused by passing through the phantom, was taken to be 0.5.

3.5 In vivo results

Having demonstrated HARBUT's performance for a set of simulated data and discussed the practical limitations, I now apply HARBUT to a set of experimental data from a real human breast. This data set was acquired with the CURE (Computed Ultrasound Risk Evaluation) system at the Karmanos Cancer Institute, Wayne State University; details about the system are included in [21]. The array has 256 transducers and a diameter of 200mm. The transducers are 12mm tall and data is extracted from frequencies between 700 and 800kHz.

The arrival times were extracted by the method in [70] and a bent-ray tomography image was generated using the method outlined in [22]. HARBUT was performed taking this as a background, with each of the frequency components extracted, and the resulting sound speed images were averaged to minimise incoherent noise (following the approach in [58]).

The current state-of-the-art in breast ultrasound tomography, bent ray tomography, is compared to HARBUT in Fig. 3.15. It is clear that HARBUT provides a significant resolution improvement over BRT, with far more detail visible. Many sections

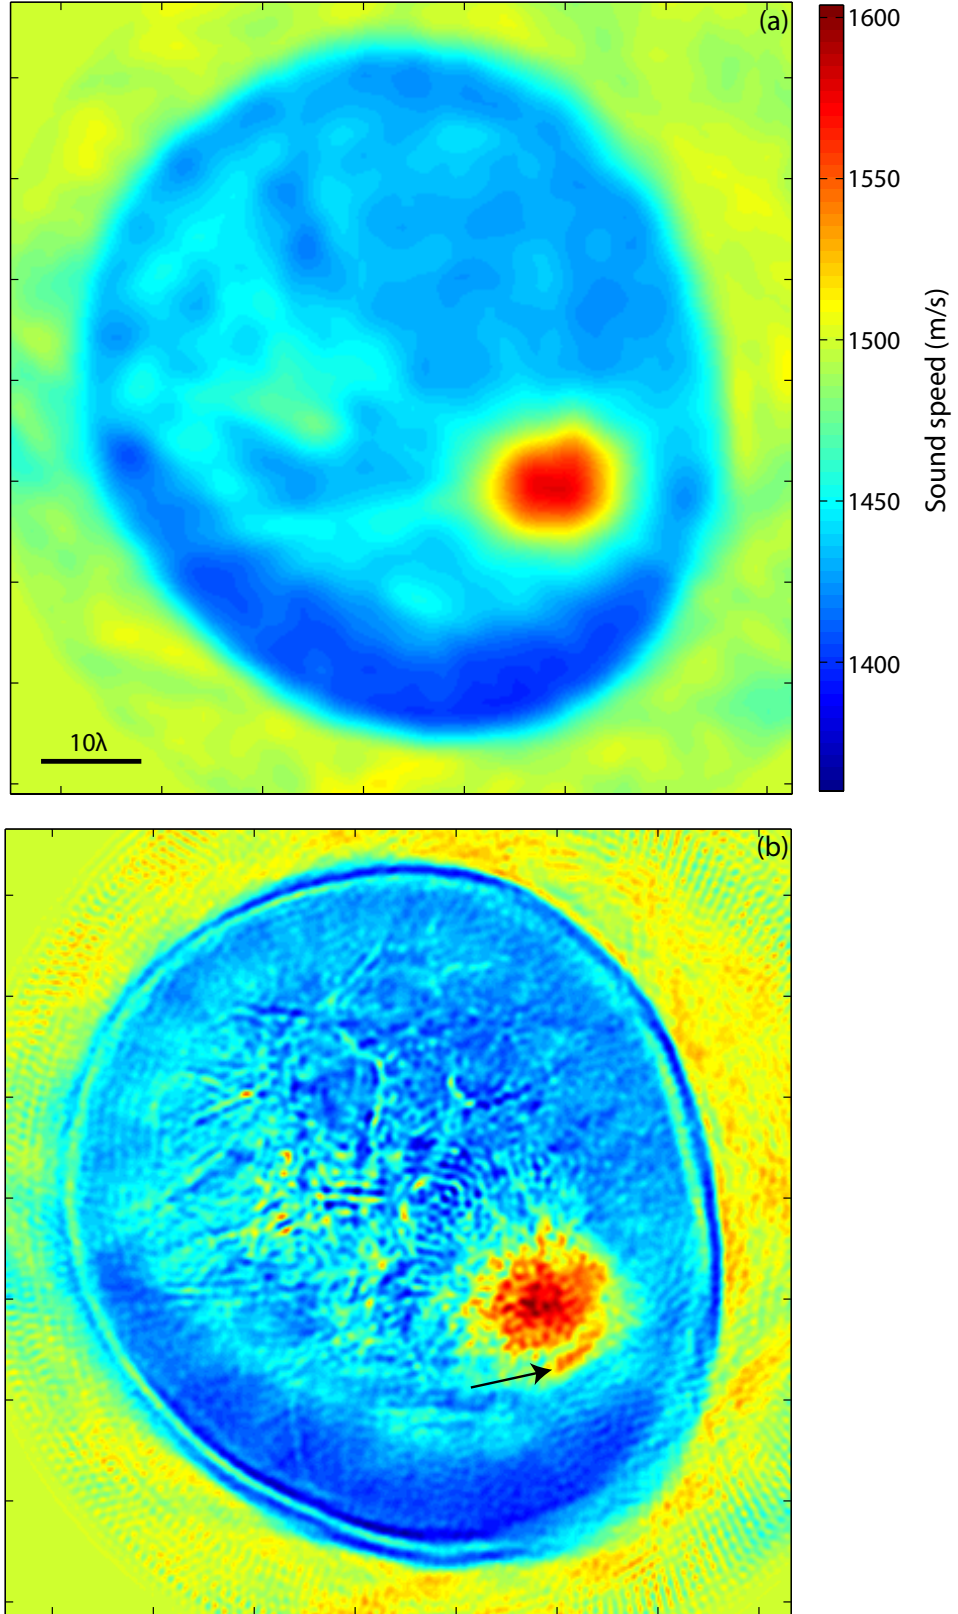


Figure 3.15: Sound speed reconstructions through the breast using in vivo data. (a) is the BRT method, the dominant method in the field of breast ultrasound tomography. This reconstruction was performed with one of the most up-to-date methods, as outlined in [22]. (b) is HARBUT, showing a significant resolution improvement over BRT.

of tree-like structures are visible, which are typical of the blood vessels, milk ducts and lymphatic vessels which can be found in the breast. Also, significantly more detail of the cancer mass itself is visible, with a notable separate section to the bottom right, marked with an arrow, which is completely missed by the BRT algorithm.

The slice thickness problem discussed in Sec. 3.4.1 will cause any features within the slice thickness to be projected into the imaging plane. This can explain some of the lack of clarity; we would expect to see many overlapping ducts and vessels within the expected $\sim 10\text{mm}$ slice thickness. Reducing the slice thickness, i.e. improving the axial resolution, is a priority to improve the reconstructions. The next chapter discusses a potential solution, using a fully 3D framework rather than the 2D approximation made here.

3.6 Summary

I have applied HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography, which provides a resolution improvement over bent-ray tomography while avoiding the convergence and speed problems of iterative methods, to a realistic simulated breast model and experimental in vivo breast data. Diffraction tomography methods are unsuitable for this purpose because the contrast and size of the breast relative to the homogeneous water background breaks the Born approximation. By reformulating the problem using an inhomogeneous background which is sufficiently close to the actual sound speed map, the relative contrast can be reduced such that the approximation becomes valid.

HARBUT is demonstrated to accurately reconstruct the sound speed through a breast phantom model from 3D simulated data, despite sampling the wavefield with an array architecture suitable for 2D imaging, and the presence of uncertainties such as transducer response that are likely to occur in real experiments. At a frequency of 750kHz, masses as small as 4mm in diameter can be clearly imaged. An in-plane resolution of 1mm was achieved, with a slice thickness of 9mm. Density contrast and randomly varying material properties with sub-wavelength coherence lengths

have little influence on the final reconstruction. 3D structures intersecting the plane of the array are partially projected onto the imaging plane due to the size of the slice thickness.

Applying HARBUT to in vivo data demonstrates that the algorithm performs well with experimental data, with a significant resolution improvement over the more widely used BRT algorithm. Small structures within the breast are visible, although as shown with the simulated data, these structures will be projected into the imaging plane because of the extent of the slice thickness. It is possible, by treating the problem as fully 3D and combining data from multiple axial locations, to improve the axial resolution and hence reduce the slice thickness; this is investigated in the next chapter.

Chapter 4

Three dimensional breast ultrasound tomography

4.1 Introduction

While the principles of two dimensional imaging can often be easily generalised to three dimensions, the complexities of adding another dimension often renders three dimensional imaging impractical. One important reason is that, while arrays for 2D imaging are formed from 1D lines or arcs of transducers, to capture the necessary information in 3D with fixed arrays requires extension of the array into an additional dimension. To obtain the same level of coverage in the third dimension as in the original two, it is necessary to square the number of transducers, which is impractical with current array technology.

One solution is to treat the 3D problem as though it was 2D, which would allow data to be captured with a 1D array of transducers. For a 3D problem to be successfully imaged in such a 2D framework, however, it must be quasi-uniform in the out-of-plane direction; this is not the case for the breast, which typically varies by similar amounts in all directions. Chapter 3 demonstrated the problems associated with simplifying 3D to 2D, showing how each point in the 2D reconstruction corresponds to an average of sound-speed values across the slice thickness. To improve this image




Figure 4.1: Signals from scatterers outside the plane of the array. Since waves scattered from A will travel further than from B, by a distance of $(a_1 + a_2) - (b_1 + b_2)$, they will encode an extra phase shift which contains information about the out-of-plane offset.

while maintaining the 2D reconstruction method, it is necessary to reduce the slice thickness, which can be achieved by reducing the transducer beam width.

The beam spread angle of the far field transducer beam will always reduce as the transducer width increases, so one approach is to go for taller transducers. However, as the transducers become taller, the near field, whose beam height is typically similar to that of the transducer, dominates the imaging region. Given this compromise, it has been found that is difficult to achieve better than the 9mm slice thickness obtained with 12mm transducers in Chapter 3.

This chapter presents an alternative solution, inspired by synthetic aperture radar techniques. As illustrated in Fig. 4.1, the signal measured from scatterer A, offset out of the plane, will have a phase shift due to the additional distance travelled compared to the measurement from scatterer B, which lies in the plane ($a_1 + a_2$ vs. $b_1 + b_2$). This phase shift encodes information about the offset of the scatterer out of the plane, and this can be exploited to help improve resolution in the out-of-plane direction.

By moving the transducer array, a synthetic aperture can be formed in the axial direction. By combining the data from multiple slices in a 3D framework, it is possible to exploit the information encoded in the out-of-plane phase distortions of Fig. 4.1 in order to obtain a significantly better axial resolution than with the 2D assumption.

Using this principle, Simonetti and Huang [14] showed that under idealised conditions, by combining data from many slices, and treating the problem as fully 3D, it

is possible to completely populate the Ewald sphere, thus obtaining a resolution of $\lambda/2$ in all directions. The formulation requires that point sources and receivers are used and are in the far field (i.e. sufficiently far away that any illumination can be treated as plane waves); also the array ring has to traverse an infinite distance to achieve maximum axial resolution.

The hardware is not available to approximate these requirements. The most significant problem is that point-source like transducers are difficult to produce; the size of the transducer must be small to avoid directional effects, yet a large amount of power must be passed through to obtain a reasonable amplitude wave. The amplitude of the wave is particularly critical in 3D compared to 2D because the energy spreads in an extra dimension, so the amplitude reduces as $1/r$ rather than $1/\sqrt{r}$ where r is the distance from the transducer. As a result of this compromise, to obtain a sufficiently strong signal, the transducers must be quite directional.

This chapter discusses how to adjust the algorithm for a realistic transducer array and applies it to a model representing typical structures within the breast. The model used to test the algorithm here is chosen to satisfy the Born approximation; this chapter therefore does not apply the full HARBUT formulation to the 3D data. HARBUT is avoided because this helps to simplify the analysis, although it would be easy to incorporate it at a later stage.

4.2 Theory

As explained in Chapter 2, the scattering problem under the Born approximation can be expressed as

$$\psi \approx \psi_u - \int_{\Omega} G_u O \psi_u d\mathbf{x}'. \quad (4.1)$$




Figure 4.2: Comparison of 2D and 3D k-space vectors. In 2D, as shown in (a), there are only two source-receiver vector combinations which correspond to a particular point in the k-space. However, in 3D, as shown in (b), there is an infinite number.

Born and Wolf [31] explain that, under the far field approximation, each measurement corresponds to a Fourier component of the object function¹

$$\psi_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}}) \approx -\tilde{O}[k_u(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})] \quad (4.2)$$

so to reconstruct the object function in 3D, the combination of send and receive vectors needs to populate the Ewald sphere. This map of the measurements in the spatial frequency domain is referred to as the ‘k-space’ through this thesis.

One theoretical approach would be to position a sphere of transducers completely around the scatterer, allowing illumination and reception to be in any direction which would completely fill up the Ewald sphere. Such an array is impractical, but also contains significant amounts of redundant data.

Figure 4.2 shows this redundancy. In the 2D case of Fig. 4.2(a), each point in the k-space corresponds to the mapping of just two sets of wave vectors; this pair corresponds to swapping the source and receiver, which is known to result in the

¹A fuller derivation is given in Sec. 6.2.1.




Figure 4.3: Mapping of measured components into the k -space Ewald sphere with the array at a fixed axial location. (a) shows a schematic diagram of scattering with the array at a particular axial location, offset from the scatterer. (b) shows how the measurements at each axial location map into a single disc of the Ewald sphere.

same value by the principle of reciprocity. In 3D, however, as shown in Fig. 4.2(b), there is an infinite continuum of source-receiver vector sets which reach the same point in the k -space. Only a single set of these vectors is needed to obtain the required value. It is therefore unnecessary to have a complete array like this to fill the Ewald sphere.

As Simonetti shows in [14], it is possible to completely fill the Ewald sphere with just a single ring array moved axially. Figure 4.3 illustrates how, at each axial location, the ring array populates a disc of the sphere; by changing the axial position the k_z component will change and the sphere can be filled. Simonetti proposed using the BF to DT approach introduced in [34] to reconstruct the data obtained from this configuration; the filter used, however, requires data that is sampled at uniform angles of elevation rather than uniform linear intervals.

While obtaining such data was straightforward for the semi-analytical model used in the paper to test the algorithm, it is not possible for practical applications of breast ultrasound tomography. However, the solution is relatively straightforward; as discussed in Appendix B, the correction in the filter from linear sampling to uni-

form angular sampling is achieved by multiplying by $\sec^2 \alpha$, where α is the elevation angle.

Another problem is the directivity of the transducers. Figure 4.4 plots the out-of-plane variation of the beam from transducers with dimensions of 4mm, 12mm and 20mm. Ideally for 3D beamforming, to obtain maximum axial resolution the beam energy would be spread equally in all directions, which is clearly not the case for the examples here. One solution is to compensate for the amplitude drop by multiplying by the reciprocal of the beam amplitude when performing beamforming, in exactly the same way the amplitudes of the Green's functions are accounted for in the formulation of Chapter 2. The problem with this approach is that the beams contain points at which the amplitude reaches zero; compensating in this way would therefore cause singularities. Also, the amplitude variations of the transducers are unknown – the models here use a simple ‘piston’ type arrangement where the transducer behaves as a line of uniform phase and amplitude sources, but in reality a number of practical uncertainties about the transducer behaviour will make estimating the beam to any usable degree of accuracy impossible.

The approach used here is to simply ignore the amplitude variations caused by the transducers in the out-of-plane direction, treating the data as if it had been obtained by point transducers. As the amplitude of the beam drops at larger elevation angles, the corresponding k-space components will also be reduced. These components are the higher spatial frequencies in the axial direction, so reducing their amplitude is, in effect, applying a low pass filter. This filtering is what limits the axial resolution for this method.

It is possible to estimate this. I define a 6dB drop as the limit of the usable k-space values. This corresponds to a drop of 3dB for both the source and receiver (considering a point equidistant from both, such as lying on the central axis of the array), which under the far field, Fraunhofer approximation [68], occurs at an angle (in radians) of

$$\theta_{lim} = \frac{0.884\lambda}{h} \tag{4.3}$$

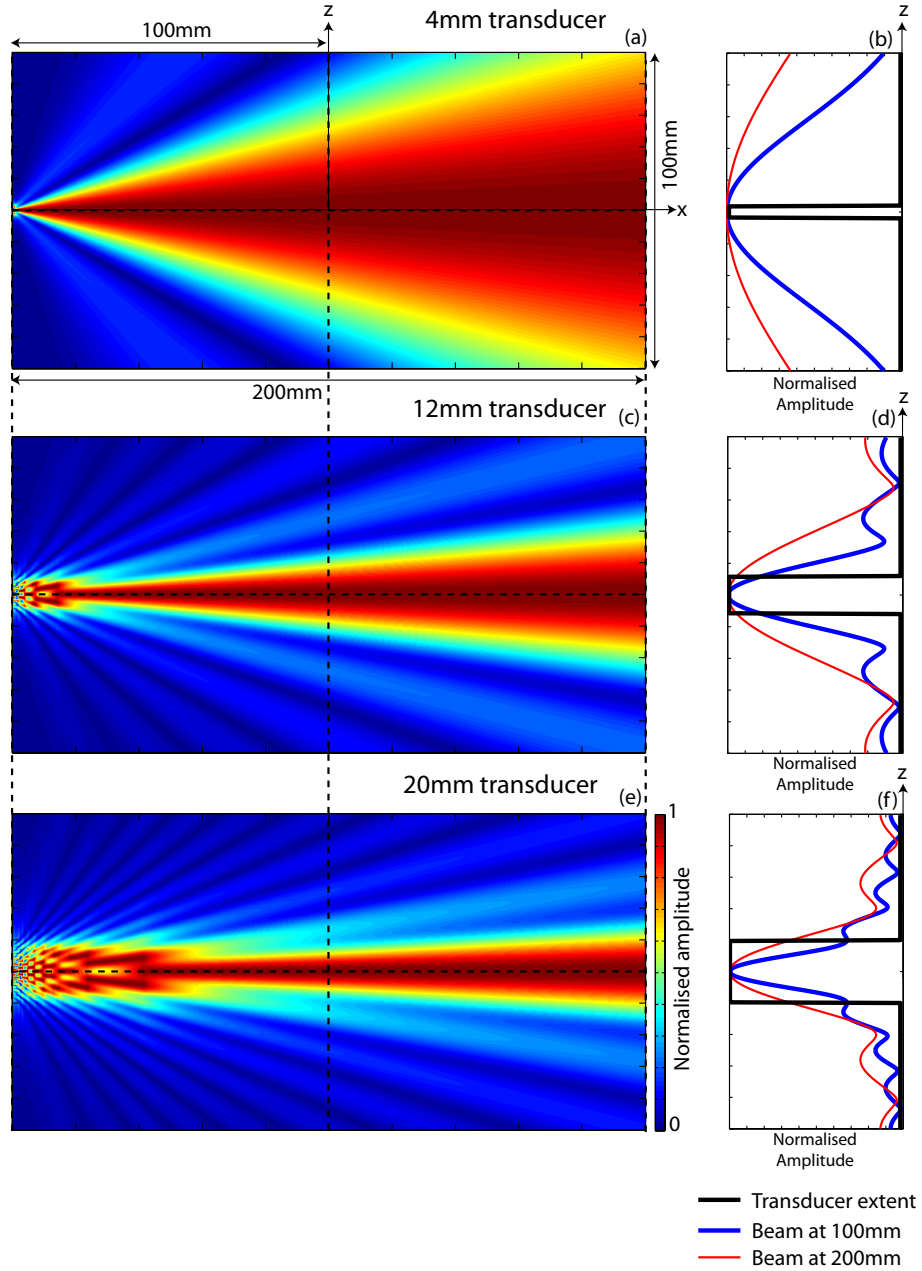


Figure 4.4: Transducer beam variation out of the plane of the array. Shown is the amplitude of beams travelling through a distance of 200mm (the diameter of a typical array) for (a) a 4mm transducer, (c) a 12mm transducer and (e) a 20mm transducer. Normalisation is performed against the peak value within each horizontal array. (b), (c) and (f) plot the amplitudes at a distance of 100mm (i.e. the centre of the array) and 200mm (the opposite side of the array) for the 4mm, 12mm and 20mm transducers respectively.

where h is the transducer height. The corresponding limit of the spatial frequency components in the axial direction is

$$k_{z,lim} = 2k_u \sin \theta_{lim}. \quad (4.4)$$

This corresponds to a resolution limit of

$$\frac{\lambda}{2 \sin \theta_{lim}} \approx 0.57h \quad (4.5)$$

where the small angle approximation of $\sin \theta \approx \theta$ has been used. Therefore, using this 3D framework leads to a resolution of roughly half the transducer height. In the case of the 12mm transducer of Chapter 3, a 7mm resolution should be possible, a slight improvement over the 9mm demonstrated under the 2D framework. However, by selecting smaller transducers better suited to the 3D technique it will be possible to improve this.

It is interesting to note that when reconstructing under the 2D model, the beam should be as thin as possible to minimise slice thickness and improve resolution, but under the 3D model, the beam should be as wide as possible for maximum resolution. One possible solution to exploit this for improved reconstructions under the 3D reconstruction method is to spread the beam wider through a divergent lens placed on each transducer.

4.3 Model

A significant challenge in testing breast imaging algorithms is to find a suitable model. The model needs to be simple enough that judgements can be easily made about the performance of the algorithm yet sufficiently complex to represent the structures that comprise the breast. Given that the model is 3D, another challenge is how the model can actually be generated.

For this investigation a ‘tree’ type model was chosen. This can represent the complex

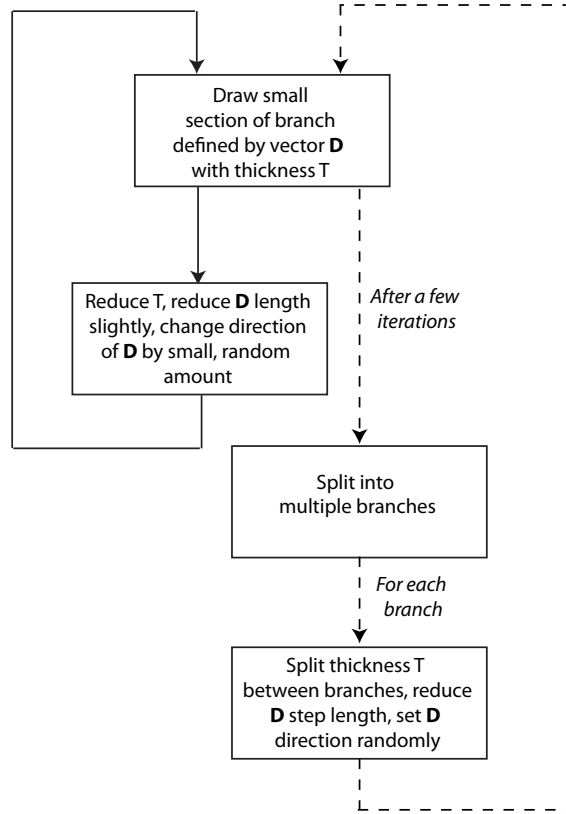


Figure 4.5: Algorithm for generating a tree-like model. Branches are drawn by stepping along, randomly varying direction and reducing thickness and step length as the algorithm progresses. After a certain number of iterations, each branch splits into several smaller branches and the process repeats for each of these.

features common in the breast such as milk ducts, lymphatic vessels and blood vessels. Such models can be generated fairly easily by computers; an outline of how the algorithm I have developed works is presented in Fig. 4.5. Branches are drawn as a series of short straight lines stepping a certain distance, each time adjusting the direction by a random amount. After a certain number of steps the branch is split, reducing the thickness and the step length by a certain amount and giving each new branch a different direction. This process repeats, with each new branch following the same path, until the branches have thinned sufficiently that they are inconsequential. The algorithm can be performed easily with a recursive function.

There are a large number of parameters involved in this process, for example the number of branches to split into (commonly two or three), or the random variation in direction at each step. In the 1980s Viennot et al. [71] demonstrated how a wide range of tree shapes could be generated by varying a few simple parameters. While

this study was done in 2D, the results easily generalise to 3D.

This model was used to generate a tree structure; the parameters were varied until a good solution was found where the branches were fairly well separated and clear to identify yet complex enough to still be representative of actual structures within the breast. Since judging the algorithm's performance will require the image to be viewed in 2D, it was decided to restrict the branches to lie within a single vertical plane in the centre of the array. Since it is necessary for the image to be reconstructed within the Born approximation, the material contrast was chosen to be low, with 1% increases in both velocity and density within the branches compared to the background.

The transducer array had 250 elements and a radius of 40mm; this is smaller than normal to reduce the computational strain. The transducers were 4mm tall line sources; these are as small as practically possible in order to obtain a good resolution. The array was positioned at 72 axial locations separated by 1mm.

A slice of 80mm by 80mm by 40mm was modelled, being meshed with $501 \times 501 \times 261$ nodes, spaced to give 12 nodes per 2mm wavelength. The simulation step time used a Courant number of 0.95 and the total time was the time for the wavefield in the water background to pass across the array multiplied by 1.5. A three cycle toneburst with a centre frequency of 750kHz was used to excite each transducer.

Since each illumination needed to be simulated separately, a total of 18 000 fully 3D simulations had to be modelled. A cluster of 40 computers of eight cores each was used for this purpose; around 64 simulations were run simultaneously on this system. Each individual simulation took around four hours, and the whole set took around eight weeks to complete.

4.4 Results

Figure 4.6 presents the results from this model. Figure 4.6(a) shows the original model, generated by the algorithm explained in the previous section and illustrated

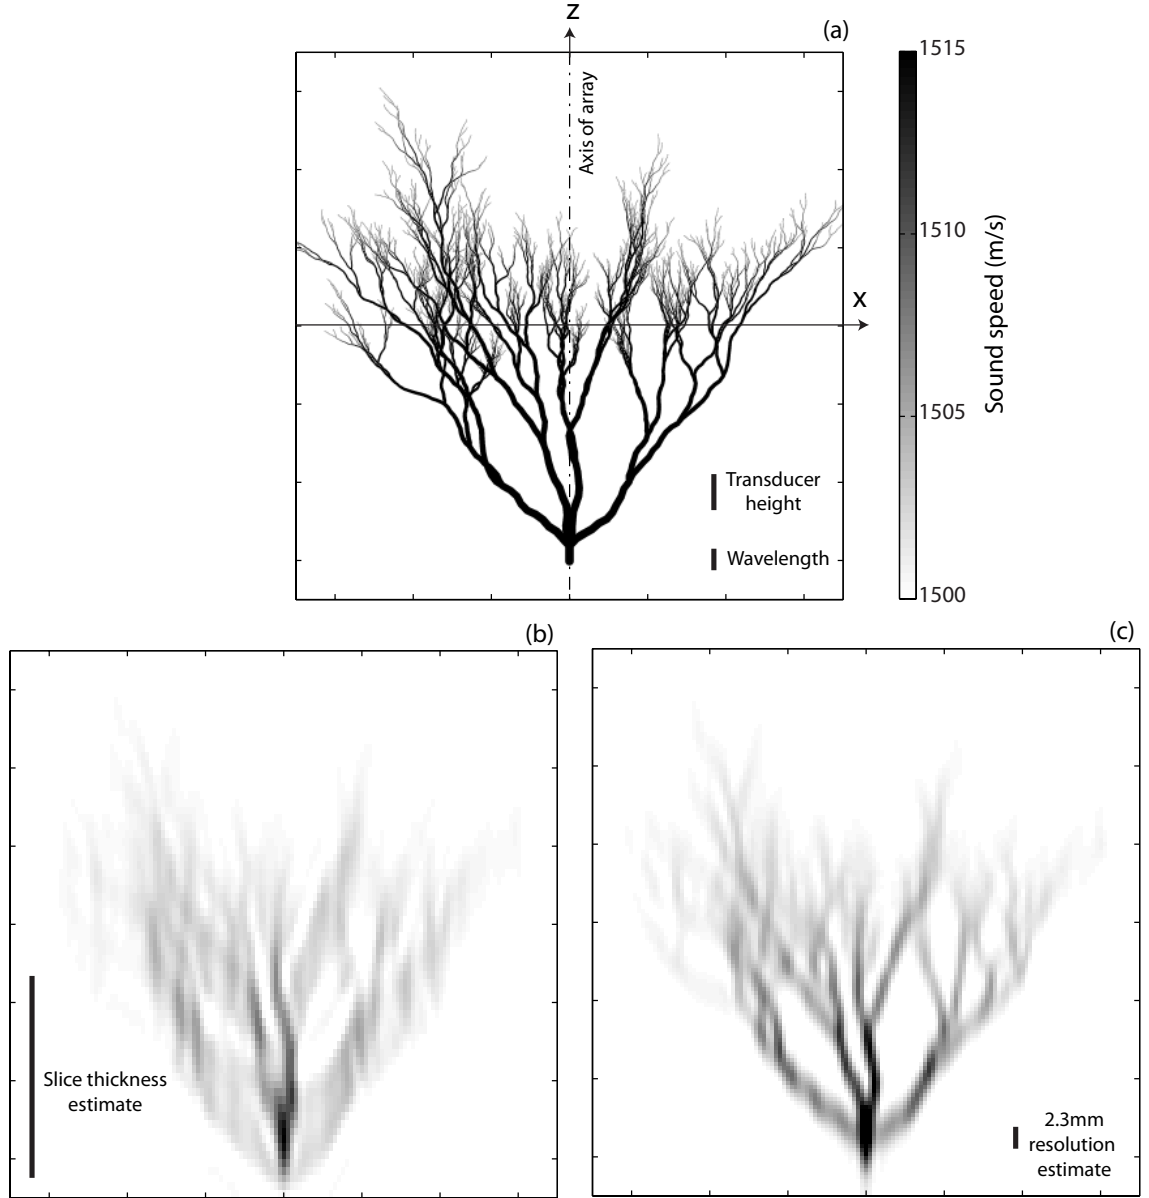


Figure 4.6: Reconstructions of a tree-like model. (a) gives the original model, generated by the method in Fig. 4.5. (b) presents the results obtained by treating each set of data separately under the 2D framework then stacking the results. (c) presents the image reconstructed using all the data sets together using the 3D framework discussed in this chapter.

in Fig. 4.5. This is in the x - z plane, which is perpendicular to the plane of the array; the plane lies along a diameter of the array. Figure 4.6(b) shows the results obtained by the standard 2D technique used in Chapter 3, making the assumption that within each slice there is negligible z -direction variation. Images were generated at each z slice then these were stacked to form the 3D image. Figure 4.6(c) presents the results using the full 3D framework outlined in this chapter. In both the 3D

and 2D reconstructions, five frequencies were averaged together: 400, 450, 500, 550, 600kHz. The wavelength shown is that for the middle of these, 500kHz.

It is clear that the Fig. 4.6(c) reconstruction has significantly better axial resolution than Fig. 4.6(b), confirming that the use of the more accurate 3D representation can improve the reconstruction. The 3D framework obtains, through visual inspection, a resolution similar to the prediction of Eq. (4.5), 2.3mm.

The resolution of Fig. 4.6(b) is significantly better than the slice thickness estimate from Chapter 3. The approximations made there are suitable for fairly directional beams, which is clearly not the case for the 4mm transducers given the plot of Fig. 4.4(a). At large elevation angles, the phase shift illustrated in Fig. 4.1 is likely to have a significant effect, which could be the cause of the resolution not being as low as would be expected from the slice thickness.

One drawback of the improved accuracy of the 3D model is the speed. As shown in (2.20), beamforming consists of performing a multiplication, then summing the results (to approximate the integral) for all transducers. In 2D, this is just done for the transducers in the plane. In the 3D method explored here, the number of operations is then multiplied by the number of slices. In this case this reduces the speed of the algorithm by a factor of 72 compared to generating a series of separate 2D images and stacking. This can be rationalised to an extent if only the transducer ring positions which actually illuminated each point were included in the calculation, but the 3D method is still expected to be at least 20 times slower than the 2D case. A C++ implementation of the algorithm took around 9 minutes to perform beamforming for a single frequency of the reconstruction in Fig. 4.6(c), although this could be improved through optimisation.

4.5 Summary

This chapter has explored synthetic aperture imaging to improve the axial resolution of breast ultrasound tomography. Rather than assuming each slice to be quasi-

uniform in the axial direction, as in the study of Chapter 3, here I have exploited the way additional information is encoded in the data as the array is moved axially. This imaging method uses the standard beamforming technique, but incorporating data from all the slices at each point.

Under this fully 3D framework, the resolution was calculated to be around 0.57 times the height of the transducer. For 4mm tall transducers, imaging a complex tree-type phantom representing some of the complex structures within the breast, the 3D framework was shown to give a significant resolution improvement over the standard 2D imaging method.

Having considered this 3D problem, the next chapter returns to the 2D framework with an ideal 2D problem: guided wave tomography.

Chapter 5

Guided Wave Tomography

5.1 Introduction

Corrosion is a significant problem in the petrochemical industry. Large, often inaccessible areas must be inspected to establish the extent of the damage caused. This is slow and difficult using conventional thickness-gauging methods which can only measure thickness values directly beneath the probe. Low-frequency guided waves have been used in Non-Destructive Testing (NDT) as a solution to long-range inspection of pipes and other structures, which allows large areas to be scanned without needing direct access [72–75]. However, while detection and location determination from these methods is good, they can generally only give rough estimates of remaining wall thickness.

Guided wave tomography has been proposed as a solution to accurately estimate the remaining wall thicknesses of corrosion patches in plates and large diameter pipes. By utilising a guided wave mode which is dispersive, if frequency is fixed then the wave velocity will become a function of thickness. One approach to guided wave tomography is therefore to reconstruct the velocity map from a series of ultrasonic measurements, then convert this back to thickness [23, 76–78].

To establish the minimum wall thickness, an accurate, reliable reconstruction method

of sufficient resolution must be used. Most of the work done in guided wave tomography uses the assumption that waves propagate as straight rays, ignoring both refraction and diffraction [23, 24, 76, 79]. By ignoring refraction, the algorithm is limited to low contrast defects, and ignoring diffraction limits the algorithm to large, smoothly varying objects; both of these manifest themselves as reducing the resolution of the reconstructed thickness map. The resolution limit is similar in scale to the size of the first Fresnel zone, $\sqrt{L\lambda}$ where L is the length from the source to the receiver and λ is the centre-frequency wavelength [25]. Slight resolution improvements are possible by accounting for refraction in bent-ray tomography [78, 80], but to increase resolution further it is necessary to account for diffraction.

The Diffraction Tomography (DT) method was proposed by Malyaranko and Hinder [80] to improve the resolution in guided wave tomography, and was investigated by Belanger et al. [77] for some simple defects. DT is, however, limited by the use of the Born approximation.

The Born approximation considers the field scattered by a defect to be a superposition of the scattered fields from many elemental scatterers. Under the approximation, these elemental scatterers behave independently, simplifying the scattering problem and making inversion straightforward. However, in reality, as the wavefield passes through a scatterer its phase will be distorted, meaning that the elemental scatterers no longer behave independently. The Born approximation is therefore only valid for defects where the phase distortion is low, such as with low contrast or limited defect size [19, 31].

In Chapter 2, HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography, was introduced, which uses the low resolution bent ray tomography (BRT) algorithm as the background for DT. This allows the phase distortion through the scatterer to be estimated and accounted for, which greatly increases the range of applicability of the DT algorithm. The algorithm was shown to provide a fast, robust and high-resolution reconstruction of sound-speed through the breast, although imaging a 3D problem using a 2D framework did cause resolution problems in the axial direction.

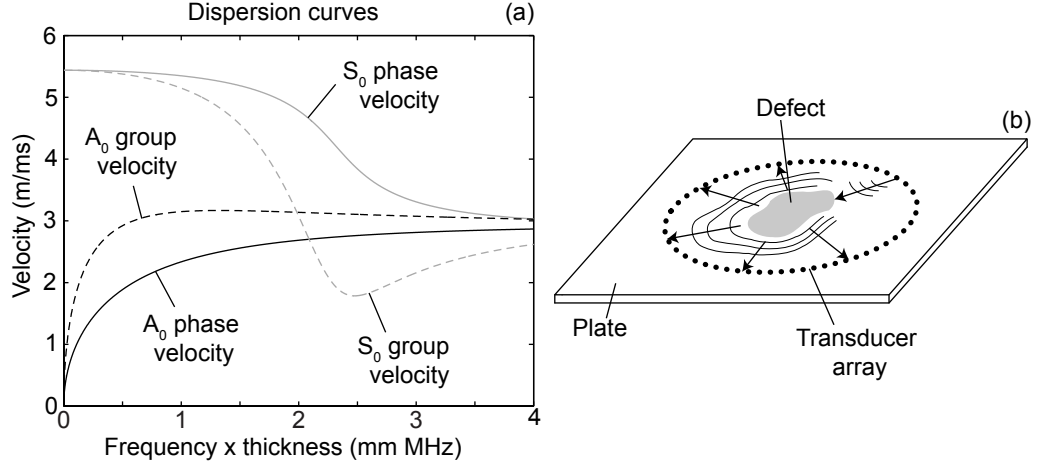


Figure 5.1: (a) shows the dispersion curves for the fundamental Lamb wave modes in aluminium, for both phase and group velocity. The relationship between thickness and velocity at a given frequency provides a means to determine thickness. (b) shows the guided wave tomography configuration considered. Waves are excited by a single source, hit the defect and are scattered in all directions. The scattered field is measured from all directions by the transducer array. The process is repeated with each transducer in turn providing the illumination, until data for all possible send and receive combinations has been obtained.

This chapter investigates the suitability of HARBUT as a practical solution to estimate the thickness of a defect in a plate or large pipe from a circular array of transducers. GWT is an ideal 2D scattering problem, so any issues with 3D resolution discussed in the previous two chapters are avoided. HARBUT's performance is established with a range of simulated and experimental data, and modifications are made to the algorithm to make it more suitable the deep, high contrast defects of interest.

5.2 Method

Figure 5.1(a) illustrates the dispersion curves for Lamb waves in an aluminium plate. These show phase and group velocity as a function of the product of frequency and thickness, providing a mapping between thickness and velocity for a given frequency. By assuming that no mode-conversion occurs at the boundaries of the defect (a reasonable assumption given that thickness variation associated with corrosion is likely to be continuous and slowly varying) the scattering of the Lamb waves can be

simplified to the acoustic wave equation [29] in 2D. Given the frequency dependence of the velocity, this is expressed for each monochromatic component of the wavefield

$$\nabla^2 \psi(\mathbf{r}, \omega) + k(\mathbf{r}, \omega)^2 \psi(\mathbf{r}, \omega) = 0 \quad (5.1)$$

where $\psi(\mathbf{r}, \omega)$ is the scalar potential of the field (equal to a single component of the Fourier transform of the pressure) at point \mathbf{r} , ω is the angular frequency and $k(\mathbf{r}, \omega) = \omega/c(\mathbf{r}, \omega)$ is the wavenumber associated with the phase velocity c , which is a function of thickness.

From this formulation, the scattering of Lamb waves at a single frequency behaves in the same way as acoustic scattering, so any imaging algorithm developed for acoustic scattered data, including HARBUT, can be used to image guided wave data. However, there are a few differences between the HARBUT implementation used for breast imaging and for GWT.

Firstly, GWT is a purely 2D problem, unlike the breast tomography studied in the last chapter, conveniently avoiding the out-of-plane resolution issues. Secondly, bent-ray tomography, using the arrival times of the waves, will produce a reconstruction of group velocity through the plate, whereas the single frequency components used in the DT stage of HARBUT will correspond to phase velocity. It is therefore necessary to convert from group velocity for BRT to phase velocity for use as a background in DT. This is straightforward to do from the dispersion curves. For consistency, in all images throughout this chapter thickness is plotted, having been converted from group or phase velocity as appropriate.

Finally, GWT needs to be able to accurately determine very deep defects; these can correspond to a much higher velocity contrast than appears in breast ultrasound tomography, although their size relative to the wavelength is typically much smaller. To better reconstruct these defects, an iterative HARBUT approach is introduced.

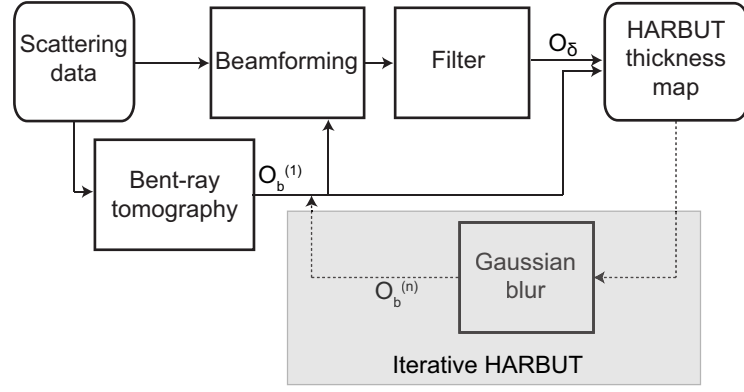


Figure 5.2: Flowchart showing the stages of iterative HARBUT.

5.2.1 Iterative HARBUT

For many applications of guided wave tomography, such as large diameter oil pipe corrosion, it is important to detect and quantify large corrosion depths of up to 80% of the wall thickness. This large thickness variation corresponds to a large velocity contrast, and it is this which causes the Born approximation to break, despite the defect's size typically being just a few wavelengths. Such defects are sufficiently small that they are unlikely to be well captured by the low resolution bent-ray tomography method, which poses a problem since HARBUT relies on having a sufficiently accurate background reconstruction.

The approach considered here is to recognise that, while it may not produce a completely accurate reconstruction because of the Born approximation not being satisfied particularly well, HARBUT will still provide an improvement over the bent-ray tomography reconstruction. This improvement is exploited by iterating HARBUT – taking an existing HARBUT reconstruction and using it as the background for another HARBUT stage in place of bent-ray tomography, as illustrated in Fig. 5.2. At each step, O_b becomes more accurate, minimising O_δ and allowing HARBUT to produce more accurate velocity maps. By iterating HARBUT many times its range of applicability is extended to make it more suitable for small, high contrast defects.

As shown in Fig. 5.2, a Gaussian filter is applied to smooth the background before the next iteration, which is a form of regularisation. This filter aims to remove as many of the artefacts from each iteration as possible, while maintaining the true

reconstruction values. Experience suggests standard deviation radii of around λ to $\lambda/2$ for the Gaussian filter are good for this purpose; as the iterations progress and the artefacts reduce this could, in theory, be reduced or even completely eliminated, however this is not something explored in this thesis.

5.2.2 Convergence criterion

There is a need to establish a suitable criterion to stop the iteration process. A parameter, Q , is therefore defined as the average thickness change between iterations over the damage area, relative to the background thickness

$$Q^{(n)} = \frac{\int |T^{(n)}(\mathbf{r}) - T^{(n-1)}(\mathbf{r})| S(\mathbf{r}) d\mathbf{r}}{T_u \int S(\mathbf{r}) d\mathbf{r}} \quad (5.2)$$

where T_u is the uniform background thickness, $T^{(n)}(\mathbf{r})$ is the thickness at position \mathbf{r} for iteration n and S is a function defining the damage area. $T^{(1)}$ is the thickness field from the bent-ray tomography algorithm. S must be defined; this is done by taking any reconstructed depth greater than 5% of the background thickness as indicating there is a defect present, so

$$S(\mathbf{r}) = \begin{cases} 1 & \frac{T_u - T^{(n)}(\mathbf{r})}{T_u} \geq 0.05 \\ 0 & \frac{T_u - T^{(n)}(\mathbf{r})}{T_u} < 0.05 \end{cases}. \quad (5.3)$$

It is not expected to obtain a thickness estimate better than about 1% of the background thickness; a stopping criterion can be based around this. One solution would be to stop when $Q \leq 0.01$, based on this being a change less than the thickness ‘resolution’. However, I choose a stricter limit, since (a) an average defect thickness change is being considered, which might mask improvements in one area of a complex defect, and (b) the cumulative effects of multiple iterations with small average changes can sum to more than 1%. Based particularly on (b) and that, from experience, I expect 10 or fewer iterations, I adjust the criterion to

$$Q \leq 10^{-3}.$$

5.3 Numerical results

5.3.1 Numerical model

To test the algorithm, a set of simulations were performed. The models were based on the plate in Fig. 5.1(b), with absorbing boundaries placed around the edge to remove edge reflections. The models were run using the commercially available Abaqus Finite Element (FE) package. The 10mm thick aluminium plate ($E = 70.8\text{GPa}$, $\nu = 0.34$, $\rho = 2700\text{kg/m}^3$) was modelled as 10 layers of 1mm thick 3D stress elements. In the in-plane directions 32 elements per wavelength at the centre frequency were used. To model the thinning of the plate at a defect, elements were removed from the surface; the element sizes were kept uniform. This allowed discrete thickness depths spaced 1mm apart to be modelled, which was sufficient to model the cases of interest.

An array of transducers was modelled by exciting and measuring nodes. Since the array is circular and the mesh is a regular grid of cubic elements, each transducer's location was rounded to the nearest node; given the refinement of the mesh this meant that the position was accurate to $\pm\lambda/64$ which will not cause significant problems. A 5-cycle Hann-windowed toneburst at 50kHz was used as a source, excited in the out-of-plane direction to produce nearly pure A_0 waves [81].

Abaqus Explicit was used, which steps in time to calculate the solution. Separate simulations were necessary for each illumination, meaning that a simulation had to be performed for each transducer to generate a full matrix of scattering data.

5.3.2 Data processing

First, the signals were processed to extract the arrival times. The signals were enveloped via a Hilbert transform, then the arrival time was taken as the point at which the enveloped signal intersected a threshold at 50% of each signal's maximum. The frequency domain data needed for the beamforming and DT stages of HARBUT

were obtained by applying a fast Fourier transform to the time-trace from each send-receive pair then extracting the frequency component(s) required.

As with the breast imaging in Chapter 3, it is important to calibrate the system to account for the uncertainties which arise in transmitting and measuring signals. While all aspects of the transduction system can theoretically be modelled to establish the cumulative effect on phase and amplitude (see for example [82]), there are sufficient uncertainties that it is usually more efficient to calibrate the system using measured data. This is true even when processing simulated data, since exciting and measuring the wavefield often introduces a complex factor into the measurements.

Belanger et al. [77] performed their calibration at the imaging stage, by applying a factor to make the DT image match the known feature depth. This factor was then reused for all similarly obtained data sets. This relies on the assumption that a perfectly calibrated image would capture the feature depth correctly. In the examples explored in the paper, however, the Born approximation was very marginal due to the high contrast of the defects. The DT-reconstructed defect depths were therefore inaccurate, making the calibration constant from this method inaccurate for general reconstructions.

The approach used in this chapter is to compare the full matrix of signals at a specific frequency to that of a theoretical incident field. This incident field is excited by a point source; the field is then given by the Green's function

$$G(r) = -\frac{i}{4}H_0^{(1)}(k_u r) \quad (5.4)$$

where r is the distance between the particular source and receiver. The target is to calibrate the measured signals such that if there was no scatterer present (i.e. the waves just passed through a homogeneous medium) they would equal the theoretical incident field of (5.4). It is rare for a set of data from the homogeneous background to be available; however, since any corrosion defect is typically confined in space, it is usually possible to visually identify large sections of the matrix where the line-of-sight between the send-receive pairs does not intersect the corrosion patch.

These sections will be relatively unaffected by the scattered field, so can be used to determine the calibration factor needed for the data. This approach was used to calibrate all the simulated and scattered data in this chapter.

5.3.3 Reconstructions

The first model produced was an axisymmetric circular defect in the centre of a 64 transducer circular array. This had a thickness reduction of 50% (i.e. 5mm remaining thickness) and a diameter of 60mm ($\approx 1.5\lambda$ at 50kHz). The original model is illustrated in Fig. 5.3(a).

Figures 5.3(b) to (f) outline the stages of HARBUT for this defect. Figure 5.3(b) presents the bent-ray tomography image of the defect. Given that the defect is only 60mm in diameter and the resolution can be estimated as $\sqrt{L\lambda} \approx 174\text{mm}$ from the width of the Fresnel zone [25], it is clear that the reason for the poor reconstruction is primarily the limited resolution. The bent-ray image is used as a background for beamforming (Fig. 5.3(c)) which is filtered to give the perturbation object function, O_δ in Fig. 5.3(d). Combining this with the background bent-ray tomography image gives the total object function in Fig. 5.3(e). This is then converted to a thickness map in Fig. 5.3(f), using the phase velocity dispersion curve.

Also, HARBUT was iterated, stopping at five iterations since $Q^{(5)} = 7 \times 10^{-4}$, less than the threshold of 10^{-3} . Figure 5.4 shows the cross section for the iterated HARBUT reconstruction along a line through the centre of the defect, comparing it to bent-ray tomography, standard DT and HARBUT. The centre of the DT reconstruction has been inverted. This is a result of the Born approximation ignoring the distortion of the illuminating waves. The bent-ray tomography reconstruction underestimates the depth reduction because the resolution is not high enough to capture the relatively small defect. By combining these two algorithms together in HARBUT, a much better estimate of the thickness is obtained. The minimum thickness taken from the HARBUT reconstruction is 5.5mm, only 0.5mm away from the true thickness. Iterated HARBUT reduces the reconstructed thickness estimate




Figure 5.3: The HARBUT algorithm applied to a single central defect. The defect is circular, 60mm in diameter, and has a thickness of 50% of the plate itself (i.e. 5mm remaining thickness) - see the original map in (a). The 50kHz frequency component was used in the beamforming and DT-based reconstructions. The λ and $\sqrt{\lambda L}$ scales shown are 37.7mm and 174mm respectively. (b) shows the bent-ray tomography reconstruction, which is low resolution so cannot capture much of the detail. This is used to correct the beamforming algorithm in (c), which, when filtered produces the O_δ of (d). This is combined with the background object function to give the total object function O in (e). Converting this to thickness via the dispersion curves is shown in (f).

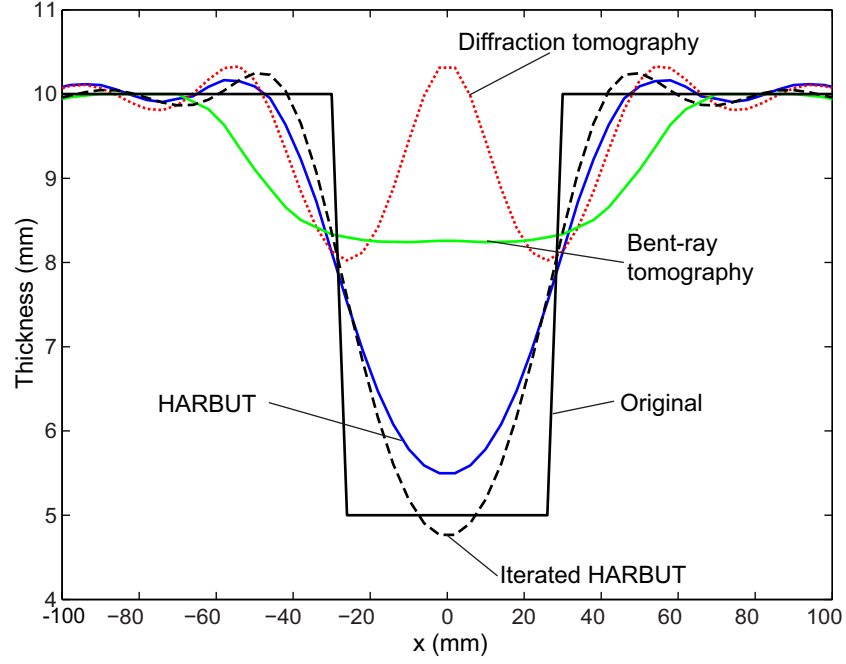


Figure 5.4: Cross sections of the thickness reconstructions from simulated data for a single circular defect.

to 4.7mm, overshooting the true value but reducing the error slightly.

Since the DT reconstruction is effectively a low-pass filtered version of the image, even a perfect DT reconstruction of a square function would be expected to produce an overshoot due to the missing high-frequency components. It could be that this is the cause of the defect depth overestimate seen here.

A more complex model, with two defects, was tested. The initial 60mm, 50% thickness defect was maintained at the centre, and an additional 30% thickness reduction (i.e. 7mm remaining) circular defect of diameter 100mm was added, offset by 200mm. This is shown in Fig. 5.5(a).

Figures 5.5(b) to (d) show the reconstructions of this model. Once again it can be seen that the bent-ray tomography method, Fig. 5.5(b), has poor resolution, failing to capture the size or depths of the defects accurately. There are also significant artefacts around the boundaries of the defects; these are caused by diffraction in the forward model, which is ignored by the imaging algorithm. The HARBUT method of Fig. 5.5(c) significantly improves the resolution of the reconstruction, however

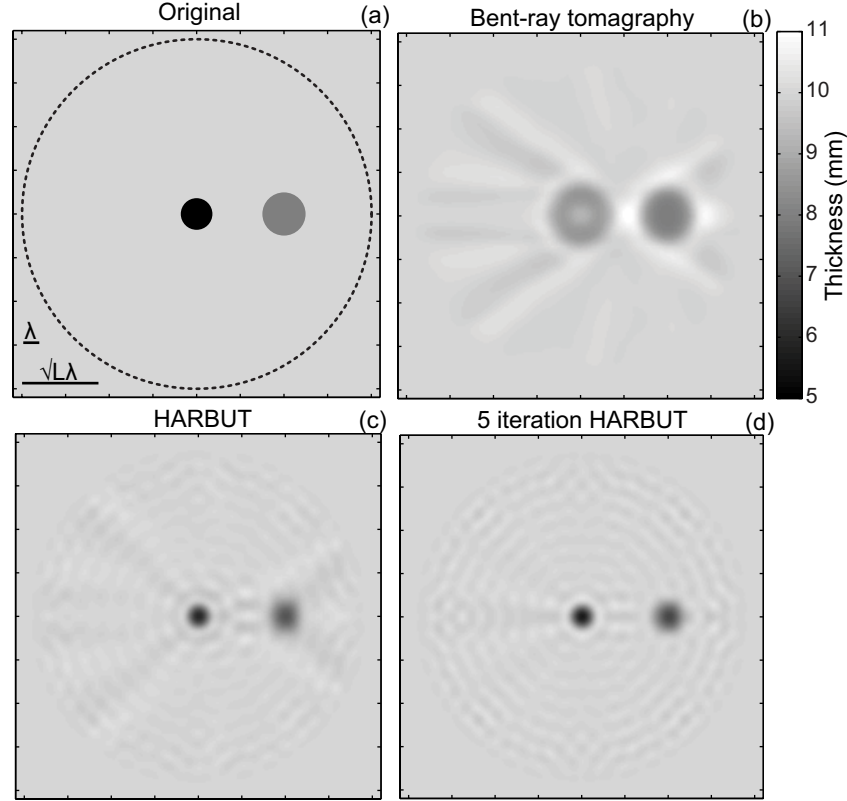


Figure 5.5: Reconstructions of a two defect model. The central defect is 60mm in diameter and 5mm thick, and the offset defect is 100mm in diameter and 7mm thick, as shown in the original map of (a). (b) shows the bent-ray tomography image and (c) gives the HARBUT image at 50kHz. (d) applies 5 iterations of HARBUT to the reconstruction.

some of the artefacts from the bent-ray tomography background do appear faintly. Figure 5.5(d) shows that iterating HARBUT has removed the remaining bent-ray tomography artefacts from the image. The resulting reconstruction is cleaner than standard HARBUT.

Figure 5.6 presents the reconstructed cross sections along the line intersecting the centres of both features. Multiple iteration HARBUT shows a slight improvement over standard HARBUT for the central inclusion, but the offset inclusion reconstruction overshoots the true value slightly. Both thickness estimates are within 0.7mm of the true values.

So far I have considered models of simple geometric shapes. While this simplicity makes judging the new algorithm's performance easier, it is important to ensure the algorithm works with more realistic cases. Figure 5.7(a) shows a more realistic

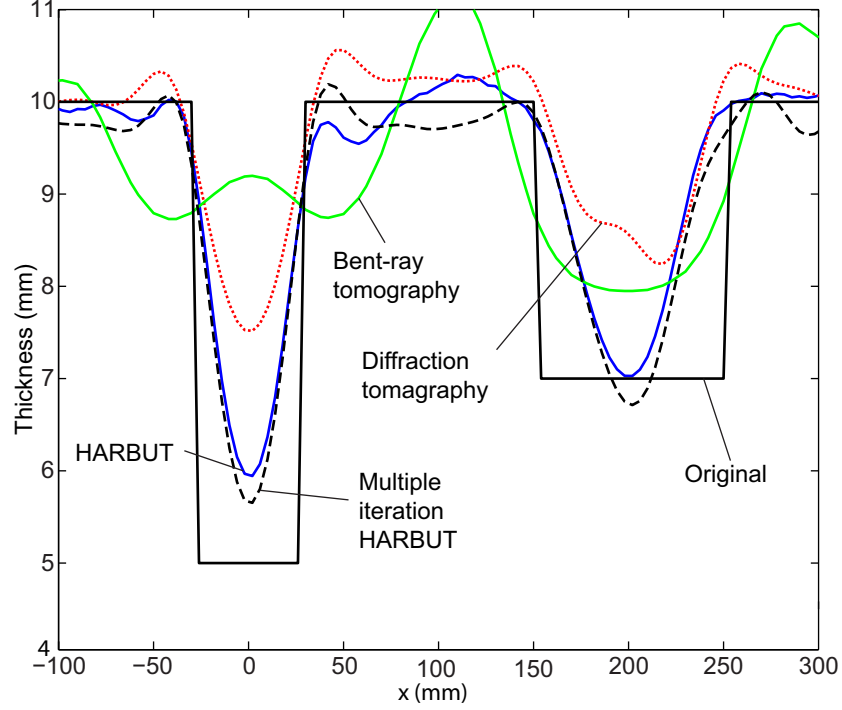


Figure 5.6: Cross sections of the reconstructions from simulated data for a plate with two circular defects.

defect. It is a random defect with thicknesses varying between the 10mm background and 5mm. The defect is challenging because of its complex shape and the large sound-speed contrast associated with the thickness variation.

The standard diffraction tomography reconstruction in Fig. 5.7(b) demonstrates clearly that the defect breaks the Born approximation. It is not possible to obtain any meaningful depth information from this data. The bent-ray tomography reconstruction in Fig. 5.7(c) is better, producing a reasonable overall estimate of the defect thinning, but failing to capture any of the details.

Figure 5.7(d) presents the HARBUT image, using Fig. 5.7(c) as a background to improve Fig. 5.7(b). There are artefacts which result from the Born approximation breaking, which indicates that O_δ is too big, meaning that the bent-ray tomography background is not accurate enough. However, despite this, some of the features of the true thickness map are clearly visible in this reconstruction.

I applied iterative HARBUT to the data from this model. At each stage a Gaussian filter of $\lambda/2$ (18mm) standard deviation radius was used to blur the background.

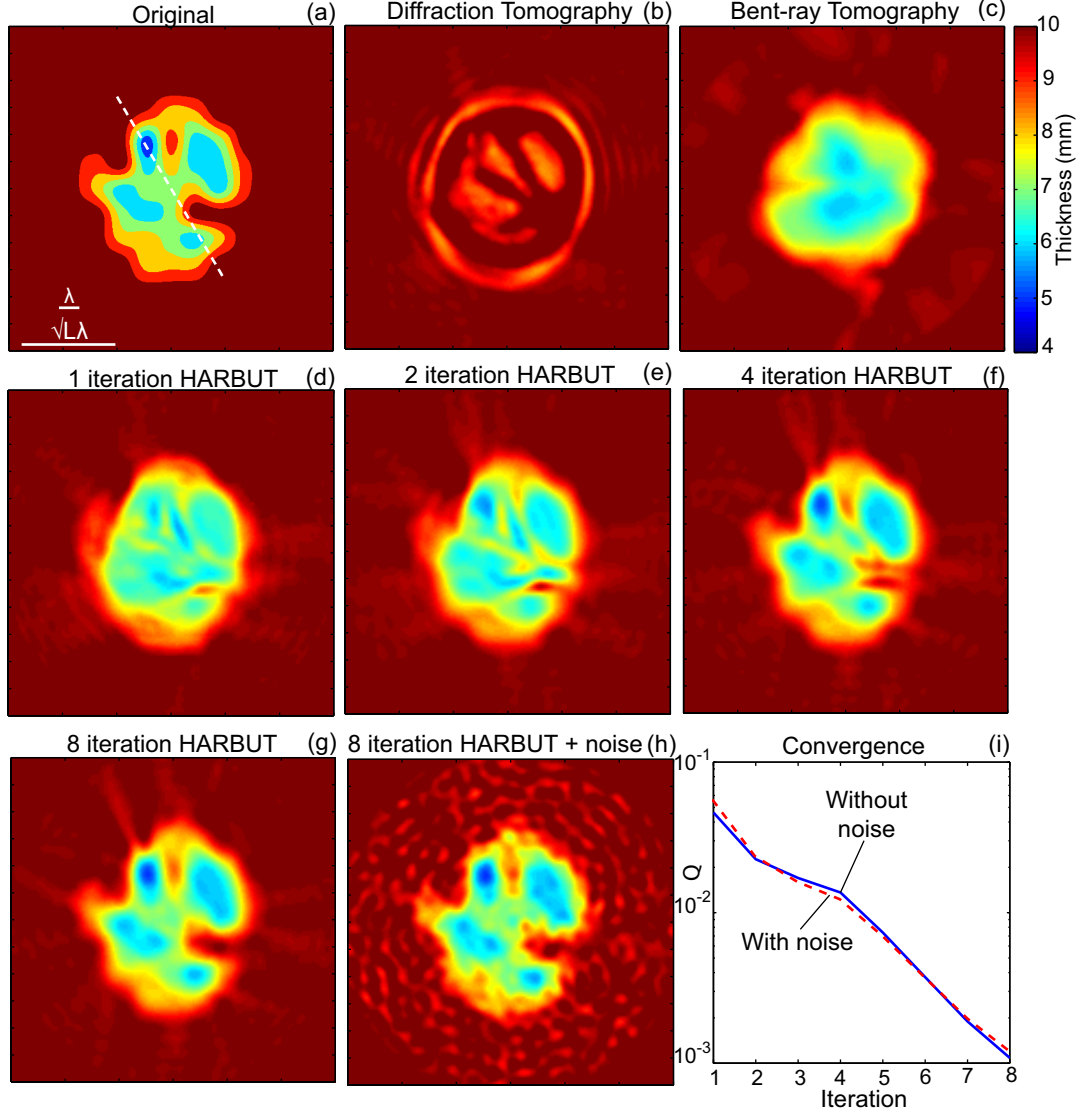


Figure 5.7: Reconstructions of a complex corrosion patch. (a) is the original thickness map. (b) is the standard DT reconstruction, and (c) is the bent-ray tomography reconstruction. (d) gives the standard HARBUT reconstruction, showing an improvement over both the bent-ray and DT. (e), (f), (g) give the results after 2, 4 and 8 iterations. (h) gives the reconstruction after 8 iterations with artificial noise added to the data. (i) presents the convergence parameter Q .

Figures 5.7(e), (f), (g) show iteration 2, 4 and 8. The image meets the criterion outlined in Sec. 5.2.2 by 8 iterations as shown in Fig. 5.7(i).

The reconstruction of Fig. 5.7(g) captures well the features of the original model, Fig. 5.7(a). This confirms that the reason for errors in the standard HARBUT reconstruction was that the initial background, O_b , failed to capture all the detail well enough, so O_δ broke the Born approximation.

The non-iterative HARBUT method, having at its core the beamforming algorithm, has been shown to be robust in the presence of noise [83]. Iterative methods, however, typically perform poorly when noise is present, often attempting to ‘overfit’ the solution to match the errors in the data. I have tested how the presence of noise affects the convergence of iterative HARBUT by adding a noise matrix to the input data prior to imaging. The noise matrix has a normally distributed amplitude with RMS 50% that of the original matrix, and phase uniformly distributed from 0 to 2π . This corresponds to a very poor signal to noise ratio of 6dB.

Figure 5.7(h) gives the reconstruction of thickness with eight iterations in the presence of noise. It is clear that, while the noise does have an inevitable effect on the image, there is no evidence of overfitting or convergence to an incorrect solution. It is possible to see all the features of the defect and obtain a good estimate of the minimum thickness.

Figure 5.7(i) plots the parameter $Q^{(n)}$ for each iteration. Also plotted is the $Q^{(n)}$ line for the case with noise. This provides additional verification that the algorithm converges well even when significant noise is present. By 8 iterations the parameter has reached the 10^{-3} level for both sets of data.

Figure 5.8 compares the cross sections of 1 and 8 iteration HARBUT, bent-ray tomography and the original along the dotted line in Fig. 5.7(a). Iterating HARBUT shows a significant improvement over bent-ray tomography and standard HARBUT. The figure also verifies the ability of the iterated method to establish the minimum thickness accurately; in this case the 5mm minimum thickness is estimated as 5.1mm.

5.4 Experimental models and results

5.4.1 Experimental setup

The experimental configuration was based around Fig. 5.1(c), with an aluminium plate of $1200 \times 1200 \times 10$ mm. One significant difference between the experiments




Figure 5.8: Cross sections of reconstructions of the complex defect along the line marked in Fig. 5.7(a).

and the FE simulations is the boundary conditions. The FE mesh is terminated in absorbing boundaries to minimise any reflections; the experiment cannot include such boundaries so the measured signals will include boundary reflections. For signals transmitted through the scatterer, it is possible to simply gate off the reflections since these arrive later, but the signals reflected from the scatterer will be corrupted by the boundary reflections. Therefore just the transmission subset is available for imaging, which will result in a loss of resolution in the reconstructed image [31].

A single A_0 source transducer was used, as described in [77]. This provides an illuminating wavefield, which is then measured around the transducer array by a Polytec OFV-505 laser Doppler vibrometer as in [84]. A 64 transducer array of 400mm radius was modelled, by moving the source to the first position and taking measurements with the laser around each receiver location in the array, then repeating the process to build up a full matrix of data containing all the send-receive combinations. A 5-cycle Hann-windowed toneburst at 50kHz was used as the input signal.




Figure 5.9: Experimental time trace from an axisymmetric central defect. Taking the defect as the origin, the source was at (400mm, 0) and the measurement was taken at (-324mm, 235mm) – corresponding to transducer 24 of 64. Noise was removed from the time trace by multiplying by the gating function. The gating function was defined as 1 between t_{arr} (the arrival time, $248\mu s$ in this case), and $t_{arr} + 120\mu s$. Outside this region the function smoothly returned to zero via a sine curve over a time of $40\mu s$.

5.4.2 Data processing

Arrival times were extracted in the same way as in Sec. 5.3.1. Whereas the relatively noiseless simulated data was directly Fourier transformed to obtain frequency domain data, the experimental data was gated beforehand to remove as much of the experimental noise as possible. Inevitably unwanted modes will be excited and measured, but since data is obtained from a regime where these travel at different speeds, gating helps to remove these components.

The gating is illustrated in Fig. 5.9 for a single experimental time trace. The signal start time was taken as the measured arrival time, and the end time $120\mu s$ after. A gating function was set to one between these times. Outside the region, the gating function reduces from one to zero via a sine function over a time of $40\mu s$. At all other times the gating function is zero. This gating function is multiplied by the signal to remove unwanted noise.

A fast Fourier transform is then applied and the 50kHz signal component extracted.




Figure 5.10: Imaging a single central defect, diameter 60mm, depth 50%, as in Fig. 5.3(a). (a) gives bent-ray tomography, (b) shows diffraction tomography. (c) is a HARBUT reconstruction, and in (d) four HARBUT iterations have been applied.

As explained in Sec. 5.3.1, these frequency domain signals are then calibrated by matching them to the field from a point source.

5.4.3 Experimental results

Figure 5.10 presents the images for a single defect at the centre, with the same parameters as the FE model in Fig. 5.3: 60mm diameter, 50% (5mm) depth. Since the model is axisymmetric, receiver measurements were only taken around half the array for a single source and were then copied to generate all the send-receive combinations.

In Figs. 5.10(a) and 5.10(b), bent-ray tomography and diffraction tomography both underestimate corrosion depth. As discussed in Sec. 5.3.1 this is because of the limited resolution in bent-ray tomography and the Born approximation limit for




Figure 5.11: Cross sections of the reconstructions from experimental data for a single circular defect.

DT.

As shown in Fig. 5.11 and Fig. 5.10(c) HARBUT shows an improved estimate of the defect depth, around 6mm thickness, 1mm out from the true depth. Four HARBUT iterations do not provide a significantly better estimate of thickness, as shown in Fig. 5.10(d). Both HARBUT reconstructions (and the multiple-iteration version in particular) show noticeable concentric ring artefacts, however.

A likely cause is errors in the measured data. Normally these errors would be incoherent and would therefore vary between the send-receive pairs, and when imaged they would be minimised through the summing process of the beamforming algorithm. In this example, however, by reconstructing the entire 64×64 matrix of send-receive pairs by repeating just 33 measurements, there is much less error cancellation. In this case, any error in each signal would be expected to form an axisymmetric ring, which is the sort of artefact seen in the reconstruction of Fig. 5.10(d).

To verify that this is the source of the errors, a non-axisymmetric model is considered. As with the FE models, a second defect of depth 30% (i.e. 7mm remaining)

radius 100mm, offset from the centre by 200mm, is added. In this case measurements from the full 64×64 send-receive combinations are required in order to reconstruct the image, so the ring artefacts visible in the axisymmetric case should not be seen.

Figure 5.12 gives the reconstructions of the thickness from various methods. The bent-ray tomography and DT methods both fail to provide accurate depth information because of the resolution and Born limitations explained previously. Single iteration HARBUT provides a significant improvement in the depth information. With multiple iterations of HARBUT, improvements in the reconstruction are seen, with the artefacts relating to the Born approximation breaking (for example between the two defects) being reduced.

It should be stressed that while bent-ray tomography uses the broadband nature of the signal to estimate the arrival time, the beamforming and DT steps in HARBUT are monochromatic, i.e. they use the information contained in a single frequency of the signal spectrum. In particular, Figs. 5.12(b)-(d) are obtained at 50kHz. Belanger et al. [77] explored the idea of using all the data from the bandwidth by averaging together images from many frequency components. This averaging can help to minimise the effects of noise. However, when using an iterative method such as HARBUT, one must decide on a strategy for how to use the available frequency components at each iteration.

A typical approach (see for example [85]) for using multiple frequencies with iterative methods is to begin with the low frequency components (which are lower resolution but more robust) and move up to the higher frequency components (which provide higher resolution details). In the case considered here, only a relatively narrow bandwidth is available (reasonable images can be generated from around 30-60kHz) which limits the possible gains from starting at low frequencies. Instead, the monochromatic image of Fig. 5.12(d) is taken as a background for a series of HARBUT images with all the available frequencies, which are then averaged together to give the final reconstruction. The advantage of this, rather than, for example, using every frequency component at every iteration, is that the computation time is significantly reduced – images for all frequencies only need to be generated at the

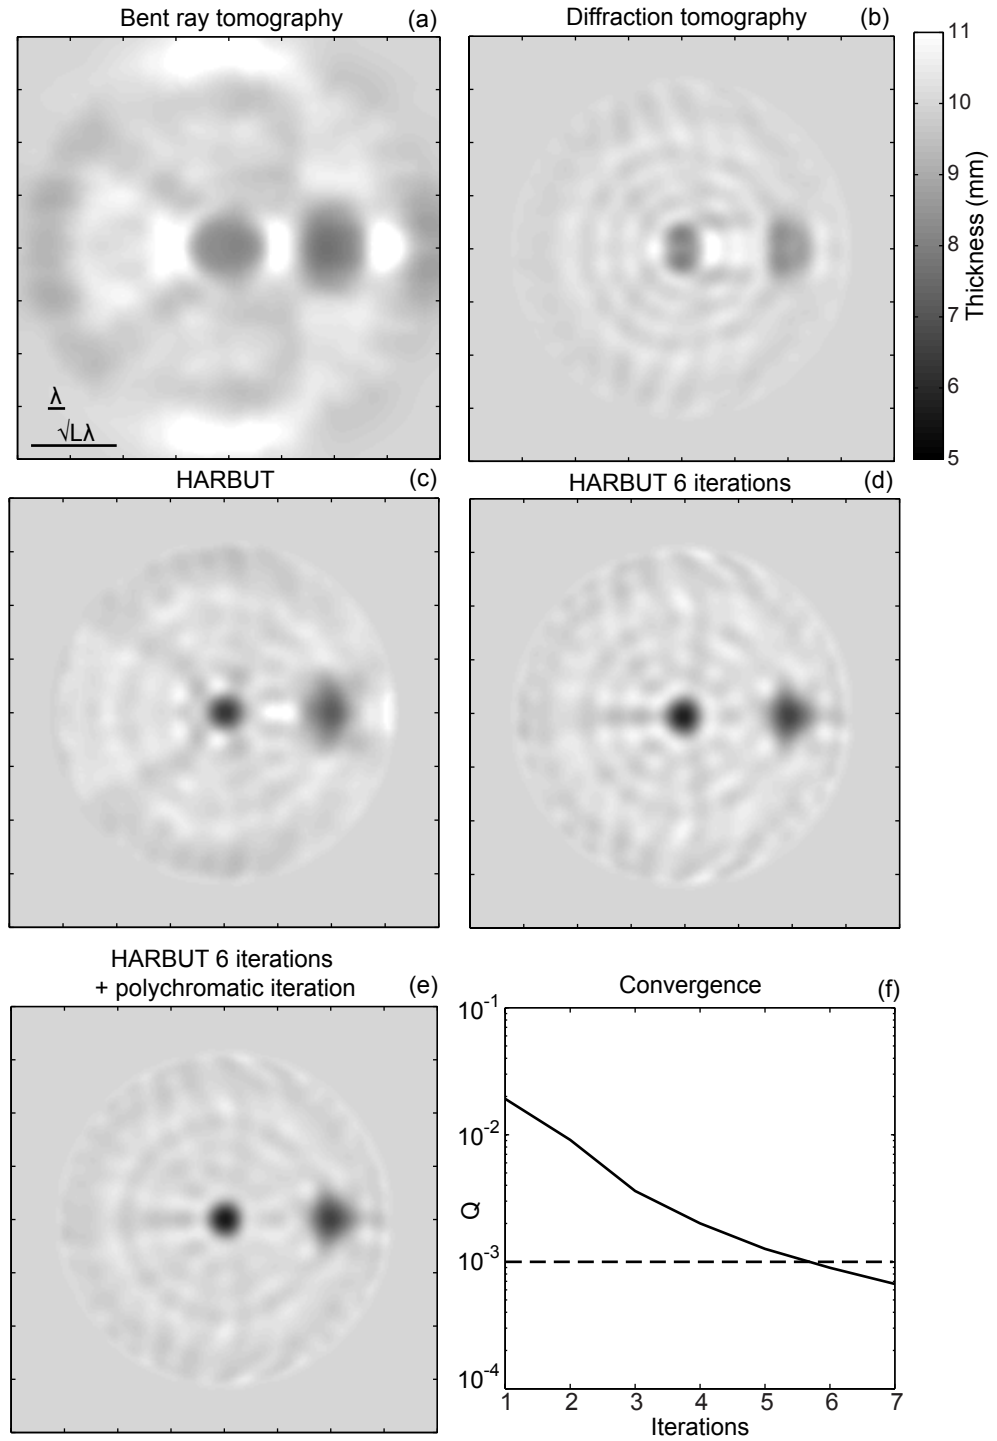


Figure 5.12: Reconstructions from experimental data for a two defect model. The central defect is 60mm diameter and 50% (5mm) deep, and the offset defect is 100mm diameter and is 30% deep, as shown in Fig. 5.5(a). (a) is the bent-ray tomography method, (b) is diffraction tomography, (c) is single iteration HARBUT, and (d) is HARBUT with 6 iterations. In (e), I have taken (d) as a background, and averaged together HARBUT images from all the frequency components between 30 and 60kHz. This makes better use of the bandwidth to minimise noise in the data. (f) plots the convergence parameter Q .




Figure 5.13: Cross sections of the reconstructions from experimental data for two circular defects in an aluminium plate.

final step. The resulting image has visibly fewer artefacts than the monochromatic images, although the defect reconstructions themselves are not significantly closer to their true values.

Figure 5.13 shows the cross sections of the images in Fig. 5.12. It confirms that bent-ray tomography and DT underestimate the defect depths, while HARBUT gets a better estimate. From iterated HARBUT, the defect thicknesses of 5mm and 7mm are reconstructed as 5.6mm and 6.6mm respectively, i.e. both are within 0.6mm of the true values.

As explained by Belanger et al. [77], the 64-transducer array is undersampling the field according to the criterion of [64], which leads to a grating lobe at a radius of 192mm. This grating lobe is visible in Figs. 5.12(b)-(e), and intersects the 70% defect. This is likely to have an effect on the thickness estimate of the feature, which could be the cause of the defect depth overestimate in Fig. 5.13. With the exception of this undersampling artefact, there are no concentric rings as seen for the axisymmetric case, confirming that these were caused by the process of copying data.

The convergence parameter is plotted in Fig. 5.12(f), once again showing a monotonically decreasing function. This confirms that, even with experimental data, iterative HARBUT remains stable and well-behaved. In this case, 6 iterations were needed to reach the convergence criterion.

The HARBUT algorithm is very fast. The bent-ray tomography method can be completed in less than 10 seconds, and each HARBUT stage for a single frequency with the 64 transducer array took around 5 seconds for a 210×210 grid on an HP Z600 8-core workstation. The latter is using the Matlab implementation; as discussed in Chapter 3 for breast imaging a C++ version of the algorithm runs about 10 times faster. Even when using multiple iterations and averaging many frequencies, the algorithm will be able to produce images within a practical time-frame.

5.5 Summary

I have introduced HARBUT (the Hybrid Algorithm for Robust Breast Ultrasound Tomography) to the field of guided wave tomography. The algorithm aims to reconstruct the sound speed through a plate, which is linked, via the dispersion curves, to the thickness; this provides a means to accurately map the residual thickness of a plate-like structure.

In its original formulation, HARBUT provided a means to overcome the Born approximation, by using bent-ray tomography to generate a low resolution background estimate, which then corrects the Born approximation used in diffraction tomography. HARBUT therefore relied on the background estimate being sufficiently accurate, which is difficult for the relatively small, high contrast defects expected in guided wave tomography.

Here I have introduced iterative HARBUT to address the problem of high contrast. By taking the standard HARBUT reconstruction as a background then repeatedly applying HARBUT to step towards the solution, the algorithm becomes more suit-

able for dealing with small, high contrast features. Accuracy was improved and artefacts minimised by using the iterative HARBUT method; it was shown that minimum depths could be estimated to within 1mm for a 10mm thick plate for a range of different defects for both numerical and simulated data. HARBUT avoids the background field subtraction that is normally necessary for diffraction tomography methods, and which is known to be a significant source of errors due to the instability of baseline signals.

Having demonstrated HARBUT's performance with an ideal 2D imaging problem, with a circular, full view array, in the next chapter I begin investigating how it performs under less ideal circumstances, with non-circular arrays providing limited viewing angles.

Chapter 6

Limited view quantitative imaging

6.1 Introduction

This thesis has so far considered data sets taken from a full view array – an array which completely surrounds the object and allows illuminations and measurements from all possible angles. In many cases, however, only a limited area is available around a component on which to place transducers, which greatly limits the range of viewing angles and hence the amount of information which can be retrieved.

In a large number of ultrasonic limited view applications, the source and receiver locations will be the same, so it is only possible to measure waves which have been reflected back towards the sources. Such reflection data is often used in structural imaging to reconstruct sharp impedance changes (see for example in medical beamforming [8] and most NDT applications [5]), but since quantitative imaging is frequently interested instead in low contrast, often smoothly varying features, very little reflection occurs and instead the information is encoded in the transmission components, which will not be directly measurable by the array.

Given the difficulty of placing transducers under the ground, this limited view problem is of great interest in geophysical imaging. Mora [86] attempted to quantitatively image beneath the earth's surface using a linear array on the surface; the key to his

algorithm was to consider the reflection from a sharp discontinuity (such as a layer of rock). Waves travelling down through the earth will encode quantitative information as they travel, and the discontinuity, by subsequently reflecting these waves back to the array, provides a method to retrieve this information without having to place receivers beneath the surface.

Mora proposes an iterative algorithm to utilise this principle to reconstruct a quantitative image in the presence of a large discontinuity. This iterative algorithm works by obtaining a reconstruction of the locations of the reflective scatterers via standard reflection imaging with an inhomogeneous velocity background, then, with this knowledge, recognises that the waves reflected off these scatterers will encode transmission information about the medium, using that to update the velocity background. This process is repeated until convergence.

As with many iterative methods, there is a significant problem with convergence to local minima; the assumed initial velocity background must therefore be sufficiently close to the true one for the algorithm to converge. There is also a requirement that the array must be large enough that ambiguity in scatterer position due to velocity errors is limited; with a larger array, more angles are obtained and a better estimate of the position is possible when the background velocity is uncertain. However, in many cases it is impractical to have such a large array.

A more recent, very similar example which uses the reflection from such a discontinuity to aid quantitative reconstructions is the work of Natterer [87]. He too, however, recognises that the initial velocity field must be close to the true field, and his images all use a large array in order to obtain sufficient data. Natterer applies his algorithm to ultrasonic imaging of a simplified numerical breast phantom from a linear array with a backwall reflector, obtaining good estimates of sound speed. However, there are significant practical uncertainties associated with this problem, as discussed in Chapter 3, which are not addressed, and there is no evidence that the algorithm works with experimental data.

A more representative study of limited view ultrasound breast imaging, which does

use experimental data of a simple phantom, was performed by Huang [88,89]. Similar to previous approaches, a single linear array was placed on the surface and the transmitted components were reflected back from a parallel backwall reflector. Quantitative imaging was achieved by using reflection imaging to obtain the outlines of several inclusions, which are then used to segment the image. The signals are then gated so only the components which reflect off the backwall are included; this data contains the transmission information needed for quantitative reconstruction. A bent ray tomography (BRT) reconstruction method uses this data to reconstruct velocity, helped by specifying constant velocity through the segments defined from reflection imaging. The significant drawback to such a method, however, is that the breast is unlikely to be made of discrete, separable inclusions in this way – certainly there is no clear way that the in vivo data reconstruction in Fig. 3.15 could be segmented. If the features cannot be segmented then the method is largely redundant.

Even though a backwall reflector can provide useful transmission information about a scatterer, particularly if its position is known, a finite length array will still only be able to illuminate from a limited number of angles, restricting the amount of information available. Often, these reconstructions can be improved through the use of *a priori* assumptions. Examples include using wavelet transforms [90], maximum entropy methods [91], restricting to positive contrast [92], and minimising total variation [93,94], but it is often challenging to incorporate these methods into the algorithm.

This chapter investigates the effects that a limited view will have on the HARBUT method developed throughout this thesis. I primarily consider a configuration with a linear array opposite a parallel backwall, as studied by Huang; the data obtained from such an array (neglecting or gating out the weak direct reflection components) can be treated as if it had been obtained from a pair of parallel arrays. This is illustrated in Fig. 6.1(a) and (b).

Despite taking advantage of the backwall, some components are unavailable because of the finite length of the array. I develop a solution to improve the imaging algorithms for limited view configurations by taking advantage of *a priori* assumptions




Figure 6.1: Scattering from an object with a backwall reflector. (a) shows two scattering components. (b) shows how the backwall reflection component, expected to be much larger than any other component due to the low contrast, can be treated as though it was obtained from pair of parallel arrays.

about the component. However, while focused on this specific limited view problem, the techniques developed here are general and can be applied to a range of limited view configurations; as a demonstration of this, the method is applied to a limited view configuration for guided wave tomography.

6.2 The limited view problem and regularisation

This section outlines the effects a limited view configuration will have on the information available and how improvements to the image can be obtained by using regularisation.

6.2.1 Consequences of the limited view

To analyse the problem I begin from a Born model and assume that the transducers are in the far field. Consider the scattering integral under the Born approximation,

as introduced in Chapter 2

$$\psi_s \approx - \int_{\Omega} G_u O \psi_u d\mathbf{x}'. \quad (6.1)$$

As shown by Born and Wolf [31], with the far field assumption, the incident field and Green's functions become

$$\psi_u = e^{ik_u \hat{\mathbf{s}}_0 \cdot \mathbf{x}'} \quad (6.2)$$

$$G_u = e^{-ik_u \hat{\mathbf{s}} \cdot \mathbf{x}'} \quad (6.3)$$

where $\hat{\mathbf{s}}_0$ is the direction of the incident illumination and $\hat{\mathbf{s}}$ points towards the receiver. Therefore, (6.1) becomes

$$\psi_s \approx - \int_{\Omega} e^{-ik_u \hat{\mathbf{s}} \cdot \mathbf{x}'} O(\mathbf{x}') e^{ik_u \hat{\mathbf{s}}_0 \cdot \mathbf{x}'} d\mathbf{x}' \quad (6.4)$$

$$\approx - \int_{\Omega} e^{ik_u (\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}) \cdot \mathbf{x}'} O(\mathbf{x}') d\mathbf{x}' \quad (6.5)$$

which is effectively a Fourier transform of the object function O

$$\psi_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}}) \approx -\tilde{O}[k_u(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})]. \quad (6.6)$$

The signal from each send-receive pair corresponds to the Fourier component of the object function at $k_u(\hat{\mathbf{s}}_0 - \hat{\mathbf{s}})$. As shown in Figs. 6.2(a) and (b), by sweeping the illumination and reception angles through 2π radians, it is possible to fill a $2k_u$ sized disc, the Ewald limiting disc [31], thus allowing a resolution-limited version of the object function to be determined via the inverse Fourier transform.

If the illumination and reception angles are restricted because of the array configuration, only limited sections of the disc can be determined from the measurements. A case with a single linear array measuring in reflection is shown in Figs. 6.2(c) and (d). This reflection data clearly only contains high spatial frequencies, verifying

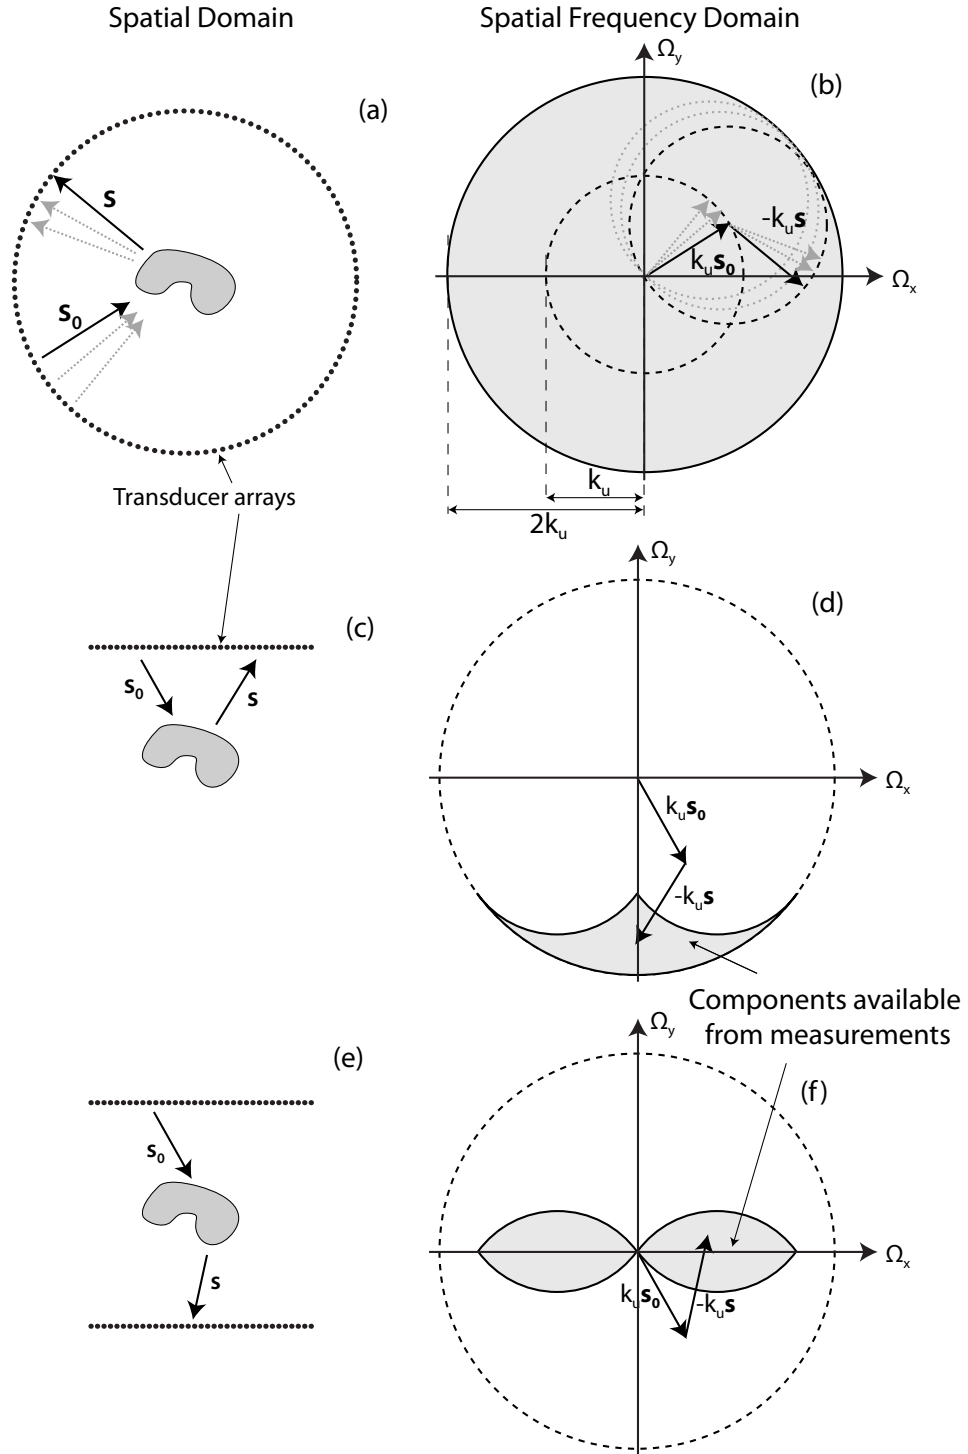


Figure 6.2: The mapping of far field measurements to the 2D Fourier transform of the scatterer under the Born approximation. (a), (c) and (e) show schematic diagrams of the transducer layouts and (b), (d) and (f) show the corresponding Fourier components which can be obtained from each layout. (a) and (b) show a full view array as considered in earlier chapters; (c) and (d) are for a linear array obtaining data in reflection and (e) and (f) show a pair of parallel linear arrays measuring transmission data.

that the reflection components are of little benefit to quantitative imaging, but can be used in structural imaging where only the outline of the structures are needed.

Finally, a case with a source array transmitting through the scatterer to the receiver array on the other side is shown in Figs. 6.2(e) and (f). Note that this configuration corresponds to the backwall reflection configuration being considered in this chapter, provided there is negligible direct reflection back from the scatterer to the array, as illustrated in Fig. 6.1. The data obtained fills an ∞ symbol in the k-space with known Fourier components.

The challenge of limited view imaging is to reconstruct an image when only a limited number of Fourier components, as in the examples here, are available.

6.2.2 Regularisation to obtain additional information

One solution to deal with the missing information is to simply perform a reconstruction with the unknown k-space components set to zero; in this way the image will simply correspond to a filtered version of the true image. This filtering will, however, introduce artefacts into the image. Sung et al. [92] showed how – in optical diffraction tomography, which approximates a far field problem – better reconstructions are possible by using a regularisation assumption to obtain more suitable values for the unknown k-space components. In that paper, the regularisation assumption is made that the scatterer has a refractive index greater than or equal to the background.

This assumption is applied via an iterative algorithm. Figure 6.3 presents an example of the process for a cylinder model. A set of known Fourier components is obtained directly from the data because of the far field assumption, and is assembled into the k-space image, as shown in Fig. 6.3(a). In this example a source array is above the scatterer and the receiver array is below; this is the same as the example of Fig. 6.2(e). Anything outside the known set is zero at first. This is inverse-Fourier-transformed to give the image (Fig. 6.3(b)), then regularisation is applied (negative values are set to zero) in Fig. 6.3(d). When subsequently forward-Fourier-transforming back to the k-space, as shown in Fig. 6.3(c), the unknown components

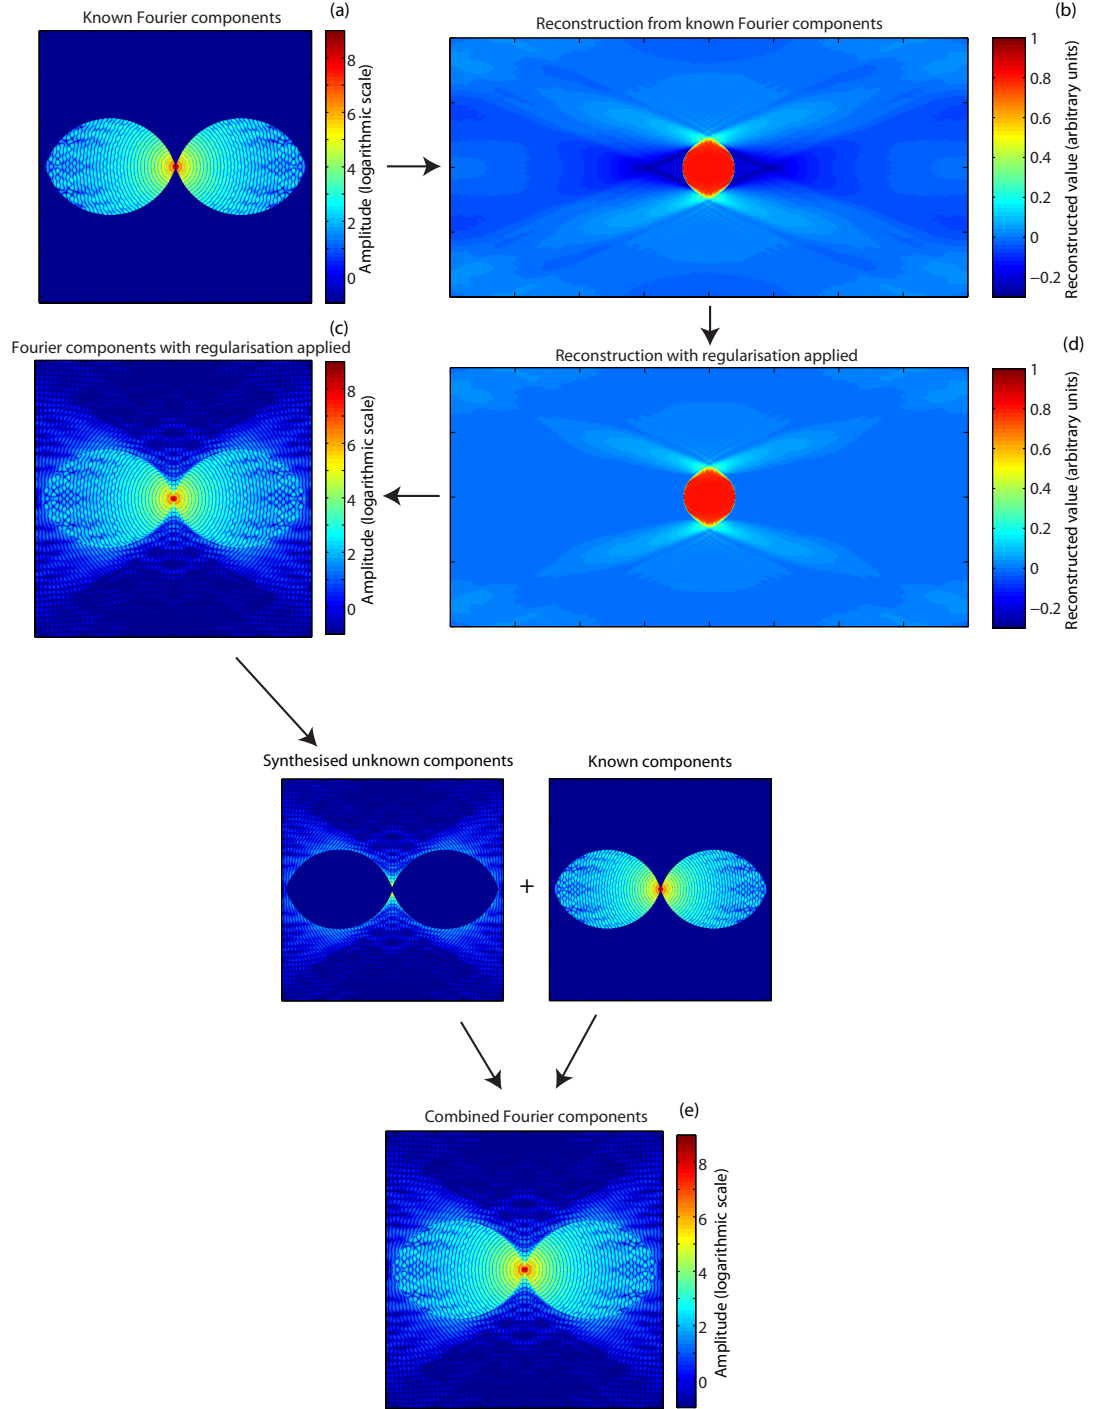


Figure 6.3: Schematic example of the approach used by Sung et al. [92] to obtain values for the unknown Fourier components via a regularisation method. The original model was a cylinder with an object function of 1, in a homogeneous background of object function 0. Fourier components from a limited view configuration with an array of sources above and receivers below are assembled in the k-space in (a). By Fourier transforming, the reconstruction can be obtained, as in (b), then regularisation is applied in (d), by setting any value less than zero to zero. (c) demonstrates that this regularisation process gives values to the unknown Fourier components. The estimates for these unknown values are combined with the known values in (e). By replacing (a) by (e), the iterative process is repeated until the unknown values converge such that they conform with both the regularisation method and the known values.

will now have non-zero values, which will better represent the knowledge about the object. The estimates of the unknown values are then combined with the known values in Fig. 6.3(e).

This process is iterated, with repeated application of the regularisation method then correcting the known Fourier components. Upon convergence, the image will contain all the known Fourier components and the unknown Fourier components will be set to the necessary values to best meet the regularisation requirement.

As Sung et al. [92] demonstrated, it is possible to obtain very good reconstructions of refractive index from limited view experiments with this method, provided data is obtained in the far field. However, the near field problem, where the data does not neatly map to Fourier components, needs an alternative technique.

6.3 A new approach to the limited view problem

So far, only the far field has been considered. This requires that the transducers are far enough away that the field incident on the scatterer must be approximately a plane wave, which is overly restrictive for the majority of ultrasonic imaging applications. Introduced here is a generalisation of the regularisation application method outlined above, making it suitable for a much wider range of configurations. It is stressed that the definition of ‘near field’ used here is simply any problem where the far field plane wave assumption cannot be used.

6.3.1 Virtual transducers

I begin, for simplicity, by considering the system using the Born approximation. To solve the problem in the near field, I introduce ‘virtual transducers’ which complete the full view. These virtual transducers are placed such that when combined with the existing transducers, the array surrounds the scatterer. In the problem being considered in this chapter, being treated as parallel arrays as in Fig. 6.1(b), the

virtual transducers join the ends of the two arrays, forming a full view rectangular array.

These virtual transducers perform the same role as the unknown Fourier components in the far-field case. Following this, the virtual transducer data is populated by repeated application of regularisation to the image, such that they ultimately conform to both the existing known data and to the regularisation.

It is necessary to have some form of forward modelling algorithm which can be used to populate the virtual transducers. In the far field case the forward modelling was implicitly performed via the Fourier transform, based on the Born approximation. Paralleling this, the method chosen here is the Born forward model of eq. (2.9), where the scatterer is split into many small component parts, each of which is assumed to behave independently. Under this, the field measured can be calculated by superposing the independent components together via an integral

$$\psi \approx \psi_u - \int_{\Omega} G_u O \psi_u d\mathbf{x}'. \quad (6.7)$$

This is fast, and, importantly, uses the same approximation as the inversion method. The integral is calculated by discretising it over a mesh with a spacing of less than $\lambda/4$; this is to avoid aliasing. Usually it is convenient to maintain the same grid for imaging and the forward Born integral to avoid resampling.

The theory explained so far relies on the Born approximation, which is only applicable to small, low contrast scatterers. This thesis has introduced HARBUT, the Hybrid Algorithm for Robust Breast Ultrasound Tomography, which extends the Born approximation by using an approximate inhomogeneous background to greatly increase the range of problems which can be solved. It is possible to generalise the virtual transducer method to include such a background. For this discussion I keep the background fixed; attention as later given to updating the background as performed in the GWT study of Chapter 5. The virtual transducer concept here is broadly similar to before; additional arrays of virtual transducers are introduced to complete a full-view array, and are populated using data from a regularised image.

There are two changes over the Born model. Firstly, the background needs to be accounted for when performing the imaging method – this follows the standard HARBUT method introduced in Chapter 2.

Secondly, the forward model used to populate the transducers needs to be adjusted to account for the background. Because of this, the integral used must be the inhomogeneous version, eq. (2.15),

$$\psi_\delta = \psi - \psi_b \approx - \int_{\Omega} G_b O_\delta \psi_b d\mathbf{x}'. \quad (6.8)$$

Note that ψ_δ is produced by the integral equation, which is convenient since this is the parameter required by the HARBUT imaging algorithm. By using the modified versions of the forward and inverse solvers as outlined here it is possible to account for a fixed inhomogeneous background and hence significantly increase the range of problems which can be imaged while using virtual transducers.

6.3.2 Storing virtual components in the image space

The forward model integral of eq. (6.8), for a fixed background (i.e. constant G_b and ψ_b terms) is linear. The same is true of the inhomogeneous beamforming integral of eq. (2.22) which forms the basis for the HARBUT inversion algorithm. This linearity provides an opportunity to greatly optimise the algorithm.

These processes can be expressed as operators. As in Chapter 3, I define a linear operator, \mathfrak{I} , that maps the measurements or data, d , onto an image, i , i.e.

$$i = \mathfrak{I}(d). \quad (6.9)$$

This operator can be used when a background is present, including the specific case when this background is homogeneous, which corresponds to the standard Born approximation. The first step is to image the limited view data ψ_{lim}

$$O_{lim} = \mathfrak{I}(\psi_{lim}) \quad (6.10)$$

which then has regularisation applied to give O_{reg} . I then introduce the forward operator \mathfrak{F} , which is the inverse of (6.9)

$$d = \mathfrak{F}(i). \quad (6.11)$$

Under the virtual transducer method, data for the full view array can be obtained using this operator on the regularised image

$$\psi_{reg,full} = \mathfrak{F}(O_{reg}). \quad (6.12)$$

The next stage is then to replace the appropriate sections of this data need with the true, known values. This can be expressed as

$$\psi_{full} = \psi_{reg,full} - \psi_{reg,lim} + \psi_{lim} \quad (6.13)$$

where $\psi_{virt,lim}$ corresponds to the components of $\psi_{reg,full}$ which overlap the known values in ψ_{lim} . Finally, the new object function value is obtained with

$$O_{full} = \mathfrak{I}(\psi_{full}) \quad (6.14)$$

which can then be used as the starting point for the next iteration. While the approach in the previous sections has attempted to explain this in a physical, intuitive way, expressing the process mathematically as I have done here presents opportunities for significant optimisation.

Equation (6.12) can be substituted into (6.13), and the result can be substituted into (6.14), giving

$$O_{full} = \mathfrak{I}[\mathfrak{F}(O_{reg}) - \psi_{reg,lim} + \psi_{lim}]. \quad (6.15)$$

Since the operators are linear, each term can be treated independently. The first term on the right hand side is simply a forward model then an inversion; these two

operators cancel and can be removed. The last stage has been calculated in (6.10) and can therefore also be simplified, which leads to

$$O_{full} = O_{reg} - \mathfrak{I}(\psi_{reg,lim}) + O_{lim}. \quad (6.16)$$

There are several things to note here. Clearly, removing the forward and inversion steps in the first term leads to the algorithm being much faster; it is also more robust since I am eliminating an imaging process which can introduce artefacts. The method does still require a forward model to be run, however, but only to populate $\psi_{reg,lim}$, which is much less computationally demanding. A second optimisation appears because the O_{lim} term only depends on the measured data (provided the background does not change if using HARBUT) so only needs to be calculated once for all iterations.

The virtual transducers are now redundant.¹ The information from the regularisation process, which was previously included in the virtual transducers, is stored in the image space, i.e. in the remainder $O_{reg} - \mathfrak{I}(\psi_{virt,lim})$. The method is therefore referred to as the Virtual Image Space Component Iterative Technique, or VISCIT.

6.3.3 Regularisation methods

VISCIT relies on a regularisation method to synthesise the unknown components; to emphasise the generality of this technique, little attention has so far been given to a suitable regularisation method. Here, I outline such a method.

As discussed earlier, there are a wide range of regularisation approaches which can be used to aid the reconstruction of images with limited view data. One of the most common is to assume that the reconstructed values must all lie either above, or below, the background value. This is appropriate in many practical examples,

¹The previous sections outlining the idea of virtual transducers are included for two reasons. Firstly they provide a convenient starting point for the explanation given here. Secondly, while I do not directly use virtual transducers in this thesis, the concept may well be of interest in other areas, such as when using non-linear algorithms which cannot take advantage of the simplifications I have described in this section.

such as guided wave tomography, where it is known that corrosion can only reduce the thickness, so the thickness throughout the domain must be less than or equal to the background.

Introduced here is a refinement of this. It is observed (see for example the image in Fig. 6.3(b)) that the peak contrast features of a limited view image are often well reconstructed, but that artefacts spread from these features which can swamp the lower contrast features. Figure 6.4(a) illustrates this. However, it should be possible to define a threshold which lies below the peak contrast features, but above the artefacts which accompany it. By setting everything below the threshold to zero contrast and maintaining the values above the threshold, it should be possible to eliminate most of the artefacts while maintaining the peak features of the reconstruction; this is shown in Fig. 6.4(b).

This image of the peak features can be used as O_{reg} . By containing only features which should be true, the additional data from the virtual image space components should help the next reconstruction, which should accordingly have fewer artefacts associated with the missing data, as shown in Fig. 6.4(c). Because of the lower artefacts, the next application can reduce the threshold slightly. Accordingly, more of the peaks in the image are above the threshold, but without their associated artefacts. At this point the process repeats, each time reducing the threshold until all the true features are reconstructed.

This thresholding method makes the assumption that the velocity is either always above or always below the background. Because this is the case for many applications, all the models considered in this chapter use this assumption.

6.4 Algorithm

A flowchart of the complete VISCIT algorithm is included in Fig. 6.5. Initially, let us consider a case which lies fully within the Born approximation, and hence ignore the steps corresponding to the inhomogeneous background (illustrated in blue). An

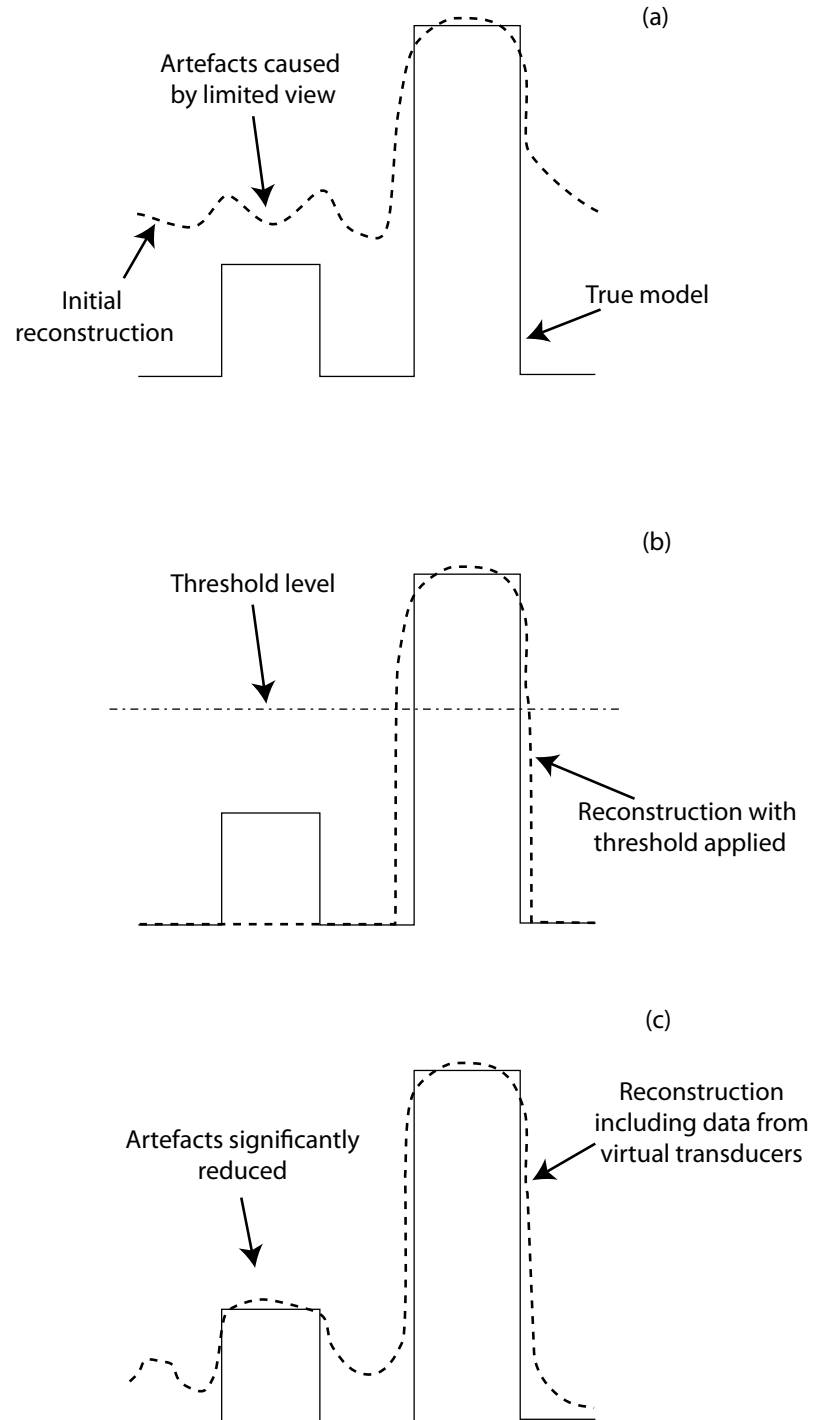


Figure 6.4: Schematic outline of the thresholding regularisation method. The initial reconstruction is shown in (a), with a threshold applied in (b) and the subsequent VISCIT reconstruction is shown in (c).

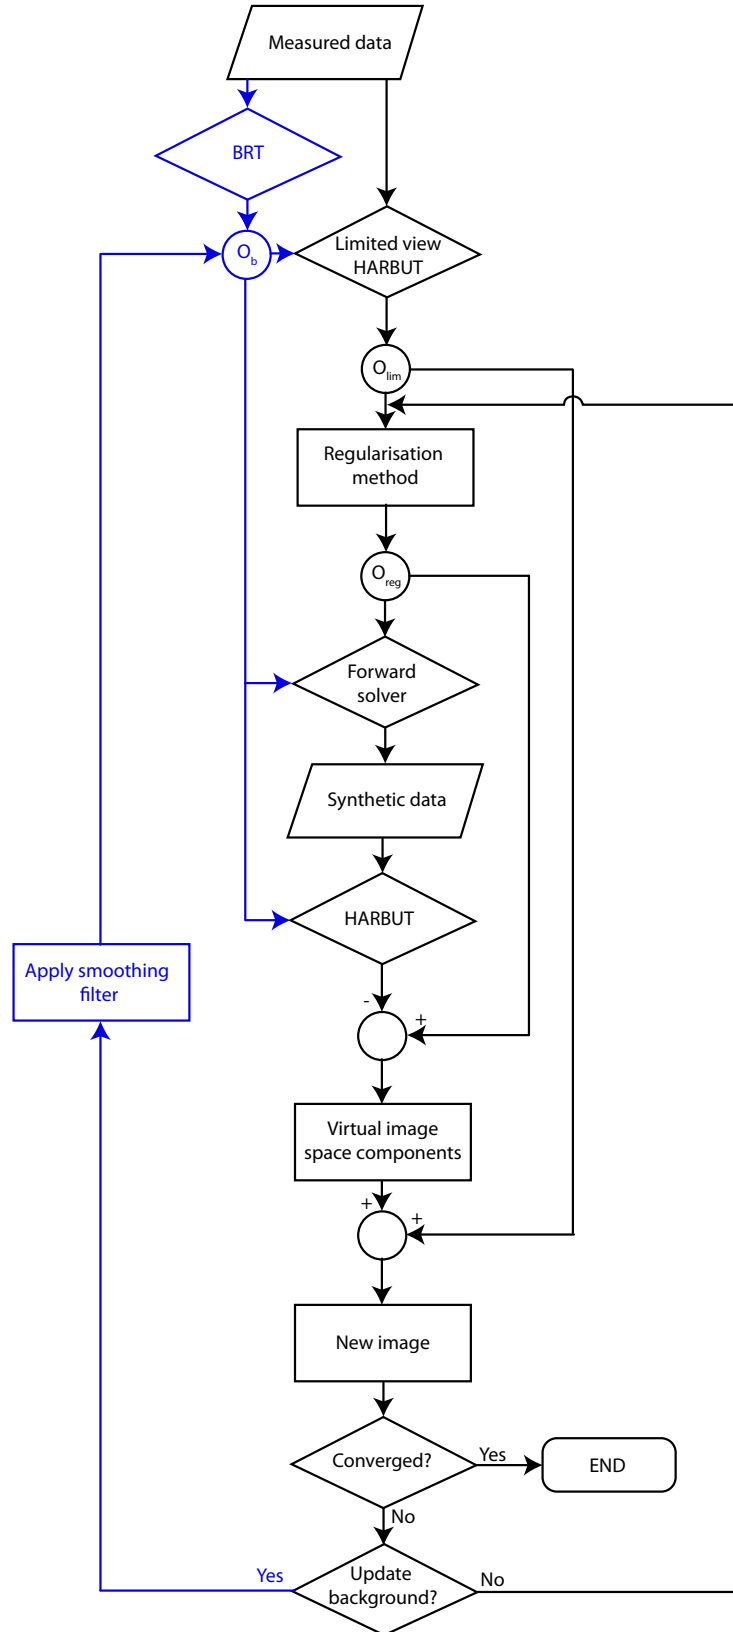


Figure 6.5: HARBUT for a limited view configuration using VISCIT. The stages in blue are needed to incorporate a background into the algorithm, which is necessary to extend the range of the algorithm beyond the Born approximation.

initial image is generated from the measured, limited view data, and regularisation is applied. I then run a forward model of this regularised image using eq. (6.7), followed by an inversion, to obtain an image of the effect the limited view array would have on the regularised image. This is then subtracted off from the regularised image to obtain the remainder $O_{reg} - \mathfrak{I}(\psi_{virt,lim})$, i.e. the virtual image space components. This can be added to the original limited view reconstruction of the object function to obtain the new image.

At this stage I check for convergence; an automated convergence criterion has not yet been developed but the algorithm is presently terminated when the image has visibly stopped changing between steps. If it has not converged, the image can be used to replace the limited view object function in the next iteration of the algorithm.

Many limited view configurations do not form convenient uniform-radius arcs, so the standard circular beamforming to diffraction tomography filter is therefore not suitable. Appendix B outlines the alterations needed to make the algorithm suitable for a general array.

It is possible to use an inhomogeneous background in the algorithm to help overcome the limitations of the Born approximation. The background stages are illustrated in blue in the flowchart. An initial background can be obtained from the BRT algorithm. Firstly, this is used in the generation of the initial limited view HARBUT image.

When the forward solver needs to be run, O_b is subtracted from O_{reg} to give the O_δ which is then run through the integral of eq. (6.8). The resulting synthetic data can be inserted straight into the HARBUT algorithm, which uses the same background again to reconstruct a new O_δ , which is then added to O_b again. This is then combined with O_{reg} and O_{lim} as in eq. (6.16).

After the convergence check, there is the option to use the improved image to update the background, O_b ; this is particularly useful since, because of the limited view configuration, the initial background image from BRT will often be poor. It is possible, therefore, to repeatedly update the background as in the iterative HARBUT

method introduced in Chapter 5 for guided wave tomography (GWT).

A problem now lies in how to perform both the VISCIT and the background update iterations in a logical manner to converge to the right solution. From the GWT study in Chapter 5, it was seen that typically around 10 iterations of HARBUT were needed. By contrast, around 100 regularisation iterations were needed in [92] for the known information to propagate fully to the unknown spatial frequency components. While this is obviously dependent on the nature of the regularisation used, it is likely that many more iterations will be needed to populate the unknown virtual transducer values than for HARBUT to converge. My approach is therefore to apply several VISCIT iterations with a fixed background, then use the resulting image as the new background and repeat the process. One advantage of this is that each VISCIT iteration requires fewer steps and hence less computational power than the HARBUT iterations.²

Experience has shown that the best approach is to update the background more frequently early on since this is fundamental to obtaining the correct phase distortions through the object. In this chapter I therefore use a geometric progression, performing 2^n regularisation iterations before updating the background, and repeating the process for $n = 0, 1, 2, \dots$

6.5 Numerical model parameters

Figure 6.1(a) presented the backwall model considered in this chapter. All the images here are reconstructed by considering the problem as a set of two parallel arrays, reconstructing a mirrored object, as shown in Fig. 6.1(b).

The model is 200mm wide and the array to backwall distance is 50mm. This aims to represent a typical backwall problem. A variety of acoustic scatterers were placed

²One useful optimisation is that the distorted Green's functions G_b need only be calculated once for any particular background. They can then be re-used multiple times by both the forward solver and the HARBUT algorithm. Calculating the distortions is the slowest stage in the algorithm; this is therefore a significant optimisation.

in a homogeneous background medium of water, although the theory is equally valid regardless of the background medium.

The array has 128 elements spaced by 1.5mm. The centre frequency was 1.5MHz and a 5-cycle Hann window was used. The array was modelled as a series of point sources; these were generated by exciting individual nodes at the appropriate location. Because of the spacing of the elements and the grid, no rounding was needed in the positioning of the array elements.

A 2D acoustic model is run to simulate the problem, using the finite difference, time domain (FDTD) method. Convolutional perfectly matched layers [66] surrounded the model on three of the four sides, with the exception being the backwall which reflected as though it was sound-soft. The grid was 3200 by 910 nodes, spaced by $\lambda/16$.

Because I am treating the data as though it had come from two parallel arrays rather than the backwall reflection, any component outside the main backwall reflection will correspond to the direct reflections of Fig. 6.1(a). This effectively forms noise in the data and is removed by gating. To do this for each send-receive pair, the approximate arrival time of the wave is estimated by calculating the distance and assuming a homogeneous background. Since I am imaging low contrast scatterers this gives a good estimate of where in the time trace the useful signal lies, making the gating straightforward.

The actual arrival times, used for BRT, were estimated by enveloping the signal and taking the point at which it exceeded a threshold of 50%. The frequency components were obtained by performing a Fast Fourier Transform (FFT) on the gated signals.

6.6 Numerical results

Three numerical models were run; one of a simple case to test the algorithm with five inclusions of a variety of sizes and contrasts, one of a more realistic Gaussian type velocity variation and one of a more complicated scatterer.

6.6.1 Five circular inclusions model

Figure 6.6 presents the results for the first model. The original model is shown in Fig. 6.6(a) and consists of a uniform medium containing five circular inclusions. Three have 2% contrast (1530m/s) and are sized 5mm, 10mm and 15mm in diameter and the remaining two both have diameter 10mm with 1% and 5% contrast (1515m/s and 1575m/s).

It is clear that the standard reflection imaging in Fig. 6.6(b) can detect the discontinuous boundaries of the objects, but it is not possible from this to quantitatively determine the sound speeds within each inclusion. Figure 6.6(c) shows the BRT reconstruction, showing the expected low resolution; in this case the resolution can be estimated as $\sqrt{L\lambda} \approx 10\text{mm}$ [25]. Interestingly the image shows diagonal, ‘X’, shaped artefacts which extend to the ends of the transducer array; these are a common result of this limited view configuration (see for example [95]). Since the inclusion is spread over a larger area by the artefacts, a lower contrast is needed to obtain the same arrival time offset, which is why the reconstructed contrasts are lower than the original.

Diffraction tomography is shown in Fig. 6.6(d). The Born approximation is violated for all but the 5mm and 1% inclusions. The artefacts are a result of the limited view and the Born approximation. Using the BRT image as a background, as is done in the HARBUT reconstruction of Fig. 6.6(e), improves the reconstruction slightly, but the BRT background is not accurate enough to avoid Born approximation artefacts completely for the high contrast inclusion. The limited view is shown to have an effect here, with the ‘X’ shaped streaks visible.

Finally, VISCIT is presented in Fig. 6.6(f). The virtual image space components are populated using the thresholding regularisation method. The threshold fraction starts at 0.8 (i.e. the threshold is set to a contrast of 80% of the maximum) and is reduced by a factor of 2 at each iteration. 31 iterations are performed to obtain the final reconstruction. As explained in Sec. 6.4, background updates are performed after 2^n iterations, where $n = 0, 1, 2, \dots$; in this case the background was updated five

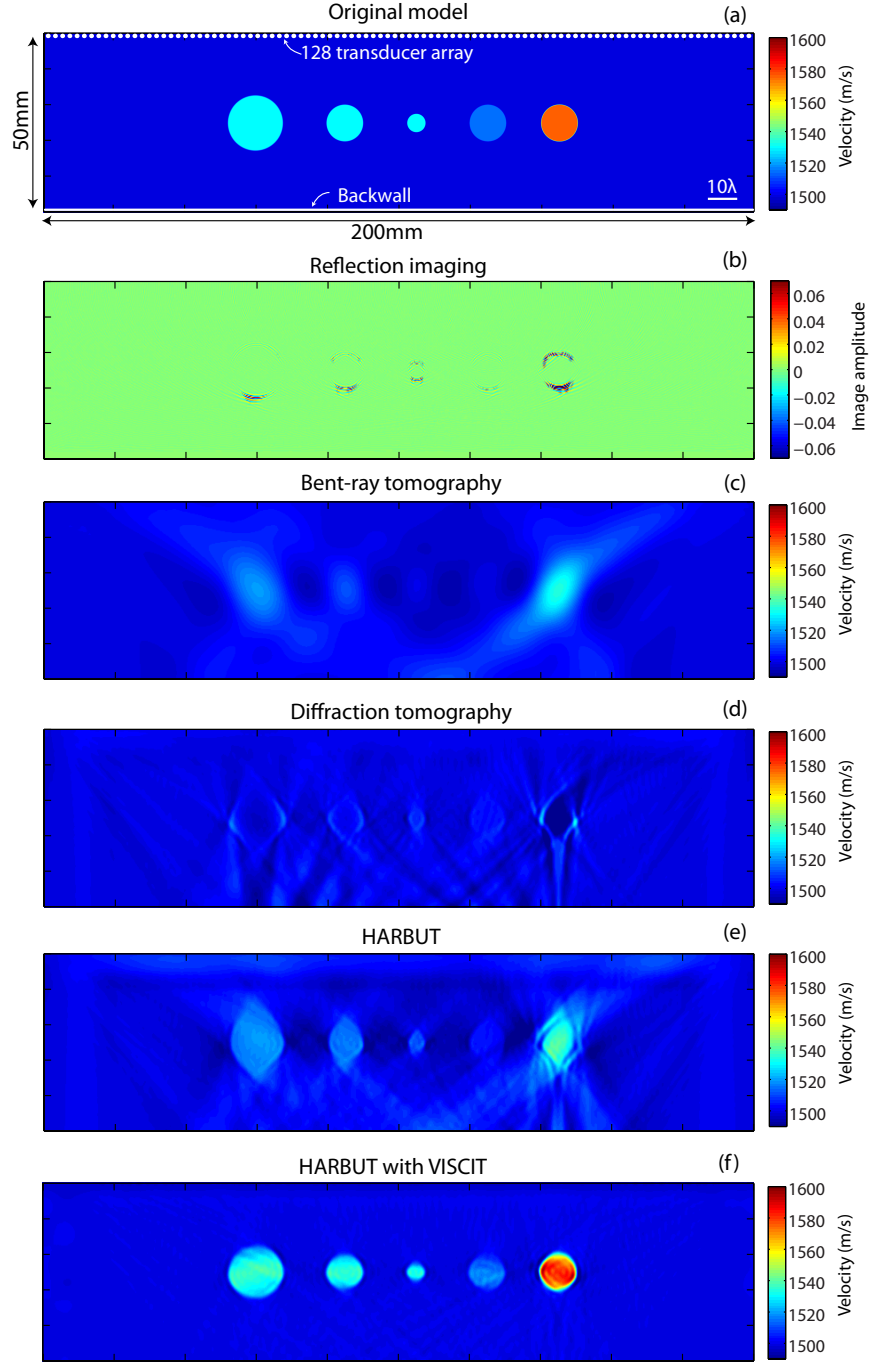


Figure 6.6: Algorithm performance with different sizes and contrasts. The wavelength is 0.8mm. The transducer array lies at the top surface of the image and the backwall coincides with the bottom surface. (a) gives the original sound speed map, consisting of three 2% contrast (1530m/s) inclusions sized 5mm, 10mm and 15mm in diameter, and two 10mm inclusions of 1% (1515m/s) and 5% (1575m/s). Reflection imaging in (b) shows no quantitative information; (c) shows that BRT is poor resolution and exhibits diagonal, ‘X’ shaped artefacts. The Born approximation is violated for most of the inclusions, as is clear from (d), the DT reconstruction, which is also poor because of the limited view. It is clear that HARBUT does provide an improvement, as shown in (e), but the high contrast inclusion in particular is poorly reconstructed with many artefacts. (f) presents HARBUT with VISCIT, which has few artefacts and is able to accurately capture the sound speeds and shapes of all five inclusions.

times during the 31 VISCIT iterations.

The shapes of the inclusions are all well defined and there are few artefacts. Importantly for quantitative imaging, the sound speeds in each of the inclusions are very close to the original model. The top and bottom surfaces of the inclusions are poorly defined, which is to be expected given that the transmission data does not encode any information about these surfaces. There is the possibility that this aspect could be improved by combining it with the reflection data of Fig. 6.6(b), which captures well the top and bottom surfaces.

Overall, VISCIT works well, but it could be that the example used here is well suited to it. Since the model has discrete inclusions containing constant sound speeds, it would be expected that the thresholding approach would be particularly appropriate. I now consider a more realistic example to test the algorithm more thoroughly.

6.6.2 Gaussian profile example

Early stage stress damage is unlikely to form discrete circular patterns as in the previous example. Here, I use a Gaussian function to define a more realistic velocity change corresponding to stress damage. The Gaussian has a peak of 1505m/s against a 1500m/s background, and a 5mm radius. In metallic components, it would also be expected that the grain boundaries would cause further scattering, which would manifest itself as noise in the data. To model this I add a randomly varying background, as used in the breast model of Chapter 3 – in this case varying from 1499 to 1501m/s with a correlation length of 1.5mm – to the Gaussian sound speed variation. This combination of slowly varying sound speed and small scale random variation is completely different to the previous example and should be more challenging to the algorithm. The results are shown in Fig. 6.7.

Figure 6.7(a) shows the original model. Figure 6.7(b) shows that reflection imaging is totally unsuitable for detecting this form of damage, due to the lack of discontinuous boundaries to reflect waves back to the array. BRT, DT and standard HARBUT

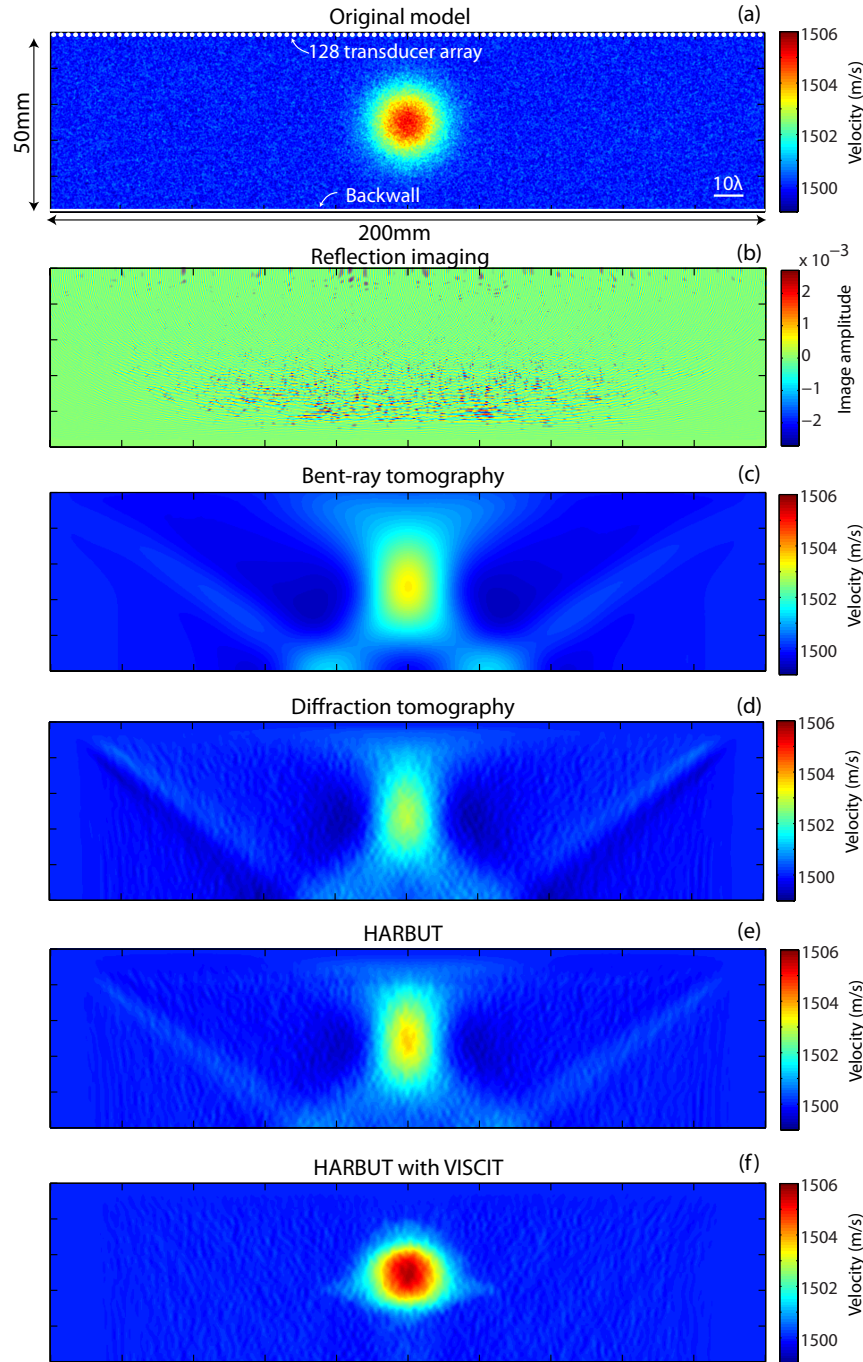


Figure 6.7: Reconstructions from a model representing stress damage within a metal component. The original model (a) is formed by summing a 5m/s Gaussian of radius 5mm to a random background. This background had a correlation length of 1.5mm and varied between 1499 and 1501m/s. (b) shows that without the presence of sharp boundaries, reflection imaging is completely unsuitable for detecting this early stage damage. (c) gives the BRT image, which, as in Fig. 6.6(c) features diagonal artefacts. (d) shows diffraction tomography, again with the diagonal artefacts. Because of the low contrast of the scatterer, the Born approximation works well so there is no sign of the Born artefacts seen before. HARBUT, in (e), shows a slight improvement over (d) although the diagonal artefacts are still present and the reconstructed values are underestimates of the true values. (f) shows HARBUT with VISCIT, confirming that the method is suitable for more realistic sound speed maps.

are shown in Fig. 6.7(c), (d) and (e); all perform reasonably well although do underestimate the sound speed and contain the diagonal artefacts seen before. The Gaussian also appears elongated in the vertical direction. HARBUT does not show a significant improvement over DT because the low contrast means that the scattering is well modelled by the Born approximation.

Figure 6.7(f) shows HARBUT with VISCIT, obtained using the same parameters and same numbers of iterations as in the previous example. This performs well, capturing the peak sound speed of the original and producing a far more accurate representation of the Gaussian. The regularisation method of thresholding to minimise artefacts seems to work well, despite the threshold not being well suited to the slowly varying nature of the Gaussian.

6.6.3 Complex scatterer

The Gaussian model provides a good representation of a typical area of stress damage, and it has been demonstrated that the algorithm performs well with this model. To demonstrate the suitability of the technique for a wider range of problems, I now consider a more complex scatterer than the simple shapes considered so far.

The sound speed map chosen is derived from the complex defect used in the guided wave tomography study of Chapter 5. This is scaled to vary between the 1500m/s background and 1510m/s at its peak, and is shown in Fig. 6.8(a).

Figure 6.8(b) gives the reflection image, which, as shown in the examples before, shows little meaningful information. Figures 6.8(c), (d) and (e) show BRT, DT and HARBUT reconstructions, which as expected show the diagonal limited view artefacts. The Born approximation is quite well suited to the low contrast, so little gain is achieved by using HARBUT over DT. As in previous cases, none of these three methods capture the sound speed particularly accurately.

Figure 6.8(f) shows the VISCIT HARBUT image (again generated with the same parameters), capturing all the details of the original, and obtaining good estimates

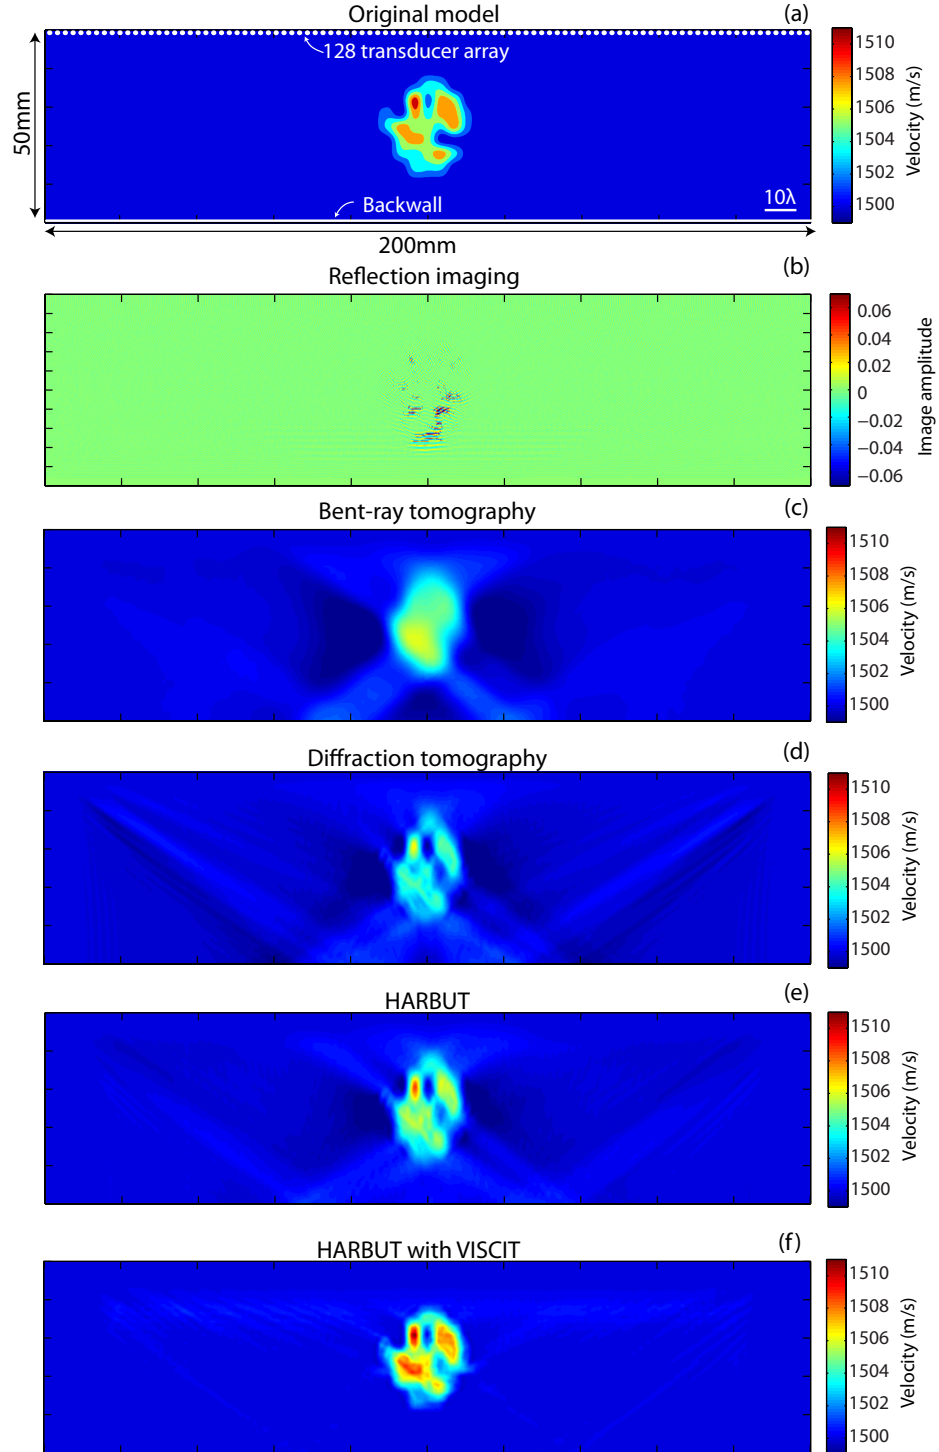


Figure 6.8: Reconstructions of a complex feature. (a) shows the original, complex scatterer. Little detail is visible in standard reflection imaging (b). The resolution of BRT, as shown in (c) is too poor to be able to identify any of the features. (d), showing diffraction tomography, does capture some of the features, but there are significant artefacts and the sound speeds are not accurate. (e) is the HARBUT reconstruction, which obtains a more accurate reconstruction, but again shows significant limited view artefacts. (f) shows how the VISCIT method obtains better sound speed estimates, although a few artefacts are visible.

of velocity throughout the scatterer. As before, VISCIT combined with the thresholding approach has minimised all the artefacts while maintaining the true features.

6.7 Experimental results

Experiments contain a wide range of uncertainties which can act to corrupt measured data. By their nature these effects are difficult to accurately model using numerical methods; it is therefore vital to verify that any new imaging method works with experimental data.

This chapter has so far primarily considered imaging using a linear array opposite a reflecting backwall. The limited view problem for quantitative imaging exists in other fields too; here I consider guided wave tomography (GWT), the problem investigated in Chapter 5 with a circular, full view array.

While the circular array is convenient for imaging, in practice in GWT it is rarely feasible to encircle the defect in this way. A more common approach is to have two parallel rings of transducers placed around the circumference of the pipe. By unwrapping the pipe into a flat plate, the problem becomes that of reconstructing a quantitative image using data from two parallel arrays.

Since any waves passing off the top of the unwrapped ‘plate’ will reappear on the bottom, due to the circular nature of the pipe, this unwrapping is not a completely accurate representation. However, with suitable gating to remove superfluous waves this is a good approximation.³

Here, I evaluate the effects of the limited view problem with experimental GWT data and test VISCIT’s performance to solve the problem. I reuse the full view, two defect data from Chapter 5 and utilise only limited sections of this data to synthesise

³Interestingly these helical waves can be exploited to improve the reconstructions, since they provide additional illumination angles over those contained in the standard parallel arrays and hence effectively increase the size of the array. The challenge is how to extract and separate each wavepacket from the measured signals. Volker and Bloom [78] claim to do this to extract arrival times for ray tomography, but to avoid such complications here I only consider direct waves.

a pair of limited view arrays.

Figure 6.9 gives the results from this for a variety of limited view array sizes. Figure 6.9(a) shows the original thickness map and Fig. 6.9(b) gives the full view image. The first column of the remaining images shows standard limited view HARBUT, and the second column shows HARBUT with VISCIT. All VISCIT reconstructions are performed with 63 iterations while updating the background six times, using the geometric progression ordering discussed earlier.

In the second row, Figs. 6.9(c) and (d), the transducers in the top half of the array are taken as sources and the bottom half are set as receivers. This results in the removal of certain send-receive combinations from the data set. There are noticeable artefacts in Fig. 6.9(c), using the standard reconstruction method, which result from (in effect) setting the unknown k-space components to zero. Importantly for quantitative imaging, setting the unknown values to zero results in a significant underestimate of the contrast of the defects. Using VISCIT, as shown in Fig. 6.9(d), achieves a much better result.

The next row shows what happens if only 75% of the transducers are used. Figure 6.9(e) shows how the artefacts expand as the arrays reduce in size. By contrast, VISCIT in Fig. 6.9(f) remains relatively free of artefacts and maintains good estimates of thickness. There are slight distortions in the shape of each defect, but these do not affect the results significantly.

The results are similar when 50% of the transducers are used in Figs. 6.9(g) and (h), although by this point the relatively small apertures are significantly distorting the defects in VISCIT, and the increased reconstruction size is accompanied by a reduction in contrast.

The 25% view shown in Figs. 6.9(i) and (j) takes this trend further. While the standard reconstruction does not allow any features to be identified, VISCIT separates out the two defects. These are extremely elongated, however, and accordingly underestimate the contrast. This trend is to be expected; by considering a signal from just a single send-receive pair, it is possible to reconstruct the velocity between

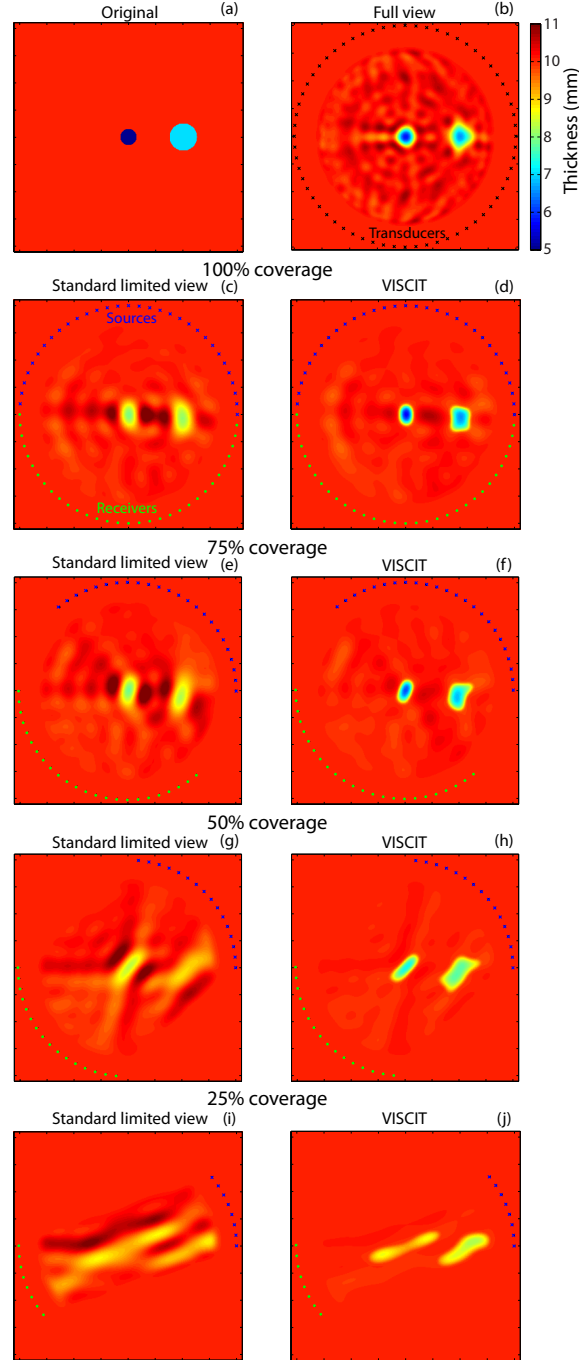


Figure 6.9: Limited view GWT results; these correspond to the full view images of Fig. 5.12. All images shown here are monochromatic and use the 50kHz component from a set of experimental data. (a) shows the true model, a 10mm plate with two circular defects; the central one has thickness 5mm and diameter 60mm and the offset one has thickness 7mm and diameter 100mm, and is offset by 200 mm from the array centre. (b) shows the standard full view result for comparison. For the remaining images, the first column, i.e. (c), (e), (g) and (i) shows standard limited view imaging with HARBUT and the second, i.e. (d), (f), (h) and (j) shows VISCIT. In (c) and (d) the top half of the array is used as sources and the bottom half as receivers. In (e) and (f) the transducer arrays have been reduced; only 75% (i.e. 24 sources, 24 receivers) are used. (g) and (h) reduce that to 50% and (i) and (j) only use 25%, i.e. eight sources and eight receivers. All the VISCIT reconstructions used six HARBUT iterations and a total of 63 VISCIT iterations in the geometric progression discussed in Sec. 6.4.




Figure 6.10: Cross sections through the VISCIT images in Fig. 6.9, taken along a horizontal line passing through the centre of the array. As the amount of available information is reduced, the reconstructed contrast is reduced.

as simply an average speed between the source and receiver, which would both underestimate the peak contrast and stretch the defect. Despite this trend, it is clear that VISCIT produces a significant improvement over HARBUT on its own.

Finally, Fig. 6.10 shows the thickness plotted along the line intersecting the centres of the two defects. The 100% limited view reconstruction shows little deviation from the full view reconstruction, but as the aperture angle reduces, the reconstructed contrast reduces too. The 75% coverage reconstruction obtains an estimate of minimum thickness within 1.2mm of the true value; this compares to around 0.8mm for the full view configuration.

6.8 Summary

This chapter has introduced VISCIT, the Virtual Image Space Component Iterative Technique, to aid imaging with limited view arrays. Extra information, obtained

by applying regularisation, is stored in the image space and is combined with data from the limited view array to improve the reconstruction.

The limited view configuration introduces diagonal, ‘X’ shaped artefacts spreading from each feature. The regularisation method exploits the fact that these artefacts are of lower contrast than the dominant feature. It reconstructs the peak contrasts first, by applying a threshold to the image to remove the lower contrast artefacts, and stores this in the virtual image space components. The data from these components then helps the next iteration of the image, reducing the artefacts, so the threshold can be reduced accordingly. This process is iterated until the image converges.

VISCIT was demonstrated to outperform standard HARBUT, DT, BRT and reflection imaging for three models: a set of five circular inclusions of various sizes and contrasts, a low contrast Gaussian representing early stage stress damage, and a more complex scatterer with several peaks. When testing it with limited view experimental guided wave tomography data, a similar improvement was observed.

Chapter 7

Conclusions

7.1 Thesis Review

New algorithms for quantitative subsurface imaging with mechanical waves have been developed through this thesis. The techniques introduced can provide significantly improved information about the material properties beneath an object's surface through improved resolution, robustness and accuracy.

Chapter 2 outlined the theory behind HARBUT (the Hybrid Algorithm for Robust Breast Ultrasound Tomography). It began by introducing the forward scattering model used to simplify the wave equation. This was then used in the inversion method allowing a velocity map to be obtained. Under this model, the velocity field was split into two components; a slowly varying, low-resolution background, which could be obtained by bent-ray tomography (BRT), and a low contrast, high-resolution remainder which was suitable for diffraction tomography (DT). The distortion of the wavefield as it passes through a scatterer is mostly accounted for by the low-resolution background, and can be incorporated into DT by adjusting the focusing used in the beamforming (BF) algorithm, which forms the basis of the DT image. The chapter explained how this allows a high-resolution reconstruction to be obtained without being constrained by the overly restrictive Born approximation.

Chapter 3 applied HARBUT to a representative numerical model of the breast to evaluate its performance, and also tested it with experimental in vivo data. The method was shown to perform well, achieving significantly higher resolution reconstructions of sound speed than the widely used BRT algorithm, allowing cancer masses to be more clearly identified. One drawback to the array configuration used, however, was the axial, out-of-plane resolution, which was demonstrated to be poor because of the large slice thickness. The effect of this was to project any features within a thickness of around 9mm into the imaging plane, potentially obscuring the details of the structures within the breast. Chapter 4 showed how this slice thickness problem can be reduced by treating the problem as fully 3D.

HARBUT was applied to guided wave tomography in Chapter 5, improving the accuracy of reconstructions from this 2D problem. Lamb waves were transmitted in the wall of a plate-like structure, and measured from a circular array. Velocity reconstructions from HARBUT were converted to thickness via the dispersion curves, obtaining accurate estimates of minimum thickness even for complex defects, allowing remaining component life to be accurately estimated. This was demonstrated for both simulated and experimental data.

Imaging when only a limited view array was available was considered in Chapter 6. Quantitative imaging is particularly sensitive to the limited view problem, especially when there are no transmission components available. This chapter solved the problem by using waves reflected from a parallel backwall to obtain transmission components and using regularisation methods in conjunction with VISCIT, the Virtual Image Space Component Iterative Technique, which can be used to reduce the limited view artefacts.

7.2 Main findings

Reconstructions in the field of breast ultrasound tomography are usually performed using the BRT algorithm, which ignores diffraction and therefore is limited to a resolution of around 10mm. By comparison, cancer masses of interest are typically

5mm and up, and hence there is significant potential to miss important, early stage tumours. The new method developed in this thesis, HARBUT, by combining the complementary strengths of BRT and DT, has been shown to significantly improve resolution in the image to around 1mm without needing to resort to slow, unstable, iterative full-wave inversion methods. The performance was demonstrated for a realistic 3D simulation (as published in [P5]) and verified with experimental, in vivo data (to be published in [P9]).

A significant challenge when reconstructing breast images from a toroidal array surrounding the breast is the large slice thickness; any features within a finite slice will be projected into the plane to form the image, which manifests itself as a limited resolution in the axial direction. By treating the data as fully 3D, rather than 2D as in the majority of breast ultrasound tomography algorithms, it was possible to improve the axial resolution from about 9mm to around 2.3mm. This has huge potential to improve breast cancer detection using ultrasound.

The field of guided wave tomography uses guided waves to determine the thicknesses of defects within a plate-like structure; this is commonly achieved by using the dispersion curves to map thickness to wave velocity. The defects are usually too small for the widely used, but limited resolution, ray tomography methods to detect. Methods based on the Born approximation are also unsuitable because of the very high contrast associated with the wall-thickness reductions of interest, which can be up to 80%. By combining BRT and DT together, HARBUT again helps this situation, using BRT as a background to improve DT. However, because the BRT reconstruction used as a background is usually so poor, further improvements are possible by taking the resulting HARBUT image as a background and repeating the process. By applying several iterations of HARBUT it was possible to reliably and robustly reconstruct minimum thicknesses of defects to within about 1mm for a 10mm plate, whereas the other methods rarely produce an estimate accurate to more than ± 3 mm for representative defects. This was demonstrated for both simulated and experimental data, and has been submitted for publication [P8].

Quantitative imaging is particularly sensitive to the limited view problem, when the

object to be imaged has restricted access preventing a full view array from being placed around the scatterer. While a number of solutions have been proposed to deal with this problem, few have been shown to work with a representative, realistic data because of the assumptions necessary to image with the missing data. This thesis has introduced VISCIT, the Virtual Image Space Component Iterative Technique, which was shown – in conjunction with a thresholding regularisation technique – to allow accurate velocity reconstructions to be obtained from limited view arrays. This was demonstrated to be stable for a range of models of varying complexity, providing far more accurate reconstructions compared to standard methods. The robustness of the algorithm was confirmed when using it to image experimental guided wave tomography data with a variety of limited view arrays.

7.3 Areas for future work

Within the field of breast ultrasound tomography, there are several future avenues which could be pursued. One of the main areas for improvement is in selecting a suitable array configuration for breast ultrasound tomography. Chapter 4 demonstrated the limitations of the traditional ring array because the height of the transducers causes a very directional beam. Given that the alternatives for 3D breast ultrasound imaging are generally limited to mechanically scanning a standard linear array across the breast [96], or the undersampled system of [15], there is a clear need for further study to develop an improved array configuration. This is not a straightforward problem; Chapter 4 showed how there is significant redundancy in collecting data from all possible directions, so there is a significant challenge of collecting suitable data such that a good resolution in all directions can be obtained without collecting unnecessary information. The expense of building arrays is high, so such a study would need to be performed numerically prior to investing in hardware.

The use of computational simulations is underdeveloped in ultrasonic breast imaging, largely being restricted to overly simplified, unrealistic 2D simulations. While this thesis addressed this with two sets of data from sophisticated 3D models, given

the inevitable advances in processing power there is massive potential to develop more realistic models which could greatly aid the understanding of this area.

Current computer hardware limitations meant that a suitable data set to test the full HARBUT algorithm in 3D was unavailable. As detailed in Chapter 4, while the theory can be logically extended such that HARBUT should work well with such a model, it is important to obtain some validation with a numerical simulation. Following on from this, the verification of the method with suitable experimental data is important.

One of the important discoveries described in this thesis is the VISCIT method of Chapter 6, which was shown to have huge potential to improve limited view quantitative imaging. At present, the method does not utilise any of the reflection data from the array. Depending on the application, this might contain useful information, particularly if a wide bandwidth can be extracted from the data. The algorithm could therefore be further developed to include wideband reflection data to improve the quantitative images.

One aspect not considered in the application of VISCIT to the early stage stress damage detection problem is that the problem is generally not acoustic as assumed in Chapter 6, but will instead be elastic. As a result, mode conversion will be likely at the backwall reflector so studies will have to be performed to help understand how this will affect the reconstructions.

The iterative HARBUT method introduced in Chapter 5 included a stopping criterion to determine when the algorithm had converged. VISCIT has not yet had such a convergence criterion developed. Since the change between iterations is going to significantly depend on the thresholding level and will be further complicated by the fact that HARBUT is iterated too, a careful study will need to be done to find a stopping criterion which will be suitable for all problems.

At present, the thresholding method contains the assumption that the sound speeds of the image features all lie one side or the other of the background. However, it is possible that the method could be extended to a more general case by applying

the threshold both sides of the background, at equal distances from it. Anything between the two thresholds could be set to the background value and anything outside maintained. Whether this (or a similar) approach would work needs to be investigated.

Appendix A

Ultrasound transmission through a penetrable cone

A.1 Introduction

The scattering of waves from simple 3D objects has received considerable interest. Morse [29] calculated the solution for scattering from an impenetrable cylinder, and this approach has been extended to scattering from finite length cylinders [97, 98], cylinders within penetrable cylinders [99, 100], two cylinders in close proximity [101, 102] and random configurations of cylinders [103]. The problem of scattering from a sphere and variations of this has similarly been extensively addressed (e.g. [104–109]).

Despite this large body of work on scattering from cylinders and spheres, the majority of work on cone scattering is limited to impenetrable boundary conditions (see [110] for a comprehensive list), with little attention paid to the penetrable cone. The latter is of practical interest in breast ultrasound tomography. A typical configuration is illustrated in Fig. A.1. A toroidal transducer array is placed horizontally into a water bath, which acts as a coupling medium. The patient lies prone with the breast in the water bath, in the middle of the transducer array. The array illuminates and measures ultrasound transmission through the bulk of the breast, and




Figure A.1: A practical example where understanding scattering from a penetrable cone is useful: breast ultrasound tomography. The bulk effect of the waves passing through the breast is approximated by acoustic transmission through a cone.

these measurements are used to reconstruct tomographic images of sound speed in a similar way to X-ray CT [47].

The breast tapers in the axial direction, causing a deflection of the ultrasonic beam outside the array plane and a subsequent drop in the transmission signals [58,69]. It is therefore important to understand how much the ultrasound transmission is affected by the shape of the breast and how much it is affected by its internal structure [111]. In the first approximation, this effect can be studied by assuming a conical shape for the breast.

The main papers which consider the penetrable cone are by Lyalinov [110,112,113] and Jones [114]. In these works, the solution is derived in the form of an integral or infinite sum, which can then be calculated numerically by a semi-analytical approximation. Although these methods provide valuable computational tools, physical approximations relating characteristic parameters of the problem to the transmitted amplitude are not available. Besides giving an insight into the scattering problem, physical approximations are instrumental in the design of scattering experiments such as those used in breast ultrasound tomography technology.

The aim of this appendix is to understand the physics that underpins the complex phenomenon of refraction of an ultrasonic beam through a penetrable cone and




Figure A.2: The cone distorts and scatters the field produced by the source transducer. The resulting field is measured by the receiver. The output plane is used to plot the output of the FDTD simulations in Sec. A.3.

to provide a physical approximation that links the characteristics of the scattered field to the parameters of the problem. In breast ultrasound tomography, arrays of rectangular transducers – both linear and toroidal – are used to provide the illuminations and measure the resulting field. Therefore, in this appendix I consider such transducers rather than plane wave illuminations and point measurements.

Section A.2 formulates the acoustic cone scattering problem, which is then used in Sec. A.3 to simulate ultrasonic transmission numerically with the Finite Difference Time Domain (FDTD) method. Section A.4 presents the physical approximation considering the case when the transducers are diametrically opposed. The approximation is validated against the FDTD results and its range of applicability is defined. The results are then generalised to the case of off-axis transmission in Sec. A.5.

A.2 The cone transmission model

Figure A.2 presents the model considered. A source transducer provides the illuminating wavefield which interacts with the cone and is subsequently measured by the receiver transducer. Both the cone and the background medium are acoustic

materials, i.e. waves in them obey the acoustic wave equation [29]

$$\rho(\mathbf{r})\nabla\cdot\left[\frac{1}{\rho(\mathbf{r})}\nabla p(\mathbf{r})\right]-\frac{1}{c(\mathbf{r})^2}\frac{\partial^2 p(\mathbf{r})}{\partial t^2}=0 \quad (\text{A.1})$$

where $p(\mathbf{r})$ is the pressure at point \mathbf{r} , $\rho(\mathbf{r})$ is the density and $c(\mathbf{r})$ is the sound speed.

The transducers are modelled by Huygen's principle [82, 115], with the surface of the source transducer modelled as a series of point sources of identical phase and amplitude. Similarly, the measured signal from the receiver transducer is the sum of the signals taken from all points on its surface. Here, the spacing between points is set to be less than $\lambda/10$, where λ is the wavelength, which ensures the resulting field or measurement is sufficiently close to the limit where the discrete sum is replaced by a continuous integral. I consider rectangular transducers, so the parameters that define the beam are transducer width and height.

The cone itself is defined by its diameter in the plane at $z = 0$, its sound speed and density, and the angle of the boundary relative to the z axis. The central axis of the cone is assumed to be fixed at $x = y = 0$; since both transducers can move relative to it generality is maintained.

In this appendix I consider the conditions encountered in breast ultrasound tomography. The diameter of the cone at $z = 0$ is of the order of 100λ at a typical $\lambda = 1\text{mm}$ which is required to achieve full breast penetration due to attenuation. In this regime, refraction dominates. By contrast, the transducer dimensions are of the order of 10λ , which allows diffraction of the beam to occur. I investigate contrasts of up to 10% where the acoustic equations are valid. Larger contrasts are likely to correspond to solid objects, where the elastic equations are more suitable for describing the material behaviour than the acoustic equations.

The primary effect of a density contrast between the cone and the background is to increase the impedance contrast and hence the reflection of the wave's energy which results in a drop in transmission amplitude by energy conservation. For a density contrast of about 10%, the reflection coefficient will remain fairly small so this loss is considered negligible. The density is therefore kept uniform at 1000kg/m^3

throughout, simplifying the analysis without losing generality.

A.3 Numerical study

This section presents a series of numerical simulations of a transducer beam scattered by a cone.

This simulation outputs the amplitude of the beam on the output plane of Fig. A.2 in the previous section. A source transducer produces a beam which is transmitted through the cone then projected onto the output plane. The source transducer height, h , is 20mm, the width, w , is 4mm, and $r_s = 80$ mm. The output plane is at $x = 80$ mm. The background sound speed is 1500m/s and the density is 1000kg/m³. The cone radius is $r_0 = 40$ mm and the aperture of the cone is 90°. β is defined as the angle of the boundary relative to the vertical; since this is half the aperture, $\beta = 45^\circ$.

The Finite Difference Time Domain (FDTD) method is used to perform the simulation. The explicit time stepping nature of the method means that memory usage is $O(N)$, where N is the number of degrees of freedom, whereas for implicit solutions this typically becomes $O(N^2)$. This allows simulations with many more degrees of freedom to be performed, which is vital for the 3D simulations considered here. A staggered Yee grid discretisation method [65] is used on the coupled first order acoustic equations of pressure and velocity [29]. The mesh is terminated with perfectly matched layers [66, 116] to minimise reflections from the boundaries of the domain.

The domain is $-81\text{mm} < x < 81\text{mm}$, $-60\text{mm} < y < 60\text{mm}$, $-28\text{mm} < z < 28\text{mm}$. Using a spacing of $\lambda/12 = 1/6\text{mm}$, this corresponds to 993 by 741 by 357 grid points including ten-node-thick boundary layers. Using a Courant number of 0.95 to determine the time step, 4345 time steps were needed, which gives enough time for the wave to pass across the domain and halfway back again. The calculations were performed using single precision floating point numbers. The simulation took around 6

hours to run on a single core of an AMD Opteron 8384. The time domain data were recorded for a grid of points on the output plane; subsequently Fast Fourier Transforms were performed to convert each time-based signal to the frequency domain.

Figure A.3 plots the amplitude of the beam on the output plane at 750kHz for 7 different cone sound speeds: 1650, 1560, 1530, 1500 (i.e. a homogeneous field), 1470, 1430 and 1350m/s which correspond to sound speed contrasts of +10%, +4%, +2%, 0%, -2%, -4% and -10% respectively. Also plotted in Fig. A.4(a) is the amplitude along the $y = 0$, $x = 80\text{mm}$ line for all the cones.

The presence of the cone causes the beam to be deflected in the z direction. The results show that for cone sound speeds below the background, the beam is deflected upwards, and for cone sound speeds above the background the beam is deflected downwards. The extent of deflection is dependent on the contrast level. Both of these effects are a result of refraction at the surfaces of the cone, as predicted by Snell's Law.

The lower sound speed cones cause the maximum amplitude of the measured beam to increase relative to that measured without the cone present. This corresponds to the focusing effect of the cone; the width of the beam is less and its energy is concentrated into a narrower area. Conversely the high sound speeds defocus the beam and spread the energy out more, resulting in lower amplitude. This result is similar to that observed for an acoustic cylinder [29].

The shape of the beam is affected by the contrast level. The lower sound speeds show a clear upturn at the edge of the beam, forming a 'smile' shape. This is less pronounced at lower contrasts and at the higher sound speeds a downturn is visible at the edges. These results suggest that the beam deflection increases away from the plane $y = 0$.

To model the effect of a receiver transducer diametrically opposite the source, the complex values of the wavefield in the output plane across the area corresponding to the receiver, in this case $-2\text{mm} < y < 2\text{mm}$, $-10\text{mm} < z < 10\text{mm}$, can be summed. The modelled responses at each cone sound speed are plotted in Fig. A.4(b); the

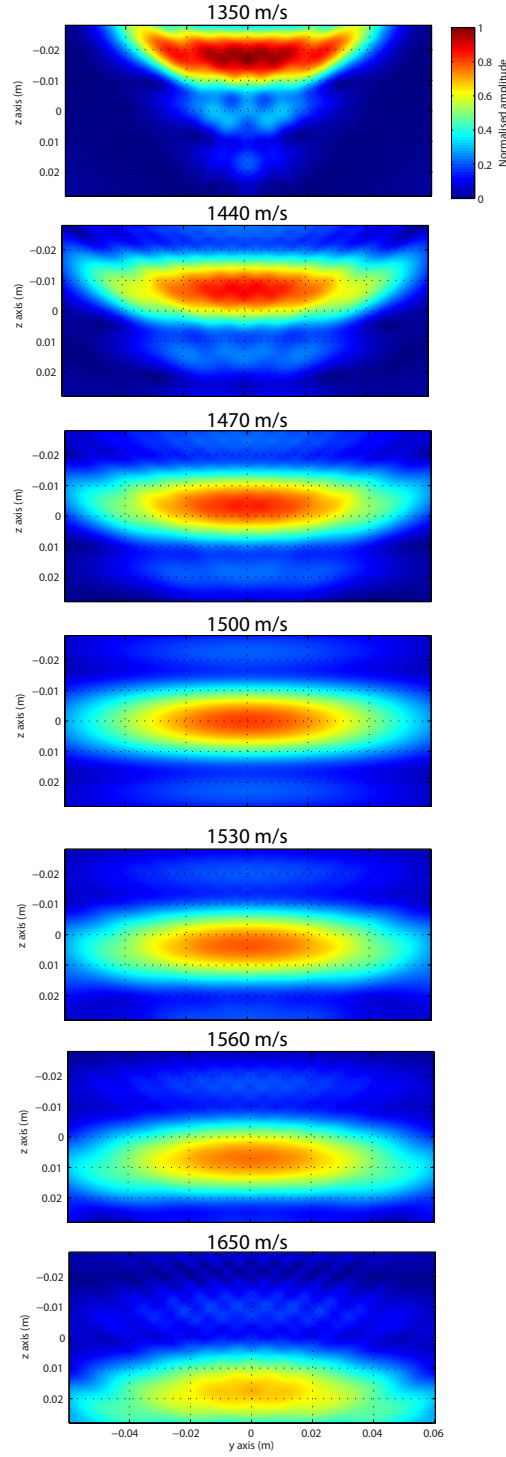


Figure A.3: 3D transmission through the cone, as calculated by a FDTD simulation. Each shows the amplitude of the beam across the output plane of Fig. A.2 for different cone velocity values at 750kHz.

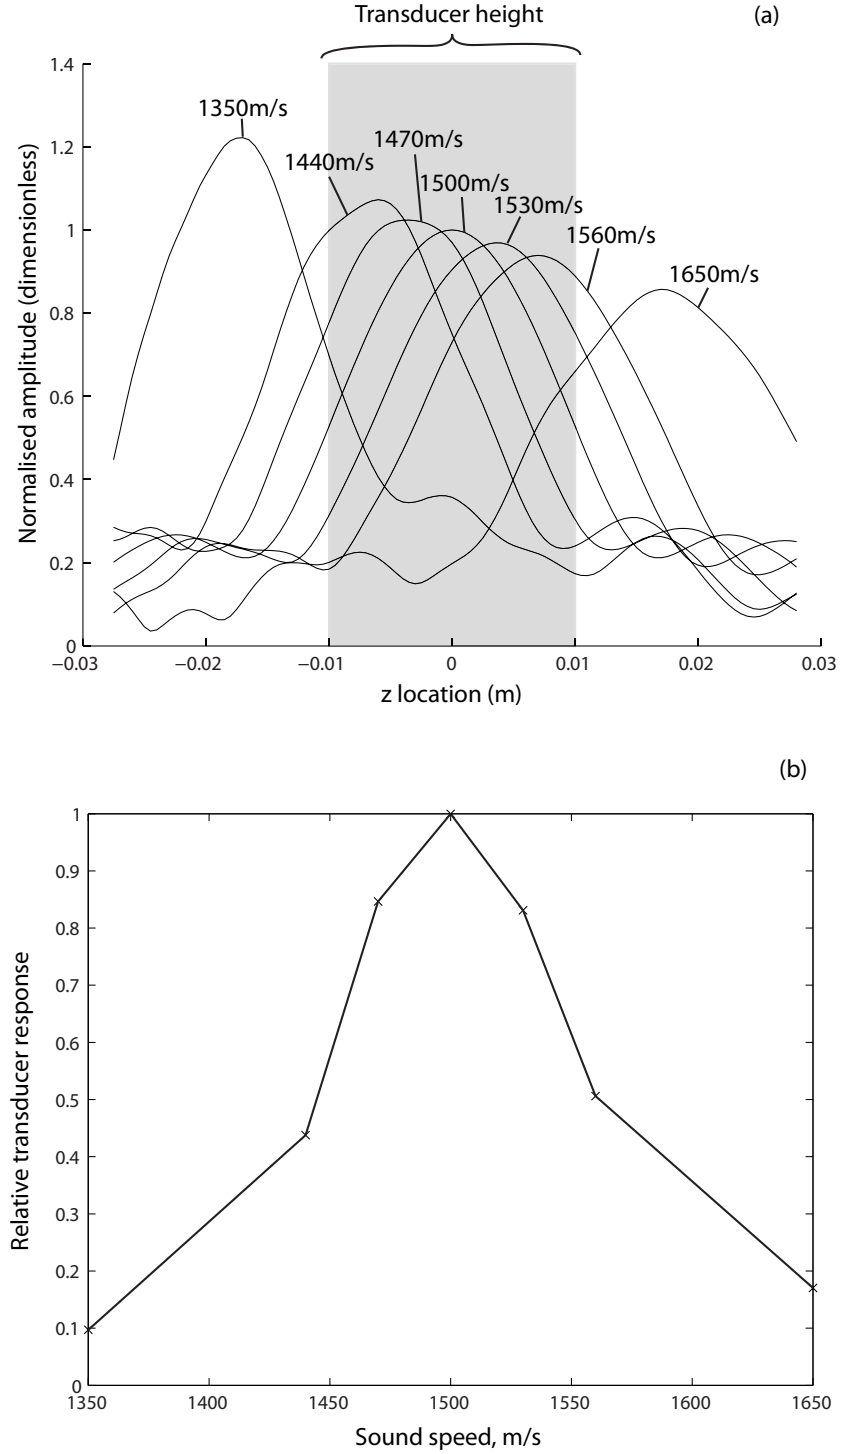


Figure A.4: (a) gives the amplitudes along the $y = 0$ lines for each of the simulations in Fig. A.3, showing how the deflection increases with contrast. (b) estimates the receiver transducer response through a Huygens approach of summing across the surface of the transducer. Values in (b) are normalised against the value that would be generated if no cone was present.

plotted values are relative to the equivalent measured incident field. There is a clear amplitude drop as the contrast increases and the beam is deflected more.

Performing repeated simulations to establish trends as the parameters change is impractical because a) the large number of independent parameters, including sound speed, boundary angle, cone diameter and transducer height, mean that many simulations would be necessary for a full understanding and b) the simulations take a large amount of time and resources to run. I therefore develop a simplified approximate analytical model to help generalise the effects of the various parameters.

To proceed with the analysis, I define the ‘transducer plane’ as being the vertical plane passing through the centre of the cone and the centre of the face of the source transducer. I also define the ‘array plane’ as being the horizontal plane at $z = 0$. I start the analysis by considering the behaviour of the beam within the transducer plane.

A.4 Amplitude drop within the transducer plane

By firstly studying transmission in the transducer plane, it is possible to develop an understanding of the trends that occur in the measured signal, which can subsequently be expanded to the full 3D problem. I firstly demonstrate how the wavefield within this plane behaves like the wavefield scattered in 2D by a prism of the same properties. For this purpose I use a 2D FDTD simulation, which had half the mesh spacing of the 3D model and was run on a graphics card using NVidia’s CUDA technology. In effect, the 2D simulation corresponds to a model that is uniform in the y direction, meaning that the cone is treated as a trapezoidal prism of infinite extent in the y direction and both the source and receiver transducers have infinite width.

The output along $x = 80\text{mm}$, $y = 0$ is plotted for both the 2D and 3D simulations in Fig. A.5, for the 1350m/s sound speed cone from the previous section. Each curve is normalised such that the maximum amplitude of the beam without the cone present




Figure A.5: Comparison between the 2D and 3D amplitude on the central plane. Each is normalised such that the maximum value of the incident beam is 1 in the corresponding 2D or 3D model.

will be 1.

The 2D result has a peak amplitude greater than 1. This amplitude increase relative to the incident beam is due to the wedge focusing the wave field as it deflects it upwards. The 3D result shows an additional focusing, which is due to the circular cross section of the cone in the x - y plane and is similar to what is seen for a penetrable cylinder [117]. However, this cylinder focusing effect is limited, so that the two curves are within a few percent, which means that this effect can be treated as negligible. The overall trends of the 3D model are captured by the 2D model, so it can be assumed that the behaviour on the central plane of the 3D model can be predicted by 2D simulations.

A.4.1 Plane wave illumination with finite-sized receiver

Now consider a plane wave illumination in the plane of the generatrix. While such an illumination is not achievable in practice, it does allow diffraction at the source to be neglected, providing a convenient starting point for the analytical approach. Because refraction dominates, Snell's Law can be used to model the deflection. The deflection at the first surface, δ_1 [see Fig. A.6(a)], therefore depends on the boundary

angle β as

$$\frac{\sin(\beta - \delta_1)}{\sin \beta} = \frac{c_1}{c_0}, \quad (\text{A.2})$$

where c_1 is the sound speed of the cone and c_0 is the sound speed of the background.

The total deflection δ_2 after the second interface becomes

$$\frac{\sin(\beta + \delta_2)}{\sin(\beta + \delta_1)} = \frac{c_0}{c_1}. \quad (\text{A.3})$$

By solving these two equations, the deflection angle of the wavefield incident on the receiver transducer can be determined.

Figure A.6(a) shows how the waves are rotated by the cone and consequently hit the transducer at an angle. Given a particular deflection angle δ_2 , the z component of the rotated wavenumber is $k_z = k_0 \sin \delta_2$ and the x component is $k_x = k_0 \cos \delta_2$, where $k_0 = 2\pi f/c_0$ is the wavenumber in the background at frequency f . The wavefield is integrated across the finite surface of the transducer to give the response relative to that of the undeflected wavefield

$$\begin{aligned} T &= \frac{\int_{-h/2}^{h/2} \exp[i(k_x x + k_z z)] dz}{\int_{-h/2}^{h/2} \exp(ik_0 x) dz} \\ &= -\text{sinc}(k_z h/2) \exp[ix(k_x - k_0)] \end{aligned} \quad (\text{A.4})$$

where the transducer height is taken as h and the sinc function is defined as $\text{sinc}(x) = \sin x/x$. Figure A.6(b) compares this analytical expression to a 2D FDTD simulation of the transducer plane.

The result demonstrates that as the contrast increases, the deflection of the refracted waves increases and the measured amplitude decays. This reduction occurs due to the transducer integrating the wavefield across its surface. If the wave fronts are parallel to the surface, as is the case when there is no cone present, the phase at all points on its surface is constant, leading to constructive interference and a strong response. As the direction changes, a phase difference occurs across the

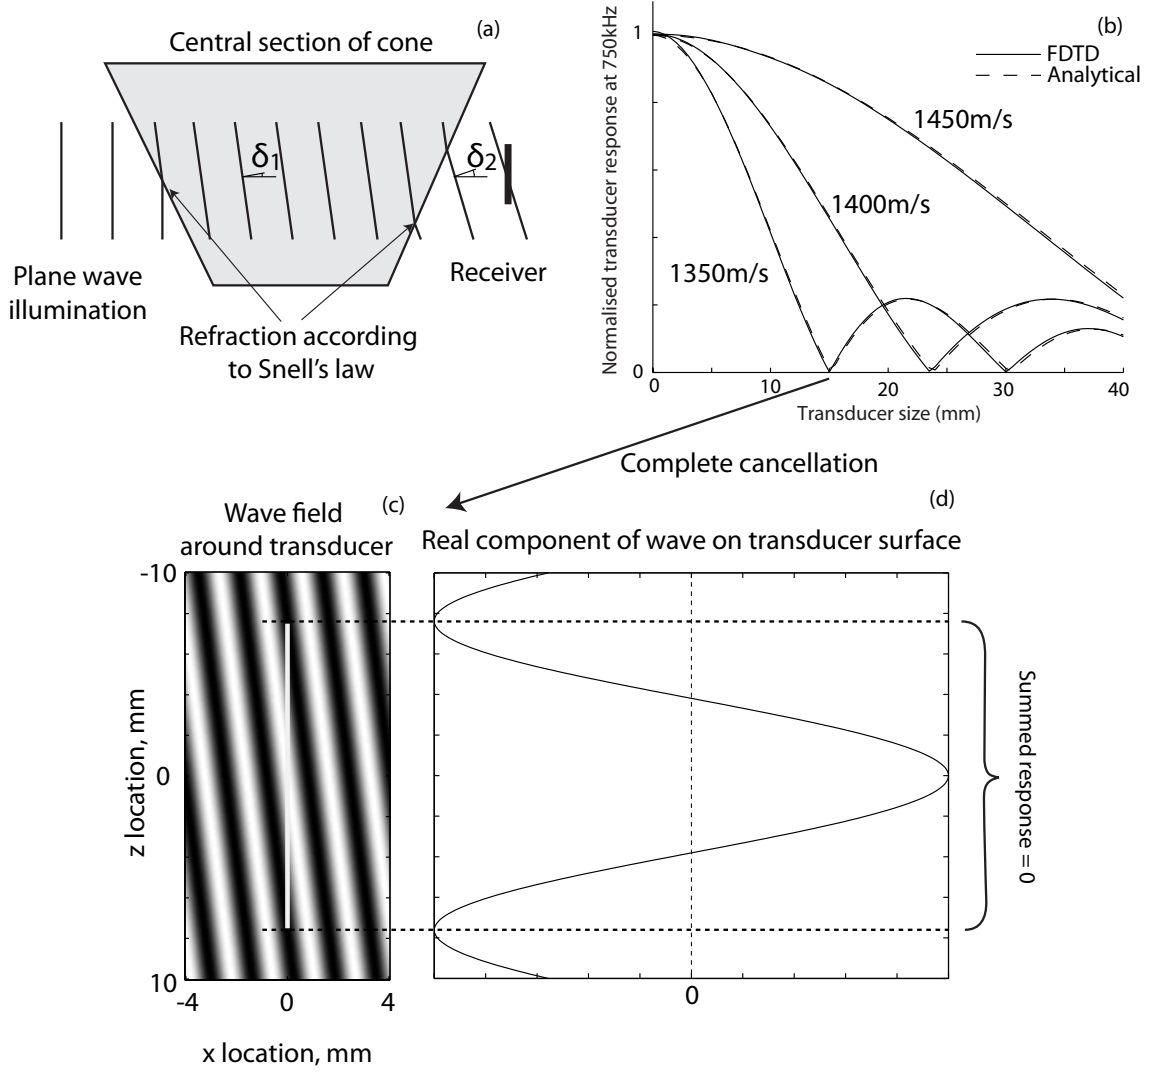


Figure A.6: Plane wave interaction with the cone, where $\beta = 30$ degrees (a) The refraction causes the wave field to be rotated. The effect of the transducer is to integrate the wavefield across its surface. (b) The estimated measurements from the transducer relative to the equivalent incident wavefield, for both the FDTD (dashed lines) and the analytical result (solid lines) of eq. (A.4). As the field gets rotated more, the integration causes the measurements to cancel, ultimately resulting in the zero for the 15mm tall transducer with the 1350m/s wedge. This is shown in (c) and (d).

surface, causing a degree of cancellation and reducing the amplitude of the integral. Complete cancellation is achieved at the 15mm transducer size for the 1350m/s cone. The wavefield interacting with the transducer when this cancellation occurs is shown in Figs. A.6(c) and A.6(d). In this case it is clear that the Huygens sum is performed across a complete cycle of a sine wave, which cancels completely.

Between the cancellation points, maxima occur, which decay as the transducer size

increases. At each successive maximum, an additional wave cycle will be present on the surface of the transducer, meaning that the remaining non-cancelling section is constrained to a increasingly small fraction of the surface. Relative to the measured amplitude of the incident, undeflected wave, which will increase linearly as the surface area increases, each successive maxima will therefore have a lower amplitude.

One of the assumptions made in eq. (A.4) is that all energy is refracted by the boundary of the cone; i.e. there is no reflection. The close matching between the FDTD and the analytical values demonstrates that this assumption is reasonable for the levels of contrast used here. This is one area where a density contrast could have an effect on the result by increasing the boundary impedance and hence increasing the reflection coefficient. However, assuming that density contrasts are low this will not have a significant effect on the results.

When a wave travels from a slower medium to a faster one, total internal reflection occurs whenever the critical angle $\theta_c = \arcsin(c_1/c_0)$, is exceeded. For a high sound speed cone, this can only occur at the first interface, when the wave enters the cone. Conversely for cones with sound speeds less than that of the background, total internal reflection can only occur when the wave leaves the cone. For the fairly low contrasts and boundary angles being considered here, the critical angle is unlikely to be reached. However, it could be encountered when larger angles occur in the full 3D model, as will be considered later.

A.4.2 Finite transmitter and receiver

The plane wave illumination used in the previous section is experimentally difficult to reproduce, so now the analysis is extended to account for a finite-sized transducer.

Here I demonstrate that the main effect the cone has on the beam from a finite-sized transducer is to alter its direction while leaving its shape and amplitude unchanged, provided the cone sound speed is relatively low contrast compared to the background. In other words, the beam is rigidly rotated by δ_2 , the deflection angle calculated in

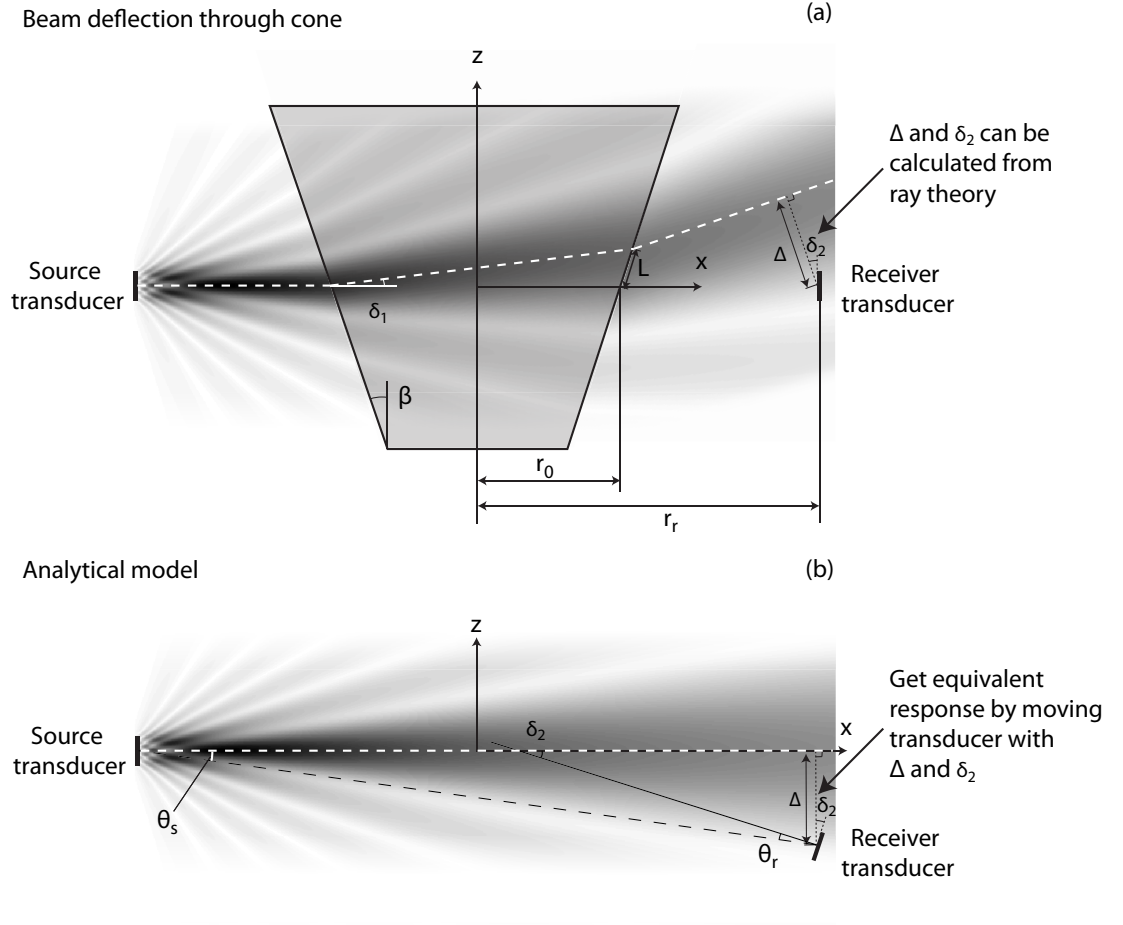


Figure A.7: Analytical model for estimating transmission response. The effects of the refraction of the cone (a) can be estimated from the straight beam as shown in (b).

eq. (A.3). Therefore, it is possible to decouple ray-based refraction caused by the cone, and beam diffraction from the transducer aperture.

To predict the amplitude at the receiver one can proceed as illustrated in Fig. A.7. The effect on the measured signal of the transducer beam being deflected upwards is modelled by having a straight transducer beam and moving the receiver transducer to the equivalent position relative to the beam, as shown in Fig. A.7(b). Huygen's summation can then be used across the surface of the receiver to estimate the measurement that would be taken.

δ_2 was calculated for the plane wave case in the previous section. When accounting for the beam, it is necessary to calculate the location offset, Δ . Through trigonometry, this becomes

$$\Delta = [L \cos \beta + (r_r - L \sin \beta - r_0)] \cos \delta_2 \quad (\text{A.5})$$

where r_0 is the radius of the cone in the array plane, r_r is the distance between the centre of the cone and the receiver transducer and

$$L = \frac{2r_0 \sin \delta_1}{\cos(\beta + \delta_1)} \quad (\text{A.6})$$

is the distance between the point the ray leaves the cone and the intersection of the cone with the x axis.

To verify that the model of Fig. A.7 produces the correct trends, it is compared to the output from a 2D FDTD simulation. Although most parameters are fixed for a single FDTD simulation, by recording time traces at many points in the z direction at the line of the receiver transducer, many different sized transducers can be modelled at the post-processing stage by summing across particular sections of the line. Of course this means that – in general – the receiver transducer will not be the same size as the source, which is unlikely to be the case in reality, but the data produced will enable a comparison to be made between the analytical model and FDTD to determine accuracy. Figure A.8 performs this comparison.

The results show excellent agreement. Even at the higher contrast (where the approximations are less valid) the analytical model is close to the FDTD response and captures the same trends. This indicates that for the ranges being considered here, the underlying assumptions of the analytical model, that the refraction causes the beam to deflect without changing shape or amplitude, are reasonable. It is therefore possible to assume that the trends predicted by the model are valid.

The combined ray/source diffraction approach is similar to that presented in [118, 119]. Here, instead of a single ray, however, separate rays are drawn from each elemental area of the source transducer to the point of interest. Each of these




Figure A.8: Comparison of the analytical model of Fig. A.7 and the FDTD data for a range of different receiver transducer sizes. Normalisation is calculated relative to the same measurement without the cone present. The source transducer height is 12mm and $\beta = 30$ degrees.

elemental areas is considered a point source and the response at the point of interest corresponding to each is determined by the distance calculated from the ray path. By summing all these responses the field can be calculated at the point of interest.

This approach is more accurate than the one given here, since it does not assume that the beam will rigidly rotate upon being refracted. The benefit of this approach is that the refraction effect is separated completely from the diffraction of the beam; this allows the parameters relating to the transducer to be separated from the parameters which control the path of the beam. As well as being easier to generalise, this method is also faster since only a single ray path needs to be determined.

The approach introduced in this section and illustrated in Fig. A.7 provides an important insight into the mechanism of ultrasound transmission through the cone. In fact, depending on the extent of beam deflection, it is possible to define two regimes according to whether the equivalent receiver transducer location lies inside the main lobe of the source beam or not. If the receiver transducer lies within this region, the amplitude of the transmitted beam varies little across the receiver

aperture, and the receiver response depends on the phase distribution as in the case of the plane wave illumination. On the other hand, outside the main lobe, the amplitude of the transmitted field drops significantly and experiences very rapid variations in space. As a result, the receiver response will be much weaker and strongly dependent on transducer position as well as the phasing effect. Clearly, this condition is not suitable for measurements aimed at probing the interior of the cone as in breast ultrasound tomography because signals would be highly sensitive to the position of the breast relative to the transducers. Next the condition that separates the two regimes is defined.

Using the Fraunhofer approximation [68] to model the far field of the beam, it is possible to estimate, based on Δ and the distance between the transducers, whether the receiver is within the central lobe. The equivalent angle between the centre line of the source transducer and the receiver of Fig. A.7(b) can be calculated as

$$\theta_s \approx \frac{\Delta}{r_s + r_r} \quad (\text{A.7})$$

taking the distance from the source transducer to the centre of the cone as r_s . The limiting angle of the central lobe – i.e. the first zero – can be established as

$$\theta_s^{lim} = \frac{\lambda}{h} \quad (\text{A.8})$$

with h defined as the height of the transducer and λ as the wavelength [68]. Therefore, the condition for the receiver to be within the main beam of the source is

$$\frac{\Delta}{r_s + r_r} \frac{h}{\lambda} < 1. \quad (\text{A.9})$$

It is assumed that the model of Fig. A.7 is accurate provided the condition in eq. (A.9) is satisfied. Outside this region, the amplitude will depend on the amplitudes of the grating lobes. An estimate of the maximum expected amplitude can be obtained by considering the peak of the first grating lobe. This has amplitude 0.217 in the far field using the Fraunhofer approximation [68], so it is assumed that

the maximum amplitude measured outside the region defined by eq. (A.9) is around 20% of the undeflected value, i.e. a significant fraction of the energy is lost. Given the low amplitude, the results in this region are not extensively studied further.

A.4.3 Parametric study

Under the decoupled beam diffraction/cone refraction model, the large number of parameters of Sec. A.2 can be significantly simplified and reduced. In particular, it can be concluded that the refraction is controlled purely by the material properties of the cone whereas the beam shape is controlled purely by the transducer characteristics.

To study the transmission phenomenon, the total angular deflection δ_2 is first determined from eqs. (A.2) and (A.3). It is possible to calculate Δ exactly with eq. (A.5), but this would mean maintaining a large number of parameters, making the results difficult to generalise. Assuming that the transducers are at a large distance from the origin of the Cartesian coordinate frame, it follows that $\theta_s = \theta_r = \delta_2/2$ where θ_s and θ_r are defined as in Fig. A.7(b).

Using the Fraunhofer approximation [68], the field produced by a transducer is

$$u(\mathbf{p}) = \frac{\exp(ik_0 R)}{\sqrt{R}} \text{sinc}\left(\frac{k_0 h \sin \theta_s}{2}\right) \quad (\text{A.10})$$

where R and θ_s are the polar coordinates of point \mathbf{p} relative to the centre of the source transducer, taking $\theta_s = 0$ perpendicular to the surface. The response for a receiver transducer can be calculated by integrating the field in eq. (A.10) across its surface in the same way as for the plane wave model in eq. (A.4). Since this case lies in first regime of transmission, the amplitude remains approximately constant across the surface, so only the $\exp(ik_0 R)$ term varies. If the Fraunhofer approximation is then applied again to approximate this term in the far field as a plane wave, it can




Figure A.9: The effect of transducer size on the response with different δ_2 values. The transmitter and receiver have the same size.

then be integrated across the receiver surface, giving the relative response as

$$T = \text{sinc}^2 \left[\frac{k_0 h \sin(\delta_2/2)}{2} \right], \quad (\text{A.11})$$

where it is assumed $\theta_s = \theta_r = \delta_2/2$ from before.

Under this approximation, it is possible to simplify the validity criterion of eq. (A.9). The first zero of the beam function in eq. (A.10) is taken as the boundary of the beam, giving the criterion as

$$\frac{h}{\lambda} \left| \frac{\delta_2}{2} \right| < 1. \quad (\text{A.12})$$

Figure A.9 plots the response as a function of transducer size at different values of δ_2 . This is done using the far field approximation of eq. (A.11) and, for comparison, using the more exact Huygens approach for the receiver positioned at $r_r = 100\lambda$.

The approximation shows very good agreement for most of the range. For each δ_2 curve, however, there is a slight divergence just before the validity limit of eq. (A.12) breaks. This suggests that the receiver has become sufficiently tall that it is no longer entirely contained within the main beam and the locally plane wave

Table A.1: Actual responses for the cone simulations of Sec. A.3 compared to the estimated responses from eq. (A.11). Responses are given as percentages of the signals measured without the cone present. Both the transmitter and receiver have width 4mm and height 20mm.

Cone sound speed (m/s)	δ_1 (rad)	δ_2 (rad)	Estimated response (%)	Actual response (%)
1350	0.0956	0.2442	< 20	10
1440	0.0392	0.0855	53	44
1470	0.0198	0.0413	87	85
1530	-0.0202	-0.0389	88	83
1560	-0.0408	-0.0757	61	51
1650	-0.1058	-0.1772	< 20	17

assumption breaks down. It is therefore important to recognise the limitation of eq. (A.12) as the transducers become larger.

The response from eq. (A.11) is compared to the 3D simulations of Fig. A.3(b). For each of the cones simulated, the δ_2 values are calculated by solving eqs. (A.2) and (A.3). The response is then calculated using eq. (A.11). These estimates are compared to the results of the 3D simulation in Tab. A.1.

A good estimate of the response for a variety of sound speeds above and below that of the background is obtained. The estimated responses are typically within around 10% of the 3D simulations; more importantly the trends predicted by the analytical model are shown to be valid.

The process for estimating the transducer response can be summarised as follows:

1. The total deflection δ_2 is calculated from the material properties and geometry of the cone according to eqs. (A.2) and (A.3).
2. The criterion $h |\delta_2/2| / \lambda < 1$ is applied to check whether the receiver lies within the main beam. If not, it can be assumed that there will be a significant – at

least 80% – amplitude drop.

3. If the receiver does lie within the main beam, the response can be estimated from eq. (A.11), making the assumption that the receiver is in the far field.

A.5 Deflection outside the transducer plane

The previous section has considered the deflection of the transducer beam as it passes through the centre of the cone. However, as can be seen by the ‘smile’ shapes of Fig. A.3 in Sec. A.3, the deflection caused by the cone increases outside this plane. Here the ray theory is extended to explain this trend for a fixed source transducer, then the effect this will have on the measured amplitude is considered.

A.5.1 The ‘smile’ effect

The transducer is treated as a point source with a series of rays diverging from it, each of which is then projected through the cone.

The previous section demonstrated how the deflection caused by the cone along the central plane could be predicted by ray theory. In Snell’s Law, the incident angle is the angle of the incident ray relative to the normal of the surface separating the two materials. For the interaction with the cone, the incident angle has to be treated as a combination of two angles in different planes, as shown in Fig. A.10. The resulting compound incident angle, θ , will therefore be larger so a greater deflection will occur.

By using trigonometry to calculate the incident angles and utilising Snell’s Law, it is possible to calculate the path within the cone, how it interacts with the second interface and where it ultimately intersects the output plane. Figure A.11 shows three such ray paths.

As the incident angle in the x - y plane increases from lines 1 to 2 to 3, it can be seen that the overall deflection increases because of the compound incident angles. This




Figure A.10: 3D incident angles. The angle θ is a combination of θ_{xy} and θ_{rz} such that $\sin \theta = \sin \theta_{xy} \sin \theta_{rz}$.

results in the rays' intersection with the output plane forming a 'smile' pattern, which matches the behaviour of the 3D FDTD simulation, as was visible in the higher contrast transmitted fields of Fig. A.3.

A.5.2 Amplitude drop outside the transducer plane

Having used Snell's law to determine the path of the ray through the cone, the corresponding amplitude drop is now calculated. The decoupling approach used before, separating beam diffraction from cone refraction, can be applied here too. The additional dimension, however, adds significant complexity, since it is necessary to account for two angular deflections and two displacement offsets in the horizontal and vertical planes.

It can be observed that the far field of a rectangular transmitter transducer in 3D




Figure A.11: Paths through the cone plotted for three different rays. The deflection is clearly increased as two incident angles are combined. A ‘smile’ pattern can be seen in a similar way to the result of the 3D FDTD model in Fig. A.3.

is the product of the fields of two linear apertures in 2D [68]

$$U(x, z) = \frac{e^{iky}}{i\lambda y} A \left[\text{sinc} \left(\frac{wx}{\lambda y} \right) \exp \left(\frac{ikx^2}{2y} \right) \right] \left[\text{sinc} \left(\frac{hz}{\lambda y} \right) \exp \left(\frac{ikz^2}{2y} \right) \right] \quad (\text{A.13})$$

where $A = wh$ is the area of the transducer. The field is the product of the field of a linear aperture in the vertical plane and that of a linear aperture in the horizontal plane. This suggests that it is possible to treat the amplitude drop problem by considering two separate 2D problems. First the transmission amplitude due to a prism equal to the cross section of the cone in the vertical plane joining the




Figure A.12: Estimates of the transducer response as the transducer is moved away from the transducer plane. The solid lines are the values derived by summing sections of the FDTD responses of Fig. A.8. The dashed lines are the analytical estimates.

transmitter and receiver is estimated. Then the transmission amplitude due to a circular cylinder is calculated. In each case the amplitude is normalised with respect to the transmission measurement without the scatterer. Finally the two transmission amplitudes are multiplied to obtain the 3D normalised transmission coefficient.

Figure A.12 gives this response for transducers in different y locations in the $x = 80\text{mm}$ output plane of Fig. A.2. The response is relative to the result which would be obtained if the transducers were held in the same locations and the cone was removed. Compared are the transducer responses calculated from the ray-based approximation outlined above and the responses calculated by summing the results of the 3D FDTD simulation across the rectangular aperture of the receiver at different positions along $\hat{\mathbf{y}}$. In all cases the receiver transducers remain orientated to face in the $-\hat{\mathbf{x}}$ direction, with their faces matching the size of the source: 20mm tall and 4mm wide.

The FDTD lines and the ray approximations are shown to match reasonably well – even the 1350m/s case which is outside the ‘central beam’ threshold, causing the significant relative offset at $y = 0$. In all cases the normalised signal can be seen to drop as the receiver moves away from the transducer plane. This is caused by the additional ray deflection which occurs towards the edge of the cone, resulting in a lower relative signal. It should be emphasised that normalisation is performed relative to the transducer measurement of the undisturbed beam in the *same position*, so this amplitude drop is not a result of transducer beam spreading in the horizontal plane. The drop shown in Fig. A.12 as y increases is therefore purely a result of the presence of the cone deflecting the waves.

The compound angle effect present in the 3D model increases the possibility of the critical angle being reached. This is particularly true for rays hitting the cone boundary close to grazing incidence. For cone sound speeds greater than the background, total internal reflection can occur as the ray enters the cone, but for cone sound speeds of less than the background, this can happen as the ray attempts to leave the cone. In the model performed here, such rays are disregarded. It would be expected at these low contrasts that these effects will only be present in a limited section of the parameter space, so in general it is possible to consider them negligible; the FDTD results of Fig. A.3 show no evidence of extensive total internal reflection.

A.6 Conclusions

The problem of transmission through a penetrable acoustic cone has been investigated to determine the effects on practical ultrasound measurements.

I have studied the combined effect of diffraction occurring at the aperture of the transducers and refraction through the cone. The beam emerging from the cone can be approximated by rigidly rotating the beam in the free background medium through an angle that can be estimated by Snell’s law.

There are two regimes that characterise the strength of transmission. In the first

regime the receiver is contained within the main lobe of the beam radiating from the transmitter. Here the incident beam intersecting the surface of the receiver can be approximated as a plane wave and the strength of transmission measurement depends on the phase distribution along the height of the receiver. In the second regime, the receiver is outside the main lobe and the phase oscillations along the receiver are accompanied by large amplitude variations that lead to a significant transmission amplitude drop.

I have provided an analytical expression that marks the transition between the two regimes. Moreover, by using the plane wave approximation within the first regime, a surprisingly simple solution can be derived, allowing one to predict the amplitude drop for the many parameters of the problem directly. Outside the first regime, the amplitude drop is in excess of 80% of the incident field.

Transmission with two diametrically opposed transducers can be studied considering a 2D problem based on the vertical cross-section of the cone. However, off-axis transmission requires the consideration of an additional 2D problem defined by the horizontal cross-section of the cone. Combined, these account for amplitude drop in the vertical and horizontal directions respectively; the total drop is given by the product of these two values.

The analytical trend and physical approximation introduced reported will be of significant practical benefit in the design of future breast ultrasound tomography systems.

Appendix B

Beamforming to diffraction tomography for a non-circular array

Many configurations do not correspond to the circular array (or a section of a circular array in the case of a limited view array) considered throughout this thesis. An alternative is therefore needed to deal with more general array configurations. Here I address the modifications which must be made to make the HARBUT implementation suitable for a general array, starting by considering the BF to DT filter discussed in Chapter 2.

Simonetti and Huang [34] derived the filter by analytically comparing the point spread function of the BF algorithm in the far field for a circular array with that of DT. Here, I introduce a more general approach to deriving the filter which can be extended to other array configurations. I begin with the BF equation from Chapter 2,

$$I^{BF}(\mathbf{z}) = \int_{\mathbb{S}} \int_{\mathbb{S}} \frac{\psi_s(\mathbf{x}, \mathbf{y})}{G_u(\mathbf{x}, \mathbf{z}) G_u(\mathbf{z}, \mathbf{y})} d\mathbf{x} d\mathbf{y}. \quad (\text{B.1})$$

If the problem can be treated as far field, plane wave Green's functions can be

substituted

$$I^{BF}(\mathbf{z}) = \int_{\mathbb{S}} \int_{\mathbb{S}} \frac{\psi_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}})}{e^{-ik_u \hat{\mathbf{s}} \cdot \mathbf{z}} e^{ik_u \hat{\mathbf{s}}_0 \cdot \mathbf{z}}} d\hat{\mathbf{s}}_0 d\hat{\mathbf{s}} \quad (\text{B.2})$$

$$= \int_{\mathbb{S}} \int_{\mathbb{S}} e^{ik_u(\hat{\mathbf{s}} - \hat{\mathbf{s}}_0) \cdot \mathbf{z}} \psi_s(\hat{\mathbf{s}}_0, \hat{\mathbf{s}}) d\hat{\mathbf{s}}_0 d\hat{\mathbf{s}}. \quad (\text{B.3})$$

This has similarities to the standard inverse 2D Fourier transform

$$F(\mathbf{x}) = \int e^{i\mathbf{k} \cdot \mathbf{x}} \tilde{F}(\mathbf{k}) d\mathbf{k} \quad (\text{B.4})$$

where the function is also multiplied by the exponential term and integrated. The difference lies in how the integration is performed. The standard inverse Fourier transform does this linearly across the \mathbf{k} -space, however, the BF algorithm instead performs the integration along the circular arcs mapped out in the \mathbf{k} -space as the source ($\hat{\mathbf{s}}_0$) and receiver ($\hat{\mathbf{s}}$) vectors are varied. This introduces a weighting into the transformed data. This weighting causes the distortion in BF, which is removed by Simonetti and Huang's filter to convert into DT.

It is possible to perform a coordinate transform in the integral, changing from integrating around the different directions of the transducer array, as in eq. (B.3), to integrating across the \mathbf{k} -space, as in the standard 2D Fourier transform (B.4). The Jacobian determinant of the transform will define the weightings applied to each spatial frequency component, and hence should equal the BF to DT filter. The derivation is straightforward and will not be included here for brevity; the weighting becomes

$$w(k_r) = \frac{1}{k_r k_u \sin \beta} \quad (\text{B.5})$$

where β is the angle between the source and receiver vectors and $k_r = k_u |\hat{\mathbf{s}}_0 - \hat{\mathbf{s}}|$, with k_u being the wavenumber in the uniform background as defined in Chapter 2.




Figure B.1: Coordinates for calculating BF weightings for a linear transducer array

k_r is related to β

$$\cos \beta = \frac{k_r}{2k_u} \quad (\text{B.6})$$

which allows the weighting to be rewritten as

$$w(k_r) = \frac{1}{k_r k_u \sqrt{1 - \left(\frac{k_r}{2k_u}\right)^2}}. \quad (\text{B.7})$$

This matches the result from [34], although there are some constants which result from the Fourier transforms and far field normalisations which are not considered here. This method of deriving the filter is more flexible than the original technique, making it suitable for a variety of configurations.

The previous analysis assumes that measurements are uniformly taken at equiangular intervals around the scatterer, which is only true for a circular, far field array. If the angular spacing is no longer uniform, the integral of eq. (B.3) will be further distorted, so additional weightings would need to be incorporated into the filter.

This is demonstrated for a linear array. The parameters are defined in Fig. B.1, such that

$$l = D \tan \alpha. \quad (\text{B.8})$$

In order to incorporate this into the weighting, the rate of change along the length of the array as the angle, α , changes needs to be included in the Jacobian term. This is calculated as

$$\frac{dl}{d\alpha} = D \sec^2 \alpha. \quad (\text{B.9})$$

This can be easily incorporated into the filter in the k-space by first calculating α for the linear source array, calculating $dl/d\alpha$, doing the same for the linear receiver array, then multiplying them both by the standard weighting of eq. (B.7) to obtain the final overall weighting.

This filtering approach is built on the far field assumption: since the array is distant from all points in the image, the illumination angle corresponding to a particular transducer will remain constant as the imaging point changes. In the near field the angles of the transducer array vary with position. In order to correctly account for this, each pixel in the image would, theoretically, need its own filter, which is impossible. A better solution is to incorporate some or all of the weightings into the beamforming integral of (B.1) rather than applying a filter as a post-processing stage. Since each pixel is calculated separately, the correct weightings for each position can be included at this stage rather than having to use a single set of weightings for the whole image, as is the case for the filter.

It is noted that in the earlier chapters, successful images were reconstructed from a circular array without accounting for the positional variation of the filter. Figure B.2 illustrates why this is possible. Figures B.2(b) and (c) show how the filter remains relatively constant for a circular array as its location is changes, explaining why the images of the earlier chapters did not exhibit any significant filtering artefacts. For a square array, however, the filter is much more complex, and can no longer be approximated as constant in space, as shown in Figs. B.2(e) and (f). The offset filter of Fig. B.2(f) becomes discontinuous, and the discontinuity magnitudes and locations depend on the position, making the use of a single filter for all parts of the image inappropriate. In this case it is more appropriate to include angular weightings directly in the beamforming integral.

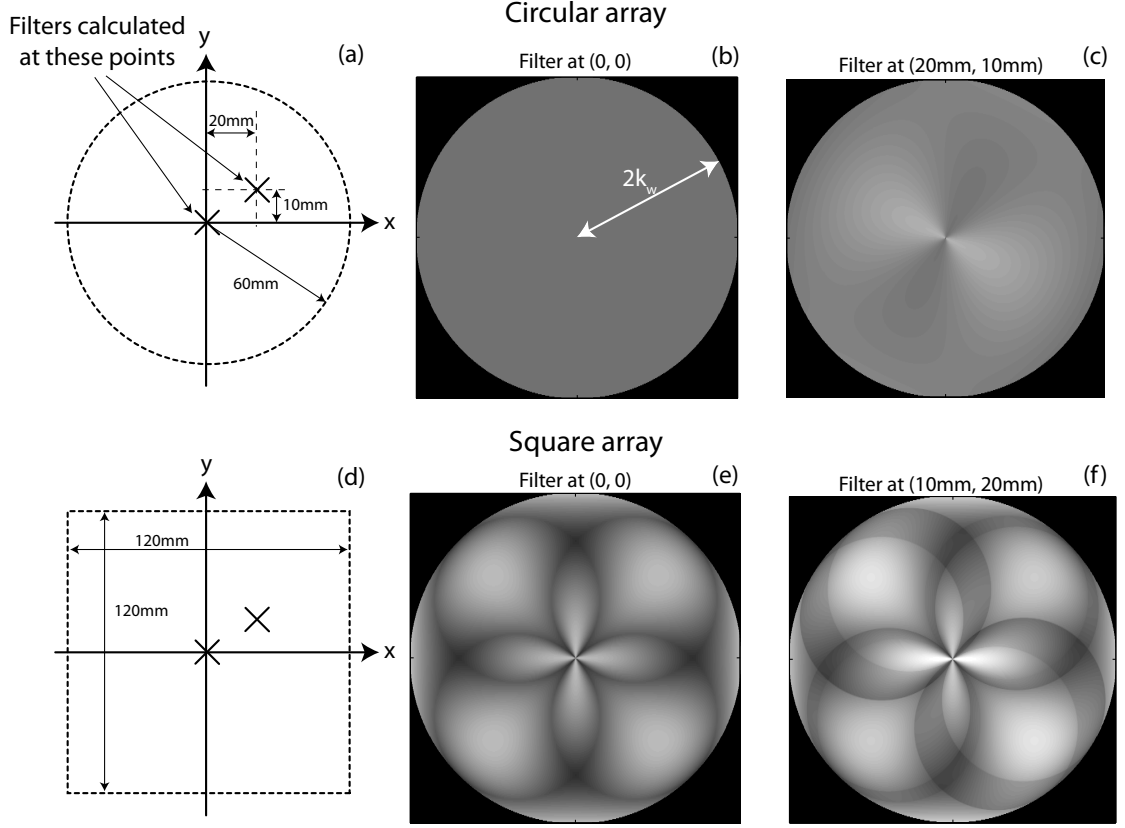


Figure B.2: Near field position dependency of filters. Note that just the angular spacing component is present; the weightings of eq. (B.7) are not included for clarity. (a) and (d) show the circular and square array configurations respectively, and mark the two points at which the filter is calculated. (b), the circular array filter at the centre, and (c), the circular filter at the offset location, show that the filter remains relatively constant when the position changes for the circular array. (e) shows that the square array filter is far more complex, and (f) shows that as the position varies, discontinuities appear in the filter which move as a function of position; as such the filter for the square array cannot be approximated as constant in space.

Finally it should be noted that the implementation of the algorithm relies on discretising the BF integral of (B.1). This discretisation provides a convenient mechanism to correct for irregular angular spacings for general arrays. When the integration is discretised, each component value is multiplied by the angular ‘contribution’ of each transducer and summed to integrate around 2π radians

$$I^{BF}(\mathbf{z}) \approx \sum_{i=1}^N \sum_{r=1}^N \frac{\psi_s(\mathbf{x}_i, \mathbf{y}_r)}{G_u(\mathbf{x}_i, \mathbf{z}) G_u(\mathbf{z}, \mathbf{y}_r)} \delta\theta_i \delta\theta_r \quad (\text{B.10})$$




Figure B.3: The angular ‘contribution’ of a transducer, used in the BF integral discretisation. Note that the extent of the transducer here is considered to be from halfway from the previous transducer to halfway to the next transducer; this definition is such that the sum of all angular contributions should sum to 2π if a full view configuration is used, or to the angular extent of a limited view array.

where N is the number of transducers, i and r are illumination and reception transducers respectively, and $\delta\theta_i$ and $\delta\theta_r$ are the angular contributions.

The angular contributions are defined as in Fig. B.3. The surface of the array is divided into lengths associated with each transducer, importantly with no gaps between; as such the boundary of each length is defined as the halfway point between two transducers, rather than being the physical limit of the transducer. The angle subtended for each length defines each $\delta\theta$ term. If a full view configuration is used, the sum of all these angles should be 2π ; if not then the sum will be equal to the angular extent of the array.

By using eq. (B.10), the irregular angular spacing of the transducers is automatically accounted for. Accurate sound speed maps can then be obtained from this image via the standard BF to DT filter.

References

- [1] A. Stanton, “Wilhelm conrad röntgen on a new kind of rays: translation of a paper read before the würzburg physical and medical society,” *Nature*, vol. 53, no. 1369, pp. 274–276, 1896.
- [2] I. Donald, J. MacVicar, and T. Brown, “Investigation of abdominal masses by pulsed ultrasound,” *Lancet*, vol. 1, no. 7032, p. 1188, 1958.
- [3] G. Hounsfield, “Computerized transverse axial scanning (tomography): Part 1. description of system,” *British Journal of Radiology*, vol. 46, no. 552, p. 1016, 1973.
- [4] H. Carr and E. Purcell, “Effects of diffusion on free precession in nuclear magnetic resonance experiments,” *Physical Review*, vol. 94, no. 3, p. 630, 1954.
- [5] B. Drinkwater and P. Wilcox, “Ultrasonic arrays for non-destructive evaluation: A review,” *NDT & E International*, vol. 39, no. 7, pp. 525 – 541, 2006.
- [6] Rapiscan, “Rapiscan secure 1000 dual pose.” http://www.rapiscansystems.com/en/products/item/productsrapiscan_secure_1000_dual_pose, 2012. Accessed 6/1/12.
- [7] M. Skolnik, *Radar handbook*, vol. 2. McGraw-Hill New York, 1990.
- [8] P. Fish, *Physics and instrumentation of diagnostic medical ultrasound*. Wiley, 1990.
- [9] R. Stewart, *Exploration Seismic Tomography*. Society of Exploration Geophysicists, 1991.

-
- [10] P. Docherty, "A brief comparison of some kirchhoff integral formulas for migration and inversion," *Geophysics*, vol. 56, no. 8, p. 1164, 1991.
- [11] J. Ophir, I. Cespedes, H. Ponnekanti, Y. Yazdi, and X. Li, "Elastography: a quantitative method for imaging the elasticity of biological tissues," *Ultrasonic imaging*, vol. 13, no. 2, pp. 111–134, 1991.
- [12] P. Barbone and N. Gokhale, "Elastic modulus imaging: on the uniqueness and nonuniqueness of the elastography inverse problem in two dimensions," *Inverse Problems*, vol. 20, p. 283, 2004.
- [13] J. Bercoff, M. Tanter, and M. Fink, "Supersonic shear imaging: a new technique for soft tissue elasticity mapping," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 51, no. 4, pp. 396–409, 2004.
- [14] F. Simonetti and L. Huang, "Synthetic aperture diffraction tomography for three-dimensional imaging," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, vol. 465, no. 2109, p. 2877, 2009.
- [15] H. Gemmeke, L. Berger, M. Birk, G. Gobel, A. Menshikov, D. Tcherniakhovski, M. Zapf, and N. Ruiter, "Hardware setup for the next generation of 3d ultrasound computer tomography," in *Nuclear Science Symposium Conference Record (NSS/MIC), 2010 IEEE*, pp. 2449–2454, IEEE, 2010.
- [16] M. Riccabona, T. Nelson, and D. Pretorius, "Three-dimensional ultrasound: accuracy of distance and volume measurements," *Ultrasound in Obstetrics & Gynecology*, vol. 7, no. 6, pp. 429–434, 1996.
- [17] J. Blackall, D. Rueckert, C. Maurer, G. Penney, D. Hill, and D. Hawkes, "An image registration approach to automated calibration for freehand 3d ultrasound," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2000*, Springer, 2000.
- [18] T. Nelson, D. Pretorius, A. Hull, M. Riccabona, M. Sklansky, and G. James, "Sources and impact of artifacts on clinical three-dimensional ultrasound imaging," *Ultrasound in obstetrics & gynecology*, vol. 16, no. 4, pp. 374–383, 2000.

- [19] A. Kak and M. Slaney, *Principles of computerized tomographic reconstruction*. New York: IEEE Press, 1998.
- [20] J. Greenleaf, S. Johnson, S. Lee, G. Herman, and E. Wood, “Algebraic reconstruction of spatial distributions of acoustic absorption within tissue from their two-dimensional acoustic projections,” *International Symposium on Acoustical Holography and Imaging, 5th, Palo Alto, Calif*, vol. 5, pp. 591–603, 1973.
- [21] N. Duric, P. Littrup, L. Poulo, A. Babkin, R. Pevzner, E. Holsapple, O. Rama, and C. Glide, “Detection of breast cancer with ultrasound tomography: First results with the computed ultrasound risk evaluation (cure) prototype,” *Medical physics*, vol. 34, p. 773, 2007.
- [22] S. Li, M. Jackowski, D. Dione, T. Varslot, L. Staib, and K. Mueller, “Refraction corrected transmission ultrasound computed tomography for application in breast imaging,” *Medical Physics*, vol. 37, p. 2233, 2010.
- [23] D. Jansen and D. Hutchins, “Lamb wave tomography,” in *Ultrasonics Symposium, 1990. Proceedings., IEEE 1990*, pp. 1017–1020, IEEE, 1990.
- [24] E. Malyarenko and M. Hinders, “Fan beam and double crosshole lamb wave tomography for mapping flaws in aging aircraft structures,” *The Journal of the Acoustical Society of America*, vol. 108, p. 1631, 2000.
- [25] P. Williamson, “A guide to the limits of resolution imposed by scattering in ray tomography,” *Geophysics*, vol. 56, no. 2, pp. 202–207, 1991.
- [26] F. Natterer and F. Wubbeling, *Mathematical methods in image reconstruction*. Society for Industrial Mathematics, 2001.
- [27] J. Wiskin, D. Borup, S. Johnson, M. Berggren, T. Abbott, and R. Hanover, “Full-wave, non-linear, inverse scattering,” *Acoustical Imaging*, pp. 183–193, 2007.
- [28] D. Colton and R. Kress, *Inverse acoustic and electromagnetic scattering theory*, vol. 93. Springer Verlag, 1998.

- [29] P. Morse and K. Ingard, *Theoretical Acoustics*. New York, London: McGraw-Hill Book Company, 1968.
- [30] S. Pourjavid and O. Tretiak, “Numerical solution of the direct scattering problem through the transformed acoustical wave equation,” *The Journal of the Acoustical Society of America*, vol. 91, p. 639, 1992.
- [31] M. Born and E. Wolf, *Principles of Optics*. Cambridge: Cambridge University Press, 1999.
- [32] T. Mast, “Empirical relationships between acoustic parameters in human soft tissues,” *Acoustics Research Letters Online*, vol. 1, p. 37, 2000.
- [33] W. Chew and Y. Wang, “Reconstruction of two-dimensional permittivity distribution using the distorted born iterative method,” *IEEE Transactions on Medical Imaging*, vol. 9, no. 2, pp. 218–225, 1990.
- [34] F. Simonetti and L. Huang, “From beamforming to diffraction tomography,” *Journal of Applied Physics*, vol. 103, p. 103110, 2008.
- [35] S. Li, K. Mueller, M. Jackowski, D. Dione, and L. Staib, “Fast marching method to correct for refraction in ultrasound computed tomography,” in *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006*, pp. 896–899, 2006.
- [36] J. Sethian, “A fast marching level set method for monotonically advancing fronts,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, no. 4, p. 1591, 1996.
- [37] W. Schneider, “Integral formulation for migration in two and three dimensions,” *Geophysics*, vol. 43, p. 49, 1978.
- [38] F. Kamangar, G. Dores, and W. Anderson, “Patterns of cancer incidence, mortality, and prevalence across five continents: defining priorities to reduce cancer disparities in different geographic regions of the world,” *Journal of Clinical Oncology*, vol. 24, no. 14, p. 2137, 2006.

- [39] P. Autier, M. Boniol, C. LaVecchia, L. Vatten, A. Gavin, C. Hery, and M. Heanue, “Disparities in breast cancer mortality trends between 30 european countries: retrospective trend analysis of who mortality database,” *British Medical Journal*, vol. 341, p. 3620, 2010.
- [40] M. Silverstein, M. Lagios, A. Recht, D. Allred, S. Harms, R. Holland, D. Holmes, L. Hughes, R. Jackman, T. Julian, *et al.*, “Image-detected breast cancer: state of the art diagnosis and treatment,” *Journal of the American College of Surgeons*, vol. 201, no. 4, pp. 586–597, 2005.
- [41] T. Kolb, J. Lichy, and J. Newhouse, “Comparison of the performance of screening mammography, physical examination, and breast us and evaluation of factors that influence them: An analysis of 27,825 patient evaluation,” *Radiology*, vol. 225, pp. 165–175, 2002.
- [42] P. Stomper, D. D’Souza, P. Dinitto, and *et al.*, “Analysis of parenchymal density on mammograms in 1353 women 25-79 years old,” *AJR Am J Roentgenol*, vol. 167, pp. 1261–1265, 1996.
- [43] D. Saslow, C. Boetes, W. Burke, S. Harms, M. Leach, C. Lehman, E. Morris, E. Pisano, M. Schnall, S. Sener, *et al.*, “American cancer society guidelines for breast screening with mri as an adjunct to mammography,” *CA: A Cancer Journal for Clinicians*, vol. 57, no. 2, p. 75, 2007.
- [44] M. Kriege, C. Brekelmans, C. Boetes, P. Besnard, H. Zonderland, I. Obdeijn, R. Manoliu, T. Kok, H. Peterse, M. Tilanus-Linthorst, *et al.*, “Efficacy of mri and mammography for breast-cancer screening in women with a familial or genetic predisposition,” *The New England journal of medicine*, vol. 351, no. 5, p. 427, 2004.
- [45] P. Carson, C. Meyer, A. Scherzinger, and T. Oughton, “Breast imaging in coronal planes with simultaneous pulse echo and transmission ultrasound,” *Science*, vol. 214, no. 4525, p. 1141, 1981.
- [46] J. Schreiman, J. Gisvold, J. Greenleaf, and R. Bahn, “Ultrasound transmission computed tomography of the breast,” *Radiology*, vol. 150, no. 2, p. 523, 1984.

-
- [47] M. Andre, H. Janee, P. Martin, G. Otto, B. Spivey, and D. Palmer, “High-speed data acquisition in a diffraction tomography system employing large-scale toroidal arrays,” *Int. J. Imag. Syst. Tech.*, vol. 8, pp. 137–147, 1997.
- [48] J. Greenleaf, S. Johnson, W. Samayoa, and F. Duck, “Algebraic reconstruction of spatial distributions of acoustic velocities in tissue from their time-of-flight profiles,” in *Acoustical Holography* (P. S. Green, ed.), vol. 6, pp. 71–90, Plenum Press, New York, 1975.
- [49] J. Greenleaf, S. Johnson, and R. Bahn, “Quantitative cross-sectional imaging of ultrasound parameters,” in *Ultrasonics Symposium, 1977*, pp. 989–995, 1977.
- [50] H. Schomberg, “An improved approach to reconstructive ultrasound tomography,” *Journal of Physics D: Applied Physics*, vol. 11, p. L181, 1978.
- [51] M. Krueger, A. Pesavento, and H. Ermert, “A modified time-of-flight tomography concept for ultrasonic breast imaging,” *IEEE Ultrasonics Symposium*, pp. 1381–1385, 1996.
- [52] Y. Quan and L. Huang, “Sound-speed tomography using first-arrival transmission ultrasound for a ring array,” in *Proc. SPIE*, vol. 6513, 2007.
- [53] A. Hormati, I. Jovanovic, O. Roy, and M. Vetterli, “Robust ultrasound travel-time tomography using the bent ray model,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7629, p. 17, 2010.
- [54] K. Tanabe, “Projection method for solving a singular system of linear equations and its applications,” *Numerische Mathematik*, vol. 17, no. 3, pp. 203–214, 1971.
- [55] A. Louis and F. Natterer, “Mathematical problems of computerized tomography,” *Proceedings of the IEEE*, vol. 71, no. 3, pp. 379–389, 1983.
- [56] S. Hung, F. Dahlen, and G. Nolet, “Wavefront healing: a banana–doughnut perspective,” *Geophysical Journal International*, vol. 146, no. 2, pp. 289–312, 2001.
-

-
- [57] A. Devaney, "A filtered backpropagation algorithm for diffraction tomography," *Ultrasonic imaging*, vol. 4, no. 4, pp. 336–350, 1982.
- [58] F. Simonetti, L. Huang, N. Duric, and P. Littrup, "Diffraction and coherence in breast ultrasound tomography: A study with a toroidal array," *Medical physics*, vol. 36, p. 2955, 2009.
- [59] A. Tarantola, *Inverse problem theory and methods for model parameter estimation*. Society for Industrial Mathematics, 2005.
- [60] R. Lavarello and M. Oelze, "Density imaging using inverse scattering," *The Journal of the Acoustical Society of America*, vol. 125, p. 793, 2009.
- [61] C. Li, N. Duric, P. Littrup, and L. Huang, "In vivo breast sound-speed imaging with ultrasound tomography," *Ultrasound in medicine & biology*, vol. 35, no. 10, pp. 1615–1628, 2009.
- [62] S. Goss, R. Johnston, and F. Dunn, "Comprehensive compilation of empirical ultrasonic properties of mammalian tissues," *The Journal of the Acoustical Society of America*, vol. 64, p. 423, 1978.
- [63] E. Franceschini, S. Mensah, D. Amy, and J. Lefebvre, "A 2-d anatomic breast ductal computer phantom for ultrasonic imaging," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 53, no. 7, p. 1281, 2006.
- [64] F. Simonetti, L. Huang, and N. Duric, "On the sampling of the far-field operator with a circular ring array," *J. Appl. Phys.*, vol. 101, p. 083103, 2007.
- [65] K. Yee, "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media," *IEEE Transactions on antennas and propagation*, vol. 14, no. 3, pp. 302–307, 1966.
- [66] 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, pp. 155–167, 2007.
- [67] G. Guennebaud, B. Jacob, *et al.*, "Eigen v3." <http://eigen.tuxfamily.org>, 2012.
-

-
- [68] J. Goodman, *Introduction to Fourier Optics*. New York: McGraw-Hill, 1996.
- [69] F. Simonetti, L. Huang, and N. Duric, "A multiscale approach to diffraction tomography of complex three-dimensional objects," *Appl. Phys. Lett.*, vol. 95, p. 061904, 2009.
- [70] 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, no. 1, pp. 61–72, 2009.
- [71] X. Viennot, G. Eyrolles, N. Janey, and D. Arques, "Combinatorial analysis of ramified patterns and computer imagery of trees," in *ACM SIGGRAPH Computer Graphics*, vol. 23, pp. 31–40, ACM, 1989.
- [72] D. Alleyne, B. Pavlakovic, M. Lowe, and P. Cawley, "Rapid long-range inspection of chemical plant pipework using guided waves," *Insight*, vol. 43, no. 2, pp. 93–96, 2001.
- [73] R. Dalton, P. Cawley, and M. Lowe, "The potential of guided waves for monitoring large areas of metallic aircraft fuselage structure," *Journal of Nondestructive Evaluation*, vol. 20, no. 1, pp. 29–46, 2001.
- [74] P. Cawley, M. Lowe, D. Alleyne, B. Pavlakovic, and P. Wilcox, "Practical long range guided wave testing: Applications to pipes and rail," *Materials Evaluation*, vol. 61, no. 1, pp. 66–74, 2003.
- [75] J. Michaels and T. Michaels, "Guided wave signal processing and image fusion for in situ damage localization in plates," *Wave Motion*, vol. 44, no. 6, pp. 482–492, 2007.
- [76] J. McKeon and M. Hinders, "Parallel projection and crosshole lamb wave contact scanning tomography," *The Journal of the Acoustical Society of America*, vol. 106, p. 2568, 1999.
- [77] P. Belanger, P. Cawley, and F. Simonetti, "Guided wave diffraction tomography within the born approximation," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 57, no. 6, pp. 1405–1418, 2010.
-

-
- [78] A. Volker and J. Bloom, “Experimental results of guided wave travel time tomography,” in *AIP Conference Proceedings*, vol. 1335, p. 215, 2011.
- [79] J. Pei, M. Yousuf, F. Degertekin, B. Honein, and B. Khuri-Yakub, “Lamb wave tomography and its application in pipe erosion/corrosion monitoring,” *Research in Nondestructive Evaluation*, vol. 8, pp. 189–197, 1996. 10.1007/BF02433949.
- [80] E. Malyarenko and M. Hinders, “Ultrasonic lamb wave diffraction tomography,” *Ultrasonics*, vol. 39, no. 4, pp. 269–281, 2001.
- [81] P. Belanger and P. Cawley, “Lamb wave tomography to evaluate the maximum depth of corrosion patches,” in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 27, pp. 1290–1297, 2008.
- [82] L. Schmerr Jr and S. Song, *Ultrasonic Nondestructive Evaluation Systems*. Springer, 2007.
- [83] P. Huthwaite and F. Simonetti, “High-resolution imaging without iteration: a fast and robust method for breast ultrasound tomography,” *Journal of the Acoustical Society of America*, vol. 130, no. 3, pp. 1721–1734, 2011.
- [84] P. Belanger and P. Cawley, “Feasibility of low frequency straight-ray guided wave tomography,” *NDT & E International*, vol. 42, no. 2, pp. 113–119, 2009.
- [85] R. Pratt, “Seismic waveform inversion in the frequency domain, part 1: Theory and verification in a physical scale model,” *Geophysics*, vol. 64, no. 3, pp. 888–901, 1999.
- [86] P. Mora, “Inversion = migration + tomography,” *Geophysics*, vol. 54, no. 12, pp. 1575–1586, 1989.
- [87] F. Natterer, “Reflectors in wave equation imaging,” *Wave Motion*, vol. 45, no. 6, pp. 776–784, 2008.
- [88] S. Huang and P. Li, “Computed tomography sound velocity reconstruction using incomplete data,” *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 51, no. 3, pp. 329–342, 2004.
-

-
- [89] S. Huang and P. Li, "Ultrasonic computed tomography reconstruction of the attenuation coefficient using a linear array," *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, vol. 52, no. 11, pp. 2011–2022, 2005.
- [90] M. Rantala, S. Vanska, S. Jarvenpaa, M. Kalke, M. Lassas, J. Moberg, and S. Siltanen, "Wavelet-based reconstruction for limited-angle x-ray tomography," *Medical Imaging, IEEE Transactions on*, vol. 25, no. 2, pp. 210–217, 2006.
- [91] S. Gull and J. Skilling, "Maximum entropy method in image processing," *Communications, Radar and Signal Processing, IEE Proceedings F*, vol. 131, no. 6, pp. 646–659, 1984.
- [92] Y. Sung, W. Choi, C. Fang-Yen, K. Badizadegan, R. Dasari, and M. Feld, "Optical diffraction tomography for high resolution live cell imaging," *Optics express*, vol. 17, no. 1, pp. 266–277, 2009.
- [93] M. Persson, D. Bone, and H. Elmqvist, "Total variation norm for three-dimensional iterative reconstruction in limited view angle tomography," *Physics in medicine and biology*, vol. 46, p. 853, 2001.
- [94] E. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Physics in medicine and biology*, vol. 53, p. 4777, 2008.
- [95] G. McMechan, J. Harris, and L. Anderson, "Cross-hole tomography for strongly variable media with applications to scale model data," *Bulletin of the seismological society of America*, vol. 77, no. 6, pp. 1945–1960, 1987.
- [96] Y. Ikedo, D. Fukuoka, T. Hara, H. Fujita, E. Takada, T. Endo, and T. Morita, "Computerized mass detection in whole breast ultrasound images: Reduction of false positives using bilateral subtraction technique," in *Proc SPIE*, vol. 6514, p. 65141T, Citeseer, 2007.
-

- [97] T. Stanton, "Sound scattering by cylinders of finite length. i. fluid cylinders," *Journal of the Acoustical Society of America*, vol. 83, pp. 55–63, 1988.
- [98] Z. Ye, "A novel approach to sound scattering by cylinders of finite length," *The Journal of the Acoustical Society of America*, vol. 102, p. 877, 1997.
- [99] J. Sinai and R. Waag, "Ultrasonic scattering by two concentric cylinders," *The Journal of the Acoustical Society of America*, vol. 83, p. 1728, 1988.
- [100] J. Roumeliotis and N. Kakogiannos, "Acoustic scattering from an infinite cylinder of small radius coated by a penetrable one," *The Journal of the Acoustical Society of America*, vol. 97, p. 2074, 1995.
- [101] J. Hersch and E. Heller, "Observation of proximity resonances in a parallel-plate waveguide," *Physical Review Letters*, vol. 81, no. 15, pp. 3059–3062, 1998.
- [102] T. Tsuei and P. Barber, "Multiple scattering by two parallel dielectric cylinders," *Applied optics*, vol. 27, no. 16, pp. 3375–3381, 1988.
- [103] V. Twersky, "Multiple scattering of radiation by an arbitrary planar configuration of parallel cylinders and by two parallel cylinders," *Journal of Applied Physics*, vol. 23, no. 4, pp. 407–414, 2009.
- [104] L. Rayleigh, *Theory of sound*, vol. II. Dover Publications, New York, 1945.
- [105] V. Anderson, "Sound scattering from a fluid sphere," *The Journal of the Acoustical Society of America*, vol. 22, p. 426, 1950.
- [106] A. Rudgers, "Separation and analysis of the acoustic field scattered by a rigid sphere," *The Journal of the Acoustical Society of America*, vol. 52, p. 234, 1972.
- [107] G. Gaunard, H. Huang, and H. Strifors, "Acoustic scattering by a pair of spheres," *The Journal of the Acoustical Society of America*, vol. 98, p. 495, 1995.

-
- [108] S. Olsson, "Scattering of acoustic waves by a sphere outside an infinite circular cylinder," *The Journal of the Acoustical Society of America*, vol. 88, p. 515, 1990.
- [109] A. Aden and M. Kerker, "Scattering of electromagnetic waves from two concentric spheres," *Journal of Applied Physics*, vol. 22, no. 10, pp. 1242–1246, 2009.
- [110] M. Lyalinov, "Acoustic scattering of a plane wave by a circular penetrable cone," *Wave Motion*, vol. 48, pp. 62–82, January 2010.
- [111] F. Denis, O. Basset, and G. Gimenez, "Ultrasonic transmission tomography in refracting media: reduction of refraction artifacts by curved-ray techniques," *Medical Imaging, IEEE Transactions on*, vol. 14, no. 1, pp. 173–188, 1995.
- [112] M. Lyalinov, "On an integral equation in the problem of diffraction of a plane wave on a transparent circular cone," *Journal of Mathematical Sciences*, vol. 132, no. 1, pp. 56–68, 2006.
- [113] M. Lyalinov and N. Zhu, "Acoustic scattering by a circular semi-transparent conical surface," *Journal of Engineering Mathematics*, vol. 59, no. 4, pp. 385–398, 2007.
- [114] D. Jones, "Rawlin's method and the diaphonous cone," *The Quarterly Journal of Mechanics and Applied Mathematics*, vol. 53, no. 1, p. 91, 2000.
- [115] G. Kino, *Acoustic waves: devices, imaging, and analog signal processing*. Englewood Cliffs, New Jersey: Prentice-Hall Inc., 1987.
- [116] J. Roden and S. Gedney, "Convolution pml (cpml): An efficient fdtd implementation of the cfs-pml for arbitrary media," *Microwave and optical technology letters*, vol. 27, no. 5, pp. 334–339, 2000.
- [117] J. Van Bladel, *Electromagnetic fields*. McGraw-Hill New York, 1964.
- [118] C. Griffice and J. Seydel, "Spherical wave decomposition approach to ultrasonic field calculations," *Journal of Nondestructive Evaluation*, vol. 2, no. 3, pp. 241–247, 1981.
-

- [119] J. Johnson, N. Carlson, and D. Tow, “Ray trace calculations of ultrasonic fields,” *Research in Nondestructive Evaluation*, vol. 3, no. 1, pp. 27–39, 1991.

Publication list

- [P1] P. Huthwaite, F. Simonetti and M.J.S. Lowe. On the Convergence of Finite Element Scattering Simulations. in *Review of Progress in Quantitative Nondestructive Evaluation*, Kingston, Rhode Island, USA. Volume 29, pages 65–72. Edited by D. Chimenti, B. Thompson. 2010.
- [P2] 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 *SPIE Medical Imaging 2010*, San Diego, California, USA. Volume 7629, pages 76290L-1-8. Edited by J. D’hooge, S. A. McAleavey. 2010.
- [P3] P. Huthwaite and F. Simonetti. A Practical, Robust Approach to High Resolution Breast Ultrasound Tomography, in *SPIE Medical Imaging 2011*, Orlando, Florida, USA, 2011. Volume 7968, pages 79680S-1-8. Edited by J. D’hooge, M. M. Doyley. 2011.
- [P4] P. Huthwaite and F. Simonetti. On the Measurement of Ultrasound Transmission Through a Penetrable Acoustic Cone, in *SPIE Medical Imaging 2011*, Orlando, Florida, USA, 2011. Volume 7968, pages 79681H-1-8. Edited by J. D’hooge, M. M. Doyley. 2011.
- [P5] P. Huthwaite and F. Simonetti. High-Resolution Imaging Without Iteration: a Fast and Robust Method for Breast Ultrasound Tomography. *Journal of the Acoustical Society of America*, 130(3):1721-1734, 2011.
- [P6] P. Huthwaite and F. Simonetti. Damage Detection Through Sound Speed Reconstruction, in *Review of Progress in Quantitative Nondestructive Evaluation*,

Burlington, Vermont, USA, 2012. In press.

- [P7] P. Huthwaite and F. Simonetti. On the Measurement of Ultrasound Transmission Through a Penetrable Acoustic Cone. Submitted to *Journal of the Acoustical Society of America*.
- [P8] P. Huthwaite and F. Simonetti. High-Resolution Guided Wave Tomography. Submitted to *Wave Motion*.
- [P9] P. Huthwaite, F. Simonetti and N. Duric. Application of the Hybrid Algorithm for Robust Breast Ultrasound Tomography (HARBUT) to in vivo data. In preparation for submission to *Journal of the Acoustical Society of America*.