

University of Strathclyde
Department of Electronic and Electrical
Engineering

**2D ultrasonic phased arrays
for quantitative characterisation
of complex defects**

Jerzy Dziewierz

A thesis presented in fulfilment of the requirement for
the degree of Doctor of Philosophy

2015

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under terms of United Kingdom copyright acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material in, or, derived form of this thesis.

Signed:

Date:

Acknowledgements

I would like to express my sincerest thanks to my supervisors, Prof. Tony Gachagan and Dr. Richard O'Leary, for accepting me as a PhD student. Thank you for your useful guidance, patience, keen interest and efforts at every time.

I would like to give my special thanks to Dr. John Mackersie, Thomas McCunnie and Grant Smillie for their technical and moral support during my studies.

Finally, I express my deep thanks to my entire family for their unfaltering support.

Abstract

Ultrasonic inspection can be considered as one of many ways the technical system - installation or vehicle - can be made economical and safe. Contemporary ultrasonic systems are capable of detecting a wide variety of mechanical defects and flaws that may or may not affect the operation of a given product. Ultrasonic testing techniques are widely accepted for quality control and material testing. Moreover, the technology is proven, well-understood and widely used.

Upon detecting a flaw, a decision has to be made to ensure the component is fit for the purpose: is the flaw acceptable or is repair of the given part or its replacement? Here, 2D ultrasonic phased arrays hold promise to quickly deliver detailed, 3D resolved information about the extent and nature of the flaw. This information can then be used to develop and justify the technical and economic decision concerning the existing state of the product. In effect, an opportunity exists for significant cost savings by using ultrasonic 2D phased array systems for defect characterisation.

The interest of the work is to establish a process of designing and manufacturing of piezoelectric, 2D phased array ultrasound probes for application in non-destructive evaluation of materials. Furthermore, implementation of practical signal processing method is investigated.

In the first part of the work the sensor mechanical and electrical design is addressed. The properties of piezoelectric ceramic composite materials are studied. Detailed numerical models have been used to analyse conditions required for crafting materials of desirable properties. A novel technique has been demonstrated that allows design of well-behaved triangular cut piezoelectric composite. Built into a single array element (of hexagonal shape by taking 6 triangular pillars) this new composite exhibits properties comparable to a reference rectangular composite (sensitivity of 0.60nm/V for hexagonal, 0.62nm/V for square; and inter-element crosstalk of -21.2dB for hexagonal and -21.9dB for square element). This composite then allows building of compact, dense-layout 2D phased array transducers with hexagonal or sparse element layout. The benefits of hexagonal element layout over classic, rectangular layout have been analysed theoretically and showed to be beneficial. Importantly, using hexagonal elements enables increasing the aperture of

individual array elements by approximately 10% without the corresponding drop in acceptance angle. This in turn allows a commensurable rise in the sensitivity of the sensor or alternatively, reduction in array element count for a given overall array aperture by over 20% without the corresponding drop in the image quality measurements/levels.

In the second part of the work, the problem of high output impedance of the miniature ultrasonic sensor is addressed by means of an in-probe miniature signal conditioning circuit. This improved the response amplitude of the element by 36dB and shortened its impulse response by a factor of 1.6. The novelty and practical benefit in this case lies in the fact that no high power components are needed in the probe body.

In the third part of this work, an emerging General-Purpose Graphics Processing Unit (GP-GPU) computer architecture is considered for the opportunities it offers to rethink the implementation of algorithms typically used in ultrasonic signal processing. Single-way beamforming, and two-way TFM and PCF beamforming have been developed for execution on the new platform, and show increase in performance of over 930 times compared to CPU processor.

This software platform has been further enhanced by a new approach to solving the refracted ray Timeof Flight problem in a way that is particularly well suited for this architecture. This resulted in a further increase of performance, i.e. 56x over the best published result found in the literature. The unprecedented performance and low cost of this new approach enables industrial deployment of advanced beamforming methods, as well as building of practical CAD tools for engineering and education.

Table of contents

Table of contents.....	6
Table of figures.....	14
Table of symbols.....	22
Glossary of terms and acronyms	24
Chapter 1. Introduction.....	26
1.1 About Non-Destructive Evaluation	26
1.2 Ultrasound as an NDE technique	27
1.3 Motivation and scope of interest of the work.....	28
1.4 Guide to the Thesis	29
1.5 Introduction to Ultrasonic NDE inspection technology	30
1.6 Construction of an Ultrasonic Phased Array Probe - overview.....	31
1.7 Key performance metrics for a 2D phased array ultrasonic probe.....	33
1.7.1 Introduction	33
1.7.2 Responsivity	33
1.7.3 Sensitivity.....	34
1.7.4 Resolution	34
1.7.5 Contrast.....	35
1.7.6 Coverage	35
1.8 Key construction considerations for a 2D phased array ultrasonic probe.....	35
1.8.1 The Piezoelectric composite	36
1.8.2 Overall aperture shape and size.	37

1.8.3	Element layout	37
1.8.3.1	Dense arrays.....	37
1.8.3.2	Sparse arrays	38
1.8.4	Transducer element size and shape.....	38
1.8.5	Acoustic impedance matching layers.....	39
1.8.6	Acoustic backing layers	39
1.8.7	In-probe electrical interconnect	40
1.8.8	Out-of-probe electrical connector (cable)	40
1.9	Probe construction considerations summary.....	40
1.10	Typical signal acquisition and processing methods	41
1.11	List of contributions to knowledge	41
1.12	Publications associated with the Thesis.....	42
Chapter 2. Design of piezoelectric ceramic composite		45
2.1	Introduction to piezoelectric ceramic - polymer composite material	45
2.2	Properties of 1-3 connectivity composite.....	47
2.3	The research question.....	48
2.4	Literature review.....	49
2.5	About 1D Linear systems modelling	54
2.6	Some known design parameter interactions.....	55
2.7	Simplified 3D model of regular 1-3 composite with load medium.....	56
2.7.1	Exploration of the problem space.....	57
2.7.2	Running the parameter sweep	61
2.7.3	Discussion of results.....	61
2.8	Summary	67

Chapter 3. Properties of hexagonal elements	69
3.1 Introduction	69
3.2 Relation of image quality and measurable quantities	69
3.3 Element layouts for dense phased arrays.....	70
3.4 Literature review.....	72
3.5 Comparison of the packing of squares and hexagons	72
3.6 Acceptance angle of a hexagonal array element.....	77
3.6.1 Rationale	77
3.6.2 Equations relating acoustic field and radiating aperture.....	77
3.6.3 Simulation 1 – constant aperture; acceptance angle versus wavenumber... 80	
3.6.3.1 Description of the model	80
3.6.3.2 Simulation 1: Results.....	83
3.6.4 Simulation 2 – constant acceptance angle	87
3.6.4.1 Description of the Problem	87
3.6.4.2 Simulation 2: Results.....	87
3.7 Probe aperture overall shape and size considerations.....	89
3.7.1 Lambert azimuthal equal-area projection as a space sampling method.....	90
3.7.2 Simulation 1. Comparison of probes with equal element count and different probe apertures	95
3.7.3 Simulation 2. Comparison of probes with different element count and similar apertures97	
3.8 Summary	99
Chapter 4. Triangular-Cut Piezoelectric Composite	100
4.1 Introduction	100
4.2 Review of recent literature	101

4.3	Modelling the vibration modes of the material.....	103
4.3.1	Introduction	103
4.3.2	Model Visualization Run	104
4.3.3	Model simulation study 1: varying kerf pitch	107
4.3.4	Model simulation study 2: varying material thickness	109
4.3.5	Discussion of results.....	111
4.4	Modelling the mechanical crosstalk	111
4.5	Experimental validation	113
4.5.1	Manufacturing the test samples	113
4.5.2	Evaluating the test samples	114
4.5.2.1	Electrical impedance analysis	114
4.5.2.2	Scanning Laser Vibrometry	115
4.6	Conclusions and future work	116
Chapter 5. In-probe active impedance matching module.....		118
5.1	Introduction	118
5.2	The effect of a cable on a response of an element of 2D phased array	119
5.2.1	Twisted pair cable properties	120
5.2.2	Modelling the effect of the cable capacitance on the probe sensitivity	121
5.3	Optimal receiver circuitry resistance	126
5.3.1	Instrumentation signal input impedance.....	126
5.3.2	Resistive damping for shortening the impulse response.....	126
5.3.3	Transducer model in transmit mode	127
5.3.4	Model parameter sweep of transducer in receive mode	128
5.3.5	Results analysis	131
5.3.6	Summary	131

5.4	Active In-Probe Line Driver	132
5.4.1	Introduction	132
5.4.2	Circuit operation principle	134
5.4.3	Component selection considerations	135
5.4.4	Example implementation.....	137
5.4.5	Evaluating the prototype	138
5.5	Case study	140
5.6	Conclusions and future work	141
Chapter 6. GP-GPU accelerated ultrasonic CAD.....		142
6.1	Introduction	142
6.2	The cueBEAM acoustic field pressure calculation package	143
6.2.1	Introduction	143
6.2.2	<i>cueBeam</i>	144
6.2.3	Implementation overview.....	144
6.2.4	Performance of CUDA Implementation	145
6.2.5	Applications.....	146
6.2.5.1	ArrayEdit3: Sparse array editor.....	146
6.2.5.2	BeamDemo: Interactive beam forming and focussing demonstration ...	149
6.3	Computationally efficient solution of Snell's Law of refraction.....	150
6.3.1	Introduction	150
6.3.2	Direct solution for the ray entry point (ray refraction point), 2D case.....	152
6.3.3	Solution for the ray entry point (ray refraction point), 3D case	155
6.3.4	Implementation	156
6.3.4.1	New direct solver	156
6.3.4.2	Reference Newton iterative solver and simplex iterative solver.....	157

6.3.5	Benchmarking the code	159
6.3.6	Accuracy of the solution	160
6.3.7	Remaining notes	161
6.3.8	Case study	161
Chapter 7. GP-GPU accelerated ultrasonic signal processing.....		164
7.1	Motivation for work	164
7.2	Implementation of the TFM algorithm, 1st generation.....	164
7.2.1	Introduction.	164
7.2.2	The distance function cost	165
7.2.3	Overall algorithm cost.....	166
7.2.4	Implementation benchmark, 1 st generation.....	167
7.3	TFM algorithm: 2 nd generation with support of refraction through arbitrary 3D surface	169
7.3.1	Implementation considerations.....	170
7.3.2	The proposed method	174
7.3.2.1	Phase 1: Prototype Time of Flight (ToF) points calculation	174
7.3.2.2	Phase 2: Calculation of interpolant coefficients	175
7.3.2.3	Phase 3: Signal accumulation.....	176
7.3.3	Notes on the proposed algorithm.....	180
7.3.4	The non-branching polynomial interpolant coefficient solver	181
7.3.4.1	Introduction	181
7.3.4.2	Implementation	181
7.3.4.3	Selection of interpolant order and error analysis.....	182
7.3.4.4	Discussion.....	184
7.3.5	Implementation benchmark	185
7.3.6	Experimental validation	187

7.3.7	Remaining implementation notes	191
7.4	Case studies.....	191
7.4.1	Case study: cueMAP: probe CAD tool.....	191
7.4.2	Case study: 3D ultrasonic inspection of an industrial specimen	197
7.5	Summary and conclusion	201
Chapter 8. Concluding remarks and future work.....		202
8.1	Thesis summary	202
8.2	A personal perspective.....	204
8.3	Proposed future work in probe technology.....	204
8.3.1	Electrical interconnect by soldering.....	204
8.3.2	Electrical interconnect by thin printed circuit board.....	206
8.3.2.1	Introduction	206
8.3.2.2	Electromechanical crosstalk.....	209
8.3.2.3	Electric field transfer.....	213
8.3.2.4	Future work.....	213
8.3.3	Manufacture of graded acoustic impedance material using additive manufacturing principle.....	213
8.3.4	Active array element with differential output.....	214
8.3.5	Stricter definition of the 'near field' extent.....	216
8.3.6	Surface resolving beam forming for improved defect detection in difficult areas	218
8.4	Closing remarks.....	222

References 225

Appendix A. Snell's law of refraction and polynomial fit equations processed using Wolfram Mathematica	231
Appendix B. About Matlab's fminsearch.....	232

Table of figures

Fig. 1.1 Schematic of an inspection system with ultrasonic phased array.	31
Fig. 1.2. An overview of the probe construction. (a) Sensor layer; (b) electrical interconnect wires; (c) casing; (d) cable strain relief; (e) cable;	32
Fig. 1.3. Details of probe construction around the piezoelectric composite material. (a) - first acoustic matching layer; (b) - second acoustic matching layer(optional); (c) common ground electrode; (d) piezoelectric composite (piezocomposite) material layer; (e) array element electrode; (f) ground electrode wire; (g) signal wire; (h) acoustic backing block ...	32
Fig. 2.1. An example of 1-3 connectivity piezoelectric composite. Ceramic pillars embedded in a polymer resin support.....	46
Fig. 2.2. An example of 0-3 connectivity piezoelectric composite. Ceramic particles embedded in a polymer resin support	46
Fig. 2.3. A single piezoelectric pillar can expand more than a piezoelectric slab because it is less constrained in the lateral direction. (a) piezoelectric slab, unexcited state. (b) piezoelectric slab, excited state. (c) piezoelectric pillar, unexcited state. (d) piezoelectric pillar, excited state.....	47
Fig. 2.4. A top-down view of the model. The model boundaries are marked with black. The pitch is the array element pitch (not the pillar pitch).....	58
Fig. 2.5. A model view – isometric. Side boundaries have symmetric condition applied. The top boundary has absorbing condition applied. The bottom boundary has a free surface condition applied.	58
Fig. 2.6. Material properties: longitudinal wave damping in polymer vs. polymer stiffness .	59
Fig. 2.7. Material properties: characteristics of damping the polymer used in the simulation.	60
Fig. 2.8. GBP vs. crosstalk for all test cases.....	63
Fig. 2.9. Gain-Bandwidth product vs. composite thickness.	63
Fig. 2.10. GBP vs. polymer stiffness	64
Fig. 2.11. GBP vs. ceramic volume fraction.....	64
Fig. 2.12 Best crosstalk vs Polymer shear wave-damping ratio	65
Fig. 2.13. Shear wave damping characteristic of the polymer range used.....	66
Fig. 2.14. Polymer stiffness vs. CVF for selected near-optimal cases.	68
Fig. 2.15. Composite thickness vs. CVF for selected near-optimal cases.....	68

Fig. 3.1. Left: Layout of 1D Phased Array elements. Middle: 2D phased array layout created by subdividing 1D phased array elements. Right: Proposed 2D phased array layout with hexagonal elements.....	71
Fig. 3.2 Optimal sphere packing (2D case). Left: square packing. Right: Hexagonal (optimal) packing. If we select such a, b, that $a=b=\lambda/2$, then $R_2>R_1$. The same applies for packing corresponding rectangles and hexagons, which fully cover the available area.	73
Fig. 3.3. A space-filling placement of squares and hexagons.	73
Fig. 3.4. Rectangle and hexagon dimensioning symbols	74
Fig. 3.5. Comparison of area of hexagon and rectangle of the same major radius R. Note the large patches of hexagon that extend beyond the perimeter of the rectangle.	76
Fig. 3.6. Illustration of Rayleigh integral	78
Fig. 3.7. Integration limits for hexagonal element.....	79
Fig. 3.8. A comparison of rectangle and hexagon shapes of an equal area. Units are arbitrary distance units.	81
Fig. 3.9. The location of the radiating element and the acoustic field probe points. Blue crosses depict the points at which the field pressure is calculated.....	82
Fig. 3.10. Example field amplitude at $\beta=0$, area of 1 and at significantly high wavenumber of 8.8 to illustrate the difference in acceptance angle and the appearance of the null in sensitivity.	83
Fig. 3.11. Acceptance angle of hexagonal element	84
Fig. 3.12. A comparison of the acceptance angle for hexagonal and square element at the same wavenumber, and at the same plot scale.	84
Fig. 3.13. Comparison of acceptance angle for hexagonal and square element versus wavenumber. The difference is bigger for wavelengths close to the size of the element (which are of most interest)	86
Fig. 3.14. Area of the element adjusted to obtain desired acceptance angle at desired amplitude drop	88
Fig. 3.15. Dense array layouts with circular overall aperture. Left: Example 98-element layout of an array with hexagonal elements. Right: 128-element layout of an array with square elements. The layouts have been selected to be nearly equal in aperture, equalizing focal spot size of the generated acoustic beam, while keeping the $\lambda/2$ sampling.	90
Fig. 3.16 Scene setting for cueBEAM. Green: element locations (actually, points); red cross: "probe points" where field is calculated.....	92

Fig. 3.17 Lambert azimuthal equiareal map field probing point distribution. Left: perspective view. Right: top-down view.	93
Fig. 3.18. A hexagonal-element probe overlaid with a square element probe of the same pitch and element count.....	96
Fig. 3.19 Comparison of cross-sections of beam shapes and side lobes of 128-element equal-pitch probes. Left: Hexagonal element probe. Right: square element probe. Small red circle depicts location of peak in side lobe.	96
Fig. 3.20. Cross-section of beam shape generated by respective probe layouts. Left: Hexagonal-element probe. Right: square-element probe. The small red circle shows location of the peak side lobe.....	99
Fig. 4.1 Left: Layout of 1D Phased Array elements. Middle: 2D phased array layout created by subdividing 1D phased array elements. Right: Proposed 2D phased array layout with hexagonal elements.....	101
Fig. 4.2 Triangular pillar, 1-3 piezoelectric composite schematic, top-down view. Array elements are defined by placing an electrode pattern over selected group of pillars. Left: 6 pillars per electrode (shaded); middle: 13 pillars per electrode; right: 24 pillars per electrode. Other configurations are possible.	101
Fig. 4.3 Triangular-cut piezoelectric ceramic composite a) composite overview; b) single-element model. Dark grey – ceramic pillars; light grey - polymer.....	105
Fig. 4.4 A visualization of the displacement of the surface of the model at t=0; light blue - ceramic pillars; dark blue – kerf (polymer filler); lemon yellow – excited array element; red / orange –neighbouring, grounded elements, used to calculate electromechanical crosstalk; dark red – remaining neighbouring elements.	105
Fig. 4.5 An example of the desired displacement mode. Visualization of the model surface at t=297ns after the excitation pulse. The centre element is being excited and protrudes from the surface of the material, radiating a coherent wave of motion into the front half-space; while neighbouring elements remain at rest.	106
Fig. 4.6 An example of undesirable behaviour: the apices of the triangles vibrate at a different frequency than the ceramic pillars; such motion modes do not convert into radiated energy and are absorbed into internal material friction instead.....	106
Fig. 4.7 An example of a simulated impedance plot of triangular-cut single array element for platelet thickness of 450 μ m and kerf pitch of 450 μ m. Green dot shows frequency of the	

electrical resonance, and the red dot shows the frequency of the mechanical resonance of the device.....	107
Fig. 4.8 Surface plot of admittance as a function of frequency and kerf pitch.....	108
Fig. 4.9 Trace of peaks in the admittance of composite platelet versus kerf pitch. Red dotted lines mark the kerf pitch range for which only one resonant mode is detectable.....	108
Fig. 4.10 Surface plot of admittance as a function of frequency and kerf pitch.	110
Fig. 4.11 Frequency of peaks in admittance versus piezoelectric composite thickness. Red lines mark the thickness range for which only one resonant mode is detectable	110
Fig. 4.12. PZFlex model of surface displacement amplitude for hexagonal element (left) and square element (right). Colour scale is relative.	112
Fig. 4.13. Left: triangular-cut composite; right: square-cut composite. The side shown contains the stock fired-on electrode through which the respective cuts have been made.	113
Fig. 4.14. Ground electrode side of the square cut piezoelectric composite. The electrode is made by depositing chromium vapour and then gold vapour in a vacuum chamber.....	114
Fig. 4.15. Impedance plots of the manufactured devices: square and hexagonal single element.....	115
Fig. 4.16. Measurement of the displacement amplitude of the hexagonal and square composite at their resonant frequencies. The Red line denotes the element boundaries used for displacement and crosstalk calculation.....	116
Fig. 5.1 Typical electrical schematic of connection between probe elements and the Phased Array Controller	119
Fig. 5.2. Schematic of the PZFlex model used in the example.....	122
Fig. 5.3. Electrical impedance of the modelled system assuming no influence from the capacitance of the cable.	123
Fig. 5.4. Electrical transmit impulse response of the modelled system assuming no influence from the capacitance of the cable.	123
Fig. 5.5 Simulation parameter sweep: insertion loss versus capacitance of the cable.	125
Fig. 5.6 Impedance plot of the ceramic slab in the study and the characteristic points marked out.....	128
Fig. 5.7 Phase plot of the ceramic slab in the study and characteristic points marked	128
Fig. 5.8 Received voltage decay time vs load resistance	130

Fig. 5.9. Centre frequency of the received signal vs load resistor. The frequency spectra used for the analysis has been normalised versus the spectra of the excitation signal.....	130
Fig. 5.10: Line driver schematic overview - single channel.....	136
Fig. 5.11 Line driver schematic with additional long line driver	136
Fig. 5.12: visualization of PCB Layout of the proposed line driver, 4 channels per board. .	138
Fig. 5.13: Photo of assembled PCB with 4 channels	138
Fig. 5.14. The results of an analysis of pulse-echo time signals acquired with and without the line driver. The blue line represents -6dB boundary.....	139
Fig. 5.15. Design view of 32x PCBs (green) inside the body of $\varnothing 120$ mm probe case.....	140
Fig. 6.1 Densely packed 122 element phased array layout defined over triangular grid in the ArrayEdit3 software package.....	148
Fig. 6.2 Screenshot of the ArrayEdit application with a probe design loaded. Array elements have been manually placed. Note the changed performance figures compared to the case in Fig. 6.1: narrower beam width and higher side lobe amplitude for the sparse array case.	148
Fig. 6.3. Example screenshot of <i>cueBeam</i> demonstrator. Top: two 2D cross-sections through 3D space; Lambert equiareal mapped hemisphere (top) and XZ plane (middle). Bottom – 1D cross-section of the beam along the white curve presented in the top views. From this last view, two beam characteristics are instantly calculated: main lobe width and energy leakage factor.....	150
Fig. 6.4. Depiction of symbols used in the application of Snell's Law.....	151
Fig. 6.5. Ray travel time from a point at (0.01,-0.02) in medium 1 (the lower part of the figure) to all other points in space. There are numerical singularities in the direct solution when calculated using single-precision numbers.	154
Fig. 6.6. Example solution as in Fig. 3, but with a corrected algorithm - there are solutions for all possible locations and no singularities are present. This correction is obtained by switching to an iterative algorithm, only for the point where singularity is detected. Note that a critical angle of refraction effect is visible, as there is an area on the interface where travel time is continuous and at one point it becomes discontinuous.....	154
Fig. 6.7. Points P_1 , P_2 , P_i of the 3D refraction case and their respective helper points, P_{1hr} , P_{2hr} and P_{ih} . Note the rotated coordinate system.	155
Fig. 6.8. Screenshot of <i>cueBeam</i> R4 in operation. Top: simulated probe element layout. Middle: sphere section. bottom: planar section.	162
Fig. 6.9. Example outputs of <i>cueBeam</i> R6 for a linear phased array probe.	163

Fig. 7.1. Comparison of performance of TFM algorithm on CPU and GPU platforms, January 2011. The synthetic FMC data (8-element probe) has been produced in COMSOL.....	169
Fig. 7.2. Elements of ultrasonic imaging scenario.....	170
Fig. 7.3. Functional diagram of memory architecture of CUDA capability 3.5 processor (Kepler architecture). The Stream Processors (SP) can benefit from broadcast mechanism (loading of value from particular memory location to all SP's simultaneously) if the programme is written to take advantage of it. This applies to Constant Cache and Shared Memory. The L2 cache can broadcast its data across multiprocessors.	173
Fig. 7.4. Location of probe, refracting material interface, image buffer, and z-line of pixels inside the buffer. The linear memory locations progress down-first, then right. The third dimension is X.	174
Fig. 7.5. Schematic of per-pixel computation kernel and memory access pattern. There are 10 algorithm steps per pixel. SP - Streaming Processor (also known as CUDA core). TEX – texturing unit. Shared – shared memory module. TEX cache – texturing unit local cache, if present. L2 – L2 on-chip cache buffering the device global memory access. Constant cache – the on-chip, per-Multiprocessor constant cache (also known as the code cache). Global memory – device global memory.	177
Fig. 7.6 Comparison of interpolation error when using different interpolation coefficient solvers	184
Fig. 7.7. The experimental specimen and probe holder. Irregular geometry of the top plate has been used to avoid spurious reflections.	188
Fig. 7.8 Schematic drawing of the specimen and probe aperture.....	189
Fig. 7.9 The image of the flat bottom hole inside the specimen. Image assuming that the rays do not undergo refraction.....	190
Fig. 7.10 The image of the reflector as in Fig. 7.9, but with correct refracting surface taken into account. The amplitude of the reflector is 4.57dB higher and the reflector is correctly positioned. Importantly, the shape and orientation of the reflector is imaged correctly (flat bottom hole).	190
Fig. 7.11 The image from Fig. 7.10, but with wider dynamic range. The correctly flat back-wall is visible at -42dB from the top surface. The three black straight lines exemplify calculated ray paths between the three probe elements and a pixel in the image.	190
Fig. 7.12 The cueMAP process - overview	193
Fig. 7.13 The cueMAP process – per pixel	193

Fig. 7.14 Calculated image quality maps – example	195
Fig. 7.15 Map of acceptable quality of image assuming example inspection requirements.	196
Fig. 7.16 pump casing showing location of phased array transducer and scan geometry ..	197
Fig. 7.17 Definition of reference planes in sample	198
Fig. 7.18 Internal crack penetration and depth visualisation	198
Fig. 7.19 A-A reference plane image of pump casing	199
Fig. 7.20 C-C reference plane image of pump casing.....	200
Fig. 7.21 D-D reference plane image of pump casing	200
Fig.8.1. General view of the PCB layout. There are 6 30-way 1mm pitch edge male edge connectors. In the centre there is a 2D electrode array area with 128 pads. Dimensions are in mm.	207
Fig. 8.2. Detail of the track routing over the piezoelectric composite area and the connection of tracks to the element electrodes.....	208
Fig. 8.3. View of the alignment and deposition rig. The thin flexi PCB is suspended under a bridge; underneath the bridge, the composite is manipulated with 4 degrees of freedom. The flexi PCB, being semi-transparent, allows observing and guidance of the composite below through a 3D microscope.....	209
Fig. 8.4. Final inspection of the align accuracy of the electrode (dark patch at the front) over the ceramic pillars (light triangles at the back). Note that the element in this view is NOT the one in the laser vibrometry in the Fig. 8.7	209
Fig. 8.5 Electrical field distribution (colour) and deformation of the piezoceramic pillar (left)- polymer kerf (right) under 1V electric field if the field is applied over the piezoceramic pillar. The electrical field gradient is distributed across the thickness of the piezoelectric ceramic phase. Deformation scale: $2 \cdot 10^5:1$	211
Fig. 8.6 Electrical field distribution (colour) and deformation of the piezoceramic pillar(left)- polymer kerf(right) under 1V electric field, when the electrical field is applied over the polymer kerf only(right). Portion of the electrical field does enter the ceramic and pulls the polymer in; however, the effect is very weak. Deformation scale: $1 \cdot 10^8:1$. Note that deformation scale is 1000x of the scale in Fig. 8.5.	211
Fig. 8.7 Top-down view of displacement of the hexagonal element under thin-PCB electrode, as measured by laser vibrometry. Left: amplitude view; right: phase view. Dark hexagonal patches correspond to electrode locations. The colour indicated	

amplitude/phase and has been laid over a photo acquired through the laser vibrometer viewfinder. The PCB track delivering the charge to the hexagonal element does not excite pillars that do not belong to the element of interest. 212

Fig. 8.8 Pulse-echo signal obtained with the prototype device. There signal is well-behaved and there is no signs of a crosstalk in the spectrum (no ‘secondary pulse’). Please note that this was acquired with a device that has no acoustic backing, matching or electrical damping; therefore pulse length is not representative of a final device and serves to show that the device is operational only. 212

Fig. 8.9 Proposed circuit of the Piezo transistor for application in phased arrays. TN – transistor with N-doped channel; TP – transistor with P-doped channel. 215

Fig. 8.10. Proposed circuit of the Piezo transistor for application in phased arrays – single type of transistor..... 215

Fig. 8.11. Example of narrowing the side lobes when impulse response of the probe is shortened. Cross-section through the space in plane where the radiating elements are located. The radiating elements are to the left (red points). Top: continuous radiation; middle: intermediate impulse response length; bottom: short impulse response length. Left figures: pressure field view; right figures: interference count view (0, 1, or 2 times interfered)..... 218

Fig. 8.12 Example TFM image of a planar reflector and a point reflector – synthetic data. The point reflector sought is buried within imaging artefacts. 220

Fig. 8.13 Same situation as Fig. 8.12 but with the planar reflector removed. The point reflector is clearly identifiable. 220

Fig. 8.14 Geometric construction to show the error of location when planar reflector is considered. Note that the ellipse intersection (TFM reflector solution) does not coincide with the real reflector..... 220

Fig. 8.15 SR-TFM image..... 221

Fig. 8.16 Residue TFM image on the FMC data with the surface reflector suppressed, based on the SR-TFM image data. The red line indicates the recognized, suppressed surface. 222

Table of symbols

$A_{triangle}$ – Area of a triangle

a_H - radius of a hexagon

a_R – edge length of a square

$A_{Hexagon}$, A_H – Area of a hexagon

A^{square} - area of a square

d – diameter of a hexagon

δ – dirac delta

h - height of a triangle

k – wave number

k_t - thickness mode electromechanical efficiency coefficient

λ - wavelength of interest

ω – wave frequency

$p(\vec{r})$ - pressure at a point

R – radius of a shape of interest

\vec{r}_0 – a spatial vector – between two points of interest in a space

Rx – receive, wave receiving element

t - time

Tx – transmit, wave transmitting element

$U(\vec{r})$ – displacement vector at point

v – velocity of wave of interest

X, Y, Z – spatial axis names

Glossary of terms and acronyms

- *Acceptance Angle* – angle between array element axis and line or surface at which the sensitivity or amplitude of the radiation drops by a specific ratio to the peak sensitivity or radiation amplitude. The ratio is typically -3dB, and the angle is expressed in linear degrees.
- *BINDT* - British Institute for Non-Destructive Testing
- *CMAP* – Conformally Mapped array layout
- *CVF* – Ceramic Volume Fraction, a ratio of the volume of ceramic material to the volume of entire composite material.
- *Decay time / ring down time* – the time that it takes for an oscillation of the signal to die out. The limits are marked by the envelope of the signal crossing a specified threshold, e.g. -20dB re the peak of the signal amplitude.
- *Electrical resonance frequency* – a frequency at which a piezoelectric ceramic material will resonate under condition of electrical **short-circuit** between the electrodes of relevant mode of vibration
- *element* - refers to a single phased-array element, connected electrically to a single channel of the PAC
- *GBP* – Gain–bandwidth product
- *GP-GPU* – General Purpose Graphics Processing Unit
- *L1 Cache, L2 Cache* – Level 1 (closest to the core) and Level 2 (on-die, larger) Cache memories inside the computer’s processor. These are very fast, low latency memory regions that are used to accelerate repeated retrieval of recently accessed data from the main memory; it is also used to buffer and coalesce writes to main memory. In modern processors, the size of the cache is typically 16kB - 4MB. Due to its relatively high silicon cost, it is not included at all in some low-cost processors.

- *Mechanical resonance frequency* – a frequency at which a piezoelectric ceramic material will resonate under condition of electrical **open-circuit** between the electrodes of relevant mode of vibration
- *MPAR* – Characteristic value for certain pillar geometry in 1-3 piezoelectric composite. It is defined as such ratio of the size of the linear dimension of pillar base to the pillar height, at which the first lateral mode frequency occurs at first harmonic frequency of the base thickness mode.
- *NDE* – Non-destructive evaluation
- *NDT* – Non-destructive testing
- *PAC* – Phased Array Controller
- *PCF* – Phase Correlation Factor, a beamforming technique
- *Pixel* – Picture Element
- *ToF* – Time of Flight
- *TFM* – Total Focussing Method, a beamforming technique
- *Q, Q-factor* – mechanical vibration system quality factor; inversely proportional to amplitude loss per oscillation; the higher the Q factor, the longer the ring-up and *ring-down time* and higher selectivity of the system – more narrowband system.

Chapter 1. Introduction

1.1 About Non-Destructive Evaluation

Before introducing the topic of this Thesis, a short introduction to Non-Destructive Testing, and the related Non-Destructive Evaluation is appropriate.

Non-Destructive Evaluation / Testing (NDE / NDT) is a term used to encompass a wide range of techniques used to control and verify the quality and safety of components.

The test essentially consists of a measurement of some physical quantity and an analysis of the measurement results. Most often, some sort of limited energy packet is introduced into the specimen to excite it and promote the quantities to be measured. The limited quantity of energy used ensures that the test is non-destructive.

A longer definition is provided by the British Institute for Non-Destructive Testing (BINDT) website [1] as:

Non-destructive testing (NDT) is the branch of engineering concerned with all methods of detecting and evaluating flaws in materials. Flaws can affect the serviceability of the material or structure, so NDT is important in guaranteeing safe operation as well as in quality control and assessing plant life. The flaws may be cracks or inclusions in welds and castings, or variations in structural properties which can lead to loss of strength or failure in service.

Non-destructive testing is used for in-service inspection and for condition monitoring of operating plant. It is also used for measurement of components and spacings and for the measurement of physical properties such as hardness and internal stress.

The essential feature of NDT is that the test process itself produces no deleterious effects on the material or structure under test.

The subject of NDT has no clearly defined boundaries; it ranges from simple techniques such as visual examination of surfaces, through the well-established methods of radiography, ultrasonic testing, magnetic particle crack detection, to new and very specialised methods such as the measurement of Barkhausen noise and positron annihilation.

1.2 Ultrasound as an NDE technique

As described in [2], Ultrasonic methods of NDT use beams of mechanical waves (vibrations) of a short wavelength and high frequency, transmitted from a probe and detected by the same or other probes. Such mechanical waves can travel large distances in a material, in the form of a divergent wave with progressive attenuation. The vibrations are typically aperiodic (impulse with reverberation of a Q-factor close to 1/2) with a centre frequency of 0.1 to 20MHz and a wavelength in the range of 0.1mm to 10mm. The velocity of the sound depends on the material's density and stiffness properties, and is typically in range of 1000-6000m/s, higher for longitudinal (compression) waves and slower for shear waves. The technique detects internal, hidden discontinuities that may be deep below the surface. Most non-porous, resilient materials used for structural purposes – steel, aluminium, titanium, magnesium and many ceramics - can be inspected. Large cross-sections can be tested successfully for minute discontinuities. Compared to X-ray, ultrasound can detect defects that do not change the material's local density, e.g. closed cracks.

Ultrasonic testing techniques are widely accepted for quality control and material testing. The technology is proven, well understood and applied to industries including production of materials, fabrication of airframes, ship building, jet and rocket engines, electric power generation, and others.

The basic technology starts with a single-element probe, which can be used for thickness gauging, flaw detection and basic sizing, and when combined with a mechanical scanner, for imaging too. However, inspection using a single-element probe is time consuming. Its field of view is very limited, and consequently has to be mechanically scanned throughout the area of inspection.

The evolution of single-element probes are 1D, or linear arrays of transducers; these can be divided into multi-transducer arrays and phased arrays; the latter being of interest of this work. In all of the remaining work, 'array' refers to a 'phased' array and not a multi-transducer array.

Phased array ultrasound is based on multi-element ultrasonic probes where elements are in close proximity. Importantly, the relative phases of transmission and reception are carefully controlled. This precise timing enables exploitation of the phenomenon of constructive and destructive interference of the waves in the specimen.

In transmission mode, the interference is used to direct the bulk of the energy in a preferential direction into the specimen. In reception mode, interference is used to discriminate the direction of arrival of any echoes from the specimen's features.

Electronic control of the phases enables spatial scanning of the specimen at a typical rate of up to a few kilohertz, effectively producing an image of the specimen's internal features.

It is important to note that the image produced is an image of acoustic contrast, and its features are not similar to features of an optical image. In particular, one should note that beams of ultrasound are more similar to 'spotlight' or 'flashlight' beams of light and unlike ambient illumination known from the optical world. Somewhat counterintuitively, the indication of a reflector appears at its (reflecting) surface, and not at its geometrical centre of mass. Secondly, the reflection only appears at the surfaces that can reflect the beam energy back to the sensor, and flat reflectors can often be completely invisible on the image if they reflect the beam energy away from the sensor.

1.3 Motivation and scope of interest of the work

Technology for production of 1D arrays can be described as matured. There are numerous commercial systems exploiting such probes, and they are still gaining ground in the non-destructive testing market. Mechanical scanning is not eliminated completely in such a system though.

Compared with 1D phased arrays, 2D phased arrays can generally deliver more information about the reflectors/defects, and the potential benefits include:

- A larger volume covered in a single image
- An enhanced defect characterisation, and classification
- 3D volumetric imaging may be easier to interpret by the operator than 2D slices
- The ability to access inspection angles otherwise impossible to obtain with typical 1D array, especially in components of complicated geometry
- A reduced need for mechanical scanning,
- A reduced overall cost of inspection

This approach however, comes with problems. The main issues with current 2D probe technology include:

- Low probe sensitivity. The active element size is limited by its directivity and layout requirements. A single array element cannot be much larger than half of the wavelength of interest. Total aperture of the probe is also practically limited. This also implies problems with the noise level in the reconstructed image – as high amplifications have to be used in receiver stage.
- A large number of individual array elements. Each element has to be connected to a transmit/receive module. Connecting them by wires is difficult, because elements are small, and there are many wires required. In addition, the cost associated with a multi-channel Phased Array Controller (PAC) is considerable.
- Crosstalk between array elements. Signals can go through electrical and mechanical paths, deteriorating image quality, and causing problems with image interpretation. There exist some potential de-blurring techniques that could be used, but in general, this remains a consideration when designing the probe.

1.4 Guide to the Thesis

In this Thesis, a range of topics relating to the problem of 3D ultrasonic imaging are addressed, from core probe technology to final image generation.

The main objective of the work is to address some of the current problems with the construction of 2D ultrasonic array probes for non-destructive evaluation (NDE) applications.

Chapter 1 provides an introduction to the scope of the thesis and basic information about the ultrasonic probe technology.

Chapter 2 reviews the technology of the fundamental element of the modern phased array probe – the 1-3 piezoelectric composite. The classic square cut 1-3 composite material design is numerically optimised, and a pareto line between sensitivity and crosstalk is found.

In Chapter 3, properties of the aperture of the 2D ultrasonic array are explored, and the properties of its constituent elements aperture are investigated. Hexagonal aperture elements are compared to square aperture elements and found to be beneficial.

In Chapter 4, a design method is given that yields a piezoelectric composite material that fulfils the needs of the hexagonal aperture elements proposed in Chapter 3.

Having the array layout and the constituting materials, the next step is to convey the signals from the transducer to the digitizer of the probe controller – the phased array controller. The problem of preserving the signal fidelity as it travels through the cable is addressed in Chapter 5.

The recent advances in the computer architecture – GP-GPU - has opened a window of opportunity to re-implement some popular algorithms in a completely new way.

Monochromatic beam simulation, refraction, and advanced beamforming problems are tackled in Chapter 6. These are then extended by treating the problem of calculation of ray Time of Flight through refracting interfaces in a novel way in Chapter 7.

There is more interesting problems that have been touched upon during the work on this thesis; some of them are listed with a brief introduction in Chapter 8.

1.5 Introduction to Ultrasonic NDE inspection technology

A schematic of an ultrasonic inspection system relevant to this work is presented in Fig. 1.1.

It consists of the following elements arranged in a chain topology:

(A) Mechanical scanner: moves the probe relative to the specimen. Alternatively, the specimen can be moved to present relevant sections to the probe.

(B) The specimen under inspection.

(C) The ultrasonic phased array probe.

(D) The PAC (Phased Array Controller). Internally, it consists of several sections:

- Pulsar (probe excitation) section
- Receiver section (low noise amplifier and Analogue-to-digital converter)

- Optional signal processing section (filtering, beamforming, and other algorithms)
- Uplink communications section.

(E) Desktop or laptop computer with Graphical User Interface software and data storage.

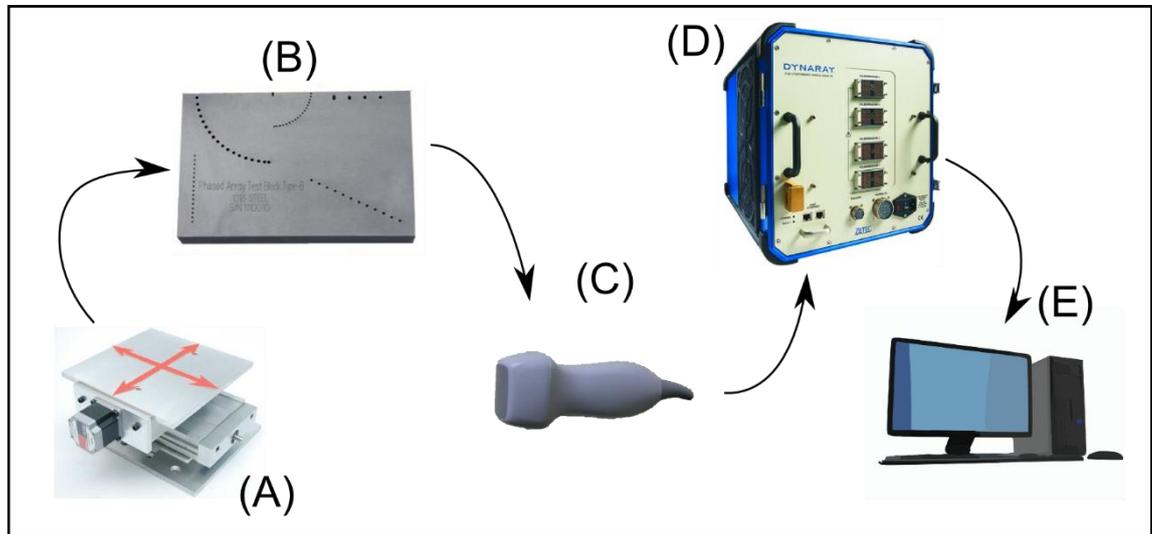


Fig. 1.1 Schematic of an inspection system with ultrasonic phased array.

Construction of the ultrasonic probe (C) is the main topic of this work, with additional attention given to the signal processing algorithms (E).

1.6 Construction of an Ultrasonic Phased Array Probe - overview

Probe assembly consists of these basic elements:

- PAC connector
- Cable
- Sensor housing
- Acoustic support materials - acoustic damping, acoustic impedance matching, wedge(s), housing
- Acoustic transmitter and sensor element - piezoelectric material

Typical elements of an ultrasonic 2D phased array probe are illustrated in Fig. 1.2. The details of the construction around the active piezoelectric material are illustrated in Fig. 1.3.

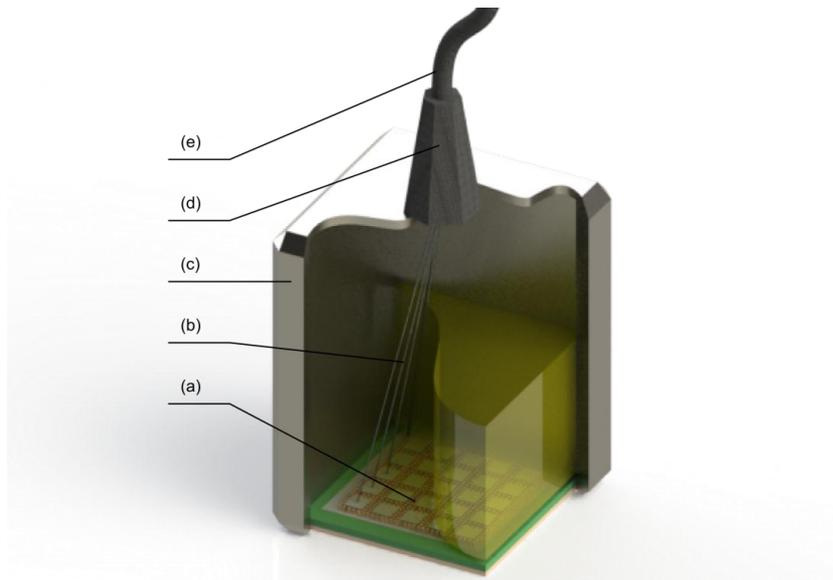


Fig. 1.2. An overview of the probe construction. (a) Sensor layer; (b) electrical interconnect wires; (c) casing; (d) cable strain relief; (e) cable;

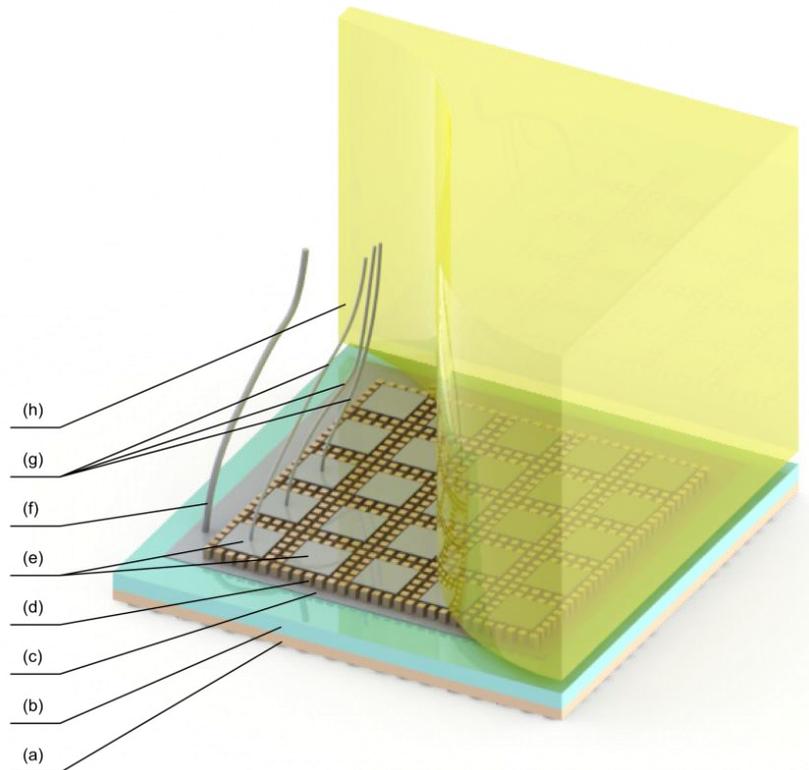


Fig. 1.3. Details of probe construction around the piezoelectric composite material. (a) - first acoustic matching layer; (b) - second acoustic matching layer (optional); (c) common ground electrode; (d) piezoelectric composite (piezocomposite) material layer; (e) array element electrode; (f) ground electrode wire; (g) signal wire; (h) acoustic backing block

Practical probe engineering is always searching for a balance between performance, technology limitations, and the cost of the solution.

The following Section is devoted to the means of defining the probe performance.

1.7 Key performance metrics for a 2D phased array ultrasonic probe

1.7.1 Introduction

An ultrasonic imaging system can be characterized in terms of the properties of the final generated image. The main characteristics of the probe and the resulting image are summarized as follows.

1.7.2 Responsivity

For this work, responsivity is analogous to its use in electronics; that is, a measure of input-output gain of the probe system. The piezoelectric layer in the probe is typically excited with a single pulse or a pulse train of limited energy; the energy is transduced into a mechanical wave, which then travels into the acoustic medium. Depending on the imaging case, the returned echoes may carry only a small fraction of the original energy. In an example laboratory scenario, the voltage produced by the echo returned from a small reflector might be from -20dB to -80dB lower than excitation pulse voltage, in fact, it's often close to the system noise floor. Therefore high probe responsivity is desirable as it contributes to the detectability of small reflectors.

In case of a perfect reflector (one that would reflect all incoming energy back towards the probe), responsivity of the probe is primarily a function of the active piezoelectric material used, acoustic path matching, and geometry of the probe elements.

Responsivity of the probe to a given reflector type in a given inspection scenario is mainly affected by probe frequency, probe aperture, and the imaging algorithm (signal processing) used.

Reflectors smaller than approximately one wavelength in the specimen will appear much weaker than larger ones; therefore higher frequency waves will be more responsive to these. However, at the same time, it is easy to understand that the bigger the sensor is, the more energy it can send and receive; a higher operation frequency will typically imply smaller transducers.

In general, the bigger the aperture, the better the overall responsivity due to more of the echo energy being captured in the probe. However, one of the practical limits is that the sound field has to be coupled into the specimen and complex geometry of the specimen might pose a problem for large planar probes.

1.7.3 Sensitivity

Sensitivity is the minimum magnitude of input signal required to produce a specified output signal having a desired signal-to-noise ratio. Since the signal-to-noise ratio can be defined as many different things for a given application, and also depends heavily on the used PAC quality, firing mode, and signal/ image processing algorithm, it is important not to confuse sensitivity with responsivity of the probe.

For example, sensitivity of the probe to a given reflector type is a complex function of an element layout and an element count, but also reflector and probe relative location in space. Either dense or sparse element layouts (see Section 1.8.3) can be favoured depending on the imaging scenario.

In general, using high quality, defect free materials and careful mechanical and electronic design of the probe can improve sensitivity.

1.7.4 Resolution

Resolution of the image produced using a given probe is a measure of the low limit of distance between two small reflectors before they appear indistinguishable from a single reflector.

In general, the resolution of an ultrasonic probe is limited by the sensitivity to short wavelengths (high frequency signals) probe aperture, and the focussing power of a given signal-processing algorithm.

The higher the operational frequency, the better the absolute resolution of the image can be, however, as will be shown in this Thesis, the actual governing parameter for that is impulse response duration (after a compression algorithm has been applied, if any).

Advanced signal processing algorithms can increase resolution beyond the diffraction limit, at a cost of high computational complexity. Such algorithms, called 'super-resolution' [3] are beyond the scope of this work.

1.7.5 Contrast

Image contrast is a measure of how different an object in an image is from its immediate surroundings. Often times, this is the major measure of applicability of a given inspection system to a given inspection scenario.

In general, the bigger the aperture, the better the contrast produced, except when a sparse array concept is applied.

1.7.6 Coverage

Array probes' primary advantage over single-element probes is that the volume of specimen available (without mechanical scanning) to the inspection can be significantly bigger than for single-element probes. The exact volume where the image quality is acceptable for a given inspection scenario depends on a number of factors.

The primary factor affecting the coverage is wavelength fractional element size. If the probe element size is less than approximately $3/4$ wavelength of interest, the probe will have a wide field of view. If the element size is less than $1/2$ wavelength, it will have a nearly 180° field of view, and it will also be sensitive to shear waves and surface skimming waves.

For most materials, it is typical that the higher frequency waves are attenuated more; therefore typically, the lower the frequency, the further into the specimen the wave can be sent, reflected, and received. This however does not imply sensitivity to any particular type or size of the reflector.

1.8 Key construction considerations for a 2D phased array ultrasonic probe

In the following Section, some of the key design choices needed to construct a 2D ultrasonic phased array probe are given with discussion. Most of these points apply equally to 1D phased array probes. However the emphasis is placed on what is important for 2D array considerations. Where appropriate, it is assumed that classic beam forming algorithm, (or its close extension, Total Focussing Algorithm[4]) is used to process the signals acquired with the probe into an image, and the comments are given accordingly. It is appreciated that there is a number of advanced imaging algorithms available but detailed analysis of these is outside of the scope of this work. Only mechanical and electrical

engineering parameters are considered in this work; improving on material properties is also outside of the scope of this work.

The guidelines given are primarily a result of general experience in the field and as such, are given here without proof.

1.8.1 The Piezoelectric composite

The active sensing element of the ultrasonic array probe is a piezoelectric ceramic – polymer composite material. Its location in a typical probe is depicted in Fig. 1.3. An important distinction from other electro-acoustic devices is that, at frequencies above $\sim 100\text{kHz}$, a piezoelectric ceramic is used in resonant mode. The mechanical resonance occurs within the body of the ceramic tiles; where other transducers may vibrate as a stiff, solid body following the crest of the mechanical wave, the piezoelectric transducers vibrate by extending and contracting the overall shape of the sensing part itself. Typically, a quarter of the wave ‘fits’ inside the sensing element, momentarily changing its shape.

Discussion of details of piezoelectricity is beyond the scope of this work and reader interested in the physical phenomena underlying the practical effects are encouraged to look them up in a library of their choice; e.g. [5]

However, it is of interest to this work how the shape of the sensing element affects practical outcomes of the piezoelectric phenomenon. In particular, the sensing element(s) have to be designed in such a way, that useful vibration modes are promoted, and parasitic (not beneficial) vibrations are not affecting the transmitted or received signals.

Any solid body, including piezoelectric ceramics, can vibrate in one of a number of modes. Out of these, the thickness-extension mode is the most useful for manufacturing longitudinal wave imaging probes. Shear waves can also be produced in the specimen using this vibration mode. Thickness of the piezocomposite is the most important parameter that determines the electrical and mechanical resonant frequency, at which the transducer will be most sensitive in transmit and receive modes respectively. The thicker the material, the lower the resonant frequency. For PZT type materials, the frequency constant is typically around $2000\text{Hz}\cdot\text{m}$. The electrical and mechanical loading of the piezocomposite layer influences the centre frequency.

Piezoelectric materials are useful from a low kilohertz range up to a high gigahertz range; however for NDE imaging applications, 500kHz to 20MHz are typically the most useful. The low limit is due to low spatial resolution of the image; the high limit is due to limited sensitivity and penetration range. There is a complex interaction from this parameter to many other design parameters of the probe.

1.8.2 Overall aperture shape and size.

In general, the bigger the probe aperture, typically expressed in wavelengths, the better the focussing power at that wavelength, and the further away from the probe surface the focal point can be placed. This is also closely related to the definition of the 'near field' region. However, the bigger the probe (assuming constant element density), the more expensive it is, and the economic effect of NDE inspection should not be ignored. Over time, various aperture shapes have been researched; in general, circular-symmetric apertures will give a circular-symmetric quality of image across volumes of interest; a non-symmetric aperture will give non-symmetric image quality, which can be acceptable or desired in some cases. In such cases, as a rule of thumb, the resolution is better in the plane where the cross section of the aperture is bigger.

Much research, with some success, has been committed into developing probes with apertures that conform to the variable curvature of the specimen, e.g. [6]. This primarily allows for avoiding intermediate acoustic coupling materials, and in effect better coupling of ultrasonic energy between the probe and the specimen. Such probes are beyond the scope of this work.

1.8.3 Element layout

Probe aperture is populated by probe elements – individually controllable transducers, of which some can be dedicated for exclusive generation or reception of ultrasonic signal, or more typically, they can be used for both transmission and reception. A review of typical element layouts can be found in [7, p. 102] or [8, p. 65]. More layouts have been researched over time.

1.8.3.1 *Dense arrays*

Array is said to be “dense” when the elements of the probe fill the entire aperture of the probe, and the maximum distance between neighbouring elements is equal or smaller than

half of the dominant wavelength of the probe in the acoustic medium it directly interfaces into. According to the Nyquist–Shannon sampling theorem, such an arrangement is enough to arbitrarily shape the transmitted acoustic field (at given frequency/wavelength), and measure full information about the impinging acoustic field.

However, full information is often not necessary to obtain sufficient imaging quality; therefore the concept of sparse arrays is of practical importance. In dense arrays, regular element placement is most often used, as it allows for optimal sensitivity and directivity of the elements.

1.8.3.2 *Sparse arrays*

Array is said to be “sparse” when in a given aperture there is less elements than required for the probe to be “dense”. Elements can still fill an entire aperture (being bigger than half-wavelength on average, sacrificing directivity) or they can also not fill the entire aperture of the array (sacrificing sensitivity to obtain better directivity).

The primary advantage of sparse arrays over dense arrays is that the aperture can be bigger for a given number of elements, yielding improvements in focussing ability and coverage. However, such array cannot measure full information about the impinging acoustic field, and which results primarily in deteriorated contrast of the image, and sometimes reduced sensitivity of the probe.

Since there can be many ways of laying out elements in a sparse array, much research has been committed to finding optimal locations. It is currently known that the best layouts are aperiodic[9][10][11].

1.8.4 Transducer element size and shape

The area of the probe aperture taken by a given element primarily affects its directivity (acceptance angle). The bigger the element, the more acoustic energy it can transmit and receive. However, for the array to be regarded as a phased array, the acceptance angle of the element must be at least comparable with the desired steering capability of the probe. The bigger the elements, the higher the on-axis sensitivity, but this may impact off-axis sensitivity.

1.8.5 Acoustic impedance matching layers

Acoustic impedance is an inherent property of a given homogenous material, and is related to its physical density and acoustic wave speed in that material. If an acoustic wave impinges on an interface between two materials of different acoustic impedances, some of the energy will be reflected, and only part of the energy will be transmitted into the second material.

Often, acoustic impedance of the specimen is significantly different than the one of the transducer. The transmission of energy can be improved by using mediating materials of intermediate impedance. Typically, two layers of acoustic matching are used; ideally, a continuous, profiled transition between impedances would be desirable. However a detailed discussion of this topic is beyond the scope of this work.

Sensitivity is typically improved when using appropriate impedance matching layers, because more energy can leave and reach the transducer in transmission and reception respectively.

Resolution of the image is most often improved when using appropriate impedance matching due to shortened impulse response of the transducer.

Contrast of the image is most often improved when using appropriate impedance matching layers due to shortened impulse response of the transducer.

Coverage of the image is only remotely affected by the acoustic impedance matching.

1.8.6 Acoustic backing layers

It is important to consider the material that exists on the backside of the piezoelectric transducer. There are generally 5 options for the material: no backing (air/vacuum), light backing (low impedance), medium backing (impedance of up to $\frac{3}{4}$ in value of the impedance of the transducer), heavy backing (impedance equal to the impedance of the transducer) or virtually fixed backing (impedance much higher than the impedance of the transducer). The impedance of this material has an important effect on the performance of the probe. Detailed considerations of the backing design have been thoroughly studied and are available in separate publications, e.g. [12]. The following paragraphs briefly describe the effects of medium backing versus no backing.

Sensitivity of the probe is reduced by medium and heavy backing because part of the energy that has been radiated into the backing is not radiated into the load.

Resolution of the image is significantly improved with backing due to shortened impulse response.

Contrast of the image is significantly improved with backing due to shortened impulse response.

Coverage of the image is only remotely affected by the acoustic backing layer.

1.8.7 In-probe electrical interconnect

In-probe electrical interconnect is a critical consideration for 2D ultrasonic phased array probe construction. An ideal interconnect should present low capacitance, low impedance (in non-waveguide regime), and most importantly, should not affect mechanical vibrational behaviour of the probe by being massless and volumeless. If any of these requirements are not met, the interconnect will have detrimental influence on the image quality.

1.8.8 Out-of-probe electrical connector (cable)

The cable is of similar consideration as the in-probe electrical interconnect. The use of a non-ideal cable can only have detrimental influence on image quality.

1.9 Probe construction considerations summary

To summarize, an ideal 2D phased array probe is characterised by:

- Appropriate element thickness and unimodal vibrational behaviour with a very short impulse response time
- Large circular aperture with dense element layout – implying lots of elements of $\lambda/2$ diameter
- Transducer layer acoustic impedance well matched to its load medium
- All elements electrically connected in an acoustically ‘transparent’ way.

Naturally, many of these postulates cannot be met in practice simultaneously; probe design is a matter of defining optimality criteria and engineering compromise. The following Chapters will discuss some of the avenues for pushing the boundaries of what is possible.

1.10 Typical signal acquisition and processing methods

The phased array probe can be operated in a variety of ways. The most typical ways are described in detail in the following publications accessible online from Olympus (Shinjuku Monolith, 3-1 Nishi-Shinjuku 2-chome, Shinjuku-ku, Tokyo, Japan) : [7, pp. 16–20] and [8, Ch. 2]. The signals acquired by the probe are amplified, filtered and digitized in the PAC and undergo digital signal processing. For the classic linear, azimuthal e.t.c. scans, that result in A,B,C e.t.c. views, the delay-and-sum processing is relatively simple and is typically implemented in the FPGA chip set inside the PAC. Alternatively, high throughput DSP processors are sometimes used in low-channel-count devices.

The increasing demands on data acquisition speeds call for advancements in signal processing techniques. Recently, compute-intensive techniques such as flash focusing, also known as volume focusing[8, Ch. 7.7], TFM[13] are finding their way into industrial use. Still, even more demanding methods such as PCF, CAPON require ever more signal processing power, typically not available for on-line, low-latency application.

1.11 List of contributions to knowledge

Sensor mechanical and electrical design and manufacturing technology.

Firstly, the sensor mechanical design has been addressed. The properties of piezoelectric ceramic composite materials have been studied. Detailed numerical models have been used to analyse conditions required for crafting materials of desirable properties. A novel technique has been demonstrated that allows for the design of a well-behaved triangular cut piezoelectric composite.

This composite then allows building of compact, dense-layout 2D phased array transducers with a hexagonal or sparse (triangular grid) element layout. The benefits of hexagonal element layout over classic, rectangular layouts have been analysed theoretically and shown to be beneficial.

Secondly, the problem of a high output impedance of the miniature ultrasonic sensor is addressed by means of an in-probe miniature signal conditioning circuit. The novelty and practical benefit in this case lies in the fact that no high power components are needed in the probe body.

Signal processing and CAD tools.

An emerging GP-GPU computer architecture offers opportunities to rethink the implementation of algorithms typically used in ultrasonic signal processing. Single-way beamforming, and two-way TFM and PCF beamforming have been developed for execution on the new platform.

These have been then enhanced by a new approach to solving the TimeofFlight problem in a way that is particularly well suited for the new platform. The unprecedented performance and low cost of this new approach enables industrial deployment of advanced beamforming methods, as well as building of practical CAD tools for engineering and education. A number of CAD tools have been implemented.

1.12 Publications associated with the Thesis

2009

Dziewierz, J., Gachagan, A., O'Leary, R., & Ramadas, S. N. (2009). Numerical optimisation of piezocomposite material properties using 3D finite - element modeling. Paper presented at 16th International Congress on Sound and Vibration (CSV16), Krakow, Poland.

2010

Dziewierz, J., Ramadas, S. N., Gachagan, A., O'Leary, R. L., & Hayward, G. (2010). *A 2D ultrasonic array design incorporating hexagonal-shaped elements for NDE applications*. 422-425. Paper presented at 2009 IEEE Ultrasonics Symposium, Rome, Italy.10.1109/ULTSYM.2009.5441447

Ramadas, S. N., Dziewierz, J., O'Leary, R. L., Gachagan, A., Velichko, A., & Wilcox, P. D. (2010). *An annular array with fiber composite microstructure for far field NDT imaging*

applications. 2557-2560. Paper presented at 2009 IEEE Ultrasonics Symposium, Rome, Italy.10.1109/ULTSYM.2009.5441460

Dziewierz, J., Ramadas, S. N., Gachagan, A., & O'Leary, R. (2010). Hexagonal array structure for 2D NDE applications. In D. O. Thompson, & D. E. Chimenti (Eds.), *Review of progress in quantitative nondestructive evaluation: AIP Conference Proceedings*. (Vol. 29, pp. 825-830). (AIP Conference Proceedings; Vol. 1211). Melville: American Institute of Physics.
10.1063/1.3362496

2012

Dziewierz, J., Gachagan, A., Lord, N., & Mulholland, A. (2012). An application-specific design approach for 2D ultrasonic arrays. In *51st Annual Conference of the British Institute of Non-Destructive Testing 2012 (NDT 2012): Proceedings of a meeting held 11-13 September 2012, Northamptonshire, UK*. Northampton: British Institute of Non-Destructive Testing.

Dziewierz, J., & Gachagan, A. (2012). Enhancing 2D phased array sensitivity and bandwidth using in probe electronics. In *2012 IEEE International Ultrasonics Symposium (IUS) Proceedings*. (pp. 2400-2403). 9781467345620: IEEE. 10.1109/ULTSYM.2012.0600

2013

Dziewierz, J., & Gachagan, A. (2013). Computationally efficient solution of Snell's Law of Refraction. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 60(6), 1256-1259. 10.1109/TUFFC.2013.2689

Dziewierz, J., & Gachagan, A. (2013). *A design methodology for 2D sparse NDE arrays using an efficient implementation of refracted - ray TFM*. Paper presented at 2013 Joint UFC, EFTF and PFM Symposium, Prague, Czech Republic.

Gachagan, A., & Dziejewicz, J. (2013). *2D ultrasonic array transducer design to maximise coverage in composite material structures*. Paper presented at NATO Workshop on Advanced Non-Destructive Evaluation Techniques for Polymer Based Composites in Military Vehicles, Riga, United Kingdom.

2014

McGilp, A., Dziejewicz, J., Lardner, T., MacKersie, J., & Gachagan, A. (2014). Inspection design using 2D phased array, TFM and cueMAP software. In D. E. Chimenti, L. J. Bond, & D. O. Thompson (Eds.), *40th Annual Review of Progress in Quantitative Nondestructive Evaluation: Incorporating the 10th International Conference on Barkhausen Noise and Micromagnetic Testing*. (Vol. 1581). Melville, NY: AIP Conference Proceedings. 10.1063/1.486480

Ramadas, S. N., Jackson, J., Dziejewicz, J., O'Leary, R., & Gachagan, A. (2014). Application of conformal map theory for design of 2-D ultrasonic array structure for ndt imaging application: a feasibility study. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, *61*(3), 496-504. 10.1109/TUFFC.2014.2933

Publications submitted or in preparation as of May 2015

J. Dziejewicz, S.N. Ramadas, R.L. O'Leary, A. Gachagan, Triangular-cut piezoelectric composite for 2D ultrasonic phased array probes.

J., Dziejewicz, T., Lardner & Gachagan, A Computationally efficient method of obtaining Time of Flight for 3D arbitrary curvature refracted-ray Total Focusing Method on GP-GPU processors.

Chapter 2. Design of piezoelectric ceramic composite

2.1 Introduction to piezoelectric ceramic - polymer composite material

The aim of combining materials into macroscopic composites is to amplify the desirable properties of constitutive materials and suppress their undesirable properties at the same time. Materials that are composite, that is, consist of more than one separable phase are ubiquitous both in nature and in engineering.

In previous research, a range of methods to mix piezoelectric ceramic and polymer have been investigated [14][15][16], [17][18][19][20][21]. Over time, 4 of these methods have found wider applications: 1-3 connectivity composite (Fig. 2.1), 2-2 connectivity composite, 0-3 connectivity composite (Fig. 2.2), and 3-0 connectivity composite.

The connectivity refers to the way material particles of each phase are connected. "0" means "not connected", "1" means "connected along 1 axis", "2" means "connected along two axes" and "3" means "connected along 3 axes". The first number corresponds to the ceramic, and the second number corresponds to the polymer.

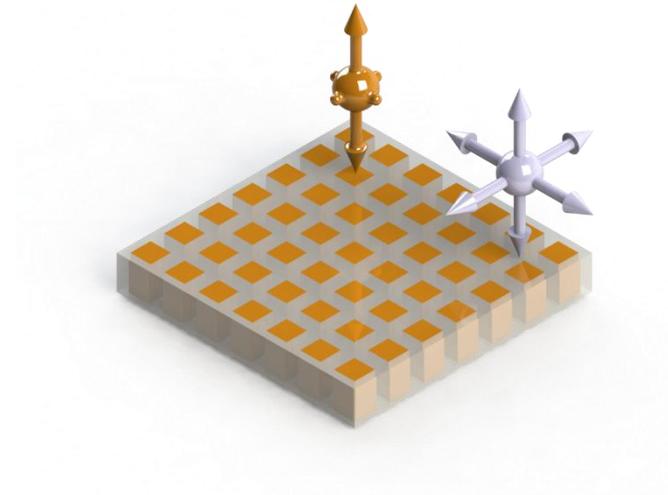


Fig. 2.1. An example of 1-3 connectivity piezoelectric composite. Ceramic pillars embedded in a polymer resin support.

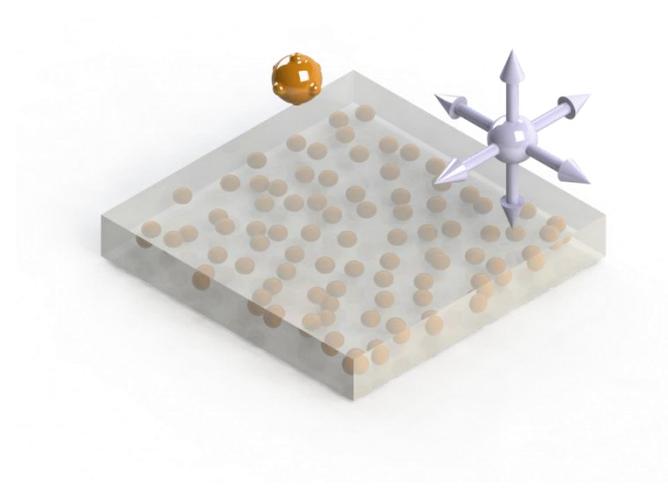


Fig. 2.2. An example of 0-3 connectivity piezoelectric composite. Ceramic particles embedded in a polymer resin support

Piezoelectric composite is used for the active part of the phased array probes at frequencies ranging from 30kHz to 40MHz. For frequencies below 30kHz, piezoelectric ceramic transducers are often less efficient and more expensive than electromagnetic coil 'speaker' type transducers. At high frequencies, the grain size of the piezoelectric ceramic material limits machinability of the material. In such cases 0-3 composites (ceramic powder directly mixed with polymer) are used, or alternatively, materials like PVDF can sometimes be easier to employ.

This work considers a frequency range of up to 40MHz, and therefore focus will be placed on the 1-3 connectivity composite.

2.2 Properties of 1-3 connectivity composite

The 1-3 connectivity composite (Fig. 2.1) comprises of *pillars* and *kerf*. This is a configuration that allows the composite to work more efficiently than solid piezoelectric ceramic (Fig. 2.3). This is due to the effect of the Poisson ratio in solid ceramic. When an electric field is applied to the poled ceramic, it attempts to extend in the poled direction and at the same time, contracts in the perpendicular directions. In solid materials, contraction is hampered by the fact that the material has to remain solid; the interior parts of the ceramic resist squeezing. By cutting the slab into tall pillars, the contraction in the XY direction is enabled, releasing the material to expand in Z direction, and increasing transduction efficiency.

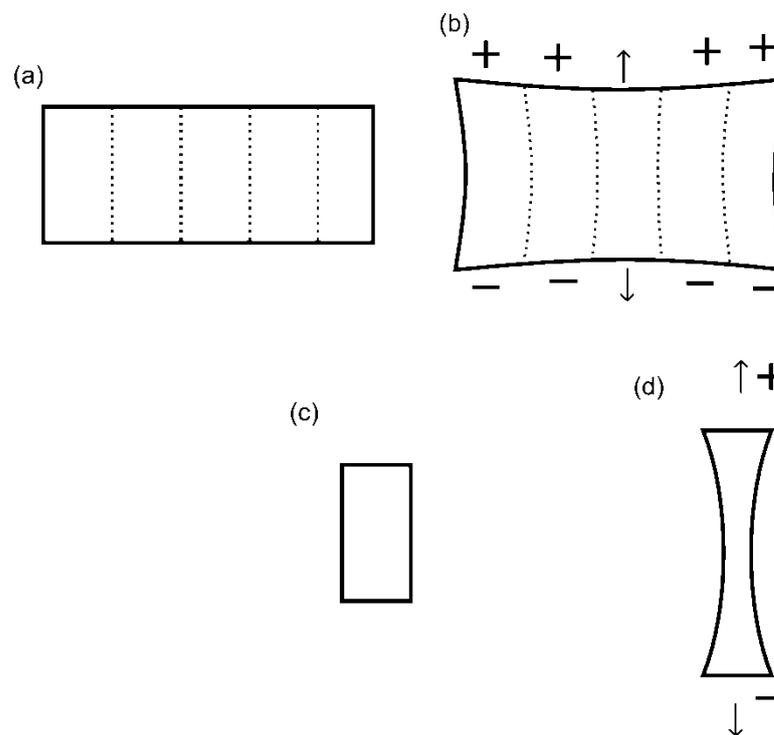


Fig. 2.3. A single piezoelectric pillar can expand more than a piezoelectric slab because it is less constrained in the lateral direction. (a) piezoelectric slab, unexcited state. (b) piezoelectric slab, excited state. (c) piezoelectric pillar, unexcited state. (d) piezoelectric pillar, excited state.

The kerf filler material serves several roles:

- It eases the shape change of the piezoelectric pillars.
- It helps to preserve the mechanical integrity of the composite.

- It provides mechanical damping and lowers the mechanical Q of the resonance.
- It helps to transfer the vibration energy to the external materials by effectively lowering the overall acoustic impedance of the composite.

2.3 The research question

The purpose of this work is to establish guidelines for design of the piezoelectric composite material with application specifically to 2D phased ultrasound arrays.

The following properties are desirable in the final material:

- A high sensitivity in pulse-echo mode, or pitch-catch mode, depending on application.
- A short impulse response of transduced signals - this corresponds to wide bandwidth.
- A low mechanical, electromechanical, and electrical crosstalk between neighbouring phased array elements.
- Unimodal behaviour (body resonant frequencies well separated from each other) – this contributes to a short impulse response and a high transduction efficiency.
- Ease of manufacture. This term covers both cost and reliability of the manufacturing process.
- Uniform element behaviour to make signal post-processing straightforward and reduce the need for calibration

However the above requirements are conflicting. Specifically:

- The bandwidth can be enhanced by increasing damping - this however reduces sensitivity.
- The low crosstalk between materials calls for a soft kerf filler, this however also reduces sensitivity and makes the unimodal behaviour condition difficult to satisfy.

- Unimodal behaviour calls for fine pillar pitch, this however increases the manufacturing complexity.
- If the design is driven by the array element size, the range of pillar pitch that can be used is restricted, potentially missing favourable dimensional configurations.

The basic parameters that can be altered, and have to be designed, are:

- The type of the ceramic used.
- Geometric Shape of the pillar.
- The type of filler material.
- The thickness of the material.
- The pillar pitch.
- The kerf width.

2.4 Literature review

The following review attempts to capture the history of development in piezoelectric composite design. Particular attention is given to developments that take into account the application of piezoelectric composite configuration for arrays.

Optimal pillar shape, dimensions, and way of selecting the kerf filler have been the focus of extensive research.

Gururaja et al. in [22] offers an early analysis of properties of 1-3, 3-3, and 3-1 composites. In 1-3 composites only cylinder rods are considered. In [23], he offers extended analysis and concludes that in case of highly damping polymer fillers, linear theory models are not reliable. Wang and Auld in [24] analyse the 1D case of a periodic composite and give expressions for tracing the dispersion curve and predicting lateral resonance frequency. They do not take longitudinal nor shear wave attenuation into account.

An even more detailed study is presented by Hossack et al. in [25]. Devices of operational frequency of 500kHz were analysed. Using simple 2D FEM, the authors analyse the influence of the number of pillars per array element on the thickness-mode

electromechanical efficiency coefficient k_t . In general, the more ceramic pillars per array element, the better the k_t . The reason for such behaviour is that when the pillars are thin and tall, the unavoidable lateral contraction forces in the ceramic become better absorbed by the compliance of the polymer filler, therefore leaving more energy for thickness mode extension. However, visible in this study is the fact that increasing the number of pillars increases the experimental deviation from simulation predictions. One explanation for this phenomenon might be that the ceramic becomes depoled during sawing - the thinner the pillars, the heavier the depoling. Although quantification of this effect is difficult, it is clear that there is a practical limit on how thin the pillars can be. This effect is much more exacerbated in high frequency composites, where pillar thickness is reduced below 0.5mm.

An alternative approach to manufacturing the piezoelectric composite is offered with success by M. de Espinosa et al. [26]. Instead of sawing or moulding ceramic pillars, he fractures the ceramic and then fills the voids with polymer resin under vacuum. This ensures that uniform poling of such a material remains difficult.

Since 1989, thanks to increasing computational power available, and numerical optimisation, the design has become possible. Ceramic volume fraction, saw pitch, and polymer selection and topology have all received due attention with computer models of increasing accuracy.

A representative set of modes and their frequencies for a single elementary piezoelectric block is given by Chalande [27]. The attempt is made to generate 'look-up tables' that would be suitable for all combinations of construction parameters. Unfortunately, results of other authors show strong influence of the polymer phase properties on the overall composite behaviour; hence, Chalande's results are of limited practical application.

Hossack in his analysis in [28] includes an assessment of equilateral, triangular-cut topology. Only modal analysis was used, and the damping effect of the polymer was not taken into account. A parametric sweep over width to height ratio of the ceramic pillars was performed and plots of coupling coefficient k_t were obtained. A modest increase of k_t for triangular-cut composite is predicted theoretically, but not observed experimentally. The research concludes with: "*...It is possible that suitable triangular structures will provide improved performance at lower volume fractions*". However, improvement is only obtained for triangular pillars arranged in rectangular patterns. The lower k_t of the main mode of

vibration of the 'regular' triangular cut composite might stem from the fact that it was significantly bi-modal. Pertinently, for all options (including circular and square pillars), for a given resonant mode, the maximum in the plot is observed when coupling of other modes is low. No design for unimodality was considered and k_t of up to 0.60 was reported.

Rapidly increasing computational power available for research allows Hossack [29] to perform a useful parameter sweep study on the effects of pillar geometry and sizing and spacing by using increasingly accurate 3D eigenvalue FEM. Although eigenvalue solutions are later found to be systematically off the real values and limited in the results they can provide (eigenvalue problem formulation assumes continuous oscillation mode only and no damping), the general trends shown are true for a wide range of parameters and geometries. Regular triangular pillar geometry is considered, but found not to be unimodal, and alternative arrangement of triangular pillars is proposed to remedy the issue. Triangular pillars arranged in a rectangular matrix show unimodal behaviour and high k_t in a range of 0.68 is reported in an experimental 600kHz device.

A very interesting approach to suppressing lateral resonances is again presented by Auld and Hossack in [30]. Instead of using a regular grid of pillars, they propose to apply Gaussian distributed spacing of pillars in each direction. This method works well for a uniformly electroded piezoelectric composite slab, visibly smearing out spurious resonances and resulting in a measurable increase in k_t . However, results for a 1D array nor 2D array are unknown. Since the smearing effect only works over a large slab, it is anticipated that non-uniform distribution of pillar sizes would result in a distribution of array element sizes, with corresponding problems arising in terms of per-element sensitivity and beam simulation.

Hamilton and Hayward obtain interesting result in [31], as in their analysis of various pillar shapes one of the conclusions is that regular triangular pillars can result in the best surface dilatation quality for high aspect ratio pillar composites.

In their work, the new "MPAR" factor is defined as a ratio of the size of the pillar base to the pillar height, at which the first lateral mode frequency occurs at the first harmonic frequency of the base thickness mode.

In 1993 the effect of loss in polymer phase has been incorporated into the numerical study by Bennet in [32] and [33], using harmonic analysis rather than eigenvalue analysis, but still

on a single symmetric cell. (Lossy materials have been previously only covered by 1D models, e.g. like in [34]). Importantly, he notes: “(...) *the MPAR for the high-loss composite is not constrained by the lateral composite modes, but by the modes contained within the pillar itself. This implies that the effective MPAR can be larger **than** for the low-loss devices.*” He then shows the effect of polymer loss on two devices, 250kHz and 700kHz, having non-fundamental modes of vibration effectively suppressed. In this work the polymer damping properties were estimated by matching simulated results to the experimentally obtained impedance. The general understanding is that the loss matrix coefficients obtained this way are only valid for the particular device and not transferable to a different model, and in particular, not valid at a different frequency. One of the results obtained is that MPAR is higher for stiff polymers, so the polymers of low loss counteract the effect of damping on MPAR. At the same time, the uses of softer polymer results in a significant drop in absolute transmit sensitivity (0.9 vs 5.8 nm/V).

Hyslop et al. in [35], briefly summarizes his analysis of the effect of the composite topology on the MPAR (Maximum Pillar Aspect Ratio). MPAR for triangular-cut composite is significantly higher at 0.39 than for any of the square-cut composites simulated (up to 0.31). In this work a model larger than a single pillar is analysed, having the full single element of the array along with the area around it simulated. He also notes that one of the positive effects of lossy polymer fillers is that the secondary vibration modes are heavily damped, up to the point where they are undetectable. Pertinently, the coupling factor k_t , along with bandwidth, is higher when only a single element of the composite plate is excited, when compared to entire plate excitation. This is due to both less restrictive movement of the separated array element and additional damping provided by the surrounding, unexcited material. This effect should be kept in mind when comparing efficiency of various array designs.

Gachagan et al. in [36] performs a 1D parameter sweep search on Pillar Aspect Ratio for a probe to be operated at 1.8MHz. Bi-modal behaviour is utilized to extend the bandwidth of the probe in transmit mode. Combining optimal coupling case with backing and dual matching layer, he obtains 90% bandwidth. Importantly, phase of the two modes seem to be aligned to such an extent that a nearly unidirectional pulse shape is obtained, which is an excellent result. This occurs at an expense of overall coupling efficiency, however one can easily imagine applications in which such approach is beneficial. This shows that FEA

can be utilized to find cases of a combination of parameters that not only maximises certain metrics (typically k_t) but also results in special properties.

Certon et al. [37] again uses FEA and a parameter sweep to establish the influence of stiffness of the polymer phase on the lateral resonances. Importantly, they use a 1/8th symmetric cell of a square cut composite. Their calculations show that increasing transverse wave velocity of the polymer pushes the lateral resonance frequency up, and they suggest the use of 22% boron carbide powder loaded epoxy as a means of obtaining the effect. (Boron has a very low density and a very high stiffness, resulting in extreme acoustic properties)

E.C.N. Silva et al. attempts to numerically optimise topology of the piezoelectric composite cell by discretizing the volume and trying out various ceramic/void microstructures [38]. However, manufacturability of such composites has not been discussed.

J.A. Brown et al. [39] modify the classic, square cut composite design with a third, 45° cut to obtain a 40MHz piezoelectric composite for use in a 1D phased array. Apart from the expected effect of removing unwanted resonances and making the impedance look unimodal, the measured pulse width becomes significantly shorter, improving the final probe axial resolution, albeit by expense of reducing the sensitivity by approximately 5dB.

When low ceramic volume fractions are acceptable, J. Yin et al. proposed a hexagonal-only pillar geometry composite for frequencies of up to 20MHz [40]. Good results were obtained for single-element devices with k_t up to 0.62, however only for highest frequency composites. At lower frequencies, the effects of intra-pillar resonances are clearly visible and result in a much lower k_t . Due to this, and despite the use of 100% matched impedance backing, distortion to the pulse shape is visible in single-way transmission.

Alternative manufacturing methods do exist. Recently, moulding techniques [41] are being developed into micro moulding techniques [42], which offer completely arbitrary pillar shape for frequencies up to 100MHz.

Summary

Although various shapes are possible this literature review shows that for 1-3 piezoelectric composites, square base pillars offer the best sensitivity with manageable unimodality, and are also possible to manufacture for frequencies up to 20MHz.

The final properties of the piezoelectric composite depend non-linearly on the properties of its constituting materials, and the way that they are composed. There is a number of materials that can be combined, and this leads to a design problem - what materials to combine, and in what quantities, to obtain the best possible result?

Both experimental and modelling techniques are used to answer this question. Modelling is the preferred start point as given an accurate model it becomes cheap to evaluate multiple scenarios. However, accuracy of the model depends on how detailed it is, what assumptions and approximations have been made, and most importantly, on the accuracy of the relation between modelled material properties and actual material properties, which are often hard to measure and may change over time.

2.5 About 1D Linear systems modelling

With current desktop computer computational power, simple mathematical models derived from first principles offer near instant solutions to forward modelling problems. There exists a number of modelling schemes for piezoelectric ceramic materials, supporting insertion of the transducer into electrical and mechanical environments, for example, 1D Krimholtz Leedom Matthaei (KLM) model [43] or linear systems (LSM) model [44]. These linear models of a piezoelectric transducer use 'lumped' material parameters. In case of a piezoelectric ceramic-polymer composite, these can be obtained from another model, by Smith and Auld[17] which has been validated with FEM and experimentally by numerous researchers, e.g. [45][46].

However, for purposes of this work, 1D models do not model the internal behaviour of the composite material itself, in particular:

- Modes of vibration of the piezoelectric composite structure itself, including the desired fundamental thickness resonances and parasitic lateral modes, as considered in [16][19].
- The dilatation quality of the transducer surface, which takes into account possible anti-phase behaviours.

- The behaviour of the polymer phase and its detrimental effect on the overall acoustic performance of the transducer, as considered in [33].
- The mechanical crosstalk between array elements, which must be minimised to ensure good directional performance from the linear array structure, as considered in [19], [47], [48].

In order to address these design considerations whilst retaining low model complexity, a simplified 3D model will now be proposed.

2.6 Some known design parameter interactions

When changing the ceramic volume fraction or polymer composition, the effective longitudinal and shear wave speeds travelling through the composite structure also change [17]. Therefore, the thickness of the active layer must be adjusted to compensate for any changes in these parameters to match the desired frequency of operation.

Gain-Bandwidth Product (GBP) is a nonlinear function of the ratio of piezoelectric composite acoustic impedance to load material acoustic impedance.

The acoustic impedance of the composite depends on the ceramic acoustic impedance, ceramic volume fraction and the polymer properties in a piezoelectric composite structure. The higher the impedance mismatch between device and load results in a reduction in terms of device bandwidth and sensitivity, as less energy can be transferred between the device and the load. However, when lowering the ceramic volume fraction to improve impedance matching, the volume of active piezoelectric ceramic also reduces what impacts on the electro-mechanical efficiency.

The softer the kerf filler used, the easier it is for the ceramic to expand, because the polymer absorbs lateral contraction of the ceramic pillars as they elongate. However, softer polymer materials have a corresponding higher longitudinal wave damping, which also limits GBP. Additionally, with very soft polymers and high ceramic volume fractions, the kerf filler surface tends to vibrate in anti-phase, being pushed out by the mechanically stiff ceramic [33]. Similar effects may occur at very low ceramic volume fractions, when the polymer does not bond well to the ceramic pillars and tends to vibrate independently.

Accordingly, these effects reduce GBP and may give rise to parasitic surface Lamb waves. Therefore stiffer polymers have to be used at large ceramic volume fractions.

The impact of material thickness is more straightforward: the thicker the device, the greater the volume of piezoelectric ceramic, which results in a device with enhanced sensitivity. However, thickness cannot be changed freely because the operational resonant frequency defines the material's thickness.

All three parameters have a complex influence on mechanical crosstalk. Softer polymer materials exhibit higher shear wave damping, and hence will result in lower mechanical crosstalk; however, longitudinal wave damping also rises and effectively reduces GBP. Moreover, at certain thickness - wave speed combinations, Lamb wave modes of the periodic 1-3 connectivity piezoelectric composite configuration may come close to the operational frequency of the device. This would cause a significant increase in mechanical crosstalk and is in fact observed in the simulated results found in this work, as explained later in the Chapter. However, careful design can lead to particular ceramic volume fractions which may not be affected by this phenomenon because their periodic structure creates a stop band filter for these Lamb wave modes [19].

Given all these dependencies, it is clear that a global optimisation of design parameters using an accurate modelling strategy would enhance understanding of the design process. For this work, a 3D FE model of a piezoelectric composite array structure has been used to predict important ultrasonic material parameters: GBP and mechanical cross talk. These are evaluated for a range of input parameters as described later in this Chapter.

2.7 Simplified 3D model of regular 1-3 composite with load medium

From the literature review it is apparent that there are a large number of design variables that influence the physical behaviour of a 1-3 connectivity piezoelectric composite material. In this work, the variation of three design parameters and their impact on the performance of the final array device, is considered:

1. Kerf filler material composition (indexed by polymer stiffness).
2. Kerf size to pillar size ratio – indexed by Ceramic Volume Fraction.

3. Overall composite thickness.

A design rule is adopted - the size of each linear array element is determined by the half-lambda design rule, and an integral number of pillars must constitute a single array element. These assumptions reduce the number of geometric design parameters.

2.7.1 Exploration of the problem space

A 3D FE model of a 1-3 connectivity piezoelectric composite array structure has been developed using PZFlex® (Weidlinger Associates, CA). Fig. 2.4 depicts a top down representation of a full 1-3 connectivity piezoelectric composite configuration, in which only the area identified by the solid box has been modelled. This constitutes the maximum reduction in the size of the FE model through symmetry and comprises two adjacent array elements, each consisting of $1\frac{1}{2}$ piezoelectric ceramic pillars. Hence, when symmetrical boundaries are applied to all four sides within the FE code (as illustrated in Fig. 2.5), the model will simulate alternating linear array elements comprising of 3 piezoelectric ceramic pillars. It is important to note that one half of the model, representative of a single linear array element ($1\frac{1}{2}$ ceramic pillars) is electrically excited, with the other half always in receive mode (connected to a 50Ω load resistor). This approach will provide an estimate of mechanical crosstalk between array elements.

The crosstalk values returned from this model are not directly comparable to real-world values, because in this symmetrical model every other element in an infinite array is simultaneously excited. Importantly, this assumption does not impede optimising the design towards low crosstalk.

The array elements are mechanically interfaced with the load medium as illustrated in Fig. 2.5. The load medium for this work is Poly (methyl methacrylate) - commonly referred to as Perspex – a typical NDE wedge material. The load extends for 3 wavelengths to facilitate surface wave and near-field wave interaction. The top face of the load medium is terminated by an absorbing boundary to prevent energy reflecting from this surface back to the active piezoelectric composite layer.

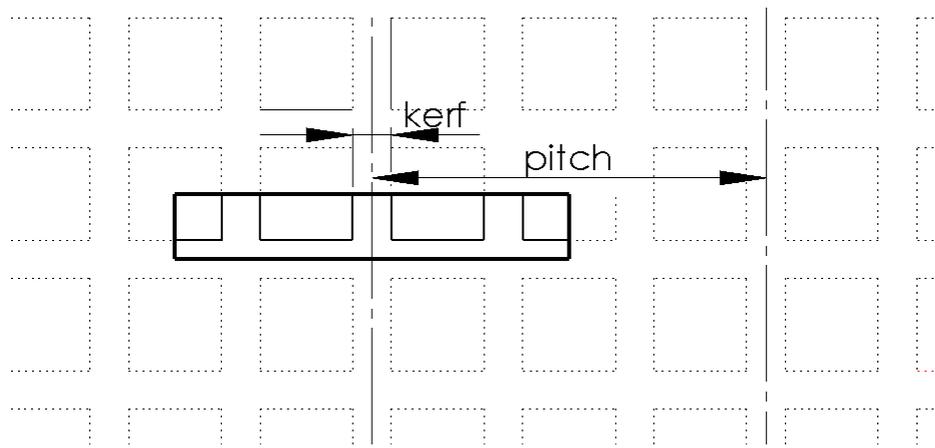


Fig. 2.4. A top-down view of the model. The model boundaries are marked with black. The pitch is the array element pitch (not the pillar pitch).

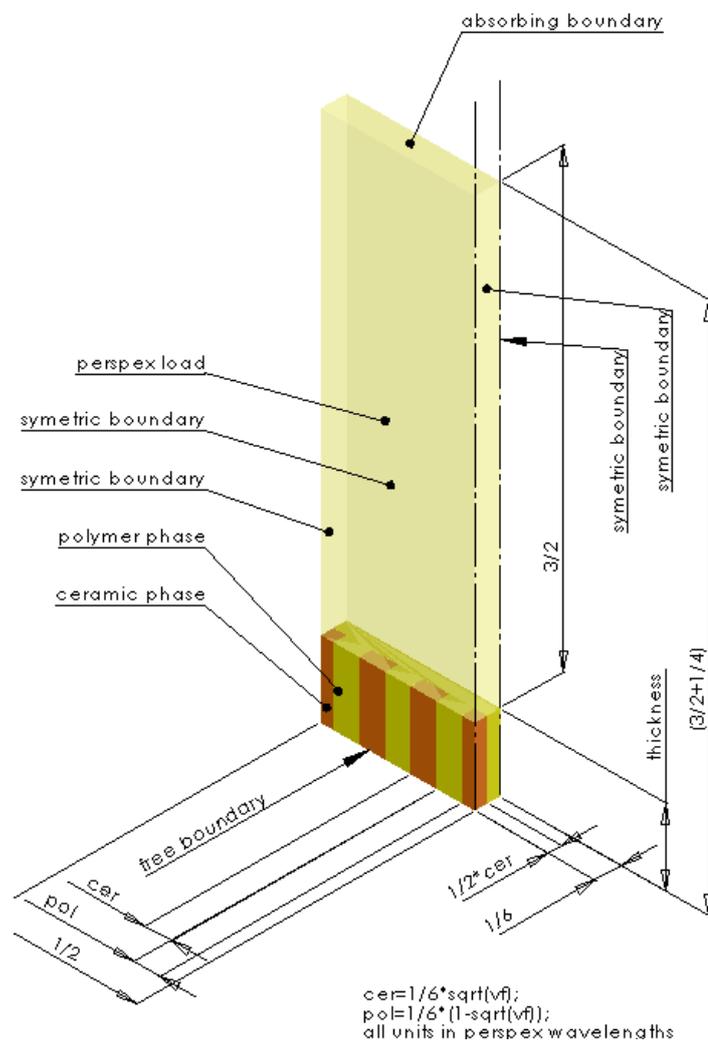


Fig. 2.5. A model view – isometric. Side boundaries have symmetric condition applied. The top boundary has absorbing condition applied. The bottom boundary has a free surface condition applied.

In order to facilitate a parameter sweep over a minimal, but meaningful problem space, the size of the array element width is fixed to half-wavelength in perspex at the transducer operating frequency. The piezoelectric composite material consists of PZ29 piezoelectric ceramic (Ferroperm Piezoceramics A/S, Denmark) for the piezoelectric ceramic phase, and for the kerf filler phase, a special polymer material developed by the University of Strathclyde for implementation in piezoelectric composite array configurations has been used [46]. The chemical composition of the polymer is indexed by its stiffness. The properties relating polymer stiffness to the longitudinal and shear damping of the material used are illustrated in Fig.2.6 and Fig.2.7 respectively.

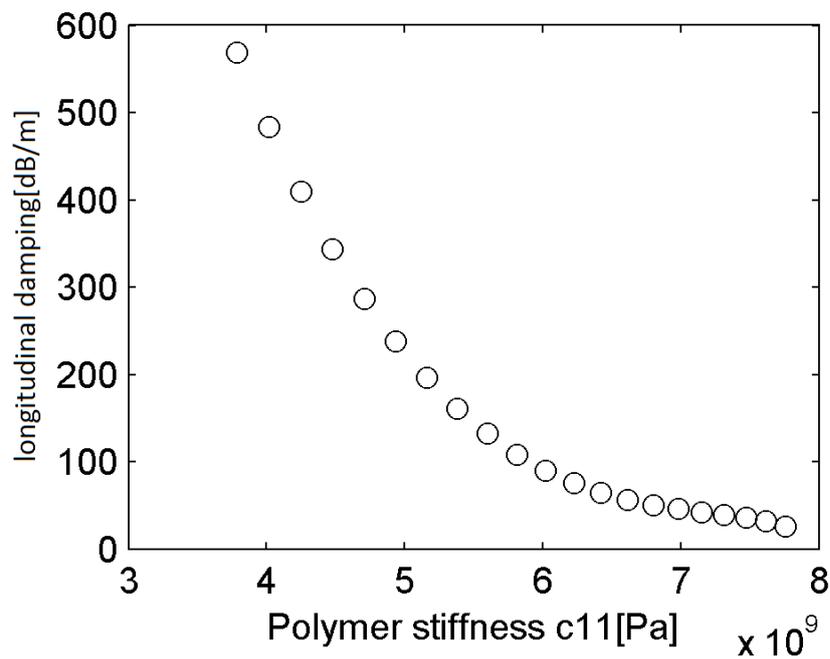


Fig.2.6. Material properties: longitudinal wave damping in polymer vs. polymer stiffness

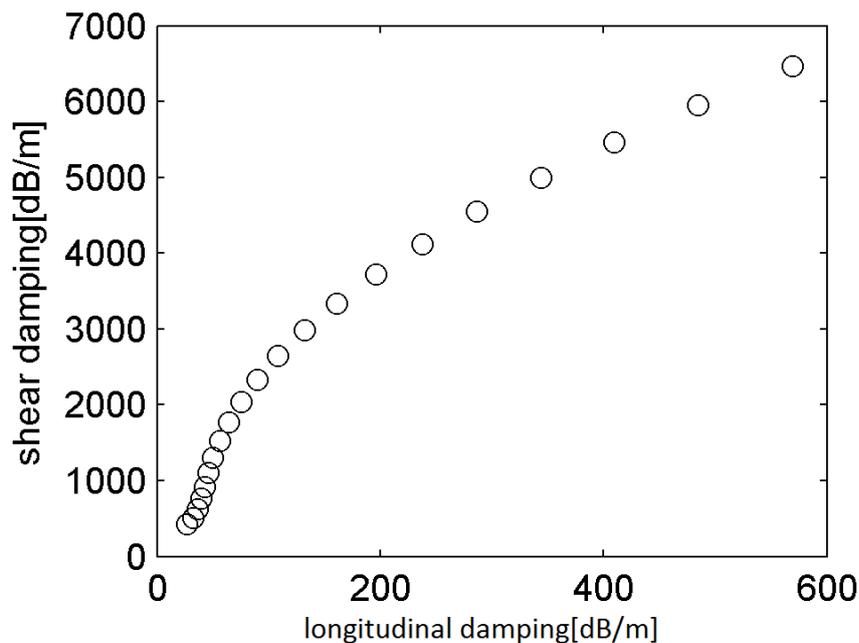


Fig.2.7. Material properties: characteristics of damping the polymer used in the simulation.

The simulation run comprises of three phases: Transmit Phase, Receive Phase, and Post - Processing. In the Transmit Phase, one of the simulated array elements is set up in a transmit mode, and excited by a wideband electrical pulse. The adjacent element is set up in receive mode. The pressure history is recorded at the absorbing boundary of the load medium for use in the next phase of the simulation. Additionally, the displacement of the surface of the piezoelectric composite is recorded to provide an estimate of the mechanical crosstalk between excited and neighbouring elements.

In the Receive Phase, both modelled array elements are switched into receive mode, and their electrodes are connected to a load resistor (50Ω). The pressure history acquired during the Transmit Phase is used to simulate a returning echo/signal from an imaginary reflector located in the far field in the load medium. The voltage output from the receiving array elements is recorded. In the Post-Processing phase, two figures of merit are calculated: the Gain-Bandwidth Product of the pulse-echo signal, and the RMS value of mechanical crosstalk (in dB) between the array elements when operating in transmit mode. These results, and the corresponding waveforms, are saved for further analysis.

In this work, GBP was calculated over a windowed frequency range between 2.0MHz to 3.0MHz to promote devices operating with a desired centre frequency of 2.5MHz.

2.7.2 Running the parameter sweep

It takes less than 15 minutes to run a test case using this FE model for a single case of design parameters on a 1GB RAM, dual-AMD Opteron™ 246 Processor PC. At first, running a classic optimisation of the parameters was considered. However, as each case takes considerable time to calculate, and the MATLAB (The MathWorks Inc., Natick, MA, USA, 2015) optimisation toolbox does not support restarting an aborted process, a decision on exploring the entire constrained problem space was made. 3995 test cases have been generated to cover a range of thicknesses, ceramic volume fractions and polymer stiffness combinations. Specific ranges have been selected to be realistic and realisable using conventional manufacturing techniques.

Multiple computers, working in parallel, have been used to speed up the calculation process in the following manner: a parameter set for each test case has been written to a numbered file accessible from a shared location in the local network. Then, each computer participating in calculations picks a file containing a test case description, runs three phases of the calculations, stores the results into a zip file, and returns it to a shared location. Finally, a single node reads the zipped files and collects all the results for analysis. This method allows for simple parallelisation of the computations. It is important to note, that such distributed mode of operation is now known as 'cloud computing' and it is commercially available as a service.

As the FE model results are cached in the last stage, any optimisation algorithm may be used to search inside the pre-calculated result set. Test cases were processed in a randomized order to enable an early preview of the results.

2.7.3 Discussion of results

Fig. 2.8 depicts the relationship between GBP and mechanical crosstalk for selected test cases. Results with GBP lower than 80000 were automatically considered unsuitable for implementation and discarded.

It is desirable for any transducer array to have a good sensitivity and a low crosstalk. A range of results in the upper-left corner (i.e. shaded box) of the plot can be interpreted as satisfying this requirement. It is clear that there is a compromise to be considered, as it is

possible to select either a case with high GBP, or one with a very low crosstalk, or a blend of the two.

From Fig. 2.8 it is also apparent that higher mechanical-crosstalk can also result in lowering of the GBP. It is hypothesised that this is because a fraction of the mechanical energy generated by the active element is absorbed by the neighbouring (passive) element, and it is effectively damped instead of being transmitted into load.

Fig. 2.8 demonstrates that it is indeed viable to conduct numerical optimisation of the piezoelectric composite array transducer design using a 3D FE modelling approach. From these results it is possible to identify a region of optimal solutions that minimizes crosstalk and maximizes GBP. The FE results will now be examined in more detail and the crosstalk and GBP results discussed in terms of the optimiser variables: polymer stiffness, CVF, and composite thickness.

These results can be related back to the identified grouping in Fig. 2.8 (shaded box) and will be marked as “acceptable solutions” in the following figures.

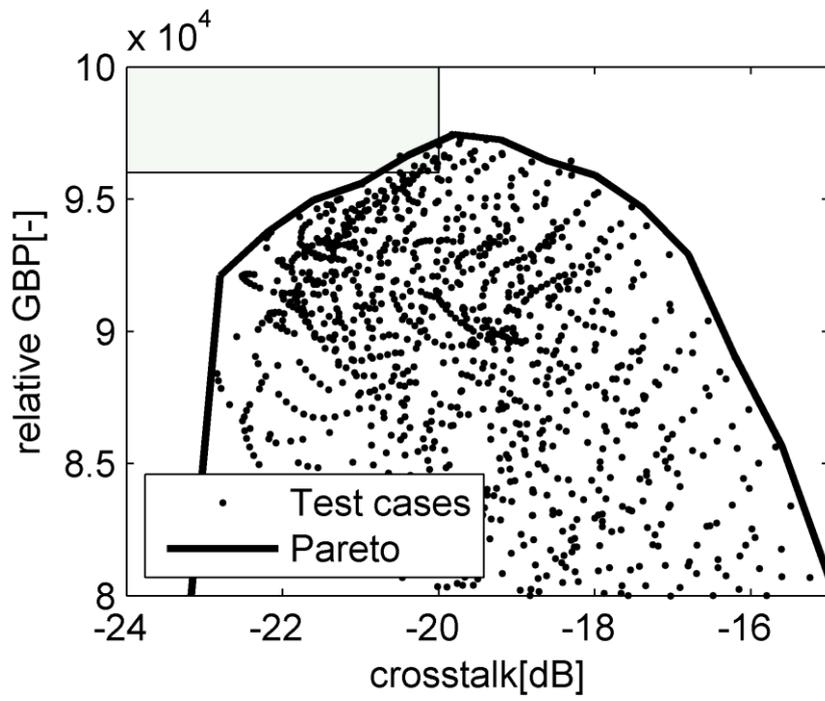


Fig. 2.8. GBP vs. crosstalk for all test cases

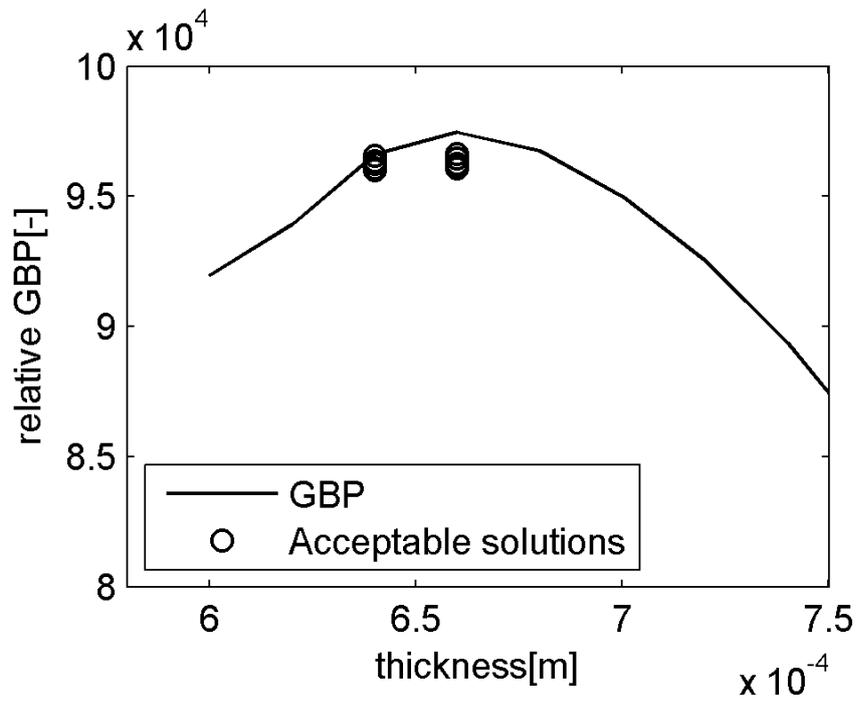


Fig. 2.9. Gain-Bandwidth product vs. composite thickness.

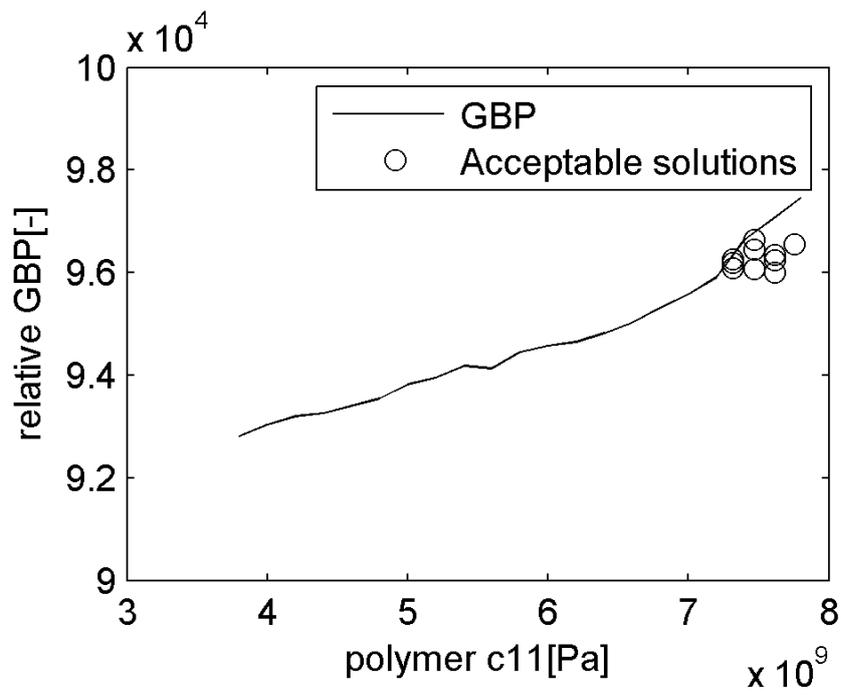


Fig. 2.10. GBP vs. polymer stiffness

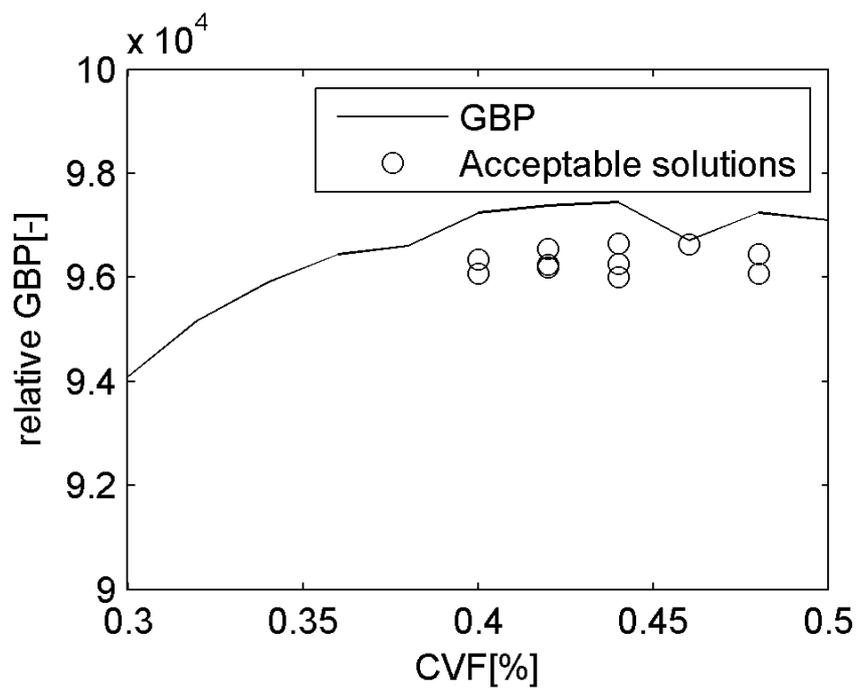


Fig. 2.11. GBP vs. ceramic volume fraction

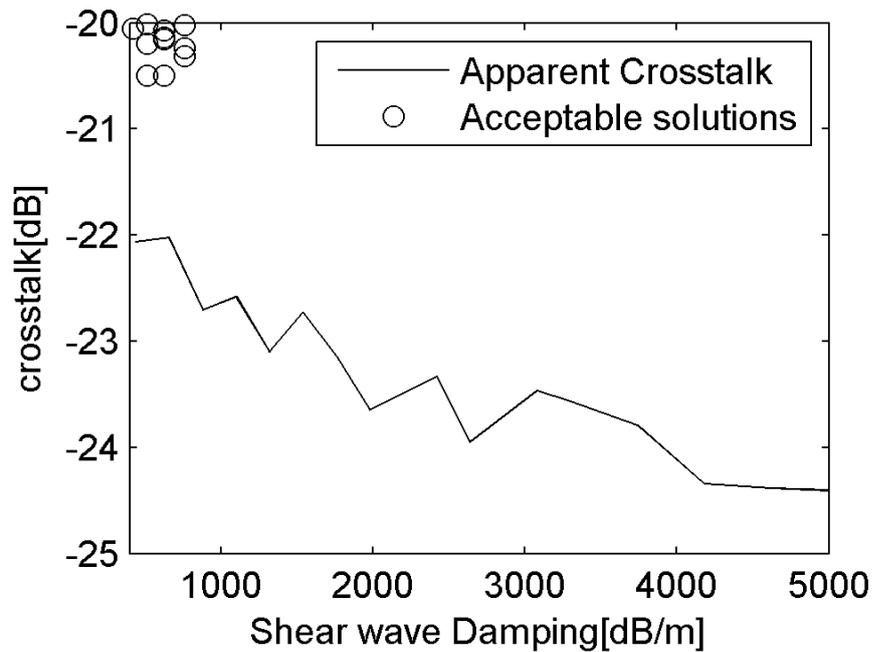


Fig. 2.12 Best crosstalk vs Polymer shear wave-damping ratio

The relationship between GBP and piezoelectric composite thickness, polymer stiffness, and CVF are presented in Fig. 2.9, Fig. 2.10, and Fig. 2.11 respectively. The continuous line in these figures depicts the best possible GBP at a given parameter, and small circles indicate the solutions for cases identified from Fig. 2.8 potential design solutions (that is, located within the shaded box). It is important to note that the GBP calculation was weighted to promote devices that are most effective around the desired 2.5MHz centre frequency.

The results shown in Fig. 2.9 predict that there are two piezoelectric composite thicknesses from the acceptable solutions that produce a high GBP value. In fact, it can be concluded that there are acceptable results for devices in the thickness range of 0.64mm to 0.68mm.

This shows that although there may be a specific thickness that gives an ideal centre frequency, a thicker device will comprise more active piezoelectric ceramic material and produce a stronger electric field, even if the piezoelectric composite is not resonant at the frequency of interest.

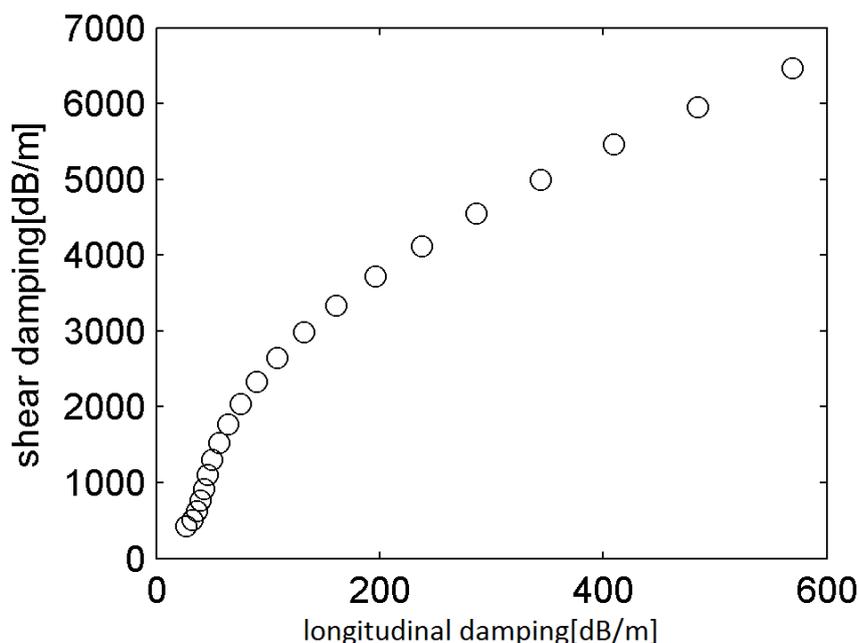


Fig. 2.13. Shear wave damping characteristic of the polymer range used.

The dependence of GBP on polymer stiffness is shown in Fig. 2.10, and there is a general trend for the GBP to increase with an increase in polymer stiffness. However, mechanical crosstalk tends to increase as well and this is illustrated in Fig. 2.12. The basis for this behaviour can be explained by examining the variation of polymer stiffness with both longitudinal and shear-damping properties, as illustrated in Fig.2.6 and Fig.2.7. The polymer material of a higher stiffness is also the one with significantly reduced shear damping. Therefore there will exist a specific optimal polymer type of given stiffness and of shear wave damping properties that will maximise the overall desired quality factor and push the pareto line.

Moreover, the acceptable result values are more spread out in the GBP versus CVF relationship as illustrated in Fig. 2.11, when compared to the results for thickness and polymer stiffness in Fig. 2.9 and Fig. 2.10. This indicates that, given the pareto line of GBP/crosstalk as illustrated in Fig. 2.8 CVF may be a function of thickness and/or polymer stiffness.

It is apparent that the properties of the passive polymeric material employed in the device will have a profound influence on performance of the device. This has been further evaluated through this work by considering a wide range of polymer materials that differ in their stiffness and damping properties (see Fig.2.6 and Fig.2.7). Fig. 2.10, and Fig.

2.11 consider the relationship between mechanical cross talk in array structures and the polymer stiffness and shear wave damping respectively. From the acceptable solutions depicted in these figures it can be determined that the ideal formulation of polymer, in terms of reduced mechanical cross-talk, would comprise low stiffness with a shear damping value towards the lower end of this data set.

Finally, to investigate the distributed nature of the GBP versus CVF (Fig. 2.11), the modelled results have been utilised to determine the interdependence between CVF and the two other design variables in this study: polymer stiffness (Fig. 2.14) and piezoelectric composite thickness (Fig. 2.15). The relationship between CVF and polymer stiffness is again fairly distributed and indicates that there are other significant parameters involved in determining these predictions.

Pertinently, a simple relationship between CVF and device thickness can also be extracted. However, it should be noted that this is a relatively small dataset and it is not certain if such a relationship will remain valid for all piezoelectric composite array designs.

2.8 Summary

A distributed computing environment has been used to perform a brute-force search on several design parameters associated with 1-3 connectivity piezocomposite array structures commonly used for ultrasonic transducer array devices.

The parameters optimised were material thickness, ceramic volume fraction, and polymer stiffness. The Gain-Bandwidth product in pulse-echo operation, and inter-element crosstalk were considered as performance indicators by an optimisation process.

Notwithstanding the relatively complex results from this process, this optimisation approach has shown that it is viable to conduct such a search to aid the design process associated with ultrasonic array transducers.

Importantly, it has been demonstrated that an excessive amount of inter-element crosstalk may, apart from its inherent undesirability, lead to a reduction of effective gain-bandwidth.

The final recommended design parameters have been collected in Fig. 2.14, Fig. 2.15, and Table 1.

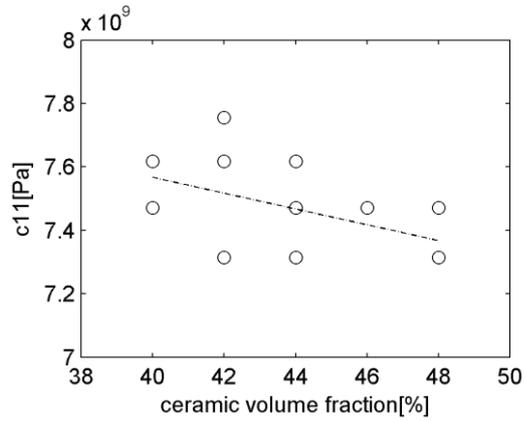


Fig. 2.14. Polymer stiffness vs. CVF for selected near-optimal cases.

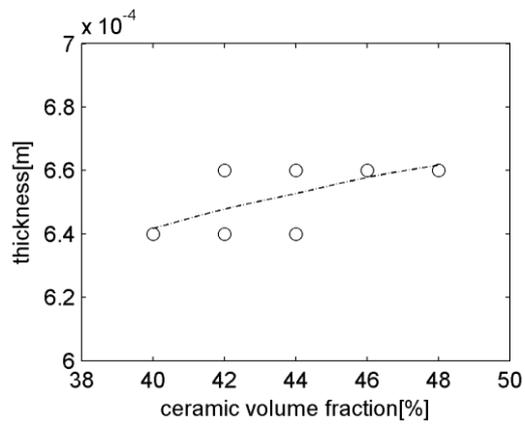


Fig. 2.15. Composite thickness vs. CVF for selected near-optimal cases.

Table 1 Final piezoelectric composite design parameters for a 2.5MHz array

Piezoelectric ceramic	Polymer Stiffness	CVF	Thickness
PZ29	7.75GPa	44%	0.66mm

With this work a modelling tool set has been developed that enables detailed study of the properties of piezoelectric composite materials in application to 2D ultrasonic NDE array probes. This tool can now be used to study novel piezocomposite structures, such as one discussed in Chapter 4.

Chapter 3. Properties of hexagonal elements

3.1 Introduction

2D phased array probes when compared to now well-established 1D array probes promise to significantly reduce the time of inspection. This is by reducing the complexity of mechanical scanning, and acquiring more information from a given probe location. Ultimately, they promise a higher plant reliability and a lower total cost of ownership.

They generate large volumes of information but with decreasing costs of data storage and signal processing, it may soon become cost-effective to apply the technology of 2D phased ultrasound arrays into the field of Non-Destructive Evaluation (NDE).

For the NDE inspection to be successful, the ultrasonic images obtained should present high resolutions and contrasts to the operator. However, currently, even state of the art technology doesn't provide the probe sensitivity and the image contrast needed for many types of inspection of interest.

The purpose of this work is to compare hexagon and square shapes and their suitability for building densely sampled phased array ultrasonic probes.

For the purposes of this Chapter, all quantities are expressed in terms of wavelengths, so that the results presented are easily interpreted across a range of operating frequencies. Also, the effect of the reflector characteristic is not considered. Consequently, only the monochromatic single-way beam forming simulation is utilised as this is a classic method and representative of array performance, and an easy way to compare array element shapes and layouts. It should be noted that the effect of special imaging algorithms on the image quality is outside of the scope of the work presented in this Chapter.

3.2 Relation of image quality and measurable quantities

When looking for a probe that would deliver the best possible image quality one has to define what the image quality is and how to measure and compare it. For the purposes of this study it is assumed that the image quality consists of 4 measurable factors:

- Beam spot size with a smaller quantity implying better image resolution. Where possible, the measure is the volume of the space bounded by an isoamplitude surface measured at -3dB relative to the peak amplitude inside the volume. Where 2D or 1D simulations are conducted, area or length of the focal region can be considered as suitable metrics respectively.
- Image contrast. This is defined by two interconnected quantities:
 - Side lobe to main lobe ratio. Typical phased array beam forming artefacts are called “side lobes” – that is, volumes/areas/directions outside the main lobe, at which the sensor exhibits sensitivity and may produce spurious features (artefacts) in the processed array image.
 - Energy leakage. This is a measure of a relative energy ratio received from the volume/area/direction of the focussed beam spot versus energy received from everywhere else in the vicinity of the probe
- Relative image intensity. It is natural that the sensitivity of the probe to the returning echo depends not only on the reflector type, but also on its spatial location versus the probe elements location. Some probe designs may be able to capture more of the echo energy than others. In order to equalize image intensity, a spatial gain correction is often applied. However, higher gain implies higher noise and there exist a certain practical limit of how much gain can be applied. Therefore this metric is justified for comparing probe designs.

In the following work the above measures will be used to find ways of designing a probe.

3.3 Element layouts for dense phased arrays

A 2D phased array is said to be dense if it fulfils the Nyquist Criterion of sampling for the acoustic field at the wavelength of interest. The elements are placed in such a way that the

acoustic field is sampled with adjacent elements located at a distance of $\lambda/2$, or less at the wavelength at which the probe is specified (typically the frequency of the highest sensitivity). This way full information about the impinging acoustic field of the specific wavelength can be acquired.

One way to increase the sensitivity of the probe is by increasing the aperture of the sensing elements so that each channel can acquire more signal energy, and breaking the $\lambda/2$ limit. The next limit here stems from the requirement of obtaining a sufficiently wide directivity (acceptance angle) for each element, as the phase of the impinging signal is integrated (and smeared) over the sensing area of the probe element.

A typical modern dense 2D phased array element layout is created by a natural extension of 1D phased array design, which is to sub-divide the elements to create a 2D matrix, as illustrated in Fig. 3.1.

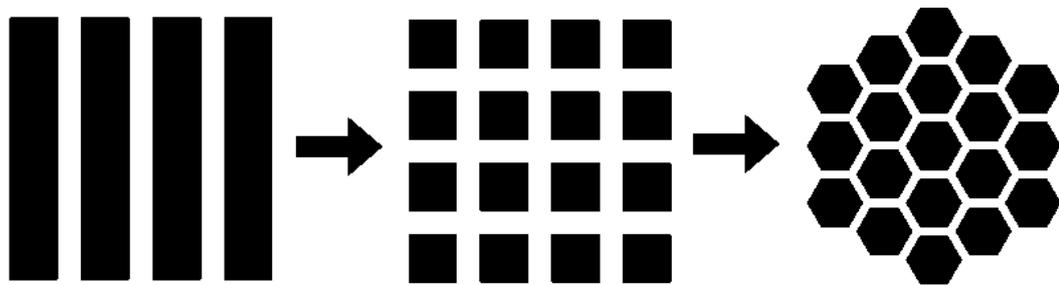


Fig. 3.1. Left: Layout of 1D Phased Array elements. Middle: 2D phased array layout created by subdividing 1D phased array elements. Right: Proposed 2D phased array layout with hexagonal elements.

Other element configurations include sparse layouts, that is, where probe aperture is not fully covered with active elements, or element-to-element centre-to-centre distance is significantly larger than half of the dominant wavelength [1]. Such configurations are outside of the scope of this Chapter.

Here it is proposed to replace classic, square elements with hexagonally shaped elements, as shown in the right hand image in Fig. 3.1. The rationale for doing so is presented in this Chapter.

3.4 Literature review

Hexagonal array element layouts are generally well known in the radar and optical telescope literature. For example, the author of a 1961 paper [49] analyses in a great detail the improvements in the radio antenna array element count, under assumption of not forming the grating lobes, and using non-focussed far field beam amplitude equations. Importantly, he notes that the improvement depends on the range of steering angles requested, as the steering range is limited by the angle where the first grating lobe appears. He finds that for a design requirement of a spherical 45° steering angle, the reduction in the needed element count is 13.4% when using a hexagonal element layout.

Relationships for designing the hexagonal planar (all elements residing on a single plane) sensor for far field operation and specific directivity and side lobe level are presented in [50]. This approach has since found a wide uptake in radar design [51].

For miniature ultrasonic arrays, the topic requires special consideration. Radar arrays are typically used in far field mode, while ultrasonic arrays are used both in far field, transition, and near field imaging modes. Additionally, in contrast to military radar applications, in ultrasonic inspection scenarios the probe is designed for a specific application. As of 2014, the cost of the electronic excitation and data acquisition system is much higher than the probe itself, and mainly dependent on the number of simultaneously active channels. Since it is generally true that increasing the channel count improves the resulting image quality, this comes at the cost of using a more expensive phased array controller. Therefore typically, it is a requirement to use a minimal count of channels to achieve a specific image quality.

In this Chapter the primary intention is to maximize the active element area, thus maximizing the probe sensitivity whilst keeping the image quality consistent with the conventional $\lambda/2$ element spacing requirement.

3.5 Comparison of the packing of squares and hexagons

As a first approximation to the benefit of hexagonal rather than square elements, a comparison of the area of the square and hexagonal elements embedded in a regular grid is considered.

Placing the element centres at $\lambda/2$ spacing fulfils the requirement of dense sampling. At the same time, the requirement of a high sensitivity of the probe dictates that the elements should be as big as possible. This implies placing them next to each other to fill the available space. This is known in mathematics as the problem of sphere packing or optimal space filling. Fig. 3.2 depicts properties of square packing (left) versus optimal packing (right). The same concepts can be used for packing of regular shapes like squares and hexagons as shown in Fig. 3.3.

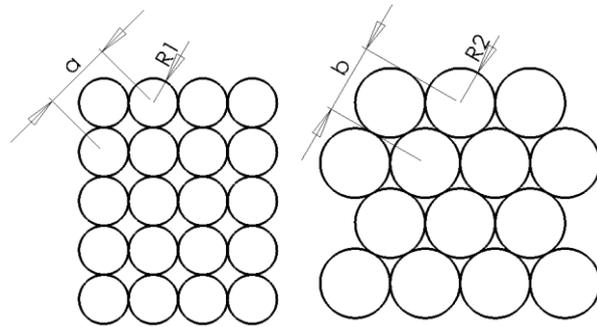


Fig. 3.2 Optimal sphere packing (2D case). Left: square packing. Right: Hexagonal (optimal) packing.

If we select such a, b , that $a=b=\lambda/2$, then $R2>R1$. The same applies for packing corresponding rectangles and hexagons, which fully cover the available area.

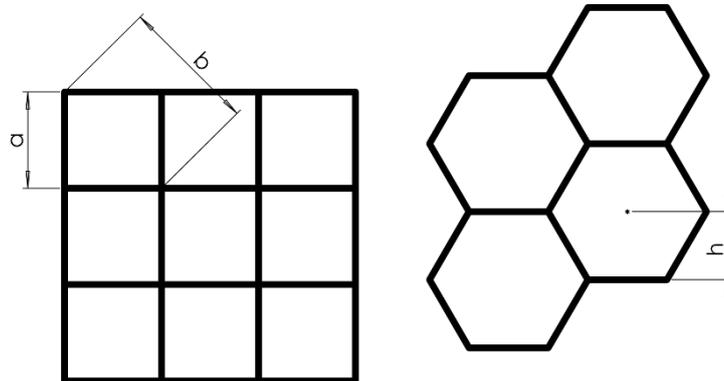


Fig. 3.3. A space-filling placement of squares and hexagons.

Fig. 3.4 shows the symbols used in the following discussion about rectangle and hexagon properties.

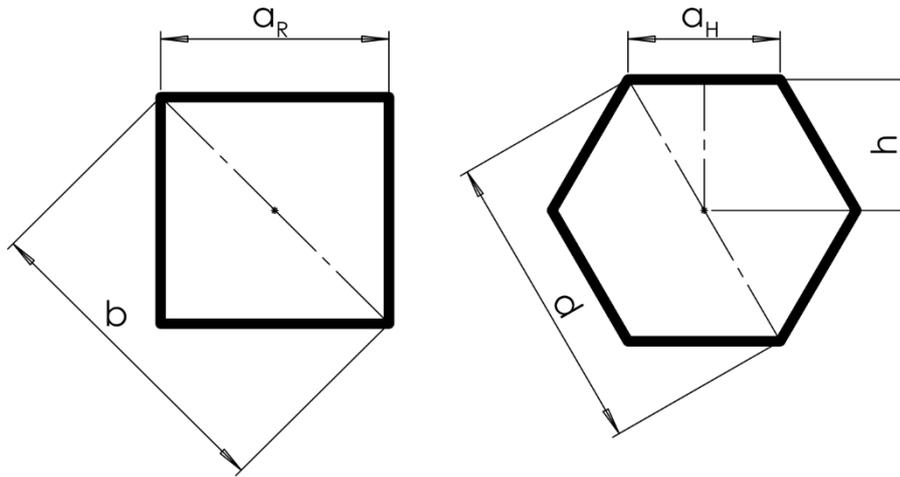


Fig. 3.4. Rectangle and hexagon dimensioning symbols

It is understood that the grating lobe (spatial aliasing) can appear if the array is not sampled densely enough in any given direction - this is known as Nyquist-Shannon sampling theorem. Therefore the question becomes: assuming the distance between the elements in the matrix is equal between hexagonal and square packing, which packing scheme provides an array footprint, or aperture, of a bigger area?

When comparing the size of the rectangle and hexagon elements in a dense packing situation, the question becomes: what constitutes the distance between the two elements of the grid?

As a first approximate, one can take the circumscribed circle radius of the element as the least dense direction, therefore assuming that the effective distance between neighbouring elements is the corner-neighbour distance. In such a case, derivation of the ratio is as follows.

Area of an equilateral triangle in function of the edge length is given by

$$A_{triangle} = \frac{\sqrt{3}}{4} a_H^2 \quad \text{Equation. 3.1}$$

Noticing that the diameter of the hexagon $d = 2a_H$, area of the hexagon created by 6 of such triangles is

$$A_{Hexagon} = \frac{6}{4} \sqrt{3} a_H^2 = \frac{3}{8} \sqrt{3} d^2 \quad \text{Equation. 3.2}$$

Now take the height of the triangle instead of the edge length as an input to get

$$h = \frac{\sqrt{3}}{2} * a_H \quad \text{Equation. 3.3}$$

$$A_{triangle} = \frac{\sqrt{3}}{3} h^2 \quad \text{Equation. 3.4}$$

The hexagon area is given by an area of 6 equilateral triangles and of specific interest is the area as a function of distance between the two hexagon centres that are connected by edge a , and exhibit radius $d=2a_H$, so that

$$h \rightarrow 1/2 * d \quad \text{Equation. 3.5}$$

$$A_{hexagon} = 6A_{triangle} = \frac{3}{4}\sqrt{3}d^2 = \frac{3}{2}\sqrt{3}a_H^2 = \frac{6}{3}\sqrt{3}h^2 \quad \text{Equation. 3.6}$$

The area of the square can be given as a function of distance between edge-neighbours (a), or as a function of the distance between corner neighbours (b):

$$A_a^{square} = a^2 \quad \text{Equation. 3.7}$$

$$A_b^{square} = \frac{1}{2}b^2 \quad \text{Equation. 3.8}$$

Using Equation. 3.6 and Equation. 3.8, one can obtain a ratio of areas assuming that the element to element distance q is $b = d$:

$$\frac{A_{hexagon}}{A_b^{square}} = \frac{\frac{3}{8}\sqrt{3}q^2}{\frac{1}{2}q^2} = \frac{6}{8}\sqrt{3} \cong 1.299 \quad \text{Equation. 3.9}$$

Given such distance between neighbouring polygons, the area of the hexagonal shape is larger. Hence, it is clear that with hexagonally-packed elements it is possible to increase the active area of each element while maintaining the $\lambda/2$ inter-element spacing. Fig. 3.5 shows a visual comparison of the areas of these respective shapes.

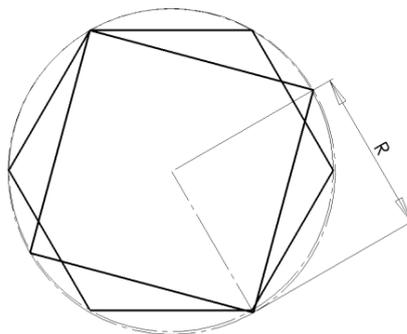


Fig. 3.5. Comparison of area of hexagon and rectangle of the same major radius R . Note the large patches of hexagon that extend beyond the perimeter of the rectangle.

3.6 Acceptance angle of a hexagonal array element

3.6.1 Rationale

As noted before, the sensitivity of a given transducer element depends on, among other factors, its size, and therefore its effective area. The bigger the area, the more energy can be transmitted and subsequently captured from the medium. However, one cannot simply have the elements larger, because the impinging signal spatial distribution is integrated over the area of the sensor. Large sensor size causes smearing of the signal phase and amplitude. A destructive interference occurs, resulting in an effect known as the “acceptance angle”. An acceptance angle is an angle, measured from the Z axis of the element, at which the sensitivity drops by a specified amount. When a wave front arrives at an angle to the element, the instantaneous amplitude at the element’s surface is integrated, and therefore diminishes with the increased element area, up to the point where grating lobes appear.

Here it is argued that the hexagonal shape of the element for a 2D phased array offers superior performance when compared against a square element. This is due to its projected shape extent being more uniform than for a square, as seen from a variety of possible skew angles.

A set of theoretical derivations and numerical simulations follow to prove this point.

3.6.2 Equations relating acoustic field and radiating aperture

Following the equation – Rayleigh integral - is for the time domain spatial impulse response of an array element, as described in [52] and illustrated in Fig. 3.6:

$$p(\vec{r}_0, t) = \int_S \frac{\delta\left(t - \frac{|\vec{r}|}{v}\right)}{2\pi|\vec{r}|} dS \quad \text{Equation. 3.10}$$

Where p - impulse response in space at point r ; \vec{r} – is the vector of distance between a point in space and a point on the surface of the array element; \vec{r}_0 – vector to the centre point of the radiating element; t -time; v -velocity of respective wave mode; S – surface area; δ – Dirac impulse function.

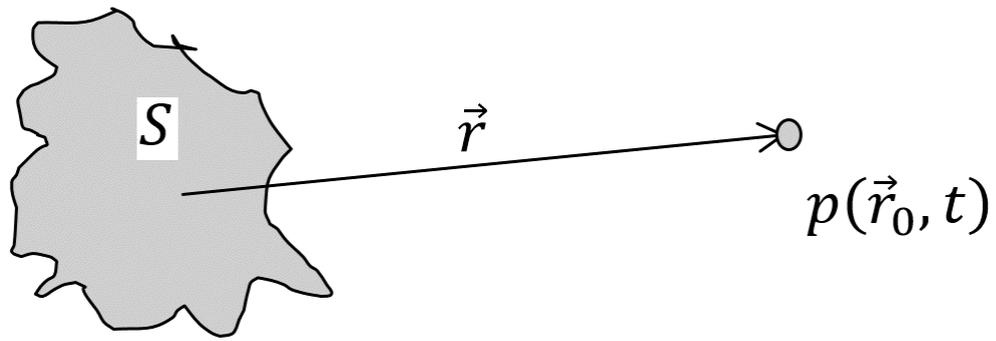


Fig. 3.6. Illustration of Rayleigh integral

For continuous wave of frequency ω , assuming surface velocity v and wavenumber k , the same can be written as [53]:

$$h(\vec{r}_0, \omega) = \frac{j\omega p_0}{2\Pi} \int_S v(\vec{r}, \omega) \frac{1}{|\vec{r}|} e^{jk|\vec{r}|} dS \quad \text{Equation. 3.11}$$

In order to calculate the acoustic field distribution several approaches can be used. One approach is to use closed form solutions for triangular apertures, as presented in [54]. An even more general approach is to use numerical integration. Here we use the numerical approach for its simplicity of implementation and wide applicability.

The integration limits for hexagonal and square apertures are as follows:

For a hexagonal element of inner radius a , one can use the composite integration limits as shown in Fig. 3.7 and Equation. 3.12 to Equation. 3.16:

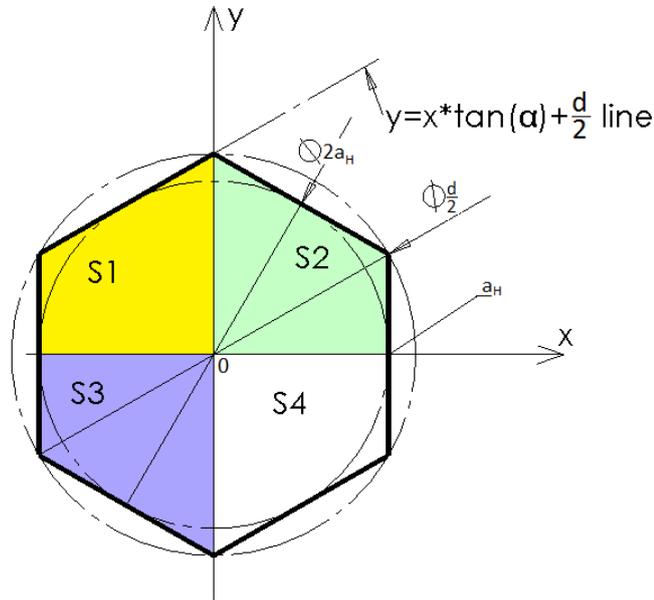


Fig. 3.7. Integration limits for hexagonal element

$$S = S1 + S2 + S3 + S4;$$

Equation. 3.12

$$S1 = \left\{ x: \langle -a_H; 0 \rangle; y: \langle 0; x \cdot \tan(\alpha) + \frac{d}{2} \rangle \right\};$$

Equation. 3.13

$$S2 = \left\{ x: \langle 0; a_H \rangle; y: \langle 0; -x \cdot \tan(\alpha) + \frac{d}{2} \rangle \right\};$$

Equation. 3.14

$$S3 = \left\{ x: \langle -a_H; 0 \rangle; y: \langle 0; -x \cdot \tan(\alpha) - \frac{d}{2} \rangle \right\};$$

Equation. 3.15

$$S4 = \left\{ x: \langle 0; a_H \rangle; y: \langle 0; x \cdot \tan(\alpha) - \frac{d}{2} \rangle \right\};$$

Equation. 3.16

Where $\alpha = \frac{\pi}{6}; \frac{d}{2}$ is the radius of the circumscribed circle of the hexagon (major radius),

and $a_H = \frac{d}{2} * \cos \alpha$, is the radius of the circle inscribed in the hexagon, and 0 is the origin of the cartesian coordinate system.

Assuming a rectangle with edge length a_R , the limits of integration for a square element are:

$$S = \left\{ x: \left\langle -\frac{a_R}{2}; \frac{a_R}{2} \right\rangle; y: \left\langle -\frac{a_R}{2}; \frac{a_R}{2} \right\rangle \right\};$$

Equation. 3.17

3.6.3 Simulation 1 – constant aperture; acceptance angle versus wavenumber

3.6.3.1 Description of the model

In order to demonstrate the difference between the acceptance angle of square and hexagonal aperture elements the following simulation is proposed.

The calculation is conducted in arbitrary distance and velocity units; this makes the results transferrable to any scale associated with a real scenario.

Assume two elements, hexagonal and square, with an equal area, $A_H=A_R=1$. Their respective dimensions will be

$$a_R = \sqrt{A_R} = 1$$

Equation. 3.18

$$a_H = \sqrt{A_H} * \sqrt{\frac{2}{3\sqrt{3}}} \approx 0.6204$$

Equation. 3.19

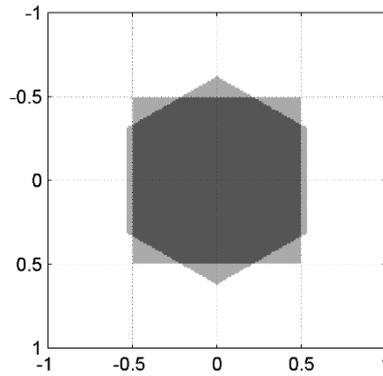


Fig. 3.8. A comparison of rectangle and hexagon shapes of an equal area. Units are arbitrary distance units.

As depicted in Fig. 3.8, elements are assumed to lie on a $Z=0$ plane, with their centres at Point $(0,0,0)$; wave velocity in the medium is 1. The simulation is conducted with a wavenumber as a parameter and the range of wavenumbers used is between 1 and 15. The radiating aperture operates as a piston with a hard boundary, and the amplitude of the field is probed at a distance of $r=10$ from the centre of the element, as depicted by the blue crosses in the Fig. 3.9 (note, in the figure the field points are at $r=2.5$ for clarity). The acoustic field intensity at the probing points is integrated numerically with machine accuracy (using double precision floating point numbers). Area-integration of smooth linear problems are known to be numerically stable. Although exact, closed-form solutions for this problem are available in the literature, their implementation is outside of the scope of this work.

The field is probed at (x,y,z) points forming a quarter-circle around the radiating element, with an indexing angle α (first rotation around y -axis), and then this quarter circle is rotated around z -axis by angle β to obtain the view of the element as seen by a wave front approaching from different angles. Point locations are given by:

$$x = r * \sin(\alpha) * \cos(\beta) ;$$

$$y = r * \sin(\alpha) * \sin(\beta)$$

$$z = r * \cos(\alpha)$$

Equation. 3.20

example scene, $\beta=0$, $r=2.5$

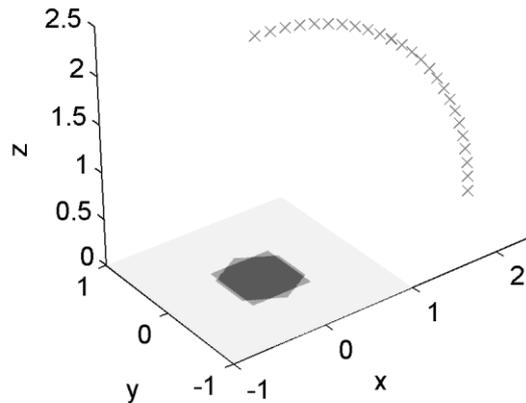


Fig. 3.9. The location of the radiating element and the acoustic field probe points. Crosses depict the points at which the field pressure is calculated.

The calculated field is normalised against an amplitude at $\alpha=0$. Normalisation is used because between the hexagonal and square aperture, despite their equal area, there exists different **average distance** that an example sound ray has to travel from the point on the element and the probe point in space. Therefore peak amplitude and phase of the field generated by hexagonal element and square element are not **exactly** the same. However, the difference is negligible for practical purposes (in range of 1ppb or -180dB).

An example amplitude recorded for a given wavenumber is presented in Fig. 3.10. Note that in this particular case, a high wavenumber of 8.8 was selected to show the location of the first null at $\alpha \approx 0.75$ [rad]

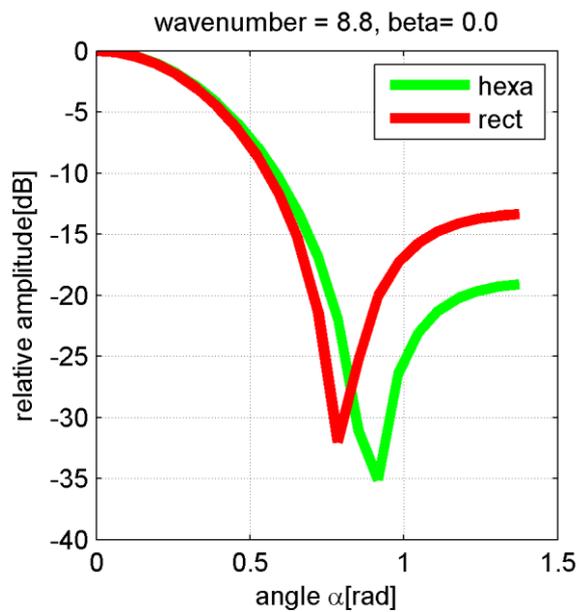


Fig. 3.10. Example field amplitude at $\beta=0$, area of 1 and at significantly high wavenumber of 8.8 to illustrate the difference in acceptance angle and the appearance of the null in sensitivity.

3.6.3.2 Simulation 1: Results

The acceptance angle is calculated as the angle at which the relative amplitude crosses a -1dB point. This value varies depending on the approach direction β , as depicted in Fig. 3.11 and Fig. 3.12. From these figures it is clear that (assuming elements of an equal area) the acceptance angle is higher for an element of hexagonal shape. Secondly, the acceptance angle is uniform with respect to the angle of incidence (variation of less than 0.1 milliradian), whereas for a square shape it exhibits a dependence on the angle of incidence (variation of ~ 10 milliradians)

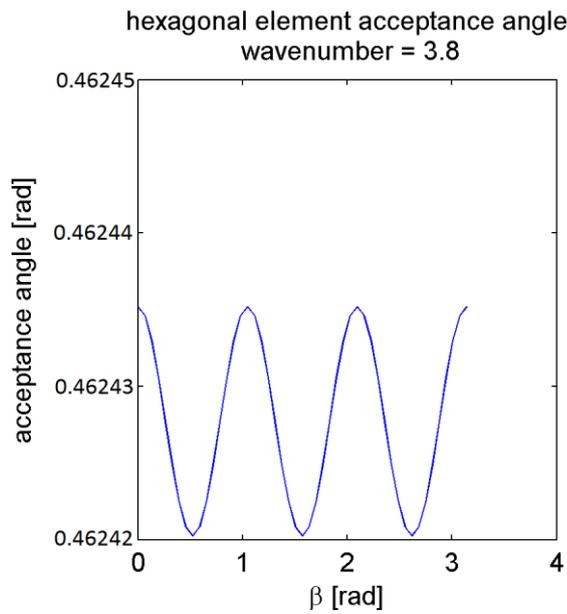


Fig. 3.11. Acceptance angle of hexagonal element

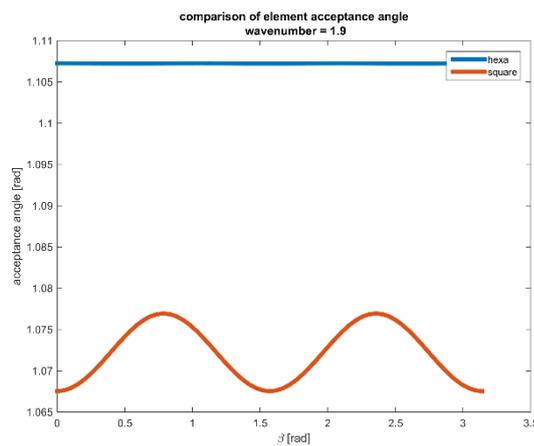


Fig. 3.12. A comparison of the acceptance angle for hexagonal and square element at the same wavenumber, and at the same plot scale.

Finally, comparison of the -1dB and -3dB acceptance angle versus wavenumber is calculated and presented in Fig. 3.13. Note that the -3dB acceptance angle could not be calculated for wavenumbers below 4, and -1dB could not be calculated for wavenumbers below 2. This is because as the wavenumber decreases, the wavelength increases, and the wavelength eventually becomes comparable with the size of the element. In such cases, the

element becomes omnidirectional and the radiated amplitude does not drop below -1dB for any angle.

The observed difference in the acceptance angle between square and hexagonal elements of the same active area is very small. However, the significance of this result is that the area of the hexagonal element can be made larger than the square element whilst keeping the acceptance angle at a specified level.

By the same token, for a big hexagonal element, the location of the first null is further away from the axis than for the square element (Fig. 3.10), which means that the ultimate limit of the element size is larger.

The next simulation rearranges the calculations to find out how much bigger the area of the hexagonal element can be to keep a given acceptance angle.

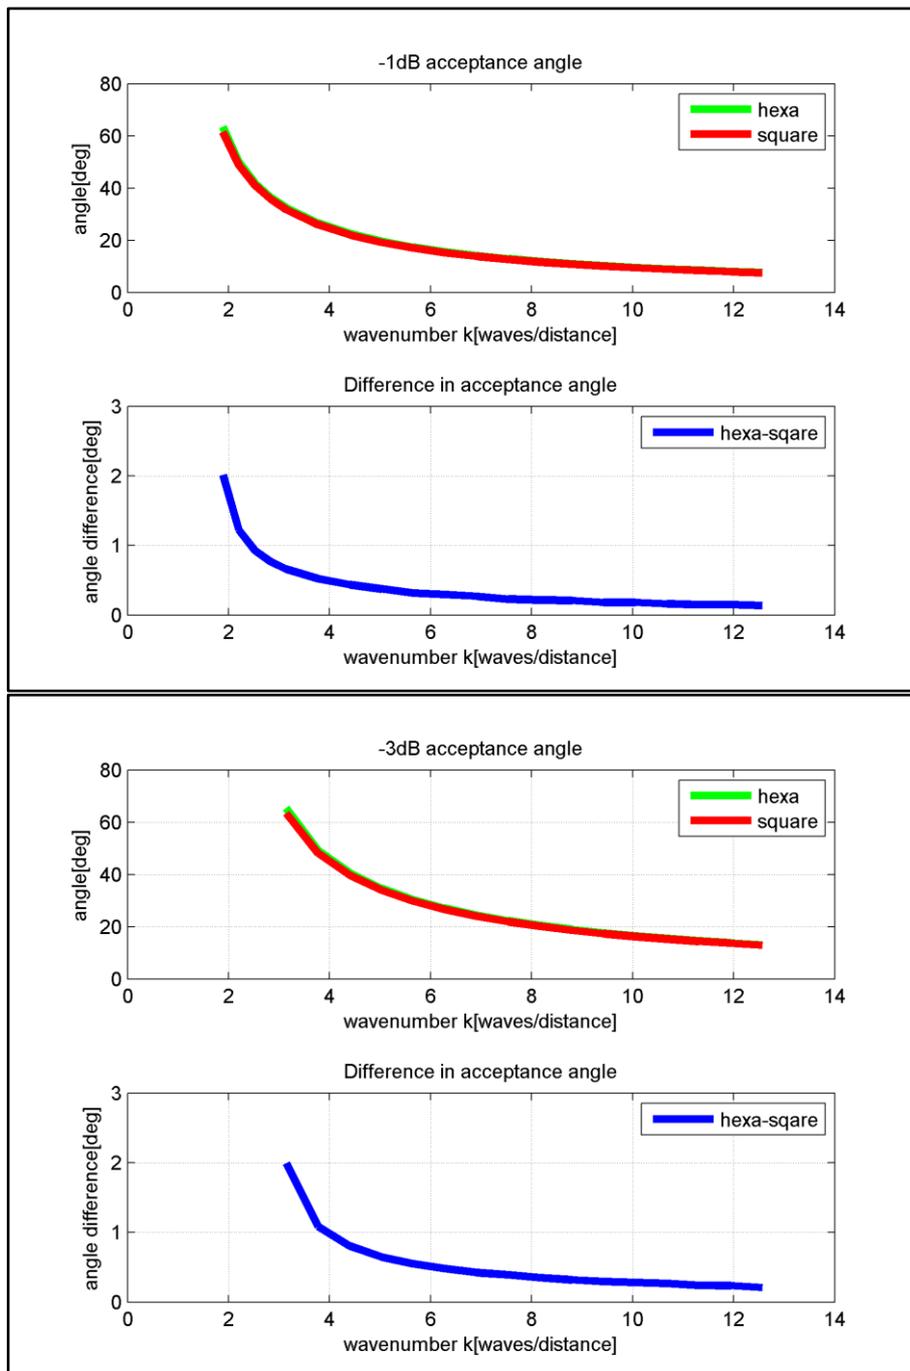


Fig. 3.13. Comparison of acceptance angle for hexagonal and square element versus wavenumber. The difference is bigger for wavelengths close to the size of the element (which are of most interest)

3.6.4 Simulation 2 – constant acceptance angle

3.6.4.1 *Description of the Problem*

The method from simulation 1 can be inverted to answer the following question: what area can be used of a given shape to obtain a required acceptance angle?

The purpose of the following calculations is to find the area of the hexagonal and square elements that fill the following assumptions: (a) constant wavenumber of 1, and (b) acceptance of -1dB, -2dB, -3dB and -6 dB sought to be placed at angles of 15, 30, 45 and 60degrees off axis. Thus it will be possible to compare the theoretical pulse-echo response amplitudes to assess the array element performance.

If the area for a hexagon is larger than for a rectangle, it is proposed that the pulse-echo sensitivity will be bigger by at least the square of the area ratio. This neglects the effect of improved electrical impedance matching.

3.6.4.2 *Simulation 2: Results*

An example result is presented in Fig. 3.14. Note that the calculated area changes depending on the beam skew angle. Obviously, practical elements can only have one, fixed dimension. Therefore, the minimal area (the amplitude is 'at least' as big for any skew angle) has been used to quantify the difference between the two element shapes. A set of results for various amplitude drops and acceptance angles is presented in Table 3.1. The difference between the pulse-echo amplitudes obtained by array elements of different shapes varies slightly at around 10%, and generally increases slightly if the required element acceptance angle is wide.

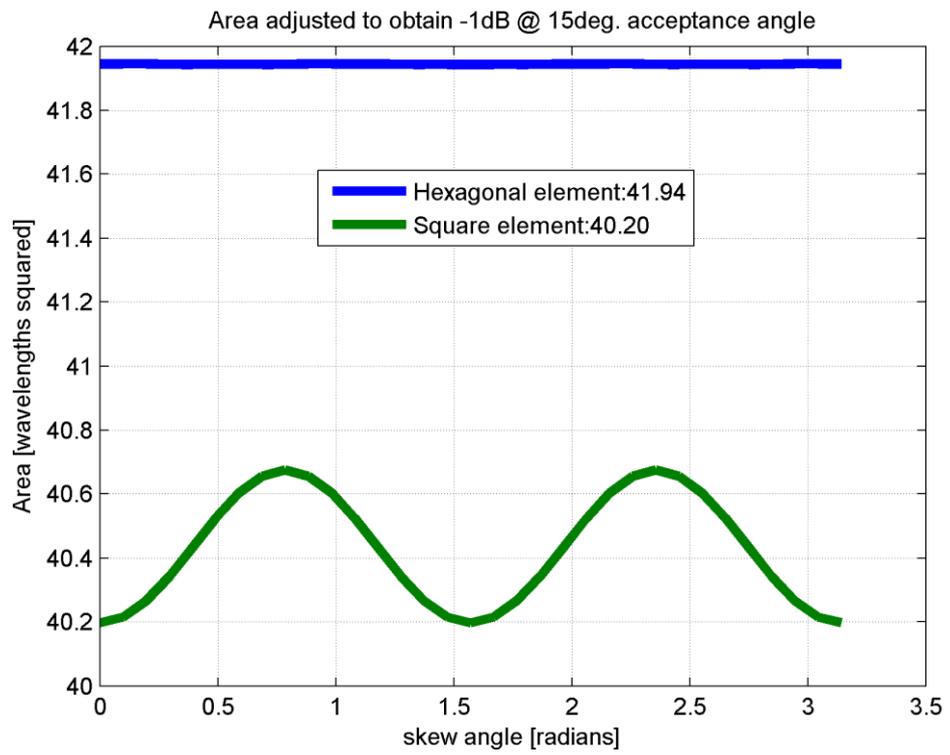


Fig. 3.14. Area of the element adjusted to obtain desired acceptance angle at desired amplitude drop

Table 3.1. Pulse-echo amplitude difference for various amplitude drops at various acceptance angles.

Acceptance amplitude	Acceptance angle	Area for hexagon	Area for rectangle	Pulse-echo amplitude change
dB	degrees	Wavelengths squared	Wavelengths squared	%
-1	15	41.9	40.2	8.84%
	30	11.3	10.8	8.88%
	45	5.64	5.4	9.09%
	60	3.76	3.6	9.09%
-2	15	82.6	78.8	9.72%
	30	22.1	21.1	9.80%
	45	11.1	10.5	9.90%
	60	7.37	7.03	9.91%
-3	15	122	115	10.5%
	30	32.5	30.9	10.6%
	45	16.3	15.4	10.7%
	60	10.9	10.3	11.5%
-6	15	228	214	12.8%
	30	61.2	57.4	13.4%
	45	30.4	28.6	13.5%
	60	20.4	19.1	13.5%

3.7 Probe aperture overall shape and size considerations

To make the beam forming side lobes as regular as possible and equal in all directions, a circular overall aperture is recommended. Oval or rectangular apertures are sometimes used if enhanced resolution is needed in one axis at expense of the resolution in the other axis, or there exist special requirements as to side lobe positioning.

In this Section, the apertures made up with square and hexagonal elements respectively are compared, assuming there is no beam steering.

For the purpose of this work, probes with circular symmetric apertures and up to 128 channels are considered. An example layout of array elements is depicted in Fig. 3.15. In

order to compare the performance of these two configurations, two simulations are conducted.

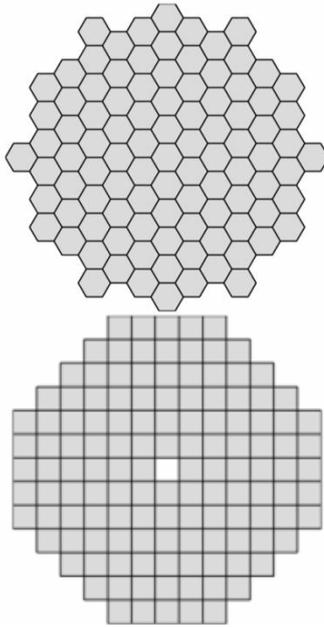


Fig. 3.15. Dense array layouts with circular overall aperture. Left: Example 98-element layout of an array with hexagonal elements. Right: 128-element layout of an array with square elements. The layouts have been selected to be nearly equal in aperture, equalizing focal spot size of the generated acoustic beam, while keeping the $\lambda/2$ sampling.

3.7.1 Lambert azimuthal equal-area projection as a space sampling method

From [55] and [56, p. 218], “The Lambert azimuthal equal-area projection is a particular mapping from a sphere to a disk (that is, a flat region bounded by a circle). It accurately represents area in all regions of the sphere, but it does not accurately represent angles”.

The application of this transform in ultrasonic array research is that it enables to create an easy to process, regular grid of points, with spatial coordinates distributed on a hemisphere. If the simulated ultrasonic transducer is then placed in the centre of the hemisphere, and the acoustic field generated by it sampled at these points, the sampled values can be then analysed using a wide range of methods and algorithms applicable to regular point grids. Knowing that each point represents an area equal to the other points,

this simplifies the problem of visualization and drawing statistical conclusions about the sampled field.

For example, the flux of the main beam can be integrated, compared to the total flux of side lobes.

Storing calculation results in a regular grid of points is beneficial from the point of view of the computational performance. cueBeam uses this principle to achieve very high calculation throughput.

Fig. 3.16 and Fig. 3.17 depict the typical (virtual) scene setting inside cueBeam. The green shapes depict location of the radiating elements; the red crosses is where the acoustic field is sampled – calculated. Although the physics calculation is fully 3D, the result can be stored in an efficient manner in a square 2D matrix. Fig. 3.16 shows planar distribution of sampling points. Fig. 3.17 shows spherical distribution of sampling points. The spherical distribution allows sampling of the field that is at an approximately equal distance from the radiator centre, and secondly, it allows taking into the account all the side lobes that would not permeate the planar sampling area from Figure 1.

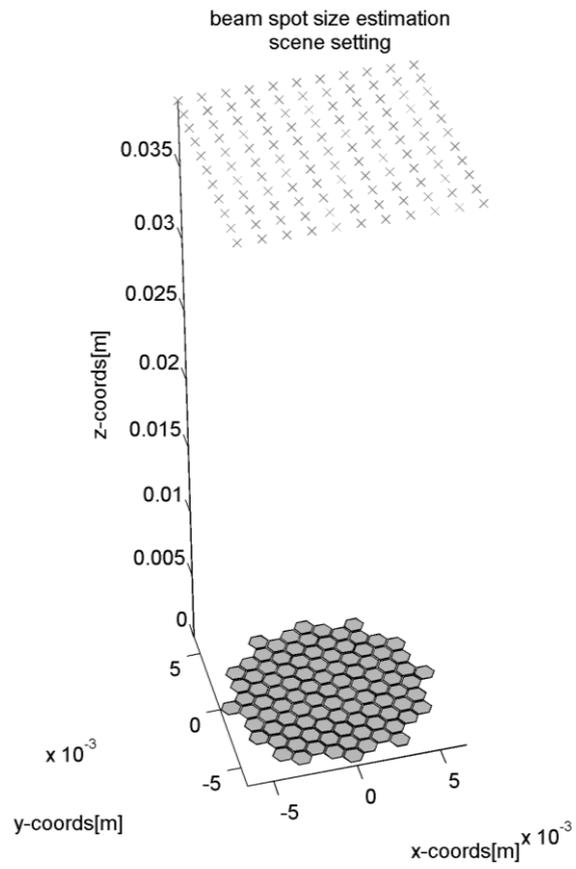


Fig. 3.16 Scene setting for cueBEAM. Hexagons: element locations (actually, points); crosses: "probe points" where field is calculated.

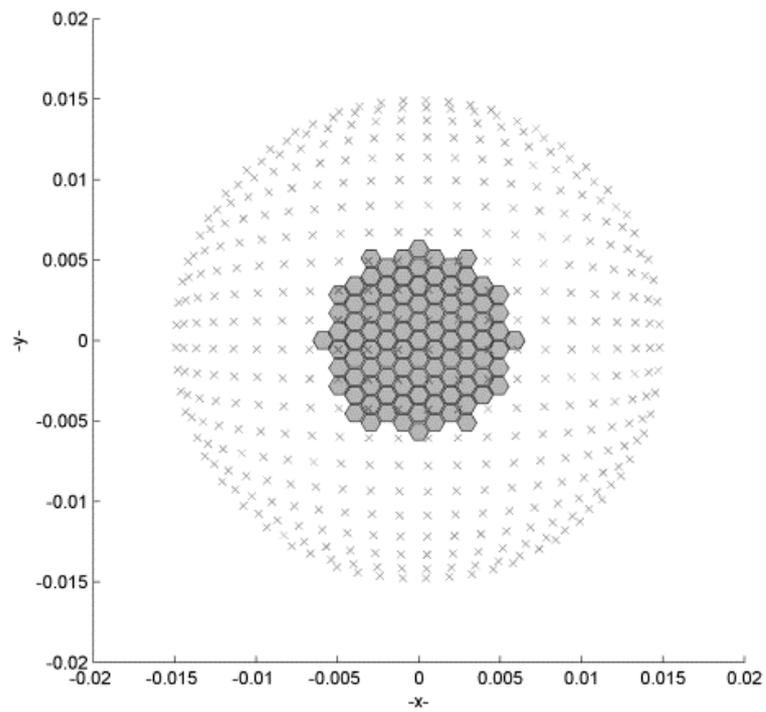
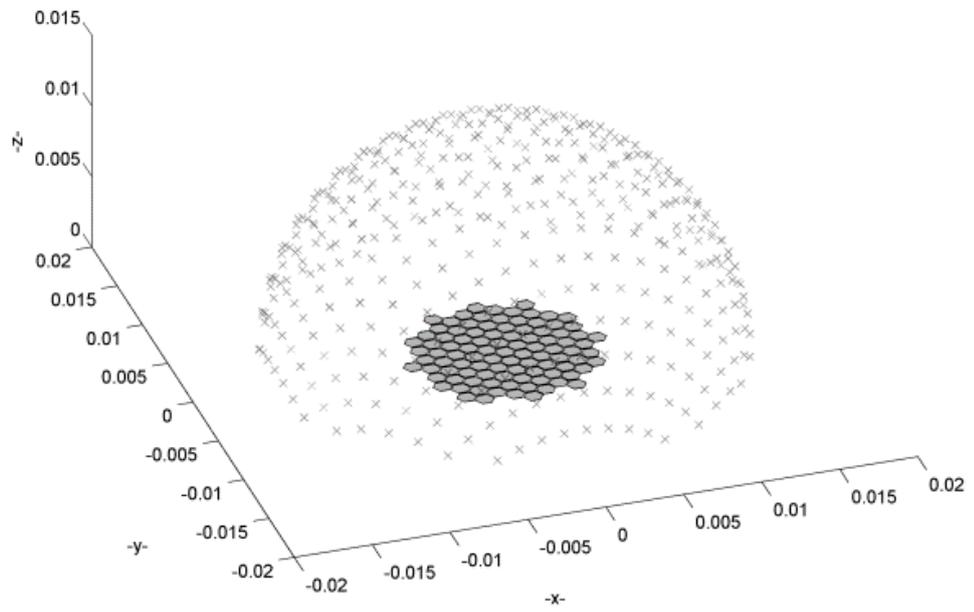


Fig. 3.17 Lambert azimuthal equiareal map field probing point distribution. Left: perspective view. Right: top-down view.

The forward transformation – from spherical polar coordinates to 3D Cartesian coordinates are:

$$\begin{aligned}x &= k' \cos \phi \sin \lambda - \lambda_0 \\y &= k' (\cos \phi_0 \sin \phi - \sin \phi_0 \cos \phi \cos(\lambda - \lambda_0))\end{aligned}$$

where

Equation. 3.21

$$k' = \sqrt{\frac{2}{1 + \sin \phi_0 \sin \phi + \cos \phi_0 \cos \phi \cos(\lambda - \lambda_0)}}$$

Where ϕ, λ are zenith and azimuth on the sphere, and ϕ_0, λ_0 is the coordinate of the centre of the projection disk.

The Z-coordinate follows the surface of the sphere. Given Cartesian X,Y coordinates on the disk, this inverse transformation rules are used to calculate standard parallel ϕ_1 and central longitude on the sphere:

$$\rho = \sqrt{X^2 + Y^2}$$

$$c = 2 \arcsin\left(\frac{1}{2} \rho\right)$$

$$\phi = \arcsin\left(\cos c \sin \phi_0 + \frac{Y \sin c \cos \phi_0}{\rho}\right)$$

Equation. 3.22

$$\lambda = \lambda_0 + \arcsin \tan\left(\sin \phi_0 + \frac{X \sin c}{\rho \cos \phi_0 \cos c - Y \sin \phi_0 \sin c}\right)$$

Of more interest are the plane-to-sphere Cartesian coordinate transforms. For (x,y,z) coordinates on the sphere, and a regular grid of (X,Y) coordinates on the plane, one can obtain:

$$(x, y, z) = \left(\sqrt{1 - \frac{X^2 + Y^2}{4}} X, \sqrt{1 - \frac{X^2 + Y^2}{4}} Y, -1 + \frac{X^2 + Y^2}{2} \right)$$

Equation. 3.23

And the inverse transform is given by

$$(X, Y) = \left(\sqrt{\frac{2}{1-z}} x, \sqrt{\frac{2}{1-z}} y \right) \quad \text{Equation. 3.24}$$

The above assumes that the transformation is centred around $(x,y,z) = (0,0,-1)$.

Details of the implementation and source code can be found in the documentation of the cueBeam software.

3.7.2 Simulation 1. Comparison of probes with equal element count and different probe apertures

In this simulation, the acoustic field distribution generated by two phased array apertures is made: one made with square elements, and the other with hexagonal elements. The qualities of interest are main beam spot size and peak side lobe level. The assumptions are that both of the candidates have an equal number of elements, and having hexagonal elements of such a size that their pitch matches the pitch of the square elements configuration.

A visual comparison of the two apertures that satisfy those conditions is given in Fig. 3.18. The hexagonal element array overall aperture is larger.

For the purpose of this evaluation, Lambert azimuthal equal area mapping¹ is used to generate a set of points where the acoustic field amplitude is evaluated, as depicted in Fig. 3.17

The beam profile is taken using the same model code base as in previous Sections. Visualization of the respective beam profiles is then provided in Fig. 3.19. It is apparent that for the hexagonal element array, the side lobe energy is distributed more evenly around the outer rim of the main lobe. It is this more even energy distribution that makes the peak side lobe level lower than for a square element array.

Finally, the numerical results are gathered in Table 3.2.

¹See section 3.7.1, page 89, for rationale and governing transformations

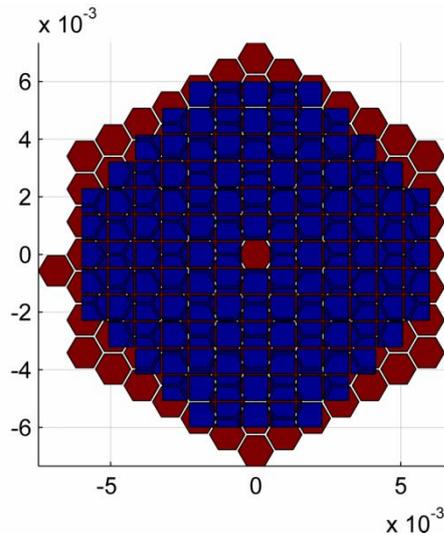


Fig. 3.18. A hexagonal-element probe overlaid with a square element probe of the same pitch and element count.

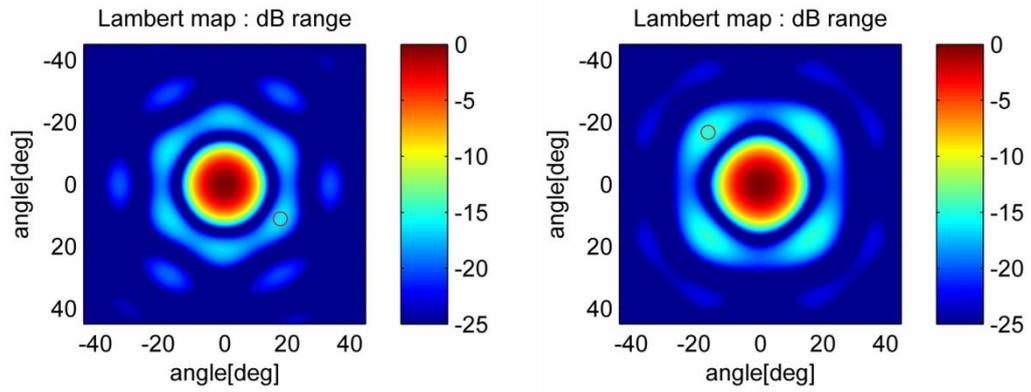


Fig. 3.19 Comparison of cross-sections of beam shapes and side lobes of 128-element equal-pitch probes. Left: Hexagonal element probe. Right: square element probe. Small red circle depicts location of peak in side lobe.

Table 3.2. Results of acoustic beam simulation.

Parameter	Hexagonal element probe	Square element probe	Unit
Number of elements	128	128	-
Element pitch	0.6556	0.6556	mm
Area of element	1.1165	0.85951	mm ²
Total aperture of probe	142.9	110.0	mm ²
-3dB spot diameter at 50mm	22.3	25.2	mm
Peak side lobe, single-way propagation	-16.2	-14.8	dB

In this comparison, the -3dB spot diameter (directly corresponding to effective image resolution) was predictably better for the probe with the larger aperture. Importantly, the Nyquist-Shannon sampling criterion was upheld for the hexagonal-element array, the result being the peak side lobe amplitude (image contrast) was not compromised.

3.7.3 Simulation 2. Comparison of probes with different element count and similar apertures

In this supplementary simulation a comparison of the focal spot size and side lobe amplitude is undertaken assuming that the overall probe apertures are equal. The count of hexagonal elements is reduced to make the hexagonal element aperture of approximately the same size as the square element aperture. A simple calculation shows that in order to satisfy this requirement, the hexagonal-element array with elements as in previous sections, should have 98 elements.

The acoustic field generated by a 98-element hexagonal-element probe configuration was compared with a field generated by 128-element square-element probe. The results have been collated in Table 3.3. The hexagonal-element array still offers better image contrast - lower peak sidelobe, by 2.11dB in single-way propagation. Pertinently, this effect is doubled

in pulse-echo operation. This is achieved even though there are fewer active phased elements in the hexagonal-element probe. This shows that the wave is properly sampled over the aperture of the probe, since the centre-to-centre distance of the elements is still less than half of the wavelength of interest.

Even though the aperture of the hexagonal element is larger than the aperture of the square element, the effect of its directivity is negligible for skew angles less than approx. 70 degrees off axis. On the other hand, the positive aspect is that the larger aperture ensures a higher sensitivity of the element, as more energy can be transduced both in transmission and in reception.

Taking another view, the pitch of the hexagonal elements could be made larger than with square elements, while keeping the side lobe level equal to the reference square-element design. This demonstrates that the hexagonal element arrangement is superior and the hexagonal-shaped elements can be used to fill an available aperture.

Table 3.3. Results of an acoustic beam simulation.

Parameter	Hexagonal element probe	Square element probe	Unit
Number of elements	98	128	-
Element pitch	0.65	0.65	mm
Aperture of element	1.11	0.859	mm ²
Total area of probe	109.4	110.0	mm ²
-3dB spot diameter at 50mm	25.4	25.2	mm
Peak side lobe, single-way propagation	-16.9	-14.8	dB

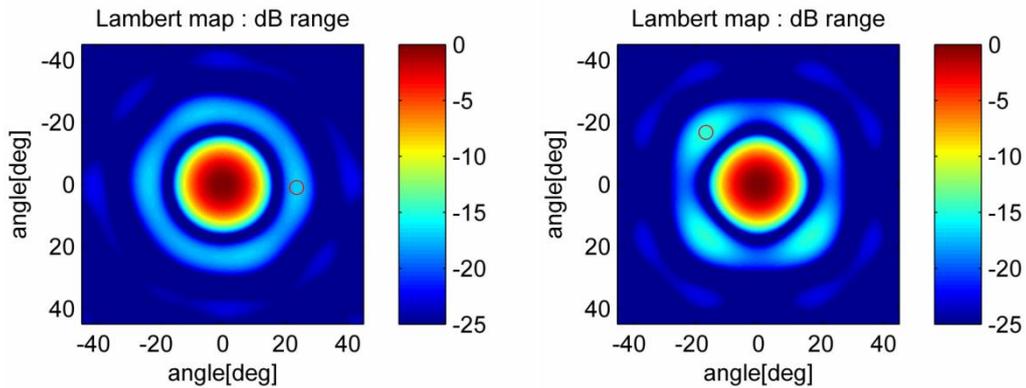


Fig. 3.20. Cross-section of beam shape generated by respective probe layouts. Left: Hexagonal-element probe. Right: square-element probe. The small red circle shows location of the peak side lobe.

3.8 Summary

The purpose of this Chapter is to compare hexagon and square geometries as shapes suitable for building densely sampled phased array ultrasonic probes.

It has been shown that hexagonal shaped ultrasonic elements possess some beneficial properties over square shaped elements.

The radiation pattern of a hexagonal shape is more uniform; therefore it can possess a bigger area whilst preserving the acceptance angle of choice.

If we disregard the acceptance angle, an aperture utilizing hexagonal shaped elements exhibits better resolution and contrast in case of either overall aperture or the element count being constrained. However, the acceptance angle of the element appears to be the limiting factor in both cases; if we consider the acceptance angle of -6dB at 60° off-axis (single-way radiation) then the area of the hexagon can only be bigger by 7% compared to a square configuration, theoretically improving the pulse-echo sensitivity by 13%.

In order to fully exploit the benefits of hexagonally shaped elements a suitable substrate material is needed; this problem is addressed in Chapter 4.

Chapter 4. Triangular-Cut Piezoelectric Composite

4.1 Introduction

In chapter 3, an argument has been made to replace classic, square array elements with hexagonal ones, as illustrated in Fig. 4.1. There remains a challenge of creating a suitably configured piezoelectric composite microstructure that will support the proposed array element layout. Here, a regular grid of equilateral triangle pillars is considered for the said purpose.

The triangular grid can be applied to create hexagonally shaped array elements as illustrated in Fig. 4.2. By collecting 6 triangular pillars one can create a hexagonal element. Other configurations are possible, for example collecting 13 pillars into an oval shape, or by using a strip of pillars to create a 1D array element.

Importantly, such a structure can be manufactured using a standard dice-and-fill method [57].

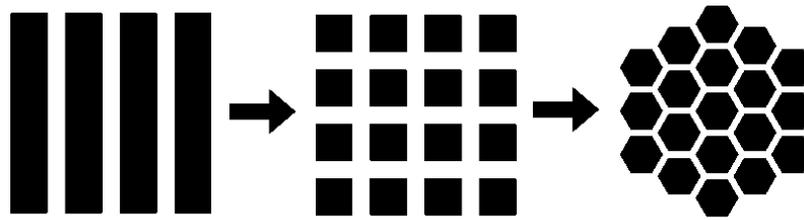


Fig. 4.1 Left: Layout of 1D Phased Array elements. Middle: 2D phased array layout created by subdividing 1D phased array elements. Right: Proposed 2D phased array layout with hexagonal elements

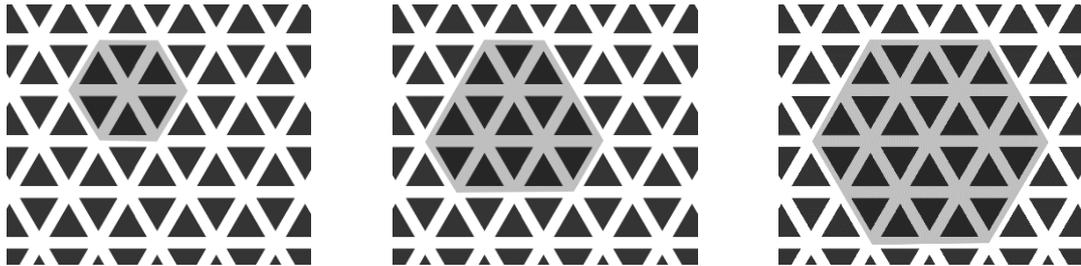


Fig. 4.2 Triangular pillar, 1-3 piezoelectric composite schematic, top-down view. Array elements are defined by placing an electrode pattern over selected group of pillars. Left: 6 pillars per electrode (shaded); middle: 13 pillars per electrode; right: 24 pillars per electrode. Other configurations are possible.

As discussed in Chapter 2, it has been reported that triangular pillar piezoelectric composite can perform as well or slightly better than the square pillar piezoelectric composite in some cases [28] , [29]. The difference appears to lie in designing the piezoelectric composite material to exhibit unimodal vibrational behaviour; that is, when excited with a broadband pulse, it mainly vibrates at the thickness mode frequency only, and no energy is wasted (i.e leaked) into other vibration modes. Importantly, secondary modes of vibration, if present, not only decrease electromechanical efficiency but also give rise to mechanical crosstalk through Lamb waves[47].

Here, a design method based on a parametric sweep is proposed to allow identification of designs, in which the resulting material behaviour is unimodal. Throughout this chapter, a variety of example results are presented.

4.2 Review of recent literature

In 2009, Brown et al. [57] found themselves needing to fabricate a high frequency piezoelectric composite for their single-element focussed probe. There is a known method

to design a square cut piezoelectric composite [58] however, they find that for the desired 40MHz operating frequency, it is un-manufacturable due to the extremely thin kerf required ($9\mu\text{m}$) to obtain the desired pillar width-to-height aspect ratio. Instead, they choose to use the $15\mu\text{m}$ kerf technology to manufacture a piezocomposite with additional cut across the square pillars to create triangular pillars. They found that this allows the piezocomposite to be significantly better behaved (shorter impulse response) than the one without the extra cut. The apparent lack of spurious resonances near to the fundamental resonant frequency of the material is explained by the additional cut modification of the periodic structure of the material.

In 2010, a publication was released on the study of triangular pillars [59] with cuts applied by different angles: 30, 45 and 60 degrees in search of a design with the shortest impulse response. In this work, the composite with 45° cut exhibits the shortest response and what can also be seen in the measurements, but has not been commented on by the authors, is that the 45° cut composite is the only one with an unimodal impedance profile. Good analysis of the fabrication results are given, but no systematic method to design the material towards this desirable property has been proposed.

A well researched 2012 paper [60] again pits regular square composite and 45deg. cut triangular pillar composite material against a novel off-periodic 45° pseudo-random pillar cut composite. The paper finds improved response of the latter and the total electromechanical crosstalk between array elements was -6.5dB which is better than for a square element array. Again, although not highlighted by the authors, it appears that it is the unimodal response of the material that implies the shortest impulse response time.

At this point, it is important to note that the undesirable long impulse response is also sometimes mentioned as existence of the 'secondary impulse' in the time domain trace of the received voltage. The 'secondary impulse' can be characterized as occurring after the primary impulse, being of lower amplitude, longer decay time and of a slightly different frequency. These characteristics are not difficult to link to a non-unimodal electrical impedance plot.

In the light of these recent new publications, the hypothesis of this chapter stands that a design method is needed that will reliably produce a material with unimodal behaviour.

4.3 Modelling the vibration modes of the material

4.3.1 Introduction

In order to explore the characteristics of the triangular cut piezoelectric composite, the PZFlex finite element package (Weidlinger Associates, Mountain View, CA) has been used.

Compared to the square pillar piezoelectric composite configuration, as described in chapter 2, equi-angle, regular, uniform (non-a-periodic) triangular pillar composite material (as needed for regular hexagonal element 2D phased arrays) can be characterized with just two parameters:

- Kerf width
- Kerf pitch

In this work, a piezoelectric composite material for operating at a frequency of around 3MHz is sought, which implies a starting material thickness of 450 μ m, to be adjusted as a function of other parameters.

For this work, the ceramic modelled is a PZT5H (Ferroperm, Kvistgaard, Denmark), and the polymer filler is Vantico CY221/HY956EN (Hunstman, Cambridge, UK). Although varying constituting material is outside of the scope of this study, the method itself is expected to be transferable to different materials.

The kerf width is governed by the technology of manufacture; at the time of this work, in the CUE, the most practical kerf making technology cut the kerf with a 100 μ m saw blade. For this work the automated MicroAce dicing saw was used (Loadpoint, UK).

Therefore, the only two variables to manipulate are kerf pitch and the thickness of the material. Fortunately, changing the kerf pitch also changes the material's Ceramic Volume Fraction (CVF). In the literature, CVF is often cited as the design parameter for particular piezoelectric composite configurations.

4.3.2 Model Visualization Run

A full 3D model of a slice of the piezoelectric composite ceramic material was created to facilitate this modelling study, as visualized in Fig. 4.3. The electrode boundary condition has been applied to the model in such way as to simulate a single hexagonal element electrically excited with a voltage pulse; the remaining electrodes were left at ground voltage level. Fig. 4.4 shows the initial state of the model, with the surface of the composite material at rest; the coloured areas represent the extent of the area that was electroded to provide an excitation pulse or measure an electrical response. The model is excited with very short electrical impulse² applied to the central electrode.

In this example run, a snapshot of the evolution of the finite element model is shown in at time $t=297\text{ns}$ after the beginning of the excitation pulse in Fig. 4.5. The central element has extended in an uniform manner, which is a desirable behaviour. However, the same model at time $t=242\text{ns}$ exhibits 'corona peaking' as illustrated in Fig. 4.6, where the apices of the triangles appear to vibrate out of phase from the kernel of the element; they in fact vibrate at a slightly different frequency (2.07MHz vs. 1.68MHz). The energy of this vibration is not converted into ultrasonic radiation in the same pattern as the first mode; instead, it will be dissipated into internal friction and re-energize the main mode over time, elongating impulse response.

The following modelling study seeks to identify unimodal designs where parasitic modes are well separated from the fundamental thickness-mode of the vibration.

² The model is a time domain, time quantized model; the excitation pulse length is equal to a single time step of the model progress. This ensures that all model-possible modes of vibrations are excited.

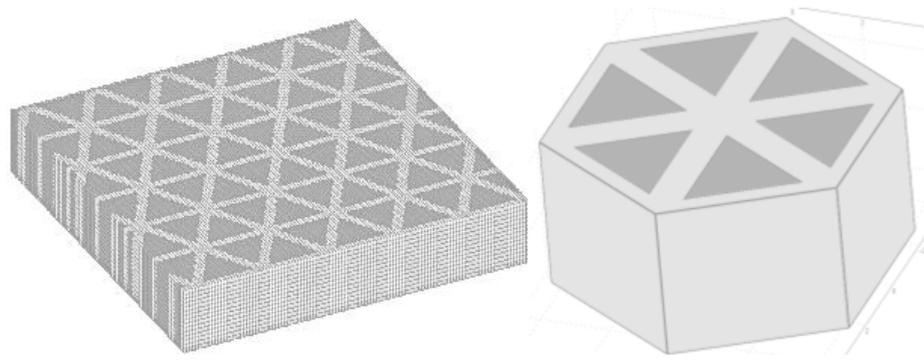


Fig. 4.3 Triangular-cut piezoelectric ceramic composite a) composite overview; b) single-element model. Dark grey – ceramic pillars; light grey - polymer.

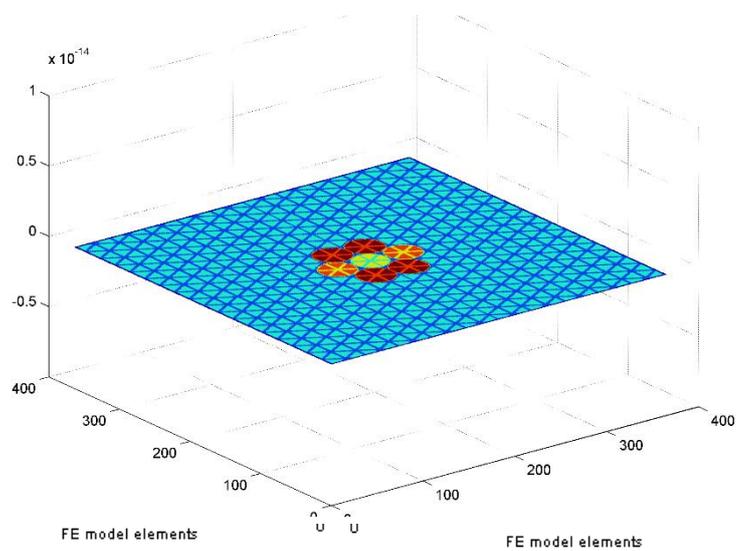


Fig. 4.4A visualization of the displacement of the surface of the model at $t=0$; light blue -ceramic pillars; dark blue – kerf (polymer filler); lemon yellow – excited array element; red / orange – neighbouring, grounded elements, used to calculate electromechanical crosstalk; dark red – remaining neighbouring elements.

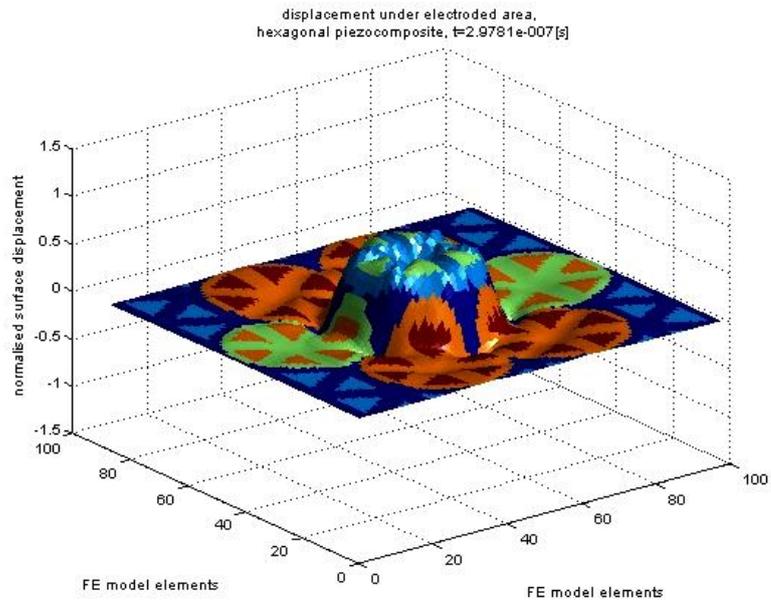


Fig. 4.5 An example of the desired displacement mode. Visualization of the model surface at $t=297$ ns after the excitation pulse. The centre element is being excited and protrudes from the surface of the material, radiating a coherent wave of motion into the front half-space; while neighbouring elements remain at rest.

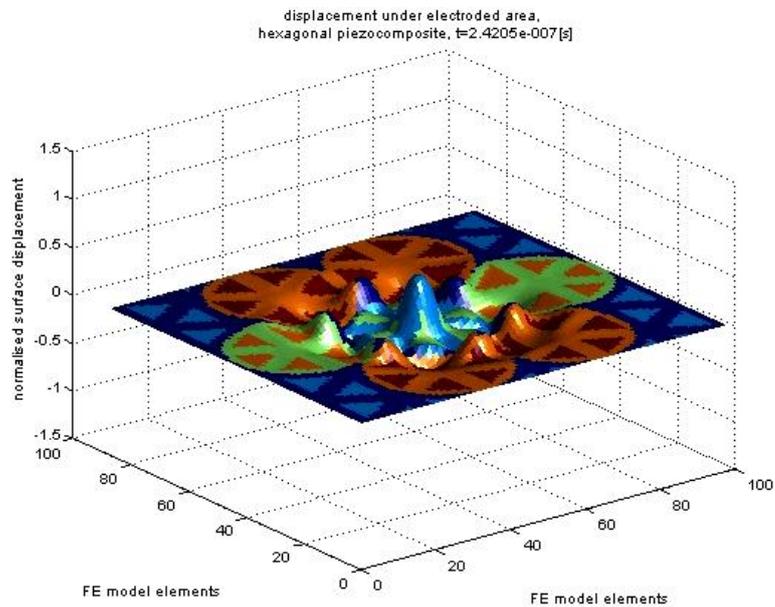


Fig. 4.6 An example of undesirable behaviour: the apices of the triangles vibrate at a different frequency than the ceramic pillars; such motion modes do not convert into radiated energy and are absorbed into internal material friction instead.

4.3.3 Model simulation study 1: varying kerf pitch

The initial modelling study predicted electrical impedance of the slice as a function of kerf pitch, assuming constant kerf size of $100\mu\text{m}$ and platelet thickness of $450\mu\text{m}$. An example impedance obtained with this model is shown in Fig. 4.7.

It is known that the frequency of resonant modes of vibration that couple through the piezoelectric effect to the electrical response of the model, can be identified as peaks and troughs in the impedance plot of the excited model. Analysis of these peaks and troughs is used to gain insight into the design of the material properties.

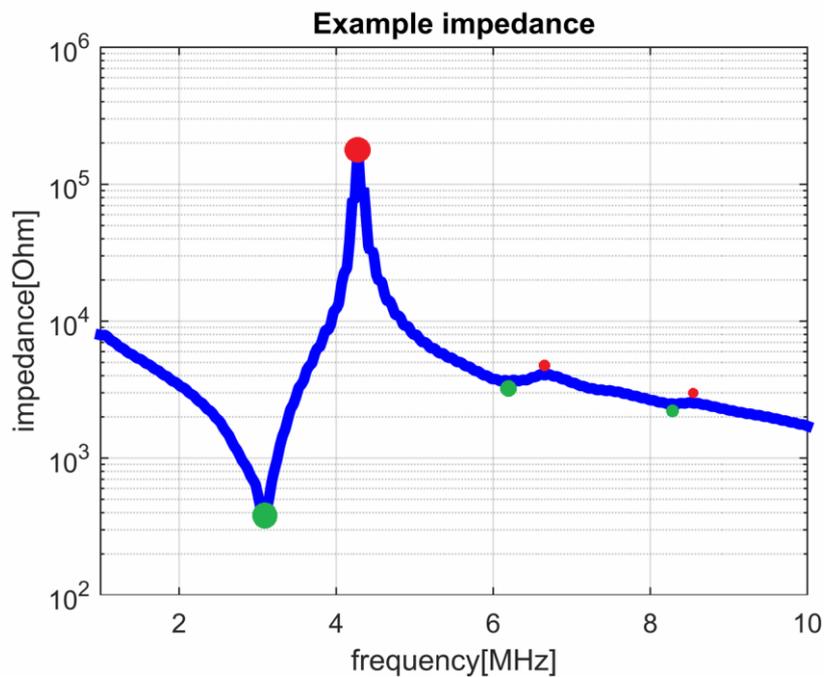


Fig. 4.7 An example of a simulated impedance plot of triangular-cut single array element for platelet thickness of $450\mu\text{m}$ and kerf pitch of $450\mu\text{m}$. Green dot shows frequency of the electrical resonance, and the red dot shows the frequency of the mechanical resonance of the device

Peaks in the calculated impedance curve have been found for each case in the parametric sweep, using a non-derivative peak detection method[61]. The waterfall plot of conductance as a function of frequency and kerf pitch is given in Fig. 4.8. Here, absolute value of admittance (inverse of impedance[62]) is used because it lends itself better to visual analysis.

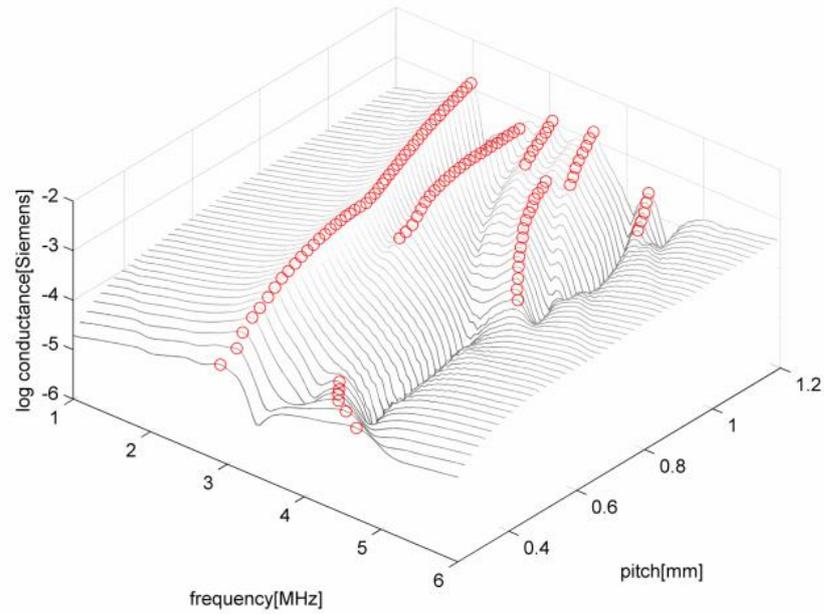


Fig. 4.8 Surfaceplot of admittance as a function of frequency and kerf pitch.

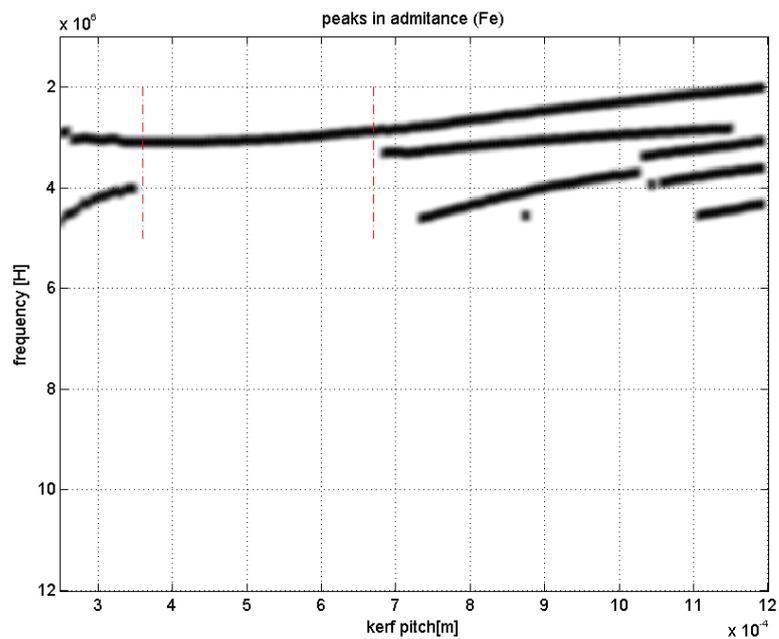


Fig. 4.9 Trace of peaks in the admittance of composite platelet versus kerf pitch. Red dotted lines mark the kerf pitch range for which only one resonant mode is detectable.

Location of the peaks in admittance have been identified and plotted against kerf pitch, as illustrated in Fig. 4.9. The thick black lines are painted at the kerf-frequency location of the crest of the admittance. These peaks correspond to vibration modes of the modelled

element. Analysing Fig. 4.9 permits criteria for unimodal design to be identified. Within a certain kerf pitch range (here, 0.35mm to 0.67mm) there is only one admittance peak visible across the frequency spectrum. Hence, it would be appropriate to fabricate a triangular cut piezoelectric composite within this range.

4.3.4 Model simulation study 2: varying material thickness

The initial modelling study considered the kerf pitch as an adjustable parameter. Alternatively, if kerf width is constrained, the material thickness can be adjusted to obtain a similar result.

In the second simulation study, kerf pitch has been held constant at 450um, and the parameter sweep was conducted on composite thickness. All of the other simulation properties have been left as in simulation study 1. Resulting admittance results are presented in Fig. 4.10. Location of the peaks in conductance have been identified and plotted against kerf pitch, as illustrated in Fig. 4.11. Similarly as in the previous run, a range of thickness values are clearly identifiable that result in unimodal operation. From the Figure, a piezoelectric composite thickness between 0.4mm – 0.67mm will produce a distinct fundamental resonance mode, with the thickness of 0.47mm predicted to produce a device operating at 3MHz.

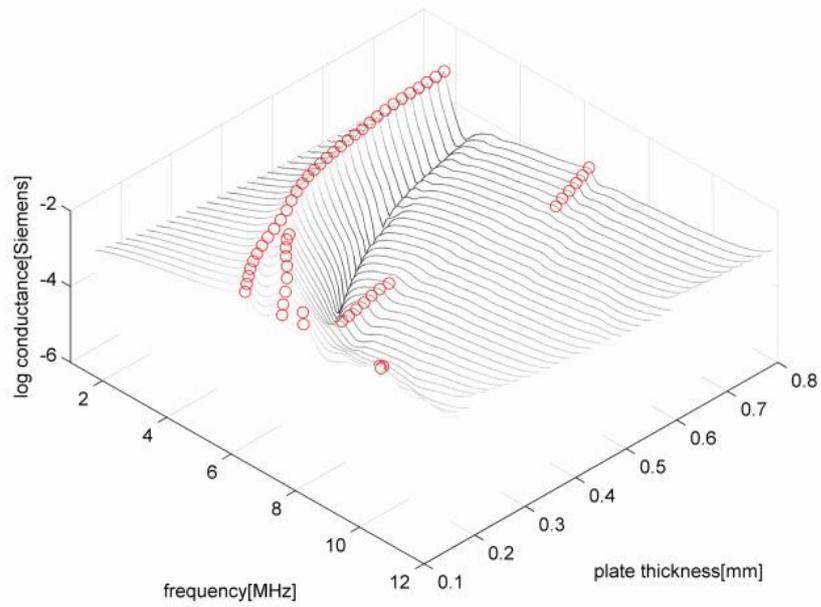


Fig. 4.10 Surfaceplot of admittance as a function of frequency and kerf pitch.

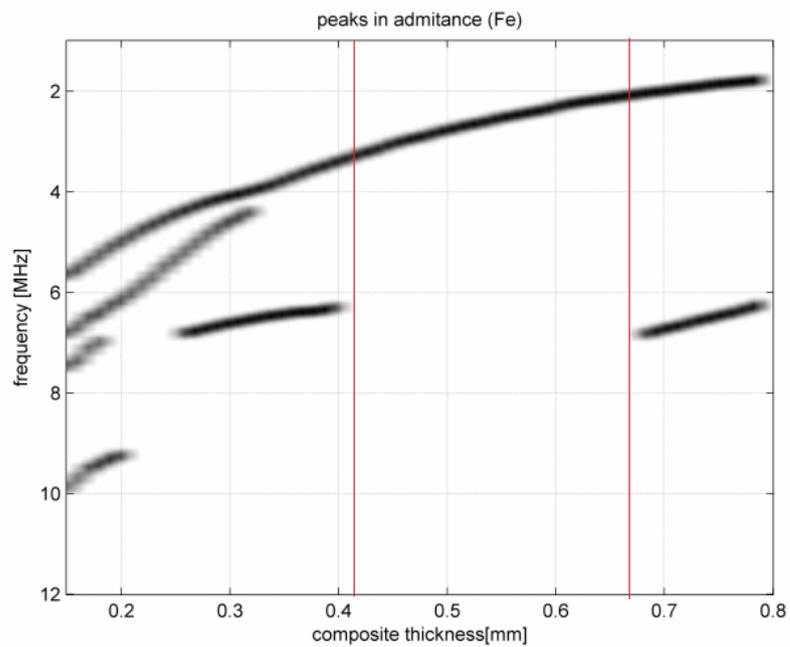


Fig. 4.11 Frequency of peaks in admittance versus piezoelectric composite thickness. Red lines mark the thickness range for which only one resonant mode is detectable

4.3.5 Discussion of results

Both simulations show that for low CVF configurations (below 30%) the polymer filler situated at the apices of the ceramic pillars tends to vibrate at a separate frequency. At high CVF (above 70%), ceramic pillars have a low height-to-width aspect ratio, and they develop internal secondary lateral vibration modes close to the desired operational frequency of the device.

It has been found that there is an intermediate CVF range at which neither of these phenomena have significant influences on the device behaviour, resulting in a unimodal behaviour of the piezoelectric composite plate. At the same time, this results in a noticeably high k_t of 0.64 - 0.75, depending on the specific configuration. The downside of the proposed approach is that there is only a narrow range of CVF where the unimodality requirement is met.

Fortunately, the middle range (~50%) CVF is favoured for many NDE applications as it represents a compromise between electromechanical efficiency and acoustic impedance matching.

4.4 Modelling the mechanical crosstalk

The FE model has been extended to evaluate mechanical crosstalk between the array elements. The mechanical crosstalk is defined here as the ratio of the spatially averaged, time-peak displacement amplitude over the area of the active (excited) element, to the same of the area of the receiving element that is immediately adjacent to the excited element.

The receiving element's electrodes are grounded [63], as explained in Fig. 4.4.

In order to establish a reference value, a square cut piezocomposite material model has been created. The material constructional details are listed in Table 4.1. The properties of the reference square cut material have been designed to have the same central frequency, ceramic volume fraction and kerf width. Unfortunately, this means that the area under the element is slightly smaller for the hexagonal element; this should be kept in mind when comparing sensitivity figures.

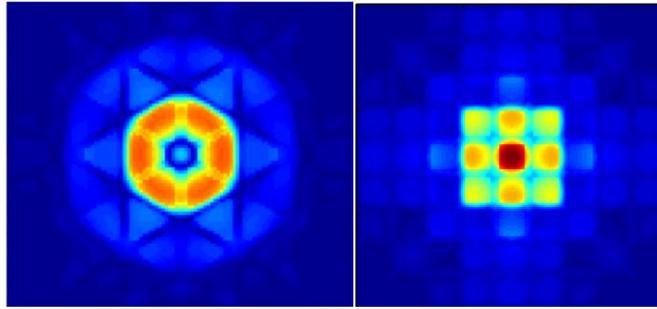


Fig. 4.12. PZFlex model of surface displacement amplitude for hexagonal element (left) and square element (right). Colour scale is relative.

Simulations show an improved level of inter-element mechanical crosstalk for an hexagonal-element composite: -16.7dB versus -14.7dB for corresponding square-cut composite. No additional techniques like leaving an un-electroded row of pillars between the elements [48] have been used to reduce the crosstalk.

At this point, it is worth noting that for a triangular-cut composite the ceramic pillars have significantly higher Pillar Aspect Ratio (PAR)[29], meaning a larger base area when compared to the square-cut composite. This facilitates manufacturing of high frequency devices and it reduces the risk of fracturing pillars during the fabrication process.

Table 4.1. Comparison of properties of piezoelectric composites used in the experimental evaluation

<i>Property</i>	<i>Triangular cut</i>	<i>Square cut</i>
kerf pitch, μm	450	200
kerf width, μm	100	100
Plate thickness, μm	450	450
Plate size, mm	10 x 10	10 x 10
Ceramic Volume Fraction	0.453	0.453
area under respective element, mm^2	0.703	0.837
Pillar Aspect Ratio	1	0.44
Number of pillars per electrode	6	9

4.5 Experimental validation

4.5.1 Manufacturing the test samples

A set of 12mm x 12mm square plate piezoelectric ceramic composites, corresponding to triangular-cut and square-cut devices as per design parameters in Table 4.1 have been manufactured to validate the simulation results. The ceramic is a PZT5H (Ferroperm, Kvistgaard, Denmark), and the polymer filler is Vantico CY221/HY956EN (Hunstman, Cambridge, UK);

On one side, the fired-on electrode (as supplied by the manufacturer) has been cut through to separate the pillars. Then, the liquid precursor of the filler material has been poured in, degassed in a vacuum and left for 24 hours of curing. The result is documented in Fig. 4.13. The other side has been lapped to desired thickness. The surface was cleaned using an ultrasonic bath, dried, and then electroded using metal evaporation in a vacuum chamber. The evaporated materials were chromium first (approximately 5-10 nm) and then gold (approximately 100-200nm). The result of the process is shown in Fig. 4.14.

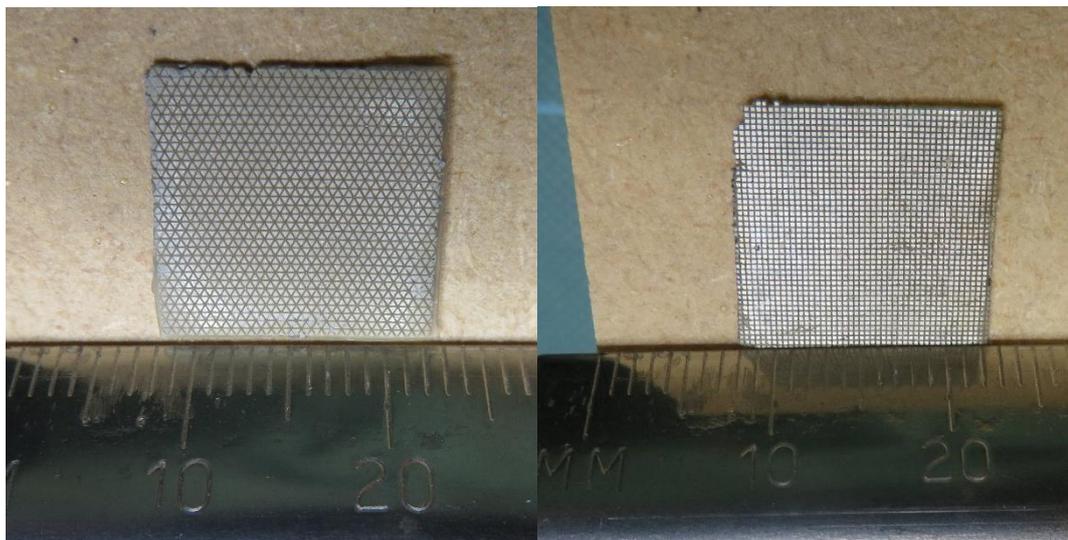


Fig. 4.13. Left: triangular-cut composite; right: square-cut composite. The side shown contains the stock fired-on electrode through which the respective cuts have been made.

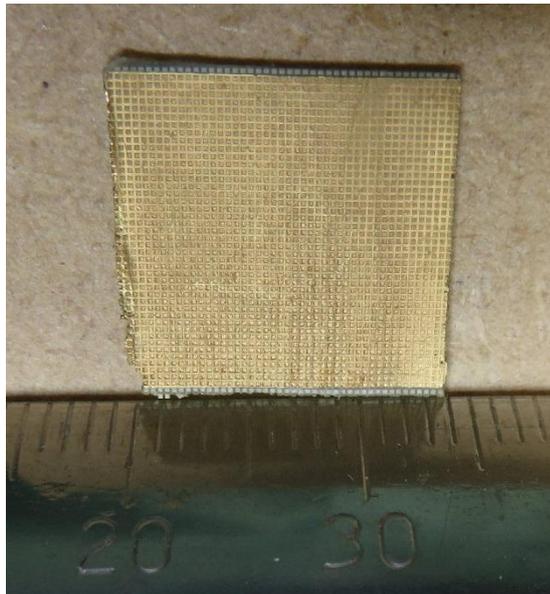


Fig. 4.14. Ground electrode side of the square cut piezoelectric composite. The electrode is made by depositing chromium vapour and then gold vapour in a vacuum chamber.

The capability to fully electrode a dense array pattern onto these substrates was not available at the time of this experiment. Instead, a single element electrode was created by carefully painting the respective pillars by hand, and then gluing a silver wire (75 μ m diameter) to deliver the electrical excitation signal. Subsequent measurements using an impedance analyser and laser vibrometer show that the approach was successful.

4.5.2 Evaluating the test samples

4.5.2.1 *Electrical impedance analysis*

Electrical impedance of both of the new triangular-cut device and the reference square-cut device was measured and the result is presented in Fig. 4.15. The hexagonal element behaviour is unimodal as designed, the k_t [64] is 0.623 for the square element and 0.628 for the hexagonal element.

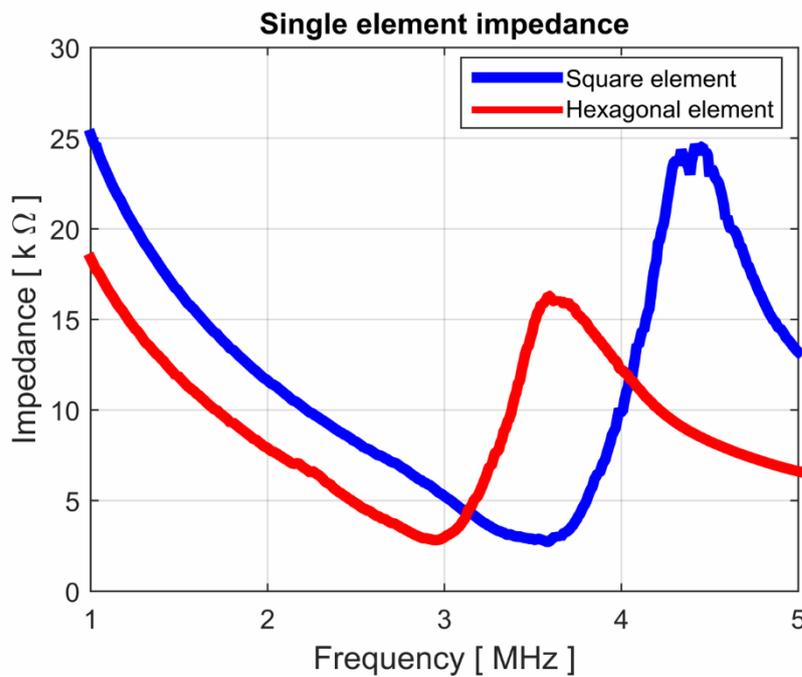


Fig. 4.15. Impedance plots of the manufactured devices: square and hexagonal single element.

4.5.2.2 Scanning Laser Vibrometry

The two representative specimens have been placed in the field of view of the Polytec OFV056 scanning laser vibrometer (Lambda Photometrics, UK). For purposes of this measurement, the specimens were excited with continuous 10V p-p sine wave at the exact resonant frequency (electrical resonance) of the respective specimen, and an out-of-plane displacement amplitude measurement was acquired across a specified surface area within each substrate.

The averaged surface displacement amplitude in air was 0.60nm/V for the hexagonal composite, and 0.62nm/V for the corresponding square composite. The mechanical crosstalk measured was -21.2dB for the hexagonal element composite, and -21.9dB for the square element. Comparing to the simulation, the sensitivity is lower and the mechanical crosstalk is lower for the real specimens than in the simulation. This can be explained by the real kerf filler material exhibiting higher damping than in the simulation³.

The measured performance figures are not better for the hexagonal element, however, given the significant variability in the quality of the manufactured material, it can be concluded that the performance values are similar.

³ For the purpose of the simulation, material properties (including damping) have been measured at 1MHz, but then used to run the simulation at 3MHz.

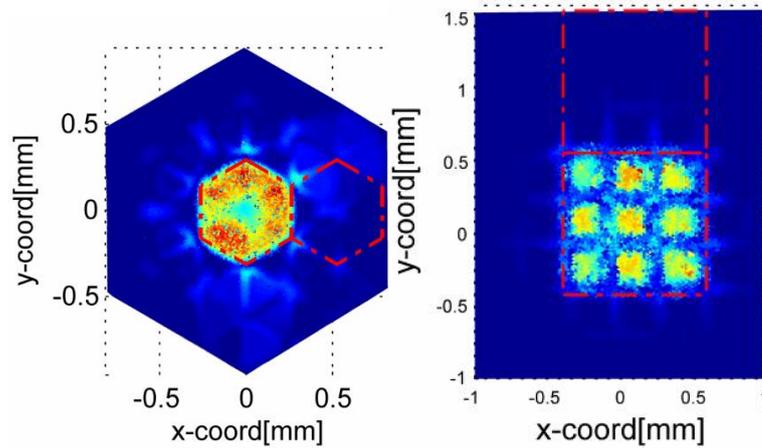


Fig. 4.16. Measurement of the displacement amplitude of the hexagonal and square composite at their resonant frequencies. The Red line denotes the element boundaries used for displacement and crosstalk calculation.

4.6 Conclusions and future work

In chapter 3 it was shown that hexagonal elements have some theoretical advantages over square elements when incorporated within a phased array system. In this chapter a triangular cut piezoelectric ceramic composite material design method has been presented that allows a selection of design parameters to ensure unimodal behaviour of the manufactured material, resulting in acoustic performance comparable to a classic square cut composite. Pertinently, this method allows for a design of a material with pillars that have an aspect ratio closer to 1. This helps the pillars to be mechanically robust and it avoids the problems associated with thin pillars collapsing during the manufacturing phase.

Unfortunately, the material was not packaged up into the final form of having the acoustic backing and acoustic matching layers and because of this, it was not possible to directly compare the impulse response of the devices manufactured here, and the ones found in the literature.

Moreover, the early array elements have been electroded by painting the element with silver paint, and in later attempts, by soldering a thin silver wire to the surface of the piezoelectric ceramic. Both these methods are very risky (in a sense, that it is done by hand, and should it not work the first time, the rework is likely to make things worse), and are

likely to be the main source of variance between the results. There exists a need to develop a method to robustly electrode the 2D array in a way that does not degrade individual 2D array element performances.

Chapter 5. In-probe active impedance matching module

5.1 Introduction

2D phased arrays consist of a matrix of small piezoelectric elements. Their small physical dimensions result in a high output impedance in range of 1-10 k Ω . A typical connection scenario is presented in Fig. 5.1. Due to the electrical impedance mismatch, most of the transduced energy is not delivered to the receiver when a typical 50 Ω receiver circuitry is used.

In the case of a high input impedance instrument being used, the next weakest link becomes the coaxial cable whose capacitance is in the range of 50-200pF/metre (depending on the cable type). For example, if an effective connector capacitance is 400pF (this consists of connector, cabling, and housing capacitances), and the probe working frequency is between 3-5MHz, an effective shunt impedance in range of 100-300 Ω is inserted in parallel with the receiver. The cable impedance loads the sensor, reducing the energy that reaches the receiver circuitry.

In both cases the sensor is effectively working in short-circuit condition, reducing voltage sensitivity and causing a long ring down time. At the receiver end, there is a practical limit on how much gain one can apply, due to the noise gain always being higher than signal gain in amplifier applications [65, Sec. 1.10 and MT-033]. The higher the receiver gain, the higher the noise floor. At the same time, it is known that the impulse response length and wavelet shape have important influence on the final image quality.

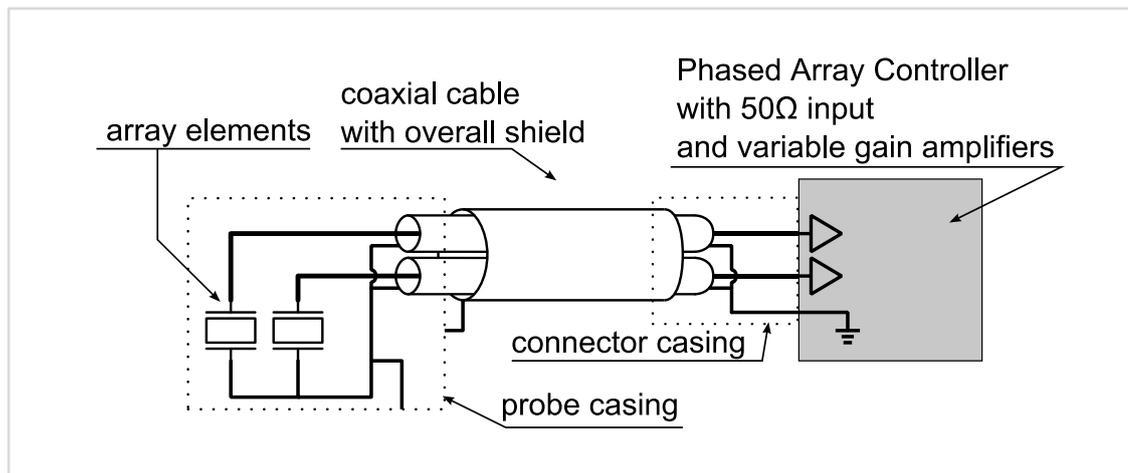


Fig. 5.1 Typical electrical schematic of connection between probe elements and the Phased Array Controller

In this Chapter, a bidirectional in-probe line driver circuit with controllable input impedance is developed to address these issues.

5.2 The effect of a cable on a response of an element of 2D phased array

For typical ultrasonic operational frequencies of up to 20MHz, the corresponding electromagnetic wavelength in copper ($v=0.7c$, where c is speed of light) is 85m for 5MHz; the same is down to 21m for 20MHz operation. Therefore, for a cable length typically in use with NDE equipment (2-5 m), the cable does not behave as a waveguide. However, it does not behave as an ideal conductor either. Due to high electrical impedance (in range of 1k-10k Ω at the operational frequency) of the ultrasonic transducer, the capacitive nature of the cable dominates signal transfer characteristics. Basic models show that inductive and resistive characteristics of the cable are of secondary importance.

For the purpose of this work, a review of commercially available cable characteristics has been performed. The lowest capacitance cable found was 38pF/m for a flat ribbon, 600 Ω TV antenna cable; this however comes at the cost of the cable being bulky and inflexible, as it is composed mostly of the spacer between two conductors. Typical capacitance for an unshielded, twisted pair cable is in range of 45pF/m; typical capacitance for a solitary, high quality coaxial cable is 55pF/m. The capacitance increases for multi-signal bundles, increases compact and miniature cables, and also increases further if the cable is

individually and/or overall shielded. For a multi-signal miniature coaxial bundle with overall shielding the capacitance reaches 150pF/m. This value applies to the cable itself, meaning that the capacitance of any connectors need to be also taken into account.

5.2.1 Twisted pair cable properties

The argument for using coaxial cable in bundles is that coaxial cable exhibits lowest crosstalk out of all cable types. However, using such cable might be an over-engineering of the problem, as shown in this Section. For a twisted pair cable bundles (e.g. Ethernet cat5 type), assuming that the currents are balanced for each pair, the front-to-end crosstalk between separate pairs is typically rated at -40.8dB per 100m at 10MHz (lower for lower frequencies). This translates to -80.8dB/metre, in other words, very low. For comparison, the maximum achievable signal dynamic range with 12bit ADC is only 72dB; in practice significantly less.

Transmit signals might be as strong as +100dB versus received signals, but they occur at different times and therefore are separated. Even if the electrical crosstalk is significant, its effect is going to be reduced by the application of the imaging algorithm. The typical delay-and-sum ultrasonic imaging algorithms reduce the influence of the time-synchronous electrical crosstalk.

Another consideration is that twisted-pair cables are easier to work with during probe assembly, and also can be potentially more flexible than coaxial types.

Regarding the EM energy radiation, it is recommended that there is an overall electrostatic shield (aluminium foil) employed over the cable. This will help to reduce any incoming or outgoing EM interference.

General wisdom dictates that a coaxial cable will provide best signal fidelity out of all cable options; however this is only true if the currents are correctly balanced between the inner and outer conductor. This might not be the case for short, non-waveguide cables with common ground as shown in Fig. 5.1.

Summary

Considering the earlier discussion in this section, the recommendation for probe design is to use a twisted pair cable bundle with an overall shield. Compared to coaxial, such cable

provides lower capacitance, it benefits from balanced conductor currents if available, it can be potentially thinner and lighter, and the electrical crosstalk is low enough for typical ultrasonic NDE imaging applications.

5.2.2 Modelling the effect of the cable capacitance on the probe sensitivity

The effects of including a capacitive element in the signal path are well understood, however, the following simulation has been conducted to demonstrate the severity of the issue.

This analysis is conducted for a 5MHz, 1λ spacing sparse 2D array operating in direct contact with steel, which is the best case in terms of obtaining low source sensor impedance. The worst case scenario would be a dense, 0.5λ spacing, 20MHz 2D array operating in water; in this instance it would be more affected by the cable capacitance. The cable is simulated by insertion of a lumped capacitance, which is a valid simplification given the non-waveguide behaviour as per the analysis given in section above. The 2D array is based on a generic PZT5H piezocomposite, which is the preferred material for NDE probes due to its relatively high sensitivity and relatively low Q factor. The device is a half-backed, single layer matched to water and has a matched damping resistor. For the above 2D array specification, the calculated element size is 1.12mm and material thickness is 0.30mm, kerf is 0.1mm, pillar width is 0.16mm, and ceramic volume fraction CVF=38%. In this example a single array element is created by placing a common electrode over 9 ceramic pillars. The element is connected in parallel with a $1k\Omega$ damping resistor to improve bandwidth.

An initial simulation of the design was created in PZFlex using a full 3D model with backing, acoustic matching, and water as the load medium, as depicted Fig. 5.2. This transmission only simulation yields an estimate of the resulting impulse response and electrical impedance of the device.

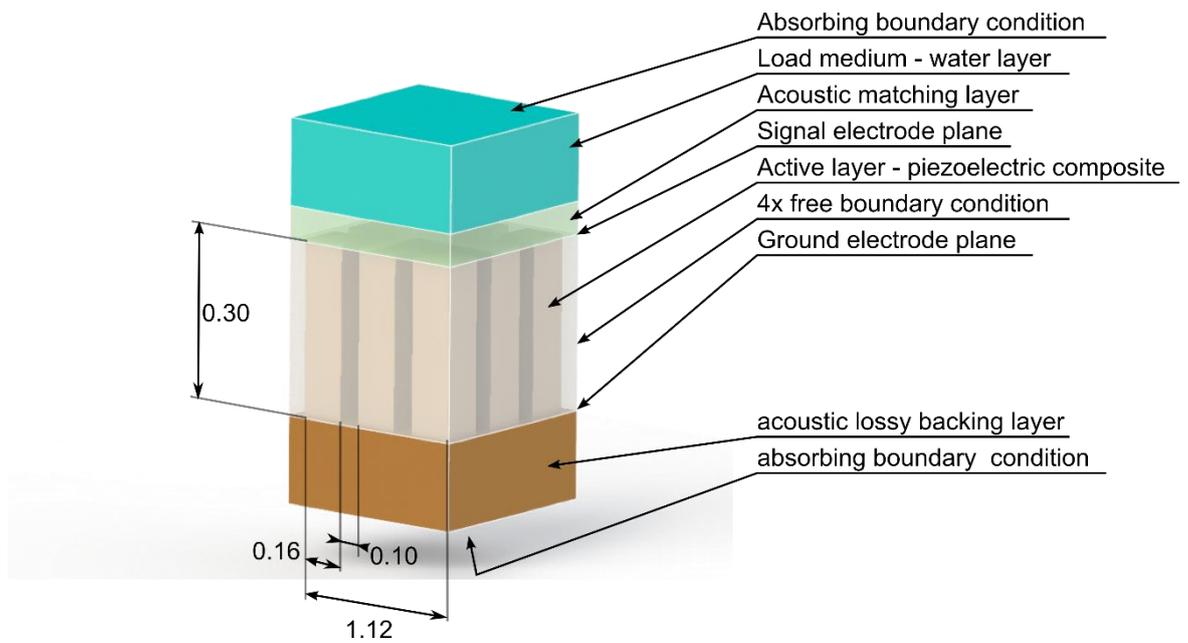


Fig. 5.2. Schematic of the PZFlex model used in the example

The impedance profile and time domain impulse response of the device obtained are presented in Fig. 5.4. The electrical impedance of a single array element is 765Ω at 4.0MHz at electrical resonance, and the impedance at mechanical resonance is 1706Ω . The impedance is consistent with the prediction based on ceramic permittivity. The simulated element has a static capacitance of approximately 45pF .

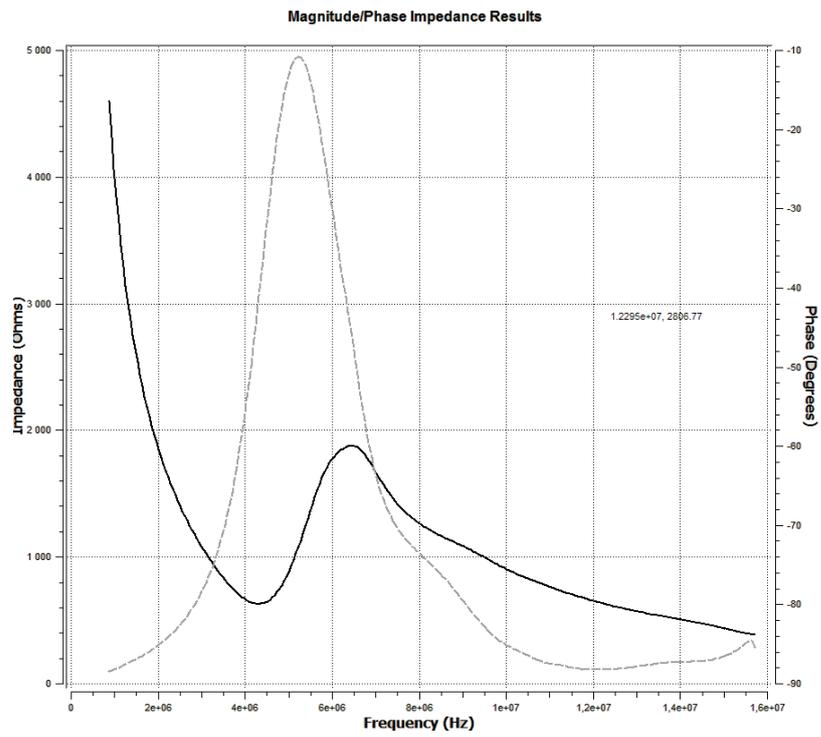


Fig. 5.3. Electrical impedance of the modelled system assuming no influence from the capacitance of the cable.

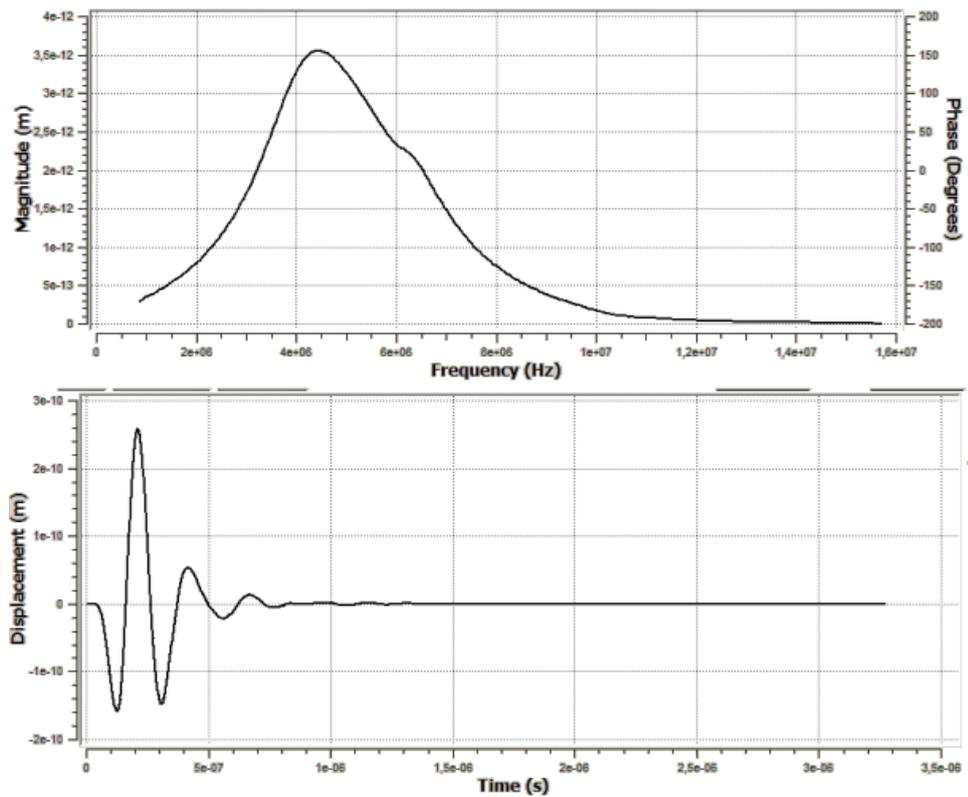


Fig. 5.4. Electrical transmit impulse response of the modelled system assuming no influence from the capacitance of the cable.

In order to demonstrate the effect of the cable, the model has been converted to pulse-echo operation mode. The receiver circuitry is of the high impedance type. The received voltage has been recorded and an effective pulse-echo insertion loss has been calculated. The cable capacitance (directly dependent on the cable length) has been varied and the insertion loss of the transducer-cable system has been calculated. Results of the simulation are gathered in Table 5.2.1 and Fig. 5.5.

Table 5.2.1. Results of the insertion loss simulation

Cable capacitance [pF]	Pulse-echo insertion loss[dB]
10	-31.1
20	-36.3
30	-39.6
40	-42.0
50	-43.9
75	-47.3
100	-49.8
150	-53.3
200	-55.8
250	-57.8
300	-59.4
350	-60.8
400	-62.1
450	-63.2
500	-64.2

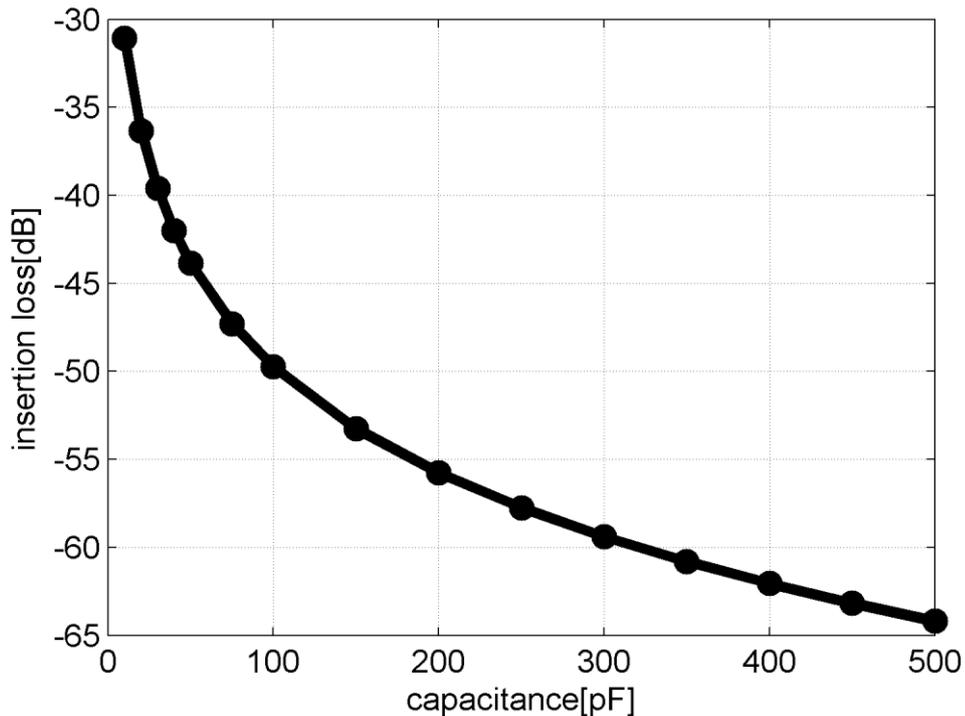


Fig. 5.5 Simulation parameter sweep: insertion loss versus capacitance of the cable.

It is clear that the higher the capacitance, the higher the insertion loss, increasing by 31dB by going from 10pF (virtually no cable, rated PAC input amplifier capacitance only) to 450pF (typical 3m micro-coaxial bundle cable).

The capacitive nature of the cable means that the cable behaves as a capacitive shunt. In the transmit phase it acts as a voltage divider. The apparent divider ratio depends on two factors: the physical area of the transducer element (the smaller footprint, the worse) and length of the cable (the longer length, the worse).

In reception the same cable acts as a shunt, lowering the apparent input impedance of the receiver electronics, when ideally, the desirable impedance is very high. In some cases it might act as an energy reservoir, resulting in resonant behaviour and reducing temporal resolution.

There are 3 known approaches to this problem:

- (a) Increasing the lateral transducer area (lowering its impedance).
- (b) Lowering the cable capacitance.

- (c) Installing an impedance transformer in the probe handle or as close to the transducer as possible.

Approach (a) is addressed in Chapter 3. Approaches (b) and (c) are addressed in this chapter.

5.3 Optimal receiver circuitry resistance

5.3.1 Instrumentation signal input impedance

There are, in general, two types of receiver electronics used: with input impedance of 50Ω , and “high impedance” (high-Z) – meaning input impedance of at least $1M\Omega$. The motivation for using 50Ω input impedance in the instrumentation is that it is matched to the impedance of coaxial cables at high frequencies⁴ (waveguide mode). It has been shown however, that the range of frequencies and cable lengths used typically in NDE systems does not justify this approach. Therefore it is recommended to always consider the high-Z input option first.

5.3.2 Resistive damping for shortening the impulse response

The intention of this study is to find the optimal resistive loading condition that would shorten the impulse response time (ring down time) of the transducer. In this study, the transducer is considered to be the source of energy.

The assumption is that the resonant behaviour of the transducer will be damped down by absorbing energy from it. According to Jacobi’s Law, the energy will be transferred from the source most efficiently when the load impedance equals source impedance. Note that in case of the complex-valued source impedance, energy extraction can be made (momentarily) faster by using a complex-valued load. This however leads to resonant behaviour of the circuit. The additional requirement is to make the circuit non-resonant; hence the load must be of resistive character.

⁴ Often, single element ‘standard’ transducers are designed to exhibit 50Ω impedance at electrical resonance frequency.

In case of a piezoelectric transducer, the effective source impedance is a function of the load impedance, as will be shown later in this chapter. Therefore a numerical sweep is given here to establish the conditions leading to optimal loading impedance.

In this study, a generic PZT5H slab with dimensions 15x15mm and a thickness of 0.4mm, suspended in a vacuum, is modelled in PZFlex. This new model is introduced for experimental convenience. The specific area and dimensions were selected to avoid any spurious resonances and make the analysis easier, and also to match the stock material available in the lab at the time of this writing. It is assumed that the dependencies found will transfer over to piezoelectric ceramic composite materials.

5.3.3 Transducer model in transmit mode

The PZFlex model operates in transmit mode, assuming electrical excitation and no electrical damping (Fig. 5.6 and Fig. 5.7). This is to establish basic parameters such as electrical and mechanical resonance frequencies, and the impedance of the slab in the electrical resonance mode.

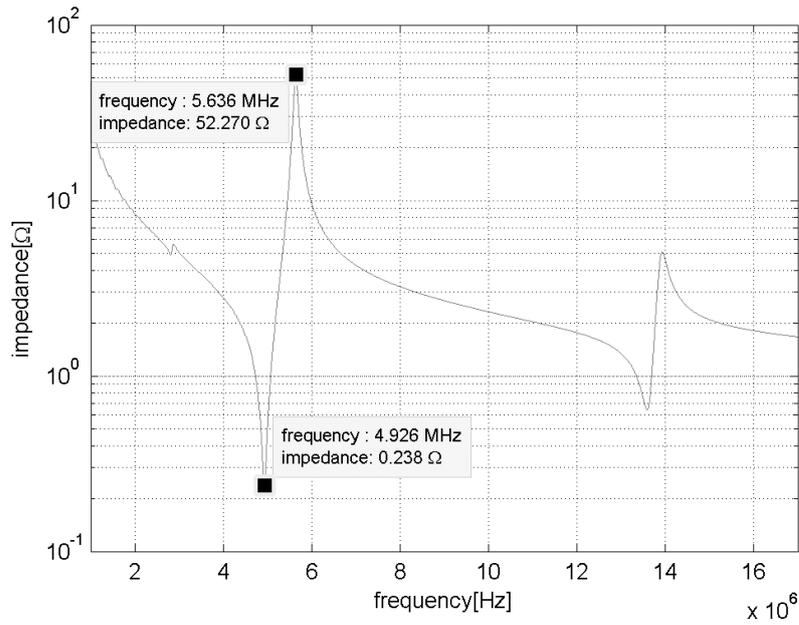


Fig. 5.6 Impedance plot of the ceramic slab in the study and the characteristic points marked out

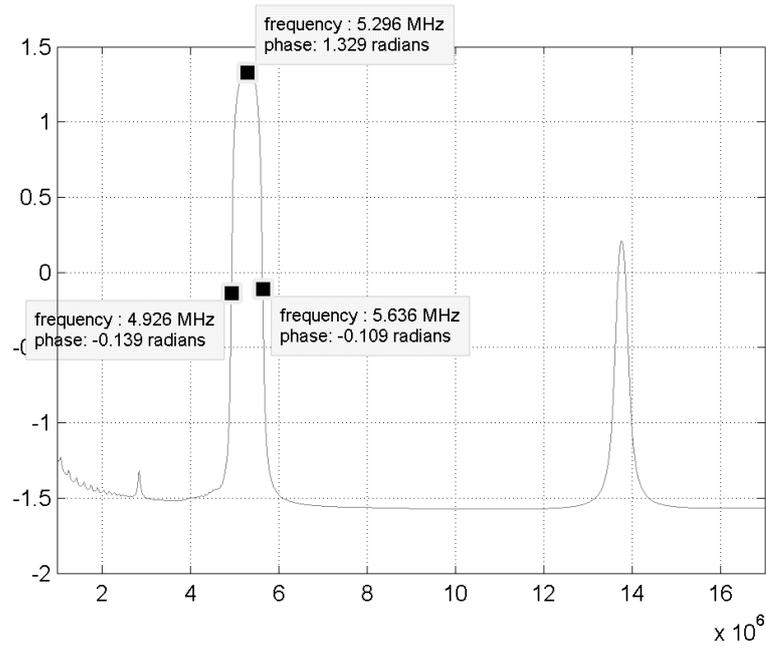


Fig. 5.7 Phase plot of the ceramic slab in the study and characteristic points marked

5.3.4 Model parameter sweep of transducer in receive mode

In this mode there is no electrical excitation, and instead, the excitation of the model is a 5MHz centre frequency, 1N amplitude force pulse, coupled mechanically to one side of the

ceramic. A resistor is connected between the two electrodes of the ceramic, and the voltage between the electrodes is measured as the receive mode sensitivity. The resistor is varied to produce a parameter sweep. Recorded parameters are:

- -20dB decay (ring down) time of the received voltage(Fig. 5.8)
- Central frequency of the received signal (Fig. 5.9)

It should be noted that the effects of the resistor in the receive mode can also shorten the impulse response in the transmit mode. This is because as soon as the transmit excitation signal ceases, the transducer is instantly in receive mode and any damping modality helps to shorten the residual oscillation.

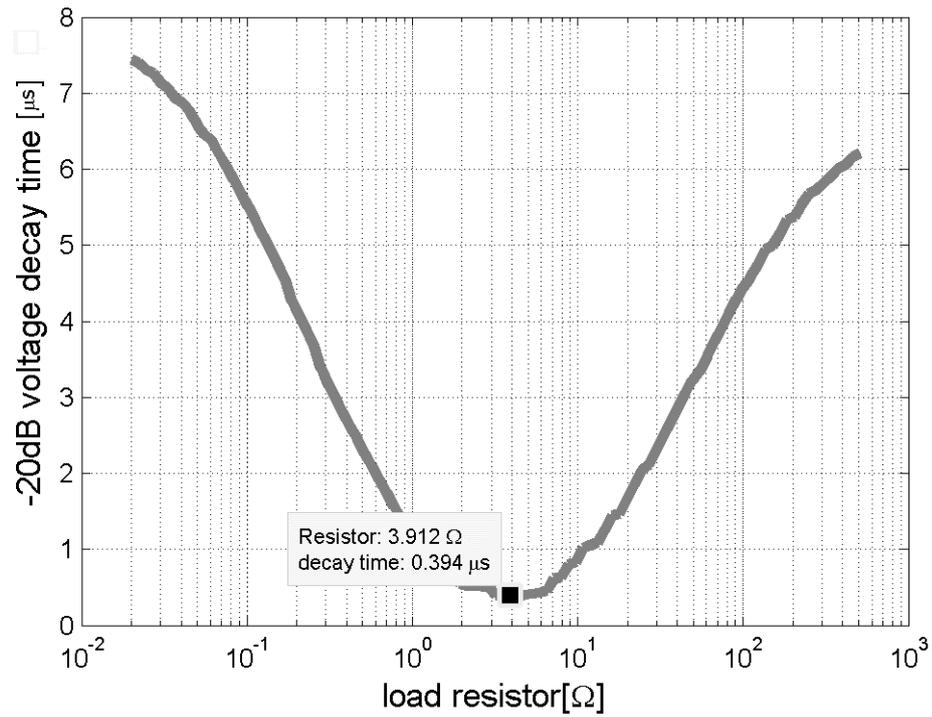


Fig. 5.8 Received voltage decay time vs load resistance

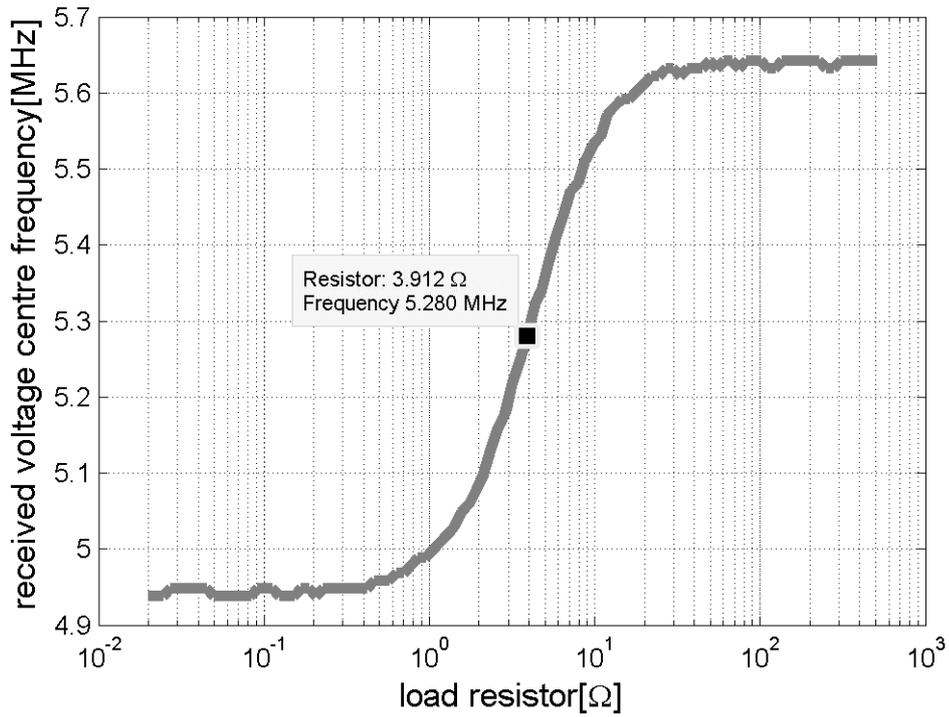


Fig. 5.9. Centre frequency of the received signal vs load resistor. The frequency spectra used for the analysis has been normalised versus the spectra of the excitation signal.

5.3.5 Results analysis

The criterion for optimality of the solution is a load resistance that allows for the shortest ring down (decay) time of the received signal. For the above simulation, the optimal damping resistor value was 3.91Ω , as shown in Fig. 5.8.

The optimal resistance in terms of absorbing the energy from the ceramic is not the one that matches the electrical f_e resonance impedance of the ceramic, nor the one that matches the mechanical f_m resonance impedance of the ceramic. Due to loading of the ceramic with the resistor, its receive mode resonant vibration frequency changes between f_e and f_m . Convergence occurs when the ceramic is loaded with impedance matching to the approximate frequency of peak phase shift between the f_e and f_m .

It is known that the mechanical loading condition also affects the resonant frequency. It is therefore anticipated that the optimal resistance value will depend on the kind of mechanical loading of the ceramic in the actual probe, for example, depending on acoustic backing, matching and load impedance. Additionally, it will change depending on the parasitic capacitance present in the system, either due to the cable, or the in-probe electronics load.

Fortunately, it appears that as long as the matching resistor value is close to the optimal value, the decay time is not overly sensitive to its exact value, changing smoothly by a factor less than 2x when resistance changes by a factor of 2x away from what is optimal. Therefore, depending on the required performance of the final probe and the budget of the project, the resistance may or may not need tuning to match the optimal value.

In the simulation case the matching resistor has shortened the impulse response from very resonant (narrowband) down to $0.4\mu\text{s}$, or 2 cycles of oscillation at 5.28MHz, which is an excellent result.

5.3.6 Summary

It has been shown that:

- For a typical 2D phased array transducer element, the character of the cable is capacitive.

- Capacitive behaviour adversely affects both the voltage sensitivity and the impulse response time of the probe.
- There exists an optimal value of the loading electrical resistance that can be connected to the transducer to minimise the impulse response decay (ring down) time.

Pertinently, depending on the electrical load resistance, the peak in voltage sensitivity in receive mode moves from electrical resonance to mechanical resonance frequency, as shown in Fig. 5.9. Taking the measurement of the transmit-mode impedance, and matching the resistor with the impedance of the ceramic where the phase shift is maximally positive can approximate the optimal value of the loading resistor.

In the design phase, if possible, a full model including effects of backing, matching and probe electronics (in particular, including the effects of the cable capacitance) should be used to optimise the resistor value.

In the manufacturing phase, depending on the budget, tuning the exact resistor value may shorten the response time up to a factor of 2; otherwise the approximate value will still yield an enhanced response when compared to no dedicated damping at all.

Note that damping resistor value considerations only apply if the input impedance of the remaining part of the circuit is high, and in particular, if the cable is short enough for its capacitance to be of a small effect.

5.4 Active In-Probe Line Driver

5.4.1 Introduction

Cable impedance issues can be avoided entirely if an in-probe operational amplifier configured as "repeater" or "line driver" can be employed. The circuit has to be carefully designed with the following guidelines in mind.

Design objectives

- Low input capacitance. Most low-cost op-amps are not suitable due to not being specified for input capacitance, or having a specified input capacitance comparable to the capacitance of a bare 5m cable. It should be noted that any Transmit/Receive(Tx/Rx) switch or multiplexer adds to the capacitive load of the transducer.
- High current drive capability, and high output slew rate, to preserve waveform shape. Some of the op-amps have bandwidth specified for small signals only; handling large signal swings only to 1/10th of the rated bandwidth due to the slew rate limit.
- Overload tolerance and fast overload recovery. Unfortunately, most fast (>10MHz) op-amps operate at low voltages of up to $\pm 5V$. There are only a few op-amps working at up to $\pm 18V$. There is however some specified for higher voltages from low-volume manufacturers, at a much higher cost. This requirement is of less importance if transmit and receive channels are physically assigned to different array elements, or if some kind of Tx/Rx switch is employed.
- Low physical volume. One has to keep in mind that for a 128-element array, there has to be 128 analogue amplifier channels, which might be too bulky and impractical to be placed inside a probe handle. Therefore minimisation of circuit complexity is of paramount importance. Use of quad amplifiers in a single package, and omission of Tx/Rx switch is suggested.
- Optionally, a band-pass filter may be considered to help to reject off-band interference. Linear phase behaviour in the pass-band is essential to preserve the pulse shape, and therefore, reflector spatial location.

- Finally, price of components has to be considered, as each component has to be repeated per channel.

5.4.2 Circuit operation principle

To solve this problem, a bi-directional line driver (Fig. 5.10) is proposed to increase the impedance seen by the sensor. The circuit is designed to be inserted in a series between the piezoelectric element and the Phased Array Controller (PAC), and allows bi-directional flow of signals. The high voltage (HV) excitation pulse from the PAC is conducted by pass-through silicon diodes, reaching the piezoelectric element while being blocked from entering output or input of the low-voltage op-amp by the automatic HV switches. When the high voltage ceases, the HV switches close, allowing faint echo signals captured by the sensor to flow through the amplifier and the coaxial cable, back to the PAC. In this mode, the silicon diodes do not conduct (voltages below 0.7V) but rather behave as high value resistors. The small non-linearity introduced by these diodes is reduced if the op-amp gain is set to high value, utilising the negative feedback effect. Importantly, the resistor R_{in} , in series with the resistance of the switches and the non-conducting diodes (voltage across diode series below the threshold), controls the apparent input impedance of the inverting op-amp, and allows matching it to the impedance of the sensor. This improves energy transfer away from the sensor as explained in Section 5.3.4.

At this point, the reader should note that despite the op-amp having an idealised high input impedance, when it is used in inverting amplifier configuration, as in Fig. 5.10, the apparent amplifier input impedance is equal to the value of the resistor R_{in} . This is due to the way the currents flow through the feedback resistors. The op-amp drives its output in such a way as to balance the voltage on the positive and negative input terminals; the positive input terminal is at ground level, so the negative input terminal is driven towards this level through the feedback resistor. In effect, the input signal sees the R_{in} as if it was connected to the ground.

This damping effect of such a circuit on the transducer means that the effective bandwidth of the sensor is extended, and this is a primary benefit of the proposed device. Low input capacitance of the line driver means that the voltage output of the transducer itself becomes much higher reducing noise. At this point the reader should note that the energy

transduced by the sensor does not change; it's just the (relatively) high apparent input impedance of the circuit that causes the voltage to go up. This higher voltage is then power-amplified before travelling through the coaxial cable, improving immunity of the system to noise sources. The capacity of the cable is charged using the energy supplied from a source other than a transducer. At the same time, the transducer is insulated from the electrical capacitance of the cable. Output impedance of the circuit can be matched to the input impedance of the PAC, further reducing noise and overshoot.

5.4.3 Component selection considerations

The resistor value R_{in} is selected using the method described in Section 5.3.4. The feedback resistor is selected to set a desired circuit gain, for example, 0dB. At this point the circuit gain ranges might be limited by the stability criteria of the op-amp used.

The diodes are of a silicon type. It is important not to use the low-threshold voltage type diodes; the threshold voltage must be high enough to allow the useful, low voltage signals to pass through the amplifier rather than through the diodes.

The HV switches used are SUPERTEX MD101 (Microchip Technology Inc., Shanghai), although different solutions for similar functionality do exist [66][67].

The op-amp used must be of a bandwidth of at least a decade wider than the intended transducer operation frequency; e.g. at least 50MHz for 5MHz transducer. It must also be specified for low input capacitance and capacitive load drive. A good starting point is op-amps used as line drivers in Asymmetric Digital Subscriber Line (ADSL) modems. An example of suitable parts would include LM7171.

Note that current feedback amplifiers are not suitable for this application due to their low input impedance in the inverting input. [65, Sec. 1.18]

The reader should note that the quality of the amplifier response would affect the apparent circuit input impedance characteristics. The band pass phase performance of the circuit should be carefully designed in order to preserve the resistive characteristic of the input. This means that the circuit is best designed with a specific cable length in mind. When considering the design of the capacitive drive output ballast resistor (**not** shown in Fig. 5.10,

but often recommended in op-amp datasheet), inherent resistance of the HV switch should be taken into account.

Alternatively, the effect of the cable capacitance can be further insulated from the transducer by using a secondary voltage follower, as shown in Fig. 5.11.

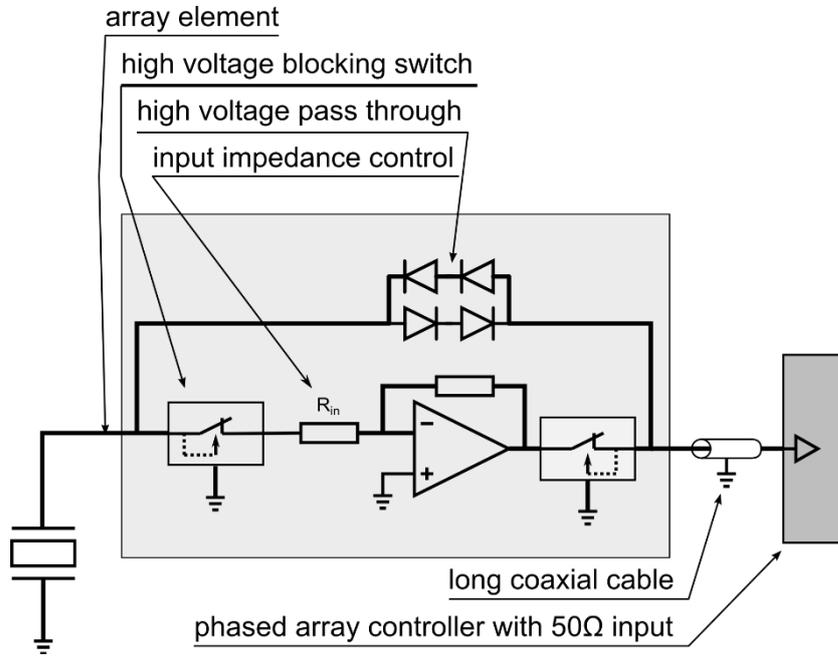


Fig. 5.10: Line driver schematic overview - single channel

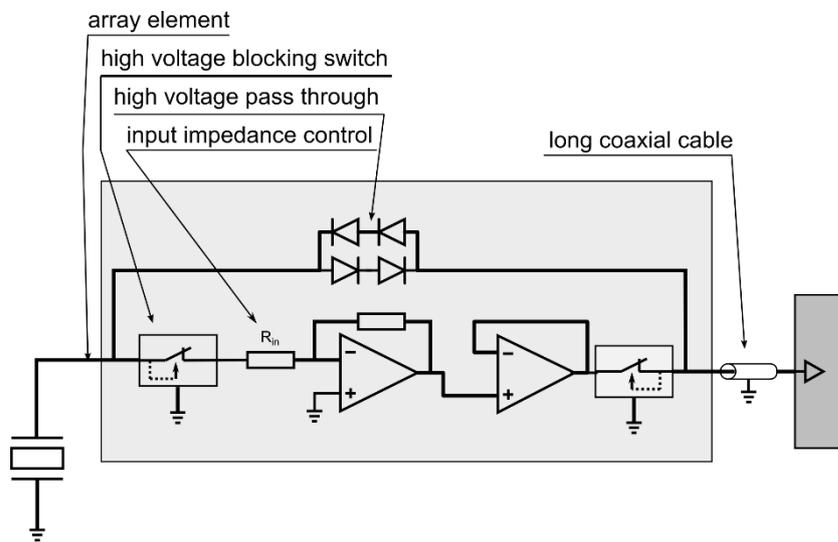


Fig. 5.11 Line driver schematic with additional long line driver

5.4.4 Example implementation

The line driver has been designed for compactness, bearing 4 channels on a relatively small PCB of size 54x17mm, as shown in Fig. 5.12 and Fig. 5.13, Further refinements are possible to reduce this size. Power consumption is in range of 0.05W/channel during active operation only (can be further reduced by powering down between signal acquisition phases), and depends mainly on the kind of the op-amp used. Since the high voltage impulses needed to excite the transducer in the transmission phase come from the phased array controller, no bulky high power high speed transistors nor capacitors are needed inside the probe case.

In the prototype board, large soldering pads intended for manual soldering were used. A production style board could use smaller pads and multiple circuits per PCB to further save space.

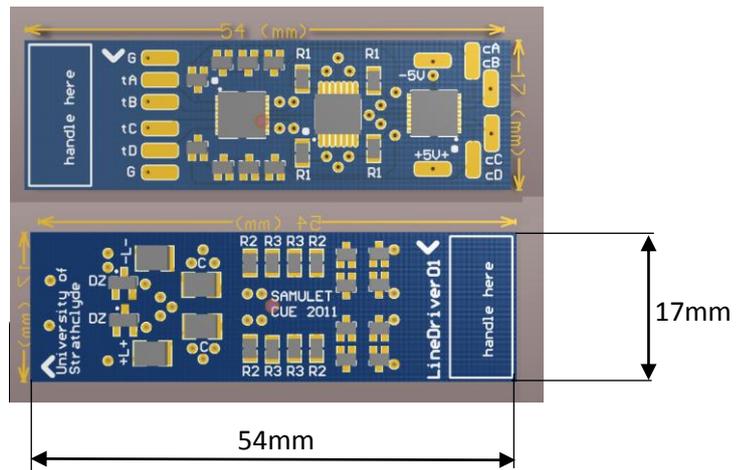


Fig. 5.12: visualization of PCB Layout of the proposed line driver, 4 channels per board.

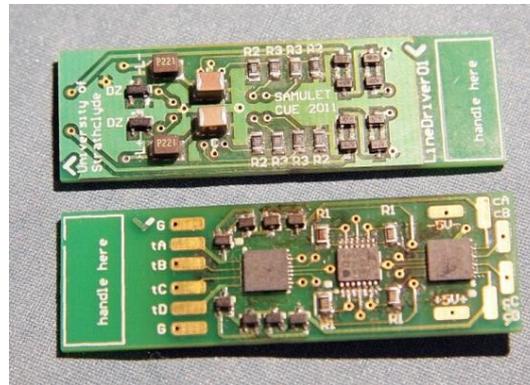


Fig. 5.13: Photo of assembled PCB with 4 channels

5.4.5 Evaluating the prototype

The prototype circuit has been connected to a ceramic-only transducer, whose electrode was trimmed to exhibit impedance of $1k\Omega$ at electrical resonance frequency. The ceramic was then placed on a glass block and pulse-echo signals have been captured. Result of their analysis are presented in Fig. 5.14.

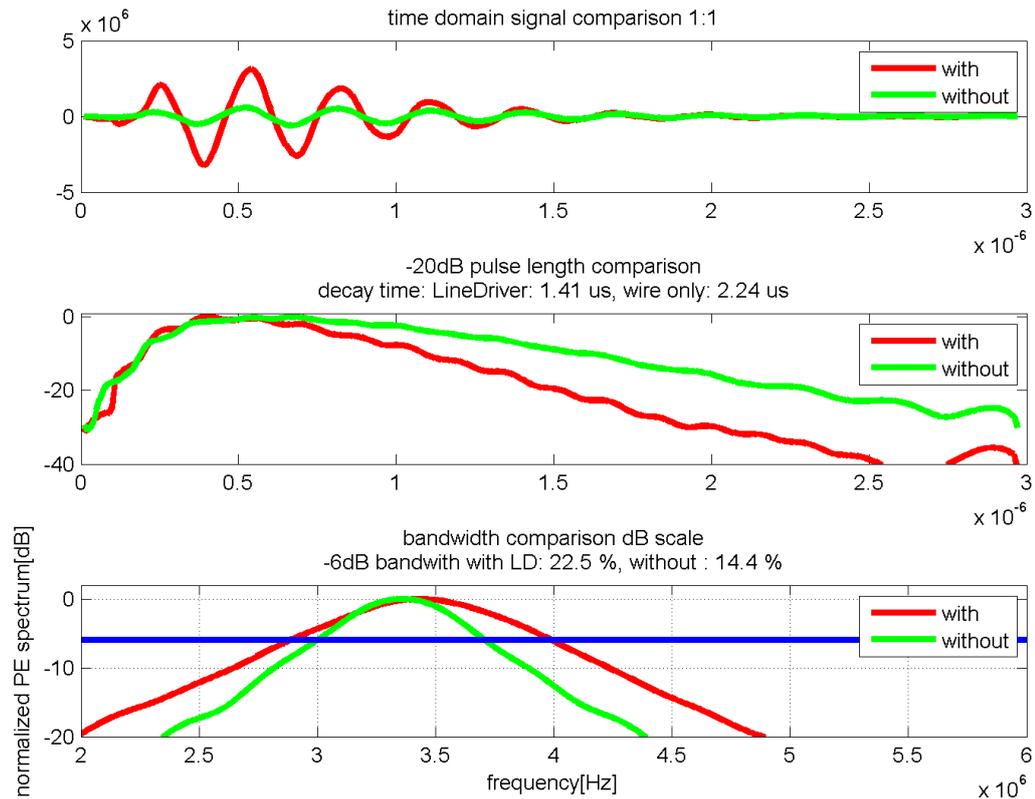


Fig. 5.14. The results of an analysis of pulse-echo time signals acquired with and without the line driver. The blue line represents -6dB boundary.

The line driver has reduced the -20dB impulse decay time from 2.24 μ s to 1.41 μ s, or by a factor of 1.6x; this is less than the expected factor of 3x. The difference can be explained by the originally unforeseen effect the cable capacitance causing the op-amp to overshoot, resulting in the input impedance not being perfectly resistive. This could be improved by utilizing an in-circuit programmable resistance, and a second line driver stage, at an additional cost.

Compared with 50 Ω PAC, the line driver's increased input impedance improved the voltage response of the transducer by 36dB, in addition to the 12dB (4:1) added by the intended fixed gain of the op-amp. This means that the gain in the PAC could be reduced by the same amount. In fact, for the test, the excitation voltage of the pulser had to be lowered; otherwise the transducer would saturate the switches (which open at $\pm 1V$ to the ground).

The noise performance has not been evaluated in this experiment. No obvious signal distortion was observed.

5.5 Case study

In order to apply the proposed circuit to a 128-element probe, 32 boards of 4 channels each are needed. The arrangement of the boards inside an $\varnothing 120\text{mm}$ probe handle is shown Fig. 5.15. Heat transfer is warranted by potting the entire interior of the can in tungsten loaded epoxy. Fig. 5.16 shows the manufactured probe with scale reference objects in a single photo. Note that the images are for feasibility evaluation only, as the boards have not been actually implemented in the final probe.

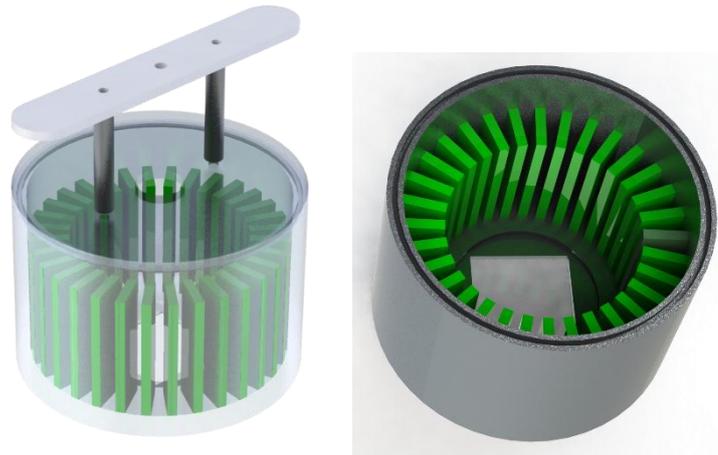


Fig. 5.15. Design view of 32x PCBs (green) inside the body of $\varnothing 120\text{mm}$ probe case.



Fig. 5.16. Photo of manufactured, 128 element sparse array probe in $\varnothing 120\text{mm}$ casing.

5.6 Conclusions and future work

The proposed line driver works by using externally supplied energy to condition the signal received by the small ultrasonic transducer. The transducer was modelled to exhibit properties typical of a 2D phased array element. The prototype device has successfully improved bandwidth by a factor of 1.56x and the sensitivity of the probe by 36dB.

It is proposed that in future versions of the device, the input impedance needs to be in-circuit adjustable so that it can be calibrated for a given transducer of the array.

Chapter 6. GP-GPU accelerated ultrasonic CAD

6.1 Introduction

In creating this thesis, it became apparent that computer-aided design (CAD) and simulation are two extremely important steps in ultrasonic probe design.

The mathematical models used range from simple, empirical closed-form equations through to finite-element models (FEM) that require equation-solving with millions of unknowns. Accuracy, and therefore the usefulness of many of the models, are often compute bound. A more accurate solution can be produced at the cost of performing a greater number of computations. However, these computations bear a cost in terms of hardware, energy and, most importantly, time. Therefore, a 'trade-off' exists between the resources applied during the engineering of the ultrasonic probe solution and the final solution's performance.

In recent years the high-performance computing industry has experienced a significant and evolutionary development. For many years, the silicon transistor technology was undergoing development at a rapid pace, meaning that the computational power one could hire at a given price-point was nearly doubling year-on-year[68]. However, since around 2010, transistors have shrunk to such dimensions that fundamental physical limitations have impeded the pace of this development.

Instead of shrinking the transistors, focus shifted towards a more efficient use of existing silicon technology. Classic Von Neumann architecture is now giving way to a plethora of new architectures.

Among these new architectures, and of particular importance to this Thesis, the General Purpose Graphics Processing Unit (GP-GPU architecture), due to its high-performance to cost-ratio and its wide commercial availability.

A detailed explanation of this architecture is beyond the scope of this Thesis and there exists a comprehensive array of good publications on the topic, starting with [69][70]. In this chapter, three major new pieces of software are described to exploit this new architecture for the general benefit of research in ultrasonic phased arrays, as follows:

- cueBEAM is an implementation of continuous-wave form of the Huygens principle of wave propagation in homogenous media, for simulation of single-way ultrasonic radiation of a phased array probe. Emphasis is placed on supporting any probe geometry, including sparse 2D array probe configurations;
- The computationally expensive problem of finding the path of an acoustic ray refracted by change of the acoustic media it travels through is addressed;
- The problem of general beam forming, known as the Total Focussing Method, is implemented in a form that is extremely computationally efficient. Through its high performance, an entire new area for research is therefore opened. The implementation of this algorithm is discussed in the next chapter.

The ‘Big ‘O’ notation for computation cost

In the field of numerical computing, the complexity or cost of an algorithm is often presented in the ‘Big ‘O’ notation, where ‘O’ symbolises a fundamental and atomic operation to be performed. The complexity of the algorithm is then the number of the ‘O’ operations required to be performed to complete the program.

For the purposes of this thesis, unless otherwise noted, ‘O’ is an operation of ray tracing a single ‘path’ of sound energy traveling between two spatial points, accumulating in a specific value relevant to the problem at hand. As detailed later in this chapter, the overall cost of beam forming computation is closely related to this basic operation, performed millions of times.

6.2 The cueBEAM acoustic field pressure calculation package

6.2.1 Introduction

Phased arrays achieve their beam forming performance through the phenomenon of constructive and destructive interference of waves which are emitted by the individual elements. It is useful to predict the properties of the acoustic beam produced as a step in designing the array. The theoretical basis for doing so is well-known and there exists a

number of software packages to execute the necessary computations, e.g. Zetec UltraVision [Snoqualmie, WA, USA], CIVA [CEA LIST, Saclay, France],[71]

Commonly across all of the software packages seen by the author, predicting the beam characteristics is a relatively time consuming process.

A part of the research thrust for this Thesis is the development of sparse arrays. It has been identified that simulating the beam shape for a wide range of possible sparse configurations of the probe can yield insight into their operational performance; therefore, there is a need to conduct a large number of these simulations quickly.

6.2.2 *cueBeam*

cueBeam is an implementation of the Huygens-Fresnel principle of waveform propagation through a homogenous media. This approach yields itself particularly well for execution on GP-GPU processors, as explained later within this chapter in section 6.2.3.

Governing equations

Consider a dimensionless, omnidirectional wave radiator emitting a continuous wave at frequency f and amplitude U_0 . In the media of wave velocity v , the radiator produces a disturbance of wavelength λ and wavenumber k . The wave produced can be described at any time t and point \vec{r} in space as a complex amplitude, decaying spherically as a function of distance from the radiator $|\vec{r}|$:

$$U(\vec{r}) = \frac{1}{|\vec{r}|} U_0 e^{i*k*|\vec{r}|} \quad \text{Equation 6.1}$$

Assuming that the medium of interest is a linear, non-viscous liquid and using the principle of superposition, the amplitude of the wave stemming from multiple sources of known amplitudes and phases can be calculated as a complex sum of the contributions.

Similar considerations apply to non-periodic (time domain) signals.

6.2.3 Implementation overview

The software has been written as a MEX [72] plug-in for MATLAB, and utilises the CUDA language[73] for the computation kernel. The algorithm is decomposed in such a way that

each thread of the GPU calculates the pressure value for a single-point in space only. This allows all the computations to progress in parallel, as the pressure in space is completely and solely defined by the properties of the radiators and the media. There is no need for communication between the threads.

The following inputs are required:

- Description of the radiating points: location, amplitude and phase of radiation;
- Description of the propagation media: wavenumber;
- Description of the points in space to calculate the resulting acoustic pressure.

The algorithm can be described in pseudocode, as follows:

```
For each pixel
  Initialize pressure=complex zero;
  Foreach radiator
    Distance=distance(radiator, pixel);
    Phase shift=wavenumber*distance+ ...
    ...radiator phaseshift;
    Amplitude decayed=radiatoramplitude/distance;
    pressure=pressure+...
    ...complex(amplitude decayed, phase shift);
  End for each radiator
  Output(pixel)=pressure;
End for each pixel
```

6.2.4 Performance of CUDA Implementation

The algorithm has been implemented in MATLAB M-script, C, and CUDA. The performance has been measured on Intel Core i7 2.8GHz CPU and NVidia GTX480 GPU by calculating a large (8E+9 pixels) pressure field. The performance can be expressed through a count of acoustic paths integrated per second. A summary of the performance figures are displayed in Table 6.1.

Table 6.1 Summary of performance figures for cueBEAM software

Hardware/ software platform	Performance in millions of paths per second	Speedup vs. reference
Intel Core i7 2.8GHz / M-script	4.99	reference
Intel Core i7 2.8GHz / C-MEX	15.95	3.2x
NVidia GTX480 / CUDA in C-MEX	14935	2993x

Two versions of the base code have been produced. The first is optimised for collecting pressure calculations on a regular, planar square grid. The second is optimised for calculations of points distributed over a hemisphere, with a grid using Lambert equal area azimuthal mapping transformation, as described in Section 3.7.1

For typical problems found in the course of this thesis, the calculation time is virtually instant. This high-performance enables a number of applications.

6.2.5 Applications

The very high-performance of the GPU implementation has enabled two key interactive applications, ArrayEdit and BeamDemo.

A range of static examples of calculated beam profiles using the cueBeam is presented in appendix B.

6.2.5.1 ArrayEdit3: Sparse array editor

A sparse array CAD package has been built to facilitate the exploration of sparse array element layouts. This package allows for the placing and moving of the array elements defined over a triangular grid. The performance of the resulting element in terms of main lobe width and side lobe amplitude is recalculated and updated interactively as the elements are added, modified and removed.

An example screenshot of the editor in use is presented in Fig.6.1 (densely packed array) and Fig.6.2 (sparse array). These two figures exemplify an exploration of the problem space in placing the elements on a triangular grid. In the first figure, the elements that cover six triangles are densely placed. In the second, the elements that cover 12 triangles are sparsely placed. The beam width is narrower for the sparse array, however the side lobe level is higher (the contrast is worse).

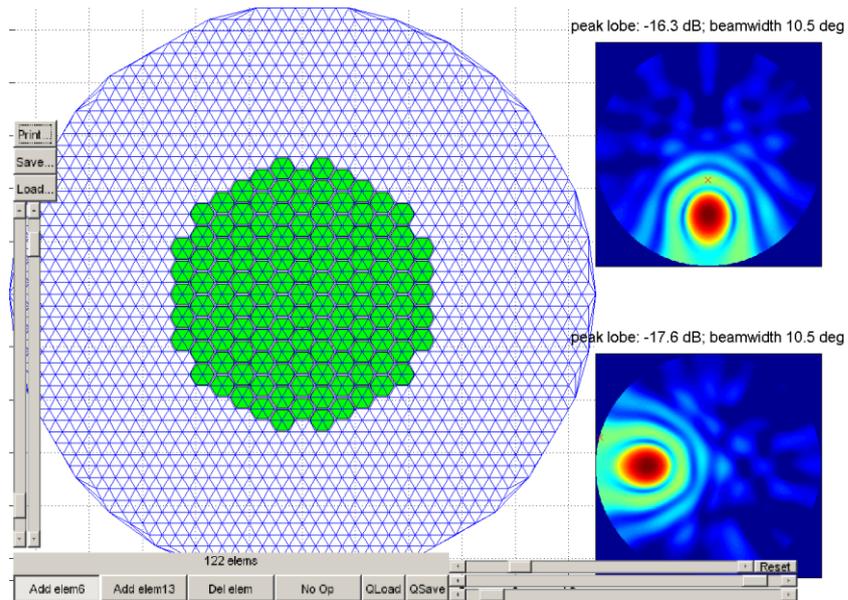


Fig.6.1 Densely packed 122 element phased array layout defined over triangular grid in the ArrayEdit3 software package.

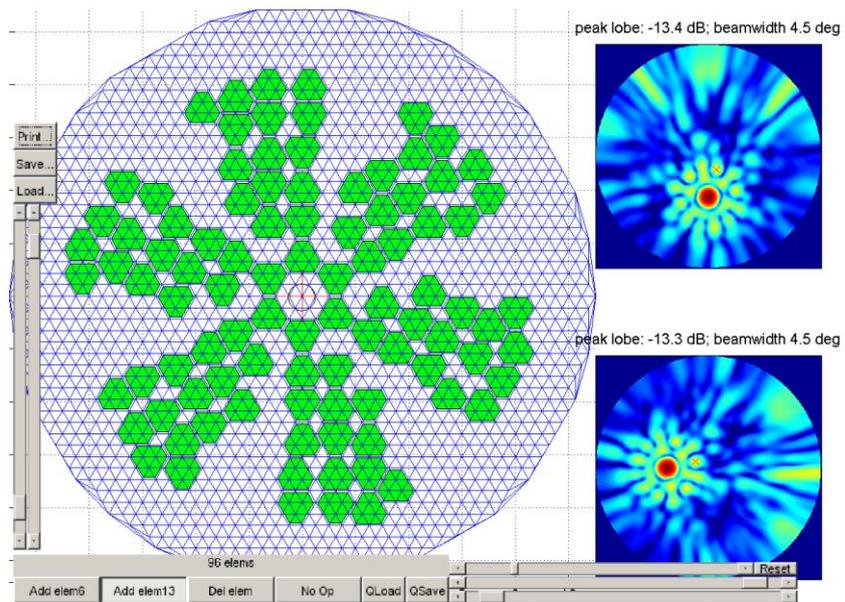


Fig.6.2 Screenshot of the ArrayEdit application with a probe design loaded. Array elements have been manually placed. Note the changed performance figures compared to the case in Fig.6.1: narrower beam width and higher side lobe amplitude for the sparse array case.

6.2.5.2 *BeamDemo: Interactive beam forming and focussing demonstration*

A software package has been built to interactively demonstrate the effects of changing the probe aperture, bandwidth, element count and the steering and focussing on the beam shape and image quality metrics. An example screenshot is presented in Fig. 6.3.

The scene settings and the probe itself are defined by means of a Matlab script.

The three main displays are interactively updated and illustrated in Fig. 6.3. The XZ cross-section display is a 2D slice of the 3D space defined by plane $y=0$. The hemisphere display is a 2D slice of the 3D space, defined by a hemisphere of radius as dictated in the settings. Both of these displays present the calculated relative amplitude of the pressure field generated by the probe. They are also used to visualise the location of a curve (semicircle, white in the figure) that is used as a final 1D cross-section of the beam. The pressure sampled at the location along this curve is displayed in the bottom plot.

Based on the 1D cross-section, example beam characteristics are calculated: main lobe width (planar angle, width of the -3dB envelope) and leakage factor (ratio of the energy outside the main lobe to the total signal energy in the measurement plane).

The user can click the XZ or hemisphere images to direct and focus the beam at that point, instantly updating the display. Keyboard shortcuts are used to modify probe characteristics, such as to make it bigger or smaller or to change the radiation frequency.

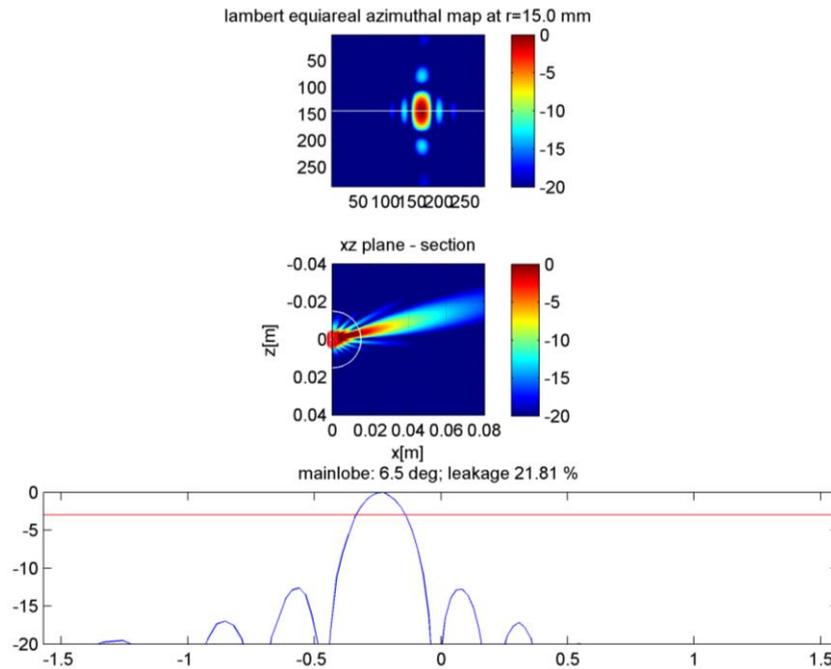


Fig. 6.3. Example screenshot of *cueBeam* demonstrator. Top: two 2D cross-sections through 3D space; Lambert equiareal mapped hemisphere (top) and XZ plane (middle). Bottom – 1D cross-section of the beam along the white curve presented in the top views. From this last view, two beam characteristics are instantly calculated: main lobe width and energy leakage factor.

6.3 Computationally efficient solution of Snell's Law of refraction

6.3.1 Introduction

The success of the *cueBeam* has prompted the following question: Can the simulation package be extended to include the effects of refraction?

Snell's Law, or the law of refraction, is a formula used to describe the bending of waves as they pass through the interface between media of differing wave velocity. In the case of acoustic waves in a solid media, the formula can be used for longitudinal, shear and mode-converted waves. The concept of a ray is used here to visualise the shortest propagation path between two given points.

For this work, (x_1, y_1) , (x_2, y_2) are the coordinates of (1) the source and, (2) the destination point of the ray respectively, (x_i, y_i) is the ray entry point on the interface between the two media, where y_i is assumed to be 0 (meaning, planar surface) and x_i has to be calculated. φ_1 and φ_2 are the angles between the normal surface and the impinging ray. Finally, v_1 and v_2 are velocities of respective waves in their respective media of propagation. The symbols are visualised in Fig. 6.4 and used in the following Equation. 6.2:

$$\frac{\sin(\varphi_2)}{\sin(\varphi_1)} = \frac{v_2}{v_1} \quad \text{Equation. 6.2}$$

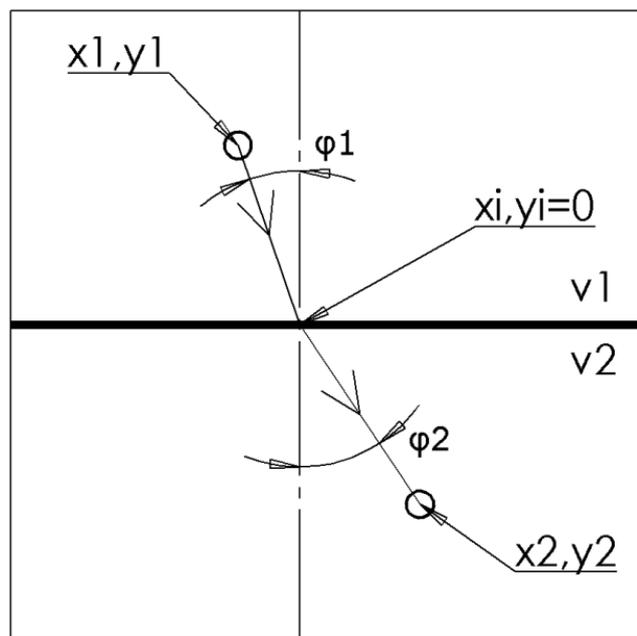


Fig. 6.4. Depiction of symbols used in the application of Snell's Law.

The challenge is to find the point where the ray refraction occurs. This is distinct from solving the 'direct' problem of calculating the refracted angle of the ray, where the location of the source and point of refraction are given.

The classical method of solving this problem is by taking an iterative approach. The procedure assumes a starting ray refraction point on the media interface, and then using the Fresnel principle, iterates it towards the minimum time of flight condition. Such a solution method is very easy to implement as it is numerically stable, supports curved surfaces and to speed up convergence it can utilise various small-scale optimisation

methods such as the Newton–Raphson method [74]. Iterations can be continued in order to achieve the desired precision. However, such an algorithm suffers from the fact that it requires a conditional loop whose length varies on a case-by-case basis. This means that on parallel hardware architecture like GP-GPU, each cluster of threads is - at best - only as fast as the slowest thread, or even slower due to serialisation. Additionally, conditional jumps incur a time penalty that is worth many floating point operations, depending on the processor architecture. This inefficiency has prompted the author to pursue an alternative algorithm, in which the equations are analytically solved for the refraction point and the result is thus obtained in a single-pass without iteration. It has been speculated that such an approach, while yielding a long instruction list, may prove more efficient due to the avoidance of conditional jumps and thread divergence, thus utilising the advantages of parallel hardware architecture and avoiding its known disadvantages.

6.3.2 Direct solution for the ray entry point (ray refraction point), 2D case

Consider a 2D case of ray refraction, as depicted in Fig. 6.4. A ray is emitted from point $P_1(x_1, y_1)$ placed in medium 1 with wave velocity v_1 , towards point $P_2(x_2, y_2)$ placed in medium 2 with wave velocity v_2 . The point of refraction $P_i(x_i, y_i)$, lying on the interface, is to be determined. The interface is assumed to be a planar surface described by equation $y_i = 0$. This simplification is realistic. For example, in the case of TFM imaging (see section Chapter 7) it can be assumed that the probe and image created is oriented with respect to the interface. The ray passes through the medium interface at point $P_i(x_i, y_i)$, the spatial location on which all remaining variables are dependent. The angle φ_1 can be represented in this scenario as:

$$\sin \varphi_1 = \frac{x_i - x_1}{\text{distance}(P_1, P_i)} \quad \text{Equation. 6.3}$$

and similarly,

$$\sin \varphi_2 = \frac{x_2 - x_i}{\text{distance}(P_i, P_2)} \quad \text{Equation. 6.4}$$

Combining Equation. 6.3 and Equation. 6.4 yields the following:

$$\frac{\left(\frac{(x_i - x_1)/\sqrt{(x_i - x_1)^2 + y_1^2}}{(x_2 - x_i)/\sqrt{(x_2 - x_i)^2 + y_2^2}}\right)}{v_2} = \frac{v_1}{v_2} \quad \text{Equation. 6.5}$$

This is a double-quadratic equation that can be solved analytically for x_i using a Computer Algebra System [75]. The solution and a detailed method of obtaining one is provided in appendix A. There are four solutions for x_i , of which two always contain an imaginary component and correspond to a non-physical situation of negative wave velocities. Out of the other two solutions, one is always valid for $x_2 > x_1$ and the other is valid for $x_2 < x_1$. There is also a special situation when $x_2 = x_1$; in such case, it is easy to see that the solution is $x_i = x_1 = x_2$.

The downside of such a direct approach is that it is poorly conditioned for some input values, as demonstrated in Fig. 6.5. This simulation result shows that when substituting single-precision numbers into the equation, a noticeable round area runs into a numerical singularity (a division by a very small number at a point in the process), and returns no solution where there should be one, as in Fig. 6.6. Interestingly, this singularity appears to be related to the spatial location at the point where the critical angle of refraction occurs.

The solution is also unstable for cases where angle $\varphi_1 \approx 0$.

This problem is alleviated, but not resolved, when using double-precision numbers. However, double-precision arithmetic is approximately two to eight times slower, which defeats the purpose of using the direct algorithm. A mixed single-double precision approach could be used [76], however as the affected area is small in comparison to the expected range of inputs required to build a typical ultrasonic image, a reasonable choice is to use the iterative solver for these singular cases. The iterative solver is triggered upon detection of the singularity at the vulnerable point inside the algorithm, so that the performance impact on other refraction cases is minimal.

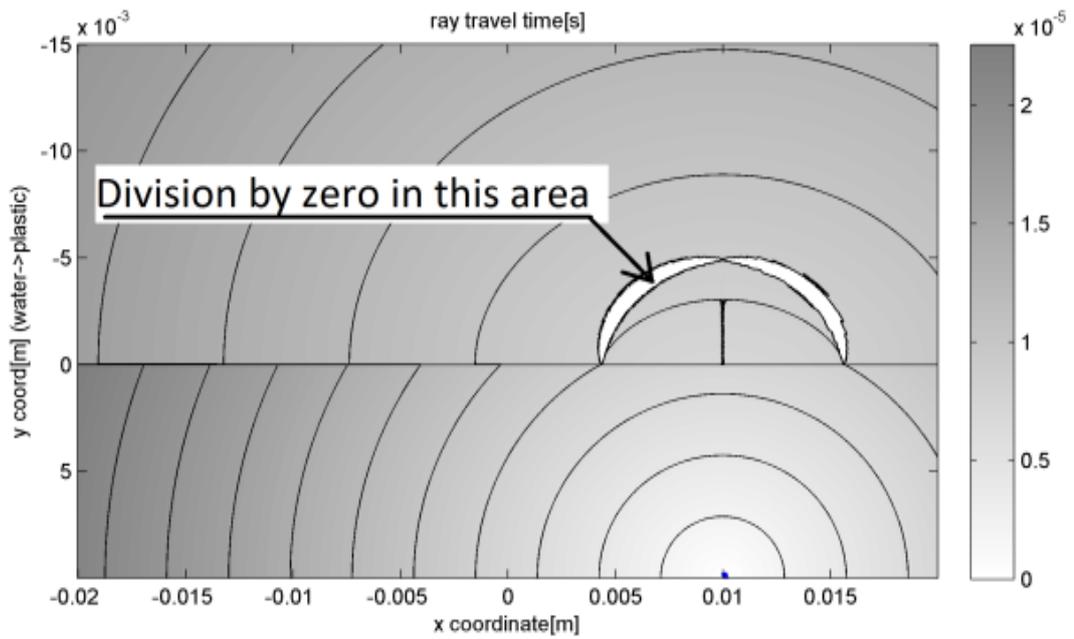


Fig. 6.5. Ray travel time from a point at (0.01,-0.02) in medium 1 (the lower part of the figure) to all other points in space. There are numerical singularities in the direct solution when calculated using single-precision numbers.

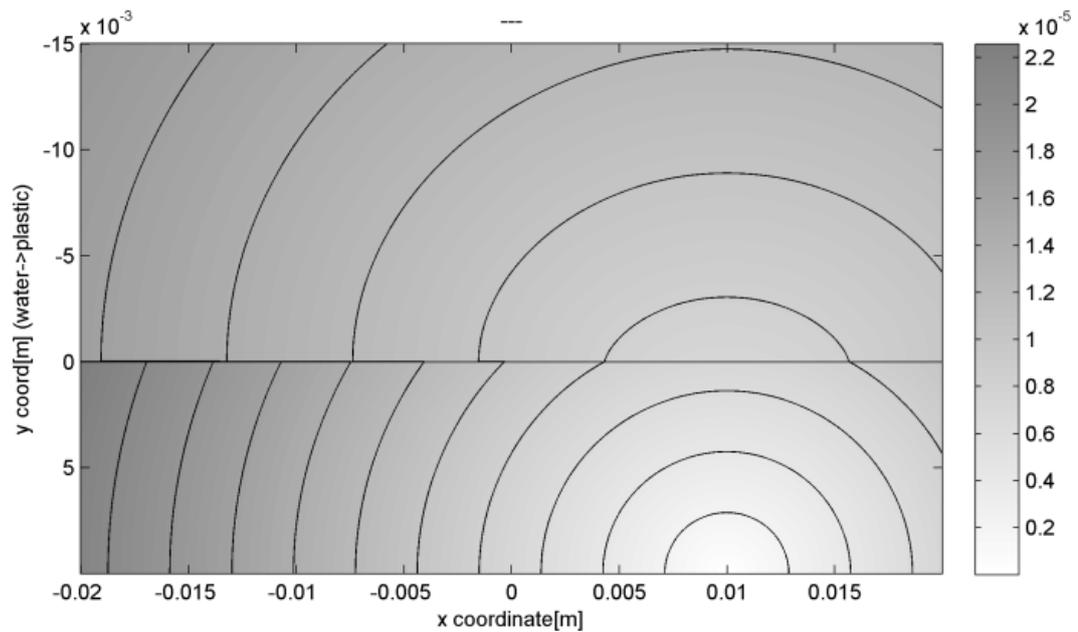


Fig. 6.6. Example solution as in Fig. 3, but with a corrected algorithm - there are solutions for all possible locations and no singularities are present. This correction is obtained by switching to an iterative algorithm, only for the point where singularity is detected.

Note that a critical angle of refraction effect is visible, as there is an area on the interface where travel time is continuous and at one point it becomes discontinuous.

6.3.3 Solution for the ray entry point (ray refraction point), 3D case

For a plane interface, laws of refraction state that the incident ray and the refracted ray lie in the same plane. The entry point must lie in the only plane that is normal to the interface and contain P_1 and P_2 . This property can be exploited to extend the previously derived equations to the 3D case. This process can be presented as follows.

With reference to Fig. 6.7, a helper coordinate system 'h' is introduced, which is translated by vector $(x_1, 0, z_1)$ and then rotated around y-axis from the origin so that the points P_1 , P_2 and P_i lie on the same plane $z_h = 0$;

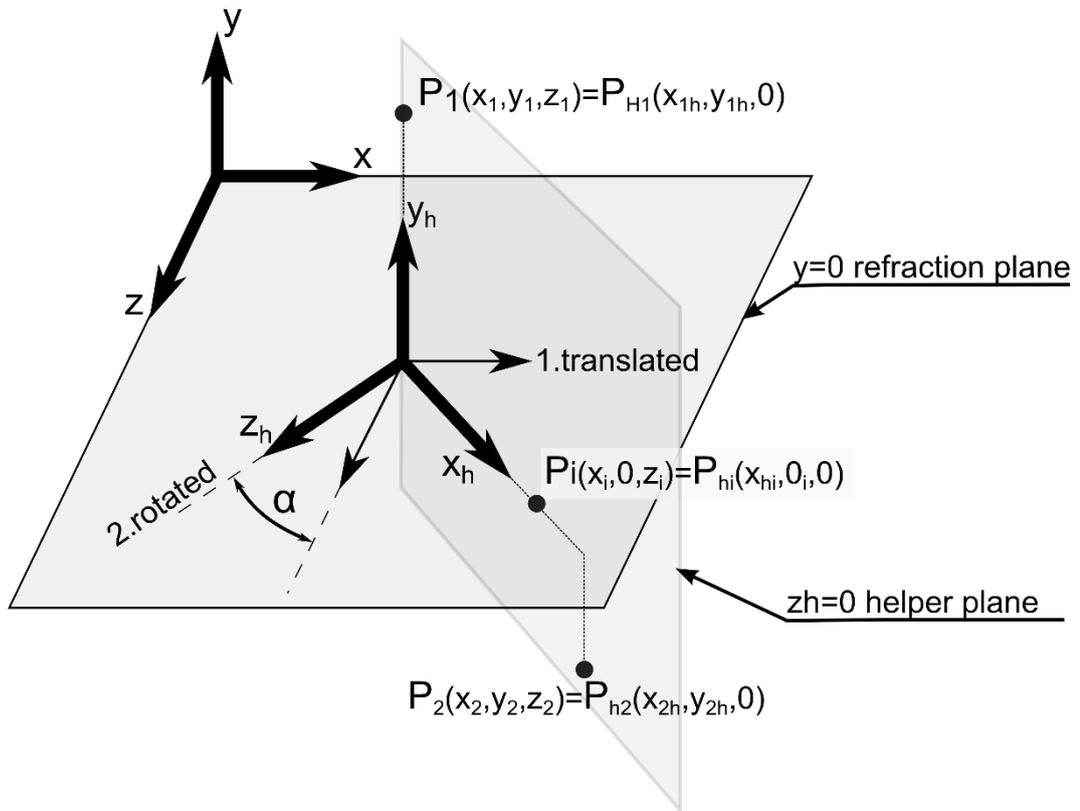


Fig. 6.7. Points P_1 , P_2 , P_i of the 3D refraction case and their respective helper points, P_{1h} , P_{2h} , and P_{ih} . Note the rotated coordinate system.

The point coordinates in the helper system become:

$$P_{1h}(x_{1h}, y_{1h}, z_{1h}) = (0, y_1, 0)$$

$$P_{2h}(x_{2h}, y_{2h}, z_{2h}) = (\sqrt{(x_2 - x_1)^2 + (z_2 - z_1)^2}, y_2, 0)$$

Equation. 6.6

In the helper coordinate system, consider the refraction point $P_{ih}(x_{ih}, y_{ih}, z_{ih}) = (x_{ih}, 0, 0)$. The only unknown remaining is the x_{ih} coordinate. Crucially, it can be found using the 2D case algorithm, as described previously in section 6.3.2.

Finally, the point P_{ih} is transformed back to the original coordinate system to become point P_i :

$$P_i(x_i, y_i, z_i) = (x_1 + x_{ih} * \cos \alpha, 0, z_1 + z_{ih} * \sin \alpha) \quad \text{Equation. 6.7}$$

Where $z_{ih} = 0$ due to the assumption of it lying in the plane $z_h = 0$

This rotates and then translates the refraction point P_{ih} back to the original coordinate system. Angle α is the reversed angle of the original rotation. It can be found as the angle between the x axis and line between P_1 and P_2 projected onto the $y = 0$ surface.

The accuracy of the result can be checked by evaluating Snell's Law using the point P_i as parameter, to find the angles, φ_1 and φ_2 .

Due to the coordinate transformation, the case-selecting condition $x_{2h} \geq x_{1h}$ is always true, so no conditional jump in the CUDA code is necessary. This means that all execution units of the GPU always follow the same code path, thus maximising parallel performance.

6.3.4 Implementation

6.3.4.1 New direct solver

Compiling the direct solution of Snell Laws equation into a machine code yields approximately 150 operations including addition, multiplication and power. The automated compiler is optimised for Wolfram Virtual Machine [75], and utilises complex-valued accumulators. Hence the code must be adapted for execution on real hardware. Most accumulators in the code do not need to be complex, but the remaining few must be expanded to become real-valued. There are six complex functions (other than add or multiply) of complex values. The code has been hand-modified to make it compatible with ANSI C and hardware execution units available in GP-GPU.

Additionally, the following optimisations that the original equation-to-instruction compiler doesn't do have been manually applied: folding constants, re-using accumulators and shuffling the order of operations to avoid read-after-write delays. The modified code contains approximately 350 real-valued operations in total. There are a minority of trigonometric functions used to handle complex functions of complex numbers, such as complex square roots, logarithms, inversions and powers.

Despite the high instruction count, register re-use is high. The implementation requires a modest 23 explicit registers, enabling integration with imaging algorithms without register spill or stack usage. Even assuming a requirement of 80 registers, in a GP-GPU with CUDA capability 2.0, there are 1024 registers available per CUDA core, meaning more than 12 threads can be run at the same time without register spill, filling the execution pipeline and also avoiding read-after-write delays altogether. Shared memory is not required and can therefore be configured as L1 Cache. All this forecasts well for the total algorithm throughput, however quirks dependent on the actual GP-GPU architecture make it very hard to predict the final performance without experimentation⁵[77].

6.3.4.2 *Reference Newton iterative solver and simplex iterative solver*

In order to compare fairly the performance of the new algorithm, a suitable baseline code is required. To provide a fair comparison, a hand-tuning effort (similar to that applied to the newly proposed algorithm) of the baseline, reference algorithm will be undertaken.

When deciding on the solution method for the iterative solver, it is important to keep the register pressure low. For this work, two algorithms have been selected for comparison; Newton-Raphson Iterative solver and simplex iterative solver.

Newton-Raphson Iterative Solver

The Newton-Raphson method has been implemented according to a classic formula:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

Equation. 6.8

⁵ The same paper, however, provides some interesting insights for operation order optimisation, for example to prefer $(\text{rsqrt}(x))*x$ operation over a solo $\text{sqrt}(x)$ for greater throughput where x is known to not be small.

Where x_{n+1} is n+1th approximation of the true value, based on n-th approximation. The refraction Equation. 6.5 has been rearranged as follows:

$$\frac{\left(\frac{x_i - x_1}{\sqrt{(x_i - x_1)^2 + y_1^2}}\right)}{\left(\frac{x_2 - x_i}{\sqrt{(x_2 - x_i)^2 + y_1^2}}\right)} - \frac{v_1}{v_2} = 0 \quad \text{Equation. 6.9}$$

Again, the value x_i is sought. The term Δx_i required by Newton-Raphson method is:

$$\Delta x_i = \frac{f(x_n)}{f'(x_n)} \quad \text{Equation. 6.10}$$

This has been obtained and optimised using algebraic manipulation. The same methodology as in the direct case has been used in the manipulation of the formula. Terms independent of x_i are precomputed before iteration, and the remaining operations are optimised to reduce the operation count. Such an approach yields 21 floating point operations in the pre-calculation step and 35 floating point operations per iteration during the iteration phase. Importantly, no complex numbers are used and apart from simple operations, just one single square root operation is used. Thus, an improvement in computational performance can be expected.

The iteration starting point is chosen to lie in the middle of the two ray end points, at $x_i = (x_1 + x_2)/2$. Tolerance on the resultant error is tied to the required sampling rate of the ultrasonic signals of interest. Iteration stopping criterion is that the iteration step value (Equation. 6.10) is smaller than the specified tolerance value.

Simplex Iterative Solver

The simplex method [78] has been implemented to solve the refraction problem directly from Fresnel's principle, that is, by minimising the ray travel time t:

$$t = \frac{\sqrt{(x_i - x_1)^2 + y_1^2}}{v_1} + \frac{\sqrt{(x_2 - x_i)^2 + y_2^2}}{v_2} \quad \text{Equation. 6.11}$$

In order to give a fair comparison, a library function from MATLAB has been rewritten in C and again, hand-optimised to achieve maximum throughput⁶. When compiled as CUDA

⁶ In the course of this work, a flaw has been found in the original Matlab's implementation of *fminsearch*, as detailed in appendix B

code, it incidentally utilises the same number (23) of registers as the direct algorithm. The stopping criterion is the same as in previous section.

The simplex method cost is approximately 50 instructions per iteration, so more than the algorithm from previous section. In the test vector, it requires between eight to 50 iterations to converge to a satisfactory result, depending on the case. The problem with this algorithm is that there are multiple execution paths (conditional jumps) which are not loops, meaning each iteration has a high chance to stall the processor's execution pipeline for many cycles until the correct branch of the code is loaded into the instruction cache.

6.3.5 Benchmarking the code

Throughput of the calculations was measured as a number of calculated propagation times, (ray delays) per unit of time. Rays were shot from a point in medium 1 to all points in an image, as illustrated in Fig. 6.6; except for that only the part of the image where refraction occurs has been used for benchmarking. In order to obtain stable results, 2^{30} representative rays were used in each case.

The hardware platform is an Intel i7-920@2.8GHz, 12GB RAM (3-channel DDR) for CPU C compiled code, and a single NVidia GeForce GTX480 for CUDA code. GPU data upload and download time is included in the measured execution time, and excluding it would favour GPU further. It can be anticipated that the refraction solver is only a part of a bigger algorithm where the data upload time could be ignored; such case is not considered in this section.

The throughput of various algorithms is presented in Table 6.2. The proposed implementation is over 3x faster on CPU, and over 6x faster on the GPU when compared to the simplex solver on the same platform. The latter is an effect of the differences in the CPU and GPU architectures, with the new implementation specifically intended for being compatible with GPU architecture.

Table 6.2. Comparison of the throughput of selected implementations of algorithms.

Algorithm - implementation	Measured Throughput [M paths /s]	Throughput, calculations only [M paths/s]	Speedup vs. same hardware	Speedup vs. reference
	Incl. data transfer time	Excl. data transfer time		
Simplex solver - C	0.42		Ref. A	Ref.
Newton - C	0.18			
Direct solver - C	1.33		3.1x Ref. A	3.1x
Simplex - CUDA	86.0	314.	Ref. B	204. x
Newton - CUDA	42.9	47.		
Direct - CUDA	411.8	2091.	6.7x Ref. B	980. x

6.3.6 Accuracy of the solution

To estimate the accuracy of the fast algorithms, an absolute difference of the relevant section (area where refraction does occur) of the images as illustrated in Fig. 6.5 was taken. The maximum time difference measured as a ratio between the difference of the new and reference results to the reference ToF was 0.0035%. This is close to the theoretical accuracy of the single-precision floating point number representation (24 bits used to represent the significand[79], corresponding to six decimal digits). The difference is due to various rounding effects and approximations of complex functions in the CUDA implementation. Such differences normally correspond to a small fraction of a sampling time of the source data and are therefore completely acceptable. Should the errors become too large, there is an option within the CUDA compiler to disable these approximations at a cost of an approximate 20% drop in performance, or to implement the use of double-precision numbers.

6.3.7 Remaining notes

It has been found that the Newton Iterative solver is not appropriate for the application of a computationally efficient implementation of Snell's Law. It can converge to the solution very quickly for some of the test rays, achieving up to 1.8M paths/second in some cases. However, for some ray paths it converges very slowly or does not converge at all, producing an invalid result and averaging at less than 200k paths/second for a reference ray set. This is because the Newton method utilizes derivative of the problem function, which happens to be poorly conditioned (very small slope) near the correct solution, causing large overshoots.

It is possible to increase its stability by introducing internal problem scaling. However, such an approach results in even slower convergence for these 'problematic' rays. It is also possible to use a hybrid approach as in [74]. Yet upon detection of slow convergence, the time has been already spent. Therefore it has been left out of the final comparison in favour of the simplex method, which reliably produces results for all test cases and achieves a much higher average rate.

Interestingly, the measured loss of compute performance with the Simplex method due to data-dependent execution path in the GPU can be given as Ratio of ratios of respective algorithms on the GPU and CPU (assuming that CPU exhibits no branching loss): $(86.0 / 42.9) / (0.42 / 0.18) = 0.859$; This means that approximately 14% of the compute performance of GPU is lost due to the branching in the simplex algorithm.

6.3.8 Case study

The algorithms described below have been implemented and used in "cueBeam" kernel. Based on this kernel, two applications have been released for public use.

- **cueBeam R4**

cueBeam R4 has been crated for exploration of the shape and characteristics of single-way beam. The probe layout is defined in Matlab script, and then in runtime, user can move the beam focus to any point on the spherical cross-section or planar cross section of the media space. Effects of probe sparsity and signal bandwidth can be interactively explored, as operation is low-latency on an entry level GPU-enabled laptop (2015). The white lines

crossing the blue background in Fig. 6.8 show the spatial location of the other respective cross-section.

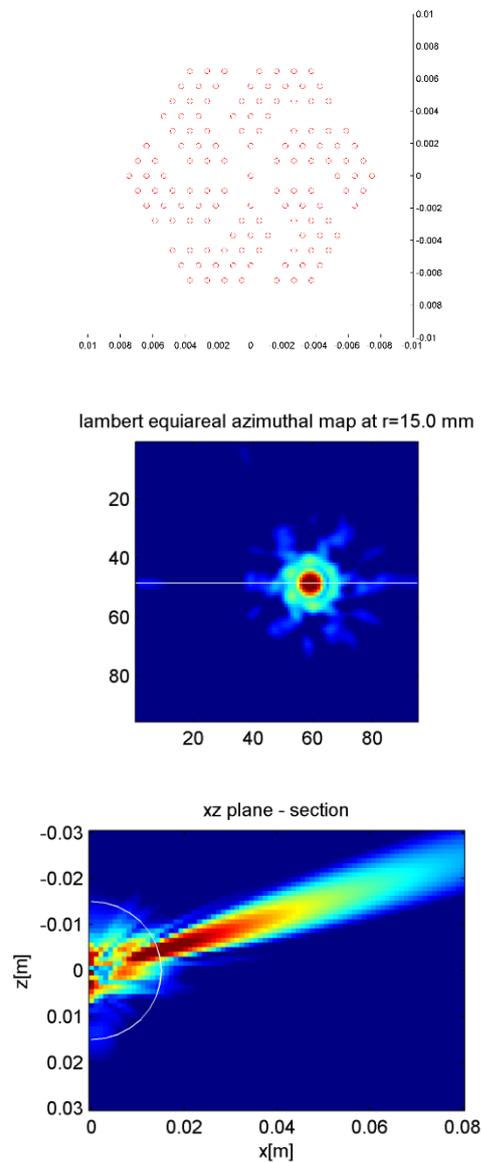


Fig. 6.8. Screenshot of cueBeam R4 in operation. Top: simulated probe element layout. Middle: sphere section. bottom: planar section.

- **cueBeam R6.**

In this public release, the core cuBeam calculation engine has been encapsulated in a series of scripts to aid 1D linear phased array design. Beam characteristics such as beam width, side lobe level, leakage e.t.c. are calculated in batch process. Additionally, 3D visualization of the probe aperture and beam shape has been implemented. Numerical results are stored

in a self-describing Matlab structure for further analysis and automation. Example outputs from the package are presented in Fig. 6.9.

This version has been utilized as a learning aid for the EPSRC “Ultrasonic Transduction” course hosted by the University of Strathclyde for EngD students.

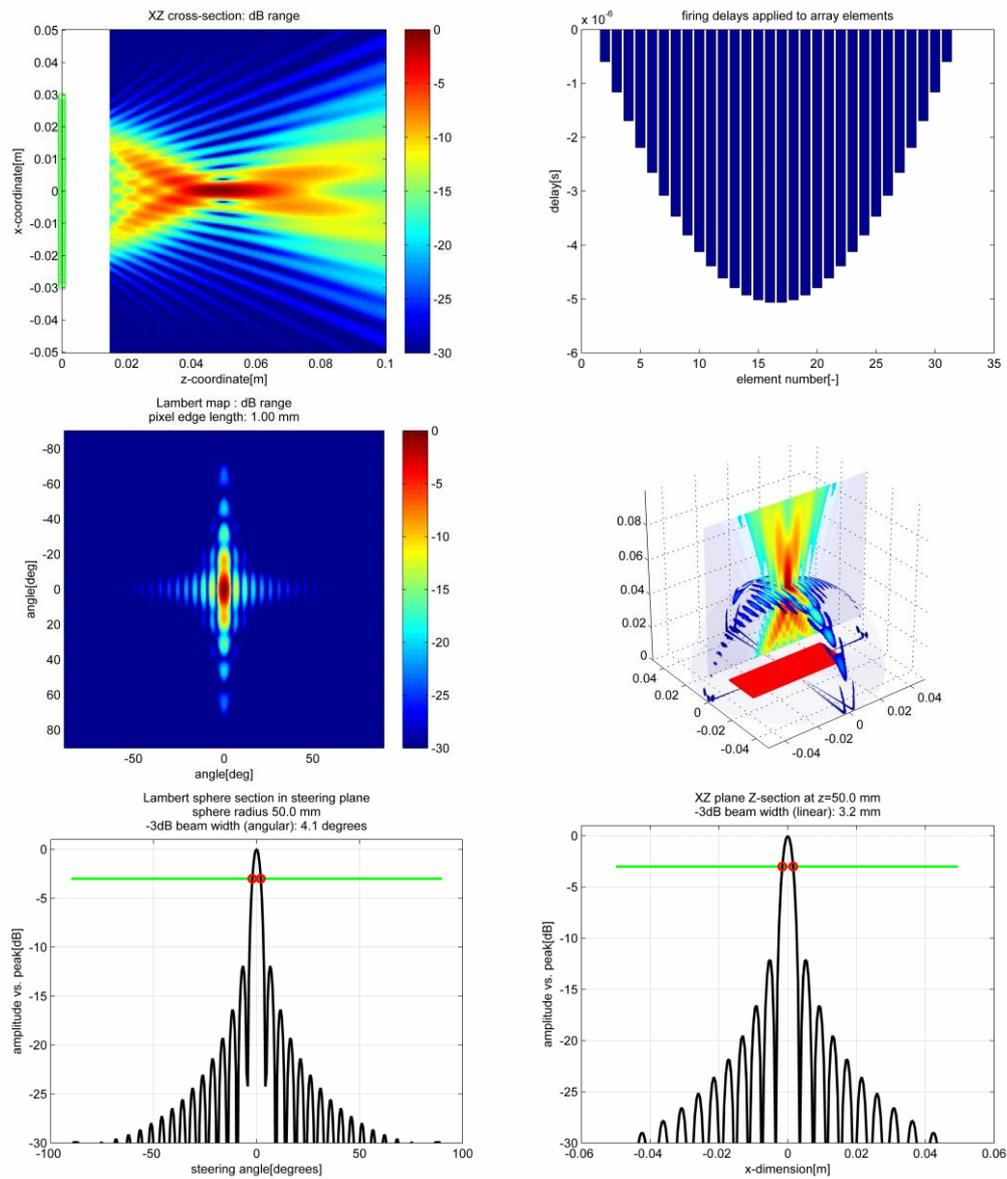


Fig. 6.9. Example outputs of cueBeam R6 for a linear phased array probe.

Chapter 7. GP-GPU accelerated ultrasonic signal processing

7.1 Motivation for work

At the time of this research, a compute-intensive ultrasonic beam forming method, the Total Focussing Method [80] (TFM), was beginning to receive attention from the research community. This method is an extension of the classic delay-and-sum beam forming in that it utilises delay laws both in transmission and reception, which are adjusted for each pixel of the image. This method offers improved and more consistent image quality, signal-to-noise ratio, detection threshold, and defect characterisation capability.

Yet a chief issue preventing its widespread adoption is its computational expense.

An effort was undertaken to implement the TFM algorithm for execution on GPU. The premises for success are that the TFM algorithm yields itself to easy parallelisation and that its compute-dense nature (many calculations, little data) makes it well suited the GPU architecture.

The intended application was to enable research in the final ultrasonic image quality, as produced by sparse phased arrays combined with the TFM signal processing.

7.2 Implementation of the TFM algorithm, 1st generation

7.2.1 Introduction.

For each pixel of the image in the TFM imaging, the following effective operations must be performed:

For each transmitter-receiver pair we must; a) calculate Time of Flight (ToF) of wave along the path between the transmitting element, through the refracting interface, to the pixel and back to the receiving element; and b) accumulate the echo value from respective A-scan memory.

For each pixel, the intensity of the pixel is defined as:

$$I(pixel) = \sum_{tx=1}^{ntx} \sum_{rx=1}^{nrx} A_{tx,rx}(ToF(tx, pixel) + ToF((pixel, rx))) \quad \text{Equation. 7.1}$$

Where:

- tx –transmitting element
- rx – receiving element
- ntx – count of transmitting elements
- nrx – count of receiving elements
- $A_{tx,rx}(time)$ – amplitude at specific time taken from the A-scan recorded using tx and rx elements, out of the Full Matrix Capture (FMC) data
- ToF(from, to) – function of Time of Flight between two points in space of interest.

If the probe elements are immersed in the same media as the reflectors to be imaged (in-contact inspection, or no-refraction case), the ToF function is given by:

$$ToF(P_1, P_2) = \frac{1}{v} distance(P_1, P_2) \quad \text{Equation. 7.2}$$

Where P_1, P_2 – starting and ending points of the straight path in the medium with an acoustic velocity v .

If the points P_1, P_2 lie in media of different acoustic velocity, the computation becomes more expensive and a refraction algorithm, as presented in section 6.3, is needed to calculate the ToF required for each pixel/probe element combination.

7.2.2 The distance function cost

2D case, homogenous environment

The distance function in Cartesian space is given by the Pythagorean Theorem:

$$distance(P_1, P_2) = \sqrt{(P_{2x} - P_{1x})^2 + (P_{2y} - P_{1y})^2} \quad \text{Equation. 7.3}$$

Where P_1, P_2 – points in homogenous, Cartesian space; subscripts x, y, z – coordinates of a given point. Point P_1 is typically associated with the ultrasonic probe element (transmitter or receiver) and point P_2 with a pixel of the image built from the signals received by the probe.

Equation. 7.3 gives a computation cost of three additions, two multiplications and a square-root per function call (assuming that number negation is free).

If the image space is tied to the probe location, the probe element's y -coordinate can be assumed to always be zero, and therefore the equation simplifies:

$$distance(P_1, P_2) = \sqrt{(P_{2x} - P_{1x})^2 + P_{2y}^2} \quad \text{Equation. 7.4}$$

Where P_1 is the probe element location and P_2 is the image pixel location. This gives a cost of two additions, two multiplications and a square-root per function call (assuming that number negation is free).

3D case

In the 3D case, the distance formula is:

$$distance(P_1, P_2) = \sqrt{(P_{2x} - P_{1x})^2 + (P_{2y} - P_{1y})^2 + (P_{2z} - P_{1z})^2} \quad \text{Equation. 7.5}$$

This gives a cost of five additions, three multiplications and a square-root per distance function call. Again, if the probe element location is fixed to the axis of the coordinate system, the cost is reduced by one or two addition operations.

7.2.3 Overall algorithm cost

From Equation. 7.1 and subsequent analysis, it is clear that the compute effort required is closely tied to the number of unique sound rays or 'paths', which have to be considered to create the entire image. Here, 'path' is an imaginary path that an acoustic impulse, generated by the excited phased array element, has to take in order to reach the location

of the given pixel of the final image, and then back to the receiving element, as illustrated in Fig. 7.2.

If the desired 3D image consists of $n_x \times n_y \times n_z$ pixels, the overall computation cost is $n_x \times n_y \times n_z \times O$, where O is (2 paths + FMC data load + pixel accumulate). For example, for a 256^3 pixel image and 128 element probe, the computation cost in this case is $O \times 1.0737e+09$.

A range of optimisations could be introduced to the basic algorithm. For example, it could be based upon exploiting various symmetries that arise if the image pixel pitch is consistent with the probe element pitch [81]. However, such optimisations are not viable here since the intention is to make the algorithm simple, flexible, and capable of sparse array imaging.

7.2.4 Implementation benchmark, 1st generation

The basic form of the algorithm has been implemented in MATLAB (The MathWorks Inc., Natick, MA, USA, 2015), C-Mex[72] and CUDA[73], and the results are collated in Table 7.1. This table gives the performance of the algorithm when executed on the various hardware/software platforms. Case one is for information only, as Matlab script is an interpreted language and as such performs poorly with such algorithms. The case two, the C-Mex case, is the algorithm compiled for execution on a CPU. This second case is treated as a reference, baseline performance figure. Case three is the code executed on the first generation of CUDA-capable cards. The performance is 21.3x better than the CPU case. Case four is the algorithm executed on the second generation of CUDA-capable cards. It is over 900x faster than the CPU case. All four of these cases do not include refraction code, but only the straight-ray code. The effect of including refraction is illustrated by case five, showing that by taking refraction into account it causes nearly 10x reduction in performance.

In order to verify the equivalence of the versions, a synthetic FMC data set has been created in COMSOL mechanical module and then processed with the above noted TFM implementations. The result is presented in Fig. 7.1. In this image, the white thick lines denote the boundaries of the solid (metal) region. These create reflections for the eight-element probe placed on the top.

In the following evolution implementation, the 3D propagation case has been used, allowing for fully arbitrary definition of probe element locations and image pixel locations, in particular enabling study of sparse arrays. The ToFs are calculated as-needed during the imaging phase instead of being pre-calculated. This is at a cost of reducing the performance when compared to the case-optimised scenario as in [82] and [81].

The improved performance marks a step change for the applicability of the TFM algorithm as it is now possible to execute many practical imaging problems in real-time.

Table 7.1. Comparison of performance of the TFM algorithm on the new hardware/software platform

Case	Hardware platform	Software platform	Performance in 10^3 paths/sec	Speedup vs. reference
1	Intel Core i7 @ 2.8GHz	MATLAB	5.1	----
2	Intel Core i7 @ 2.8GHz	C-Mex[72]	' 486.0	Reference
3	NVidia GeForce 8600GT	CUDA 1.0	10' 386.2	21.3x
4	NVidia GeForce GTX480	CUDA 2.0	455' 897.	938. x
5	NVidia GeForce GTX480	CUDA 2.0 +refraction	48' 184.	99. x

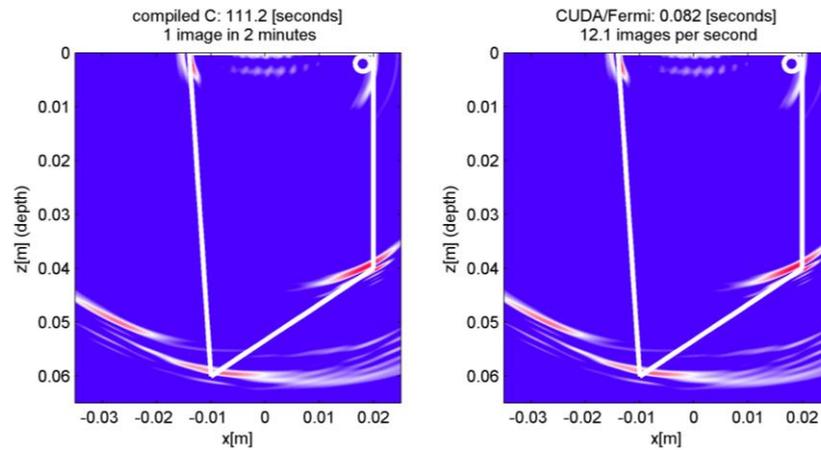


Fig. 7.1. Comparison of performance of TFM algorithm on CPU and GPU platforms, January 2011. The synthetic FMC data (8-element probe) has been produced in COMSOL.

7.3 TFM algorithm: 2nd generation with support of refraction through arbitrary 3D surface

The success of the TFM algorithm implementation as explained in 7.2 and the refraction algorithm in section 6.3, prompted further development of the two to enable imaging with refraction through curved interfaces.

However, instead of solving the refraction equation for curved interfaces, an entirely different approach has been taken here.

In this work, an imaging case is considered where the probe is positioned in, or in contact with medium 1 of wave velocity 1. The waves are sent and received through a 3D curved interface into medium 2 with wave velocity 2 and the image of reflectors in medium 2 is sought, as depicted in Fig. 7.2. The imaging area or volume is defined in terms of the corner point (x_0, y_0, z_0) , the pixel count (n_x, n_y, n_z) and the distance between pixels, (dx, dy, dz) . The refracting surface is described in terms of $z = \text{SurfaceF}(x, y)$ function or height map, allowing an arbitrary surface to be described. The probe is defined by the XYZ location of the array elements.

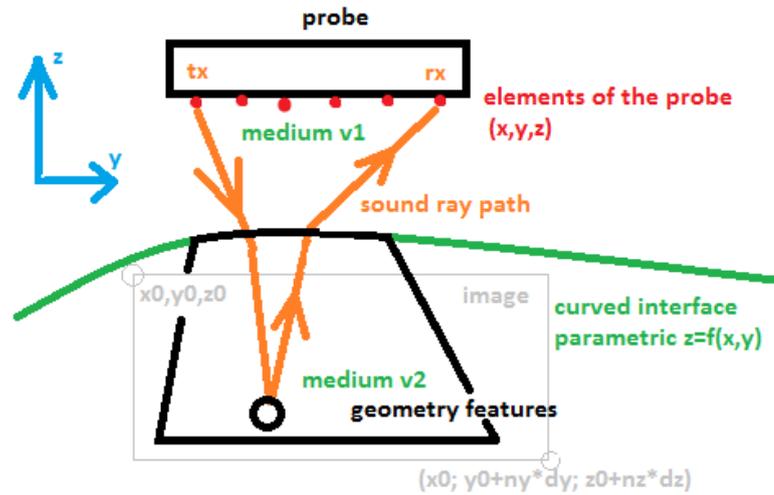


Fig. 7.2. Elements of ultrasonic imaging scenario.

7.3.1 Implementation considerations

In the simplest, 'in-place' implementation, as attempted in section Chapter 7, one would take each pixel, calculate the ToF for all combinations of transmit-receive signals, and coherently accumulate respective A-scan samples from the FMC data set. However, this results in a high computation cost for each image. On the other extreme, one could pre-compute all the ToFs and store them in a look-up table, as in [82]. In such a case, one can compute ToFs once and then re-use them multiple times with different FMC data sets. Yet this results in large memory cost and assumes the location and shape of the refracting interface does not change between consecutive data acquisitions. For example, for a 3D image of resolution of $240 \times 240 \times 512$ pixels, at 4096 A-scans, this means 120×10^9 compute operations, or 120×10^9 entries in the look-up table. Every entry of the look-up table would have to be retrieved multiple times, making memory bandwidth the performance limiting factor.

A range of intermediate solutions, balancing in-place computation, pre-computation, symmetries, calculation decomposition and reordering, pixel alignment considerations and other algorithmic and hardware-related optimisations have been already attempted [74][81][82][83][84], showing how important it is to maximise the performance of imaging algorithms.

The TFM ultrasonic imaging process is an 'embarrassingly parallel' problem, where each pixel and ToF can be calculated independently, and in parallel at the same time as others,

from a single-source data set. However, the order of calculations does matter where calculation performance is concerned, as shown later in this thesis.

Given a relatively simple set of operations required for each pixel, GP-GPU computing cards are a good candidate for the realisation of the ultrasonic imaging process. With their compute-dense architecture, the memory interface bandwidth becomes the limiting factor in utilising large look-up tables. In fact, due to their compute-dense architecture, some results are faster to be re-calculated on the chip, rather than just calculated once, stored, and then loaded from the relatively slow global memory.

When considering implementing an algorithm on the GPU, one should contemplate taking maximum advantage of various subsystems of the GPU architecture. In particular, there are a number of memory subsystems, varying in functionality, bandwidth and latency and special function units, like texturing units, that can work in parallel with the main Streaming Processors (SP) as illustrated in Fig. 7.3.

In a single GPU chip, there are a number of Streaming Multiprocessors. Multiprocessors which are mostly self-contained in that:

- they share the Global Memory with other Multiprocessors through the global memory bus only,
- they do not share any other compute resources with other multiprocessors, in particular, “shared memory”, register file and so forth,
- there can be a functional graphics card with one or more multiprocessor,
- there can be up to 16 multiprocessors per GPU chip in the Fermi architecture, more in the newer architectures.

In each Multiprocessor, there are a multiple of cores, the number of which depends on the generation of the architecture. For Fermi, this is 32 cores per multiprocessor, 192 for Kepler and 128 for Maxwell. Apart from the cores, there is a register file module, shared memory module, code cache/constant cache module, L1 cache, and texturing units. Again, the exact amount of these resources depends on the iteration inside the given generation of the architecture. These modules are bound together with interconnects that exhibit their

specific properties, for example values can be broadcast from the shared memory into all the cores in a single access cycle.

It is proposed that creating an algorithm which takes advantages of these specific subsystems will lead to dramatically improved performance of the overall system.

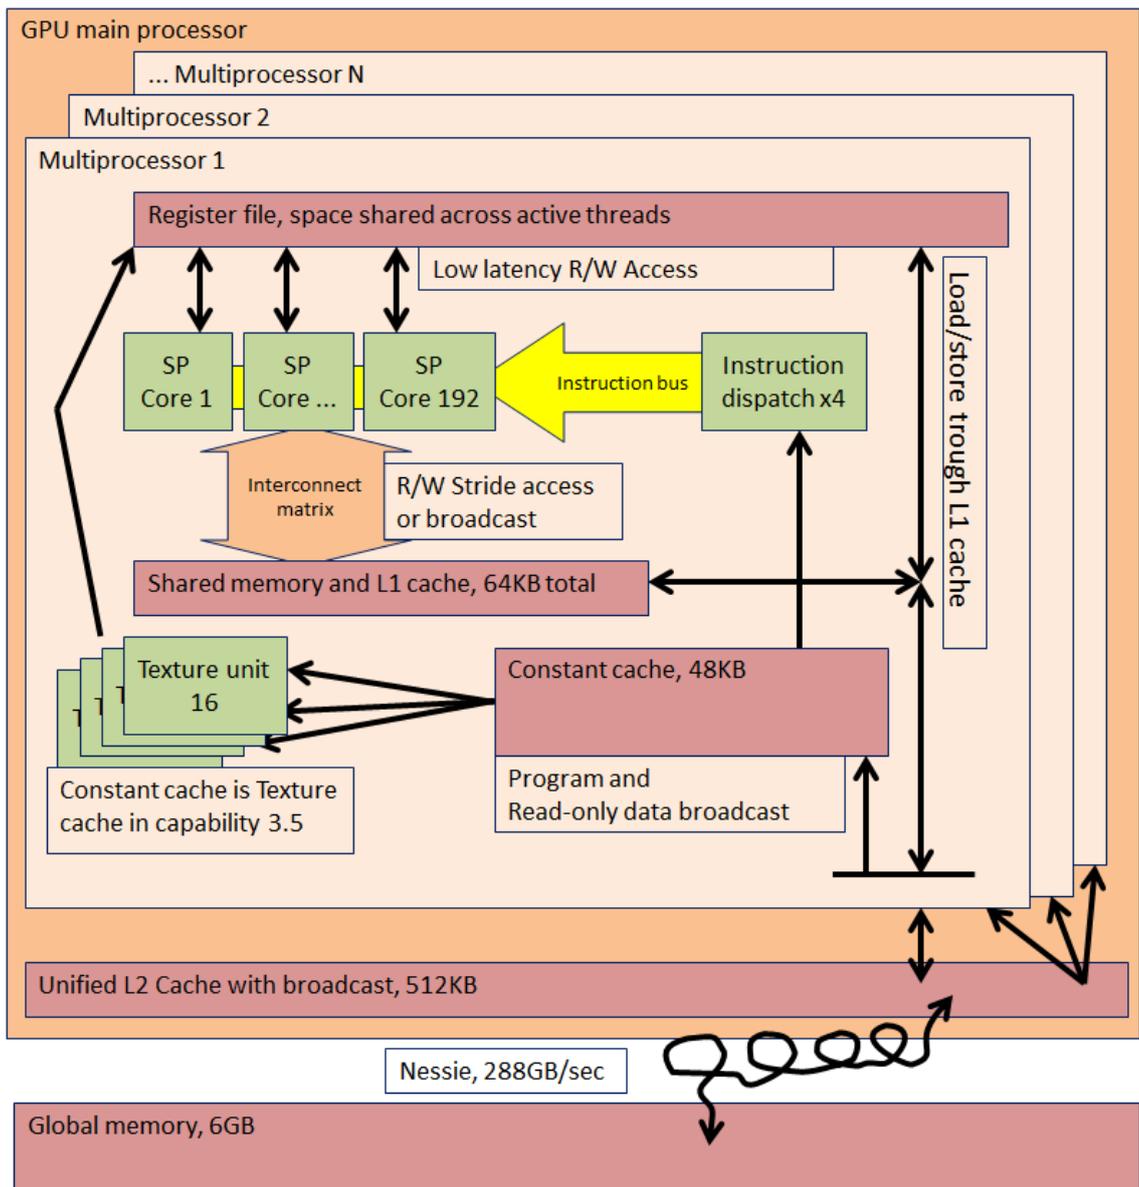


Fig. 7.3. Functional diagram of memory architecture of CUDA capability 3.5 processor (Kepler architecture). The Stream Processors (SP) can benefit from broadcast mechanism (loading of value from particular memory location to all SP's simultaneously) if the programme is written to take advantage of it. This applies to Constant Cache and Shared Memory. The L2 cache can broadcast its data across multiprocessors.

7.3.2 The proposed method

The process is divided into three phases as follows.

7.3.2.1 Phase 1: Prototype Time of Flight (ToF) points calculation

In phase one, there are three steps. These are denoted by letters to avoid confusion with steps in phase three.

Step A: The imaging space is divided into straight z-lines with z-coordinates going in a general direction away from the probe, as illustrated in Fig. 7.4. In this implementation, this is the 'depth' coordinate related to some arbitrary spatial Cartesian coordinate system in which the probe elements, refracting surface and TFM image volume are described.

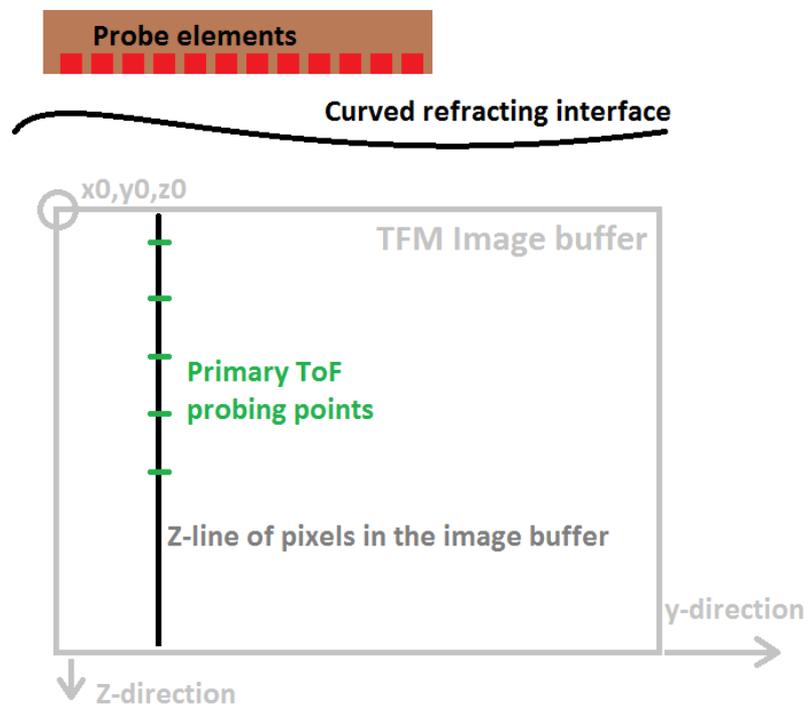


Fig. 7.4. Location of probe, refracting material interface, image buffer, and z-line of pixels inside the buffer. The linear memory locations progress down-first, then right. The third dimension is X.

Step B: For each z-line and probe element combination, propagation times for a number of points along the z-line are calculated using one of the iterative methods based on Fermat's principle. For example, there can be a 128-element probe and 1024 vertical lines in the image, giving a total of 131072 z-line point sets to be considered. The spatial distribution of

the points can be regular or generated using affine transformation[85] of Chebychev nodes[86].

The resulting propagation times are then called the primary ToF=f(z) points, where f(z) is some unknown function that describes the ToF value along the z-line.

Here, the insight begins with realising that along each of the z-lines as proposed, the consecutive ToF values are a monotonically increasing, smooth function of space. Therefore, the unknown f(z) function can be approximated with high accuracy using a polynomial of low order, as described within the next section.

Step C: The physical ToF [seconds] value calculated in step B, is scaled to a normalised value that maps into appropriate fraction of the A-scan. Crucially, this links the spatial location in the TFM image buffer with the temporal location in the FMC data buffer. This single multiplication/offset operation at this point saves millions of processor cycles that would be needed in phase 3 to execute the same operation.

7.3.2.2 Phase 2: Calculation of interpolant coefficients

It is proposed to interpolate the ToF=f(z) function with a ToF = fi(z) function in order to obtain the intermediate ToF values required for high resolution imaging.

A review of a selection of typical interpolating functions has been conducted and it has been concluded that a simple polynomial will offer the best combination of accuracy and computation cost. The polynomial is of form:

$$fi(z) = \sum_{n=0}^N a_n * z^n \quad \text{Equation. 7.6}$$

Where N is the order of the polynomial, and a_n are interpolant coefficients.

A transformation is used to process primary ToF=f(z) function samples into interpolating polynomial coefficients. This is executed using a specialised, non-branching solver as detailed in section 7.3.4.

At this point, the minimal order of the polynomial is selected which provides the requested accuracy of the approximation. To facilitate this, the prototype ToF data points obtained in phase one are divided into two groups. One group is used to fit an interpolating polynomial

of tentative order, and the second is used to estimate the error of the interpolation (if requested by the user). Selecting minimal order of the polynomial has the additional benefit of stabilising the solution.

The final interpolant coefficient database is segmented and optimised for loading into the GPU's on-chip shared memory, enabling rapid retrieval of the coefficients to the SP's using the broadcast feature, as detailed in section 7.3.2.3

For a given probe, surface and image parameter combination, a set of interpolant coefficients can be calculated once and stored for multiple reuse in phase three. If the probe-refracting surface configuration is changing (for example, when the probe is moving against the refracting surface), a re-calculation of this database is necessary. Still, this is a relatively cheap process when compared to phase three.

7.3.2.3 Phase 3: Signal accumulation

In this third and final phase, the values of pixels in the final image are calculated and stored into the output buffer.

In the GPU, during parallel processing, each thread is responsible for a single pixel of the image. The process for each pixel can be described by dividing it into 10 steps, as depicted in Fig. 7.6.

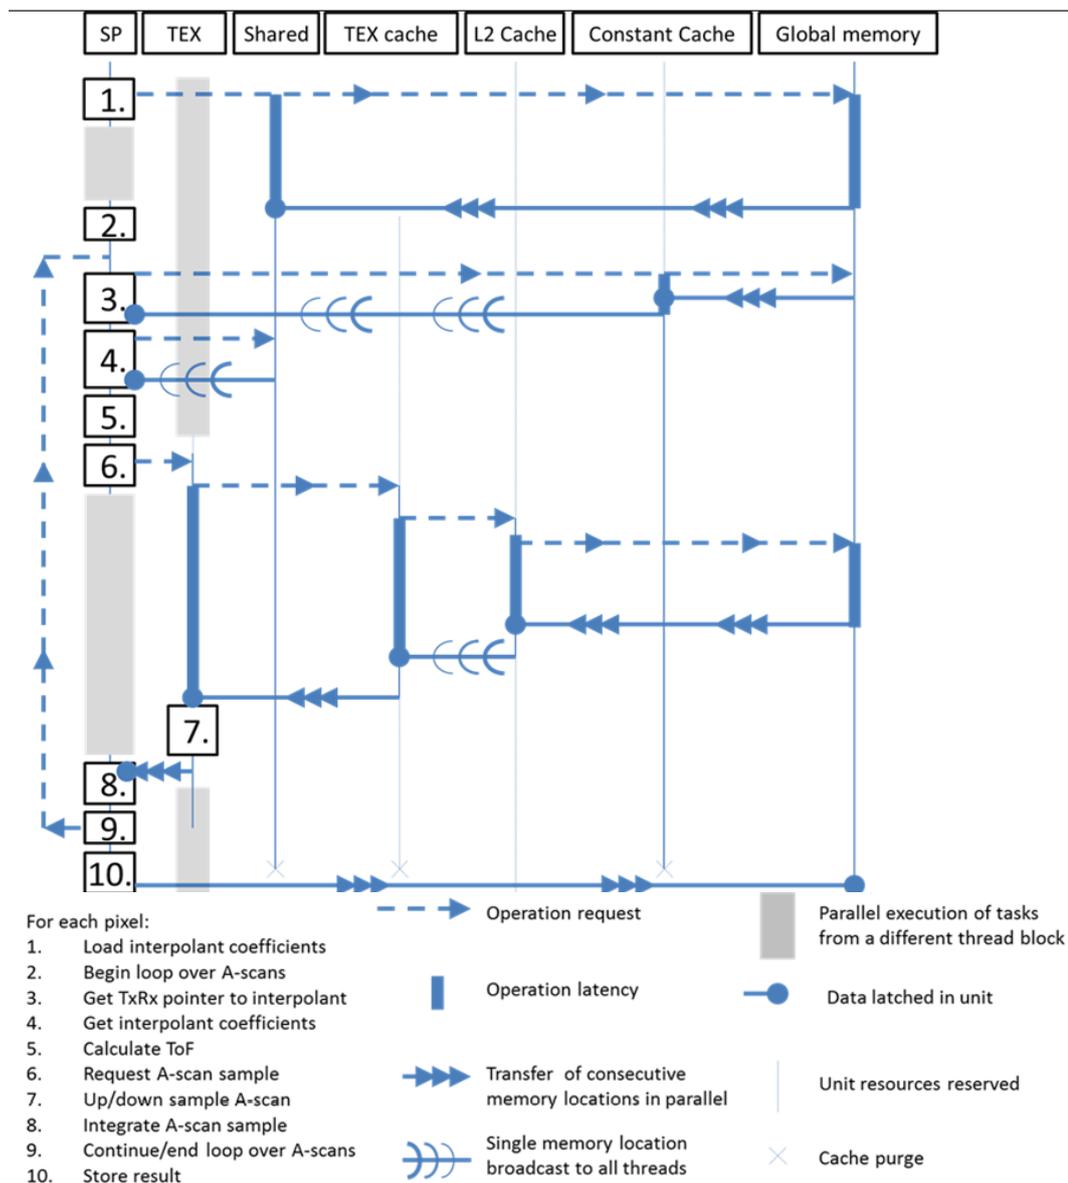


Fig. 7.5. Schematic of per-pixel computation kernel and memory access pattern. There are 10 algorithm steps per pixel. SP - Streaming Processor (also known as CUDA core). TEX – texturing unit. Shared – shared memory module. TEX cache – texturing unit local cache, if present. L2 – L2 on-chip cache buffering the device global memory access. Constant cache – the on-chip, per-Multiprocessor constant cache (also known as the code cache). Global memory – device global memory.

The thread block size is configured in such a way that all the active threads in a block process consecutive pixels that come from the same z-line of the image. This allows the threads of the thread block to take a couple of advantages from the GPU’s memory architecture, as described below.

For this implementation, a list that describes which A-scans out of the FMC buffer to use in each iteration (step B) is employed. This allows for the use of a single-code base to effectively execute ‘Sparse FMC imaging’ where some (or most) FMC signals are not acquired. This table is loaded through the Constant Cache, also benefiting from the broadcast mechanism in a similar way as the interpolant coefficient table. Additionally, this table could be used to alter the order of integration of the A-scans, which may come handy if nonlinear integration is considered.

Inside the thread for given pixel, firstly, the pixel’s z-coordinate is calculated from the thread (and pixel’s) index, as this is faster than storing the z-coordinates in a separate memory region. Then, the outermost and only loop processes the list of A-scans, to be taken into account for the pixel.

For each A-scan, the thread utilises the pixel’s z-coordinate and the polynomial coefficients obtained in phase two (which were loaded into shared memory here in step A of phase three), to calculate the respective transmitter-to-pixel and pixel-to-receiver propagation time (step 5). As explained in the section above, the resulting propagation time is already scaled to map onto the sample address in the FMC buffer.

At this point, the reader should note one particular feature of GPU architecture. That is, shared memory broadcast (Fig. 7.3). The interpolant coefficients corresponding to a single z-line of the final image are loaded from device memory into shared memory only once (step 1) at the beginning of the thread block, and from there the values are repeatedly broadcast to all threads (step 4) as needed, using the pointer loaded in step 3. This is possible because for each pixel (thread) in the z-line, consecutive A-scans are processed simultaneously in lockstep across all active threads of the thread block. One could describe this operation as ‘user-managed caching’. The first effect is that the required information regarding z-line is efficiently distributed to all the threads just as they need it, at an extremely high cumulative data-rate and extremely low latency. The second effect is that a few cycles later, the request for A-scan samples is sent simultaneously to the memory controller, which provides it with the opportunity to optimise the global memory access request, as detailed later within this section.

Again, the threads calculate the ToF using the Horner’s Rule[87] for polynomials (step 5), providing their respective pixel coordinate as a parameter. The Horner’ Rule method maps

to a short series of very efficient FMA (Fused Multiply Accumulate) instructions, with no loop. The code is specialised for a specific polynomial order so that no jumps are required. This effectively re-creates (estimates) the ToF value for each spatial location as needed. As detailed in step C of phase one, the ToF is already scaled so the resulting value is a pointer that the texturing unit can use to load the appropriate sample value from the FMC buffer (step 6).

Here, another feature of GPU comes useful to accelerate the retrieval of ultrasonic signal data. It would normally appear that a random memory access is required for the FMC sample (A-scan sample) retrieval step. However, since all the threads evaluate ToFs for a given transmit/receive pair (A-scan) simultaneously, and beginning from the top pixel down, they will request A-scan data that resides at neighbouring and progressing memory locations, albeit with non-integer stride. The texturing units can load requested A-scan data through L2 cache. This is very efficient because typically, an entire A-scan will fit in the L2 cache at once, requiring that its relevant section is loaded from the global memory only once. From there, the data block is broadcast to the texturing units that request it. This way, the data is efficiently distributed to the texturing units, which in turn operate their own, smaller texture blending cache (depending on the GPU model). The data from L2 cache is reused many times; less, depending on the effect of the L1 and texturing unit cache, or more if there are multiple thread blocks operating in parallel. In many cases, reloading the data from the global memory is avoided.

To continue, the texturing units select appropriate samples from the A-scan buffer and provide an additional service of interpolation between A-scan samples (super sampling) if necessary, thus maximising the parallel spread of the computations across the GPU chip. Some publications call this operation 'free' because as it progresses, the main CUDA cores are executing other threads, preparing a new batch of requests for the texturing units (grey blocks in Fig. 7.6 between step 6 and 8).

Finally, the threads integrate the A-scan samples they received from the texturing units. At this point, the integration operation can be summation (as in classic Synthetic Aperture beam-forming), but it can also be a different operation. For example, multiplication [88], phase coherence factor calculation [89], or CAPON adaptive apodising [90], or other. Since the ToF calculation operation uses only a few hardware registers, there is register space remaining for implementing more advanced integration methods.

Upon completion of integrating all the A-scans, threads store the computed pixel value in the final image buffer, which resides in the global memory (step 10). Again, this is an operation benefiting from the fact that neighbouring threads process neighbouring pixels in lockstep. The memory-write operation is well coalesced and parallel and exploits full width of the main memory bus.

Importantly, no synchronisation nor atomic operations are ever required because by design, no memory write race ever occurs.

7.3.3 Notes on the proposed algorithm

Although no hard mathematical proof is offered for this work, the polynomial interpolants have been found to be stable and robust across the experiments conducted. An interesting insight into polynomial as interpolant can be found in [91]. One of the most common myth sources is that: “The impression that [method] must be better than [other method] has probably also been enhanced by the fact that [method] is more difficult to compute”.

Since the image is processed line-by-line, only the interpolant coefficients related to that z-line need to be stored in the quickly-accessible on-chip shared memory. For example, for a 128-element probe and interpolating polynomial of a maximum order of seven and single-precision floating point format of the coefficients, $128 \times 7 \times 4 = 4480$ bytes of information need to be loaded from the global memory per z-line, but only once per z-line. Each thread working in this z-line re-uses the information via the broadcast mechanism, reading $7 \times 4 = 28$ bytes of coefficients per a-scan from this buffer, up to $2 \times 128 \times 128$ times per pixel. This saves a considerable amount of memory bandwidth that would otherwise be required if the broadcast mechanism and lockstep execution mechanism were not there.

However, this also means that running more than a few blocks of threads per GPU is somewhat counter-productive, as a very limited amount of on-chip shared memory is available. Therefore, this process works best if there are enough pixels in the z-line and not many z-lines processed in parallel. For small image resolutions (<196 pixels in Z-direction), the GPU resources will be under-utilised. The exact threshold of under-utilisation depends on the GPU model used.

Incidentally, it is of practical benefit to have the highest image resolution along the depth axis (z-line), because in this case the constructive/destructive interference between A-scans

produces the best TFM process gain (contrast improvement) and phase accuracy. If less than seven pixels per wavelength resolution is used, the TFM process may not achieve its peak process gain due to sampling phase error [92]. The same applies to other TFM-like algorithms, for example Phase Coherence Factor algorithm.

Overall, the process described is extremely efficient and takes approximately 48 hardware cycles per integrated A-scan (this measured value includes all exposed latencies) per pixel, and is the main source of performance for this implementation.

The core to this efficiency is the description of ToF as a simple function of space, and then, utilisation of broadcast and lock-step mechanisms of the GPU modules.

7.3.4 The non-branching polynomial interpolant coefficient solver

7.3.4.1 Introduction

To find the coefficients for interpolants required for this imaging algorithm, the well-known least-square method of fitting a polynomial into a set of points by solving a matrix-quadratic equation can be used.

Typical CPU implementations of this method utilise loops and conditional jumps to allow a single implementation of a code to solve for an arbitrary order of polynomial and an arbitrary number of data points, sometimes even reordering the data first to minimise the numerical errors. However, for this work, the solver's speed of execution is of utmost importance. It is proposed to compile a range of specialised solvers, each exclusively taking a fixed number of data points and returning a fixed order of polynomial. This approach allows for a non-branching code for each case to be obtained.

7.3.4.2 Implementation

The code has been obtained using the Wolfram Mathematica Computer Algebra System, embracing a method described in detail in appendix C. This method is similar to the one described in [93]. The detailed source code and results, along with the ports for Matlab and C, are made available within the digital version of this thesis and also on a companion website [94].

A subset of statistics on the addition and multiplication operation counts versus the number of inputs and polynomial order has been gathered in Table 5.2. In this table, for each variant of the code (variant being defined by order of the polynomial and number of input points) the count of arithmetic operations is presented. The table gives the reader an idea of how the complexity of the problem escalates with increasing fidelity of the solution.

The code obtained only requires multiply-accumulate operations and a single-reciprocal and uses no jumps whatsoever. The benefit of such an approach is that the code will execute efficiently on a GPU, solving multiple z-lines in parallel. It is appreciated that such an approach is not well suited for poorly conditioned inputs and will allow the numerical errors to surface in the results, even when using double-precision arithmetic. However, as argued in the next section, in this application the inputs are always well-conditioned and observed numerical errors are acceptable.

Table 5.2. Selected statistics of computation cost for calculating the interpolating polynomial coefficient, as a function of polynomial order and number of contributing data points. Each version of the algorithm additionally requires a single-reciprocal. FMA – Fused Multiply Accumulate operation. MUL – multiply operation. ADD, SUB – addition or subtraction operation, respectively.

Polynomial order	Number of input points	FMA operation count	MUL operation count	ADD+SUB operation count	Total operation count
4	6	120	182	44+35	381
4	10	192	260	76+36	564
5	8	330	445	76+99	950
5	12	453	571	116+98	1238
6	10	743	976	119+296	2134
6	14	935	1152	167+296	2550

7.3.4.3 Selection of interpolant order and error analysis

To establish confidence in the proposed method, it is necessary to perform a detailed error analysis of the developed algorithm. The calculated ToF errors come primarily from three sources: (a) inaccurate primary ToF solver, (b) inherent method inaccuracy of polynomial approximation, and (c) numerical inaccuracy of polynomial coefficient solver for a given numerical representation type (single- or double- precision), stemming from accumulating the round-off errors through the process.

In this work, it is assumed that (a) is exact. Errors incurred by the primary ToF solver are beyond the scope of this work⁷. Here, only the source of errors associated with (b) and (c) will be considered.

To estimate an end-to-end ToF calculation error of this process, the following method has been applied:

1. Prepare a large set of pre-calculated ToFs for a typical imaging scenario, using the reference ToF finder function,
2. Convert the subset of pre-calculated ToFs to interpolant coefficients and then back into the full grid of ToFs,
3. Calculate the peak difference between original and processed ToF data.

The summary of interpolation error results for an imaging scenario depicted Fig. 7.4, in function of the selected order of the polynomial, is gathered in Fig. 7.6. The surface is of form $z = f(y) = A \cdot \sin(B \cdot y)$, where $A=7\text{mm}$ and $B=2\pi/100\text{mm}$. The test probe element is 30mm above the $z=0$ plane, at $P = (0,0,30)\text{mm}$, which is representative of a typical ultrasonic NDE imaging scenario.

⁷ In the course of this work, a flaw has been found in the original Matlab's implementation of *fminsearch*, causing it to return incorrect results for some combination of inputs, as detailed in appendix B

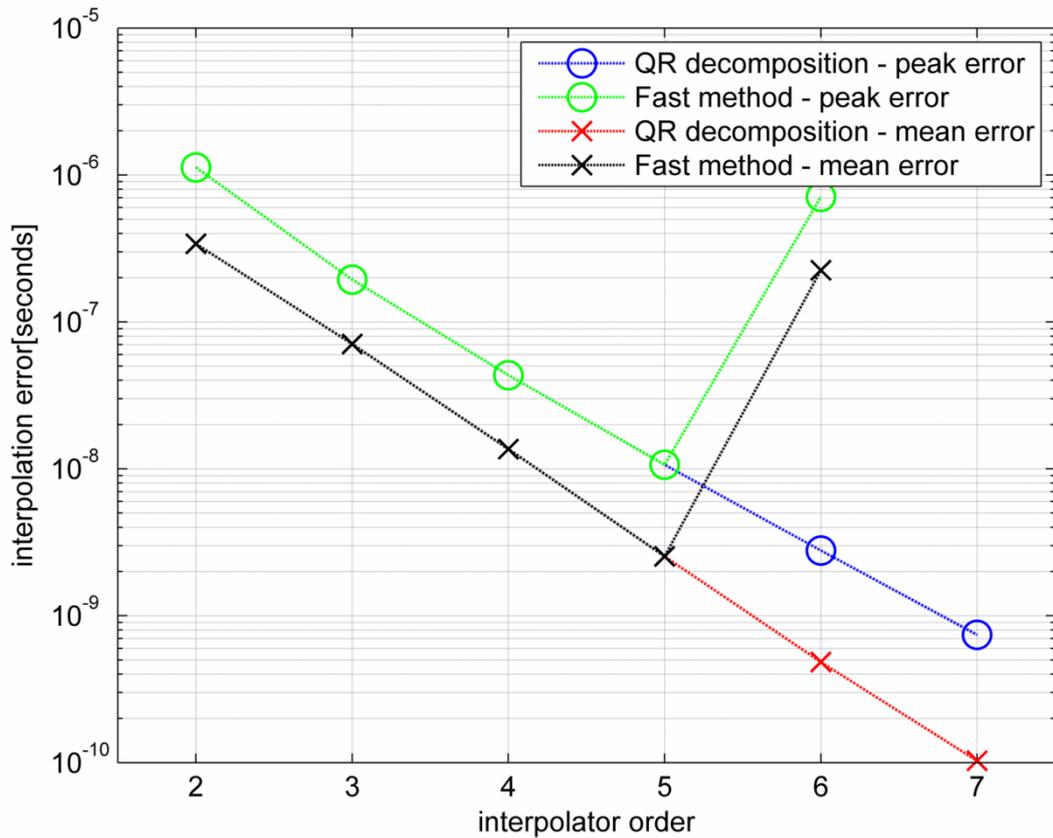


Fig. 7.6 Comparison of interpolation error when using different interpolation coefficient solvers

7.3.4.4 Discussion

The proposed direct no-branch solver returns values that are virtually equal to Matlab's QR solver up to interpolant order of five. If a higher order interpolant is needed, QR solver is recommended.

The significance of this result is as follows. For a source signal sampling of 50MHz, the useful bandwidth of the sampled signal is no better than approximately 17MHz, full sine wave cycle at that frequency is 58.8×10^{-9} seconds long. Therefore, to obtain phase accuracy better than 1/4 cycle (worst case scenario), the error must be less than 1.47×10^{-8} seconds. This implies selecting an interpolant order of at least five. In this same case, the phase accuracy for a 5MHz signal will be better than the 1/15 cycle.

The cited timing errors are peak errors as the average timing error is much lower. In any case, it is possible to obtain an arbitrarily low peak error estimate by raising the interpolant order (up to the limit of numerical representation accuracy). This is at a cost of a minor decrease in phase two performance.

It is appreciated that the evaluation method presented here does not give a strict upper bound for error. However, it can easily be repeated for any practical imaging scenario and therefore it is submitted here as satisfactory. Since phase one and two take a combined 5.6% of the total processing time, adding error validation to every image frame is practical.

7.3.5 Implementation benchmark

The proposed process consists of several phases that can be executed independently, and thus benchmarked independently. In other publications, the overall performance in a practical scenario is typically expressed in frames per second for a given specific scenario. Here however, a more synthetic approach has been taken. Each phase is benchmarked independently and then a method to estimate the overall performance for a given set of input parameters is provided.

For each phase of the process, an appropriate performance unit has been introduced. For phase one, it is ToF points / second. For phase two it is lines per second (as each atomic transform deals with entire image line). Finally, for phase three it is paths/second, as each atomic operation deals with estimating ToF over a specific path.

It is proposed that with such metrics, one can trade the number of pixels, the number of probe elements and the Tx/Rx firing scheme against the performance of the particular GPU system used and the frame rate achieved. For any TFM image resolution and number of elements of the probe, the total calculation cost can be obtained. For example, for a 1024^2 pixel image and a 64-element probe operating in full FMC firing scheme, the computational cost of basic TFM method is $2 \cdot 1024^2 \cdot 64^2 = 8.6 \cdot 10^9$ paths.

In the example above, assuming that we select a polynomial order of five and eight sampling points, there are $8 \cdot 1024 \cdot 64 = 524 \cdot 10^3$ primary ToFs to be calculated in phase one, and $1024 \cdot 64 = 65.5 \cdot 10^3$ interpolant coefficient sets to be obtained in phase two. In phase three, in the current implementation, for each pixel, Tx-to-pixel ToF is calculated

once and then for each of these, pixel-to-Rx is calculated, resulting in $1024^2 * 64 * 64 = 4.3 * 10^9$ paths.

Assuming that the FMC data is uploaded to the GPU asynchronously (which is possible to do on professional line cards), each frame will take 0.17 seconds to compute. Phase one takes 3.9%, phase two takes 1.6%, and phase three takes 94.4% of the total time. This is assuming that the best algorithm version is used in each case. For example, if one chooses not to use the QR solver instead of the newly proposed fast solver, the phase two would take much longer, for a 0.6% / 84.7% / 14.6% split on the overall processing time.

Therefore, any future improvement has to be concentrated in phase three of the process. The performances cited scale well over multiple GPUs and, for example, are halved for NVidia GTX590 GPU, and halved again if two GTX590 cards are used in a single-system.

The hardware platform is Intel i7-3820 @ 3.6GHz (all eight threads used in all cases), and NVIDIA GTX 580. Please note that this comparison is not intended to highlight differences between hardware platforms, but rather to show the effects of choosing the algorithm for a given stage of the TFM process. Table 5.3 compares QR solver performance versus the new fast solver performance (lines four and five) and then compares different implementations of the TFM algorithm (lines seven, eight and nine).

Table 5.3. Results of benchmarks of each phase

	Phase name	options	platform	Performance unit	Result
1	Phase 1 - Calculating prototype Time of Flight points	Planar interface z=0, simplex method	GPU	Points/second *10 ⁶	77.9
2		Planar interface z=0, simplex method	CPU		.42
3	Phase 2 - Transform from time points set into interpolant coefficients	8 points into 5 th order, fast solver	GPU	Lines/second *10 ³	23'538.
4		8 points into 5 th order, fast solver	CPU		2'960.
5		8 points into 5 th order, QR solver	CPU		70.
6		Double precision, 14 points into 7 th order, QR solver	CPU		68.5
7	Phase 3 - TFM integration	Nearest sample interpolation	GPU	Paths/second *10 ⁹	26.7
8	1 st generation TFM for comparison	No refraction	GPU	Paths/second * 10 ⁹	0.46
9	1 st generation TFM for comparison	Planar surface refraction	GPU	Paths/second * 10 ⁹	0.05

At the time of writing this section, the measured performance was 556x cheaper than 1st generation implementation and 56x cheaper than the best published algorithm benchmark with refraction support that we were able to access (December 2013). However, it is appreciated that this figure can change at any time.

7.3.6 Experimental validation

Validation of the developed algorithms will use the experimental set-up illustrated in Fig. 7.7 and schematically presented in Fig. 7.8. The probe is 128-element, 5MHz, linear phased array probe and the Phased Array Controller is the Zetec Dynaray [Zetec, Snoqualmie, WA, USA]. The probe is placed over a half-cylinder of solid PVC material, in which a flat bottom hole has been drilled out.

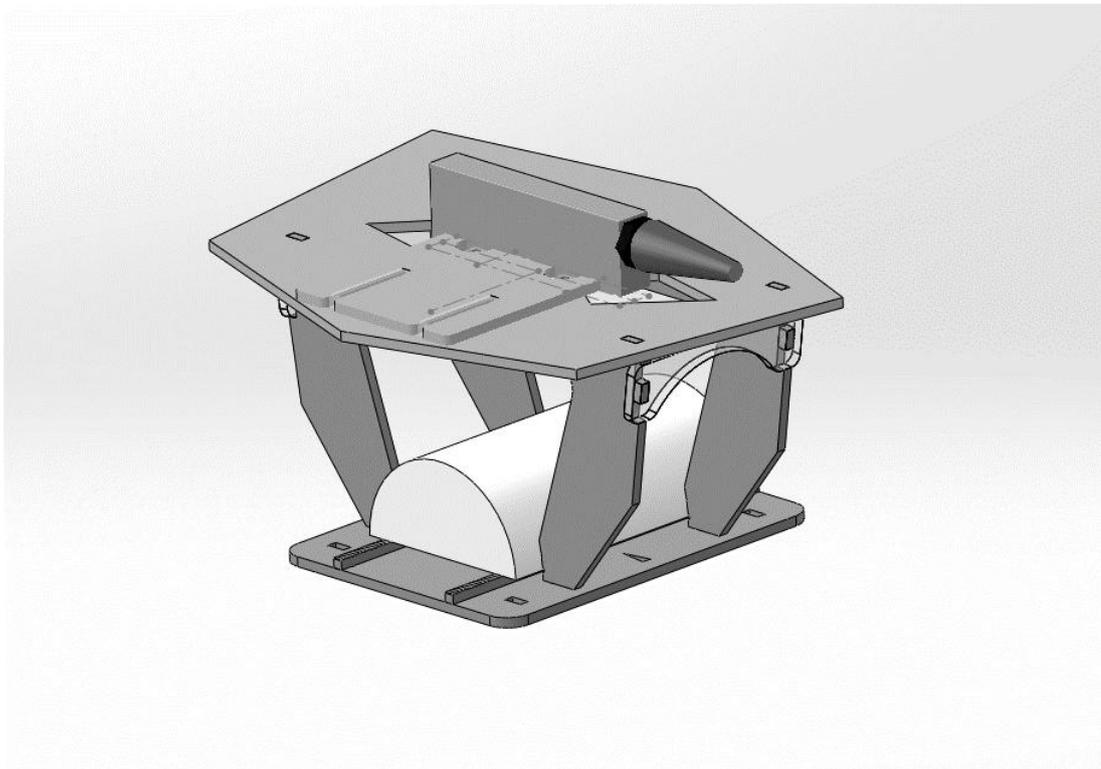


Fig. 7.7. The experimental specimen and probe holder. Irregular geometry of the top plate has been used to avoid spurious reflections.

This approach has verified that the developed algorithm returns an image as expected. The reflection coming from the hole is blurred-out when the refraction is not taken into account, as shown in Fig. 7.9. When the refracting surface is taken into account (Fig. 7.10), the reflection is properly located and focussed and the amplitude of the return is increased by 4.57dB. The large black patches on the sides of the cylinder reflection are side lobes, coming from the cylinder's surface, as expected for this probe type. In the enhanced dynamic range image in Fig. 7.11, the back-wall of the PVC specimen is correctly visible at depth of -0.1m. The back-wall being flat further confirms that the algorithm performs as desired.

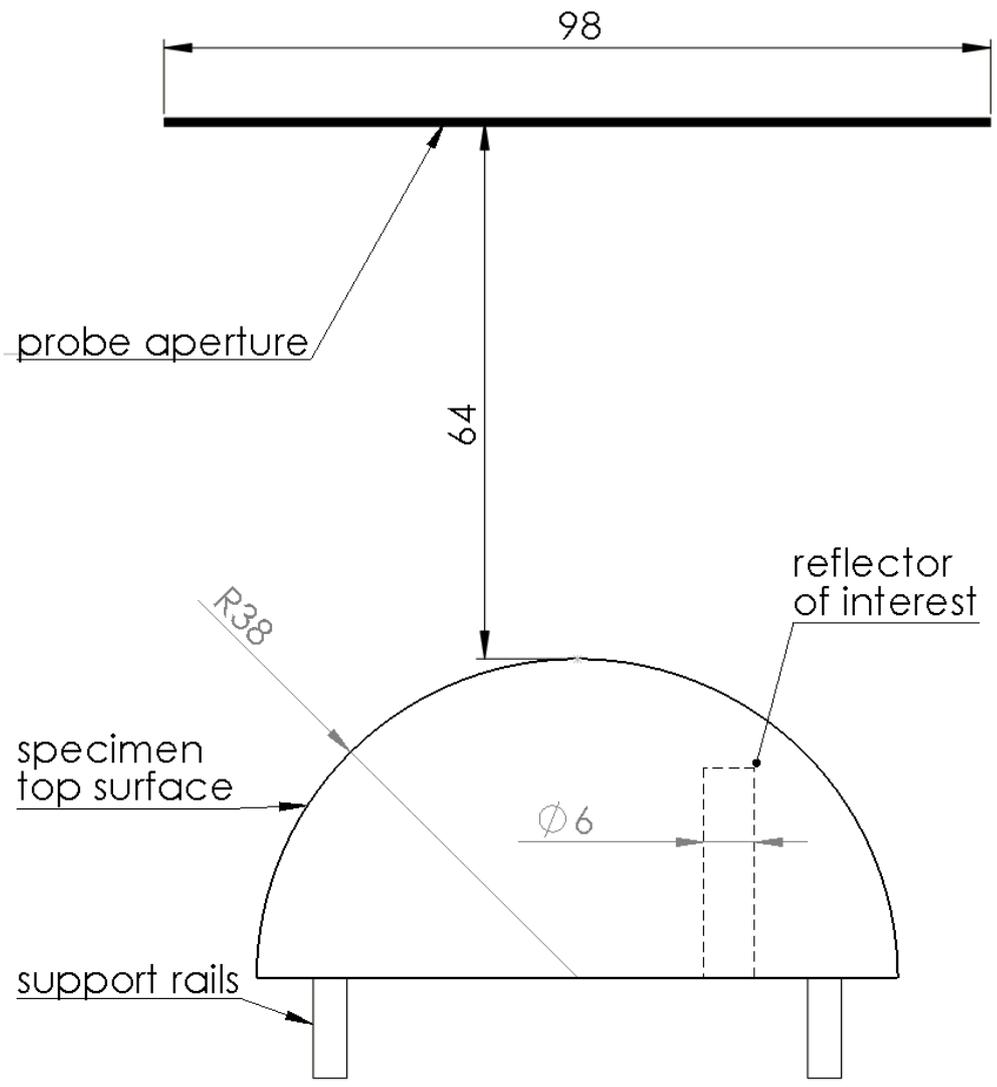


Fig. 7.8 Schematic drawing of the specimen and probe aperture

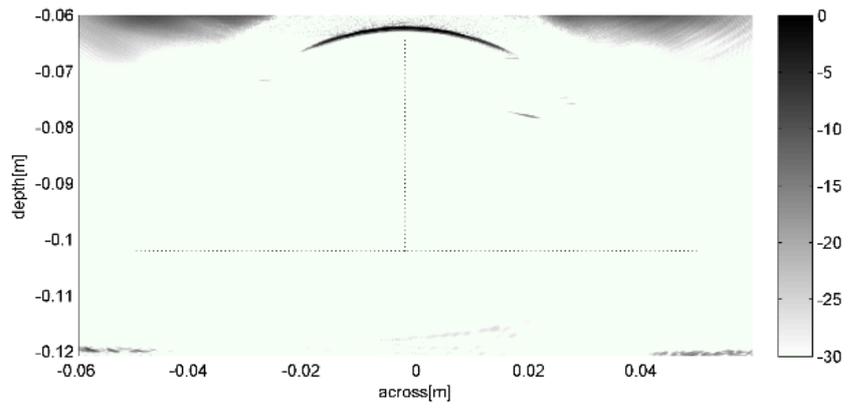


Fig. 7.9 The image of the flat bottom hole inside the specimen. Image assuming that the rays do not undergo refraction.

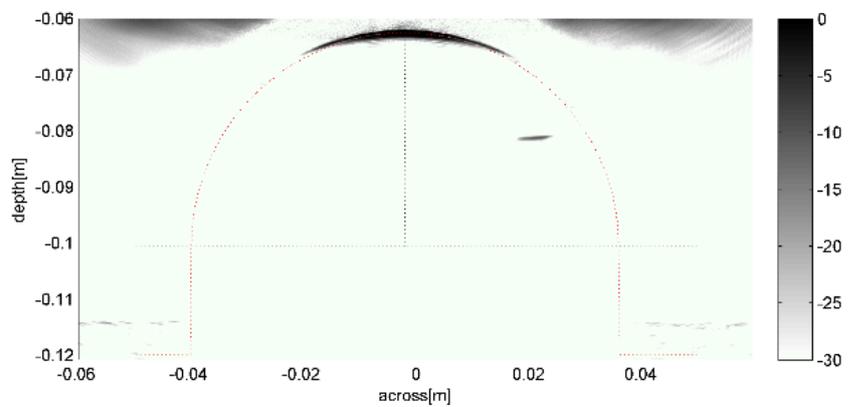


Fig. 7.10 The image of the reflector as in Fig. 7.9, but with correct refracting surface taken into account. The amplitude of the reflector is 4.57dB higher and the reflector is correctly positioned. Importantly, the shape and orientation of the reflector is imaged correctly (flat bottom hole).

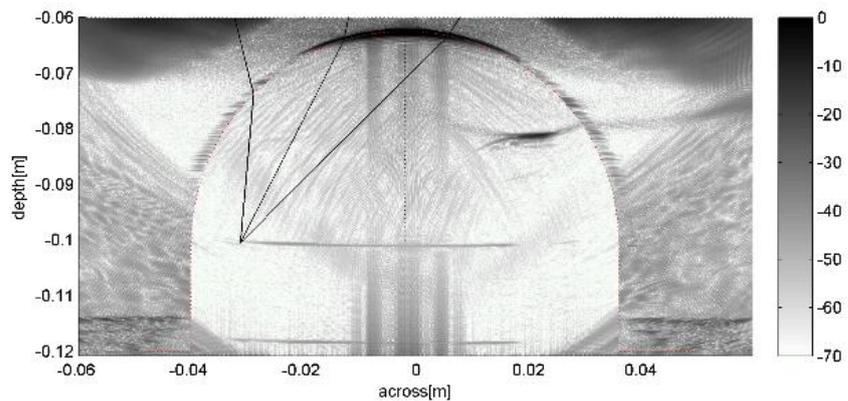


Fig. 7.11 The image from Fig. 7.10, but with wider dynamic range. The correctly flat back-wall is visible at -42dB from the top surface. The three black straight lines exemplify calculated ray paths between the three probe elements and a pixel in the image.

7.3.7 Remaining implementation notes

The storage of ToF polynomial coefficients can be viewed as a way of ‘lossy data compression’, where a ToF look-up table is replaced with a function taking carefully crafted alternate data. This new function can reconstruct parts of the original look-up table when and where needed in a way that is significantly cheaper than storing the entire look-up table.

The efficiency of this solution comes from transforming the ToF calculation problem in a way that maps well to the hardware resources available in GPU processor. It is therefore proposed that for a given desired frame rate and image size, the implemented solution should be significantly cheaper than a corresponding FPGA implementation.

Importantly, the algorithm as described in phase three can be equally used for cases where there is no refraction (single-layer media), single-refraction (dual layer media), or multiple refraction (three or more layers), with no loss in performance. The same concept can also be used to realise the classic delay-and-sum beam forming or Dynamic-Depth Focussing beam forming.

7.4 Case studies

7.4.1 Case study: cueMAP: probe CAD tool

cueMAP is a software suite that simulates full pulse-echo and pitch-catch imaging performance of the entire 3D imaging system, across a volume of interest. This package takes into account probe properties, imaging algorithm properties and load material properties. Sensitivity, coverage, and most importantly, contrast of the image offered by the system of given design can be estimated before building the probe. The specimen material can be multi-layered, and curved interfaces (eg. pipes or nozzles) are supported.

The simulation process is illustrated in Fig. 7.12 and Fig. 7.13. Firstly, the Phased Array Controller (PAC), probe, specimen, and imaging algorithm description is collated into a database. Next, the imaging volume of interest inside the specimen is defined. For each imaging point inside the specimen (voxel, volume picture element), a simulated, point-like reflector is created and the FMC (Full Matrix Capture) data simulated as transmitted and received using the previously described probe/PAC/specimen combination. An image of this

reflector is obtained using the TFM imaging algorithm. This image is a “point spread function” for a given location. This image is then analysed automatically and image quality characteristics extracted. Currently four parameters are recorded: relative sensitivity, -3dB spot size, side lobe level, and leakage factor. These image quality characteristics can be described as follows:

Spot size – represents the main lobe size at the reflector location, since the reflector itself is point-like. This corresponds to imaging resolution (smallest distance between two distinguishable reflectors) and should not be confused with the smallest detectable reflector.

Maximum side lobe level – corresponds to the maximum amplitude of a lobe that is not connected to the main lobe. This corresponds to a worst case imaging artefact – showing a reflector where there isn’t one, but the suspected region is close to other strong reflectors. One can relate this to false positive rate of detection.

Signal Leakage – integral of the image energy outside the main lobe. This corresponds to worst imaging artefact when trying to detect a lack of reflection, close to other reflectors. One can relate this to false negative rate of detection.

Relative signal sensitivity - for arrays with directional elements, sensitivity drops drastically for off-axis reflectors. Sensitivity also falls for reflectors that are far away from an array with omnidirectional elements. This measure can be related to smallest detectable reflector.

Spot size is indicative of actual **resolution** of the image, while side lobe level and leakage factor indicate image contrast. Relative sensitivity is important to establish coverage of the imaging system, since there might be places where ultrasonic energy does not reach. An example reason might be that the probe’s elements are directional, or the refraction angle is too shallow. Currently, no complex geometry effects such as shadows or re-radiating surface waves are simulated; however, curved material interfaces that refract sound are taken into account.

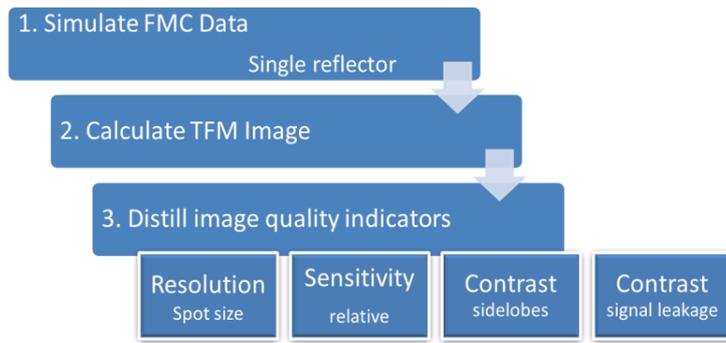


Fig. 7.12 The cueMAP process - overview

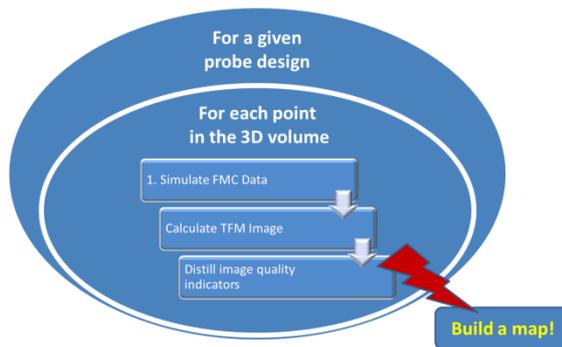


Fig. 7.13 The cueMAP process – per pixel

The Full Matrix Capture simulation module

A computationally inexpensive method of obtaining simulated FMC data was created for this work. For simplicity, only point like reflectors and hard fixed surface reflectors are supported in the module. Furthermore, a simplifying assumption has been made that the temporal shape of the reflected signal is the same irrespective of location and distance to the reflector. Directivity of the probe’s element and propagation distance is taken into account as a reduction in amplitude of the returned signal.

The FMC simulation module works as follows. A recording of the pulse-echo signal has been obtained using single element of the phased array of interest. Directivity of the probe element has been modelled using raised cosine function:

$$y(\alpha) = \cos(\alpha)^\beta$$

Equation. 7.7

Where 'y' is a fraction (reduction of) of amplitude at given α ; α is angle formed between axis of the array element and the line leading from the array element to the reflector; and β is a experimentally-measured parameter of directivity. The parameter has been obtained using curve fit to the measured amplitude of return signal in an experiment where a moving reflector has been presented to the probe.

The FMC simulation works by inserting the measured pulse-echo signal, (with amplitude modified by the element directivity model and distance), into the FMC matrix at a position that is dictated by the calculated time of propagation between the transmitting element, reflector and the receiving element. Refraction through layered media is taken into account, and the interface can be curved. All other effects are ignored. Strength of the reflector can be modified to simulate effects of shadowing of small reflectors by side lobes of larger reflectors.

Such method yields a simple, computationally cheap computer program that generates FMC data amenable to imaging using standard or novel ultrasonic imaging algorithms. Additionally, the algorithm is easily parallelisable for execution on multicore or GPU systems for extra performance, as individual lines of the FMC are essentially computationally independent.

Example map of resolution and contrast

An example, 3MHz, 128 element, sparse 2D phased array probe has been analysed using cueMAP software. Maps of resolution, contrast and sensitivity have been obtained for a 2D slice through volume of water in which the probe is immersed. The slice has dimensions 15mm x 35mm and has been sampled in a regular grid of 2mm in each direction. Full aperture of the probe has been used for both transmission and reception and the imaging algorithm is Total Focussing Method.

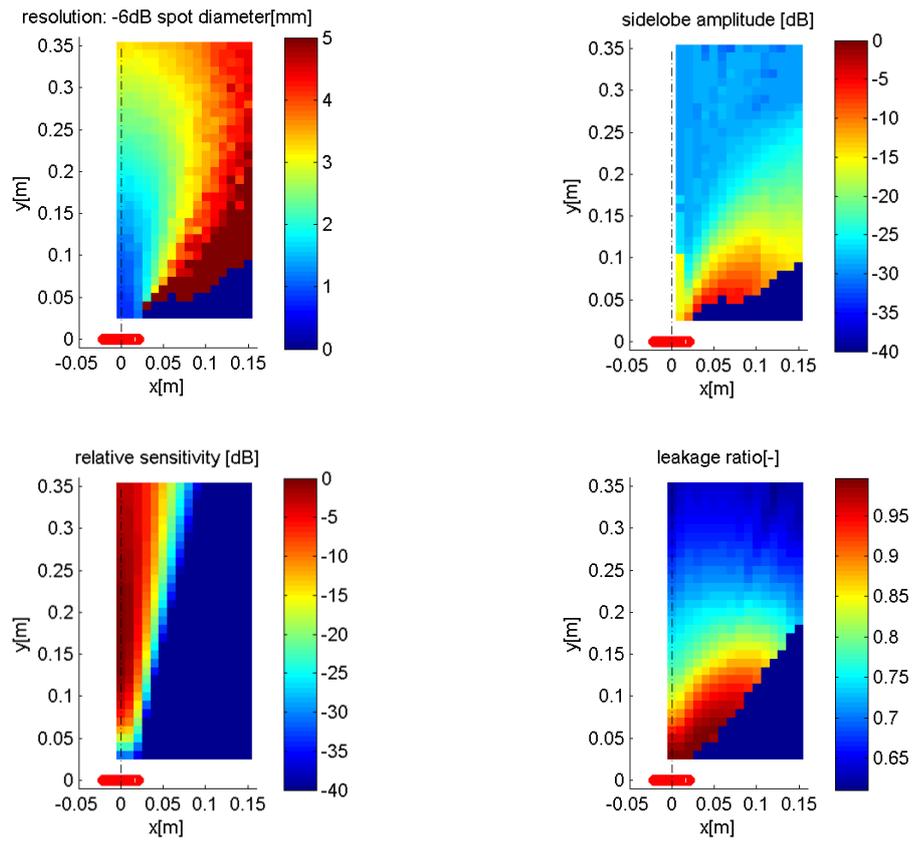


Fig. 7.14 Calculated image quality maps – example

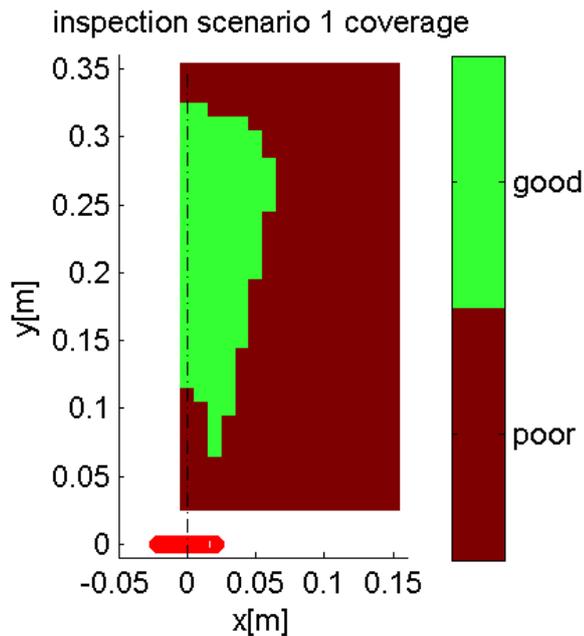


Fig. 7.15 Map of acceptable quality of image assuming example inspection requirements.

The maps reproduced in Fig. 7.14 describe variation of the calculated image quality across the volume. If the required image quality is well defined for a given inspection scenario, a combined map of coverage of the inspection can be obtained as depicted in Fig. 7.15. To produce this Figure, the following design criteria has been assumed: sensitivity $>-30\text{dB}$ of peak, spot size $\varnothing < 2\text{mm}$, side lobe amplitude $<-20\text{dB}$.

It can be noted that regular Total Focussing Method (TFM) does not produce best results in the area directly below the probe. This is not the limitation of TFM, but rather the probe itself: the limiting factor is reduction of sensitivity due to directivity of elements, where some of the elements do not 'see' the test reflector placed at that location. It is proposed that a sub aperture TFM imaging could be used in that area to resolve the situation.

Apparent 'noisiness' of the maps is a result of finite discretisation of the TFM images, resulting in quantisation error of measurement of simulated image features.

Summary

The presented approach will be used to design 2D phased array probes that fully match NDE inspection scenario requirements, while minimising the cost of both probe and related

data acquisition and processing system. Main feature distinguishing cueMAP from other similar systems is that properties of the imaging algorithms are taken into account, and the calculations are arranged in way that allows utilisation of massively parallel computers, completing calculations in practical timeframe.

7.4.2 Case study: 3D ultrasonic inspection of an industrial specimen

Ultrasonic scanning was performed using a Vermon 5MHz, 128 element, 0.7 mm pitch, 90mm aperture phased array probe. Since a flat surface was available on the pump casing, a scan was conducted using a manual scan of this probe across the flat surface face as shown in Fig. 7.16

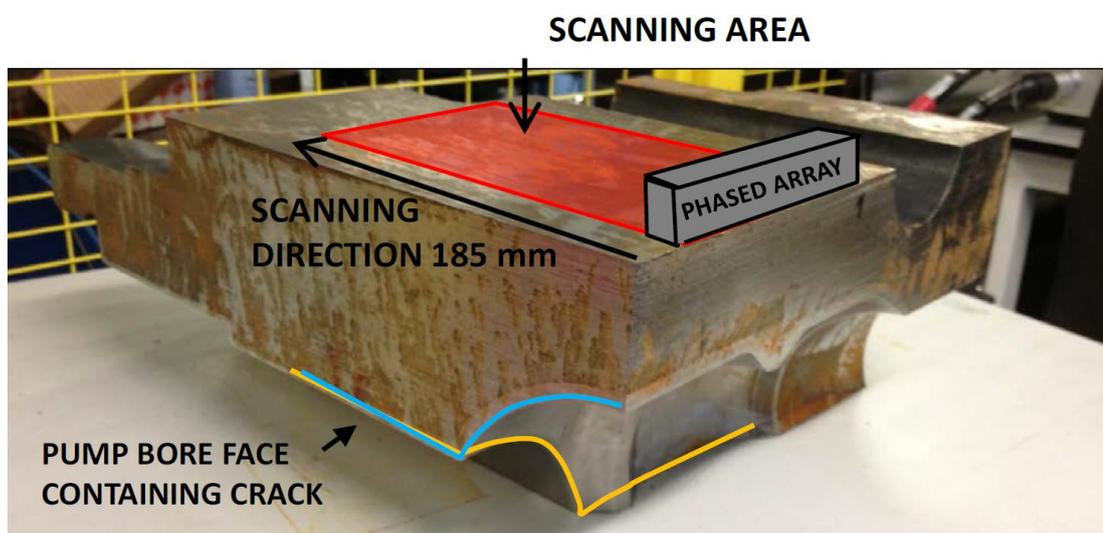


Fig. 7.16 pump casing showing location of phased array transducer and scan geometry

The array's location on the surface was monitored using a laser tracker that could locate the $[x,y,z]$ position of the probe to an accuracy of better than 0.05mm. Custom software was used to record and correlate the position and ultrasonic data from the probe. Scan time for the inspected area was 3 minutes, which included probe translation, data acquisition and signal processing using the cueTFM algorithm. The resulting data was visualized in a freeware 3D Volumetric Rendering application, Voreen[95].

In order to understand the 3D data set, it is first useful to define the inspection planes of the problem which are shown in Fig. 7.17.

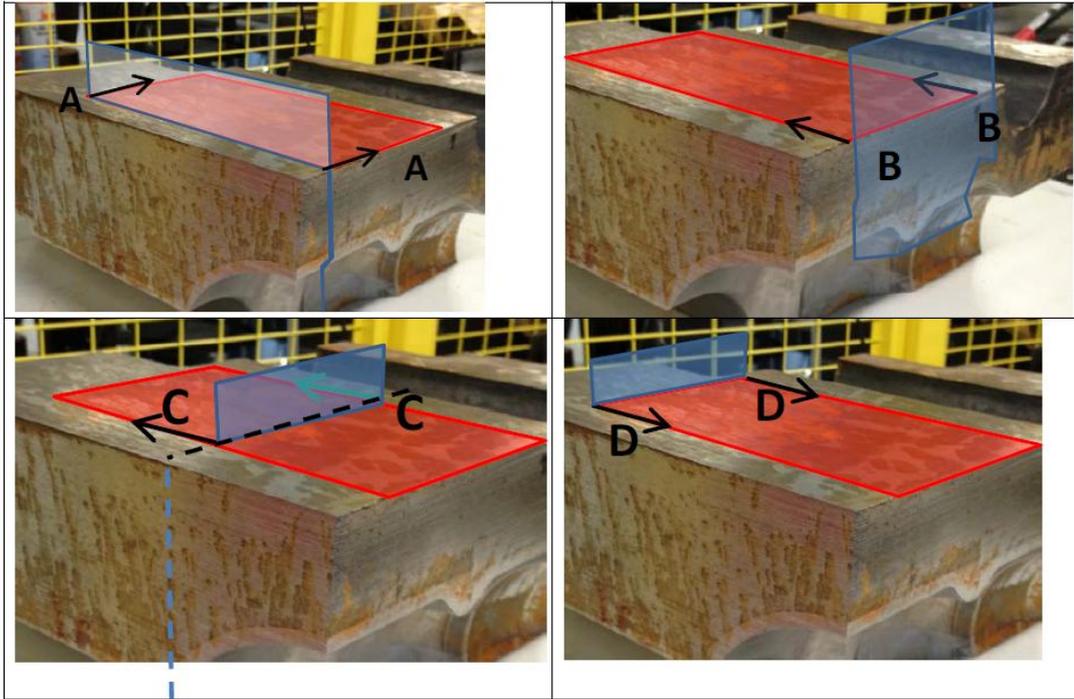


Fig. 7.17 Definition of reference planes in sample

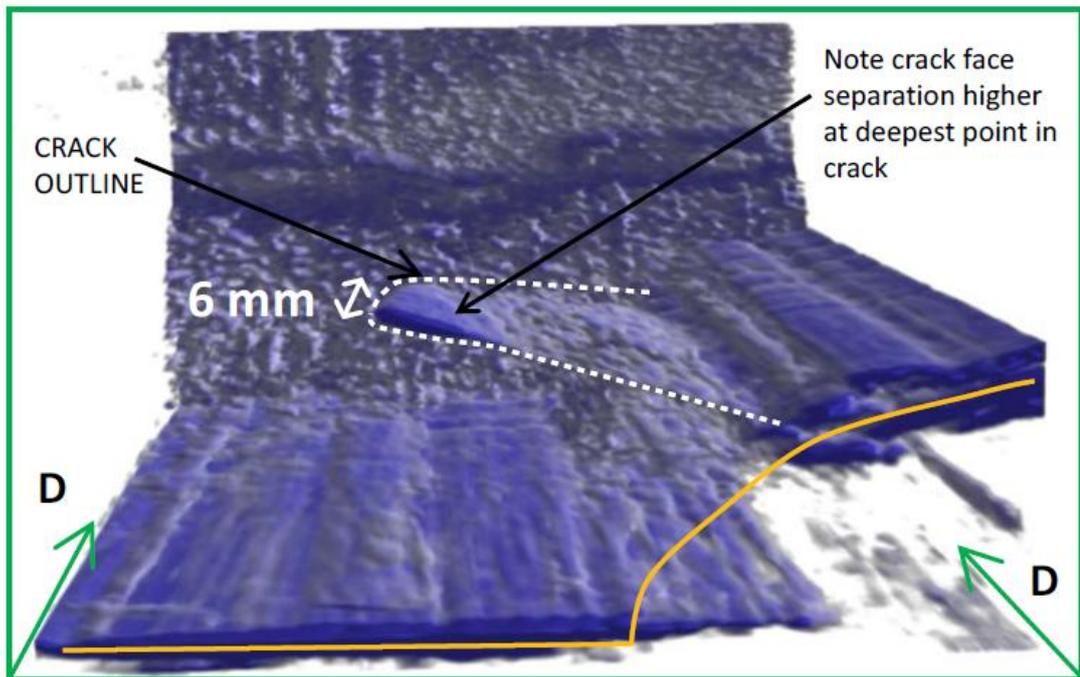


Fig. 7.18 Internal crack penetration and depth visualisation

The 3D data model was captured in a number of orientations to represent imaging along these different reference planes. Since accurate positional information was recorded at all scan points, it was possible to estimate the internal dimensions of the defect from these images. Fig. 7.18 provides perhaps the clearest overall viewpoint of the extent of crack growth into the body of the pump casing. It is evident that additional information on the crack face separation as a function of crack depth is available.

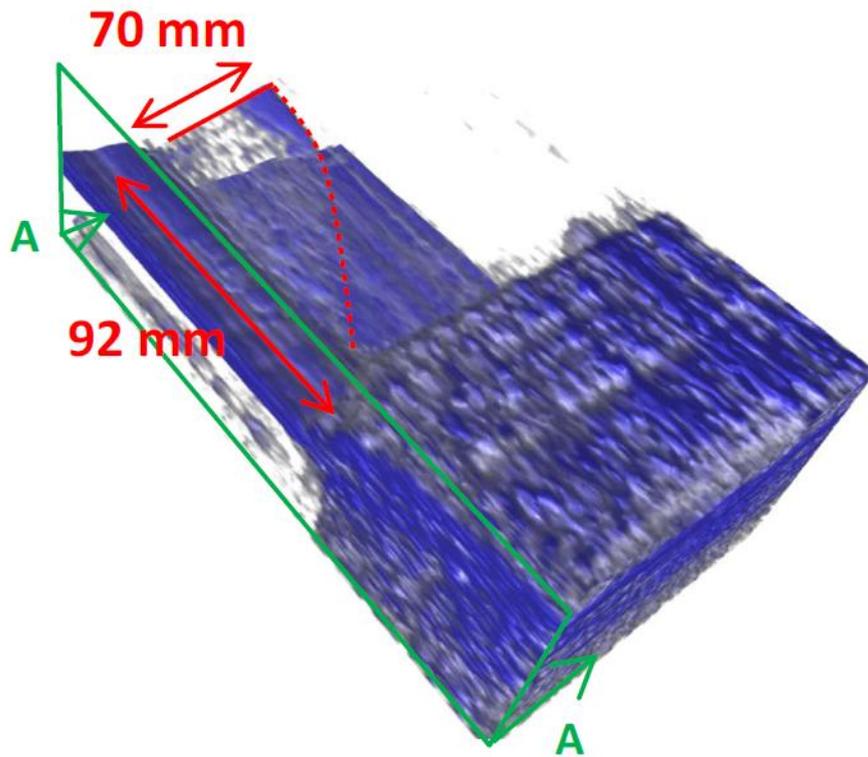


Fig. 7.19 A-A reference plane image of pump casing

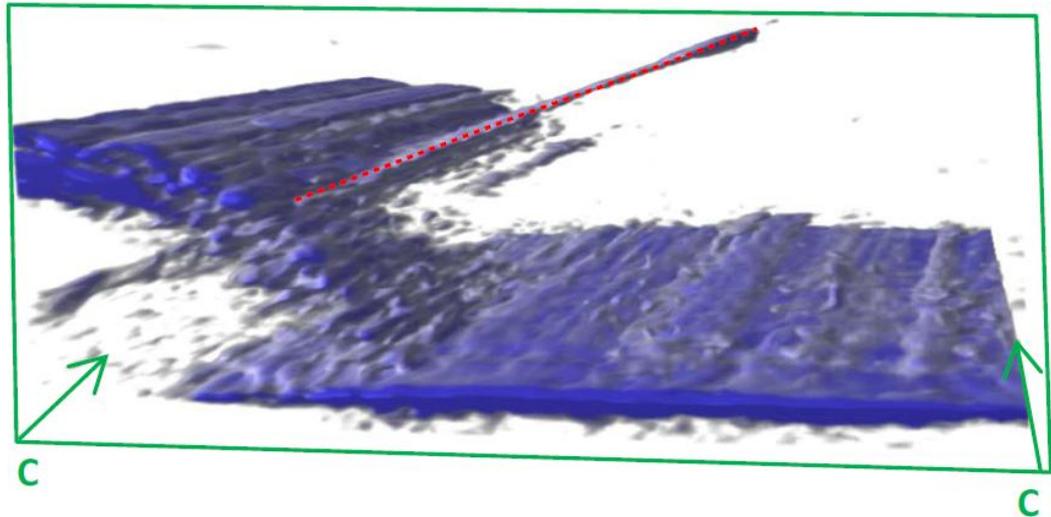


Fig. 7.20C-C reference plane image of pump casing

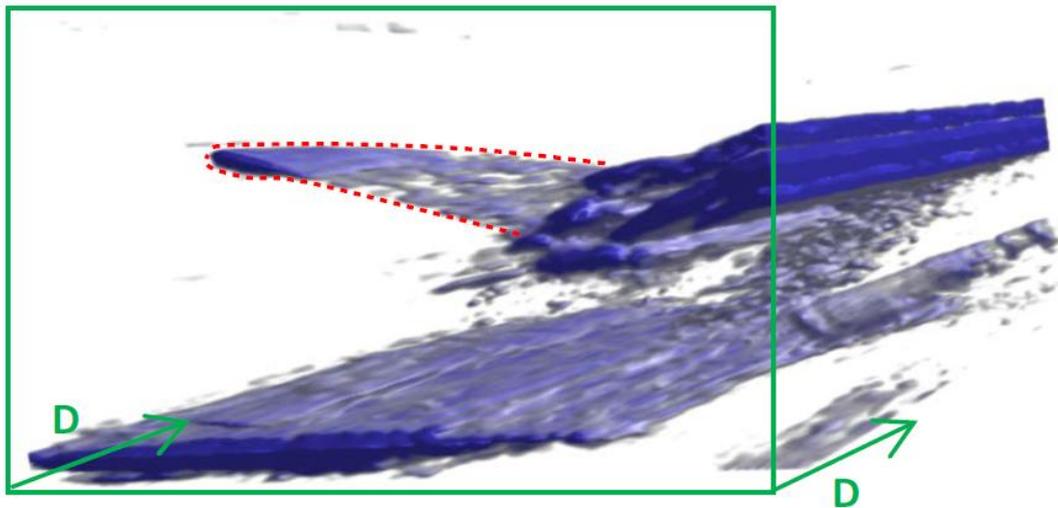


Fig. 7.21 D-D reference plane image of pump casing

Fig. 7.19, Fig. 7.20 and Fig. 7.21 show the individual AA, CC, and DD planes respectively and provides additional sizing information.

Unlike dye penetrant testing which provided an indication of surface length of crack only, the ultrasonic imaging provides detail on the evolution of the crack morphology as the defect penetrates into the volume of the sample.

Although the experimental procedure was in this case eased by the presence of a flat reference plane along which the probe was scanned, there is nothing in principle to prevent such imaging through a more complex curved surface geometry.

Summary

A 3 minute manual scan of the industrial specimen was undertaken using a 5MHz phased array ultrasonic probe. FMC and TFM data was collected from the scan and post processed to obtain a rendered 3D full volumetric image of the interior of the material.

A surface breaking crack was observed to run along the internal bore for a length of 92mm. Using the 3D ultrasound imaging, it was possible to estimate the depth penetration of the crack (70mm), and to visualise the evolution of crack shape and face separation as a function of crack depth.

7.5 Summary and conclusion

The recent introduction of inexpensive GP-GPU computing architecture presents new opportunities to develop algorithms that utilise the architecture to solve domain-specific problems.

In this chapter, the following problems have been addressed:

- Frequency domain acoustic beam simulation,
- Efficient algorithm for solving refracted ray travel time,
- An implementation of TFM algorithm with on-the-fly Time of Flight calculation,
- A second generation of TFM algorithm that uses efficient storage of precomputed Time of Flight data, enabling arbitrary refracting surface to be taken into account.

With the high practical performance of these implementations, it has been demonstrated that advanced imaging algorithms previously considered to be of research-interest only, or limited to military/nuclear grade applications due to their cost, can now be implemented in an affordable, education-grade hardware platform.

Chapter 8. Concluding remarks and future work

8.1 Thesis summary

This Thesis has presented a snapshot of a journey of the author across the field of sensor engineering for ultrasonic NDE. The research has combined computer simulation, materials science, transducer design and signal processing, as well as conventional electronic and mechanical engineering. The highlights of the work are identified below.

Firstly, the engineering method of the core material of the sensor, the piezoelectric composite material, is considered. After a detailed literature review, the decision was made to build a detailed, yet concise, 3D numerical model of a classic 1-3 connectivity piezoelectric composite. The model was used to analyse conditions required for building a composite material of desirable properties for inclusion in an array configuration. This was particularly important in providing a baseline transducer configuration at the start of my research programme.

In parallel, a literature study of basic beam forming techniques indicated that it would be desirable for the NDE sensor arrays to follow the example of advanced phased array radar systems, which use hexagonal element sampling. Hence, in this Thesis, a detailed comparison between the effects of the two sampling regimes is presented. This then led to the previously developed sensor core modelling method being extended to include a novel, triangular cut 1-3 connectivity piezoelectric composite, suitable as a substructure for such a hexagonal sampled array. This Thesis shows that, contrary to what literature suggests, it is possible to select a set of material design parameters that produces a well-behaved, uni-modal, material with triangular substructure. This new material has been manufactured and successfully tested, validating the approach. Moreover, this triangular composite configuration has been explored for the production of 2D phased array transducers with hexagonal or sparse element layouts. 2D array transducers comprising hexagonal elements have been shown to possess comparable performance metrics (sensitivity, mechanical cross-talk, acceptance angle) when compared to a conventional dense square 2D configuration. Importantly, this comes with the additional bonus of either an increase in imaging resolution due to an increase in array aperture for set number of array elements or a reduction in array element count for given overall array aperture, which has implications

for the cost effectiveness of the NDE system in terms of both hardware and data management.

One of the major problems of miniaturisation of sensors is their naturally low signal output, which is difficult to transport over a cable. One solution to this problem is to integrate the analogue front-end of the phased array controller inside the probe assembly. In this Thesis, the problem is addressed by means of an in-probe miniature signal conditioning circuit. The novelty and practical benefit in this case lies in the fact that no high power components are needed in the probe body and an enhancement to both sensitivity and bandwidth of the received echo signals is demonstrated.

In conjunction with advancing ultrasonic array techniques, this Thesis considers the processing impact of large array datasets in terms of both computational efficiency and real-time practical implementation of NDE imaging. The recent emergence of GP-GPU computer architecture offered an opportunity to rethink the implementation of algorithms typically used in ultrasonic signal processing. This was applied to single-way beamforming, but also to more advanced approaches such as TFM and PCF beamforming. It is clear that parallelisation offers significant enhancement in computational throughput, but this work highlights the importance of understanding the mapping of a software problem onto the available hardware architecture. In effect, it becomes feasible for an NDE system to deliver real-time 3D volumetric imaging from 2D ultrasonic data.

The GP-GPU platform then permitted a new approach to solving the TimeofFlight problem. The unprecedented performance and low cost of this new approach enables industrial deployment of advanced beamforming methods, as well as development of practical CAD tools for engineering and education. A number of CAD tools have been implemented taking advantage of this new approach.

It is exciting to admit that these new tools have already borne impact in the field. Within CUE, an internal team has been established and tasked with preserving and enhancing the software platform, which is marketed using the acronym "*cueART*" – CUE Acoustic Research Toolbox. Secondly, promotion of these tools has resulted in a number of industrial research contracts for the University.

8.2 A personal perspective

This thesis represents a snapshot of the journey of learning and discovery for the author. Many faces of the ultrasonic probe technology have been touched at one point in the author's progress through the field. A number of topics attempted are promising and require more work to bring the effort to fruition.

In common with all scientific work, two factors have proven to be both extremely important and initially underestimated when conducting the work. The following should be taken as an advice for a prospective or novice researcher.

Firstly, it is very important to record and preserve thoughts, ideas, methods, software code (even if it is single-use) and research artefacts in a private, but searchable database; the notes should be made readable by others only as far as practicable. Over time, it is inevitable that many of the original ideas and techniques will evolve or fall out of favour altogether: this is normal. The premise for this is that researcher is not only someone that knows a lot about a given field, but also someone that does not know quite a bit more about other things. The knowledge quickly grows deep, but often not wide.

This brings up the second, more important point. It is utterly paramount to create and then nourish a social network of contacts across the field of research interest and beyond. One has to be very careful here: it is not humanly possible to be friends with hundreds of people at a time. It is necessary to choose allies. However, at the same time, application of "social norms", professionally crafted communications, a conscious effort and financial resource allocation all help to maintain the promising contacts that may just provide an invaluable insight into one's own field of study.

8.3 Proposed future work in probe technology

8.3.1 Electrical interconnect by soldering

One of the major limitations encountered when developing the 2D ultrasonic phased array probe was the technology of electrical interconnect, that is, the interface between the piezoelectric ceramic itself and the external signal cable. At the time of the beginning of this work, the prevailing method available in the group was hand soldering of thin wires directly to electrodes evaporated onto the piezoelectric substrate. This technique is relatively cheap

and undemanding equipment wise, however, as the probe central frequency rises and element size shrinks, it encounters several issues:

- The operation heats up the piezoelectric ceramic and the polymer filler between the pillars, which is undesirable
- The operation leaves a solder blob on the surface of the element. The solder blob, being a stiff metal bonded to the surface of the vibrating element, it presents mechanical loading affecting resonant properties of the element.
- The solder blob is of uneven size causing element-to-element variation. The result is that performance is inconsistent across elements of the same probe, and across probes.
- The operation is risky in a sense that should first attempt be unsuccessful, the rework is likely to lead to further deterioration

Bonding small wires to the surface is nothing new in the microelectronics industry, so the proposed solutions draw from that field:

- **Use of robotic arm**

A robotic arm could be used to guide the wire and deliver the solder iron to exact location as needed. The time needed to program the robot should be amortised over improved quality and amount of precision solder joints that such robot is able to execute, and ability to build multiple probes quicker.

- **Use of super low melting point metals as conductive tack**

There exist metal alloys that can have temperature of melting lower than 100deg.C, e.g. Wood's metal and Field's metal. These are not typically sold as solder, because of their poor mechanical and electrical properties. Electronics industry solder is generally expected to withstand temperatures above 150degrees Celsius, as this is a typical limit of operating temperature for semiconductors. However, for purpose of the piezoelectric probe, the fact that the metal can melt at temperatures as low as 40deg.C is not detrimental because the entire assembly is typically vacuum potted with epoxy resin, and therefore mechanically supported anyway. The potting material serves as the acoustic backing material and is

mechanically sturdy enough to hold the wires in place; the solder is only needed to provide electrically conductive contact. The resistance of the metal is of secondary concern as there is only a minimal amount of it used and the capacitive nature of the piezoelectric material should dominate the circuit anyway.

8.3.2 Electrical interconnect by thin printed circuit board

8.3.2.1 *Introduction*

Instead of soldering the connections, a dense ultrasonic array could be served better by using in-plane printed circuit board as a means of delivering the charge to the electrodes. This simple approach may be cost-efficient in mass production if successful.

A prototype has been developed and is presented in Fig.8.1 and Fig. 8.2. The PCB itself is very thin, bonded to the piezoelectric composite and therefore acoustically transparent. The electrodes for outermost elements fan out from the area towards the external connector. The electrodes from inner set of elements fan with ultra-thin tracks placed between the electrodes of exterior elements, and then towards the connector. The array element arrangement is spiral and semi-dense to allow for fanning out all the connections from inside.

The triangular-cut piezoelectric composite has to be precisely placed on the PCB to prevent short-circuits and align the pillars of the composite with the electrodes on the PCB. A special alignment rig (Fig. 8.3) has been developed to perform the alignment and deposition of the PCB over the composite, and the results of manufacture are presented in Fig. 8.4.

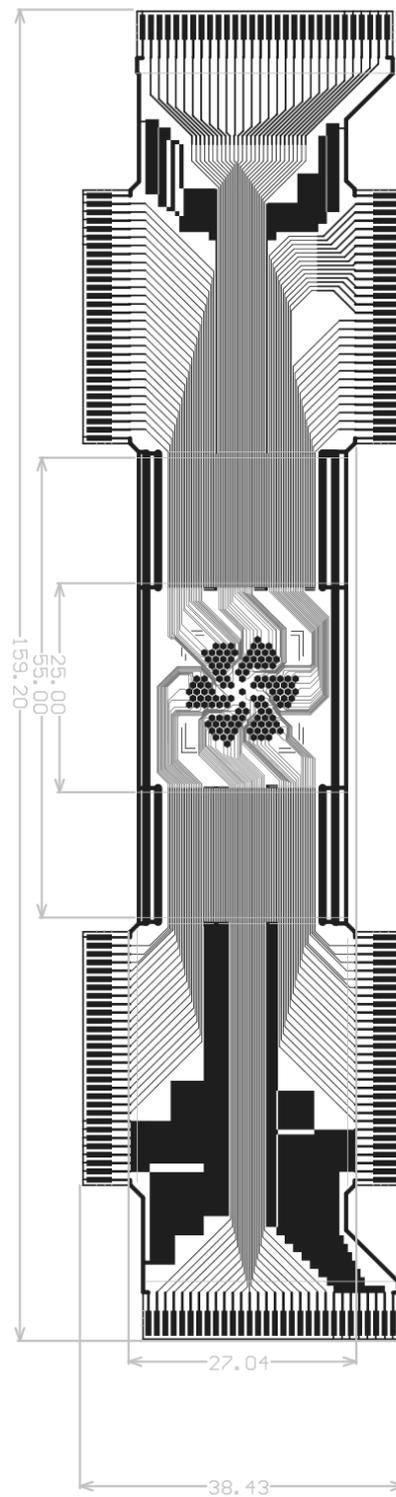


Fig.8.1. General view of the PCB layout. There are 6 30-way 1mm pitch edge male edge connectors. In the centre there is a 2D electrode array area with 128 pads. Dimensions are in mm.

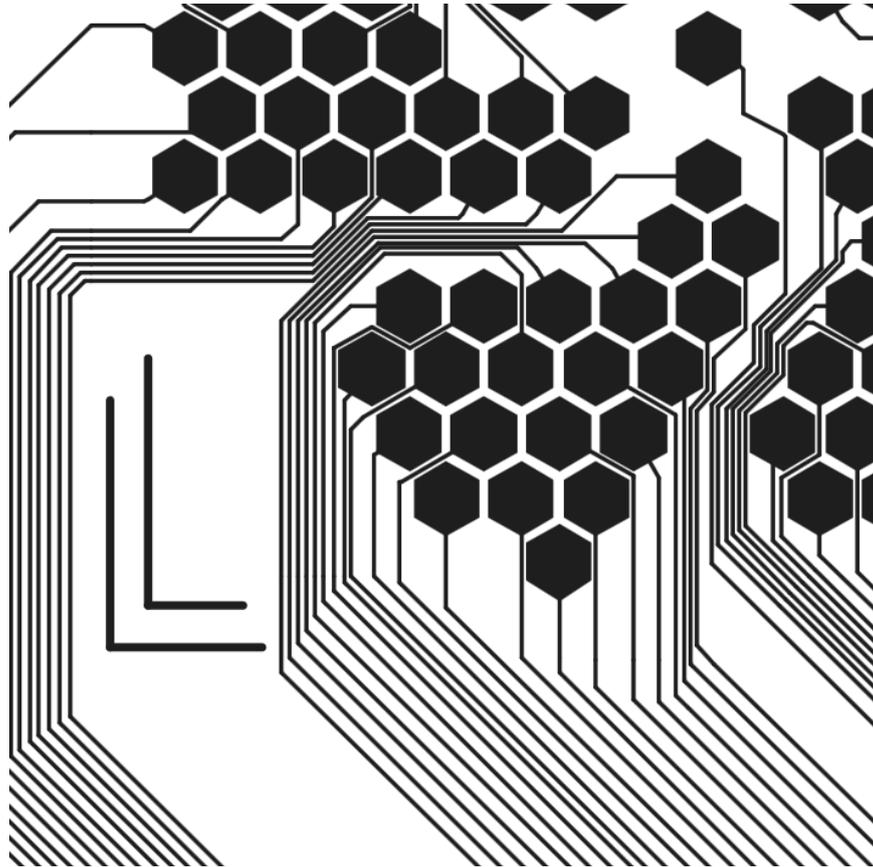


Fig. 8.2. Detail of the track routing over the piezoelectric composite area and the connection of tracks to the element electrodes.

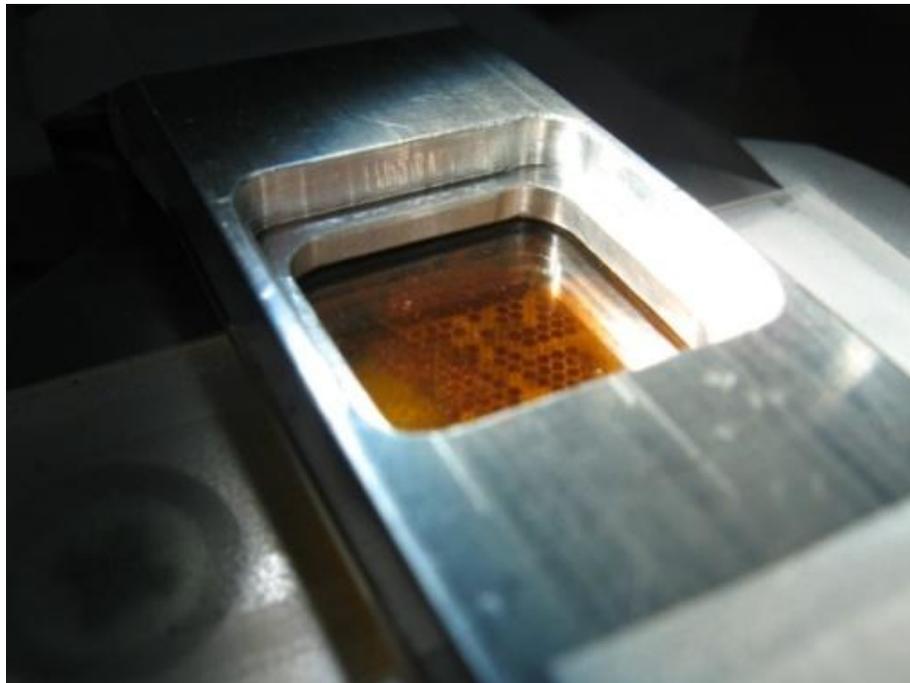


Fig. 8.3. View of the alignment and deposition rig. The thin flexi PCB is suspended under a bridge; underneath the bridge, the composite is manipulated with 4 degrees of freedom. The flexi PCB, being semi-transparent, allows observing and guidance of the composite below through a 3D microscope.

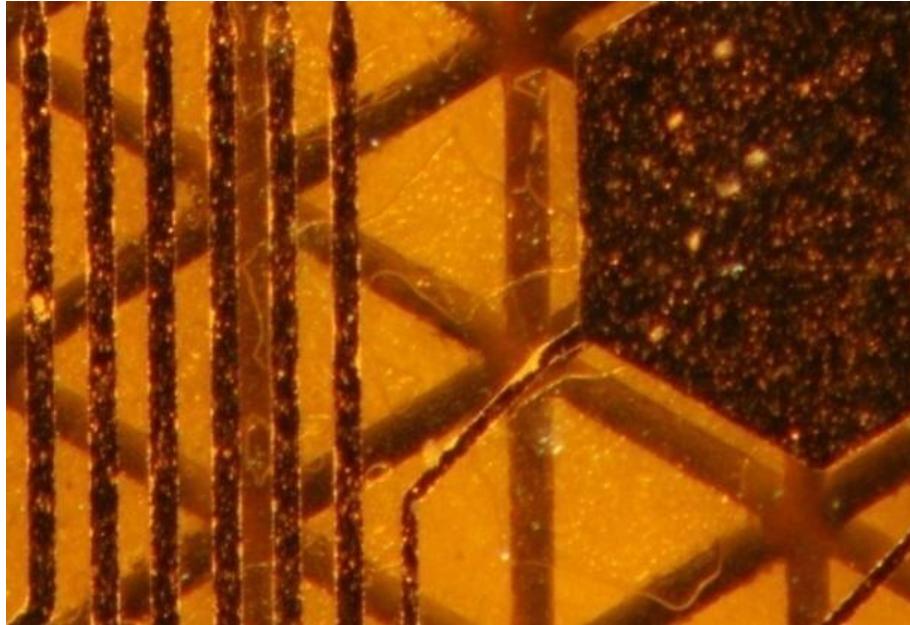


Fig. 8.4. Final inspection of the align accuracy of the electrode (dark patch at the front) over the ceramic pillars (light triangles at the back). Note that the element in this view is NOT the one in the laser vibrometry in the Fig. 8.7

8.3.2.2 *Electromechanical crosstalk*

Firstly, having electrical tracks so close to the piezoelectric ceramic itself makes them prone to electromechanical crosstalk. In the preliminary simulation in COMSOL (Fig. 8.5 and Fig. 8.6), and later confirmed by experiment (Fig. 8.7), it has been found that as long as the track is significantly thinner than the span of the piezoelectric ceramic pillar, the electrical field tends to not enter the pillar, and therefore there is no significant electromechanical crosstalk. This is further verified by the experimentally acquired pulse-echo response, as depicted in Fig. 8.8. The presented response contains a single mode vibration only, and no 'secondary pulse' characteristic of systems that exhibit mechanical cross-talk. The vibrating surface displacement under continuous excitation at resonant frequency has been measured using a scanning laser vibrometer and the result is presented in Fig. 8.7. This figure presents displacement amplitude on the left, and displacement phase (versus

excitation phase) on the right hand side of the screenshot. The measurement confirms theoretical predictions of low effect of the thin tracks. The overall benefit of the approach is that the PCB can be made extremely thin (single layer substrate, single layer copper, no vias, no solder mask; total thickness in range of 25-50 μm) and thus, can be acoustically transparent up to very high vibration frequencies.

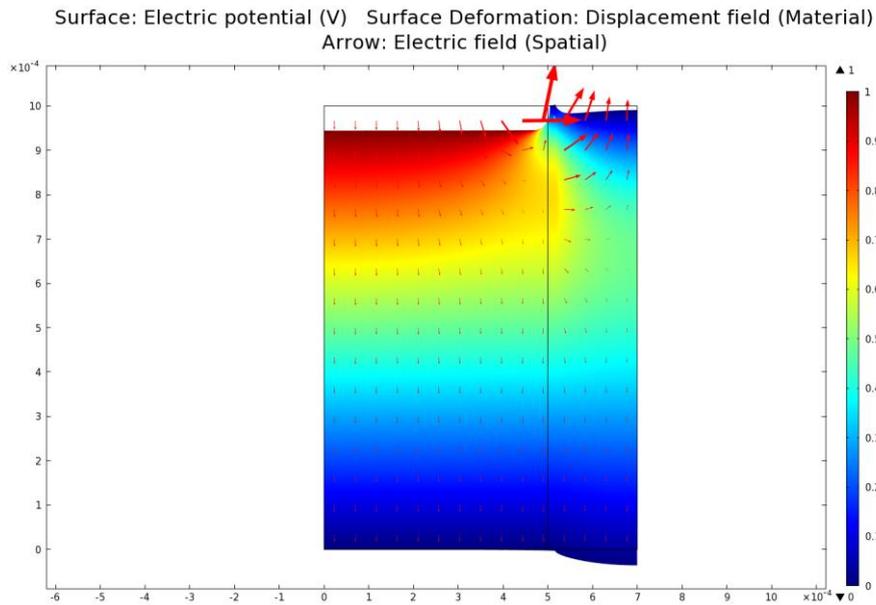


Fig. 8.5 Electrical field distribution (colour) and deformation of the piezoceramic pillar (left)-polymer kerf (right) under 1V electric field if the field is applied over the piezoceramic pillar. The electrical field gradient is distributed across the thickness of the piezoelectric ceramic phase. Deformation scale: $2 \times 10^5:1$

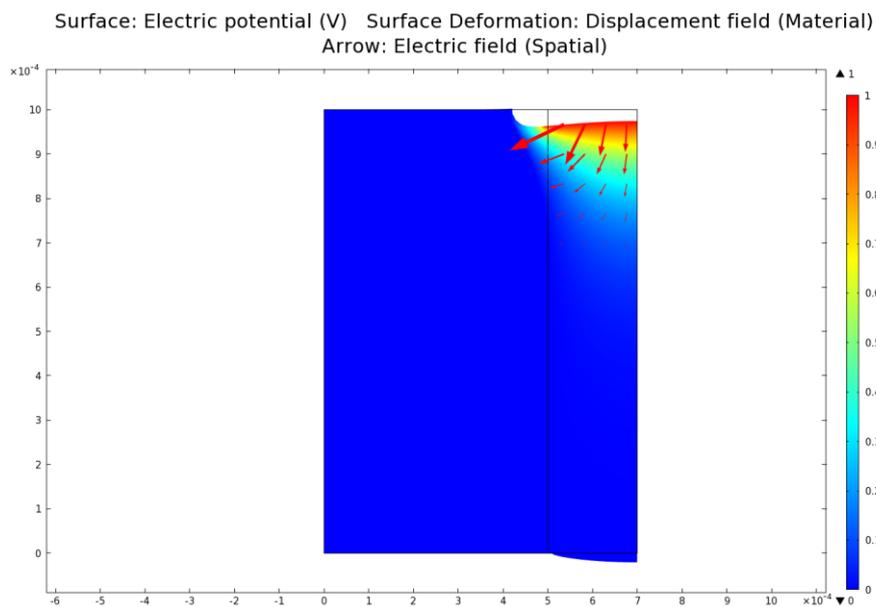


Fig. 8.6 Electrical field distribution (colour) and deformation of the piezoceramic pillar(left)-polymer kerf(right) under 1V electric field, when the electrical field is applied over the polymer kerf only(right). Portion of the electrical field does enter the ceramic and pulls the polymer in; however, the effect is very weak. Deformation scale: $1 \times 10^8:1$. Note that deformation scale is 1000x of the scale in Fig. 8.5.

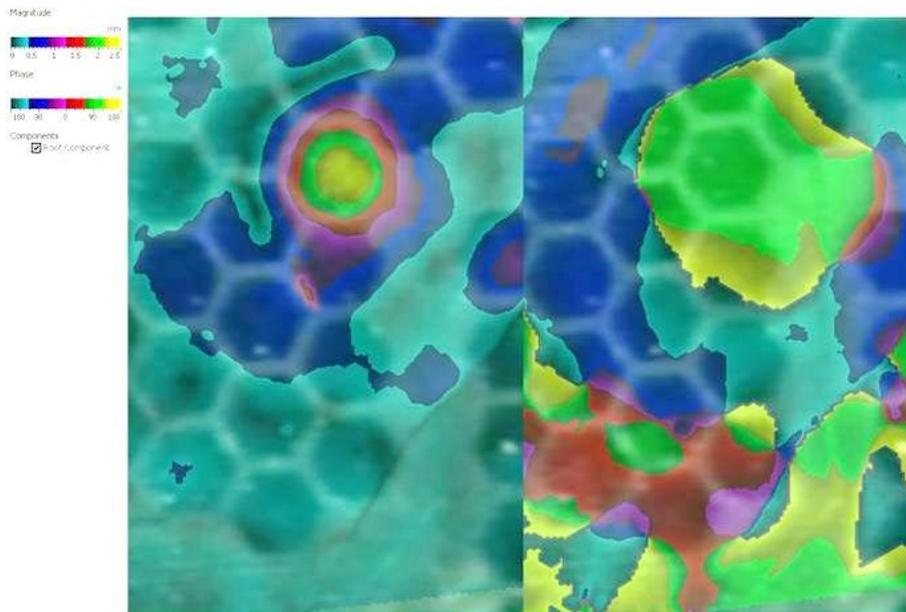


Fig. 8.7 Top-down view of displacement of the hexagonal element under thin-PCB electrode, as measured by laser vibrometry. Left: amplitude view; right: phase view. Dark hexagonal patches correspond to electrode locations. The colour indicated amplitude/phase and has been laid over a photo acquired through the laser vibrometer viewfinder. The PCB track delivering the charge to the hexagonal element does not excite pillars that do not belong to the element of interest.

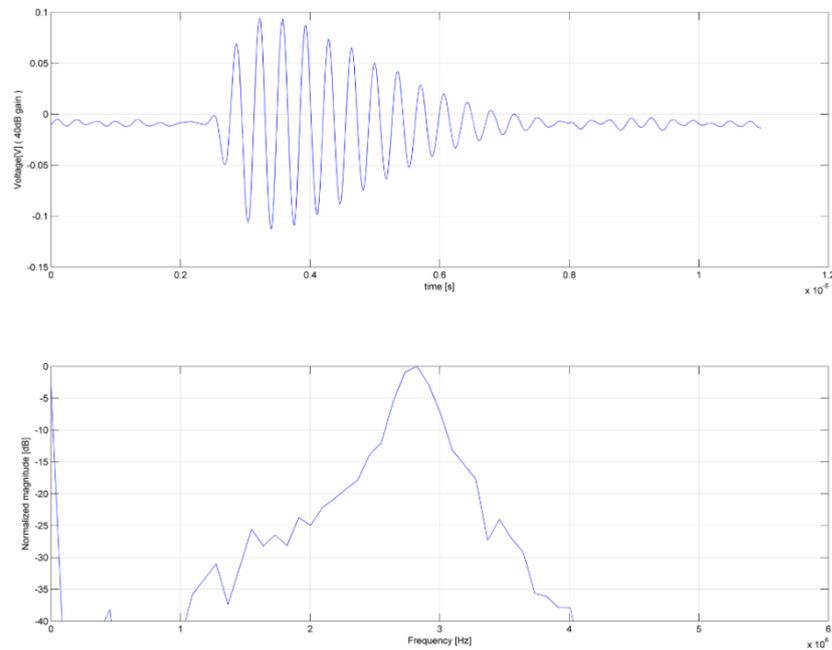


Fig. 8.8 Pulse-echo signal obtained with the prototype device. There signal is well-behaved and there is no signs of a crosstalk in the spectrum (no 'secondary pulse'). Please note that this was acquired with a device that has no acoustic backing, matching or electrical damping; therefore pulse length is not representative of a final device and serves to show that the device is operational only.

8.3.2.3 Electric field transfer

One of the issues with the prototype was to ensure the contact of the electrode with the piezoelectric ceramic pillar. The problem stems from the fact that there must be some kind of glue applied between the PCB and the piezoelectric ceramic; this glue will typically exhibit low electrical permeability, in effect acting as a series capacitor attenuating the electrical field. In the first prototype, this problem had not been addressed, yet the device was operational, although likely not at its full potential.

It is proposed that for future devices the glue should be loaded with metallic particles. Relatively low amount of well-mixed particles should be used to avoid short-circuit between the electrodes of different elements. For example, chromium metallic powder has been identified as suitable for its relatively high conductance, low acoustic impedance and the feature of sharp corners in the particles that will cause piercing and displacing of the glue layer when pressed.

8.3.2.4 Future work

The first prototype suffered from a bond line failure as a result of contact with couplant during pulse-echo testing.

It is proposed to resume the work, and build new complete devices with acoustic matching and backing to validate the concept.

8.3.3 Manufacture of graded acoustic impedance material using additive manufacturing principle

In the construction of the ultrasonic probe, the acoustic matching layer is critical for element for sensitivity and impulse response characteristic, and therefore for overall performance of the probe. It has been long established that continuously graded impedance materials are best for this purpose; however, the technology of making one is not well established. In particular, it appears that manufacturing the high impedance materials (close to the impedance of the piezoelectric ceramic) is of particular challenge.

Here, it is proposed to explore the area of robotised additive manufacturing. One can imagine a robotic cell with following effectors: tungsten dust dispenser, photo-active polymer spray (1 or 2 formulations for high and low impedance side), UV light, and mechanical press. In the process, the tungsten particles and polymer sprays are used to deposit a thin layer of material which is then compacted / compressed together using the press and cured using the UV light. Over the layers, the ratio of the tungsten to polymer mass per layer can change according to prescribed formula to form an effectively continuously varying acoustic impedance.

The robotisation of the process is needed because multiple very thin layers and precise amount of UV curing per layer are needed to obtain the desired effect; it is also a time-consuming process as low temperatures needed to obtain low epoxy shrinkage imply long cure times. It is proposed that pressing tungsten particles into half-cured polymer might help to relieve the internal stress build-up due to the polymer cure shrinkage.

8.3.4 Active array element with differential output

One problem that occurs when miniaturising the probe elements is that the smaller the piezoelectric element, the smaller the mechanical energy flux that can be converted into electrical signal power.

Here, it is proposed to not use the mechanical flux to create energy needed to deliver the electrical signal to the amplifier (passive sensing principle) but rather use the mechanical flux to modulate energy flow that is being continuously delivered to the sensor (active sensing principle). For ultrasound, this means an FET transistor with piezoelectric material in its gate; concept also known as piezoelectric transistor or piezoelectric resistor.

The attempts described in the literature have had good success with sensing ultra-small forces[96], and with very small 2D sensor arrays of 23 μ m pitch[97][98] but at the same time, limited success in sensing high frequency forces due to the high inherent capacitance of the FET transistor gate.

Here, it is proposed to arrange the piezoelectric FET transistors in such way as to transform and amplify the signal generated by the piezoelectric effect straight on the sensing die, as shown in the schematic ниже, in Fig. 8.9 and Fig. 8.10. The circuit is similar to the input stage of a differential amplifier or Wheatstone bridge. Importantly, the signals transferred

arrangement should allow combination of high operational frequency, high receiver density, low total mass (easy to operate) and utilisation of existing, relatively low cost off-the-shelf phased array controllers.

8.3.5 Stricter definition of the 'near field' extent

It has been shown that the contrast of the image depends heavily on the signal bandwidth of the probe's elements, or more specifically, on its impulse response characteristics; this is even more pronounced for sparse arrays.

As a next step in the theory of sparse array design, it is proposed to construct mathematical understanding of the extent of the near field as a function of the impulse response length. The idea is best explained by the following empirical diagram.

In the Fig. 8.11, a wave propagation has been simulated in homogenous medium of $v=1500\text{m/s}$, at $t=9\mu\text{s}$ after the wave emission. The figures represent a planar slice through the 3D space. There are two point like wave sources on the left of the figure; they emit a continuous wave or wave gated with $12\mu\text{s}$ or $2\mu\text{s}$ gate respectively. The diagram shows that the shorter the impulse response, the lower the volume of space where acoustic waves can interfere; therefore, the less opportunity for the waves to form undesired constructive interference known as side and grating lobes.

It is postulated that it should be possible to describe this phenomenon as a set of simplified analytical geometry equations, therefore facilitating such design of probes as to minimize the volume of space where undesired constructive interference occurs.

This problem is similar to the one faced in calculation of accuracy of global positioning systems. A review of literature from that area might be beneficial.

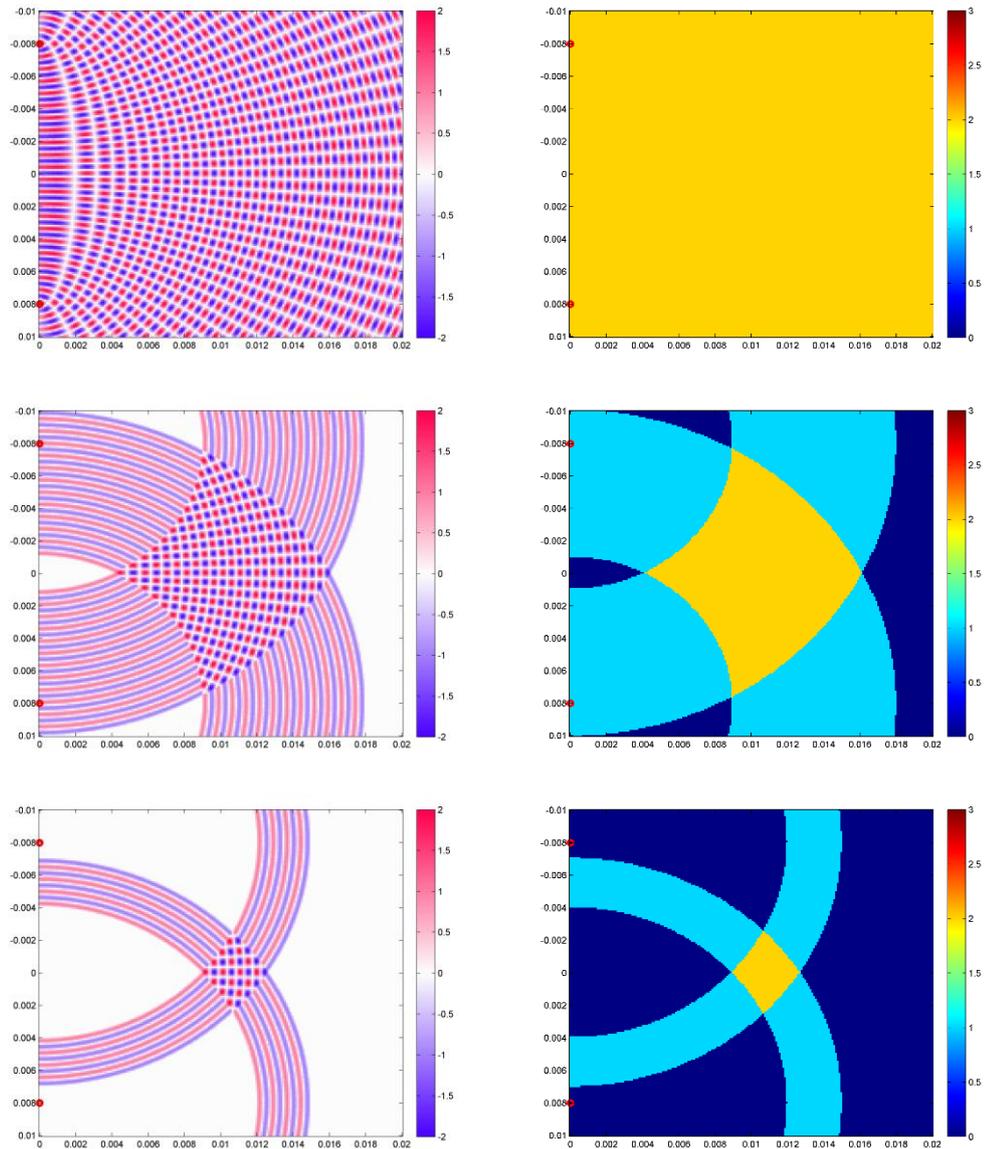


Fig. 8.11. Example of narrowing the side lobes when impulse response of the probe is shortened. Cross-section through the space in plane where the radiating elements are located. The radiating elements are to the left (red points). Top: continuous radiation; middle: intermediate impulse response length; bottom: short impulse response length. Left figures: pressure field view; right figures: interference count view (0, 1, or 2 times interfered).

8.3.6 Surface resolving beam forming for improved defect detection in difficult areas

The TFM (Total Focussing Method) is a now-established ultrasonic imaging technique, currently making its way into industrial use. However, one weakness of this technique

stems from assumption that the reflectors are point-like. Although reflectors of interest (small cracks) are often point-like, they are sometimes located near strong, non-point-like reflectors; e.g. back wall or other static geometrical features. These strong reflectors, being incorrectly resolved by standard TFM, interfere with reflectors of interest, making their detection difficult.

The problem is exemplified in a TFM image shown in Fig. 8.12 and Fig. 8.13. For this image, synthetic noise free FMC data has been created; for figure Fig. 8.12, the FMC data includes the planar reflector and a point reflector. For Fig. 8.13, the data excludes the planar reflector.

The situation is better explained by the geometrical construction shown in Fig. 8.14. In this construction, only 2 extreme and a middle phased array element have been used for clarity. Two A-scan rays are considered: The first A-scan is when transmitting on element A and receiving on element B, and the second A-scan is when transmitting on element A and receiving on element C. The rays follow the shortest path that adheres to the principle of reflection (The surface normal divides the angle between incident and reflected ray into two equal angles). The measured quantity is the Time of Flight (ToF) of the virtual ray of ultrasound between the transmitting and receiving element.

The TFM method then takes these ToFs and plots ellipses/ellipsoids (a shape where the sum of the distances between foci and the points on the ellipse is constant) in the imaging space. The intersection of the two ellipses marks the suspected location of the reflector.

This works well for point like reflector, but fails for surface like reflectors. The real surface reflector is not coincident with the intersection of these ellipses. The final result is that the surface reflectors appear blurred, not precisely placed in the final TFM image slightly above where the real reflecting surface is, as in Fig. 8.12. They also produce strong side lobes due to suboptimal destructive interference of the waves.

The effect is also clearly visible in other publications, e.g. [99] where TFM method, although it offers clearly superior Point-Spread-Function (PSF) for a point like reflector, it does not resolve the back wall location well.

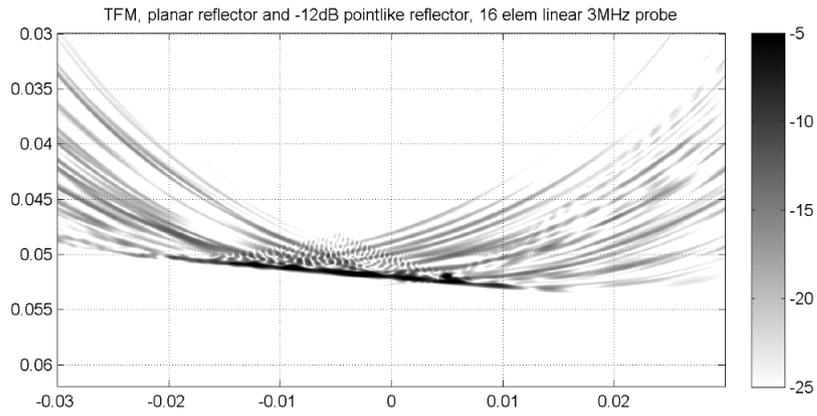


Fig. 8.12 Example TFM image of a planar reflector and a point reflector – synthetic data. The point reflector sought is buried within imaging artefacts.

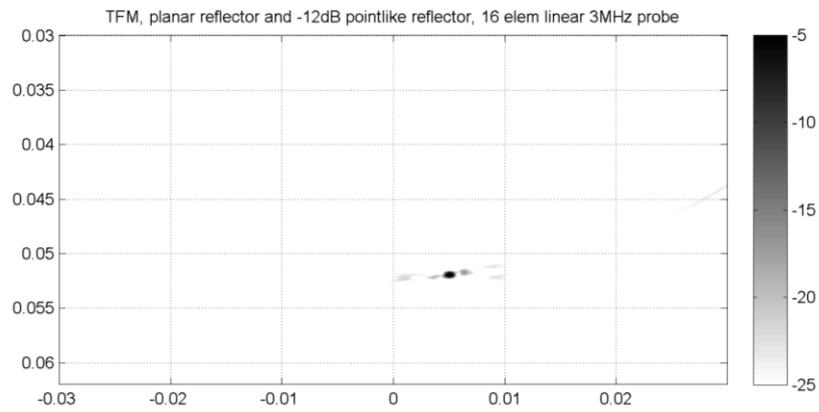


Fig. 8.13 Same situation as Fig. 8.12 but with the planar reflector removed. The point reflector is clearly identifiable.

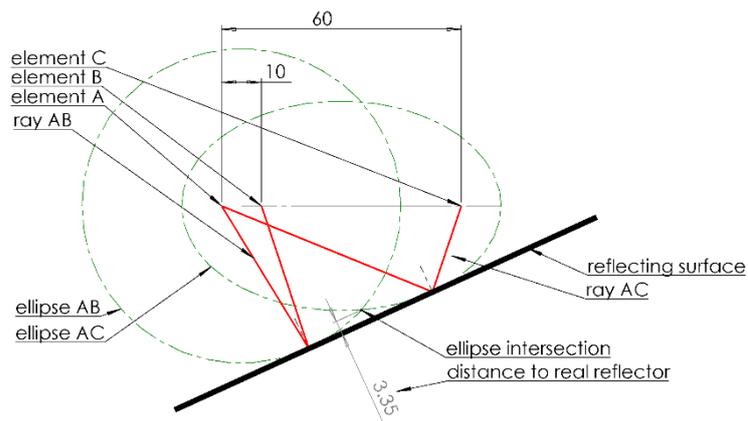


Fig. 8.14 Geometric construction to show the error of location when planar reflector is considered. Note that the ellipse intersection (TFM reflector solution) does not coincide with the real reflector.

Proposed solution

A technique is envisioned for resolving the surface-like reflectors from the Full Matrix Capture (FMC) data, and from there, separating the surface reflections from point-like reflections. The technique is based on solving an analytical geometry problem in order to find surface descriptors that fit the recorded ultrasonic reflections. This new information can be used in two ways. Firstly, the surface location can be correctly and sharply resolved. Secondly, the reflections then identified as coming from recognised surface can be separated from the FMC dataset; regular TFM is applied to the residue signals, finding point-like defect indications.

This new technique will allow improved near-surface defect detection. This is a capability that is highly sought in the NDE industry. Such a technique will be particularly useful in, but not limited to:

- In-production ball bearing inspection, where subsurface defects are of particular importance to the lifetime expectation of a bearing due to the Hertz stress fatigue,
- Inspection of multilayer composite materials, where surface-like reflectors are dominant,
- Inspection with sparse array probes, which suffer particularly bad from the non-point like reflector interference.

The proposed solution is to accumulate the signals from FMC into a reflector parameter space (b,m) similar as in Hough transform[100, p. 123], as presented in Fig. 8.15. The initial trials of prototype implementation are promising, as shown in Fig. 8.16.

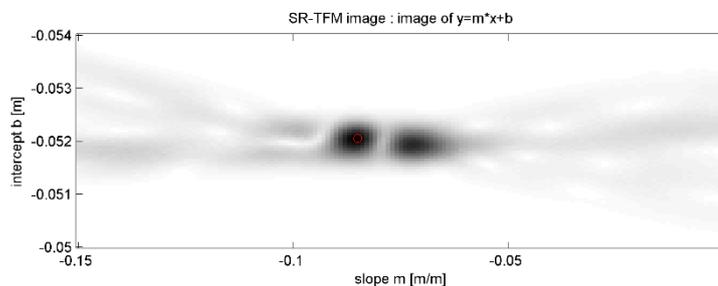


Fig. 8.15SR-TFM image.

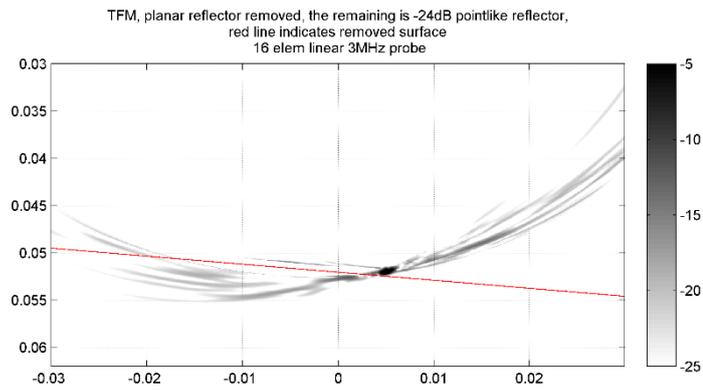


Fig. 8.16 Residue TFM image on the FMC data with the surface reflector suppressed, based on the SR-TFM image data. The red line indicates the recognized, suppressed surface.

8.4 Closing remarks

In the light of the author's experience, 2D phased arrays are not a generalisation, but rather a specialisation of linear ultrasonic phased arrays. When it comes to plethora of techniques available in NDT world, it appears that 2D phased arrays can only serve a niche not yet filled by cheaper 1D phased arrays or non-phased multi-element probes. Despite their high theoretical potential, the practical implementations fall short of the visionary 'ultrasonic holography camera' that can do 'nearly everything, instantly', as the sort of thing we come to expect from a commodity digital optical camera in 2010s. Instead, the probes and the corresponding control systems and software are still designed and specialised primarily for given application in mind, with only partial application overlap possible for a given design.

The probe element size, per-element sensitivity and bandwidth issues appear to be solvable by borrowing the manufacturing techniques from other fields, mainly from the microelectronics industry. Miniaturisation is the key. However, as it is, high initial cost of instrumentation and prototyping machinery is largely unavoidable. Therefore, companies with established markets and large research institutions are likely to lead the way in this field; for smaller establishments, cross-disciplinary networking and cooperation appears to be the way forward.

One common limitation of an ultrasonic system is the relatively large cost of a high channel count phased array probe controllers. At the time of the launch of this research, the big expectation was located in the development of sparse 2D arrays. Due to the limited capability of the CAD tools available at the time, it was believed that there exists some

exceptional element layout that will possess both the resolution of a large aperture probe and the contrast capability of regular, dense linear probe. However, over time it has been shown that even in near-optimal cases, one can only trade contrast (layout density) for resolution (overall aperture) of the image that a given probe style can provide.

Still, new nonlinear signal processing algorithms emerge (e.g. PCF[89], Capon[101], SASACI[102]) that are aware of the sparseness of the probe, and can potentially make sparse phased arrays more viable for use in NDE; undoubtedly, they will also benefit regular linear phased array probes.

One significant boundary of the development of 2D phased arrays seems to be raw quantity of data that is possible to acquire with the new sensors. It is not always the case that acquiring more data will yield a better final result. The most difficult part appears to be the making sense out of the raw data acquired. The problems include synthesizing the image out of physical signals, sorting out signals from noise, and then making a decision on what exactly a given image is saying about the specimen in question. Still, speed, or more importantly, latency of the signal processing stream is important for practical industrial applications and should not be neglected in early estimations of the cost of proposed processes.

The main problem that remains, from the industrial applicability side, is the big cost of qualifying any given solution for reliability in serving specific NDE scenario; the problem generates high costs (and associated risk of development) of a new inspection system that embodies the new and emerging signal processing techniques.

Because of this, new signal processing methods should be implemented in a portable, transferrable way. Importantly, they should be tested on a reliably comparable set of specimen and probe combinations, in a way that makes it easy to repeat the experiment for further validation and demonstration to industrial partners. For a new NDE signal processing method, the wide applicability of the method should be demonstrated as soon as possible in order to gain interest (and funding) from multiple sources.

The international researcher community must work out a common and reliable way of comparing the signal processing algorithms and probe types.

It is hoped that the work conducted for this Thesis will further the understanding of the 2D ultrasonic phased array probe technology for the betterment of the human condition.

References

- [1] "About NDT - What is NDT?" [Online]. Available: <http://www.bindt.org/What-is-NDT/>. [Accessed: 22-Jun-2015].
- [2] *NDT Yearbook 2013*. The British Institute of Non-Destructive Testing, 2013.
- [3] C. Fan, M. Pan, F. Luo, and B. Drinkwater, "Multi-frequency time-reversal-based imaging for ultrasonic nondestructive evaluation using full matrix capture," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 61, no. 12, pp. 2067–2074, Dec. 2014.
- [4] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "The post-processing of ultrasonic array data using the total focusing method," *Insight*, vol. 46(11), pp. 677–680, 2005.
- [5] L. A. International, *Piezoelectric Ceramics: Principles and Applications*. Mackeyville, PA: APC International, Ltd., 2011.
- [6] K. Shi, K. Que, and D. Guo, "Flexible ultrasonic phased-array probe," *Tsinghua Sci. Technol.*, vol. 9, no. 5, pp. 574–577, Oct. 2004.
- [7] *Introduction to Phased Array Ultrasonic Technology Applications*. .
- [8] *Advances in Phased Array Ultrasonic Technology Applications*. .
- [9] S. Ramadas, J. Jackson, J. Dziejewicz, R. O'Leary, and A. Gachagan, "Application of conformal map theory for design of 2-D ultrasonic array structure for ndt imaging application: a feasibility study," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 61, no. 3, pp. 496–504, Mar. 2014.
- [10] A. Tweedie, V. Murray, and G. Hayward, "Aperiodic and deterministic 2D phased array structures for ultrasonic imaging," in *Ultrasonics Symposium (IUS), 2009 IEEE International*, 2009, pp. 406–409.
- [11] O. Martínez-Graullera, C. J. Martín, G. Godoy, and L. G. Ullate, "2D array design based on Fermat spiral for ultrasound imaging," *Ultrasonics*, vol. 50, no. 2, pp. 280–289, 2010.
- [12] R. W. Martin, "Power Transfer Effects of Backing Impedance with Thickness Mode Ultrasonic Transducers Under Complex Conjugate Matching," in *Ultrasonics Symposium, 1977*, 1977, pp. 416–421.
- [13] G. Connolly and M. Jobst, "Demonstration of the Application of the Total Focusing Method to the Inspection of Steel Welds," presented at the ECNDT 2010, Moscow, 2010.
- [14] W. A. Smith, A. Shaulov, and B. A. Auld, "Tailoring the Properties of Composite Piezoelectric Materials for Medical Ultrasonic Transducers," in *IEEE 1985 Ultrasonics Symposium*, 1985, pp. 642–647.
- [15] L. S. Smith, W. E. Engeler, M. O'Donnell, and J. E. Piel, "Rectilinear phased array transducer using 2-2 ceramic-polymer composite," in *Ultrasonics Symposium, 1990. Proceedings., IEEE 1990*, 1990, pp. 805–808 vol.2.
- [16] G. Hayward, D. Gillies, and T. S. Durrani, "A Multidimensional Linear Systems Model of the Piezoelectric Transducer," in *IEEE 1984 Ultrasonics Symposium*, 1984, pp. 790–793.
- [17] W. A. Smith and B. A. Auld, "Modeling 1-3 composite piezoelectrics: thickness-mode oscillations," *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 38, no. 1, pp. 40–47, 2002.
- [18] W. A. Smith and B. A. Auld, "Modeling 1-3 composite piezoelectrics: thickness-mode

- oscillations," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 38, no. 1, pp. 40–47, Jan. 1991.
- [19] D. Robertson, G. Hayward, A. Gachagan, and V. Murray, "Comparison of the frequency and physical nature of the lowest order parasitic mode in single crystal and ceramic 2-2 and 1-3 piezoelectric composite transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 53, no. 8, pp. 1503–1512, 2006.
- [20] R. O'leary, G. Hayward, and V. Murray, "Finite Element Technique for the Assessment of 3-1 and 1-3 Connectivity Piezoelectric Composite Transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 54, no. 10, pp. 2024–2035, Oct. 2007.
- [21] P. Reynolds, J. Hyslop, and G. Hayward, "The influence of constructional parameters on the practical performance of 1-3 piezocomposite transducers," in *1996 IEEE Ultrasonics Symposium, 1996. Proceedings, 1996*, vol. 2, pp. 967–970 vol.2.
- [22] T. R. Gururaja, R. E. Newnham, K. A. Klicker, S. Y. Lynn, W. A. Schulze, T. R. Shrout, and L. J. Bowen, "Composite Piezoelectric Transducers," in *1980 Ultrasonics Symposium, 1980*, pp. 576–581.
- [23] L. J. Bowen and T. R. Gururaja, "High-frequency electromechanical properties of piezoelectric ceramic/polymer composites in broadband applications," *J. Appl. Phys.*, vol. 51, no. 11, pp. 5661–5666, Nov. 1980.
- [24] Y. Wang and B. A. Auld, "Acoustic Wave Propagation in One-Dimensional Periodic Composites," in *Proceedings of IEEE International Ultrasonics Symposium, 1985*, pp. 637–641.
- [25] J. Hossack, Y. Gorfou, and G. Hayward, "The modelling and design of composite piezoelectric arrays," in *Proceedings of IEEE International Ultrasonics Symposium, 1989*, pp. 793–796 vol.2.
- [26] F. R. Montero de Espinosa, V. Pavia, J. A. Gallego-Juarez, and M. Pappalardo, "Fractured Piezoelectric Ceramics for Broadband Ultrasonic Composite Transducers," in *IEEE 1986 Ultrasonics Symposium, 1986*, pp. 691–696.
- [27] P. Challande, "Optimizing ultrasonic transducers based on piezoelectric composites using a finite-element method," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 37, no. 3, pp. 135–140, May 1990.
- [28] J. A. Hossack and G. Hayward, "Assessment of different pillar geometries for 1-3 composite transducers using finite element analysis," in *Ultrasonics Symposium, 1990. Proceedings., IEEE 1990, 1990*, pp. 389–392 vol.1.
- [29] J. A. Hossack and G. Hayward, "Finite-element analysis of 1-3 composite transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 38, no. 6, pp. 618–629, Nov. 1991.
- [30] B. A. Auld and J. A. Hossack, "Distributed-period structures for suppression of spurious modes in 1-3 piezocomposites," *Electron. Lett.*, vol. 27, no. 14, pp. 1284–1285, Jul. 1991.
- [31] R. Hamilton and G. Hayward, "The design of low volume fraction 1-3 connectivity composite transducers using finite element modelling techniques," in *Ultrasonics Symposium, 1992. Proceedings., IEEE 1992, 1992*, pp. 531–534 vol.1.
- [32] J. Bennett, R. Hamilton, and G. Hayward, "Finite element modeling of 1-3 composite transducers for underwater applications," in *Ultrasonics Symposium, 1993. Proceedings., IEEE 1993, 1993*, pp. 1113–1116 vol.2.
- [33] G. Hayward and J. Bennett, "Assessing the influence of pillar aspect ratio on the behavior of 1-3 connectivity composite transducers," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 43, no. 1, pp. 98–108, Jan. 1996.
- [34] J. H. Ih and B. H. Lee, "Performance analysis of piezoelectric composite plates with consideration of the internal losses," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 35, no. 1, pp. 73–77, Jan. 1988.

- [35] J. E. Hyslop, J. T. Bennett, and G. Hayward, "An investigation into the design of high frequency two-dimensional arrays for ultrasonic imaging," in *Proceedings of IEEE International Ultrasonics Symposium*, 1994, vol. 3, pp. 1515–1518 vol.3.
- [36] A. Gachagan, J. T. Bennett, and G. Hayward, "A finite element modelling approach into the influence of mechanical matching and damping in 1-3 piezocomposites," in *1994 IEEE Ultrasonics Symposium, 1994. Proceedings*, 1994, vol. 2, pp. 995–998 vol.2.
- [37] D. Certon, F. Patat, F. Levassort, and L. Tessier, "Influence of the polymer and ceramic transverse stiffnesses in 1-3 piezocomposites on the lateral mode frequencies," in *1996 IEEE Ultrasonics Symposium, 1996. Proceedings*, 1996, vol. 1, pp. 531–534 vol.1.
- [38] R. C. . Silva, J. S. . Fonseca, F. R. . de Espinosa, A. Crumm, G. A. Brady, J. W. Halloran, and N. Kikuchi, "Optimal design of piezocomposite materials using topology optimization techniques and homogenization theory," in *Proceedings of IEEE International Ultrasonics Symposium*, 1997, vol. 2, pp. 883–886 vol.2.
- [39] J. Brown, E. Cherin, Jianhua Yin, and F. Foster, "Fabrication and performance of high-frequency composite transducers with triangular-pillar geometry," *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 56, no. 4, pp. 827–836, 2009.
- [40] Jianhua Yin, M. Lee, E. Cherin, M. Lukacs, and F. S. Foster, "High frequency piezocomposite transducer with hexagonal pillars," in *Ultrasonics Symposium (IUS), 2009 IEEE International*, 2009, pp. 2750–2753.
- [41] D. Fiore, R. Gentilman, H. Pham, W. Serwatka, P. McGuire, and L. Bowen, "Recent developments in 1-3 piezocomposite transducer fabrication," in *Proceedings of the Tenth IEEE International Symposium on Applications of Ferroelectrics, 1996. ISAF '96*, 1996, vol. 1, pp. 531–534 vol.1.
- [42] Y. Jiang, C. E. M. Demore, C. Meggs, C. Dunare, T. Stevenson, J. Bamber, S. Cochran, and T. W. Button, "Micro-moulded randomised piezocomposites for high frequency ultrasound imaging," 2012, pp. 1–4.
- [43] R. KRIMHOLTZ, D. A. LEEDOM, and G. L. Matthaei, "New equivalent circuits for elementary piezoelectric transducers," *Electron. Lett.*, vol. 6, no. 13, pp. 398–399, Jun. 1970.
- [44] G. Hayward and M. . Jackson, "A lattice model of the thickness-mode piezoelectric transducer," *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 33, no. 1, pp. 41–50, Jan. 1986.
- [45] D. M. Gordon Hayward, Z. N. Jeremmy Bennet, and Robin Hamilton, "A theoretical study on the influence of some constituent material properties on the behavior of 1-3 connectivity composite transducers," *Cell*, vol. 75, no. 2, pp. 241–251, 1993.
- [46] R. L. O'Leary and G. Hayward, "Investigation into the effects of modification of the passive phase for improved manufacture of 1-3 connectivity piezocomposite transducers," *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 46, no. 3, pp. 511–516, 2002.
- [47] G. Hayward and J. Hyslop, "Determination of lamb wave dispersion data in lossy anisotropic plates using time domain finite element analysis. Part II: application to 2-2 and 1-3 piezoelectric composite transducer arrays," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 53, no. 2, pp. 449–455, 2006.
- [48] R. L. O'Leary, A. C. S. Parr, A. Troge, R. A. Pethrick, and G. Hayward, "Performance of periodic piezoelectric composite arrays incorporating a passive phase exhibiting anisotropic properties," in *2005 IEEE Ultrasonics Symposium*, 2005, vol. 2, pp. 1073–1076.
- [49] E. Sharp, "A triangular arrangement of planar-array elements that reduces the number needed," *IRE Trans. Antennas Propag.*, vol. 9, no. 2, pp. 126–129, Mar. 1961.
- [50] N. Goto, "A synthesis of array antennas for high directivity and low sidelobes," *IEEE*

- Trans. Antennas Propag.*, vol. 20, no. 4, pp. 427–431, Jul. 1972.
- [51] R. C. Hansen, *Phased Array Antennas*, 2nd Edition. Wiley-Blackwell, 2010.
- [52] Y. Mendelsohn and E. Wiener-Avneer, “Simulations of circular 2D phase-array ultrasonic imaging transducers,” *Ultrasonics*, vol. 39, no. 9, pp. 657–666, Aug. 2002.
- [53] K. Heutschi, “Acoustics: Sound Field Calculations,” Swiss Federal Institute of Technology, Zurich, 25-Jan-2013.
- [54] J. A. Jensen, “Ultrasound fields from triangular apertures,” *Acoust. Soc. Am. J.*, vol. 100, no. 4, pp. 2049–2056, 1996.
- [55] “Lambert Azimuthal Equal-Area Projection -- from Wolfram MathWorld.” [Online]. Available: <http://mathworld.wolfram.com/LambertAzimuthalEqualAreaProjection.html>. [Accessed: 17-Jun-2010].
- [56] “‘Sphere to tangential plane’: polar (normal) aspect,” in *Map Projections*, Springer Berlin Heidelberg, 2006, pp. 161–207.
- [57] J. A. Brown, E. Cherin, J. Yin, and F. S. Foster, “Fabrication and performance of a high frequency geometrically focussed composite transducer with triangular pillar geometry,” *IEEE Trans Ultrason Ferroelec Freq Contr*, vol. Vol. 56, No.4, pp. 827–836, 2009.
- [58] G. Hayward and J. Bennett, “Assessing the influence of pillar aspect ratio on the behavior of 1-3 connectivity composite transducers,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 43, no. 1, pp. 98–108, Jan. 1996.
- [59] Jianhua Yin, M. Lee, J. Brown, E. Cherin, and F. S. Foster, “Effect of triangular pillar geometry on high- frequency piezocomposite transducers,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 57, no. 4, pp. 957–968, Apr. 2010.
- [60] H.-C. Yang, J. Cannata, J. Williams, and K. K. Shung, “Crosstalk reduction for high-frequency linear-array ultrasound transducers using 1-3 piezocomposites with pseudo-random pillars,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 59, no. 10, p. -, Oct. 2012.
- [61] “peakdet: Peak detection using MATLAB (non-derivative local extremum, maximum, minimum).” [Online]. Available: <http://www.billauer.co.il/peakdet.html>. [Accessed: 16-Dec-2014].
- [62] “Susceptance and Admittance : Reactance And Impedance -- R, L, And C - Electronics Textbook.” [Online]. Available: <http://www.allaboutcircuits.com/textbook/alternating-current/chpt-5/susceptance-and-admittance/>. [Accessed: 22-Jun-2015].
- [63] D. Certon, N. Felix, E. Lacaze, F. Teston, and F. Patat, “Investigation of cross-coupling in 1-3 piezocomposite arrays,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 48, no. 1, pp. 85–92, Jan. 2001.
- [64] “IEEE Standard on Piezoelectricity,” *ANSI/IEEE Std 176-1987*, p. 0_1–, 1988.
- [65] H. Zumbahlen and Analog Devices, inc, Eds., *Linear circuit design handbook*. Amsterdam ; Boston: Elsevier/Newnes Press, 2008.
- [66] J. Camacho and C. Fritsch, “Protection circuits for ultrasound applications,” *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 55, no. 5, pp. 1160–1164, 2008.
- [67] N. C. Chaggares, R. K. Tang, P. A. N. Sinclair, P. F. S. Foster, K. Haraierciwz, and B. Starkoski, “Protection circuitry for high frequency ultrasonic NDE,” *AIP Conf. Proc.*, vol. 509, no. 1, pp. 1987–1994, Maj 2000.
- [68] G. Moore, “Cramming more components onto integrated circuits,” *Electronics*, pp. 114–117, Apr. 1965.
- [69] “CUDA Toolkit Documentation.” [Online]. Available: <http://docs.nvidia.com/cuda/index.html>. [Accessed: 29-Dec-2014].
- [70] “Tag: Papers :: GPGPU.org.” [Online]. Available: <http://gpgpu.org/tag/papers>.

- [Accessed: 18-Apr-2015].
- [71] "Field II Ultrasound Simulation Program." [Online]. Available: <http://field-ii.dk/>. [Accessed: 29-Dec-2014].
- [72] "MEX-File Creation API - MATLAB & Simulink - MathWorks United Kingdom." [Online]. Available: <http://uk.mathworks.com/help/matlab/call-mex-files-1.html>. [Accessed: 31-Jan-2015].
- [73] J. Nickolls, I. Buck, M. Garland, and K. Skadron, "Scalable Parallel Programming with CUDA," *Queue*, vol. 6, no. 2, pp. 40–53, Mar. 2008.
- [74] M. Parrilla, J. Brizuela, J. Camacho, A. Ibanez, P. Nevado, and C. Fritsch, "Dynamic focusing through arbitrary geometry interfaces," in *IEEE Ultrasonics Symposium, 2008. IUS 2008*, 2008, pp. 1195–1198.
- [75] "Wolfram Research: Mathematica, Technical and Scientific Software." [Online]. Available: <http://wolfram.com/>. [Accessed: 15-May-2011].
- [76] M. Baboulin, A. Buttari, J. Dongarra, J. Kurzak, J. Langou, J. Langou, P. Luszczek, and S. Tomov, "Accelerating scientific computations with mixed precision algorithms," *Comput. Phys. Commun.*, vol. 180, no. 12, pp. 2526–2533, 2009.
- [77] H. Wong, M. M. Papadopoulou, M. Sadooghi-Alvandi, and A. Moshovos, "Demystifying GPU microarchitecture through microbenchmarking," in *Performance Analysis of Systems & Software (ISPASS), 2010 IEEE International Symposium on*, 2010, pp. 235–246.
- [78] J. A. Nelder and R. Mead, "A Simplex Method for Function Minimization," *Comput. J.*, vol. 7, no. 4, pp. 308–313, Jan. 1965.
- [79] "IEEE Standard for Floating-Point Arithmetic," *IEEE Std 754-2008*, pp. 1–70, Aug. 2008.
- [80] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "Post-processing of the full matrix of ultrasonic transmit-receive array data for non-destructive evaluation," *NDT E Int.*, vol. 38, no. 8, pp. 701–711, Apr. 2005.
- [81] M. Lewandowski, P. Karwat, J. Kudelka, and T. Kleczek, "GPU Implementation of the STA Algorithm on I/Q Data," in *Ultrasonics Symposium (IUS)*, 2012.
- [82] M. Weston, P. Mudge, C. Davis, and A. Peyton, "Time efficient auto-focussing algorithms for ultrasonic inspection of dual-layered media using Full Matrix Capture," *NDT E Int.*, 2011.
- [83] J. Lambert, A. Pedron, G. Gens, F. Bimbard, L. Lacassagne, and E. Iakovleva, "Performance evaluation of total focusing method on GPP and GPU," in *Design and Architectures for Signal and Image Processing (DASIP), 2012 Conference on*, 2012, pp. 1–8.
- [84] C. Fritsch, M. Parrilla, A. Ibanez, R. C. Giacchetta, and O. Martinez, "The progressive focusing correction technique for ultrasound beamforming," *Ultrason. Ferroelectr. Freq. Control IEEE Trans. On*, vol. 53, no. 10, pp. 1820–1831, 2006.
- [85] E. W. Weisstein, "Affine Transformation," *Affine Transformation -- from Wolfram MathWorld*. [Online]. Available: <http://mathworld.wolfram.com/AffineTransformation.html>. [Accessed: 22-Jun-2015].
- [86] E. W. Weisstein, "Chebyshev Polynomial of the First Kind," *Chebyshev Polynomial of the First Kind -- from Wolfram MathWorld*. [Online]. Available: <http://mathworld.wolfram.com/ChebyshevPolynomialoftheFirstKind.html>. [Accessed: 22-Jun-2015].
- [87] P. Borwein and T. Erdelyi, *Polynomials and Polynomial Inequalities*, 1995 edition. New York: Springer, 1995.
- [88] I. G. Pettigrew, K. J. Kirk, S. Cochran, and M. Sensors, "Rapid distributed data collection and processing with arrays – the next step beyond full waveform capture."
- [89] J. Camacho, M. Parrilla, and C. Fritsch, "Phase Coherence Imaging," *Ultrason.*

- Ferroelectr. Freq. Control IEEE Trans. On*, vol. 56, no. 5, pp. 958–974, May 2009.
- [90] M. Li and G. Hayward, "Ultrasound Nondestructive Evaluation (NDE) Imaging with Transducer Arrays and Adaptive Processing," *Sensors*, vol. 12, no. 12, pp. 42–54, Dec. 2011.
- [91] N. Trefethen, "Six myths of polynomial interpolation and quadrature." Oxford University.
- [92] M. F. Wagdy and M. S. Lucas, "Errors in sampled data phase measurement," *Instrum. Meas. IEEE Trans. On*, vol. 34, no. 4, pp. 507–509, 1985.
- [93] A. Dall'Osso, "Computer algebra systems as mathematical optimizing compilers," *Sci. Comput. Program.*, vol. 59, no. 3, pp. 250–273, Feb. 2006.
- [94] "Computationally efficient method of obtaining Time of Flight for 3D arbitrary curvature refracted-ray Total Focussing Method on GP-GPU processors - IEEE paper supplement website." [Online]. Available: <http://dziewierz.pl/cueTFM/>. [Accessed: 14-Mar-2015].
- [95] "Voreen." [Online]. Available: <http://www.uni-muenster.de/Voreen/>. [Accessed: 25-Jun-2015].
- [96] P. Fei, P.-H. Yeh, J. Zhou, S. Xu, Y. Gao, J. Song, Y. Gu, Y. Huang, and Z. L. Wang, "Piezoelectric Potential Gated Field-Effect Transistor Based on a Free-Standing ZnO Wire," *Nano Lett.*, vol. 9, no. 10, pp. 3435–3439, Oct. 2009.
- [97] B. Lemke, R. Baskaran, S. Ganapathysubramanian, and O. Paul, "Stress distribution under electroless nickel bumps extracted using arrays of 7 piezo-FETs," in *2010 IEEE Sensors*, 2010, pp. 2573–2576.
- [98] M. Doelle, C. Peters, P. Ruther, and O. Paul, "Piezo-FET stress-sensor arrays for wire-bonding characterization," *J. Microelectromechanical Syst.*, vol. 15, no. 1, pp. 120–130, Feb. 2006.
- [99] C. Holmes, B. W. Drinkwater, and P. D. Wilcox, "The post-processing of ultrasonic array data using the total focusing method," *Insight*, vol. 46(11), pp. 677–680, 2005.
- [100] D. H. Ballard and C. M. Brown, *Computer Vision*, First edition. Englewood Cliffs, N.J: Prentice Hall, 1982.
- [101] S. D. Somasundaram, "Linearly Constrained Robust Capon Beamforming," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5845–5856, Nov. 2012.
- [102] T. Lardner, M. Li, R. Gongzhang, and A. Gachagan, "A new speckle noise suppression technique using cross-correlation of array sub-apertures in ultrasonic NDE of coarse grain materials," presented at the American Institute of Physics Conference Series, 2013, vol. 1511, pp. 865–871.

Appendix A. Snell's law of refraction and polynomial fit equations processed using Wolfram Mathematica

Due to the sheer volume of this appendix, it is attached to this thesis as a digital version only – to be found on Thesis DVD or companion website [94]

Appendix B. About Matlab's `fminsearch`

Introduction

Matlab's `fminsearch` is an implementation of Nelder-Mead method for minimising a function of arbitrary finite dimensionality. It is sometimes referred to as simplex method, because the function minimisation operates a simplex in the domain of the function. Simplex is an N-dimensional geometric shape made up of points, with the number of points being the dimension count plus 1. Value of the function is evaluated for each corner of the simplex, and the simplex is then modified towards minimum of the function.

The main advantage of this method is that it does not require calculation of derivatives, and because of that, it performs well for poorly conditioned functions in a finite-precision calculation environment. In performance terms, this method is often not faster than derivative-based methods, however, it is typically more stable for a wide range of functions.

Matlab implementation is available with source code. It is coded in a 'defensive' way, supports arbitrary number of dimensions and can provide exhaustive diagnostics of the process.

However, all this means that it is a relatively slow implementation. If the function to be minimised is relatively simple, the implementation spends most of the time in the flow control and rare border line case checks. This prompted a re-write optimised for a specific 2D case found in the author's research, which was successful. Additionally, the following was found:

Incorrect behaviour for some inputs

For purpose of this work, the `fminsearch` is mostly used to find the spatial point where the ray is refracted when transcending the boundary between two propagation media, according to Fermat's principle. The point is found by minimising the ray flight time as a function of the location of the point on the surface between the two media.

For such a problem, it is possible to find a good first approximation for the point of refraction - it is a point lying in an intersection of the straight line between the source and destination point and the surface of the media interface.

During the course of this work, it appeared that the solution is not found correctly for some cases.

A deeper investigation led to a common programming error when dealing with floating point number:

```
in fminsearch.m, line 259:
```

```
if y(j) ~= 0 ...
```

This part of the code sets up an 'initial simplex' - a set of points that are used to evaluate the function to minimise. In this code, \bar{y} is the initial guess supplied by the user, and j is an iterator over its dimensions. If the $\bar{y}(j)$ is not zero, the simplex corners are set up to be a percentage of the initial guess point (hard coded to be $\bar{y}(j)*(1+0.05)$). If the $\bar{y}(j)$ is zero, the simplex corner is set at a fixed distance from the initial guess, the value of which being hard-coded to be 0.00025 .

The problem with such approach is that it can lead to bad initial simplex in both cases.

The first case, when $\bar{y}(j)$ is compared for equality with a fixed number, is an example of bad programming practice. This code will work as intended only if $\bar{y}(j)$ is equal to a binary representation of zero; if, however, the floating-point number is based on a real-world input, the value of $\bar{y}(j)$ could be extremely small instead, for example $1e-50$.

For engineering and scientific calculation purposes, such small value means "zero", however it is not zero for this code.

Other sources of such a result may include finite-precision operations on numbers of significantly different magnitude. For example, a series of operations in MATLAB:

```
a=single(1e-1)+single(1e6)-single(1e6)-single(1e-1)
```

yields $a=0.0250$ and not as one would expect, $a=0$. Similar or worse examples can be found for numbers stored in 'double' format or, in fact, any non-infinite precision format.

If $\bar{y}(j)$ is not equal to binary zero, the first code path is executed - the one based on a percentage approach. Since the value is extremely small, adding 5% of its value still yields

an extremely small value. Upon this new point, the cost function is evaluated, and might return numerically identical value, or a value that is less than `To1F` (the user-supplied tolerance on the cost function value) and different from the others; the simplex becomes singular, and the following code correctly decides that iteration cannot continue. In effect, the function is not being minimised at all, simply returning the initial guess.

In the second case, when one of the coordinates of initial guess, `y(j)`, is equal to binary zero, the second code path is executed. Corners of the simplex must be moved from initial guess point by a certain distance to allow exploration of the problem space. The current implementation uses a hard-coded value of `0.00025`. This value can be good for some problems, but can at the same time be too big or too small for others, leading to either unstably large or singular simplex.

In both of these cases, a better approach is to utilise the additional information about nature of the problem to be minimised provided by user-supplied tolerances on the result. For example, a good spread of the simplex initial points could be a low multiple of `To1X`; in case if this fails to achieve a non-singular simplex, a spread could be increased exponentially until the difference in function values becomes at least larger than `To1F`. Such modification would much improve the chance that the problem space would at last begin to be explored.

The proposed modification is to replace the code:

```
usual_delta = 0.05; % 5 percent deltas for non-zero terms
zerotermdelta = 0.00025; % Even smaller delta for zero
elements of x
for j = 1:n
    y = xin;
    if y(j) ~= 0
        y(j) = (1 + usual_delta)*y(j);
    else
        y(j) = zerotermdelta;
    end
    (...)
```

with code:

```
usual_spread = 10;
```

```
for j = 1:n
y = xin;
y(j) = y(j)+usual_spread*TolX
end
(...)
```

Alternatively, the original factors `usual_delta` and `zero_term_delta` could be modified to become user-selectable parameters that are optionally supplied using `optimset`. In any case, the equality operator if `y(j) == 0` should not be used due to its incompatibility with engineering representation of 'zero'.