ULTRASONIC INSPECTION
OF
HIGHLY SCATTERING
MATERIALS

by

Anton Van Pamel

A thesis submitted for the degree of

Doctor of Engineering

Imperial College London

Department of Mechanical Engineering

October 2015

Abstract

Ultrasonic Non-Destructive Evaluation (NDE) relies on the scattering of waves from discontinuities, such as fractures or voids, to probe media otherwise invisible to the naked eye. Whilst this has been industrially exploited for several decades within acoustically transparent materials, many materials maintain a microstructure that causes scattering of the propagating waves. This undermines the aforementioned premise as it becomes exceedingly difficult to discern the features of interest from the scattering inherent to microstructural features, thereby limiting the range of materials which can be reliably inspected, non-destructively.

Experimental investigations confirm the challenges and significant shortcomings for the inspection of future industrial components where such microstructures are desirable for their mechanical properties. It is demonstrated that the rapid increases in scattering with the insonifying frequency severely limit the achievable sensitivity of conventional ultrasound techniques.

A review of the latest advances in ultrasound technology, including signal processing and imaging algorithms, explore the opportunities to exceed current limitations and advance the capability of ultrasonic NDE.

Establishing these advances, and those of future approaches, requires a rigorous definition of performance. In contrast to commonly adopted strategies, a novel strategy which considers the probabilities of detection and false alarms is proposed as a valuable benchmark that can be used to make objective comparisons in terms of performance between competing algorithms.

Future progress will also rely on a better scientific understanding of scattering, which can be provided by powerful modelling tools. Here, Finite Element modelling is established to be very useful; it captures the complex scattering physics and allows an investigative flexibility which can provide extremely useful insights.

Whereas previous studies have often been restricted to weak scattering assumptions, the present FE modelling capability now enables the study of more complex, highly scattering environments. This is demonstrated by investigating ultrasonic arrays, where through optimising their engineering, especially in terms of their configuration, significant performance enhancements are shown to be possible.

These important scientific tools have enabled the assessment of the latest imaging algorithms, the optimisation of inspection configurations, and increased our understanding of scattering phenomena. Their use in the future enables wide possibilities towards further pursuing the ultrasonic inspection of highly scattering materials.

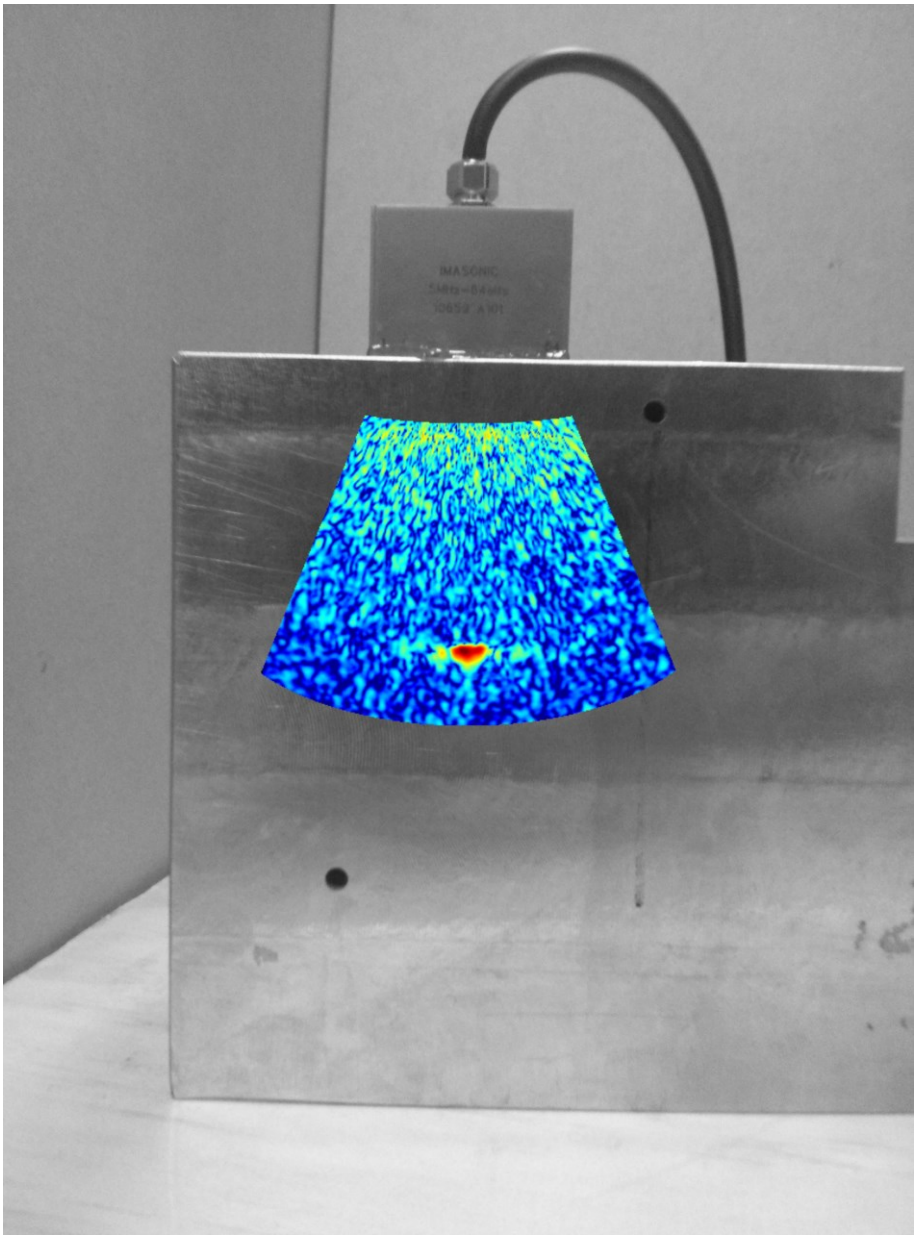

Ultrasonic array image of a cylindrical defect (red) within a highly scattering material.

Voor mijn vader en moeder.

Acknowledgements

I would like to express my gratitude to all persons who helped me realise this work.

I'm especially grateful to my supervisors: Prof Mike Lowe for his mentorship and inspiration without which this work would not have been possible; Dr Colin Brett for his continued support and guidance.

For all the impromptu meetings which were very helpful, I'd like to thank Dr Peter Huthwaite.

I have also benefited greatly from stimulating discussions with Prof Peter Nagy – an acknowledgement which now spans a generation.

Finally, I'm indebted to all the members of the NDE group at Imperial College. When I first joined the group I could hardly even recognise an oscilloscope, let alone use one. For bettering that and many other of my shortcomings, special thanks to: Dr Remo Ribichini, Dr Wonjae Choi, Dr Nicholas Brierley, Dr Bo Lan, Dr Attila Gajdacsi, Dr Xi Xiaou, Dr Vatche Attarian, Dr Eli Leinov, Matthias Seher, Fan Shi, Julio Isla, Francisco Hernando Quintanilla, Gabor Gubicza, and Joe Corcoran.

Declaration of Originality

The content of this thesis is the result of independent work carried out by myself under the supervision of Prof Michael Lowe. Appropriate references have been provided wherever the work of others has been used.

Anton Van Pamel

Copyright Declaration

CONTENTS

Chapter I

INTRODUCTION

Problem Definition

1.1 Industrial Motivation

Approximately four-fifths of world energy demand is currently supplied by thermal power stations (IEA 2014). Independent of the driving fuel - whether coal, gas, nuclear, oil, or renewables including geothermal and solar thermal electric - the majority of electricity generation worldwide is thus fundamentally limited by the same thermodynamic laws. In practise these laws restrict the highest achievable generation efficiency by the maximum power plant operating temperature. Enabling higher temperatures therefore presents an opportunity to progress our energy generation capability, universally.

During the last five decades however, the evolution of power plant operation temperature has stagnated. Scarcity of materials with sufficient high-temperature resilience have limited plants to operate around 565°C (see Figure 1.1a) and hence exceeding this barrier will require not only superior engineering materials but also economically feasible solutions.

Engineers faced a similar challenge almost half a century ago when the early development of Sir Frank Whittle's jet engine was hampered by the need to accommodate extreme temperatures. Eventually, nickel alloys proved ground-breaking and today, can make up half of a modern aero-engine. Currently, government research programs (e.g. US DOE Advanced Boiler and Steam Program and Europe's AD700) have identified similar alloys, known as Inconels (Special Metals 2013), to realise next-generation power plants. By replacing present steels, super-alloys could conceptually enable power plants to operate at "ultra-super-critical" (USC) conditions (Bugge et al. 2006), close to 700°C, raising their efficiency from approximately 36% to 50+% (see Figure 1.1b) and reducing specific emissions by almost half.

Despite this superior strength, one of the remaining obstacles is the associated increase in cost (Phillips & Wheeldon 2010). In comparison to aerospace applications, the sheer volume of nickel required for a single USC plant presents serious economic challenges. One way to reduce these costs is by prolonging the service life of components through diagnosing their structural health using Non-Destructive Evaluation (NDE).

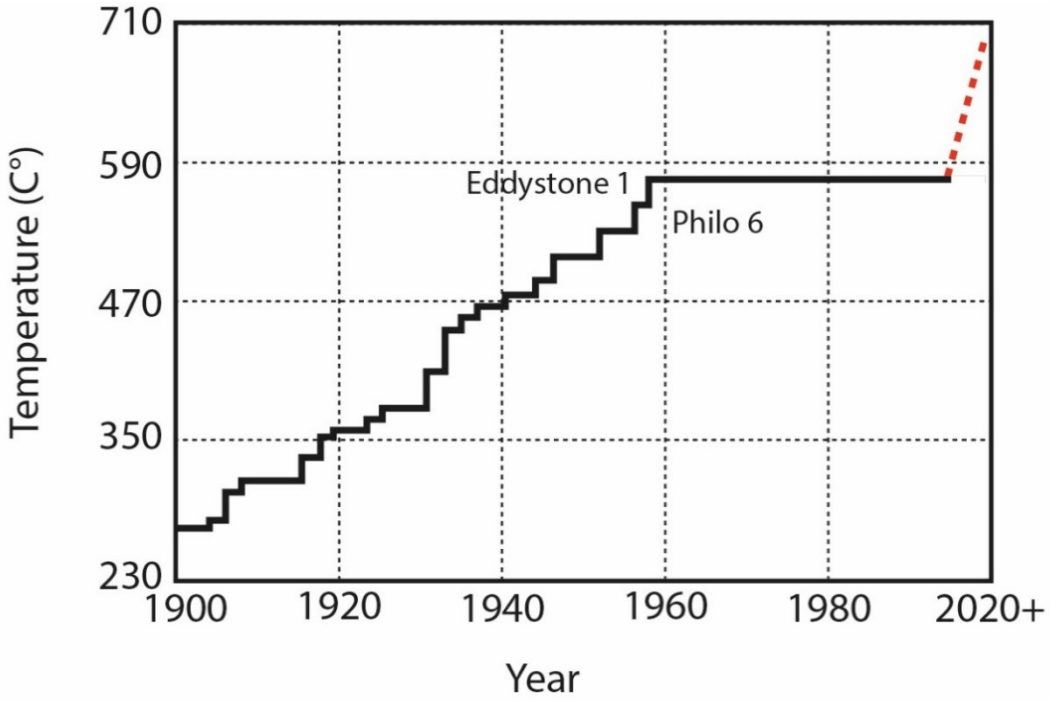

(a)

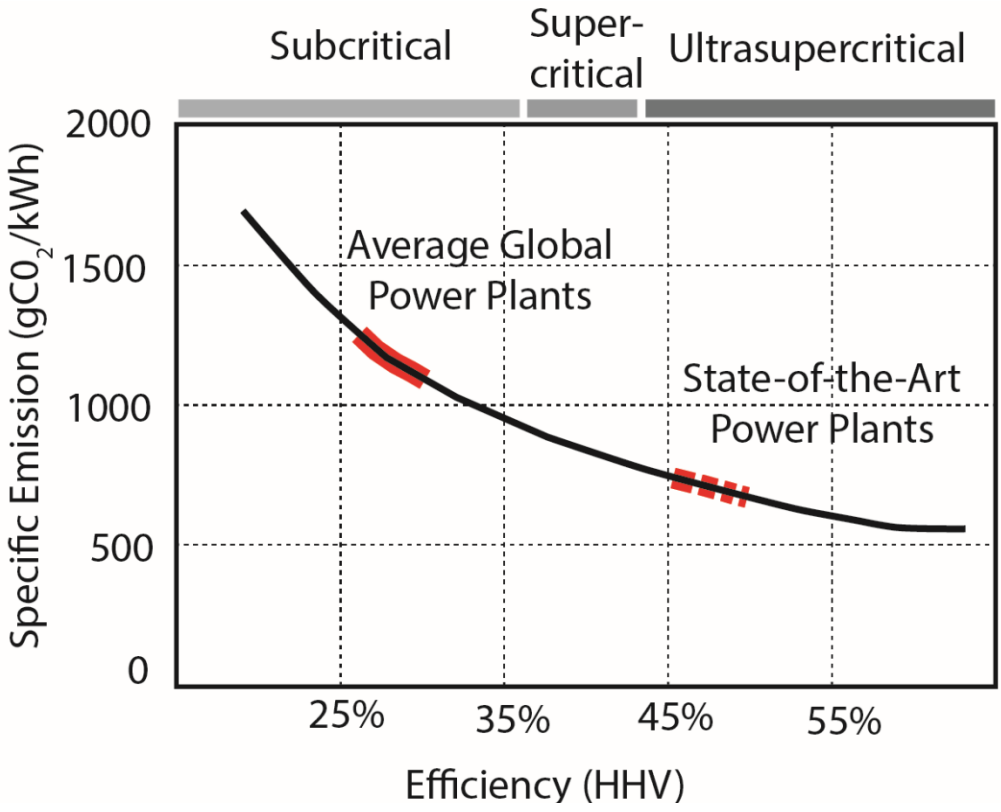

(b)

Figure 1.1: (a) Evolution of power plant operating temperature (Eddystone 1 and Philo 6 are power plants), image adapted from (EPRI 2013) (b) Concept of ultra-super-critical power plants, image adapted from (IEA 2014).

Two main damage mechanisms (Webster & Ainsworth 1994; Viswanathan 1989) threaten power plants: the long-term accumulation of creep and fatigue. The former is the static high temperature life limiting mechanism, and the latter occurs through the thermal cycling of the plant, introducing defects which, when allowed to grow beyond a critical size, can cause catastrophic failure. To safely extend the service life of components thus starts by understanding the population of flaws that might be present within. This requires the ability to detect, size and characterise any flaws accurately, and with high reliability, so that their development with time can be predicted.

Whereas surface breaking defects can be found using visual inspection, such as liquid penetrants and magnetic particle inspection, sub-surface ones require more advanced NDE technologies. Two techniques are available for the in-situ volumetric inspection of power plant components. Sub-surface defects can be found by employing either electromagnetic or mechanical waves, termed radiography and ultrasound respectively. Radiography requires costly measures to prevent any exposure to personnel within the vicinity and is poor at detecting defects whose plane is perpendicular to the direction of radiation. As wall thicknesses of future USC components are likely to increase, in order to withstand higher pressures, the need for increased penetration may require unfeasible levels of radiation. Consequently, ultrasound is the preferred method for inspection.

Ultrasound relies on the scattering of waves from any discontinuities within a material, such as fractures or voids, which are received as echoes in order to locate and identify defects. In its simplest form, this can be achieved using a single transducer which both emits and receives ultrasonic waves, known as pulse-echo operation. Other formats consist of two transducers in a pitch-catch configuration, or a combination of both, which is used in a more modern ultrasound technology: multi-element arrays. For all these configurations, typical wavelengths are of the millimetre scale, although shorter wavelengths are desirable as this improves resolution and thereby also the ability to characterise objects.

The evolution of future power plants will thus depend not only on sufficiently strong materials but also on the technical confidence, including the capability to perform reliable NDE. Inspection with ultrasound however, introduces new challenges. Aside from the appropriate alloying composition, these super-alloys would derive their superior creep strength from a coarse polycrystalline microstructure. This is where inspection difficulties arise due to scattering brought on by the material microstructure.

1.2 Problematic: Scattering

Waves propagate freely in unchanging, homogenous media. A great deal of problems in electromagnetics, acoustics, and elastodynamics have been studied in this context: antennae, guided waves, and optical fibres amongst many others (Ishimaru 1978). Nature is often more complex however, and media can vary both randomly in space and time; waves propagating in such a medium vary randomly in amplitude and phase and are therefore more easily described by probabilities rather than deterministic laws. This constitutes a form of scattering such as that found by seismic waves propagating within different layers of the Earth's sub-surface, or light scattering from particles in the atmosphere. Another form of scattering is that which arises from rough surfaces [see e.g. (Ogilvy 1991)] which is not dealt with here.

Within ultrasonic NDE, scattering is encountered when the elastic waves propagate within heterogeneous media where either density or phase velocity (collectively known as the acoustic impedance) varies - usually spatially - occurring in the aforementioned super-alloys, layered composites such as carbon reinforced plastics, and even conventional steels (e.g. cast components) can exhibit a sufficiently coarse microstructure. In fact, ultrasound is fundamentally limited by the onset of scattering once the wavelength becomes dimensionally similar to the characteristic dimension of the inhomogeneity of the propagation medium. This ultimately limits the range of materials which can be reliably inspected and restricts the maximum resolution which can be achieved.

Scattering occurs within polycrystalline materials due to a contrast in phase velocity at each grain boundary of the microstructure. As the coherent wave is scattered, it becomes exceedingly difficult to discern the signals of interest, stemming from defects such as voids and fractures, from those inherent to the microstructural features of the material. This reduced ability to distinguish signals from noise, often quantified by the Signal-to-Noise Ratio (SNR), is illustrated in Figure 1.2 by comparing typical ultrasonic measurements from (a) a transparent and (b) a scattering material. The adverse effects of scattering can be summarised by discussing (1) an increased attenuation, (2) the introduction of coherent noise, and possibly (3) anisotropic effects.

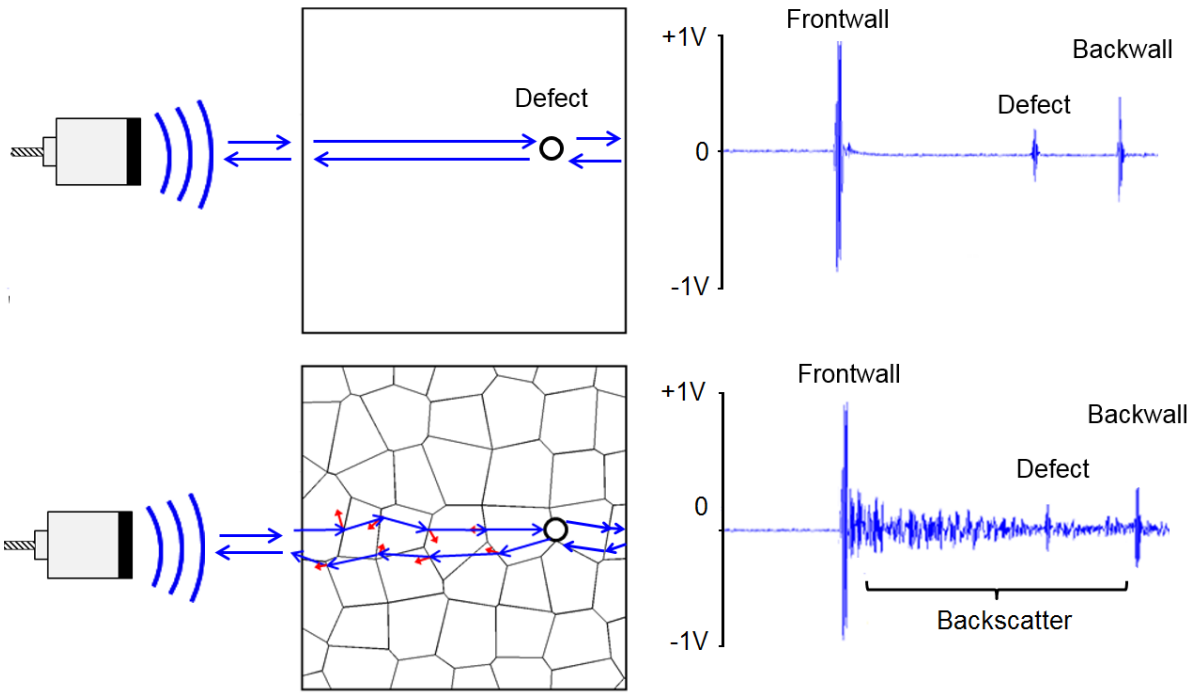

Figure 1.2 Schematic of a pulse-echo ultrasonic immersion inspection of a cylindrical defect within a (a) ultrasonically clean material and (b) a highly scattering material. Left shows the physical inspection layout and right displays the inspection information. Images adapted from (Feuilly et al. 2009).

1.2.1 Attenuation

As waves propagate their energy dissipates and spreads causing the wave amplitude to decrease; this phenomenon is termed attenuation. The mechanism by which it occurs varies according to the propagation medium, but by and large it can be classified as the summation of diffraction (also known as beam spreading), absorption, and scattering, where only the latter two are solely properties of the material.

Attenuation due to absorption effects usually follows a square dependency with frequency. In polycrystalline materials however, scattering introduces a new form of attenuation that dominates other forms, causing the overall attenuation to rise significantly. Attenuation will decrease signal amplitudes, which when maintaining the same level of noise, causes a reduction in SNR.

In many other contexts of NDE, this can often be resolved by one of two solutions. The first is to increase the signal amplitude by increasing the original excitation intensity and the second consists of averaging repeat measurements. When inspecting polycrystalline materials however, neither of these are possible solutions due to scattering induced noise.

1.2.2 Coherent Noise

Noise is a ubiquitous limitation on sensing applications which can arise from a variety of sources. Most common, in for example the ultrasonic inspection of acoustically transparent materials, are the random fluctuations in an electrical signal due to thermal agitation. Due to its temporal incoherence, repeating the same measurement and averaging it with a previous measurement, allows a reduction of the noise by a factor which is dictated by the square-root of the number of averages. Hence in the absence of time restrictions, as long as incoherence is satisfied, any SNR is attainable.

However, scattering induces temporally coherent noise in the form of backscatter and as previously hinted, averaging no longer provides benefits. Namely, repeating a measurement initiates the same scatterings events within the material as it remains unchanged, subsequently reproducing the same measurement of both noise and signal.

Equally, increasing the excitation signal amplitude in attempt to recover a loss in SNR, simply increases the noise linearly and thus maintains the same SNR.

In combination with a reduced signal amplitude due to the increased attenuation, coherent noise can cause a severe drop in SNR, and without an immediate solution, presents the fundamental problematic for the inspection of these materials.

1.2.3 Anisotropy

Besides those previously mentioned, additional difficulty can arise due to macro-anisotropy. Anisotropic materials have properties which are dependent on direction. A well-known example is composites, but metallic materials can also exhibit stiffness variations on a macro-scale (further discussed in Section 2.2). When this is the case, this can present challenges for an ultrasonic inspection as it affects wave propagation and causes deviations from the equivalent straight-ray path for an isotropic medium. The effect of anisotropy is analogous to a defocusing and skewing of the beam which reduces measured signal amplitudes rapidly.

1.3 Aims and Objectives

The work reported in this thesis explores the possibility to improve the ultrasonic inspection of highly scattering materials. This involves improving the capability to both detect and characterise - including classification and sizing - sub-surface flaws within coarse grained, polycrystalline materials. Since the ability to characterise is governed by resolution, realising this purpose rests on enabling inspections at the highest possible frequencies, despite the consequent increase in scattering, in order to maximise resolution.

The project adopts an applied-research philosophy, such that no particular technology is pre-selected at the onset. Rather, the project is problem-driven, and openly explores and adopts the available technologies to realise its aims, including the review of current standings, benchmarking of the forerunning NDE techniques, and where possible, the development of new ideas and approaches.

The scope is defined by the practical requirements for inspections that are commonplace within the power generation industry. This defines in-situ inspections which typically occur during an outage, where operators perform contact measurements with a variety of ultrasound technologies. Certain complications are not considered however, such as those which can occur due to curved geometry, welds, and complex defects, which may also represent a large portion of current inspection challenges.

1.3.1 Roadmap

The key developments of the thesis are organised as follows:

Chapter II: reviews the literature of grain scattering, including the currently most promising solutions.

Chapter III: familiarises grain scattering through three experimental investigations.

Chapter IV: outlines a benchmarking strategy for evaluating state-of-the-art de-noising algorithms, and applies it to three of the currently most promising imaging algorithms.

Chapter V: establishes Finite Element modelling for simulating wave propagation of elastodynamic scattering within polycrystalline materials in 2D and 3D.

Chapter VI: extends the aforementioned Finite Element model into a framework to investigate ultrasonic array imaging within highly scattering media.

Chapter VII: proposes a pseudo-colouring signal processing scheme for imaging highly scattering materials

Chapter VIII: explores spatial filtering to reduce the contribution of shear waves within the scattered wave field, and introduces an initial investigation of its merits to ultrasonic NDE.

Chapter II

ULTRASONIC GRAIN SCATTERING

A Review of Literature

2.1 Introduction

Our journey into ultrasonic grain scattering began more than half a century ago and remains ongoing today. The following review retraces its milestones starting with the first ultrasonic experiments with polycrystalline metals, during the 1940's when ultrasound itself was still in its infancy. These pioneering experiments led to the discovery that ultrasonic waves interact with, and hence are sensitive to, polycrystalline microstructures. This spurred the study of the scattering induced attenuation, in pursuit of material characterisation, for the next three decades - here described in Section 2.3.

The next era of research starts in 1989 when the crash of United Airlines flight 232 culminated interest into another ultrasound application, that of flaw detection. In order to avoid further such tragedies and improve existing inspections, efforts now steered towards understanding the scattering induced noise. This era of research is covered in Section 2.4, when the FAA's Engine Titanium Consortium (ETC) broke significant ground to ultimately improve the safety and reliability of aerospace engines today.

Throughout the 1980's, significant interest prevailed in signal processing techniques to supress noise sources such as clutter from radar. Frequency compounding techniques emerged and became a particularly interesting development for reducing grain noise, known today as Split Spectrum Processing, which is reviewed in Section 2.5.

More recently, the introduction of multi-element arrays to ultrasonic NDE has enabled exciting opportunities for the inspection of noisy materials. In particular the emergence of advanced coherent noise filters (e.g. DORT MSF), spatial compounding techniques, and super resolution imaging, which are outlined in Section 2.6.

Due to the abundant literature on this subject, this review is only a glimpse of approximately sixty years of research on grain scattering. For further reading beyond what is presented here, I recommend (Bhatia 1967; Papadakis 1981; Thompson 2002; Thompson et al. 2008). First however, we must define a polycrystalline material.

2.2 Polycrystalline Material

Materials exist in a crystalline state, or as a crystal, when their microstructure is atomically arranged in a regular pattern. The material is then entirely formed from a continuous repetition of a single building block: the unit cell. As the unique constituent of the material, the unit cell holds the structural information for the entire crystal, and therefore also the material anisotropy which is determined by its symmetry. A fully symmetric unit cell produces an isotropic material, and one without it produces a fully anisotropic one. The latter is known as a triclinic crystal system and requires 21 unique single crystal elastic stiffness constants (SCESC) to describe its stiffness matrix. More often however, a certain degree of symmetry exists, leading to various patterns of anisotropy, which are described by the six remaining crystal systems: monoclinic, orthorhombic, tetragonal, trigonal, hexagonal, and cubic. The latter exhibits the highest order of symmetry and can be defined by just 3 independent SCESCs.

Fully crystalline states seldom exist however, and more commonly, materials such as ceramics or metals, occur as aggregates of single crystals, termed a polycrystalline material (see Figure 2.1). A single constituent crystallite, also referred to as a grain, can vary in shape and size. The collective of these properties, known as the grain morphology, governs the macroscopic behaviour of the material. The manipulation of morphology has long enabled the creation of superior engineering materials, even dating back to early armourer's practices. A more modern example can be taken from the aerospace industry, where turbine disks (Panetta et al. 2014) are manufactured to exhibit a large grained microstructure in the creep-critical outer diameter, whilst reverting to smaller grains near the inside diameter where fatigue is the limiting failure mechanism.

Polycrystalline microstructures are governed by a number of processes, but in general are affected by the conditions during solidification of the material. Initial crystallisation occurs after the first melt during material manufacture, during processes such as casting, but it can also occur at later stages, for instance, during heat treatment where it is known as recrystallization. Most often, a prolonged cooling process allows coarser grains to form whereas quicker cooling promotes grain boundaries thereby yielding finer ones.

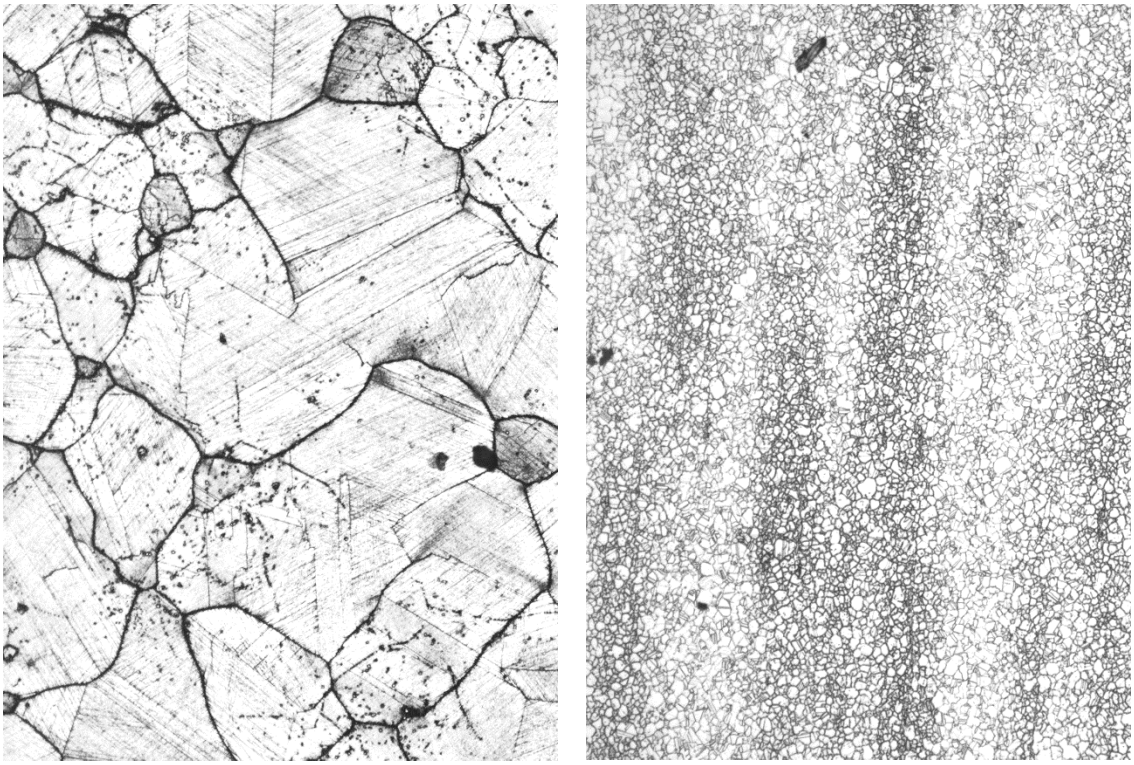

(a) (b)

Figure 2.1: Comparison of (a) coarse and (b) fine polycrystalline microstructure. Both images shown on the same scale and represent approximately 1.2mm in width.

Material	C_{11} (GPa)	C_{12} (GPa)	C_{44} (GPa)	A
Tungsten	501	198	151	1
Aluminium	108	61	29	1.2
Alpha-Iron	237	141	116	2.4
Nickel	247	147	125	2.7
Silver	124	93	46	2.9
Copper	168	121	75	3.2
Lead	50	42	15	3.7

Table 2.1: Cubic stiffness constants for several common cubic materials, sorted according to their anisotropic ratio which is related to their ultrasonic scattering strength. Source: Courtney, materials book.

Not only size, but also the grain shapes are a product of the material's manufacture. The most basic grains, and to some extent also idealised, have approximately equal dimensions in all directions, known as equiaxed, and are formed in conditions of uniform temperature gradients during crystallization. Elongated grains, ones which are columnar, arise from either non-uniform temperature distributions, like those caused during welding, or through plastic deformation by forming processes such as rolling and extrusion.

Elongated grains lead to preferential orientations, namely the crystallographic orientations within the polycrystalline material are no longer randomly distributed. The collective of these orientations can be described by Orientation Distribution Functions (ODF), which when entirely randomly distributed, denote a macroscopically isotropic material. When the material has a preferred orientation, due to for example elongated grains, the material is said to be textured as it exhibits macroscopically anisotropic behaviour. Some of the fundamental analytical tools to study textured materials were laid by (Roe 1965; Bunge & Morris 1982).

The anisotropic stiffness properties of crystallites cause the material to be heterogeneous, which forms the underlying cause for scattering. Namely, two neighbouring grains which have different orientations, create a difference in wave velocity (an acoustic impedance mismatch) at their boundary. Consequently, an incident wave upon this interface causes a scattering event to occur, the strength of which is determined by the intensity of the impedance contrast. Hence the scattering strength of a material is in part determined by the maximum possible difference in stiffness which can occur through two orthogonal orientations. This quantity is numerically represented by the anisotropic ratio A, given in Equation 2.1 for a cubic material. Table 2.1 lists this factor for some common materials and indicates their SCESCs.

$$A_{cubic} = \frac{2c_{44}}{c_{11} - c_{12}} \tag{2.1}$$

Now to illustrate the consequences of an anisotropic stiffness on elastic waves, we calculate the wave velocity versus propagation direction for a cubic single crystal material. This procedure uses the result for the dispersion relation (shown in Equations 2.2-2.4)

found in (Auld 1990) which considers a cube of dimensions X,Y,Z with propagation along the face *X Z*.

$$\left(\frac{k}{\omega}\right)_1 = \sqrt{\frac{\rho}{c_{44}}} \tag{2.2}$$

$$\left(\frac{k}{\omega}\right)_2 = \sqrt{2\rho}\left\{c_{11} + c_{44} + \sqrt{(c_{11} - c_{44})^2\cos^2(2\vartheta) + (c_{12} + c_{44})^2\sin^2(2\vartheta)}\right\}^{-\frac{1}{2}} \tag{2.3}$$

$$\left(\frac{k}{\omega}\right)_3 = \sqrt{2\rho}\left\{C_{11} + C_{44} + \sqrt{(C_{11} - C_{44})^2\cos^2(2\vartheta) + (C_{12} + C_{44})^2\sin^2(2\vartheta)}\right\}^{-\frac{1}{2}} \tag{2.4}$$

Once a solution for the phase velocity is obtained, we can insert the appropriate $\left(\frac{k}{\omega}\right)$ into Equation 2.5 and 2.6 (Auld 1990) to obtain group velocity. l_x and l_z are $\left(\frac{k_x}{k}\right)$ and $\left(\frac{k_z}{k}\right)$.

$$\left(C_g\right)_x = \frac{k}{\omega}l_x\frac{\left[c_{11}\left\{c_{11}l_z^2 + c_{44}l_x^2 - \rho\left(\frac{\omega}{k}\right)^2\right\} + c_{44}\left\{c_{11}\,l_x^2 + c_{44}l_z^2 - \rho\left(\frac{\omega}{k}\right)^2\right\} - (c_{12} + c_{44})^2 l_z^2\right]}{\rho\left[(c_{11} + c_{44}) - 2\rho\left(\frac{\omega}{k}\right)^2\right]} \tag{2.5}$$

$$\left(C_g\right)_z = \frac{k}{\omega}l_z\frac{\left[c_{44}\left\{c_{11}l_z^2 + c_{44}l_x^2 - \rho\left(\frac{\omega}{k}\right)^2\right\} + c_{11}\left\{c_{11}\,l_x^2 + c_{44}l_z^2 - \rho\left(\frac{\omega}{k}\right)^2\right\} - (c_{12} + c_{44})^2 l_x^2\right]}{\rho\left[(c_{11} + c_{44}) - 2\rho\left(\frac{\omega}{k}\right)^2\right]} \tag{2.6}$$

Figure 2.2 and Figure 2.3 graphically illustrate Equations 2.2 to 2.6 . The slowness surface and Ray surface plot the reciprocal of wave velocity and group velocity respectively. The wave velocity can be seen to vary with direction in both figures, but importantly, mostly so for the quasi-shear wave. This denotes that at a given wavelength, shear waves perceive a larger impedance mismatch than their longitudinal counterpart which results in their much increased scattering. Shear waves are thus disadvantaged for the inspection of scattering materials.

Now applying this to polycrystalline materials, where waves propagate through multiple crystals, the wave velocity consequently corresponds to an average of multiply rotated slowness surfaces. When the material is non-textured, given a sufficient propagation length, the wave speed converges to that of a macroscopically isotropic

medium. One well-known method to predict this value from the SCESCs is Voigt averaging (Voigt 1910). This enables the calculation of a Voigt velocity, shown for an orthotropic material in Equations 2.7 and 2.8.

$$E_v = \frac{(A - B + 3\Gamma)(A + 2B)}{2A + 3B + \Gamma} \qquad G_V = \frac{A - B + 3\Gamma}{5} \qquad \nu_V = \frac{A + 4B - 2\Gamma}{4A + 6B + 2\Gamma} \tag{2.7}$$

$$A = \frac{c_{11} + c_{22} + c_{33}}{3} \qquad B = \frac{c_{23} + c_{13} + c_{12}}{3} \qquad \Gamma = \frac{c_{44} + c_{55} + c_{66}}{3} \tag{2.8}$$

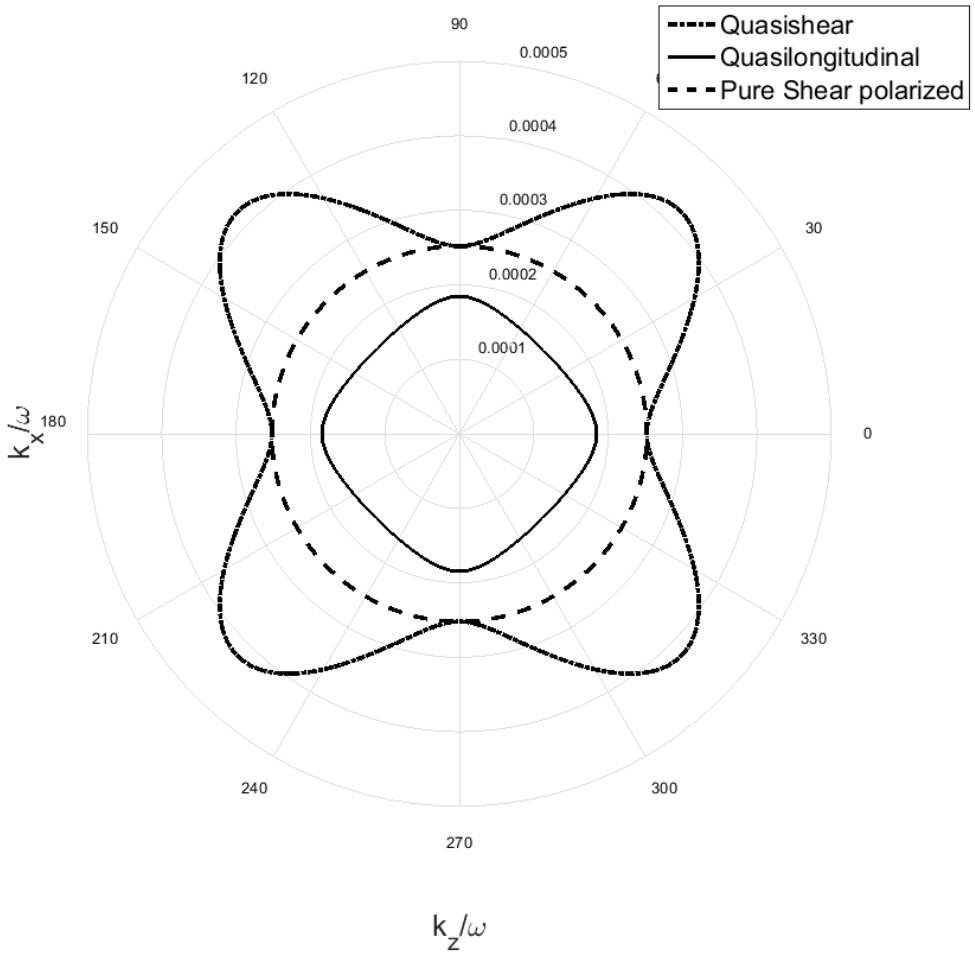

Figure 2.2: Slowness diagram for a typical cubic material.

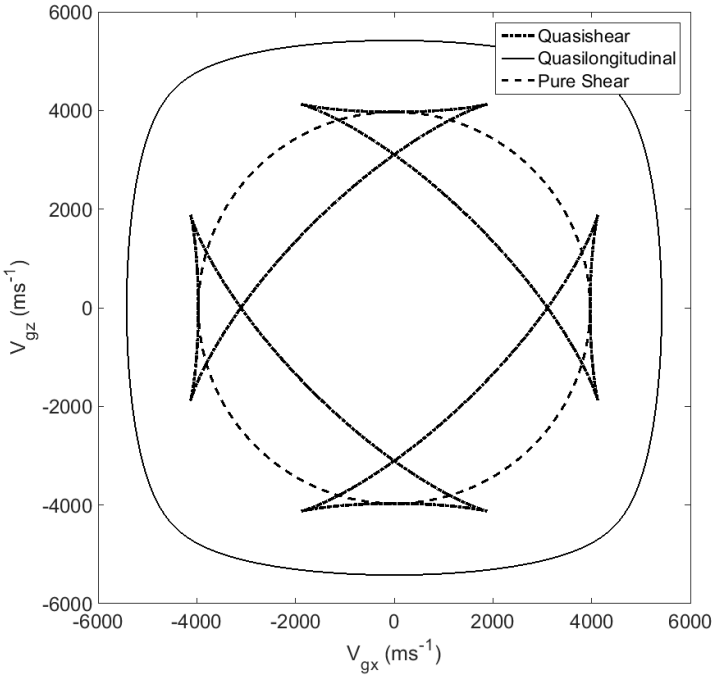

Figure 2.3: Group velocity variation for a typical cubic material.

2.3 Scattering Attenuation: Early Beginnings

In the 1940's, when pulsed ultrasound first emerged as a technique to characterise substances e.g. (Giacomini & Bertini 1939) and detect flaws buried within metals e.g. (Firestone 1946; Firestone 1942), Mason (Mason 1947) measured the scattering induced attenuation for an ultrasonic wave propagating within a polycrystalline medium. The experiment found the attenuation to increase with the fourth power of frequency, which suggested that the scattering mechanism was similar to that postulated by Lord Rayleigh (Strutt 1871a; Strutt 1871b) for visible light to explain the blue appearance of the sky. Two later experiments, which operated at higher frequencies (Roth 1948; Mason & McSkimin 1948), independently confirmed the existence of another scattering regime, a geometric one where the attenuation is independent of frequency and reduces with grain size. Lastly, (Huntington 1950) used a stochastic assumption to show that there also exists an intermediate behaviour which reduces the Rayleigh regime to a second order frequency-dependence. Hence by the 1950's, less than a decade after Firestone's patent, the three scattering regimes (Equations 2.9, 2.10, 2.11) that we know today; Rayleigh α_R, stochastic α_S, and the geometric attenuation α_G, were established.

$$\alpha_R \propto D^3 f^4 \qquad \lambda D^{-1} \gg 1 \qquad (2.9)$$

$$\alpha_S \propto D^1 f^2 \qquad \lambda D^{-1} \approx 1 \qquad (2.10)$$

$$\alpha_G \propto D^{-1} f^0 \qquad \lambda D^{-1} \ll 1 \qquad (2.11)$$

The three scattering behaviours can be intuitively explained in reverse order. Firstly, the geometric regime describes a polycrystalline material behaving as a thick-layered medium (each much larger than the wavelength). Increasing grain size here effectively reduces the number of interfaces an impeding wave has to cross, thereby reducing the number of reflections and therefore also reducing the attenuation. The stochastic regime occurs when the grain size reduces to the point that these layers dimensionally approach the wavelength. The reflections from each individual layer now begin to interfere with one another, and the attenuation starts to decrease despite adding more layers (interfaces) per unit distance. The attenuation mechanism here is dominated by the scrambling of phase

along the wavefront due to the changes in wave speed of each layer (now considering a 2D grid of layers). For example, if we think of each layer as adding a random phase, stochastic scattering denotes a regime whereby these random changes in phase do not average out even for long propagation distances. Finally within the Rayleigh regime, which occurs when the grain size is decreased further, the layers become sufficiently small compared to the wavelength (infinitesimally thin layers) that the behaviour approaches that of a homogenous medium, and any remaining scattering can be approximated by that from a random distribution of point scatterers.

Now reversing the above sequence and starting with the Rayleigh regime, as the frequency increases each grain becomes apparently larger (dictated by the wavelength) and attenuation consequentially grows rapidly until it reaches a maximum at the geometric regime. The rates at which attenuation grows, in terms of frequency, , and grain size dependence, are related across each scattering regime as shown in Equation 2.12. This result can be obtained from the rule of self-similar structures and is in part due to attenuation being defined as the loss per unit length and not per unit wavelength.

$$\alpha \propto D^{m-1} f^{m} \tag{2.12}$$

One of the practical findings from these experimental discoveries by Mason et al. was that ultrasound is sensitive to the microstructural properties i.e. grain size of a material. An attenuation measurement could therefore enable the ultrasonic characterisation of these materials. This required an explicit solution to mathematically predict the attenuation coefficient in relation to the grain size, which was independently pursued by Lifshits and Parkhomovskii (Lifshits & Parkhomovski 1950) for cubic materials which was later extended by (Merkulov 1956) to incorporate hexagonal materials (and hence hereon referred to as LMP theory), and Bhatia and Moore (Bhatia 1959b; Bhatia 1959a) for orthorhombic materials. Their equations predicted the energy lost from an incident dilation and transverse wave due to grain scattering. These models made a step improvement over their predecessors by considering the mode conversions which occur as the wave propagates across a grain boundary. These early formulations applied to single-phased, weakly anisotropic materials exhibiting no texture and spherically (equiaxed) shaped grains.

By 1965, Papadakis (Papadakis 1965; Papadakis 1968) reviewed the separate formulas for each regime of attenuation, shown here for cubic materials in the Rayleigh Equation 2.13 and stochastic regime Equation 2.15.

$$\alpha_l^R = \frac{8\pi^3\gamma^2 T f^4}{375\rho^2 C_l^3}\left\{\frac{2}{C_l^5}+\frac{2}{C_l^5}\right\} \qquad \alpha_t^R = \frac{2\pi^3\gamma^2 T f^4}{125\rho^2 C_t^3}\left\{\frac{2}{C_l^5}+\frac{2}{C_t^5}\right\} \tag{2.13}$$

$$\gamma = c_{11} - c_{12} - 2c_{44} \tag{2.14}$$

$$\alpha_l^S = \frac{16\pi^2(c_{44}+\frac{c_{12}-c_{11}}{2})^2\bar{D}f^2}{525\rho^2 C_l^6} \qquad \alpha_t^S = \frac{4\pi^2(c_{44}+\frac{c_{12}-c_{11}}{2})^2\bar{D}f^2}{210\rho^2 C_t^6} \tag{2.15}$$

Where α is the attenuation in Nepers, t and l denote longitudinal and transverse waves respectively, C is velocity, ρ is density, T is a measure of grain size with dimensions of volume, and $\bar{D}$ is the average grain diameter.

Through comparison with experiments (Papadakis 1963a; Papadakis 1963b; Papadakis 1964b; Papadakis 1965), found better agreement when refining the existing equations to also consider the distribution of grain volumes (Papadakis 1964a). This changed T in the previous Equation 2.13, the average volume of the grain, to an effective average volume T_{eff} which is slightly larger than the original T shown in Equation 2.16 where T_n represents the volume of the n-th grain for counting a total of N grains. The fact that this produced better results, was thought to explain why large grains dominate attenuation behaviour.

$$T_{eff} = \frac{\langle {T_n}^2\rangle_N}{\langle T_n\rangle_N} \text{ where } \langle T^{1,2}\rangle_N = \frac{1}{N}\sum_{n=1}^{N} T_n^{1,2} \tag{2.16}$$

Despite this early progress, available solutions shared two major assumptions. The first is that of spherical grains, where to an alternative was first presented by (Rokhlin 1972) who modelled the medium as rectangular parallelepipeds of equal size and derived the attenuation for longitudinal waves. Given its intuitive and simplistic approach, it delivered very good results (Stanke & Kino 1984) although unphysical ripples manifested due to the use of single-sized grains.

The second assumption is that of the Born approximation which reduced the solutions to single scattering. Hirsekorn (Hirsekorn 1982) was the first to amend this by incorporating multiple scattering into the derivation of attenuation and phase velocity for longitudinal waves (Hirsekorn 1982), shear waves (Hirsekorn 1983), textured materials (Hirsekorn 1985), and eventually multi-phased materials (Hirsekorn 1988). Her formalisms however, were more general than the solutions which relied on a truncated Born series and thus eventually neglected multiple scattering. Moreover, issues arose in the transient region, where the solution, similar to Rokhlin, still exhibited oscillations due to using a single-sized grains approximation.

The reliance on both these assumptions, regularly-sized grains and single scattering, ended when Stanke and Kino (Stanke & Kino 1984) adopted a second order Keller approximation (Karal & Keller 1964) and a geometric autocorrelation function to describe the grain distribution. Where previously several equations were required, Stanke and Kino's Unified Theory solved the complex propagation constant across all three scattering regimes. Within certain frequency bounds, its solutions produce the same results as those predicted by LMP theory, but it also predicts a transitional "bump" between the Rayleigh and Stochastic asymptote (see Figure 2.4).

Today, the Unified Theory remains widely accepted (+200 cit.), with good validation against experimental measurements e.g. (Stanke 1985; Yang & Rokhlin 2011). Even so, the attenuation of polycrystalline materials remains an active topic of research, with various approaches [e.g. see (Weaver 1990)] and extensions existing today. Many of these consider more complex microstructures than those discussed here including: hexagonal (Yang et al. 2011); multi-phased (Lobkis & Rokhlin 2010), and textured (Ahmed & Thompson 1996; Turner 1999) materials with elongated grains (Yang & Rokhlin 2012).

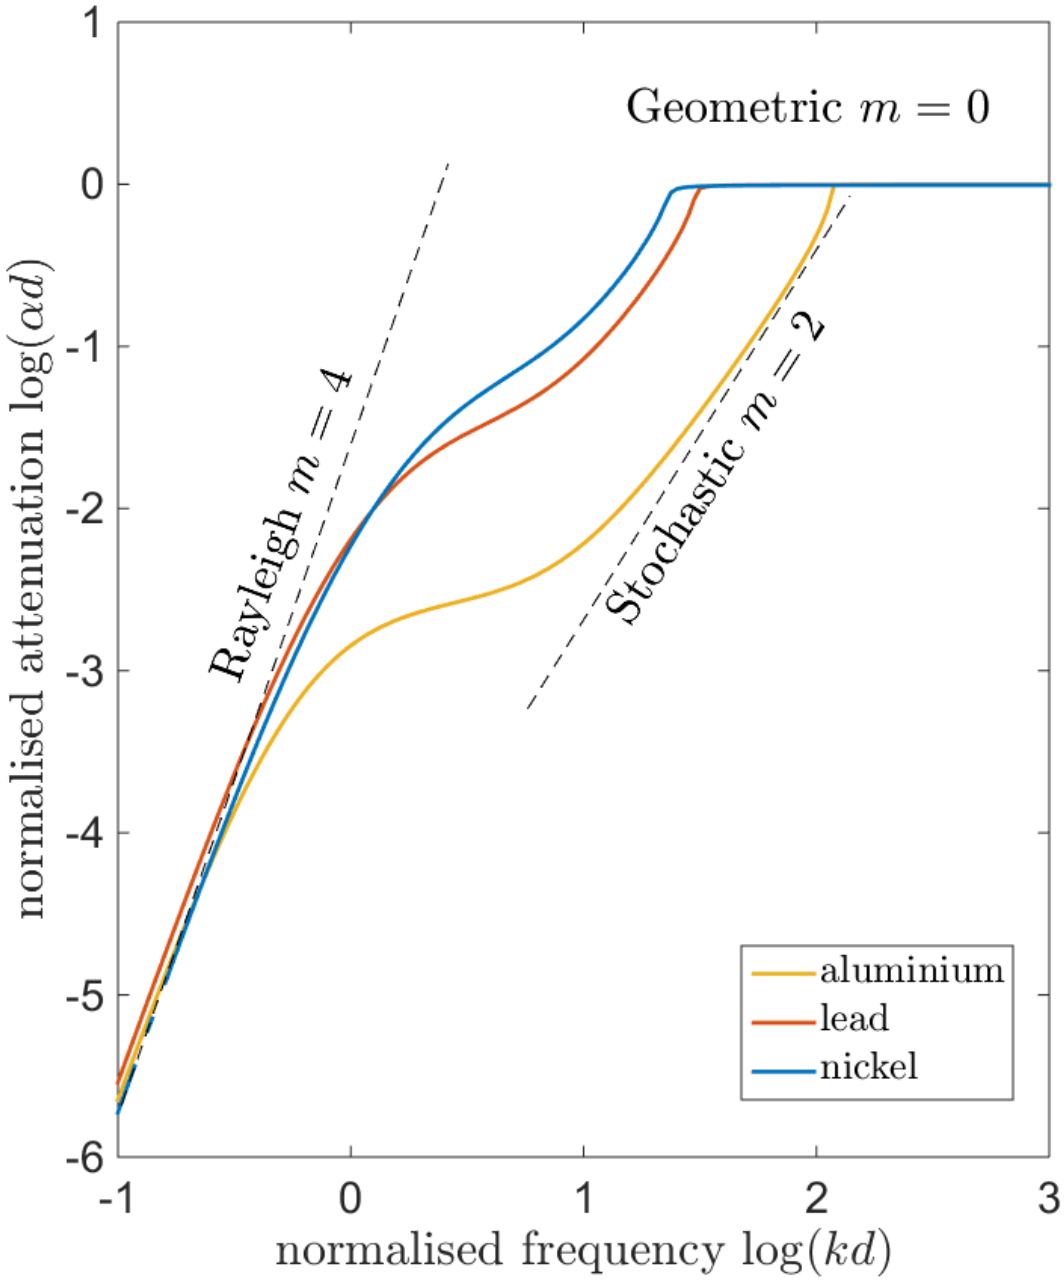

Figure 2.4: Unified theory for attenuation of three cubic materials: aluminium, lead, and nickel. The x-axis plots the normalised frequency such that results depend on the product of wavenumber and average grain size. The normalised attenuation on the y-axis is plotted such that it is independent of grain size. The three asymptotical regimes and their frequency dependence m are labelled.

2.4 Backscatter: Thompson Era

The attenuation studies described in the previous section laid an important foundation for understanding scattering and this was practically useful for material characterisation. Interest into another application of ultrasound, flaw detection, in particular the inspection of titanium billets, burgeoned in 1989 following the crash of United Airlines flight 232. Investigations into the incident concluded a hard alpha inclusion led to the catastrophic failure of the engine turbine disk. Within the same year, the Engine Titanium Consortium (ETC) was formed by the Federal Aviation Administration (FAA) which set out to improve ultrasonic inspections of scattering materials. This time, attenuation is only one piece of the puzzle, and hence these events initiated the study of ultrasonic backscatter. Across the next two decades, significant advances were made; one key contributor amongst many was Dr. R. Bruce Thompson who published several hundred articles with one of his long-time interests being ultrasonic backscatter within metallic materials (Margetan 2012).

Prior to the ETC, the majority of backscatter literature arose from the previous pursuit of material characterisation. In particular, as an alternative to relying on coherent reflections for measuring attenuation, research into backscatter initiated to instead measure the decay of the backscatter noise (Goebbels 1980; Willems & Goebbels 1981). Similarly, (Nagy et al. 1987; Nagy & Adler 1988) studied backscatter to characterise aluminium porosity and showed fundamental differences between the attenuation of backscatter and that of coherent reflections. A coefficient for backscattering was first conceived by Madsen (Madsen 1984; Hall et al. 1989; Insana 1986), albeit for biological tissue, who realised that the backscatter noise is related to the incident field and a coefficient which is a property solely of the scattering medium.

The ETC built on this foundation of backscattering literature and, similarly to the Unified Theory for attenuation, eventually developed the means to analytically predict grain noise, using what is known today as the Independent Scattering Model (ISM) (Margetan et al. 1993; Margetan et al. 1991; Margetan et al. 1994). The ISM was developed over several years, requiring several breakthroughs before its conception:

(1) One of the first challenges to study backscatter was to find a means to quantify it. The Figure-of-Merit (FOM) was introduced to this end which succeeded in experimentally measuring the inherent noise severity of a sample, independent of the inspection configuration used. Such a FOM can be directly related to the time-domain RMS noise N_{rms} observed in an A-scan. For example, for a given inspection configuration, a material having a FOM of 2 produces an A-scan with noise which is double that of a material with a FOM of unity. In the interest of brevity, a summarized Equation 2.17 is shown here.

$$N_{rms} \propto FOM\ C_1 \iiint_{-\infty}^{\infty} C^4(x,y,z)dxdydz \tag{2.17}$$

Where C_1 represents the variables associated with the inspection configuration, the volume integral of C represents the incident beam. Equation 2.18 shows the definition of FOM, where n is the number of scatterers, A_{rms} is an average of their scattering amplitude.

$$FOM = \sqrt{n}A_{rms} \tag{2.18}$$

(2) Despite having developed experimental methods to measure FOM, no means existed to predict it from known material properties. This matter was solved by (Rose 1991; Rose 1992; Rose 1993) who laid the foundation for calculating the backscattering coefficient of an untextured polycrystalline material. The frequency dependent backscatter for a cubic material, and its relation with FOM, is shown in Equation 2.19 (Thompson 2002).

$$\eta(\omega) = FOM(\omega)^2 = \left(\frac{\omega^2}{4\pi\rho C_L^4}\right)^2 \langle \delta C_{33}^2 \rangle \left[\frac{8\pi A_g^3}{\left(1+\left(2kA_g\right)^2\right)^2}\right] \tag{2.19}$$

where

$$\langle \delta c_{33}^2 \rangle = \frac{16(c_{11}^2 - 2c_{11}c_{12} + c_{12}^2 - 4c_{11}c_{44} + 4c_{12}c_{44} + 4c_{44}^2)}{525}$$

A_g is a correlation distance equal to half of the effective linear dimensions of the grain, ω is the angular frequency, and k is the wave vector. Its results have been shown to agree with experiments to an absolute factor better than 2 (Margetan et al. 1994). Rose's model was further validated experimentally

through comparison with two-phase titanium (Margetan et al. 1991; Margetan & Thompson 1992). The ISM was also extended by (Yalda et al. 1996) into a time-domain simulation using Monte Carlo simulations with 'reasonable agreement' using copper.

(3) Since steps (1) and (2) enabled the prediction of noise levels, the final keystone then was to combine this with a flaw signal model to enable the calculation of SNR. Conveniently, this had been solved earlier (Thompson & Gray 1983) by the Thompson-Gray beam model. This completed the ISM, as it now became possible to compare, in terms of SNR, the performance of different inspection configurations - thereby answering valuable questions such as optimising an inspection setup or establishing a smallest detectable defect size.

The development of ISM made several contributions to ultrasonic NDE. Firstly, it led to the well-known rule of thumb (Margetan et al. 1997), that SNR is inversely proportional to square root of ultrasonic pulse volume (see Equation 2.20).

$$SNR \propto \left[\frac{A_{flaw}(\omega)}{\sqrt{\eta(\omega)}}\right]\frac{1}{\sqrt{B^2\Delta tp}} \tag{2.20}$$

B is the average beam diameter, Δtp is the sonic pulse length or the product of velocity and pulse time length, A_{flaw} is the far field scattering amplitude of the defect. The formula also predicts the SNR to change according to the change in ratio between A_{flaw} and η with frequency. This led to the introduction of multi-zone transducers (Margetan et al. 2007), a configuration of transducers which each focus at a different depth within the billet, to improve overall SNR. The merits of focusing in this way are still used today in arrays (Wilcox 2011).

Beyond phase I (Margetan et al. 2002), ETC launched a second phase to their program which continued the work on titanium (Margetan et al. 2007) but also applied their earlier findings to nickel-based alloys (Keller et al. 2005). This nickel-phase is probably the most comprehensive study of the ultrasonic behaviour of nickel super-alloys to date. To summarise some of the major findings from Inconel 718 and Waspaloy: the velocity was

found to be relatively uniform at different locations in the samples. Attenuation and backscatter noise, which were both correlated, were found to vary significantly with position however, with up to 20-70% change in the peak noise. This was probably due to microstructural changes, which showed that both attenuation and noise increased with grain size. The transducer originally developed for titanium inspection showed improvement, and increased sensitivity from a #5 (5/64") to #2.5 and #3 (3/64") to #1 Flat Bottom Hole (FBH) in Waspaloy and In 718 respectively (Keller et al. 2005) for billets ranging from 8-14 inches.

Whist the ISM may currently represent the best noise model for grain scattering, it has certain limitations. Foremost, as its name suggests, it is built on a single scattering approximation, which breaks down at the onset of multiple scattering for stronger scattering environments. Additional sources of error are discussed in (Anxiang et al. 2003). Most literature at this time was restricted to investigating single scattering environments due to the added complexity of multiple scattering, although exceptions can be found (Russell & Neal 1997; Turner & Weaver 1995).

2.5 Signal Processing: Split Spectrum Processing

Coherent noise is a problem not unique to NDE. Similar noise limitations exist for a variety of remote sensing fields, including radar and sonar where clutter deteriorates images of complex environments, medical ultrasound where speckle reduces contrast of internal organ imaging, and seismic imaging where heterogeneity within the earth's subsurface causes backscatter. These fields have puzzled over coherent noise for over half a century, with attention that is unmatched by NDE. This has produced countless ideas and articles, one particularly interesting adaptation from this for grain noise suppression is Split Spectrum Processing (SSP), which was originally known to the radar community as frequency diversity or agility (Lind 1970).

SSP was first introduced to NDE in (Newhouse et al. 1982) to enhance ultrasonic signals belonging to flaws, above those of grains, by exploiting frequency diversity. This relies on the assumption that the received spectrum of a random distribution of point

scatterers will exhibit minima due to destructive interference. This observation is an extension from a simpler scenario, where for a regularly spaced grid of scatterers, minima of zero amplitude occur in the received spectrum at predictable frequency intervals. SSP exploits this phenomenon by decomposing a single full bandwidth time-trace into multiple sub-bandwidth time-traces, each corresponding to a different frequency window. If the underlying assumption is true, a defect will manifest a frequency coherence (be present in the several sub-signals) whereas the grain noise will be incoherent. Recombining these sub-signals, non-linearly, through for example multiplication, will amplify the defect amplitude and less so that of the noise.

One of the early challenges for SSP was to address its parameter intensiveness. Namely, dozens of parameters require fine-tuning, starting with those associated with the separation procedure. These include: the number of filters, frequency overlap, bandwidth, apodisation, and their centre-frequency. Their combined effect was investigated and answered to some extent in (Karpur et al. 1987; Karpur et al. 1988) and (Karpur & Canelones 1992).

Aside from separation, the recombination procedure presents another group of parameters which need to be determined. Here, three main categories exist: the original SSP, which utilizes minimization (Newhouse et al. 1982; Karpur et al. 1987) i.e. retrieves the minimum amplitude across the sub-signals, Polarity Thresholding (Bilgutay et al. 1989) which only accepts amplitudes when all windows have the same polarity, and the third category combines other order statistics filters such as the geometric mean (Saniie et al. 1991). Other algorithms are variations of these three categories; such as Polarity Thresholding with probability scaling (Nguyen & Jayasimha 1994) which takes the minimum amplitude multiplied by a value calculated by the number of filtered signals with the same sign; Frequency Multiplication (Karaojiuzt et al. 1998) which multiplies the windows together; Squaring and Adding Frequency components (Karaojiuzt et al. 1998) squares and adds the signals; and Phase Deviation which is similar to polarity scaling.

An example of SPP's industrial application can be found in (Baligand et al. 1986; Rose et al. 1988), but in general, its uptake has been somewhat limited. As discussed in SSP's seminal paper (Newhouse et al. 1982), this is partly due to a fundamental limit on the feasible improvement to SNR, because dividing an existing signal into sub-signals also

divides its SNR, and recombining them (linearly) would simply restore the original SNR. For instance, an evaluation of the Polarity Thresholding algorithm (Shankar et al. 1989) concluded that the algorithm shouldn't be used when SNR is close to unity, and moreover, that SNR is enhanced at the cost of the Probability of Detection (POD). It has also been suggested (Newhouse et al. 1982) that SSP's effectiveness decreases with grain size, due to multiple scattering effects which render its assumptions invalid. Another obstacle to SSPs widespread adoption, is its misrepresentation of amplitude information, and in association, the majority of industrially established defect sizing conventions. To summarize and illustrate the limitations of SSP, simply squaring a time-trace will increase SNR as long as the defect is detected (SNR is above unity), but actually increasing detection is the real challenge.

Apart from SSP, many other signal processing methods have been devised to improve inspection of scattering materials which were not mentioned here. These include Wiener filters (Izquierdo et al. 2002), Wavelets (Rodriguez et al. 2004; Pardo et al. 2006), and adaptive filtering (Zhu & Weight 1994; Kim et al. 2001) amongst many others.

2.6 Ultrasonic Arrays

In recent years, ultrasonic arrays have been widely implemented across industry, favoured for their speed of inspection and ease of interpretation over conventional NDE (Wilcox 2013; Drinkwater & Wilcox 2006). The ability to beam-steer by phase-delaying elements means that one array can reproduce an inspection which would otherwise require multiple monolithic transducers. Once ensuing advances in instrumentation facilitated the acquisition of all send-receive time trace combinations available to the array (termed Full Matrix Capture (FMC)) (Holmes et al. 2005), the additional information captured in this way enabled arrays to further surpass their monolithic counterparts by creating synthetically focused images (Wilcox 2011). Although not particular to arrays, but to the images they provide, it is also worthwhile noting that objects can be detected within a 2D image, at a much higher noise level than within the equivalent 1D time-trace (e.g. A-scan).

Since the onset of arrays, numerous algorithms are available to post-process FMC data into an image, including wave-field extrapolation (IWEX) (Portzgen et al. 2007), the wavenumber algorithm (Hunter et al. 2008), and what is regarded in the NDE community as the most established algorithm: the Total Focusing Method (TFM) (Wilcox 2013; Holmes et al. 2005). Due to the benefits of focusing as outlined in Section 2.4, arrays currently present the best chance to successfully inspect a coarse grained material.

Furthermore, arrays continue to hold the most promising opportunities to progress the inspection of scattering materials. The additional spatial information captured through FMC techniques, opens up new possibilities and has enabled the emergence of advanced imaging algorithms to further supress coherent noise sources. Three approaches which hold promise are categorised and reviewed here as follows: (1) de-noising algorithms which aim to reduce image speckle through spatial compounding. Examples of this category include; Phase Coherent Imaging (PCI) (Camacho et al. 2009; Camacho & Fritsch 2011), Spatial Coherent Imaging (SCI) (Trahey et al. 1986), and Spatially Averaged Sub-Array Correlation Imaging (Lardner et al. 2014) (SASACI). (2) a recent Multiple Scattering Filter (MSF) (Shahjahan, Aubry, et al. 2014) used in conjunction with the Decomposition for the Time Reversal Operator (DORT) (Prada et al. 1996) imaging algorithm, and finally (3) super resolution algorithms.

Other promising approaches which have not been mentioned here are those which aim to recover performance losses due to anisotropy e.g. auto-focusing (Zhang et al. 2012) and deviation correction (Connolly et al. 2009), which are mainly applicable to welds.

2.6.1 Spatial Compounding

Spatial compounding techniques share similar underlying principles to frequency compounding; both involve, instead of a linear summation, the non-linear combination of signals. Spatial compounding techniques aim to amplify spatially coherent features, which relies on the notion that flaws will exhibit a stronger coherence than grain noise. Although its potential was realised early on (Shankar 1986), its interest in both NDE and medical fields has grown significantly with the arrival of arrays. The best known examples of such imaging algorithms are Spatial Coherent Imaging (SCI) (Trahey et al. 1986; Gerig &

Varghese 2004) and Phase Coherent Imaging (PCI) (Camacho et al. 2009; Camacho et al. 2010). SCI compounds images of sub-apertures of the array, and PCI will multiply a regular delay-and sum image (i.e. TFM), with a weighting matrix which measures the coherence of instantaneous phase.

Another more recent example, based on Dual Apodisation Cross-correlation technique developed for medical imaging (Seo & Yen 2008) is the SASACI (Lardner et al. 2014) algorithm. It operates by generating two TFM images, one for each odd receiver element of the array, and another for each even. These two images are then inserted into a 2D cross correlation function which produces a correlation matrix. This matrix is then multiplied by a conventional TFM image, where any pixel amplitudes which were incoherent between the two sub-aperture TFM images, are reduced in the final image.

Similarly to frequency compounding, there has been ambiguity regarding the improvements offered by these new imaging algorithms. For example, (Jie et al. 2013) showed the improvements of SCI and PCI to be marginal. Results show that although the SNR can be increased in certain cases, PCI did not improve SNR values which were lower than unity and therefore does not affect actual detection performance.

2.6.2 Multiple Scattering Filter

DORT MSF is a very recent imaging algorithm which was originally researched in the context of the ultrasonic inspection of coarse grained nickel components found in pressurised water reactor components. It is based on the time reversal algorithm, DORT (Prada et al. 1996) which in its basic form, relies on calculating the singular value decomposition of the time-windowed T frequency f domain FMC matrix $K(T,f)$ to achieve imaging of the inspected medium. In weakly scattering conditions, each i-th scatterer in the medium is related to a singular space of the i-th singular value λ_i. The corresponding singular vector V_i is the signal of delays to apply to the array to focus onto the scatterer, termed backpropagation.

When noise dominates, such as that encountered in polycrystalline materials, significant singular values are no longer associated with scatterers of interest. Part of this noise, namely the multiple scattering portion, can be identified and reduced by exploiting the

deterministic phase relation along the anti-diagonals of the K matrix which should always occur, even for the response from a random distribution of scatterers, for single scattering. This procedure is termed a Multiple Scattering Filter (MSF) (Aubry & Derode 2009a; Aubry & Derode 2009b) and has shown great promise for detection in coarse grained materials (Shahjahan, Aubry, et al. 2014). DORT MSF involves applying the regular DORT algorithm to the filtered K matrix, K_f. The DORT MSF algorithm includes a detection criterion which aims to identify singular values which are associated with scatterers of interest and therefore merit back propagation. This is achieved by considering the statistical properties of the singular value distribution acquired from backscatter in a defect- free area for different frequencies f and depths related to time T. This method requires numerous FMCs of defect-free material which is known to be of similar microstructure. In applications of known, well-controlled, microstructure, this has been shown to generate significant improvement (Shahjahan, Aubry, et al. 2014).

2.6.3 Super Resolution

The spatial resolution of an imaging system is ultimately limited by the diffraction limit which conventionally relates the maximum resolution to half the wavelength. Super resolution imaging is devised to surpass this limit (see Figure 2.5). Whilst super resolution was first achieved some time ago, it is a relatively new arrival to ultrasonic NDE where it has since been experimentally shown that using an array it is possible to resolve points separated by $\lambda/3$ (Fleming et al. 2006).

Such techniques are particularly attractive as higher frequencies - usually required to improve resolution, also experience increased attenuation and noise, which deteriorates the quality of information (SNR) and restricts the highest practically feasible frequency. Super resolution imaging circumvents this problem, enabling inspections at a lower frequency whilst still achieving a high resolution.

This is challenging to achieve in practise, but in short, is only possible by considering second order multiple scattering effects. Traditional imaging techniques rely on the Born approximation which neglect these effects and hence tie their resolution to the diffraction limit. To explain how multiple scattering holds super resolution information, we will first discuss a simple scattering scenario.

Let us consider the interaction which would occur between an incident wave field and an object. During this interaction, the incident energy is scattered into propagating waves which can be detected at multiple wavelengths (far-field), but also non-propagating or evanescent waves (Simonetti 2006) which decay within the near-field (Fleming 2008). Whereas propagating waves are of the same wavelength as the incident field, evanescent waves are the result of the interaction with sub-wavelength spatial features of the object, and hence contain sub-wavelength information. By operating inside the near field, early pioneers broke the super resolution limit (Fleming 2008), for example with Near-field Scanning Optical Microscopes (NSOM).

For many applications including NDE, this is not a viable option which eliminates the possibility of directly extracting sub-wavelength information carried by evanescent waves. Similar to propagating waves however, evanescent waves convert to both wave-modes when scattered. Therefore an evanescent wave which is mode-converted into a propagating wave can propagate multiple wavelengths and affect far-field energy levels. In this form, this still necessitates an incident evanescent wave, and hence would require at least one of the probes, either the transmitter or receiver, to be within the near-field.

Even this limitation can be overcome by an effect known as tunnelling (Fleming 2008). If we introduce a second scattering object within the near-field of the original object, it allows the excitation of a propagating wave which generates evanescent waves during its first scattering event, which are then be converted into a propagating wave when scattered off the second object. This is a multiple scattering phenomenon and holds the key to the detection of sub-wavelength information within the far-field (Simonetti 2006).

Many approaches now exist to achieve sub-wavelength imaging (Simonetti et al. 2008); Factorisation Method (FM), Time Reversal Multiple Signal Classification (TR MUSIC) (Asgedom et al. 2011; Marengo 2007; Lehman & Devaney 2003), and Maximum Likelihood (ML) (Simonetti et al. 2008) to list some.

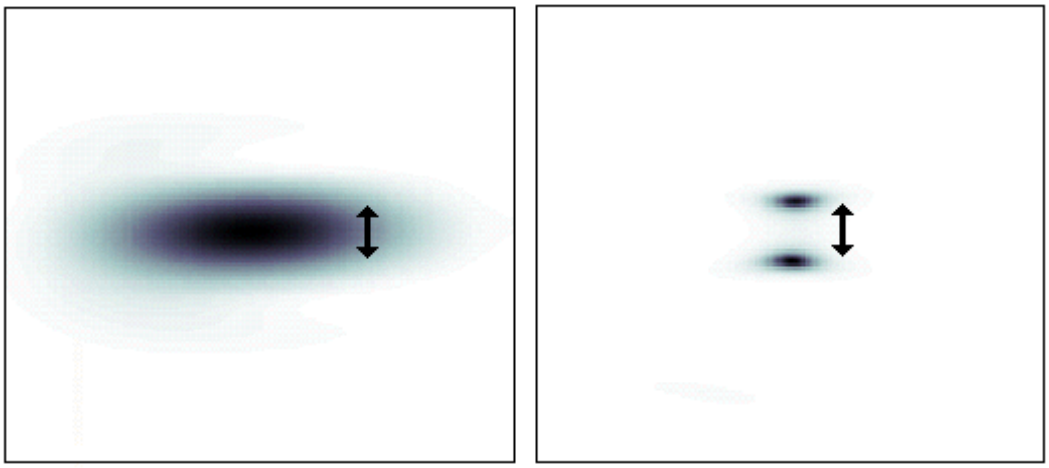

Figure 2.5: Comparing super resolution imaging to conventional imaging of two point scatterers separated by 1/3λ. Source: (Fleming et al. 2006)

Chapter III

EXPERIMENTAL DEMONSTRATIONS

This chapter pragmatically explores grain scattering through three experiments: the first replicates a conventional industrial inspection by immersion C-scanning a representative USC component; the second aims to identify scattering by analysing attenuation; and lastly, the third experiment illustrates aspects of noise by averaging both temporal and spatially incoherent noise.

3.1 Inspection of Future Power Plant Components

USC power plants largely remain a concept at this time. Many aspects including the eventual NDE challenges they will present, remain to some extent unknown. It was possible here to address this, and try to establish some of the current shortcomings, thanks to Siemens AG which loaned E.ON a sample specifically commissioned to meet this purpose.

The sample is a step wedge of Inconel 625 (see Figure 3.1) which is cut from a larger cast thick-walled cylinder. Surface etchings indicate the grain size distribution to vary from sub-millimetre to tens of millimetres in a seemingly arbitrary spatial distribution. In terms of ultrasonic targets, besides the various backwall steps, the sample contains three 5mm diameter Side-Drilled-Holes (SDH) found at various depths.

Whereas industrial power plant inspections involve contact ultrasound, it is desirable to eliminate sources of variability associated with contact measurements, and hence the data are collected in an immersion pulse-echo configuration (facility of Intertek, Derby). To investigate scattering changes with frequency, data was acquired at various frequencies by employing a combination of industrially standard, planar and focused transducers with centre-frequencies in-between 0.5MHz and 5MHz. Due to the relatively thick component, the transducer's focal zone was always in front of the backwall which is not the ideal inspection scenario. The measurement electronics are configured, where possible, such that the backwall amplitude occupies 80% of the ADC range. The data are captured by predefining gates, as is commonplace for C-scans, where the peak amplitude and its time-of-flight (TOF) within the gate is recorded. Both these metrics are stored as the probe is rastered along and across a face of the entire component, in this case with a lateral resolution of approximately 1mm. A steel cylinder of similar thickness was inserted into some of the scans, for illustrative purposes. In most figures the steel cylinder's backwall is clipped, but in some instances a variation of backwall amplitude can be seen due to the presence of a flat-bottom-hole.

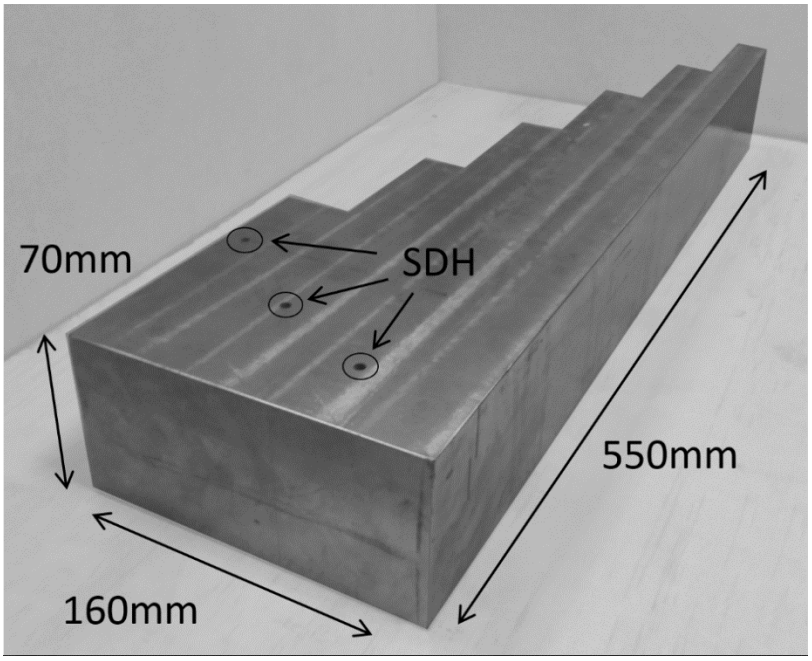

(a)

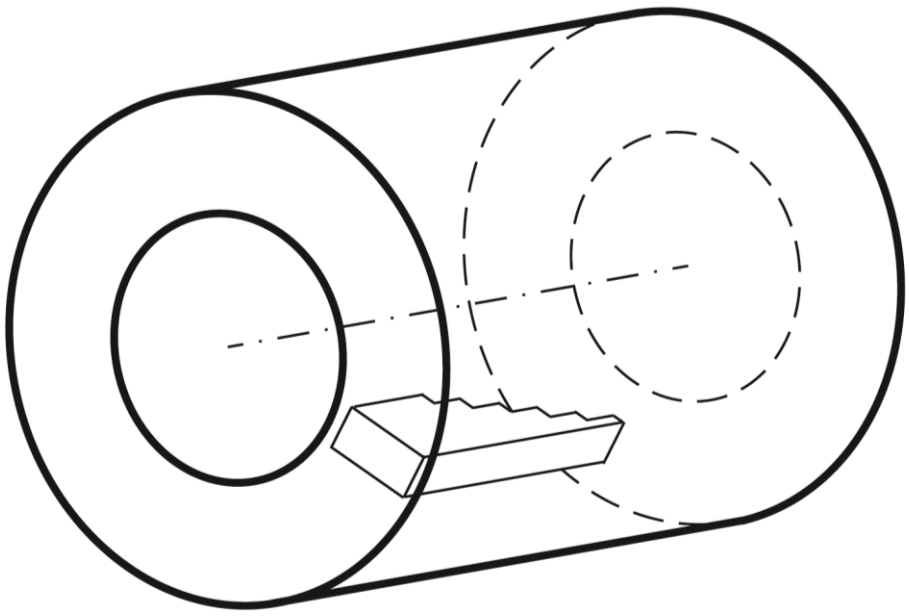

(b)

Figure 3.1: (a) Step wedge sample indicating three side-drilled-holes and a (b) schematic from its original casting.

A selection of the results are shown in Figure 3.2 to Figure 3.5 which display either the amplitude or TOF C-scan of the reflected signal from the backwall. The key findings can be summarised as follows:

- Figure 3.2 and Figure 3.3 indicate the backwall amplitude (at a depth of 70mm) to vary significantly across the side face. The variation not only manifested as a fluctuation of amplitude - which may be expected from a scattering material - it also produced distinct patterns (regions of different behaviour) within the sample. This latter observation indicates a variation of attenuation with position, and is believed to be a consequence of grain size variations shaped by the manufacturing conditions. This has previously been observed for other nickel super-alloys by Thompson et al. (Thompson et al. 2008).
- Figure 3.2 and Figure 3.3 generally indicate that higher frequencies lead to large increases in attenuation which significantly reduces backwall amplitude. This resulted in the backwall only being reliably detected at a frequency of 0.5MHz.
- Figure 3.4 and Figure 3.5 show that alongside attenuation, surprisingly large variations in wave speed (5700-6000ms^{-1}) were found throughout the block. This behaviour is believed to be evidence of preferential orientations (texture) within the material.

Given that the sample is representative of future USC components, it is likely that at least some of these effects would be reproduced in future power plants inspections. The results, albeit limited from this initial investigation, suggest that depths of up to 50mm could only be reliably inspected at 0.5MHz, when using standard single transducer techniques. This is a relatively low inspection frequency, which at a wavelength of 12mm, would be insensitive to smaller defects and provide poor characterisation ability. In summary, the sample showed significant inspection difficulties, mainly attributed to its unpredictable variation in properties.

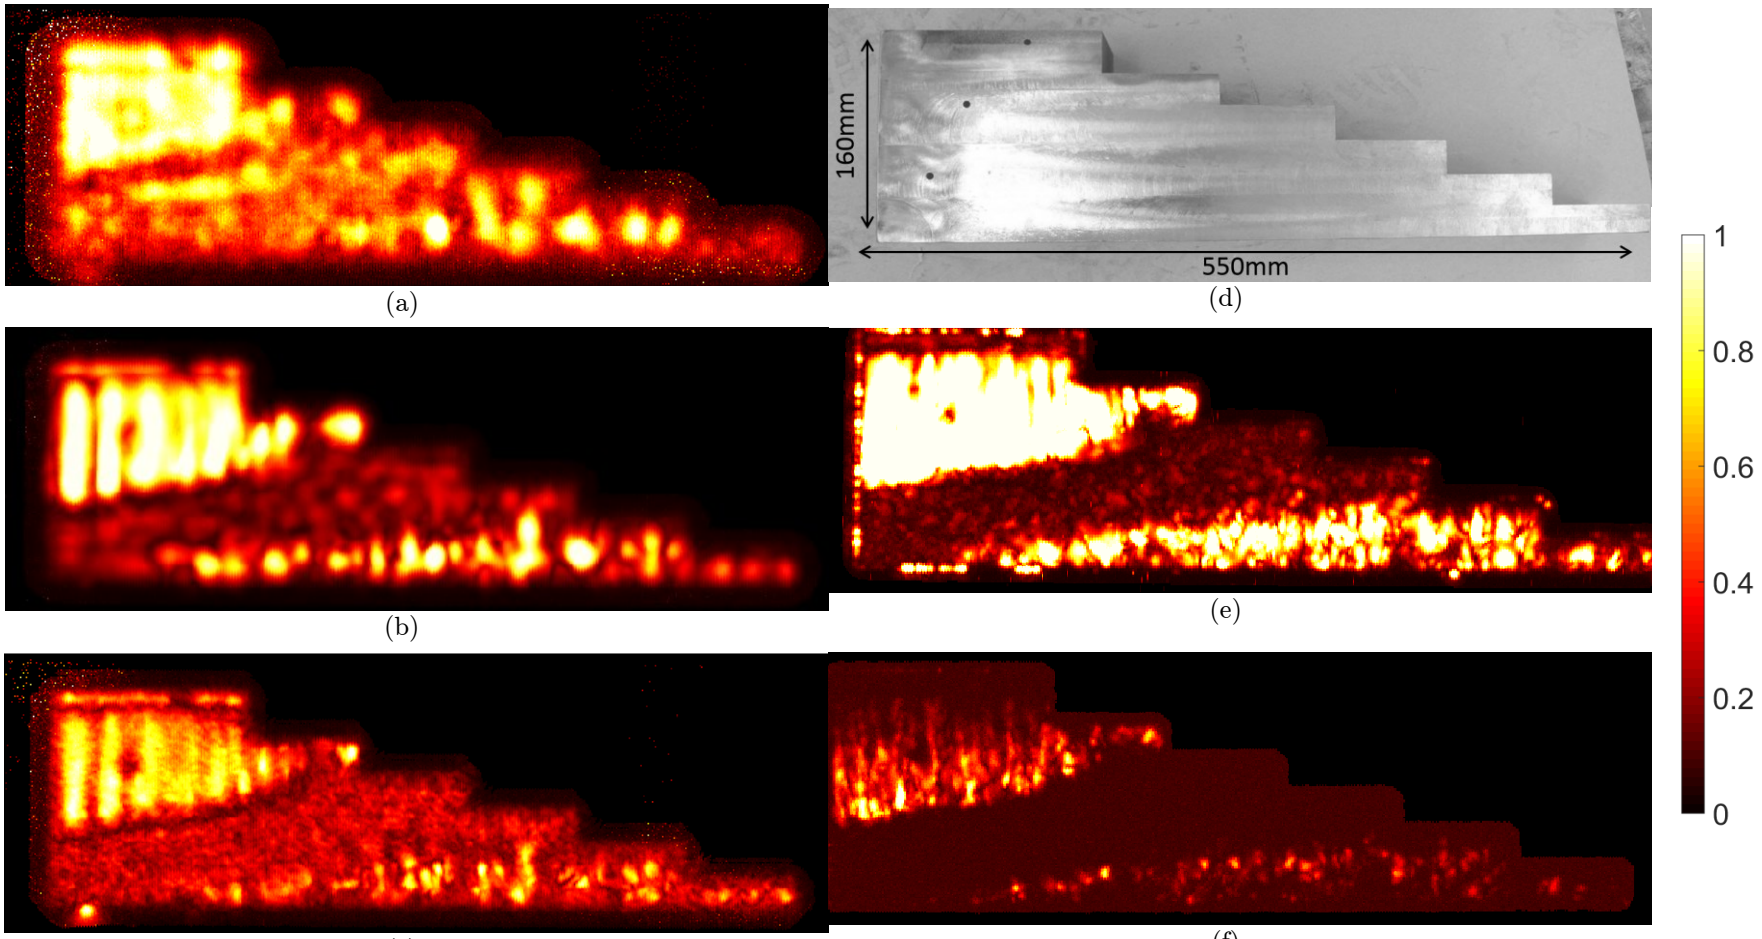

(a) (d) (b) (e) (c) (f)

Figure 3.2: Side A amplitude backwall C-scan, (a) 0.5MHz planar (b) 1MHz planar (c) 1MHz focused (d) dimensions (e) 2.25MHz focus (f) 5MHz planar.

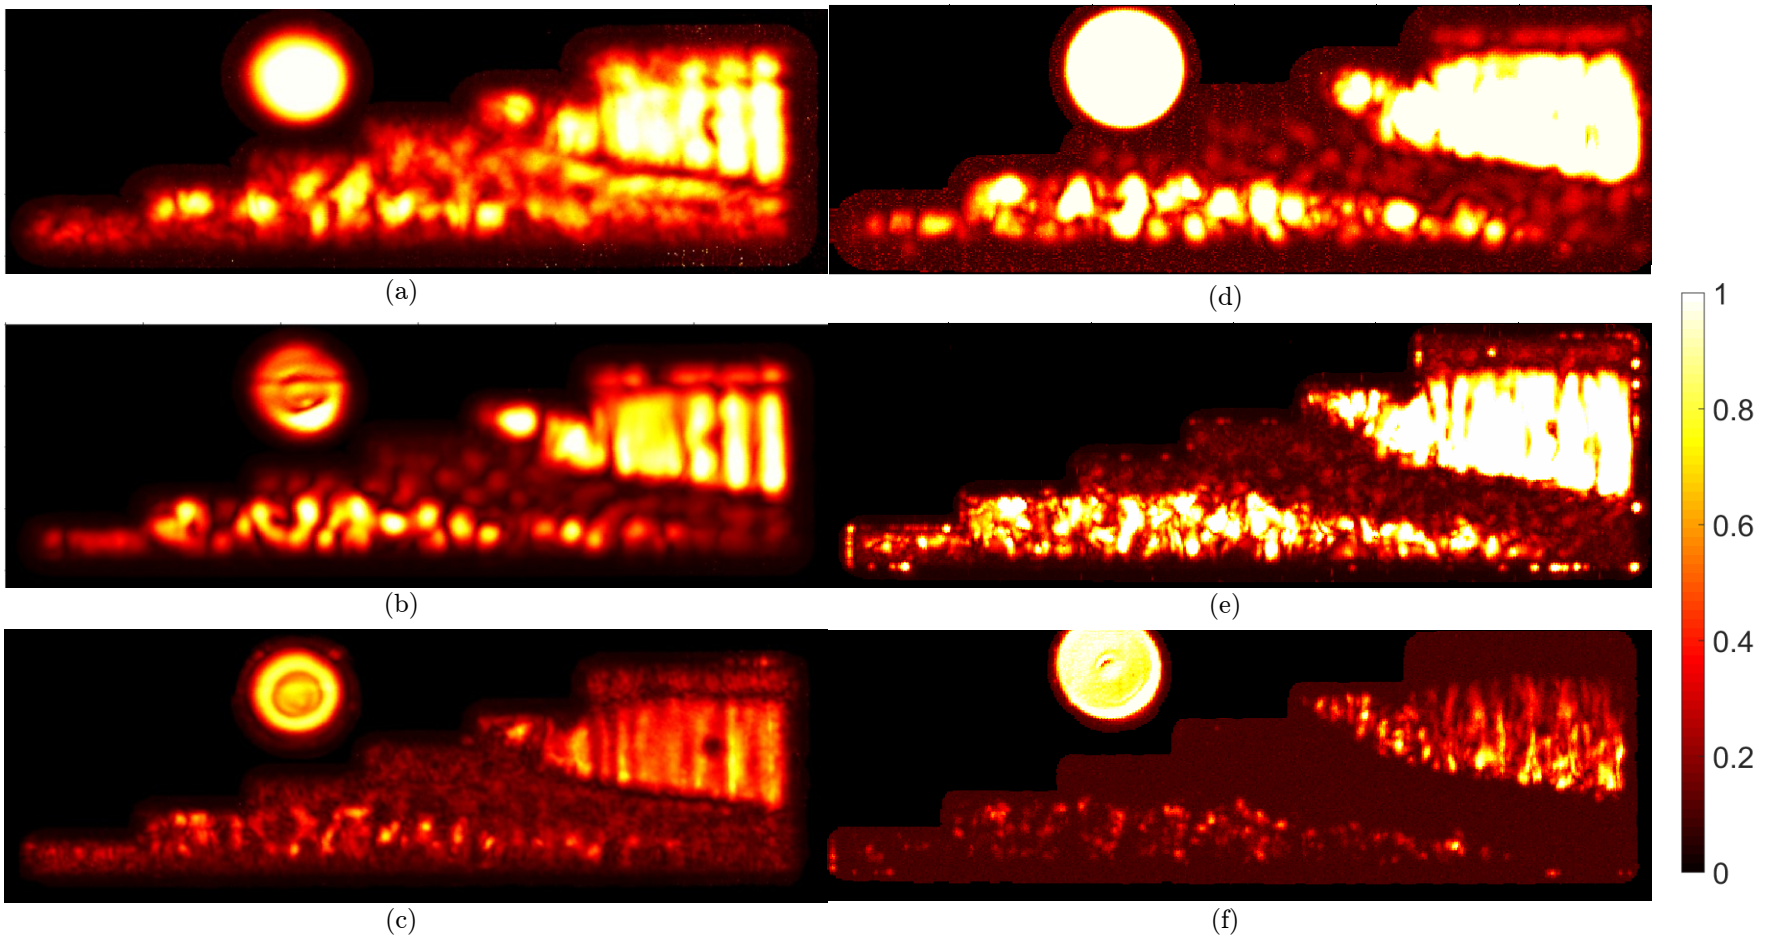

Figure 3.3: Side B amplitude backwall C-scan. (a) 0.5MHz planar (b) 1MHz planar (c) 1MHz focused (d) 2.25MHz planar (e) 2.25MHz focus (f) 5MHz planar. The circular feature present in the majority of the figures is a fine grain steel cylinder which is included for illustrative comparisons.

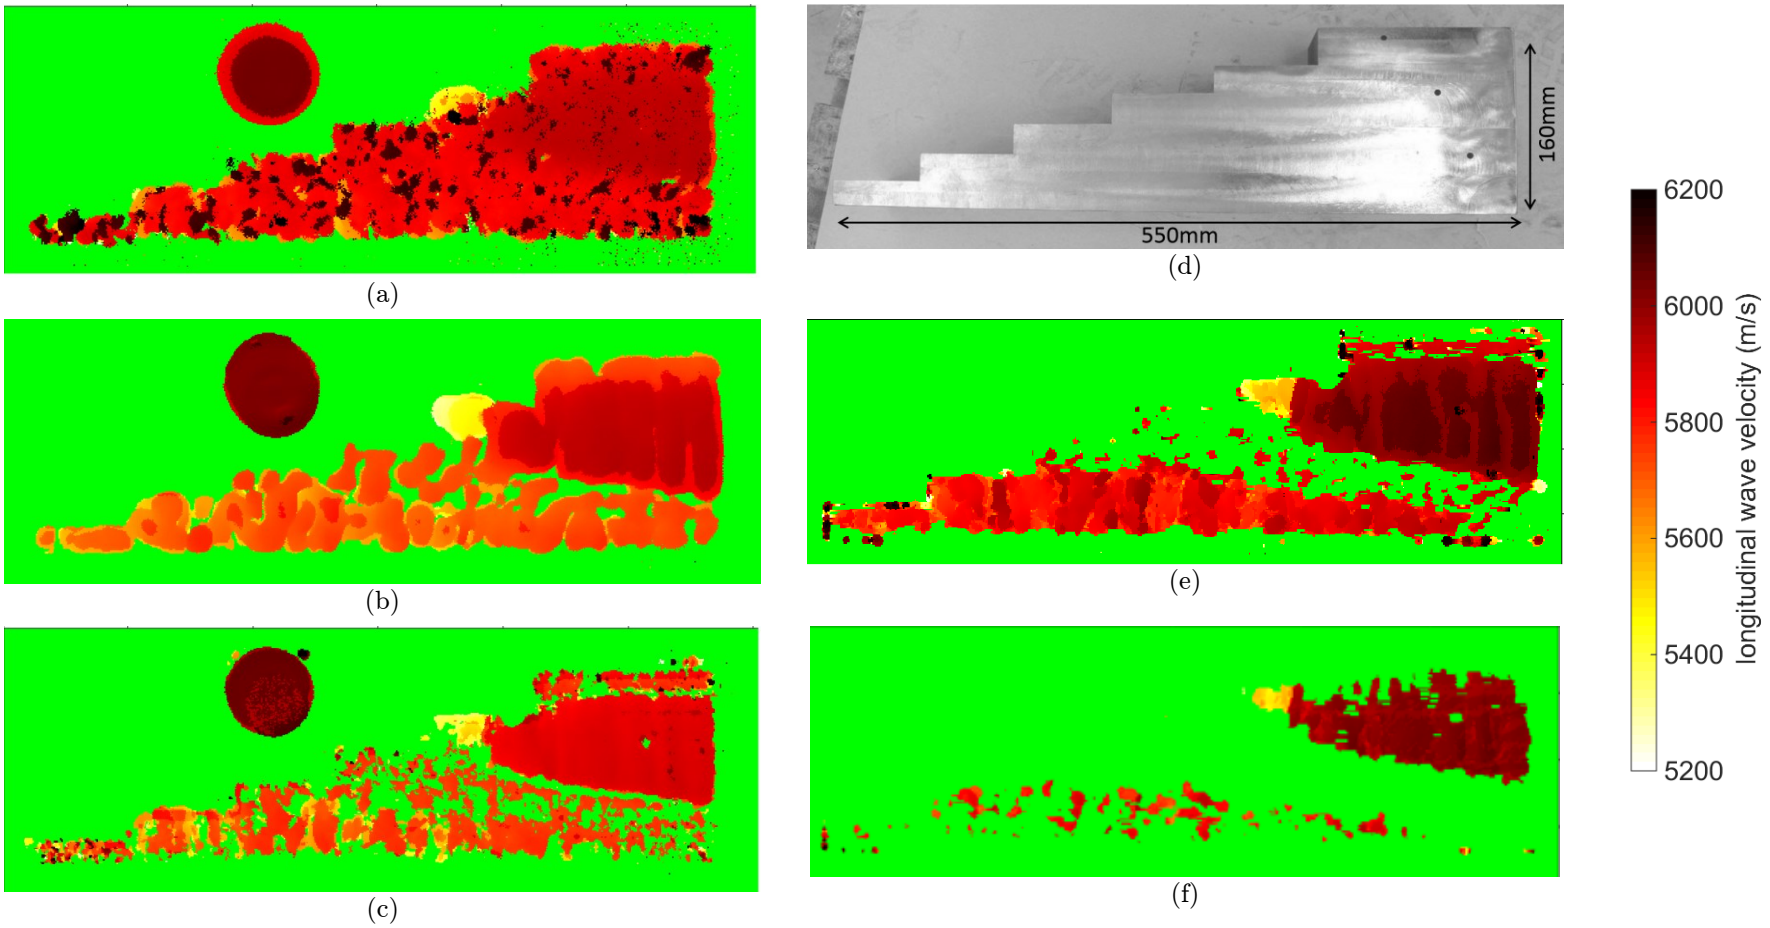

Figure 3.4: C-scan TOF side B, (a) 0.5MHz planar (b) 1MHz planar (c) 1MHz focused (d) 2.25MHz planar (e) 2.25MHz focus (f) 5MHz planar. The circular feature present in the majority of the figures is a fine grain steel cylinder which is included for illustrative comparisons.

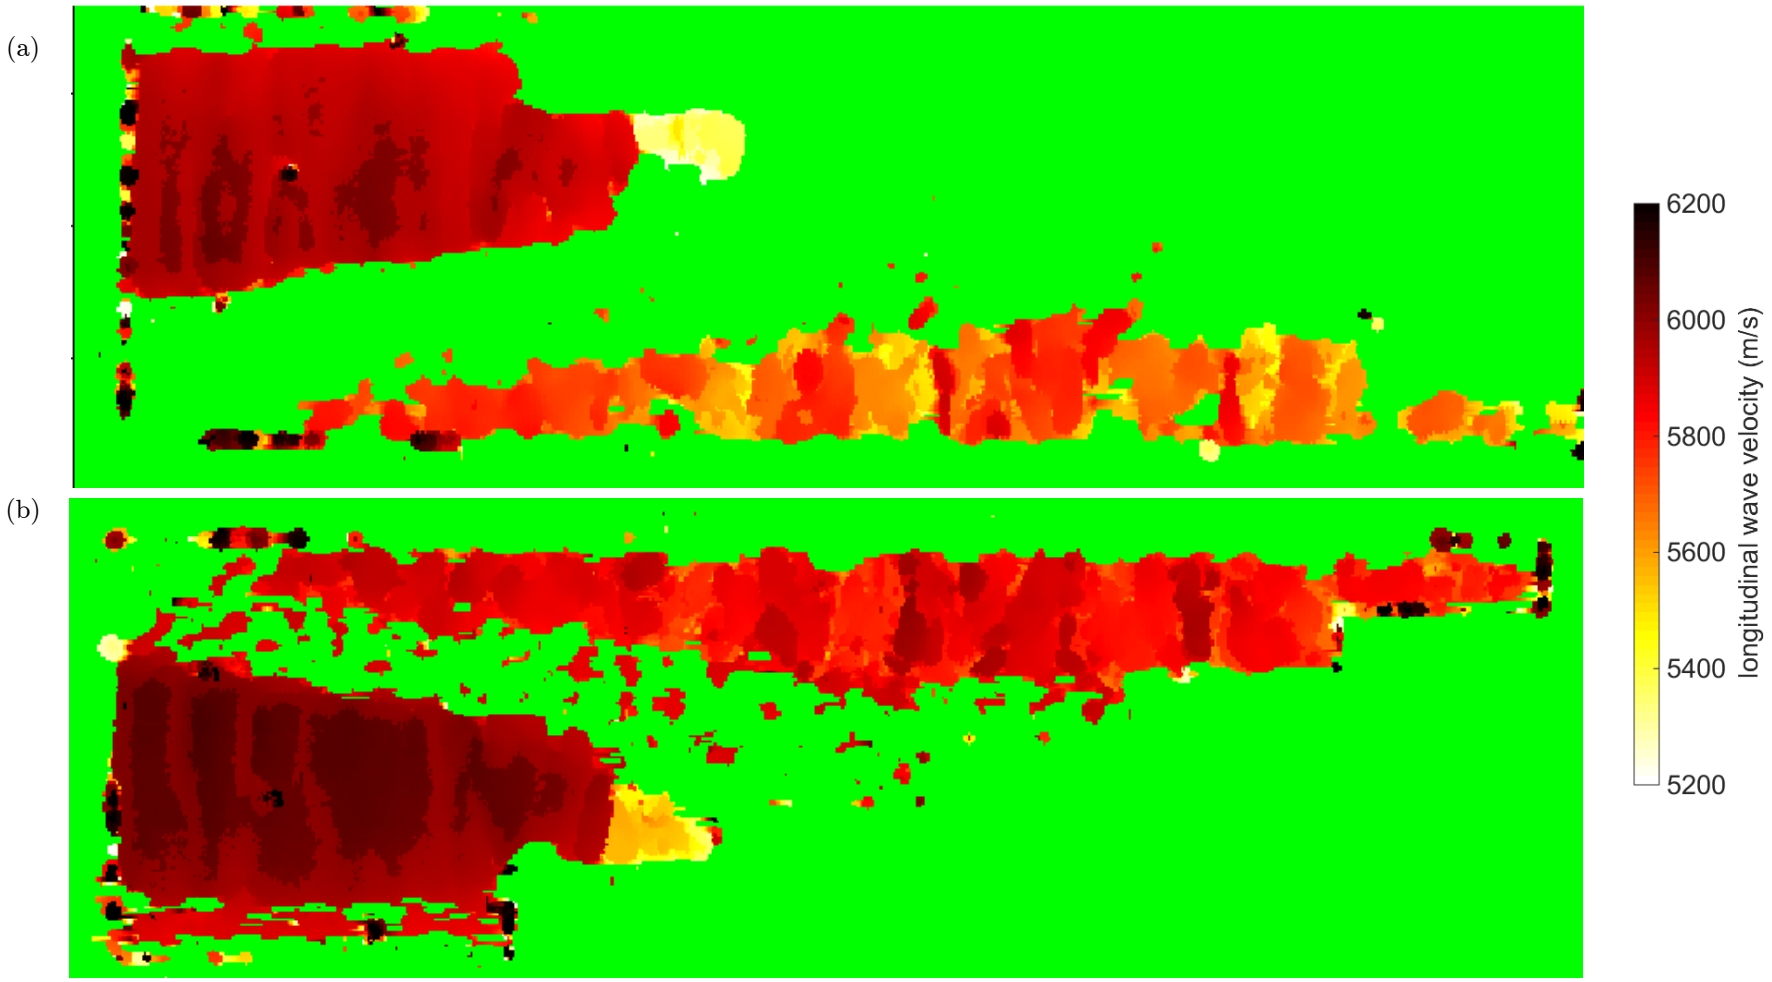

Figure 3.5 TOF variation in opposite directions at 2.25MHz (a) side A (shown in Figure 3.4(e)) and (b) side B.

3.2 Attenuation Measurements

We aim to observe scattering by experimentally quantifying the frequency dependent behaviour of attenuation within a coarse grained metallic sample. Before discussing the resulting attenuation measurements, the experimental theory and method, which consists of a pulse-echo immersion setup, follows.

3.2.1 Background Theory

Attenuation calculation

Equation 3.1 defines a wave of amplitude A_z propagating at a distance z from a finite source with an incident amplitude A_0. The forms of attenuation mentioned in Section 1.2.1, namely, scattering, and absorption are indicated by a subscript.

$$A_z = D_z A_0 e^{-(\alpha_s+\alpha_a)z} \tag{3.1}$$

The attenuation due to diffraction can be written as a coefficient D_z, and the absorption and scattering attenuation are henceforth jointly denoted by α. In a pulse-echo setup, the first and second coherent reverberations within a finite elastic medium bounded by parallel faces are introduced by Equation 3.2 and 3.3 as F and B respectively.

$$F = D_0 A_0 e^{-\alpha z_0} R_{w\leftarrow m} \tag{3.2}$$

$$B = D_1 A_0 e^{-\alpha z_1} T_{w\rightarrow m} R_{m\leftarrow w} T_{m\rightarrow w} \tag{3.3}$$

These will be referred to as the frontwall F and backwall signal B, each with their respective diffraction coefficient depending on the propagation distance, and where R and T denote the reflection and transmission coefficient where subscript m denotes metal and w denotes water and the arrow indicates the direction. The coefficients are calculated using a plane wave assumption and are independent of frequency. By dividing F and B, and rearranging for the attenuation coefficient we obtain Equation 3.4.

$$\alpha = \frac{1}{(z_1 - z_0)} \ln\left(\frac{F D_1 T_{w\to m} R_{m\leftarrow w} T_{m\to w}}{B D_0 R_{w\leftarrow m}}\right) \tag{3.4}$$

Taking Fourier transforms F' and B', and introducing the angular frequency ω:

$$\alpha(\omega) = \frac{1}{(z_2 - z_1)} \ln\left(\frac{F'(\omega) D_1 T_{w\to m} R_{m\leftarrow w} T_{m\to w}}{B'(\omega) D_0 R_{w\leftarrow m}}\right) \tag{3.5}$$

Equation 3.5 enables calculating the frequency dependent attenuation from the Fourier transform of a measured frontwall and backwall signal. The calculation for the diffraction correction D_i ensues.

Diffraction Correction

Beam spreading occurs due the finite extent of transducers and it is generally desirable to eliminate this element of attenuation as it does not reflect the material properties. A widely accepted solution is that from (Lommel 1885) and later (Rogers & Van Buren 1974) who provided an explicit solution shown in Equation 3.6 . The underlying assumptions include a piston is set in a rigid baffle which oscillates harmonically into a homogenous fluid. In addition, a normal beam is considered such refraction effects are neglected.

$$D_i = 1 - e^{-i\left(\frac{2\pi}{s}\right)}\left[J_0\left(\frac{2\pi}{s}\right) + iJ_1\left(\frac{2\pi}{s}\right)\right] \qquad s = \frac{2\pi z_i}{k_i a^2} \tag{3.6}$$

Where a is the transducer radius, z_i is the axial distance from the transducer, k_i is the wave number in the relevant medium, and J_0 represents the zeroth order Bessel function.

3.2.2 Methodology

The experiment adopts a pulse-echo immersion configuration where multiple measurements are taken at varying water paths (see Figure 3.6). The probe is maintained approximately centrally above the sample (see Figure 3.6c) to ensure no side wall reflections interfere. Each measurement is temporally averaged to sufficiently supress electrical noise sources. The probe comprises a 2.25MHz, single crystal transducer (manufactured by Olympus) which is excited with a broadband pulse. Once a time-trace is digitized, the frontwall and backwall signals are separately windowed at their zero-crossings before calculating the attenuation as described in Equation 3.5.

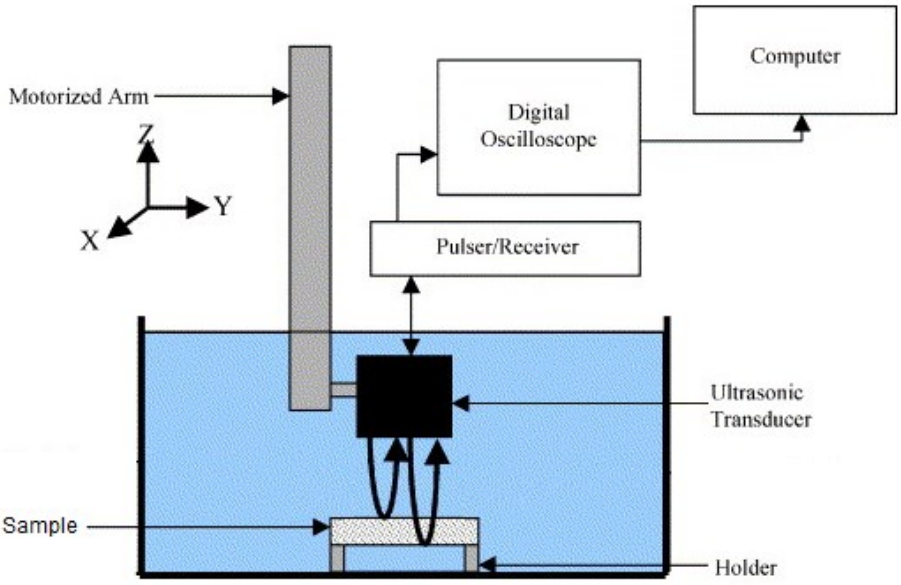

(a)

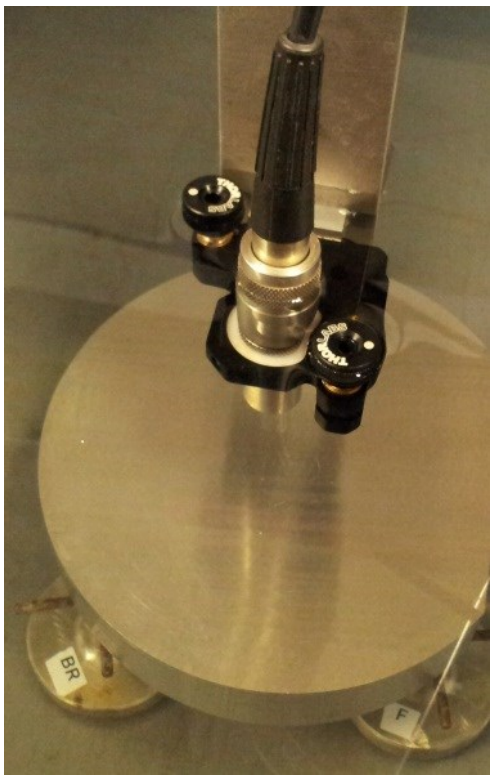

(b)

(c)

Figure 3.6: Experimental setup showing (a) schematic, (b) probe and a sample, (c) the Inconel sample.

The sample consists of an Inconel 617 block which was cut from a pipe. A macrograph of its microstructure is shown in Figure 3.7. Although no formal stereological analysis is performed here, qualitative analysis allows us to estimate the grain size of our sample in the order of 200-500μm. At the frequency used for this experiment with an approximate wavelength of 3mm, and thus with a wavelength to grain size ratio $\lambda d^{-1}=15$, suggests we should be operating within the Rayleigh regime.

Wave velocities

Before measuring attenuation, two remaining parameters need to be determined, in this case, ultrasonically. The velocity inside the sample is calculated by taking the time difference between successive backwall signals. The water wave speed is measuring the time-shift in the frontwall signal for a known depth translation of the probe (z-axis). Results shown in Table 3.1.

	Inconel	Water
Velocity (m/s)	5873.6	1494
Standard deviation (1σ)	n/a	0.6

Table 3.1: Properties for Inconel and water.

3.2.3 Results and Discussion

The measurements for the frequency dependent attenuation are shown in Figure 3.8. Attenuation can be seen to grow rapidly with frequency. As described in Section 2.3, the power law is used here to identify the attenuation mechanism; Figure 3.8 shows that a power coefficient of approximately 3 is obtained. This confirms that the attenuation is due to scattering as it is higher than the second order dependence which would be expected from absorption phenomena. The expected Rayleigh fourth order dependence is not measured however. This is probably due to the relatively high frequencies used ($\lambda >> d$ is not satisfied) which was a practical limitation of this experiment, as measuring at lower frequencies significantly reduces the attenuation values which increases the significance of experimental errors.

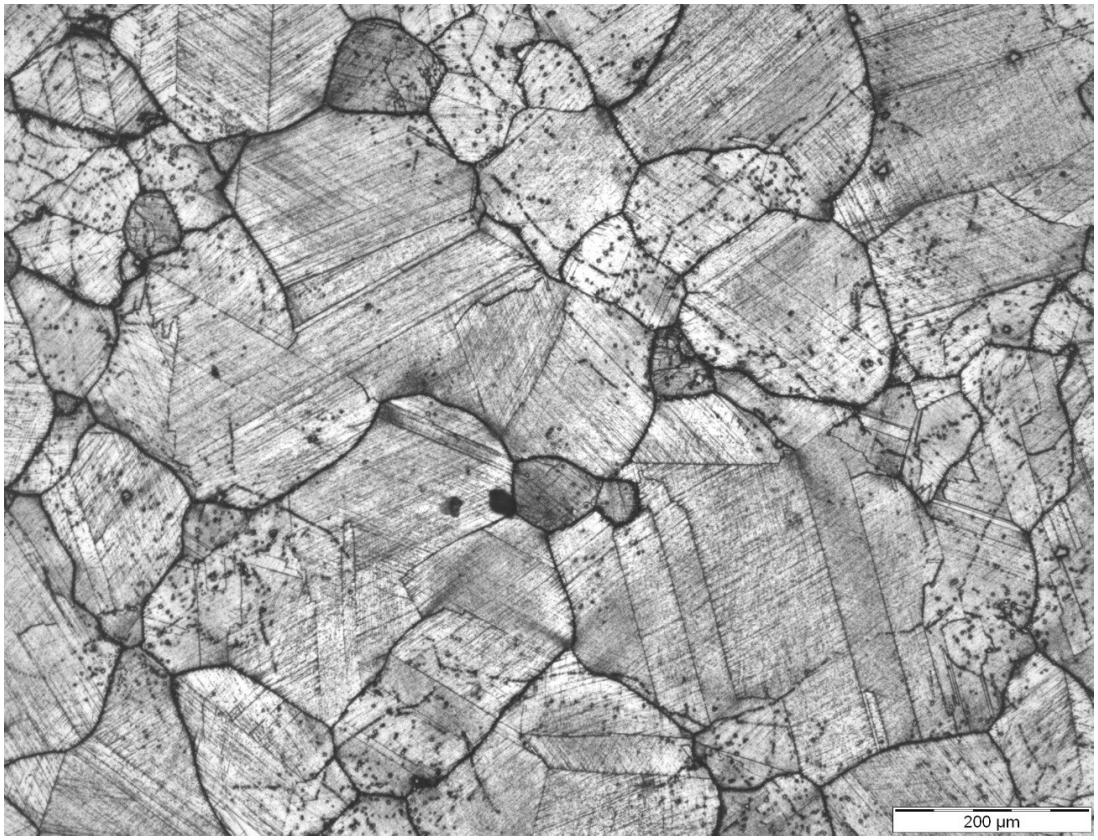

Figure 3.7: Macrograph of the Inconel 617 sample, revealing a coarse grain structure.

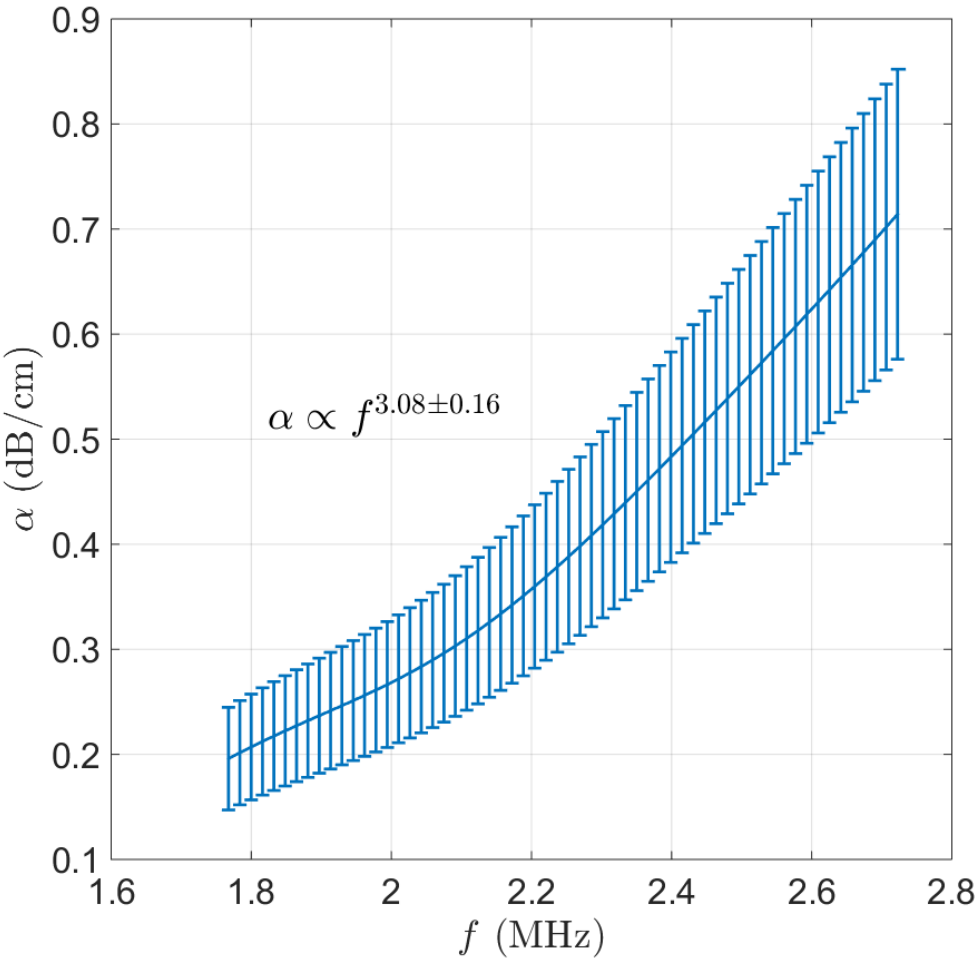

Figure 3.8 Attenuation versus frequency for the Inconel sample. The blue line connects the data points and the error bars show the standard deviation of 10 measurements.

3.3 Temporally and Spatially Incoherent Noise

Temporally and spatially incoherent noise is investigated here through two experiments. The first demonstrates conventional temporal averaging whereby temporally incoherent noise is averaged, and the second aims to replicate these results by exploiting, and thus demonstrating, the spatial incoherence of grain noise.

3.3.1 Temporal Averaging

The experimental setup constitutes a tank of glycerol in which a strand of hair (kindly provided by Misty Haith) is suspended to create a reflecting target. Glycerol is an attenuative but noiseless medium, such that any noise arises from electrical and thus temporally incoherent sources. A focused transducer with a 5MHz centre frequency, is immersed inside the fluid (see Figure 3.9) such that the target is within its beam path. In a pulse-echo setup, A-scans are recorded whilst varying the number of averages.

The data are collected by increasing the number of averages in steps from 1 to 4096 in multiples of 2, to produce 12 data points, and recording 10 measurements for each point. Figure 3.10 shows two typical A-scans obtained after 1 and 4096 temporal averages. The two largest signals comprise the main bang (MB) at around 0μs, and a dual backwall signal from the floor of the tank is found at around 105μs. The backwall consists of an initial weak reflection between glycerol and plastic, and a subsequent stronger reflection when the wave reflects off the plastic to air interface. In between these two signals, is the reflection from our object, but Figure 3.10a indicates the signal is initially embedded within the noise. Sufficient temporal averaging reduces the noise to reveal a reflection at 45μs in Figure 3.10b. The hair reflection is relatively low in amplitude in comparison with that of the backwall as it is a relatively small scatterer.

The SNR is quantified for each averaging step, by calculating the RMS of the noise, measured in an interval in the time record that is close to but but away from the hair reflection, and dividing it by the peak of the signal which is measured from the highest averaged A-scan (4096 in this case). As can be seen in Figure 3.11, the SNR improves with the square root of number of averages (or 10dB/dec in this case), as dictated by well-established theory.

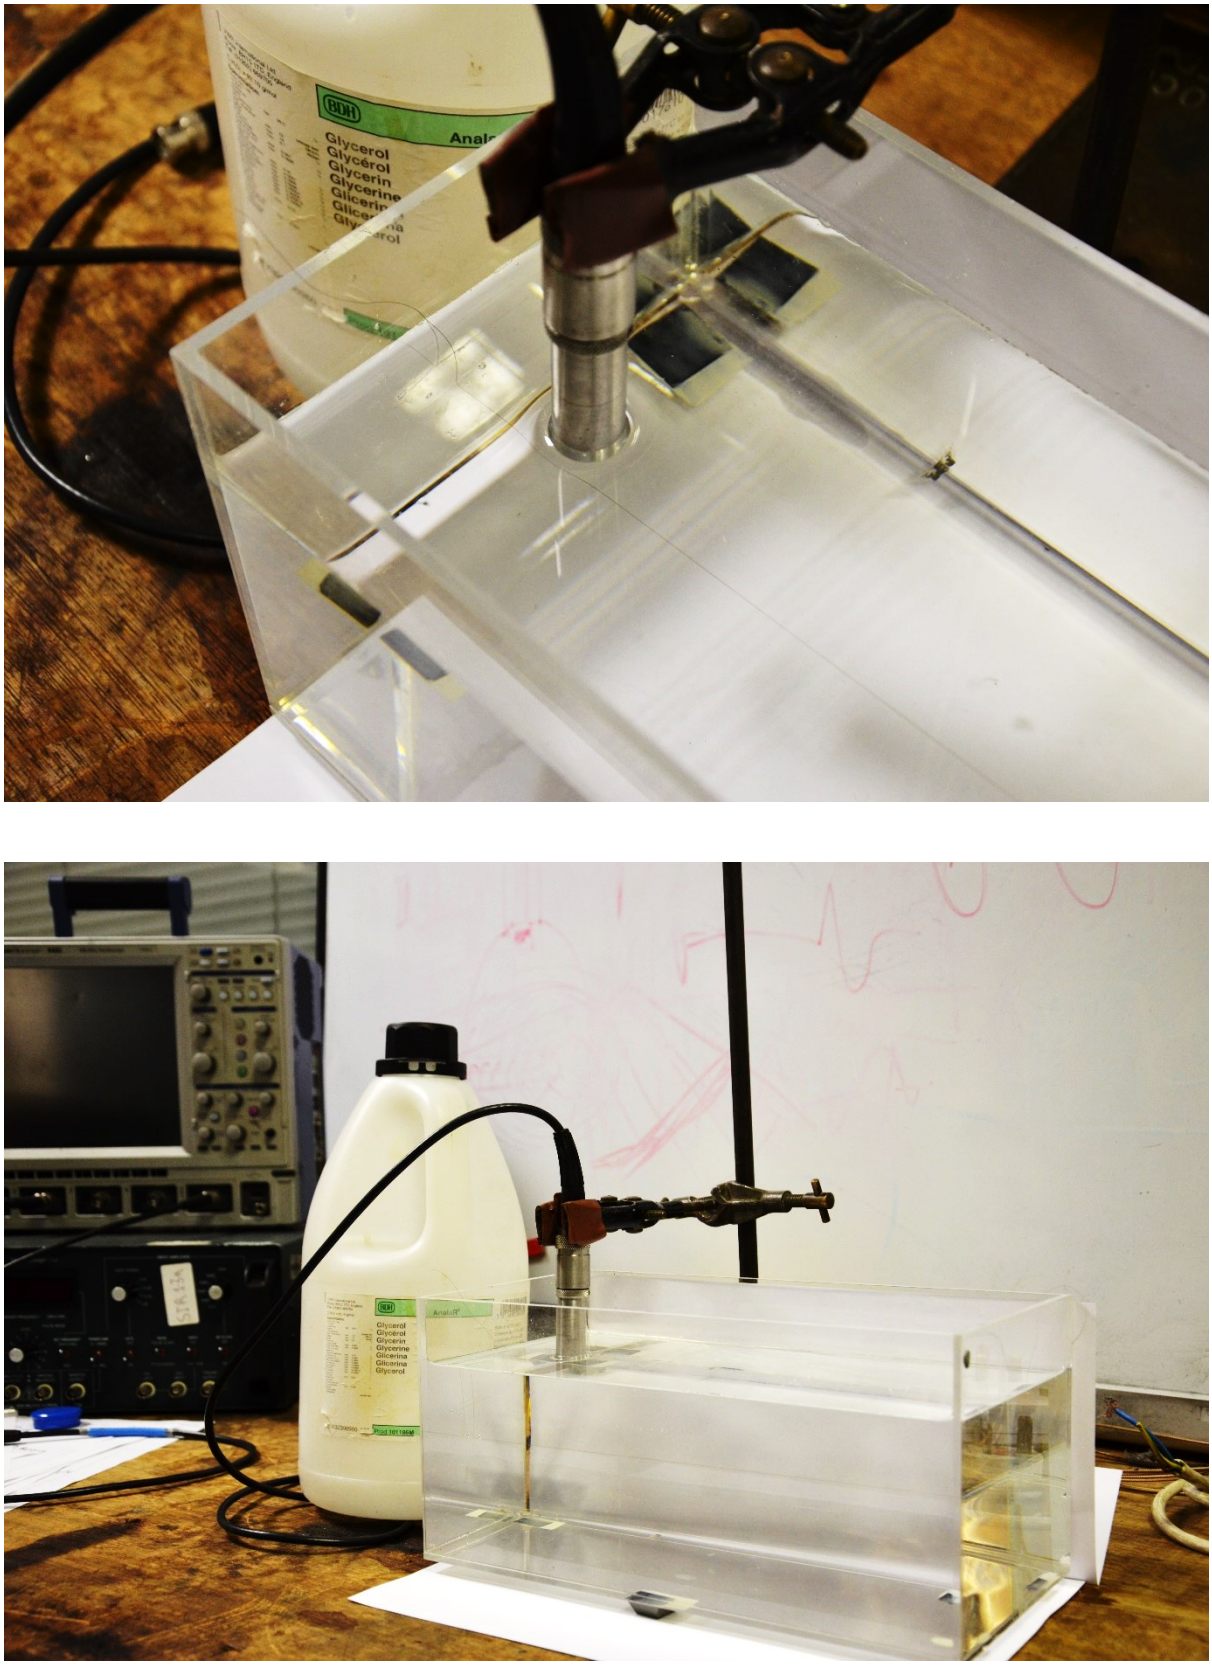

Figure 3.9: Experimental setup for a strand of hair suspended in Glycerol.

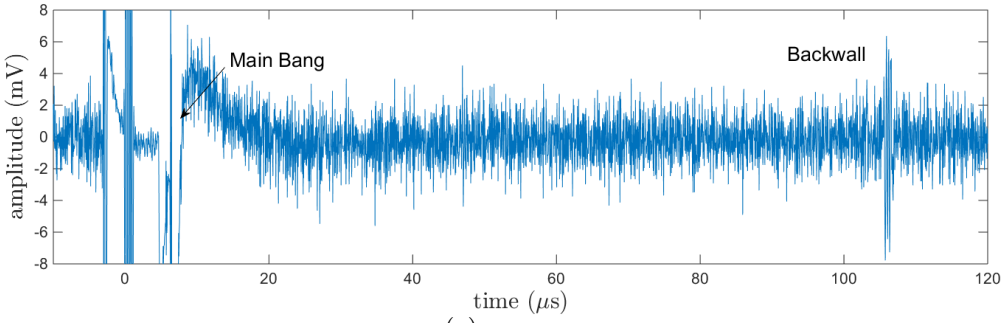

(a)

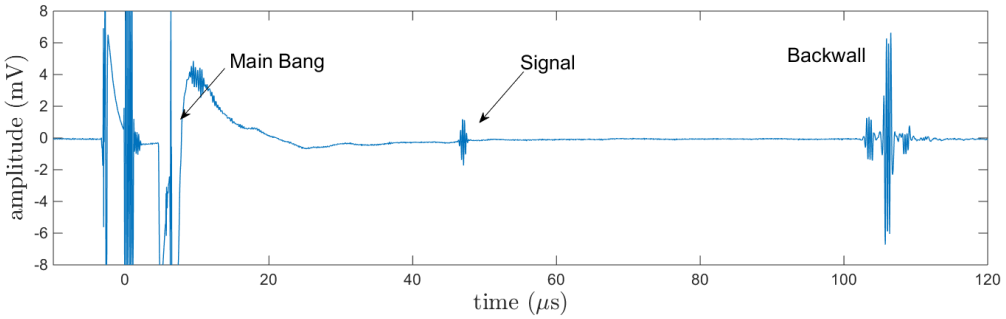

(b)

Figure 3.10 A-scans with (a) 1 average (b) 4096 averages, from a strand of hair suspended in Glycerol.

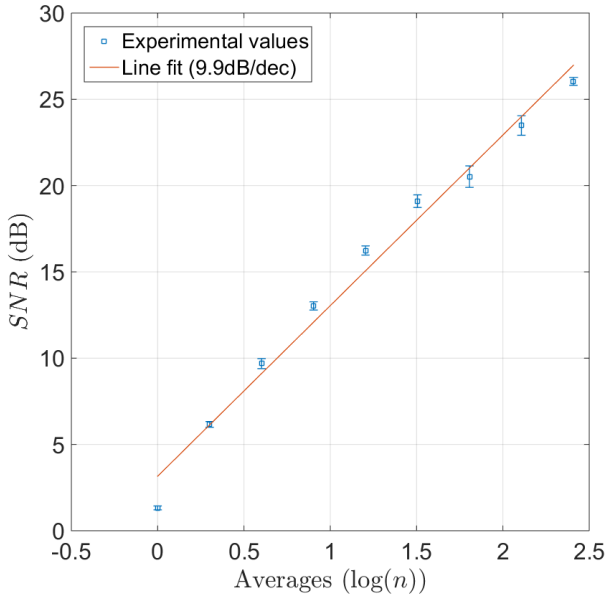

Figure 3.11: SNR versus averages, n, plotted for the measured experiment and theory for the hair in glycerol experiment.

3.3.2 Spatial Averaging

To demonstrate the spatial incoherence of grain noise, the above experiment is repeated, but this time replacing the hair in glycerol by a polycrystalline material immersed in water. The sample consists of an Inconel block of 44mm x 44mm x 200mm (shown in Figure 3.12) containing a sub-millimetre diameter SDH along its length and at 22mm depth. The transducer, with a centre-frequency of 20MHz, has a focal length of 95mm and is positioned such that it is focused at the depth of the SDH. The probe is then scanned along the length of the SDH to capture the SDH reflection embedded within independent measurement of backscatter noise. In order to supress electrical noise, 1000 temporal averages are used at each location. The success of this experiment depends on rastering the transducer by a sufficient distance, defined by the spatial correlation length of the grain noise(Yu et al. 2010), such that each noise measurement is independent. By shifting the probe at fine intervals, this was approximately calculated (correlation coefficient dropped to below 1/e) to occur for a 0.4mm shift. Figure 3.13 shows a typical A-scan (with no spatial averaging) in comparison to one which is spatially averaged 25 times. As can be seen, the initial SNR is close to unity, whereas the spatial averaging has enabled the SDH signal to emerge from the grain noise.

Similar to the previous post-processing procedure, SNR is calculated and plotted in Figure 3.14. As can be seen, the grain noise does indeed spatially average out and improves SNR approximately according to the same law as temporal averaging (10.2dB/dec was measured).

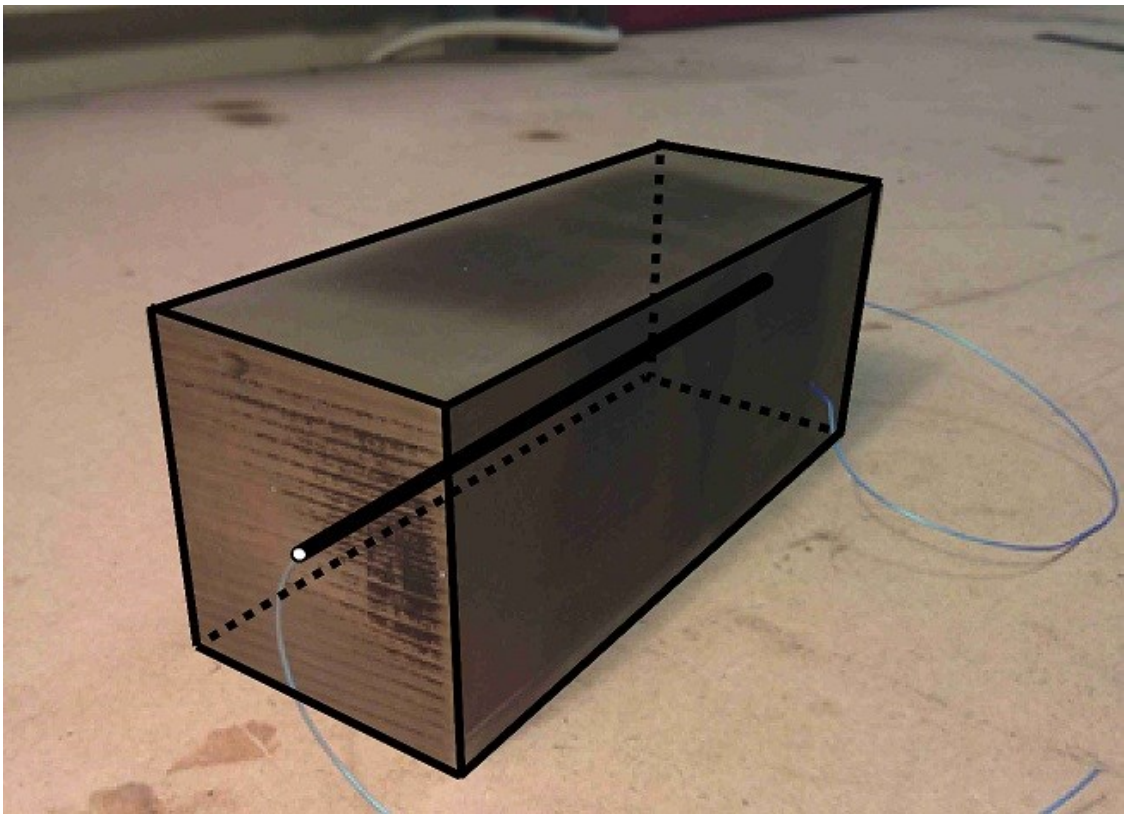

Figure 3.12: Inconel block with SDH machined along its centre.

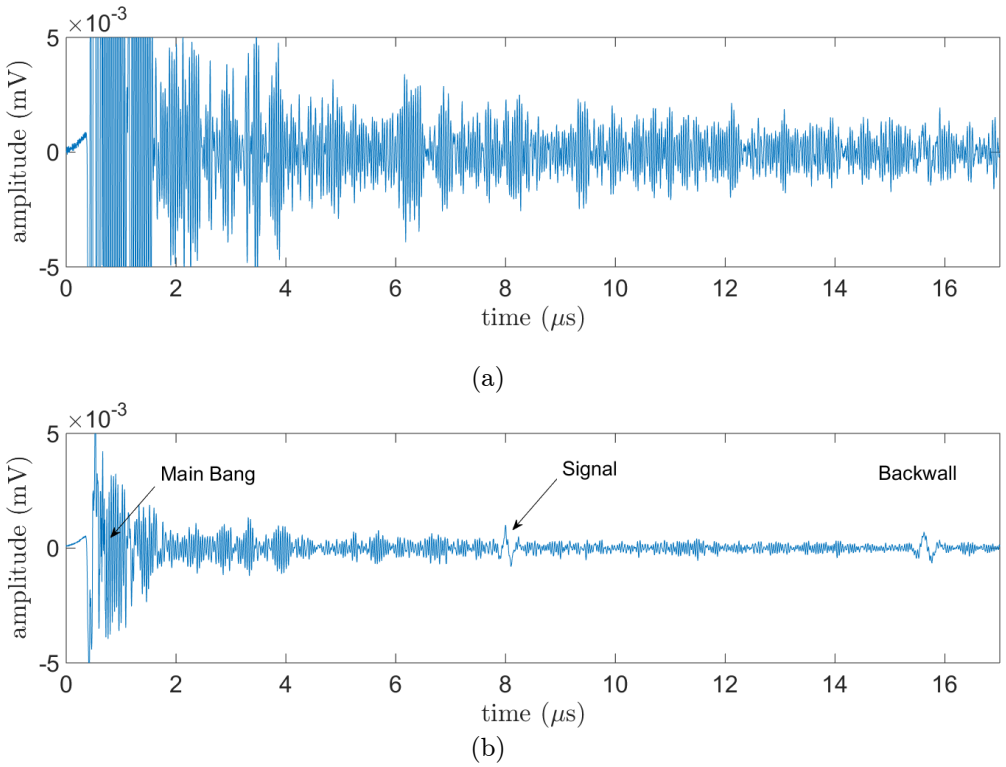

Figure 3.13: A-scans with (a) 1 average (b) 25 spatial averages, from a SDH within a scattering material, Inconel.

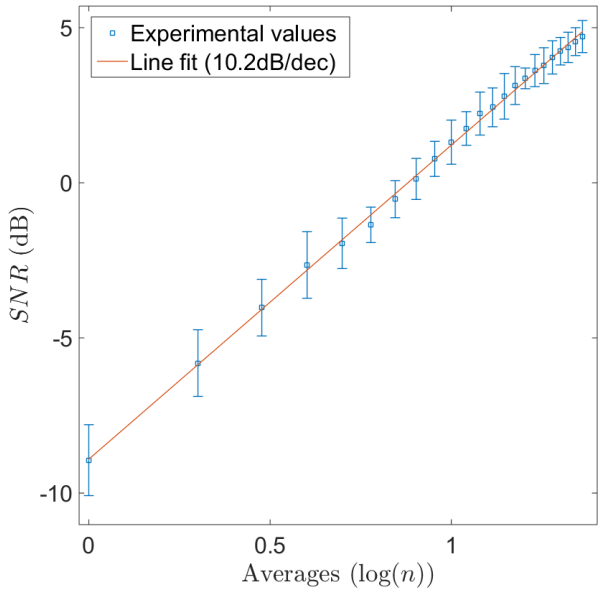

Figure 3.14: SNR versus number of spatial averages as measured from the Inconel block and compared to the theory. The measurements were taken with a focused immersion 20MHz probe.

3.4 Conclusion

This chapter set out to provide some practical evidence of scattering. Some of the inspection difficulties for upcoming USC components were established. The experiment revealed that it was surprisingly challenging to perform ultrasonic immersion tests which monitor the backwall signal whilst scanning the probe along the face of the sample. Several mechanisms contributed to this effect: a variation in velocity with position, a suspected large variation of grain sizes within a single component, and a high attenuation which varied from point to point.

Evidence of grain scattering was found by measuring the frequency dependent attenuation of a longitudinal wave propagating within a coarse grained sample. The measured power law was slightly lower than the expected fourth order Rayleigh scattering but still confirmed scattering behaviour. Well known averaging theory was demonstrated on temporally incoherent noise, but it was also shown to work when spatially averaging grain noise which is spatially incoherent.

Chapter IV

BENCHMARKING METHODOLOGY

In recent years, array probes and Full Matrix Capture (FMC) imaging algorithms have unlocked exciting possibilities for improvements, such as advanced spatial de-noising algorithms. Typically such algorithms are established by adopting a variety of methods, which renders it difficult to compare them. In order to progress and objectively compare these algorithms we must rely on robust methodologies to quantify their performance. This chapter proposes such a methodology to evaluate the detection performance of imaging algorithms, applying this by way of example to three FMC imaging algorithms; Total Focusing Method (TFM), Phase Coherent Imaging (PCI), and Decomposition of the Time Reversal Operator with Multiple Scattering (DORT MSF). The methodology considers the statistics of detection, presenting the detection performance as Probability of Detection (POD) and probability of False Alarm (PFA). A test sample of coarse grained Inconel 625, manufactured to represent materials used for future power plant components and containing some simple artificial defects, is used to illustrate the method on the candidate algorithms. The data is captured in pulse-echo mode using 64 element array probes at centre-frequencies of 1MHz and 5MHz. All three algorithms are shown to perform very similarly when comparing their flaw detection capabilities on this particular case.

This work is published in IEEE Trans. Ultras. Ferr. Freq. Control 2014 [P2].

4.1 Introduction

The emergence of advanced FMC array imaging algorithms, such as those mentioned in Section 2.6, is an exciting opportunity for NDE of scattering materials. The additional spatial information captured by FMC approaches can be exploited to ultimately improve NDE performance. If we are to progress, and adopt more advanced imaging algorithms, it becomes increasingly important to rigorously distinguish imaging algorithm performance. Answering this question however first requires a widely accepted definition of performance and a robust method of measuring it. Various methodologies have previously been adopted to compare imaging algorithms (Trahey et al. 1986; Lardner et al. 2014; Jie et al. 2013) where Signal-to-Noise Ratio (SNR) enhancement is commonly used to convey merits. As there exists no widely accepted method, to our knowledge, there remains an opportunity, and indeed an important need, to define an objective method and find a basis for future comparisons.

This chapter proposes such a methodology to evaluate, from an NDE perspective, the detection performance of different imaging algorithms. Here, the subsequent Section 4.2 discusses the importance to consider both detection and false alarm rates whereby Receiver Operating Characteristic (ROC) (Metz 1978) analysis is introduced as a tool to quantify detection performance. By way of example, three candidate algorithms (introduced in Section 4.3); TFM, PCI, and DORT MSF, are tested on a coarse grained power plant material in Section 4.6. Prior to the experimental results, the proposed approach is proven using simulated data (see Section 4.5) where complete knowledge of the inspected medium exists. Details on data collection for each method, experiment and simulation, are presented in Section 4.4.

4.2 Imaging Performance

NDE usually involves two aspects: flaw detection and flaw characterization; both involve some uncertainty and the intent of the latter is here assumed to also include sizing. Whilst imaging performance for NDE applications can be defined by these two indicators, the challenge lies in how to quantify them. Various methodologies have been adopted to

compare imaging algorithms; an alternative method is reasoned below by considering detection and characterisation.

4.2.1 Flaw Detection

Detection is purely related to signal and noise amplitudes, commonly packaged as Signal-to-Noise Ratio (SNR). Measuring SNR in cases where defects are not detected, by definition, is difficult when noise dominates signal, making it unusual for SNR values of less than one to be discussed. In a research context however, a priori knowledge of the defects often exists, especially when they are simulated or artificially created, which can be exploited to enable comparisons. Thus as a simplification, and applying this to the example of an ultrasound image, it is commonplace (Wilcox 2011; Jie et al. 2013) to assume that the pixel intensities corresponding to the predicted flaw-response location contribute as signal (although it contains signal plus noise) whilst all remaining pixels are classified as noise. The appropriate area of the image corresponding to the defect can be determined by predicting the flaw response through simulation or by observing a well detected experimental case.

Whilst SNR is a useful indicator of imaging performance, it is valuable to note that there is not always a benefit in maximising it once detection is achievable. This is illustrated in Figure 4.1 to Figure 4.3, where three candidate algorithms (defined later) image a flaw using the same raw data. Because each algorithm produces an image which allows detection of the flaw, albeit at different SNRs and perhaps characterising ability, it can be reasoned that each achieve equal detection performance. As an example, once detection is achieved it is trivial to increase SNR, as merely squaring the pixel values in an image will double its SNR (Wilcox 2013); the difficulty lies in actually improving detection rates.

Once signal and noise values are calculated, determining whether a target is detected can be achieved by setting a threshold. Using a predefined threshold for detection, however, such as the widely accepted SNR of 6dB, can be lenient towards non-linear algorithms and does not allow a trivial calculation to determine false alarms.

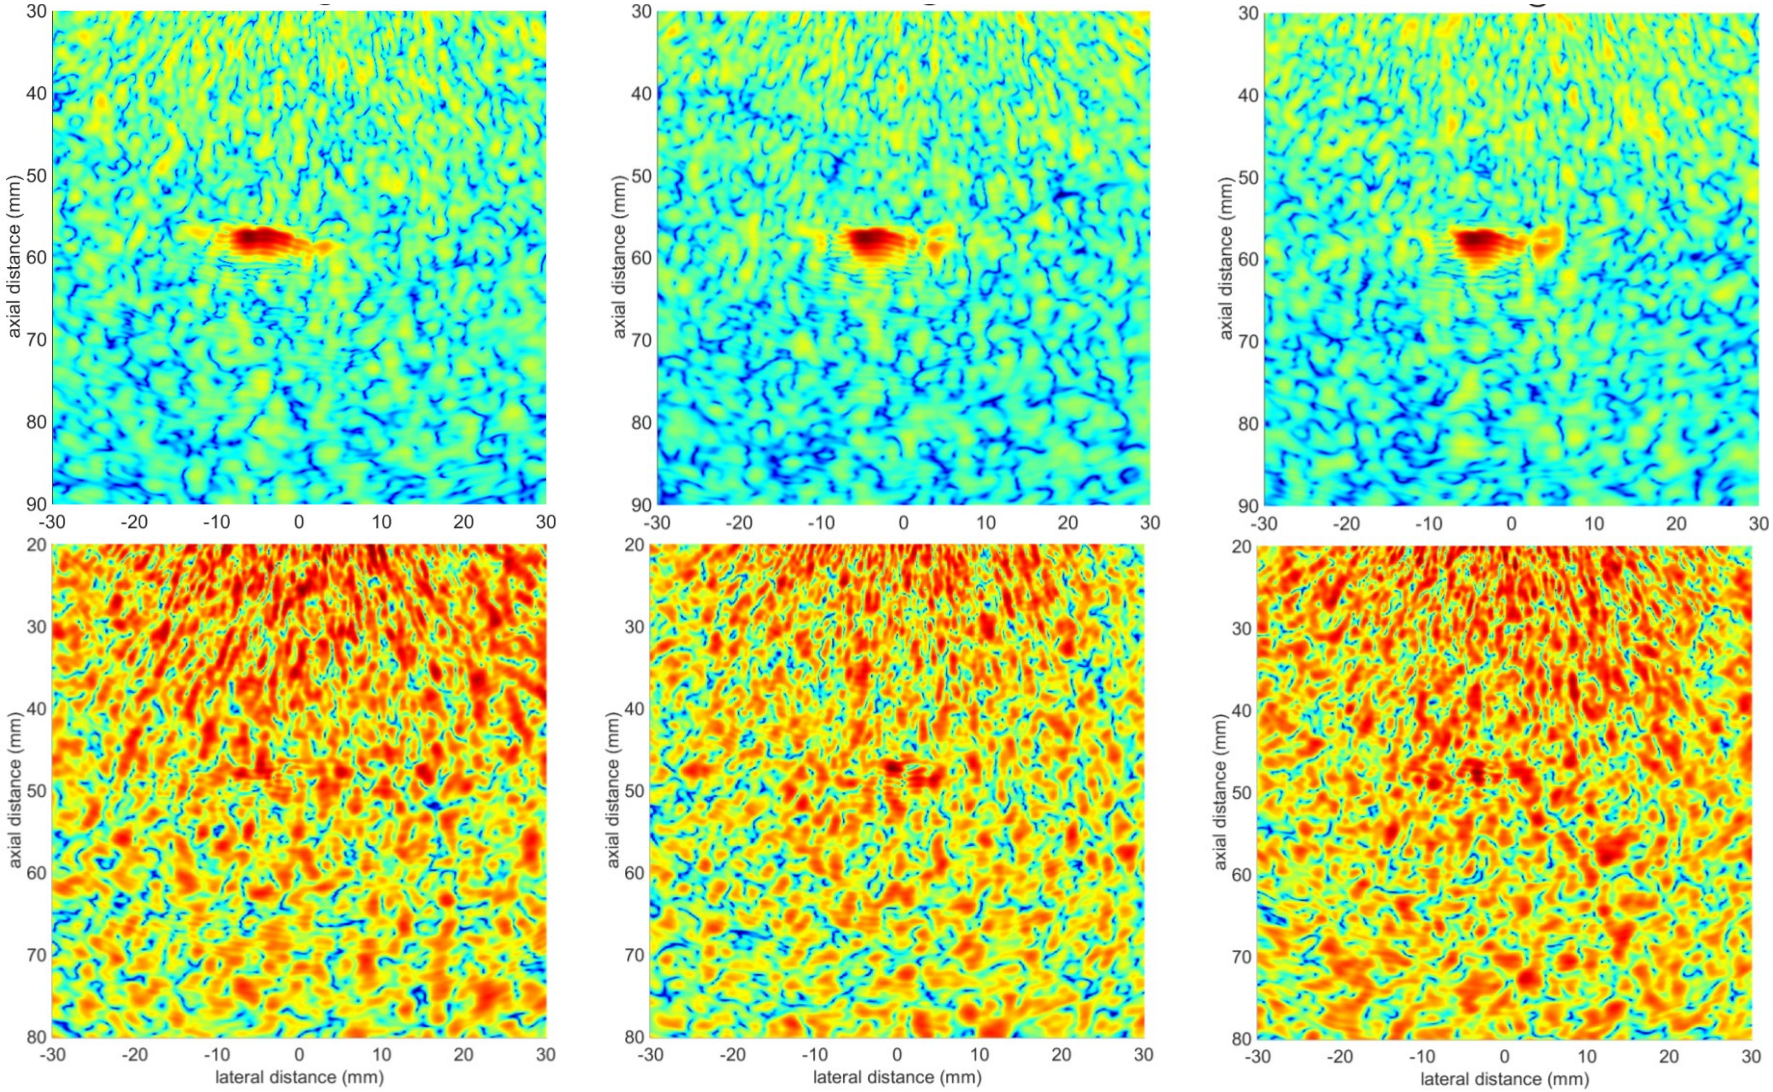

Figure 4.1: TFM images of three FMC captures along the SDH length, using a 64 element 5MHz array. TOP: SDH2@60mm. Bottom: SDH1@50mm (60dB) normalisation. The colour scale denotes increasing intensity from blue to red.

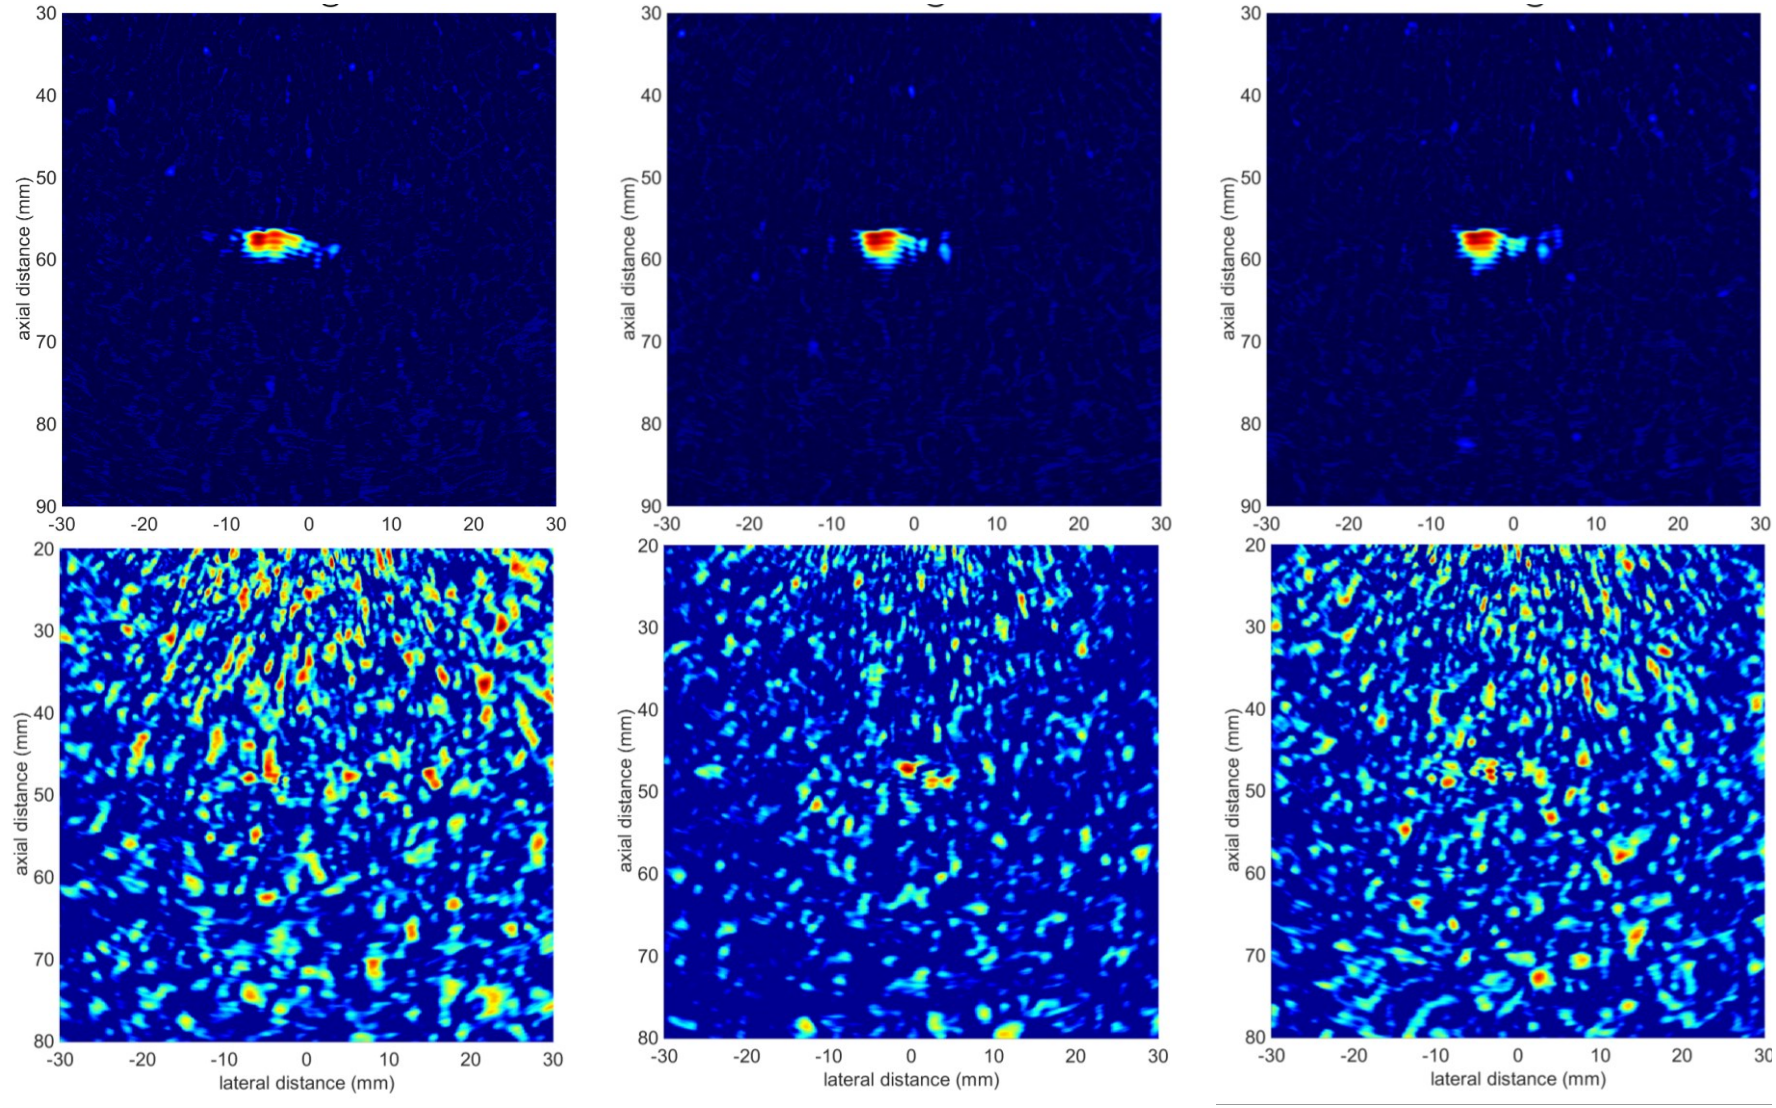

Figure 4.2: PCI images of three FMC captures along the SDH length, using a 64 element 5MHz array. TOP: SDH2@60mm. Bottom: SDH1@50mm (60dB) normalisation. The colour scale denotes increasing intensity from blue to red.

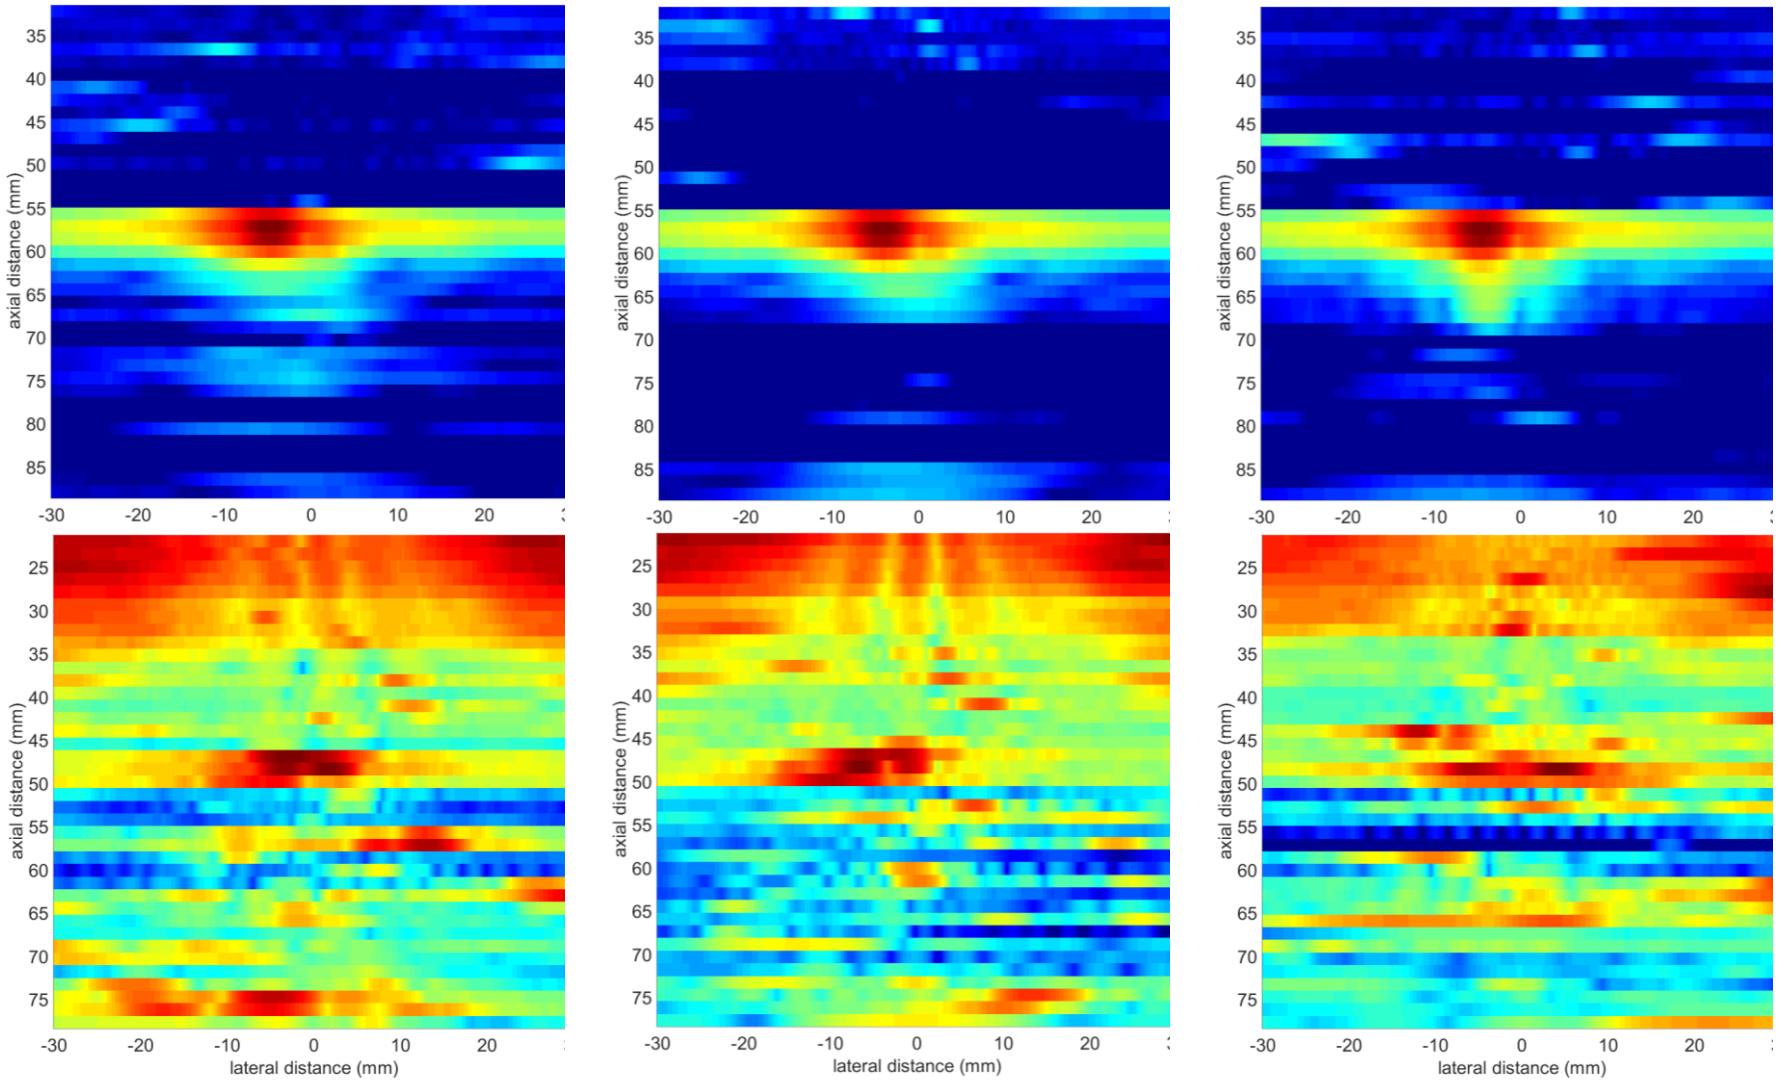

Figure 4.3: DORT MSF images of three FMC captures along the SDH length, using a 64 element 5MHz array. TOP: SDH2@60mm. Bottom: SDH1@50mm (60dB) normalisation. The colour scale denotes increasing intensity from blue to red.

False alarms are important as detection should also incorporate reliability; in addition to defects which are detected or missed, the number of false calls each algorithm generates has to be recorded. Without this consideration, an algorithm which detects only the defects, and one which classifies everything as a defect both maintain a perfect detection record. To calculate false alarms, the truthful outcome of any inspection must be known which once again requires the presumption that knowledge of the defects (true-positives) exists and thereby also the potential false-positives.

This leaves two challenges: how to set a threshold and how to include reliability. One approach, which has enjoyed success in a variety of fields concerned with detection problems, is to rely on Receiver Operating Characteristic (ROC) analysis (Metz 1978). The ROC method was originally developed for RADAR to select an optimum operating point at which to distinguish enemy targets from allies and clutter (Zweig & Campbell 1993). Later the medical research community adopted the approach to interpret diagnostic accuracy and it subsequently received plentiful attention leading to further developments of analytical tools to interpret ROCs (Metz 1978). ROC analysis demonstrates the performance of a classifier, or for instance an imaging algorithm, by calculating true-positive and false-positive rates, or in NDE terminology: Probability of Detection (POD) and Probability of False Alarm (PFA). This is an attractive approach as it circumvents the need for setting an arbitrary decision threshold by instead intentionally varying the decision threshold and observing the change in outcomes.

The ROC method is also well equipped to operate in a low SNR environment, where defects are on the verge of detection. In order to compare imaging algorithms, particularly for NDE of difficult materials, it is precisely these conditions that are of interest. Comparing two different images in terms of their SNR alone, especially when both SNRs are low, does not often allow a meaningful comparison. For instance, the SNRs may have to be situated on opposing sides of the, for example 6dB, threshold before being able to distinguish detection performance. It must be noted however that it is not suggested to use ROC analysis to redefine detection standards and select a new detection threshold; it is merely suggested as a useful tool for evaluation and comparison of candidate algorithms on cases of interest.

4.2.2 Flaw Characterisation

Once flaws are detected, their characterisation becomes important. To truly assess this aspect requires different types and sizes of defects which is usually problematic for experimental studies. This is also true for the sample that we will discuss later in this chapter, where the defects are relatively large when compared to a typical wavelength used in inspection. Inevitably, an investigation based solely on this case would quickly conclude that lower inspection frequencies are superior as grain scattering effects are reduced. It is therefore important to acknowledge that higher frequencies are in general necessary to detect smaller defects (although this can become a more complicated issue when considering scattering, this is beyond the scope of the current discussion).

A previously adopted approach has introduced a resolution parameter (Holmes et al. 2005) which measures the area of pixels which are within a 6dB (e.g.) range of the peak intensity signal-pixel. Measuring this however, commands detection with relatively high SNR as noise-pixels will otherwise contaminate the measurement. Difficult materials exhibiting low SNR, such as coarse grained materials, are therefore ill-suited for such resolution measurements. While characterisation remains an important performance indicator, when considering inspection of difficult materials, it follows secondary (in sequence) to detection. Thus this methodology will focus on detection and methods for assessing characterisation will be left to future investigation.

4.3 Imaging Algorithms

The three example algorithms considered were introduced in Section 2.6 and are hereon summarised; the Total Focusing Method (TFM), Phase Coherent Imaging (PCI), and Decomposition of the Time Reversal Operator with Multiple Scattering Filter (DORT MSF). Figure 4.1 to Figure 4.3 compares the images using the different algorithms of a relatively well-detected and weakly-detected Side Drilled Hole (SDH) in a coarse grained sample (further described in Section 4.3.1) using a 5MHz array. Figure 4.4 compares a similar image of SDH for the TFM and PCI algorithm using a 1MHz array.

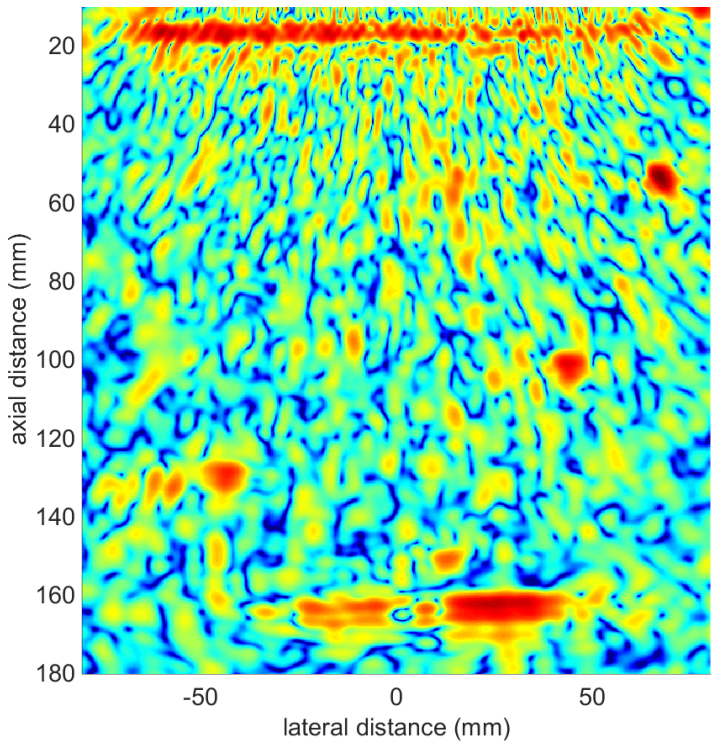

(a)

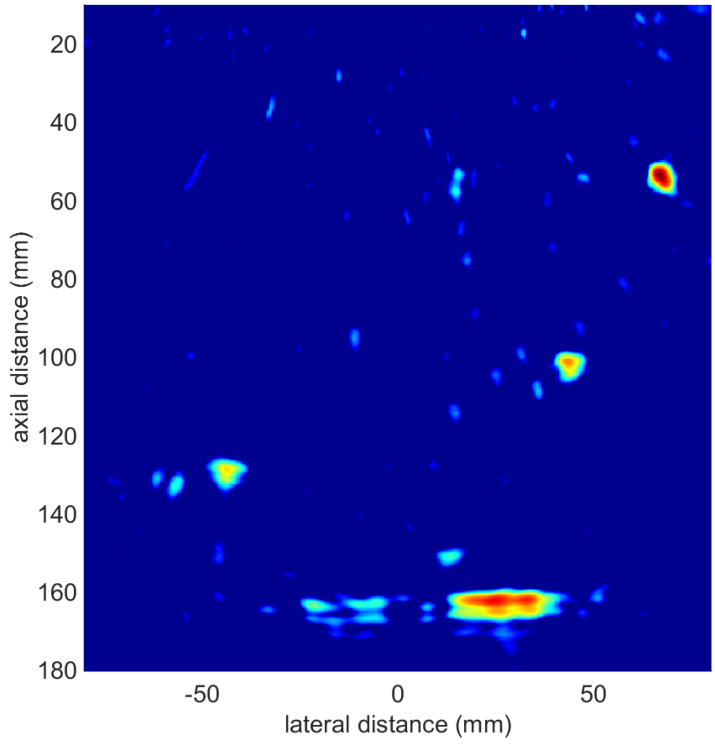

(b)

Figure 4.4: (a) TFM (b) PCI array images of a single FMC capture of three SDHs at depths (50mm,100mm,150mm) and a corner reflection from a step at (135mm,-50mm), using a 64 element 1MHz array. The SDH at (150mm,10mm) is hardly visible, but there is a clear gap in the backwall at 160mm (40dB normalisation). The colour scale denotes increasing intensity from blue to red.

4.3.1 Total Focusing Method

TFM (Holmes et al. 2005) is named for creating an image which is synthetically focused at every pixel by using the standard sum-and-delay operation, more widely known as beamforming. The transmitting *tx* and receiving *rx* array elements are denoted i and j respectively, the x-coordinate signifies lateral position, the z-coordinate is the axial dimension, c_0 is the wave speed, and the FMC matrix $H_{ij}(t)$ of dimensions $N \times N$ where N is the number of elements in the array. The image intensity, I, at location x, z, is given by Equation 4.1. The Hilbert Transform can be used to smooth the image.

$$I_{TFM} = \left| \sum_{i=1}^{N} \sum_{j=1}^{N} h_{tx,rx} \left(\frac{\sqrt{(x_i - x)^2 + z^2} + \sqrt{(x_j - x)^2 + z^2}}{c_0} \right) \right| \tag{4.1}$$

4.3.2 Phase Coherent Imaging

Phase Coherent Imaging (Camacho et al. 2009; Camacho & Fritsch 2011) amplifies the contribution of phase information by multiplying the original TFM image by a weighted matrix C according to the statistical phase φ variability of all N^2 time traces (i,j) for a particular position (x,z). When all phases are equal, $C=1$, maintaining pixel intensity, whereas when they approach a uniform distribution $C \rightarrow 0$, reducing pixel intensities. This operation has shown promise (Camacho & Fritsch 2011) to inspect structurally noisy materials as backscatter signals are expected to exhibit greater phase variability than flaw signals. In this case I is given by:

$$\begin{aligned} I_{PCI} &= I_{TFM} C \\ C &= 1 - \sqrt{\mathrm{var}(\cos\varphi(x,z,i,j)) + \mathrm{var}(\sin\varphi(x,z,i,j))} \end{aligned} \tag{4.2}$$

4.3.3 DORT MSF

DORT (Prada et al. 1996) relies on calculating the singular value decomposition of the time-windowed, T, frequency, f, domain FMC matrix $K(T,f)$ to achieve imaging of the inspected medium. In weakly scattering conditions, each i-th scatterer in the medium is related to a singular space of the i-th singular value λ_i. The corresponding singular vector V_i is the signal of delays to apply to the array to focus onto the scatterer. When structural

noise dominates, significant singular values are no longer associated with scatterers of interest. Part of this noise, multiple scattering, can be reduced by exploiting the deterministic phase relation along the anti-diagonals of the K matrix which occurs even for an unknown distribution of disorder, in this case the scatterers inside the medium, for single scattering. This procedure is termed a Multiple Scattering Filter (MSF) (Aubry & Derode 2009b; Aubry & Derode 2009a) and has shown great promise for detection in coarse grained materials (Shahjahan, Aubry, et al. 2014). DORT MSF involves applying the regular DORT algorithm to the filtered K matrix, K_f. Here G^* is the complex conjugate of the Green's function for a homogenous medium. In this case I is given by:

$$I_{DORT}(T,f) = \lambda_i \left| V_i(T,f) G^*(T,f) \right| \tag{4.3}$$

The DORT MSF algorithm includes a detection criterion which aims to identify singular values which are associated with scatterers of interest and therefore merit back propagation. This is achieved by considering the statistical properties of the singular value distribution acquired from backscatter in a defect-free area for different frequencies f and depths related to time T. This method requires numerous FMCs of defect-free material which is known to be of similar microstructure. In applications of known, well-controlled, microstructure, this has been shown to generate significant improvement in detection performance (Shahjahan, Aubry, et al. 2014). However, in circumstances where the microstructure is spatially varying, such as in this study, the backscatter of defect-free material cannot be related to that from areas which contain defects, and therefore this potential advantage cannot be realised.

4.4 Data Collection

This chapter relies mainly on experimental data obtained from a coarse grained material to demonstrate the proposed methodology whilst simulated data is used to investigate its limitations. The experimental investigation provides an interesting outcome on how the algorithms perform on a real material and the simulated data allows us to generate data with a known outcome in order to contrast the evaluation results with what is expected. Table 4.1 summarises the parameters for each method of data collection, where

both serve to acquire FMC data of imaging targets: a defect in structurally noisy material. As different imaging algorithms may exhibit varying levels of sensitivity towards defects located at large angles, the array is always placed centrally above the defect such that the most optimistic scenario is considered: the defect is located where the array has the highest sensitivity. In terms of post processing, the data is filtered about the centre-frequency of the array used. Similarly, whereas previous research (Li et al. 2013) has shown that, for certain inspections, image SNR can be optimised by using sub-apertures of the array, this investigation uses the full aperture.

Parameter	Simulation	Experiment: high freq. array	Experiment: low freq. array
Number of elements	32	64	64
Element pitch	0.6mm	0.6mm	2.1mm
Element width	0.5mm	0.5mm	1.6mm
Element length	15mm	15mm	20mm
Centre-frequency	5MHz	5MHz	1MHz
Bandwidth (-6dB)	64%	65%	50%
FMCs per defect	100	˜10	˜10

Table 4.1: Summary of simulation and experimental parameters.

4.4.1 Experimental Data Collection

The experimental data are collected from the Inconel 625 block described in Section 3.1. To summarise: the sample is representative of future power plant components and contains three 5mm diameter Side-Drilled-Holes (SDH) found at various depths as seen in Figure 3.1. The FMC data are acquired using two faces of the sample, in a contact configuration with the array orientated perpendicular to the SDH, to provide five imaging targets at different depths: 50mm, 60mm, 100mm, 110mm, and 150mm. For each target, array data sets are acquired at a series of locations along the SDH's length, although in each case centrally above it, to capture multiple FMCs of the same defect surrounded by a varying microstructure. The maximum number of FMCs which can be acquired in this way is limited by the required minimum step between the measurement data sets, and the maximum available scan area, determined by the thickness of the sample. A minimum

separation distance ensures that each measurement is independent, dictated by and approximately equal to the grain size. In this case, the constraints permitted approximately ten independent FMC measurements per imaging target. Two 64 element array probes (manufactured by Imasonic, France), with different centre-frequencies of 1MHz and 5MHz, are employed to represent a relatively low and high inspection frequency (see Table 4.1). The arrays are controlled by an OPEN System (manufactured by Lecoeur, France) and excited using a negative square pulse having a time duration equal to half the period at the probe centre-frequency.

4.4.2 Simulation of Scattering Materials

In order to control the outcome of an inspection, an existing simulation model (Jie et al. 2013) is used to generate FMC data which are representative of those obtained from scattering materials. It assumes that image speckle associated with grain scattering can be generated by calculating the acoustic response from a random sub-wavelength distribution of point scatterers. Under a single-scattering approximation, such a response can be calculated even for a large numbers of scatterers (e.g. 5000), whereas this would be more challenging for a complete scattering model.

The grain scatterers are represented by point scatterers which behave as omnidirectional scatterers with amplitude S^g. It has been shown (Jie et al. 2013) that 5 scatterers per wavelength squared is a representative and computationally efficient density of scatterers to include. At least 15 realisations are required to produce a converged SNR at this density (Jie et al. 2013).

Realistic defects such as cracks can be introduced into the model by describing their behaviour using a scattering matrix (Zhang et al. 2008), calculated for example by Finite Element models. This chapter however only considers a defect in the form of another point scatterer with a scattering amplitude S^d. The ratio, relative to the grain scattering amplitude S^g, is defined as a shown in Equation 4.4. Assuming that the scatterer density is fixed, the level of backscatter can be controlled by varying a, where larger values lead to stronger scattering.

$$a = \frac{S^g}{S^d} \tag{4.4}$$

The frequency response $E_{ij}(\omega)$ of a single point scatterer, for both grains and defects, is calculated using Equation 4.5.

$$E_{ij}(\omega) = A(\omega)B^{rx}(r_i,\omega)G^{rx}(r_i,\omega)B^{tx}(r_j,\omega)G^{tx}(r_j,\omega)S \tag{4.5}$$

$A(\omega)$ is the frequency spectrum of the transmitted signal and r is the distance from array element i or j to the point scatterer. B is the transducer directivity function (Miller & Pursey 1954) where rx and tx denote reception and transmission respectively. Similarly the Green's functions, G^{rx} and G^{tx} calculate the complex amplitude and phase received from a particular combination of receiver and transmitter element from the point scatterer for a homogenous equivalent medium. As dictated by the Born Approximation, the FMC matrix H can be obtained by super-positioning the individual contributions E^g of all n grain scatterers and that of the defect E^d as shown in Equation 4.6. The inverse Fourier transform of $H(\omega)$ is used to extract the time traces and obtain the time-domain FMC matrix $H(t)$.

$$H_{ij}(\omega) = E_{ij}^{d}(\omega) + \sum_{m=1}^{n} E_{ij}^{g_m}(\omega) \tag{4.6}$$

The model simulates a 32 element 5MHz array (see Table 4.1) with two different scattering environments representing a strong and weak detection scenario, by setting a equal to 0.1 and 0.25 respectively. Each case involves 100 FMCs for different random realisations of grain scatterer positions, each realisation contains a single defect located at an axial depth of 20mm centrally below the array. Figure 4.5 shows one TFM image for each case.

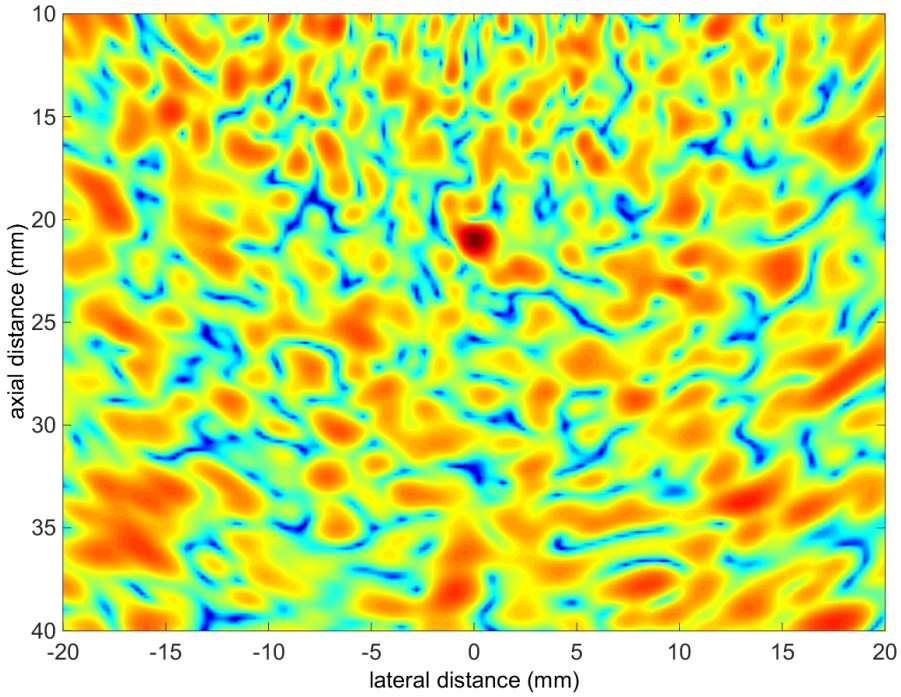

(a)

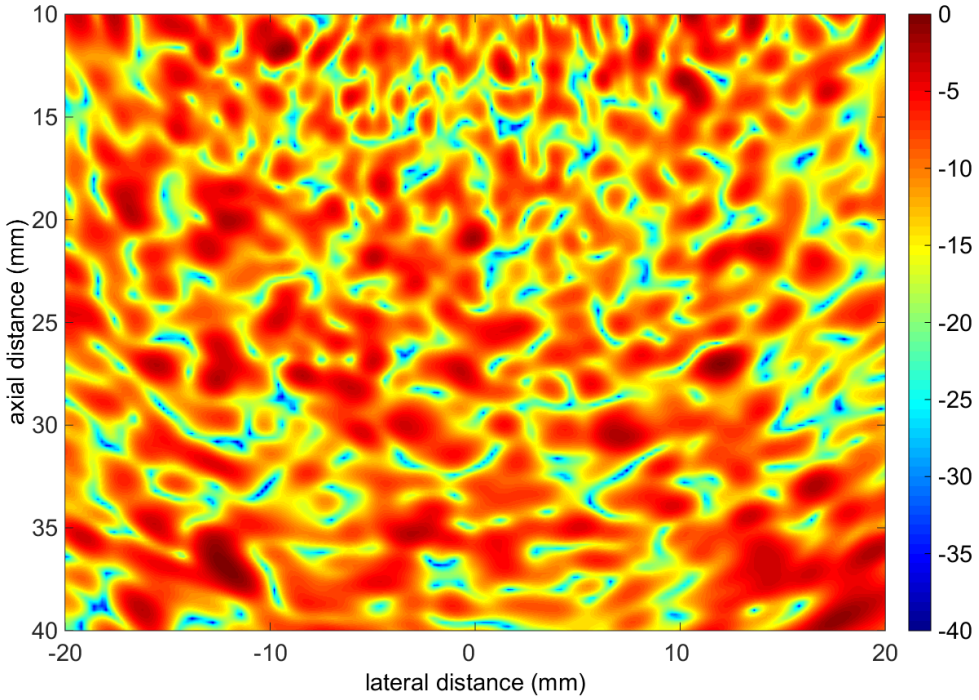

(b)

Figure 4.5: TFM image obtained from simulated data. Both images contain a point source at (20,0) with (a) a=1/10 and (b) a=1/4. The images are both normalised to their respective peak pixel intensity.

4.5 Measuring Reliable Detection

Once the FMC data are acquired, they are post-processed to produce images for each algorithm. True to the rationale defined in Section 4.2, ROC analysis is used to compare the performance of each algorithm. Whereas this analysis has been used in a qualitative manner where operators interpret images, the method outlined here relies on a basic form of image segmentation to quantitatively identify image objects from noise. This is in part possible due to the relative simplicity of the imaging targets considered which produce simple shapes (e.g. circular) which can be recognized by their increased intensity compared to the background. Hence there is no need for more complex pattern recognition, although this could be an avenue worth pursing for future investigations which consider multifaceted defects. The procedure for producing ROC plots follows these steps:

1) The images are converted into signal and noise datasets by adopting an image segmentation procedure. The procedure is possible because the defects are known, and is therefore only suited for a comparison methodology and not for real deployment with unknown defects. First, an area (a box) is mapped around the flaw response and the Root-Mean-Square (RMS) of the contained pixel intensities is calculated. This represents the signal amplitude and its area can be referred to as the signal box. For most experimental measurements, this will contain both signal and noise; a trade-off which increases and decreases depending on how well or poorly the object is detected. A TFM image of a simulated defect in a scattering material can be seen in Figure 4.6a where the outline depicts the signal box. The dimensions of the box are set to agree with the expected size of the flaw indication, determined by observing a well detected case, either experimentally or through simulation. Its shape can be modified according to the imaging target. For instance, in the case of inclined cracks, it may be more suitable to use two separated boxes which overlap with the crack tip responses.
2) Determining the noise dataset involves segmenting the remainder of the image (minus the signal box) into sections, here for demonstration purposes, of the same dimensions as the signal box. These can be considered as super-pixels which, similarly to the signal box, are calculated from the RMS value of the sub-pixels they contain.

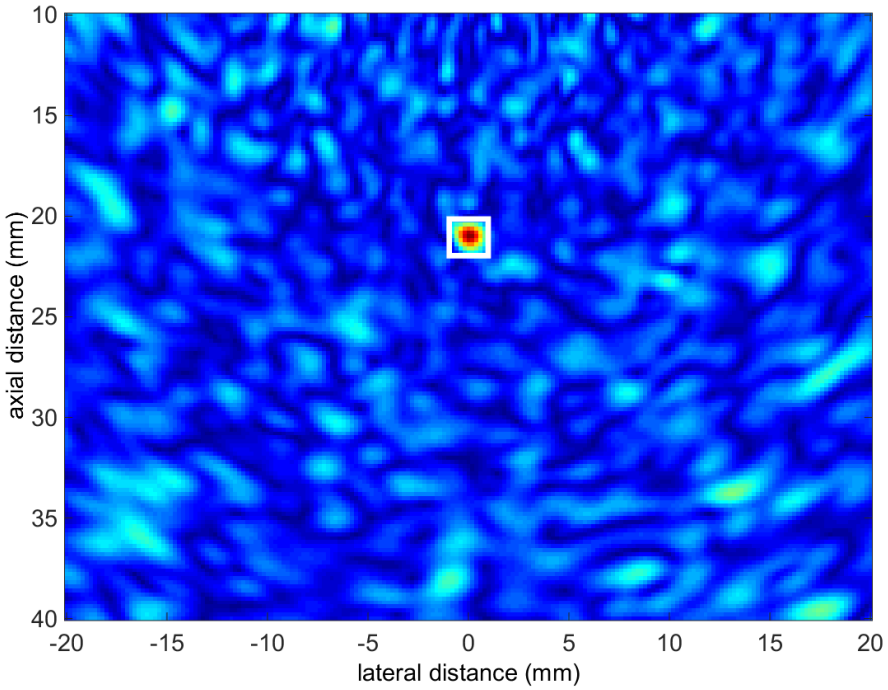

(a)

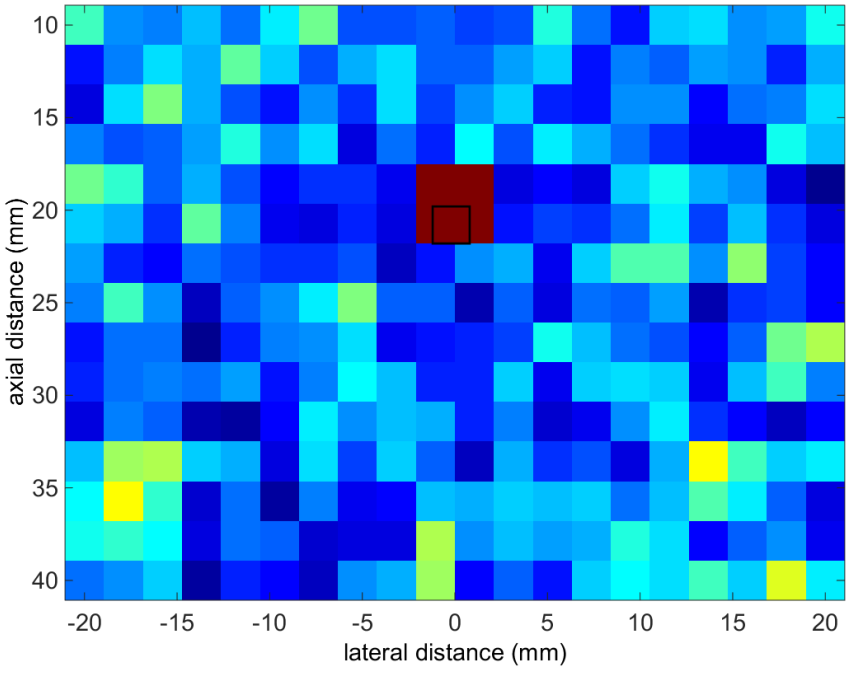

(b)

Figure 4.6 (a) Typical TFM image of a defect (outlined in white) in a scattering medium. (b) The processed version of the TFM image in (a) where the super-pixels and their RMS value has been calculated. The single super-pixel which corresponds to the defect is outlined in black. The area of omitted noise super-pixels which neighbours the signal pixel is shown in red.

The super-pixel version of the TFM image in Figure 4.6a is shown in Figure 4.6b. All the RMS values are then attributed to a noise dataset although any super-pixels which intersect with the signal box (the black outline in Figure 4.6b) are omitted from the noise distribution. The significance of the box dimensions is investigated in the following sub-section.

3) This procedure is repeated for all images of the defect to produce multiple images containing the same defect and different background noise. The use of numerous images, namely acquiring multiple FMC data of the same imaging target, is important to obtain two data distributions, one for the signal amplitudes and another containing noise values. Typical histograms for the distribution of a well and poorly detected case for the TFM algorithm, corresponding to average SNRs of 8dB and 3dB respectively, are shown in Figure 4.7. These histograms are obtained from simulated data using the model discussed in Section 4.4.2. Contrasting to real measurements, the simulated defect reflects a constant amplitude, which therefore only fluctuates according to noise in the defect box, and therefore its distribution is also similar to that of the noise. In experimental measurements however, like those of coarse grained samples, a variety of reasons can lead to signals fluctuating, including varying attenuation and anisotropy.
4) The ROC is calculated by establishing thresholds, varying from 0 to infinity, to the noise and signal distributions. Any data which exceeds the threshold is regarded as detected; the POD is calculated by the proportion of signal data which satisfy this criterion whilst the noise data contribute towards false alarms (PFA). The threshold at zero produces the point (1,1) in the ROC space whilst that at infinity produces (0,0), and any thresholds in between result in a combination of POD and PFA. The ROCs for the two cases shown in Figure 4.7 can be found in Figure 4.8. The well detected case shows a POD/PFA trade-off where almost all defects can be detected without any false alarms. For the poorly detected scenario (Figure 4.8b), however, a significant number of false alarms arise before all defects are detected.
5) The previous steps are repeated for each candidate algorithm to produce a ROC plot for each imaging target.

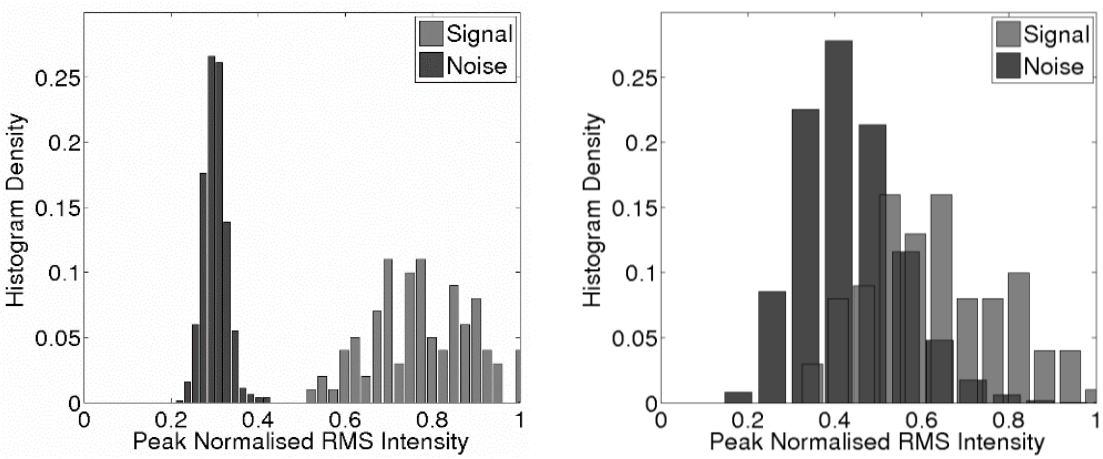

(a) (b)

Figure 4.7: Examples of the signal and noise distributions for (a) a well detected case with an average SNR of 8dB and (b) a poorly detected case with an average SNR of 3dB. The distributions were measured from simulated data imaged using TFM.

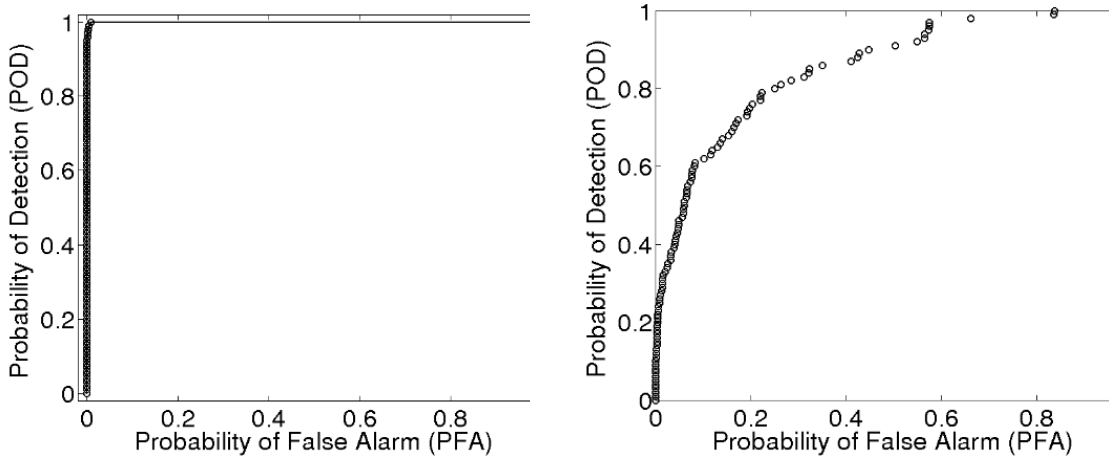

(a) (b)

Figure 4.8: Examples of ROC curves corresponding to the (a) well detected case and (b) poorly detected case of Figure 5 for the TFM algorithm.

It is important to investigate and consider the limitations of the proposed method. To enable a further discussion of ROC plots requires quantifying a parameter to describe them; the Area Under the ROC (AUC) is usually defined for this purpose (Hanley & McNeil 1982). Although there are drawbacks to reducing the ROC to a single index, it is a convenient description of the position of an ROC curve within the ROC space i.e. an AUC of unity denotes perfect detection and that of half describes the chance line corresponding to a randomised classifier.

4.5.1 Segmentation Procedure

Calculations of signal and noise amplitudes, whether this involves windowing a time-domain signal or spatially segmenting an image, are susceptible to the settings used. In this case, the signal and noise box dimensions represent such settings and affect the outcome of the ROC curve by either lowering or increasing the measured signal and noise amplitudes. Although the method was previously demonstrated by using parameters that were judged to be sensible for box sizes, the AUC is now measured as a function of varying the size for both parameters in order to investigate their impact and thus the sensitivity of the outcome to this judgement.

Figure 4.9 plots the AUC for both the well detected and poorly detected case. Both cases show that larger signal boxes increase false alarms, arising from the averaging effect which lowers the measured RMS amplitude. Noise boxes behave in the opposite way; small box sizes introduce peaks into the noise distribution which generate false alarms. From the well detected case, where it can be safely stated that the defect is detected, lower extremes for the noise box dimensions lead to a relatively sensitive result, which is undesirable. In this particular case the lower limit on the box size is 2mm for both signal and noise. For future investigations and use of this methodology, the appropriate choice of box size will depend on the imaging object and resolution of the algorithm considered. However, the results suggest that the method is largely insensitive to the box size, as long as extremes are avoided, and therefore this is not considered to be critical.

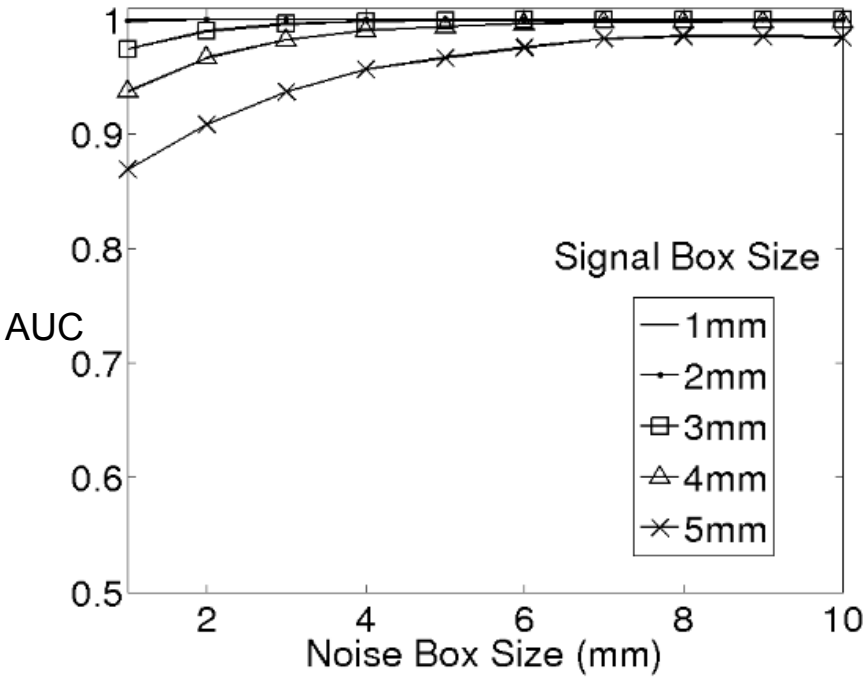

(a)

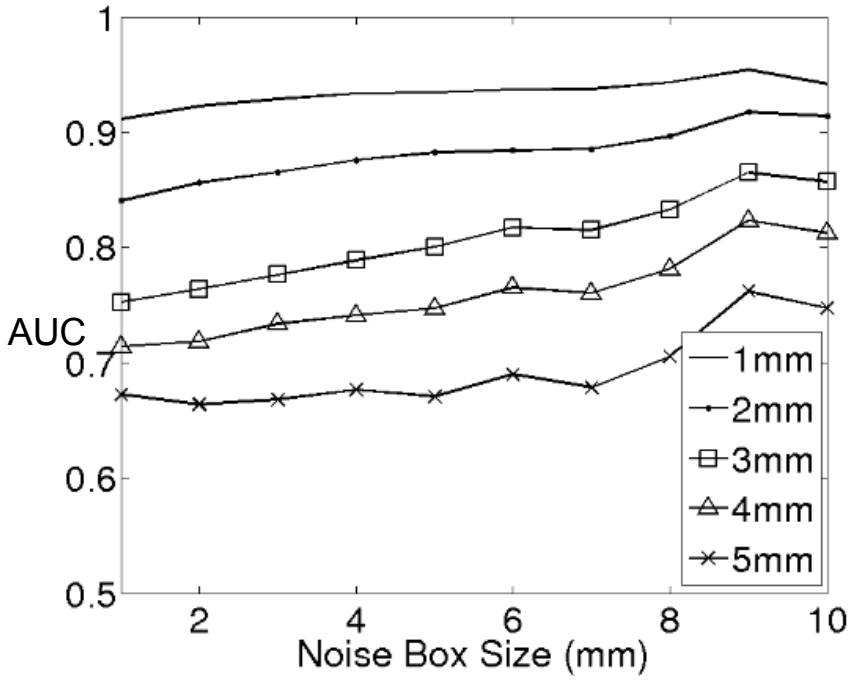

(b)

Figure 4.9: Influence of box size on ROC curve for the well detected case (a) and poorly detected case (b) of Figure 5. AUC is the Area Under the Receiver Operator Characteristics curve where 1 represents perfect detection; and 0.5 a random classifier.

The image segmentation procedure can be compared to a spatial or even a frequency filter (even though it is mathematically more akin to a decimation operation). By adopting the same box size for signal and noise, the calculation assumes only noise objects of similar size to the signal can bring rise to false alarms. Such reasoning is similar to a measure of SNR which only considers noise occupying the same frequency spectrum of the signal as anything outside it can easily be filtered. Thereby, the procedure adopts the most optimistic detection scenario, analogous to considering a maximum SNR, whereby imaging artifacts of different sizes to the imaging target are omitted.

4.5.2 Further Sampling Considerations

A point of concern for any statistical method is the required sample size. There is no simple answer as to how many cases are required to ensure a particular statistical precision of an ROC curve (Metz 1978). It is known however that the precision does increase as more cases are included and as larger AUCs are considered. Increased precision allows smaller apparent differences between resulting curves to become meaningful in order to distinguish classifier performance. Therefore, if the analysis uses only a low number of images, only broad conclusions can be drawn between the imaging algorithms. Statistical analysis of the ROC approach can be found in (Hanley & McNeil 1982). In our case here, sample size does not pose a significant concern as convergence is not necessarily required to draw relative differences between algorithm performances. It is still beneficial however to include as many independent FMC data sets as possible for each defect.

In a similar nature, curves drawn with a limited set of discrete data produces staircase effects. This can make the curve appear jagged (as we will see is the case for the experimental data in Figure 4.10 and Figure 4.11). This is investigated in (Zweig & Campbell 1993)where the ROC plots of binned and unbinned data are compared. It is shown that the ROC plot can be affected and its approximation becomes poor when the data becomes too coarsely binned. Once again, as more FMCs are included, the curves should begin to resemble their continuous equivalent.

4.6 Imaging Algorithm Evaluation

Here we present an example of the evaluation of the performance of imaging algorithms using the proposed methodology. We use the array measurements obtained from the Inconel block as described in Section 4.4.1, and perform the imaging using the TFM, PCI, and DORT MSF algorithms.

4.6.1 Results

Figure 4.10 and Figure 4.11 show the ROC results for the detection of SDHs at various depths for low and high frequency respectively (1MHz and 5MHz). The points on the ROC show the step function, whereas the lines show the optimal performance of an algorithm, in ROC terminology this is named the convex-hull (Hanley & McNeil 1982). At low frequency (Figure 4.10), all candidate algorithms detected their target at all depths but the deepest (150mm). Even though some algorithms are superior in terms of SNR (e.g. Figure 4.1 to Figure 4.3), the ROC results show their performance in this case cannot be distinguished. For the purposes of a performance evaluation, this results in a redundant outcome and more challenging detection targets are required. Higher frequencies provide this situation: the same SDHs become more difficult to detect as grain scattering increases.

At 5MHz, only the 60mm depth hole remained perfectly detected, the 150mm (Figure 4.11d) and 100mm (Figure 4.11b) depth holes were not detected at all, and the 50mm (Figure 4.11c) and 110mm (Figure 4.11a) were partially detected. When defects are not detected, the ROC curve approaches the chance line which represents a randomised classifier, shown as a dotted diagonal line in Figure 4.10a. Whilst the higher frequency results provided the desirable test case where targets are on the verge of detection, the results indicate that there is little performance difference between the algorithms. Table 4.2 summarises this by calculating the detection rates for all defects by arbitrarily setting an allowable PFA at 5% (threshold for 5% PFA shown in Figure 4.10).

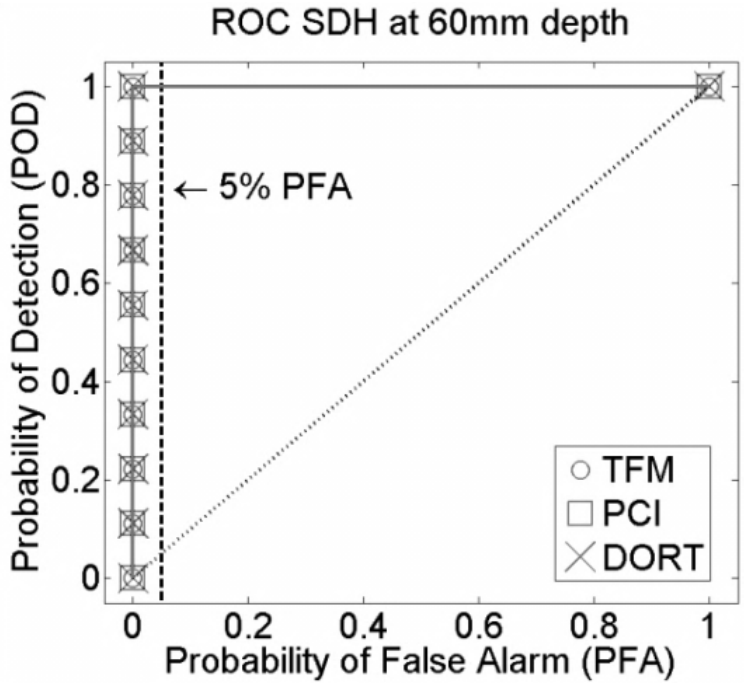

(a)

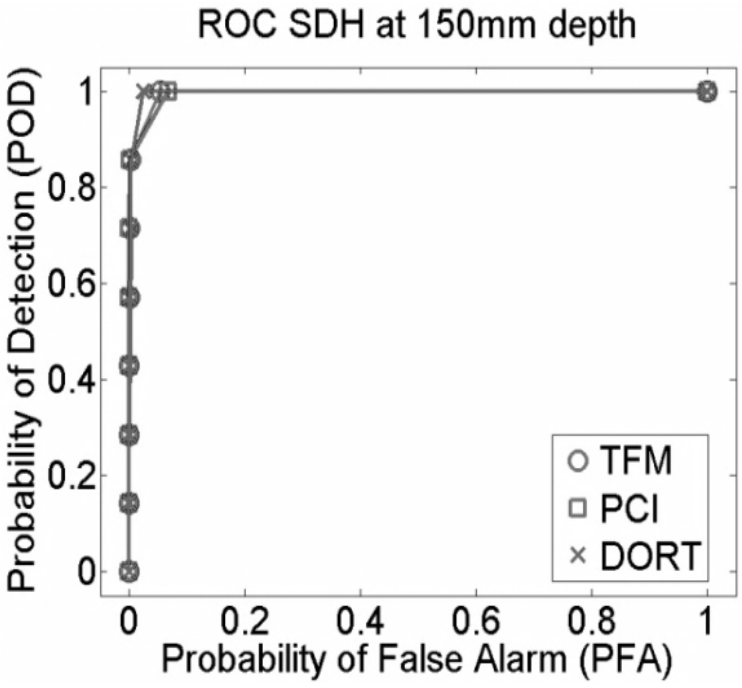

(b)

Figure 4.10: Low frequency (1MHz 64 element array) ROC results for (a) the SDH at 60mm depth which here exemplifies perfect detection (all defects are detected without raising false alarms). The diagonal dotted line represents a random classifier or chance line. The threshold for a 5% PFA is also illustrated by the dashed vertical line. (b) ROC result for 150mm depth where the detection record is no longer perfect, but the candidate algorithms produce similar results.

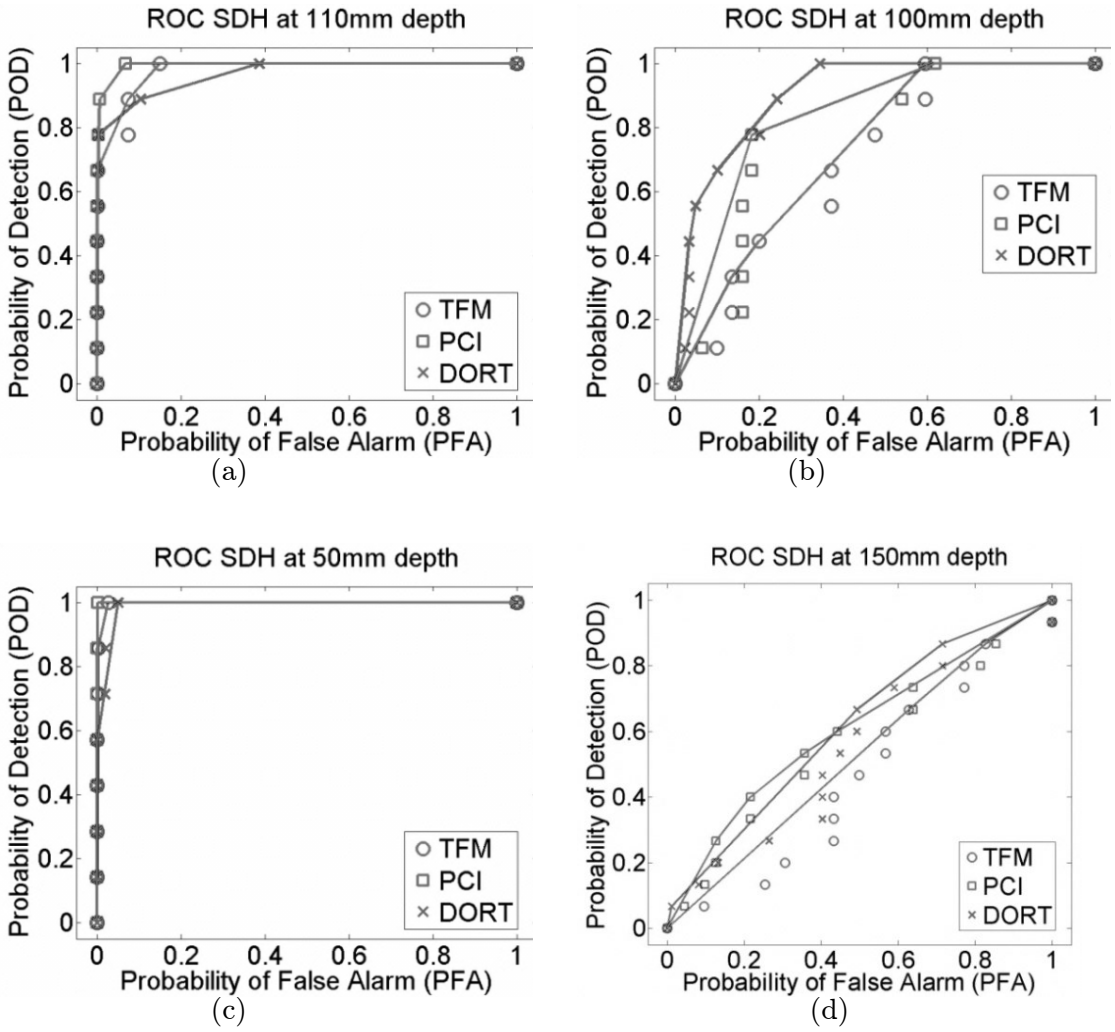

Figure 4.11: High frequency (5MHz 64 element array) ROC curves for (a) 110mm, (b) 100mm, (c) 50mm, and (d) 150mm depth SDHs. The detection rates can be seen to vary, where the defects at 100mm and 150mm were particularly difficult to detect.

Imaging Algorithm	Low Frequency	High Frequency
TFM	100%	51%
PCI	100%	55%
DORT MSF	100%	54%

Table 4.2: Summary of results: portion of defects detected by each algorithm for a set 5% PFA.

4.6.2 Discussion

The results show that even though some de-noising algorithms seem to increase SNR when flaws are detected, this had no impact on actual defect detection for this inspection setup as the considered algorithms, TFM, PCI, and DORT MSF, all performed very similarly. This observation holds true for both 1MHz and 5MHz array data, even though the majority of defects were detected at the lower frequency and therefore it was not possible to distinguish algorithms. Although both PCI and DORT MSF detected slightly more defects, this was only in the order of a few percent (see Table 4.2) when considering the defect population for a given false-call rate (5%), and hence the difference was not statistically significant.

Several possible reasons can explain the similarities in performance which were observed. Given the high levels of scattering noise, the material sample can be assumed to behave as a medium with an infinite number of scatterers, with only one of those scatterers, the SDH, being the scatterer of interest. For phase coherence to exist, and algorithms like PCI to exploit this, the scatterer of interest would have to exhibit a greater amplitude signal such that its signal and therefore also its phase dominates the signals originating from the remaining scatterers. In addition, the DORT MSF algorithm can be expected to perform better when the grains are well characterised and the threshold is implemented as is intended by its developers (Aubry & Derode 2009b; Aubry & Derode 2009a; Shahjahan, Aubry, et al. 2014). A lack of multiple scattering effects may also hinder its capacity to deliver improved imaging performance in this example.

Limitations of this particular comparison include the reliance on artificial defects. SDHs behave differently to real defects; signals received from cracks may be more phase incoherent which can cause algorithms such as PCI to behave undesirably as mentioned in

(Jie et al. 2013). Conversely, there may be characteristics of real defects that improve the performance of certain algorithms. Thus even though SDHs are an established benchmark, the conclusion for this specific case must be read with caution.

Given that the sample is representative of future power plant components, it is worthwhile to discuss some of the inspection findings. Overall, the array detection results show the sample to be difficult to inspect, regardless of which algorithm is used. The varying scattering behaviour throughout the block meant that detectability of defects did not necessarily decrease with depth, as the 50mm depth SDH exhibited a lower detection rate than one at 60mm. The sample confirmed that there is still a need to improve ultrasonic NDE of these materials before they can be reliably inspected. It is important to emphasise again that these cases are presented as examples to illustrate our comparison methodology and do not infer general conclusions about these imaging algorithms.

4.7 Conclusions

This chapter proposed a methodology to enable the detection performance of different imaging algorithms to be quantified and thereby allow comparisons of the performance of algorithms for candidate inspection applications and thus to progress ultrasonic NDE of difficult materials. By relying on ROC analysis and taking into consideration both POD and PFA, the procedure aims to deliver a robust and objective evaluation. Investigation of its limitations and the impact of the required parameters suggest its settings are relatively insensitive to the results. The only main requirement is prior knowledge of the truthful outcome for an inspection, which in a research context is often the case.

Advanced de-noising ultrasonic array imaging algorithms, PCI and DORT MSF, were matched with TFM for the detection of Side-Drilled-Holes (SDH) in a sample of coarse grained power plant material, Inconel 625. The results, applicable just to this particular case, show the algorithms to perform very similarly as no significant differences were observed in their performance. In other applications cases it may well be that there will be differences in performance, and it is proposed that this methodology would be useful to make such an assessment on arising cases of interest.

Chapter V

FINITE ELEMENT MODEL

Finite Element modelling is a promising tool for further progressing the development of ultrasonic NDE of polycrystalline materials. Yet its widespread adoption for simulating wave propagation and scattering within these materials has been held back due to a high computational cost. Within NDE, this has restricted current works to relatively small models and to two dimensions. The emergence of sufficiently powerful computing however, such as highly efficient solutions on graphics processors, is enabling a step improvement in possibilities. This chapter aims to realise those capabilities, to simulate ultrasonic scattering of longitudinal waves in an equiaxed polycrystalline material in both 2D and 3D. The modelling relies on an established Voronoi approach to randomly generate a representative grain morphology. It is shown that both 2D and 3D numerical data show good agreement across a range of scattering regimes in comparison to well-established theoretical predictions for attenuation and phase velocity. In addition, 2D parametric studies illustrate the mesh sampling requirements for two different types of mesh, to ensure modelling accuracy and present useful guidelines for future works. Modelling limitations are also shown. It is found that 2D models reduce the scattering mechanism in the Rayleigh regime in comparison to scattering in 3D.

This work is in print at the Journal for the Acoustical Society America 2015 [P7].

5.1 Introduction

Whilst modelling and simulation are becoming increasing crucial to modern day research, it can be particularly beneficial to study highly scattering materials. In comparison to experimental measurements: the interaction of multiple physics, noisy data, expensive samples, and often the need for their destructive testing, are all limitations which can be avoided through simulation. Perhaps unsurprisingly, there has been significant interest and efforts to simulate ultrasonic signals obtained from grain scattering media (see Chapter 2) - but it has proven challenging. Mathematical complexity: i.e. there are Rayleigh, Stochastic, and Geometric scattering mechanisms, and in addition the computational expense required to incorporate all effects such as multiple scattering, have both limited the benefit of modelling techniques.

Even so there has been some success, for example, models such as the Unified Theory (see Chapter 2) have proven particularly useful to characterise polycrystalline materials by inversion from attenuation measurements (Stanke 1985; Papadakis 1968). As previously mentioned however, for ultrasonic flaw detection, the scattering induced noise is also of interest (Thompson et al. 2008). To that end, the currently most complete model is the Independent Scattering Model (ISM) (Margetan et al. 1994) which has significantly benefited ultrasonic inspections (Margetan et al. 2007) today. However, the ISM neglects multiple scattering which thus limits its applicability to relatively weak scattering media (Anxiang et al. 2003). Other models, such as those mentioned in Chapter 4, whilst very useful, are more simplistic and can be regarded heuristic.

Recently however, researchers (Ghoshal & Turner 2009; Feuilly et al. 2009; Shahjahan, Rupin, et al. 2014; Lan et al. 2014) have considered numerical modelling such as the Finite Element (FE) method to overcome this limitation and confront more challenging scattering scenarios. In contrast to existing approaches, its ability to simulate time-domain signals, incorporating both attenuation and noise, whilst also including complex physics such as multiple scattering (Shahjahan, Rupin, et al. 2014), makes FE a promising candidate. Its flexibility and high fidelity will probably be instrumental to further progressing the development of ultrasonic NDE of polycrystalline materials.

Yet its widespread adoption has been held back due to a high computational cost which arises from having to numerically discretise the material's microstructure. This has restricted current works to relatively small models e.g. of the order of 1000s of grains, which, while representing impressive progress is still only sufficient for a reduced range of feasible scattering regimes, and to 2 dimensions. The latter limitation, a 2D model, obliges several simplifications including:

(1) The representation of grain size distributions of a 3D material in a 2D model. Namely, the grain cross-sections seen on a slice of a 2D material do not correctly represent the grain sizes of a 3D material.

(2) The stiffness matrix, which is reduced according to plane strain assumptions and renders the model infinite in the collapsed dimension.

(3) The scattering phenomena, which are not fully reproduced. For example, Rayleigh scattering is a 3D phenomenon which is closely linked to the scattering cross-section which is proportional to volume and therefore reduced in 2D environments where the scattering can only occur in the two dimensions.

This chapter presents recent developments of realistically large and detailed FE models of ultrasonic longitudinal wave propagation within polycrystalline materials, demonstrating and evaluating new simulation possibilities in 2D and 3D. It investigates the capability of FE to model the different scattering behaviours across regimes as predicted by the Unified Theory, and assesses the significance of 2D assumptions through comparison with 3D simulations.

This advanced modelling is now becoming possible because of the emergence of sufficiently powerful computing and new, faster modelling tools. Specifically, we make use of a highly efficient GPU based solver (Huthwaite 2014) for FE which has enabled larger studies e.g. up to 100,000 grains in 2D and 5000 in 3D. Although this approach can be suited to model a variety of microstructures, for this initial investigation, we consider a relatively simple microstructure, untextured, and comprising equiaxed grains of a single phase in a range between 100μm and 500μm. The chosen material is a relatively strong scattering medium, Inconel 600, of cubic symmetry.

As a further example of the utility of modelling such as this, recent research (Zhang et al. 2004; Zeng et al. 2010; Maurel et al. 2006) has raised interesting queries regarding our current understanding of grain scattering, including the role of grains as Rayleigh scatterers (Zhang et al. 2004) and whether it is not the material imperfections such as voids and inclusions which are contributing to that effect. FE can be useful in this matter by modelling a perfect polycrystalline microstructure, clear of flaws, and identifying the dominant scattering behaviour of the grains.

The subsequent section provides a brief step-by-step outline for FE modelling of polycrystalline materials in 2D, continued by Section 5.3 which investigates its mesh sampling requirements. Section 5.4 introduces the 3D model. The main body of results is presented in Section 5.5 where numerical simulations of a 2D and 3D model are compared to theoretical results obtained from the Unified Theory.

5.2 FE Modelling of Polycrystalline Material

Finite Element modelling of polycrystalline materials has been successfully undertaken in various fields of research (Quey et al. 2011; Kamaya 2009; Zhang et al. 2011) including NDE (Ghoshal & Turner 2009; Feuilly 2009; Chassignole et al. 2009; Shahjahan, Rupin, et al. 2014; Lan et al. 2014) where it has been limited to 2D. Although several approaches have been adopted, all of those mentioned here that consider geometrically varying grains, rely on Voronoi tessellations (Aurenhammer 1991) to numerically generate a morphology which is geometrically similar to a naturally occurring polycrystalline microstructure. This has been accepted as a good approach by researchers in crystallography and textured materials (Kocks et al. 2000). Yet for wave propagation, several questions still remain: for example whether it can accurately recreate the complex scattering phenomena, and what type and density of mesh, one consisting of regularly shaped elements (structured), or one which allows for arbitrary element shapes (free), is best to use to deliver efficient but accurate results.

This procedure was implemented, relying mainly on Matlab, although some bottlenecks were sped up using C++, and eventually Bash was used to coordinate running the models in batches. The Matlab portion involved approximately 10,000 lines of code which

eventually allowed for several model configurations, such as a plane wave mode, point sources, and adding/removing absorbing regions, and also various meshes including triangular or rectangular, or structured and unstructured. The next sub-sections provide a brief step-by-step description, and practical considerations for the aforementioned modelling approach.

5.2.1 Generating Random Polycrystals

Generating a random polycrystalline microstructure, as achieved in (Ghoshal & Turner 2009; Feuilly 2009; Chassignole et al. 2009; Shahjahan et al. 2014), starts by randomly distributing points, or seeds, in a Euclidian space. Several options are available to do this, here we adopt the approach of starting with a regular grid of seeds, and perturbing their positions independently for each dimension according to a Gaussian distribution. The original grid spacing determines the mean grain size, and the spread of the perturbation determines the grain size distribution. A 2D example of this is shown in Figure 5.1a, where the seed density will determine the resulting average grain size. The coordinates of each seed become the site for a single grain by serving as an input to the Voronoi algorithm (Aurenhammer 1991). The algorithm subdivides the original space into regions, in the form of convex polygons, whereby each polygon encloses the area which is nearest to that particular seed (see Figure 5.1b). Once a Voronoi tessellation has been generated, it requires modification to make it suitable for FE modelling. This procedure involves clipping the boundaries for instance, a step previously described as regularization (Quey et al. 2011).

Clipping to eliminate lines which are outside or partially outside the model space, is the only regularisation step required for structured meshes. This procedure becomes particularly cumbersome however (but essential) for unstructured meshes. This arises from the necessity to remove small lines and extreme angles composing the polygons, which would introduce undesirably shaped elements. Two schemes to regularise the mesh were considered during this project: (1) Undesired lines are identified and subsequently removed, leaving a gap in the tessellation, which then requires stitching. Whilst this option may be the least disruptive, it becomes increasingly complex to account for the enormous number of exceptions that can occur within a random tessellation of polygons.

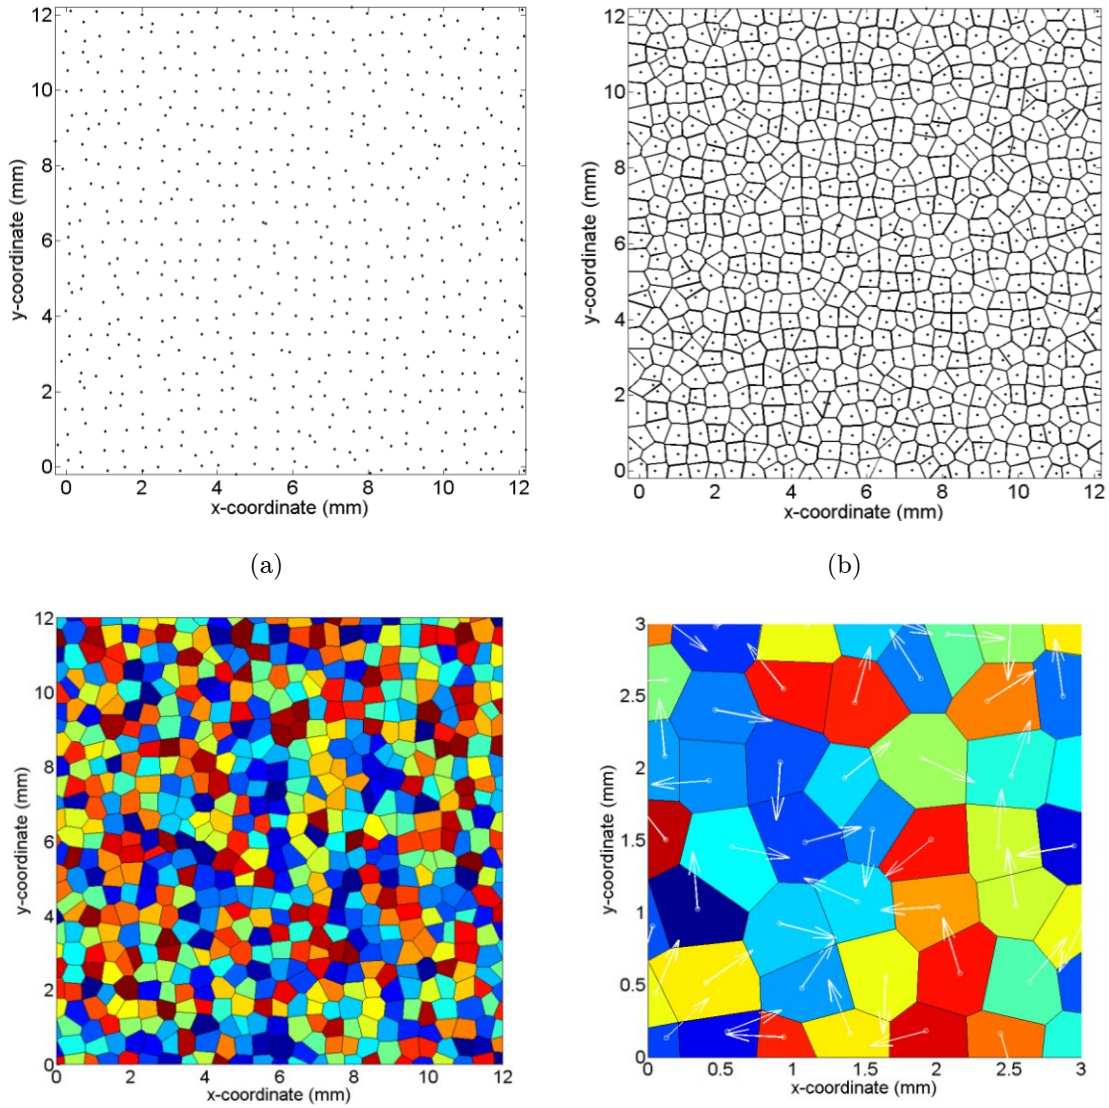

Figure 5.1: Illustration of the steps involving a Voronoi generation of polycrystals: (a) a random distribution of seeds; (b) the Voronoi tessellations produced by (a); (c) The regularized grain layout and (d) the random orientations assigned to each grain, here shown by arrows in the 2D plane for clarity (note zoomed scale of this image compared to the others). Colors are only illustrative.

For example, small lines occurring within the body or at the edges, intersections of multiple lines, or whether the alteration created new Voronoi-nodes, are all instances which need to be considered. Furthermore, during these adjustments, it is important to maintain the information of all vertices belonging to each grain, which is in part hampered by floating point errors. Although this approach was fully implemented and used for some time, ultimately an alternative approach was selected which was superior. Scheme (2) approximates the tessellation by rounding the number precision, such that all its vertices fall onto a fine grid, whose refinement is defined. This is similar to how a structured mesh would approximate the tesselation, except we can introduce a much finer grid to approximate it to. Although being more disruptive than the previous method, since this doesn't destroy the Voronoi-node database, and doesn't introduce quite the same level of staircasing effects as the structured mesh does, it is considered the more efficient scheme.

5.2.2 Mesh Generation

The minimum FE mesh discretisation for accurate modelling of wave propagation is usually constrained by the wavelength (Drozdz 2008). In this case however, whether using a structured or unstructured mesh, the objects to model, the crystallites, are often an order of magnitude smaller than the wavelength, which demands denser meshes whose refinement far exceed the said wavelength criteria. Two possibilities exist, which have previously each been adopted, either an unstructured mesh utilizing triangular FE elements (see Figure 5.2a and Figure 5.3a) to conform to the complex boundaries of the Voronoi tessellation, or an approximation of the grains with a structured mesh (Shahjahan, Rupin, et al. 2014) (see Figure 5.2b and Figure 5.3b). The hazard with a structured mesh is that it leads to " staircasing" effects (Drozdz 2008) which become a poor approximation at coarse mesh densities (see Figure 5.2d) and can lead to tip diffraction from edges, and also to disproportionately strong reflections from waves that are normally incident to the plane of the flats (Drozdz 2008). When using an unstructured mesh however, the challenge is to maintain high quality triangles, i.e. close to equilateral shapes, such that there is minimal mesh scattering.

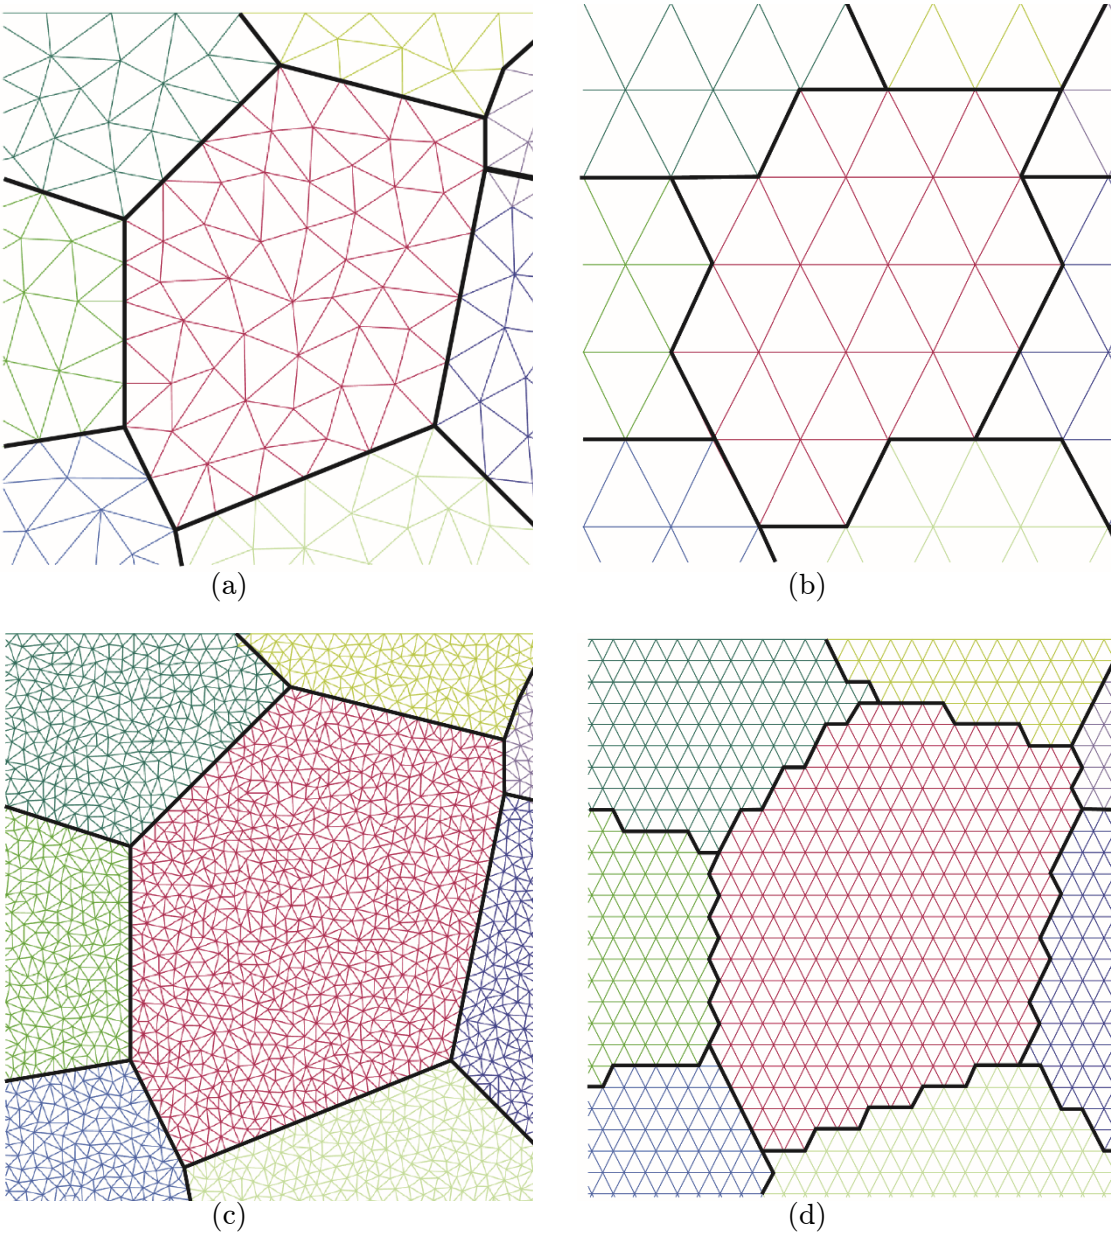

Figure 5.2: Meshed polycrystal using four different schemes: (a) free mesh (b) structured mesh (c) finer free mesh (d) finer structured mesh. Staircasing effects are illustrated in (b) and (d).

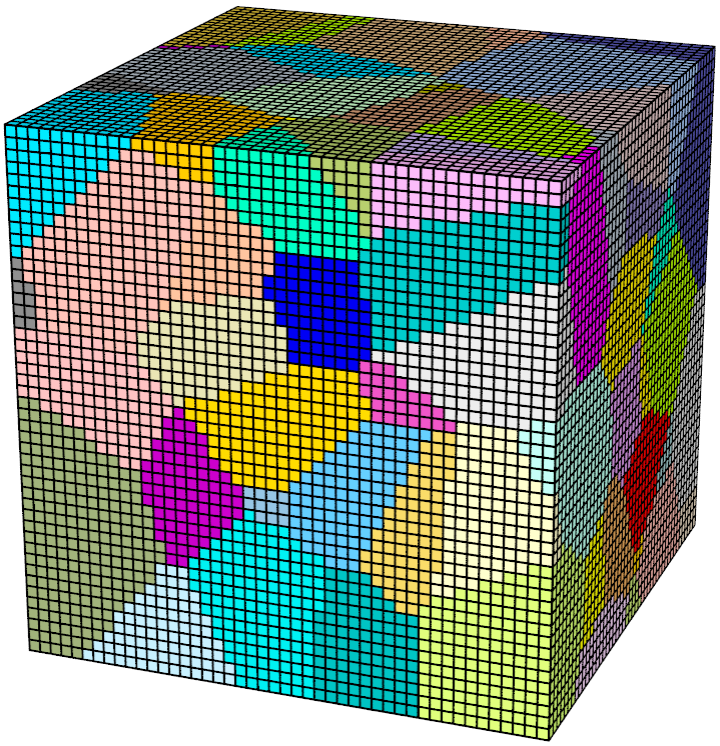

(a)

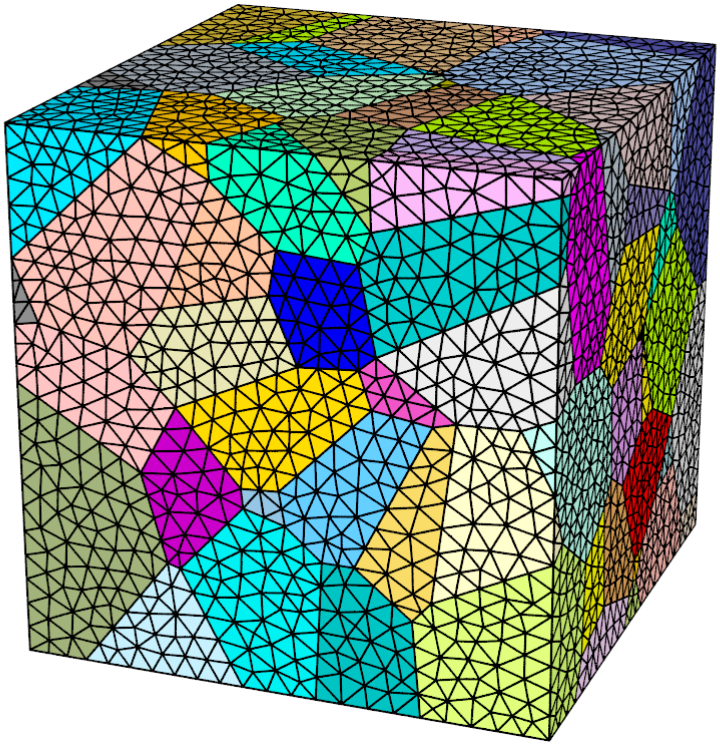

(b)

Figure 5.3:Typical grain mesh using (a) unstructured and (b) structured mesh for a 3D scenario.

For this purpose, several software solutions are available; for example good results were found, both in terms of the quality of meshes (no large deviations from equilateral, no large variations in element sizes) and the time required to generate them (e.g. order of minutes), using a Free software tool, Triangle (Shewchuk 2002). Using commercial software, such as Abaqus (Abaqus 6.14 2014), it was found that meshing can be a significant bottleneck to this type of modelling (e.g. 20 hours computation time), with the only solution to speed things up being to mesh several models in parallel.

5.2.3 Dimensional Considerations

In contrast to sub-3D modelling, fewer simplifications are necessary when all three dimensions are used. When 3D models are not feasible, reducing a polycrystalline material to a 2D model introduces certain simplifications. This includes the grain size distribution, which impacts, amongst other properties, the ultrasonic characteristics of the material. Whereas for 3D modelling approaches, the simple approach is to match the distribution of grain dimensions to that of the desired material, in 2D, this is not as trivial. Namely, a random cutting plane through a 3D tessellation of grains will not intersect every grain through its centre, rather, some intersections will occur off-centre and therefore reproduce smaller cross-sections. The study of interpreting 2D representations of 3D grains forms the basis of stereology (Underwood 1970) and is beyond the scope of this study. Here, we will assume a normal distribution of grain sizes in 2D (defined as the square root of area), as the one depicted in Figure 5.4, which assumes that our slice of a 3D material cuts every grain through its centre and therefore overestimates the grain sizes that would be seen in a proper 2D section. Whilst larger grains will increase the attenuation, we are making no claims regarding how this may compare to attenuation of a 3D material. Namely, it would be interesting as a future exercise to further investigate the opportunities and advantages of adjusting grain size distributions in 2D to better match the ultrasonic behaviour of a 3D material; this would be important for rigorous modelling in 2D and is by no means straightforward to achieve.

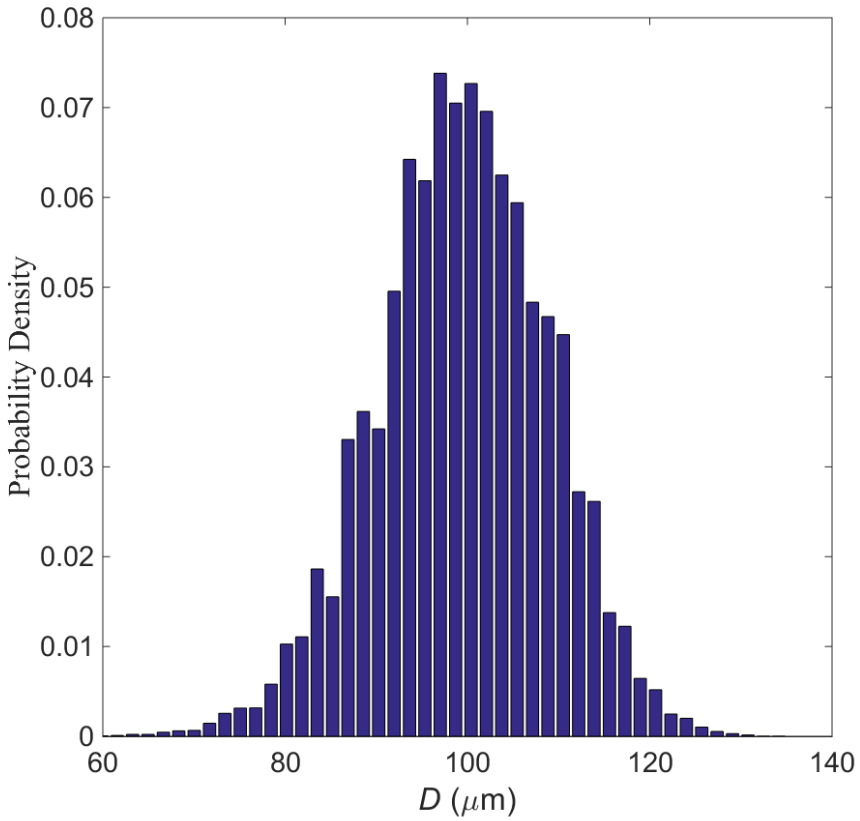

Figure 5.4: Grain size distribution for a typical random realistation of an input 100µm grain size material. The grain size D in this 2D case is defined by the square root of area.

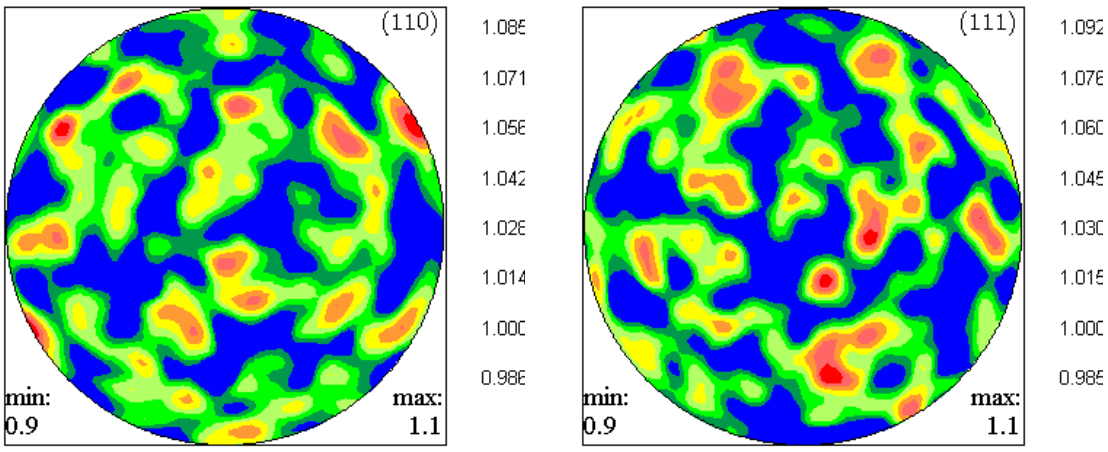

Figure 5.5: Typical pole plot (ODFs) for a randomly generated material. The distribution of grain alignments over the whole sphere shows for this example that the generated material is indeed isotropic. The scales indicate the distribution of probability density for the orientation angles of the <110> and <111> crystallographic axis.

The orientation distribution function (ODF) of a polycrystalline material is another factor which determines macroscopic properties. For a single phase material, each crystallite should be assigned the same anisotropic stiffness properties but with a random crystallographic orientation to define a macroscopically isotropic material (see Figure 5.1d). To achieve this, the three reference Euler angles, which define orientation, may be randomly distributed such that their orientations lie equally spaced on the surface of a sphere, as explained by (Shahjahan, Rupin, et al. 2014) for example. Figure 5.5 shows the result of rotating orientation angles in 3D for 2000 grains, illustrated by pole plots (the established form of display for ODFs in materials science (Lan et al. 2014)). As can be seen, as desired, a macroscopically isotropic material has been achieved.

Finally, the 2D simplification leaves two possibilities, one in which the orientation distribution is only in the plane, hence a plane strain model is possible; the other in which the distribution is in 3D, and then needs to be approximated for 2D. Plane strain assumptions then neglect the stiffness constants associated with the third dimension, when reducing the stiffness matrix from 3D to 2D. The latter is adopted here for the 2D modelling.

5.2.4 Efficient Simulations using GPU

Due to the increased mesh density, FE modelling of polycrystalline microstructure is computationally expensive. To reduce this cost and thereby enable parametric studies, the work here employs a relatively new FE solver, Pogo (Huthwaite 2014). Pogo exploits the sparsity and highly parallelizable nature of the time explicit FE method, which allows the very efficient use of graphical processing units (GPUs) instead of conventional computer processing units (CPUs) to execute the computations in parallel. It has been shown that this can result in speed improvements of up to two orders of magnitude (Huthwaite 2014) when compared to commercially established CPU equivalent software. For example, timing of a typical simulation undertaken in this chapter, when running a 6.1×10^6 degrees of freedom model, was measured to be 67 times faster using a GPU setup based on 4x Nvidia GTX Titan graphics cards when compared to 2x Intel Xeon 8-core E5-2690 2.9GHz CPUs using general purpose CPU software.

5.3 Mesh Validation Study

Here we investigate the spatial sampling requirements for both types of mesh mentioned in Section 2.3 to guarantee sufficient modelling accuracy whilst also preserving computational cost. In order to achieve this, both the mesh scattering (Section 3.1) and mesh convergence (Section 3.2) are evaluated for a plane wave model. The computational cost for such multiple runs however is far too great in 3D, and therefore no parametric studies are feasible. Instead, the knowledge gained from the 2D mesh studies, regarding the mesh requirements, will be used later to create a 3D model.

The studies in the following Sections 3.1 and 3.2 rely on three different realisations of a polycrystalline material, Inconel 600, using the material properties taken from (Shahjahan, Rupin, et al. 2014) and shown in Table 1. Each model consists of a different average grain size: 100μm, 250 μm and 500 μm. As computational costs increase for finer grains, this range was limited to keep costs manageable whilst also representing a range of grain sizes of interest to NDE.

Material Property	Inconel 600
C_{11}	234.6 GPa
C_{12}	145.4 GPa
C_{44}	126.2 GPa
ρ	8260 kg/m^3

Table 5.1: Material constants for cubic Inconel 600 (Shahjahan, Rupin, et al. 2014).

Figure 5.6 shows an example simulation by one of the models used in the study. It is a coarse-grained material represented in 2D by a strip 42mm long and 12mm wide in plane strain. A 3-cycle-toneburst with a 2MHz centre-frequency is applied to the line of nodes, at the left side where x=0mm, which forms the excitation line-source. The model uses symmetry boundary conditions at the top and bottom edges (where y=0mm and y=12mm in Figure 5.6) such that the nodes are constrained in the y-direction. This creates a plane wave which can be seen to propagate in the positive x-direction. The backscatter can be recognised from the random fluctuations in amplitude trailing the plane wave.

5.3.1 Mesh Scattering

Successful simulation of grain scattering can only be achieved if the scattering from element boundaries, here termed mesh scattering, is significantly less than the grain scattering itself. Mesh scattering arises from heterogeneity introduced by irregular element shapes, such as those encountered in unstructured meshes, and can be reduced by increasing mesh density at the cost of additional computation. In order to assess this, some unstructured mesh models are run for which the grain noise is eliminated, so that the noise is solely due to mesh scattering. This is achieved by assigning isotropic stiffness properties to the grains in this part of the study.

To quantitatively compare results for different mesh densities, the mesh noise is represented by considering the average backscatter energy received by the individual excitation nodes. This is calculated from both the temporally and spatially averaged intensity i.e. the root-mean-square (RMS) value of the time-domain backscatter received at the different nodal positions, denoted by S_{rms}. The signal is windowed such that it corresponds to a time after the excitation signal and before the arrival of the reflected signal, which represents a time window where the received energy, in absence of mesh scattering, is anticipated to be zero. For clarity, this is analogous to analysing a time window in-between the frontwall and backwall of a typical pulse-echo time trace encountered in ultrasonic NDE. A worthwhile remark here is that the noise is combined such that it corresponds to the backscatter seen by spatially infinitesimal receivers, whereas in more practical simulations, the mean displacement response across multiple nodes may be considered, which then suppresses the noise level by spatial averaging. Thus this is a relatively harsh case to present, but nevertheless allows useful comparisons.

Figure 5.7 plots the mean mesh scattering noise (in dB, with reference to the peak of the excitation signal), S_{rms} as a function of the mean element edge divided by the wavelength at centre frequency, $e\lambda^{-1}$, or elements per wavelength. As expected, the mesh scattering decreases as the mesh becomes more refined. In general, the mesh scattering is very low (i.e. all results here are below -40dB) for the range of investigated mesh densities. As the mesh is refined, the unstructured mesh results seem independent of the grain size used once an initial threshold is achieved.

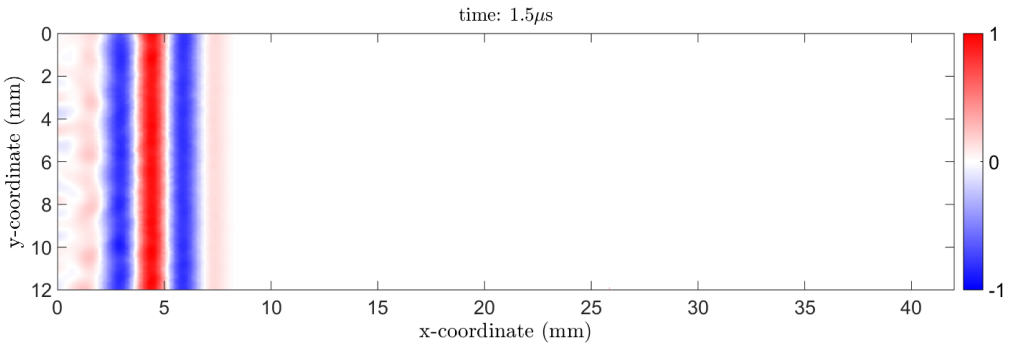

(a)

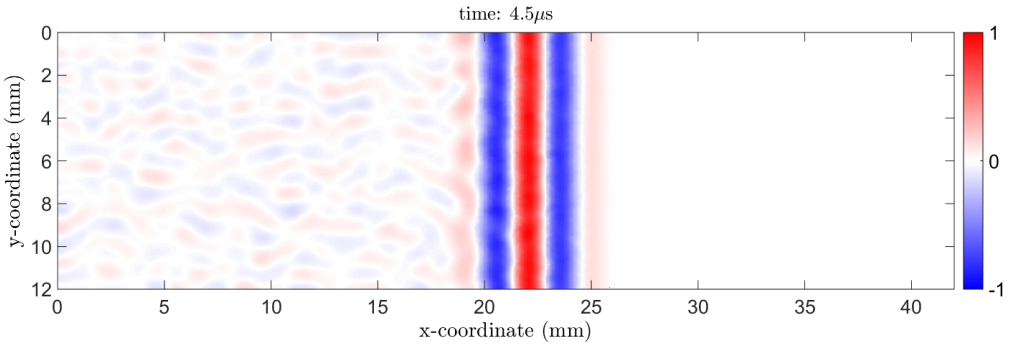

(b)

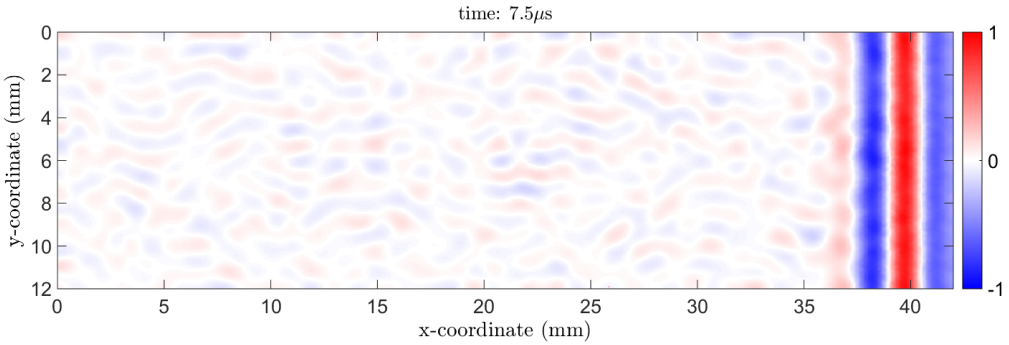

(b)

Figure 5.6: FE simulation of longitudinal plane wave propagating from left to right within a 2D slab of polycrystalline Inconel for different times after (a) 1.5µs (b) 4.5µs and (c) 7.5µs. The colour scale is the normalised displacement amplitude with reference to the peak excitation amplitude from -100% to 100%.

It is important to acknowledge that these results do not provide an all-encompassing criterion for mesh refinement. The refinement requirement will be model-specific and depend on the severity of the grain noise and on practical compromises on model size. It is crucial however, to suppress it to a controlled degree and this simple approach allows any candidate case to be evaluated.

Structured meshes, which exhibit no variation in element shape, do not require the above considerations and hence clearly outperform unstructured meshes according to this criterion. As they do not conform to the grain boundaries however, it is yet unclear whether they can correctly model the scattering behaviour, which is addressed in the next section.

5.3.2 Mesh Convergence

It is also important to achieve adequate convergence of the propagating wave pulse. The same models are used as in the previous section, namely with three different grain sizes, except that the anisotropic properties of the grains are now introduced (as described in Section 2.2) and thus the wave will be affected by grain scattering. Two metrics are employed to measure convergence, the centre-frequency attenuation, and the group velocity.

As a measure of the propagating wave, the received displacements are now spatially averaged across all the nodes which lie on the right side edge where x=42mm in Figure 5.6, emulating a pitch-catch plane-wave configuration. The centre-frequency attenuation convergence is calculated as a difference in amplitude between the peak of the received time-domain Hilbert envelope A and that of the converged solution A_c. The converged solution, A_c, is obtained from the highest available density mesh. Similarly, the measured group velocity V_g, which is calculated from the time-of-flight, as measured from the Hilbert envelope peak, is subtracted from the converged solution V_c. To clarify, an error of 0.05 would correspond to a 5% difference in group velocity from that of the converged solution.

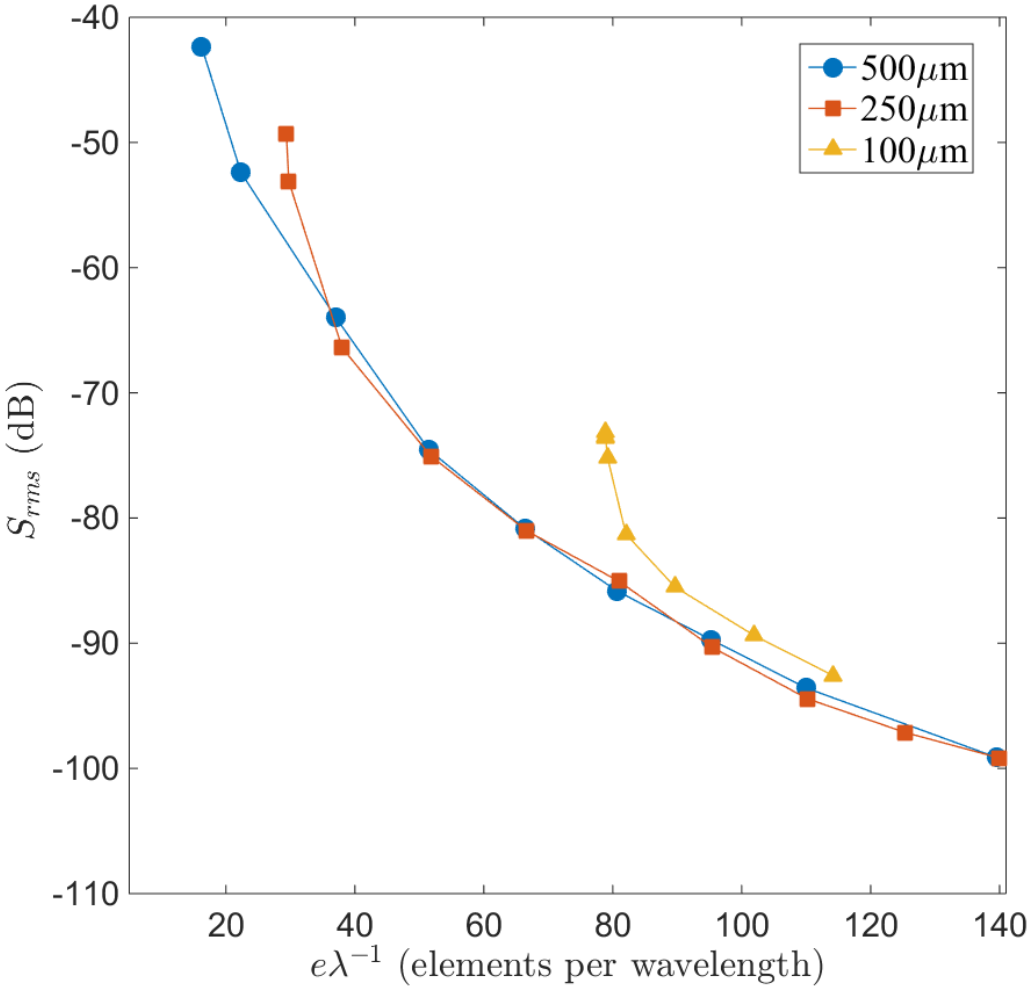

Figure 5.7: Mean normalised mesh scattering noise (in dB, with reference to the peak of the excitation signal) versus number of elements per wavelength for several unstructured meshes, each conforming to polycrystalline material with a different average grain sizes.

Figure 5.8 and Figure 5.9 plot, as a function of the mean element edge length e per mean grain size d, the convergence of the centre-frequency attenuation and group velocity respectively, for three different grain sizes, using a structured (S) and unstructured mesh (F). As can be seen, both attenuation and velocity converge as mesh density is increased, and velocity converges quickest. At ten elements per linear grain dimension, both metrics are converged to within 1% error for all grain sizes considered which agrees with the findings of (Shahjahan 2013) for another type of mesh, a rectangular structured mesh.

The progress of convergence reveals that both meshes converge at a similar rate, although the structured mesh seems to converge more monotonically. In the case for an unstructured mesh, the element size distribution can vary by several orders of magnitude within a single model which results in time stepping disadvantages in comparison to structured meshes. This is due to the need to satisfy the critical time step (Bathe 1996) throughout the model, defined by the smallest element length e_{min} in the model, which may cause oversampling for other elements which are larger, increasing their chance of accumulating numerical noise.

The results for the different grain sizes are somewhat unexpected, namely, the 100µm model seems to converge at a lower mesh density in comparison to the coarser grains. This can be explained however by a lower grain scattering induced attenuation for the grain size model of 100µm (which has a larger wavelength to grain size ratio), and hence at coarse mesh densities, the mesh scattering, in that specific case, introduces similar levels of attenuation. It can also be noted that convergence for the 500µm grain model initiates with a relatively small error which increases before eventually converging again. Comparing the results for both figures shows however, that at the lowest mesh density, the received signal peak-amplitude may be within 2% of its converged solution, the group velocity error remains unconverged and at a maximum (see Figure 5.9). The total attenuation is caused by both mesh and grain scattering (for reference, the mesh scattering induced attenuation will typically be in the order of a few percent for the models simulated here, whereas the grain scattering induced attenuation is typically an order of magnitude larger), the latter is governed by differences in velocity by adjacent grains. At very low mesh density, the velocity error is large, and the low attenuation we see here may be a fortuitous result due

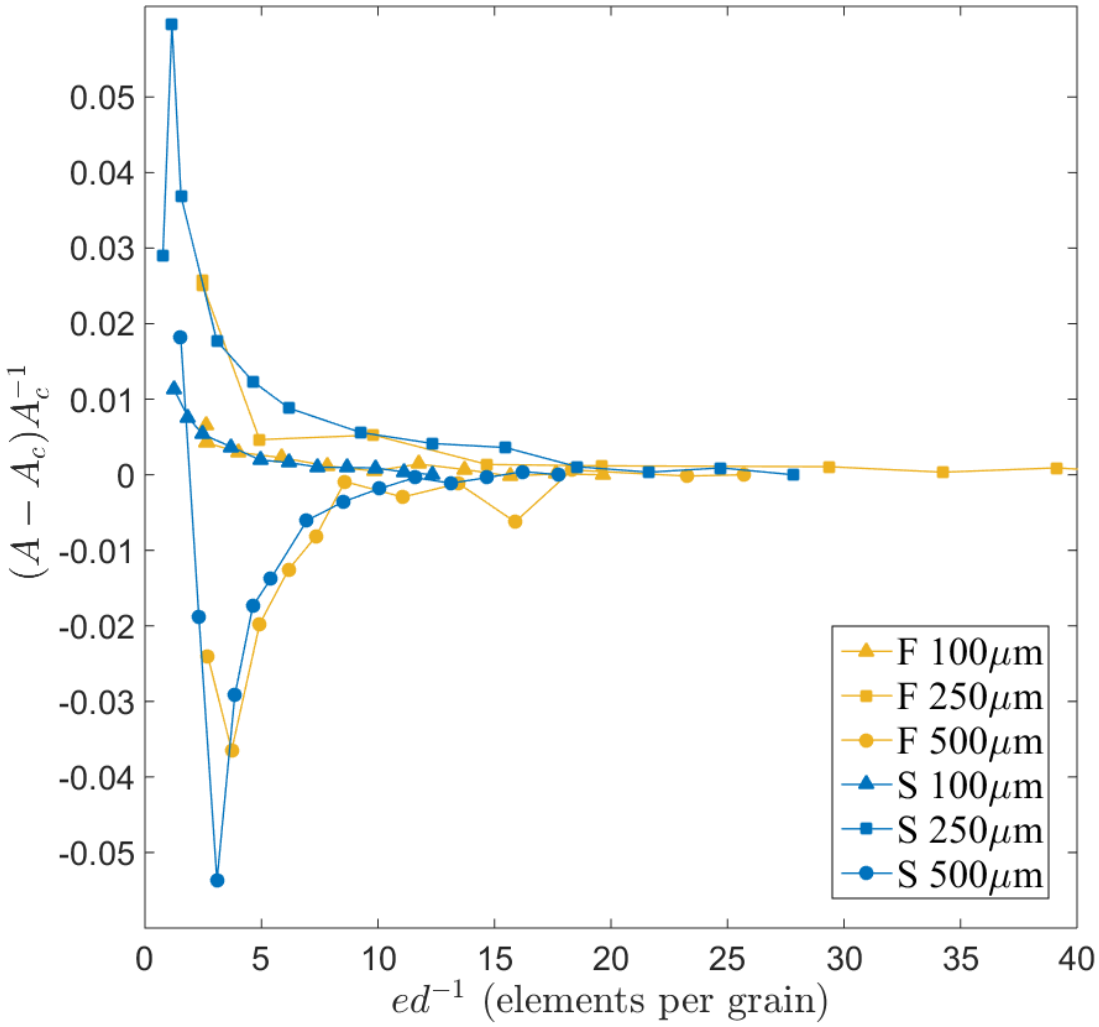

Figure 5.8: Convergence of normalised centre-frequency attenuation against elements per grain for structured (S) and unstructured meshes (F). Results are shown for three different grain size models, 100µm (triangular maker), 250µm (rectangular marker), 500µm (circular marker). The centre-frequency attenuation can be seen to converge within 1% at approximately 10 elements per grain.

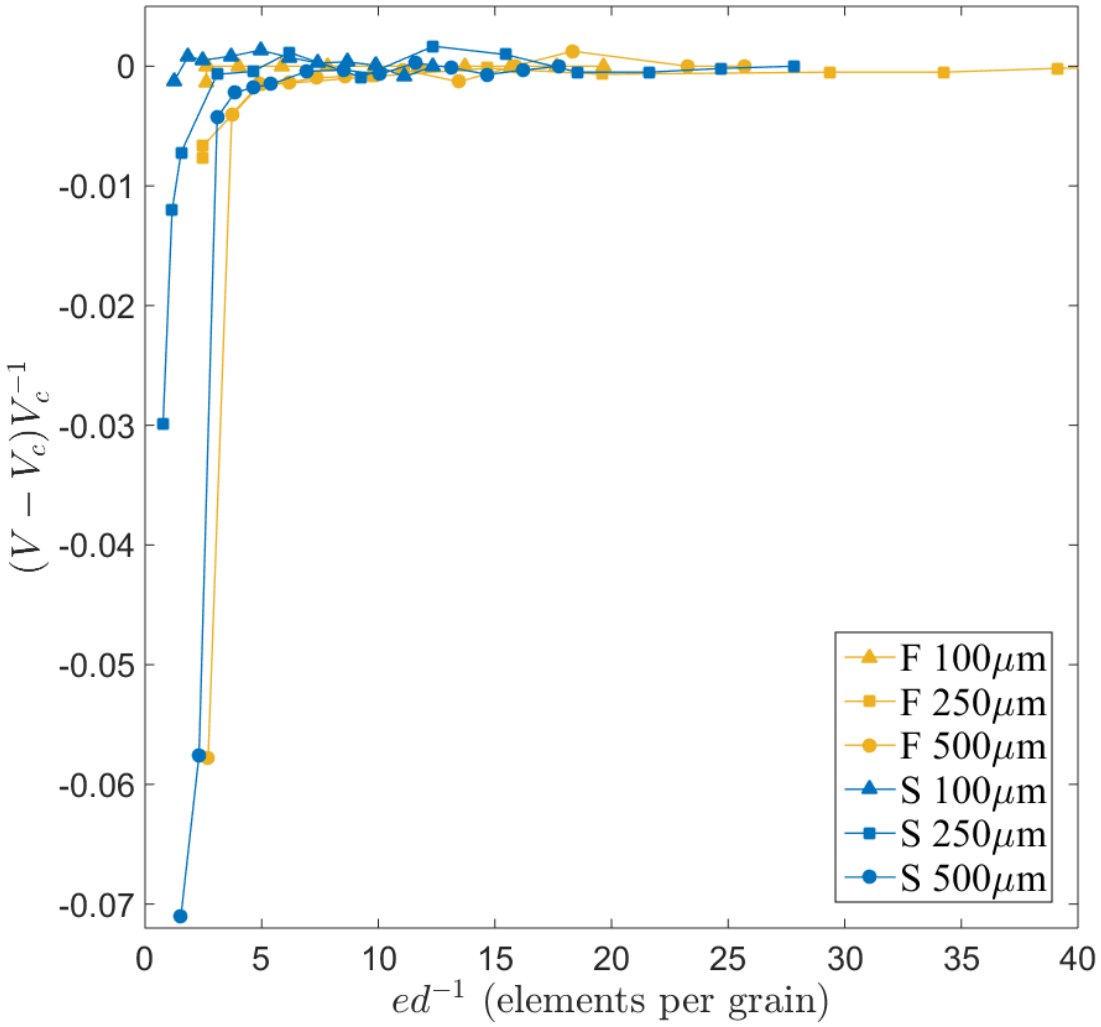

Figure 5.9: Normalised group velocity convergence against the number of elements per grain for structured (S) and unstructured meshes (F). Results are shown for three different grain sizes, 100µm (triangular marker), 250µm (rectangular marker), 500µm (circular marker). Both meshes can be seen to converge to within 1% at approximately 6 elements per grain dimension.

to an artificially increased mesh scattering and reduced grain scattering. In any case, it is clear that we need both velocity and attenuation to be converged for a useful solution.

This work refrains from advocating a particular choice of mesh, instead it has been shown that both types are viable options for modelling a polycrystalline microstructure and offer similar performance i.e. offer similar accuracy for the same computational cost. Therefore, the choice of which to use will be largely determined by the particular modelling application, which is also why within the modelling community today, both unstructured and structured meshes are in use. However, for the relatively simple models which will be considered in the subsequent sections, unstructured meshes add unnecessary complications, and hence we have selected structured meshes on this occasion.

5.4 Result Validation

The numerical results are evaluated for 2D and 3D FE models, adopting structured meshes on this occasion, and comparing their results with expectations from theory. Similarly to the mesh convergence study, both attenuation and velocity are measured, except that now both the attenuation and the phase velocity are evaluated as functions of frequency.

The theoretical values were obtained by computing the complex longitudinal propagation constant as defined by the Unified Theory (Stanke & Kino 1984) using the material properties outlined in Table 1. Our implementation of the Unified Theory calculation was validated by reproducing both results (the attenuation and phase velocity plots) for another cubic polycrystalline material, iron, presented in the original paper (Stanke & Kino 1984).

In 2D, six FE models are considered, three for each grain size, 100μm and 500μm, and each excited by a different centre-frequency 3-cycle-toneburst. The range of frequencies applied (see Table 2) are believed to represent a good range of interest, and were limited by increases in computation costs for cases outside this range.

Model	2D D=100μm			2D D=500μm			3D D=500μm
Centre Frequencies	2MHz	3MHz	5MHz	1MHz	2MHz	3MHz	1,2,and 3MHz
Number of grains	60 x10^3	100 x10^3	100 x10^3	30 x10^3	23 x10^3	25 x10^3	5 x10^3
Length (mm)	100	50	50	150	75	25	40
Width (mm)	6	20	20	50	75	250	4x4
Degrees of Freedom	12 x10^6	20 x10^6	31 x10^6	5 x10^6	8 x10^6	8 x10^6	16x10^6

Table 5.2: Parameters for three models with different grain sizes, 100 μm and 500μm for two 2D models, and 500 μm for a 3D model.

The 3D model created here measures 4x4x40mm and counts 5210 randomly orientated Inconel grains with an average grain size of 500μm. For a closer view, only a slice of the full model is shown in Figure 5.10 which was created using Neper (Quey et al. 2011). Similarly to the 2D model, a plane wave is created by imposing symmetric boundary conditions on the rectangular plane surfaces of the model and applying a 3-cycle tone burst to the nodes which lie on the end-surface, seen as a square plane surface at the end of the picture in Figure 5.11. The key statistics of the model are shown in Table 5.2. Once the model is solved, post-processing involves calculating the mean nodal displacement of the nodes which lie on the end-face opposite to the excitation plane, thereby emulating a pitch-catch configuration.

The single 3D FE model, detailed above, is solved for various centre frequency excitations in the range of 1-3MHz. Both 2D and 3D model parameters are detailed in Table 5.2. To enable comparisons to theoretical results which provide results for a mean field, analogous to an infinite plane wave, the dimensions of each FE model are adjusted to ensure sufficient spatial averaging of the received displacements and reduce the effect of phase aberrations and noise. This is a demand which grows with frequency and grain size, thereby increasing computation costs, and therefore this defined the frequency range of interest for this chapter. Similarly, although multiple realisations would ideally be considered to gather more statistics, only this one realisation is considered here. Nevertheless, since a large number of grains are considered in each realisation, significant variations in the determined quantities would be unexpected, such as the attenuation and phase velocity, if we were to repeat the calculation for further realisations.

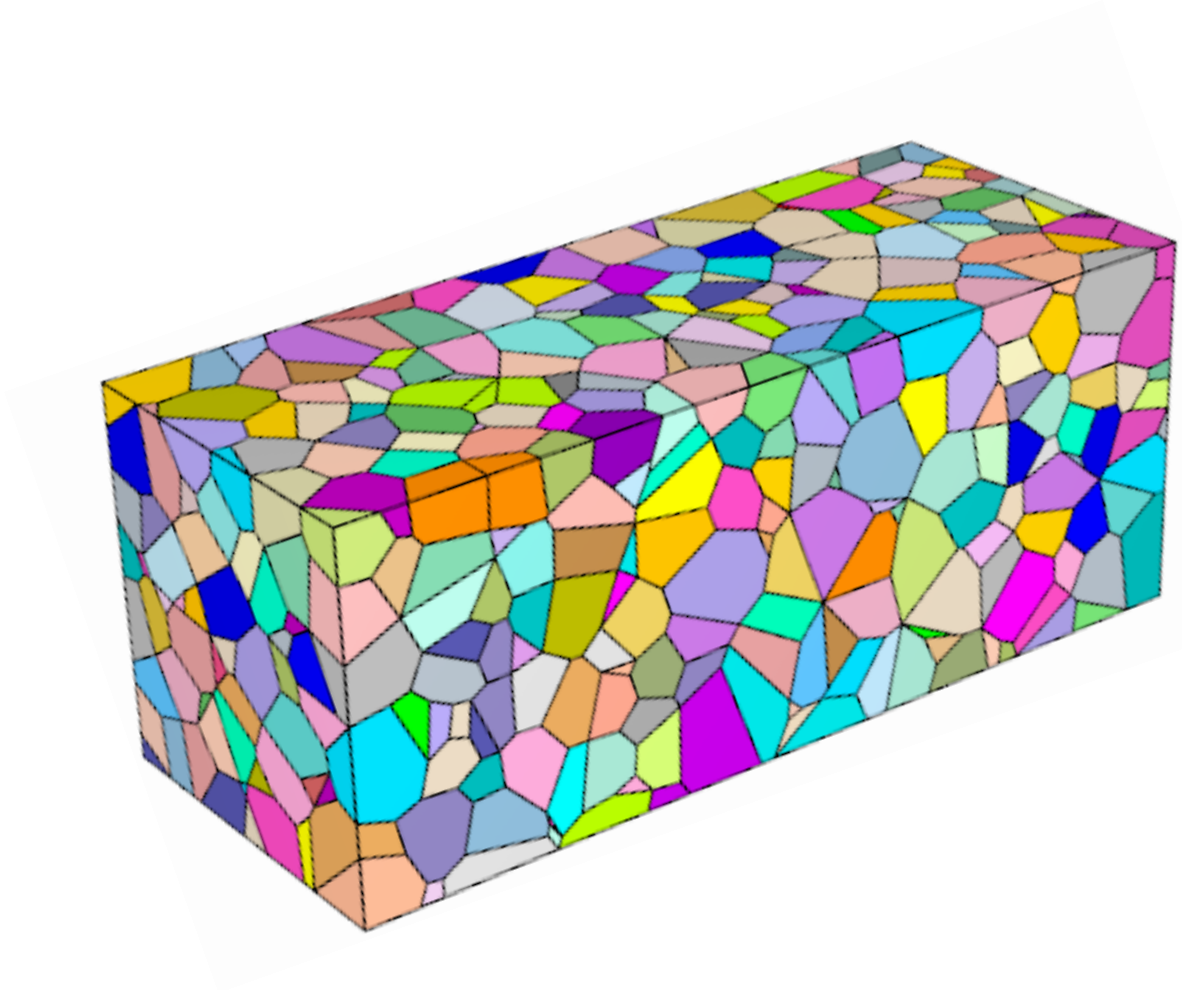

Figure 5.10: Slice (4mmx4mmx10mm) of the 3D model of a polycrystalline material with 500µm average grain size where the shades denote different grains. The full model contains 5210 grains and $16x10^6$ degrees of freedom.

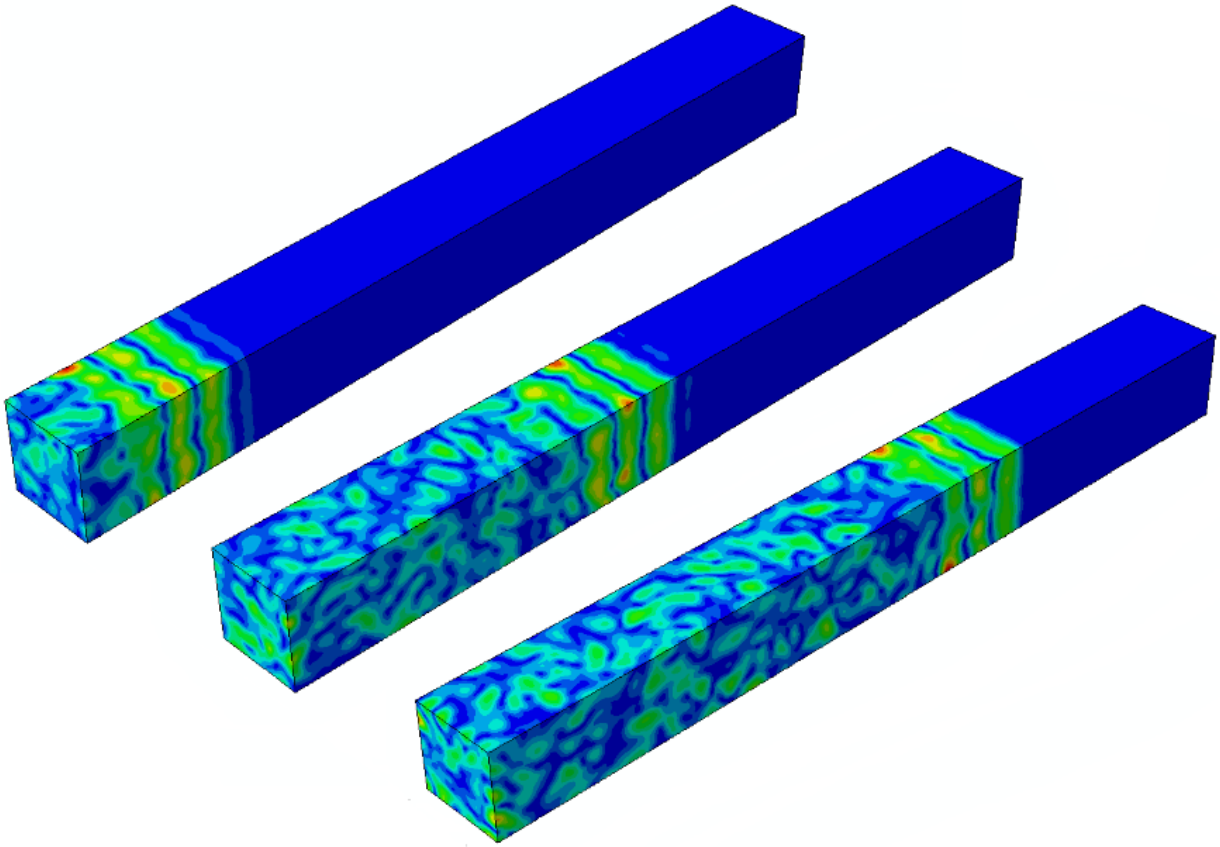

Figure 5.11: 3D FE simulation for a plane wave propagating throughout a polycrystalline material, Inconel, with an average 500μm grain size, shown at three different times: 1.5 μs, 3.5μs, and 4.5μs.

5.4.1 Attenuation

To start we compare the 2D and 3D FE results. The numerical attenuation is calculated by comparing the two frequency spectra corresponding to the transmitted signal and the pitch-catch received signal. This can be achieved by Fast Fourier Transforming the windowed time-domain signals and dividing the resultant frequency amplitudes, as explained by (Kalashnikov & Challis 2005) for example. Figure 5.12 shows attenuation against frequency for three cases. The results show that attenuation increases with both frequency and grain size, which suggests, at least initially, a good qualitative fit with the expected behaviour. By also plotting the power fitting coefficients for each simulation curve, we can further evaluate the results and determine the dominant scattering mechanism. This indicates that a fourth order frequency dependence for the Rayleigh regimes is only produced for the 3D simulation, whereas in 2D, only values close to three are produced. This might be explained by the 2D simplification, where the scattering cross-section is now proportional to the area and not volume of the grain; we expect that this would reduce the Rayleigh scattering to a third order frequency dependence in 2D, according to, for example the observations by (Chaffaï et al. 2000), although we are not aware of formal proof (see Section 5.4.4 for more details). This also confirms that the grains behave as Rayleigh scatterers and shows that, in this specific case, other scatterers, such as voids or material imperfections were not required to explain the dominance of Rayleigh scattering at low frequencies (Zhang et al. 2004).

Now we can compare the attenuation in the simulations to the theoretically predicted equivalent. According to the approach outlined by Stanke (Stanke & Kino 1984), the results are normalised such that they are independent of the mean grain size d. In Figure 5.13, the attenuation coefficient a, normalised through multiplication with d, is plotted against the normalised frequency (product of wavenumber k and d) on a log-log scale. Some ambiguity exists regarding the appropriate choice of d, as previous works (Ghoshal & Turner 2009) have used several values, namely, the mean grain size plus/minus one standard deviation of the grain size to match numerical and theoretical results. Although the choice of d significantly affects the results, in this work we have only used the mean grain size to normalise the results.

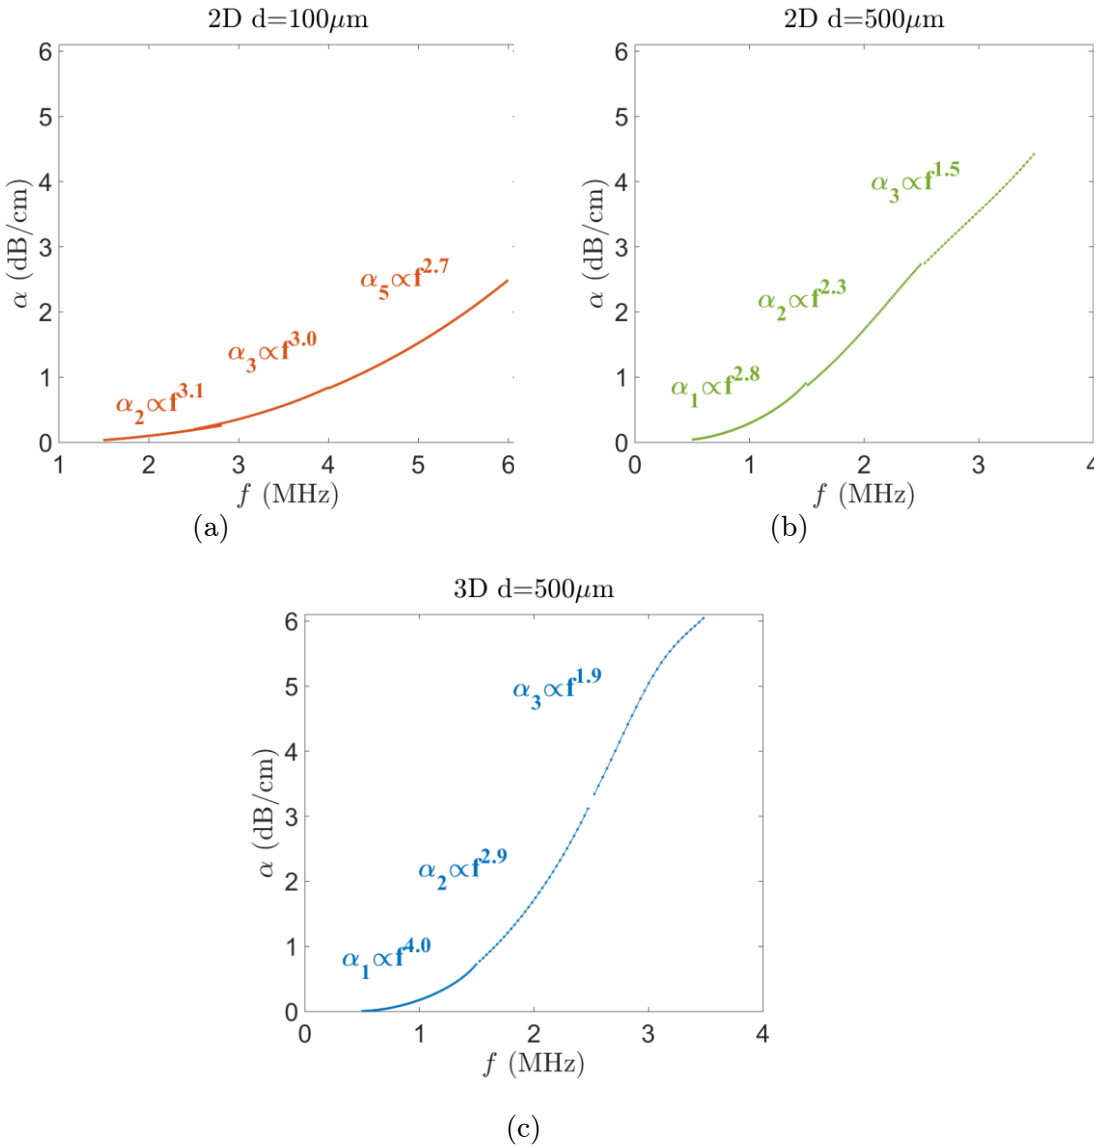

Figure 5.12: Frequency dependent attenuation in dB/cm against frequency, for (a) 100µm (b) 500µm grain sized material in 2D and (c) in 3D for 500µm. As expected the attenuation increases with frequency and grain size. The best-fit power coefficient is plotted for all nine (three per model) simulations, where the subscript denotes their centre-frequency in MHz. In the long wavelength to grain size ratios, the power coefficient approaches the Rayleigh result, whilst at higher frequencies, they converge towards the stochastic limit.

The Unified Theory, as shown in Figure 5.13, indicates the three scattering regimes; Rayleigh for $kd<<1$, stochastic $kd \approx 1$, and geometric $kd>>1$, which can each be recognised from their respective gradients, m, relative to their anticipated frequency dependence. In between the Rayleigh and stochastic regime, a transitional regime (Stanke & Kino 1984) exists where the frequency dependence can vary before converging to the stochastic asymptote.

As can be seen in Figure 5.13, the numerical results show good agreement with the established theory suggesting FE has the capacity to model the changing scattering behaviours across frequency. The match is not perfect, however, since the 3D model underestimates and overestimates at low and high frequency, respectively. In this case, the 3D model seems to agree slightly better with the theory in comparison with 2D, but the difference is marginal, and as previously mentioned, largely dependent on the choice of d. This would suggest that even with a simple assumption which overestimates the grain size, good matching with the behaviour of a 3D material is possible. Given the complex and random nature of these materials, these results are considered to be satisfactory.

5.4.2 Phase Velocity

Along with a complex frequency-dependent attenuation, propagating elastic waves in these materials exhibit small changes in phase velocity. Here we compare predictions of the Unified Theory to numerical results for phase velocity, obtained by comparing the phase angles of the transmitted and received signal. This can be achieved by Fast Fourier transforming the windowed time-domain signal and subtracting their unwrapped phase as explained by (Kalashnikov & Challis 2005) for example.

Figure 5.14 shows phase velocity as a deviation from the Voigt velocity (Voigt 1910), which is an average velocity for an equivalent macroscopically isotropic medium, calculated from the material elastic constants in Table 1. The x-axis plots the same logarithm of normalised frequency $\log(kd)$ as described in the previous section.

The results show that FE matches well with the Unified Theory, and there is good trend matching in the dispersive region, which is accurate to within 1%. The 3D results suggest a better match than 2D in this case.

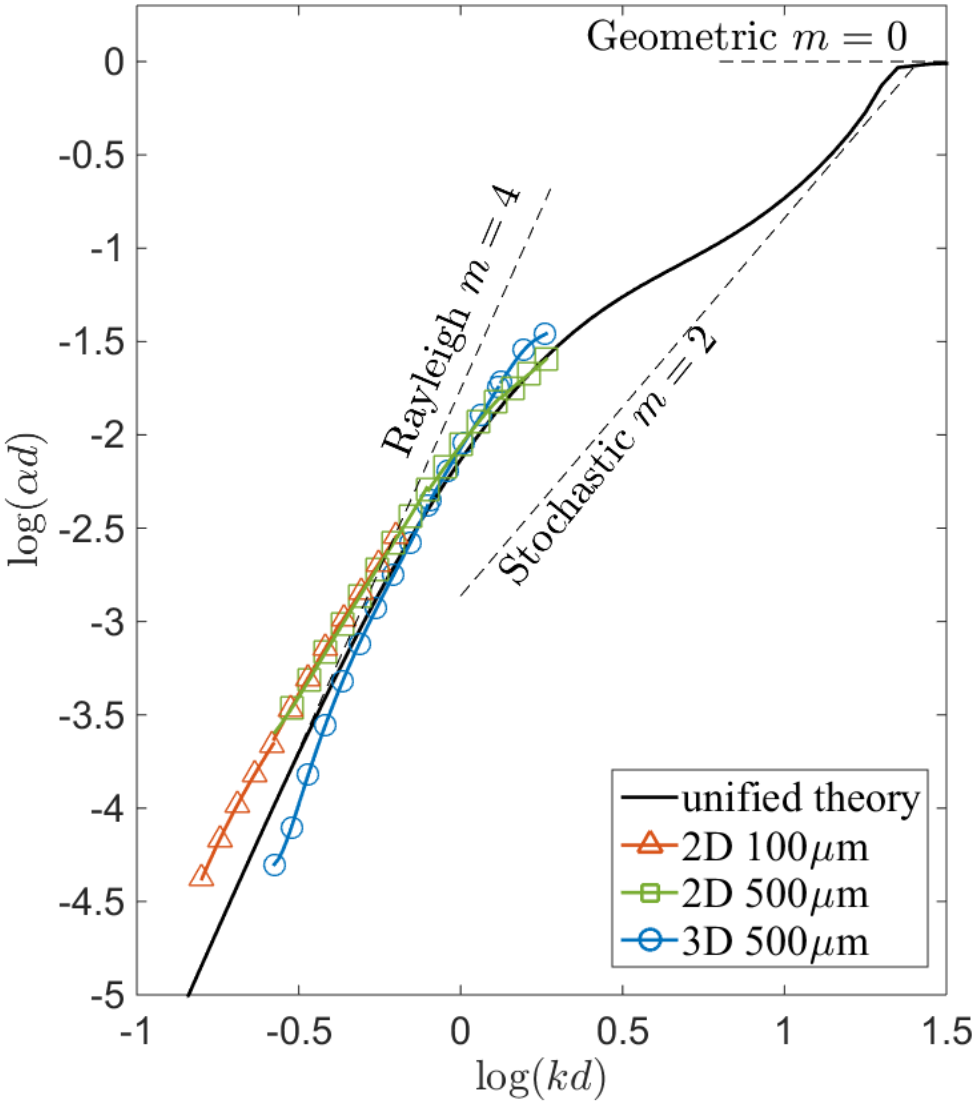

Figure 5.13: Normalised attenuation coefficient versus normalised frequency for a longitudinal wave in polycrystalline Inconel for for three different models, a 100um 2D (triangular maker), 500um 2D (rectangular maker), and 500um 3D (circular marker). The three different scattering regimes are indicated (dashed lines) with their respective gradients m. The attenuation results can be seen to compare well to the Unified Theory (black solid line). The empty markers are for labelling purposes only, and hence are not indicative of sampling.

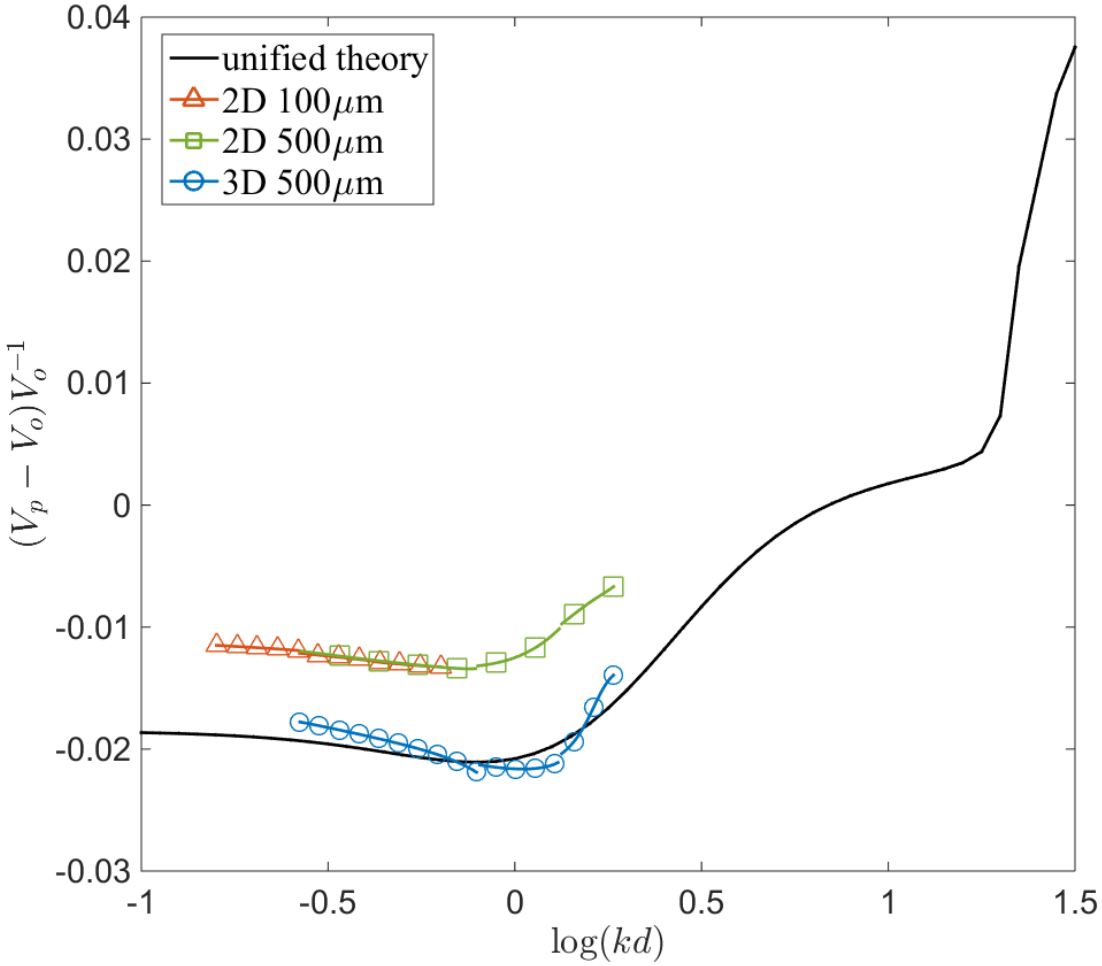

Figure 5.14: Normalised variation of longitudinal phase velocity against normalised frequency, for three different models of polycrystalline Inconel, a 100um 2D (triangular marker), 500um 2D (rectangular marker), and 500um 3D (circular marker). The results can be seen to compare well to the Unified Theory for both 2D and 3D finite element results. The empty markers are for labelling purposes only, and hence are not indicative of sampling.

5.4.3 Grain Noise Considerations

Thus far it has been shown that FE is capable of simulating the scattering physics with satisfactory attenuation and phase velocity accuracy. The remaining step then is to validate the grain noise which is also incorporated into the FE data.

As an initial assessment, attenuation and grain noise are interdependent phenomena as they are both governed by the same scattering physics. In this regard, the results for attenuation in the previous section also show promise for FE's ability to accurately model grain noise. Unlike attenuation however, noise is a spatially incoherent quantity, and hence the assumptions mentioned previously, such as reducing the model to 2D, may bear different implications. A thorough investigation of these limitations would require simulating hundreds of models to average numerous independent measurements of noise e.g. 500 has been used by (Margetan et al. 1994). This is beyond the scope of this work, and instead we will use the previously obtained simulation data to consider two simple scenarios. It is thought that this will provide an initial understanding and useful insight for future studies.

The first scenario quantifies the average backscatter received by individual nodes. As there are many receiving nodes, this enables temporal and spatial averaging of the fluctuating grain noise but since single nodes correspond to infinitesimally small transducers, it only offers relative comparisons which are not physically representative. In fact, for such a receiver, the backscatter noise will be artificially high. The backscatter noise, $N_{rms,}$ is calculated from the RMS value of the normalised time-displacement response during an 8μs time-window, corresponding to a time between but also away from the excitation and received signals. The resulting values, calculated for all the receiving nodes, are then averaged to produce the mean RMS backscatter $<N_{rms}>$. To clarify, an N_{rms} of 0.1 would signify an average grain noise level of 10% relative to the peak displacement of the incident plane wave (not attenuated).

The second scenario similarly involves calculating the backscatter, N_{rms}, but for a receiving transducer with a finite size. We investigate the implications of modelling a transducer in 2D compared to 3D, by assuming the example of a circular transducer. For

the 3D simulation, this entails averaging the displacement across a circular surface with radius r. In 2D, this reduces to an averaging of displacement along a line of length $2r$. Unlike the previous scenario, averaging across multiple independent transducers is impossible in this case, and hence N_{rms} is an un-spatially-averaged metric of the grain noise. This offers limited confidence for comparisons between simulations although some averaging is achieved by virtue of the multiple nodal results contributing to each transducer signal. Instead, the radius r is increased and the behaviour of N_{rms} for the 2D and 3D simulations is observed.

Figure 5.15 plots the mean RMS backscatter, $<N_{rms}>$, for simulations of varying centre-frequency shown in Table 2. The results confirm the expected increase in grain noise, both with frequency and grain size, for 2D and 3D simulations alike. Figure 5.16 compares the 2D and 3D results for 500μm grains by plotting the RMS backscatter N_{rms} for an increasing transducer radius. A maximum radius of 2mm was achievable for the 3D model due to size constraints (4x4mm cross-section).

Comparing the data points in 2D and 3D for the smallest radius in Figure 5.16, the results are similar to those of Figure 5.15 with the exception that the grain noise in 2D is higher for the 2MHz simulation than that of 3MHz. This is most likely a consequence of being restricted to an insufficient sample size to evaluate grain noise, which could be solved by averaging the grain noise for a batch of simulations.

Now taking into account the effect of increasing transducer radius, for both the 2D and 3D simulations, the grain noise in Figure 5.16 can be seen to decrease with receiver radius which is due to spatial averaging. It must be noted that this does not necessarily imply that SNR would also improve by increasing transducer radius, as the received signals from targets of interest may also decrease. Of the two simulations, the 3D receiver sees the sharper decline in grain noise with increasing radius which is explained by the spatial averaging progressing quicker over a 2D area than a 1D line. Similarly, spatial averaging can be expected to progress quicker at higher frequencies as the backscattered field exhibits a spatial correlation length which is inversely proportional to wavelength. The results at, for example r=2mm, in this specific case, demonstrate that reducing the model to 2D caused the grain noise level to be approximately three times higher than for a fully 3D simulation. This suggest that 2D and 3D models cannot be directly compared in terms of grain noise level, unlike for attenuation and phase velocity.

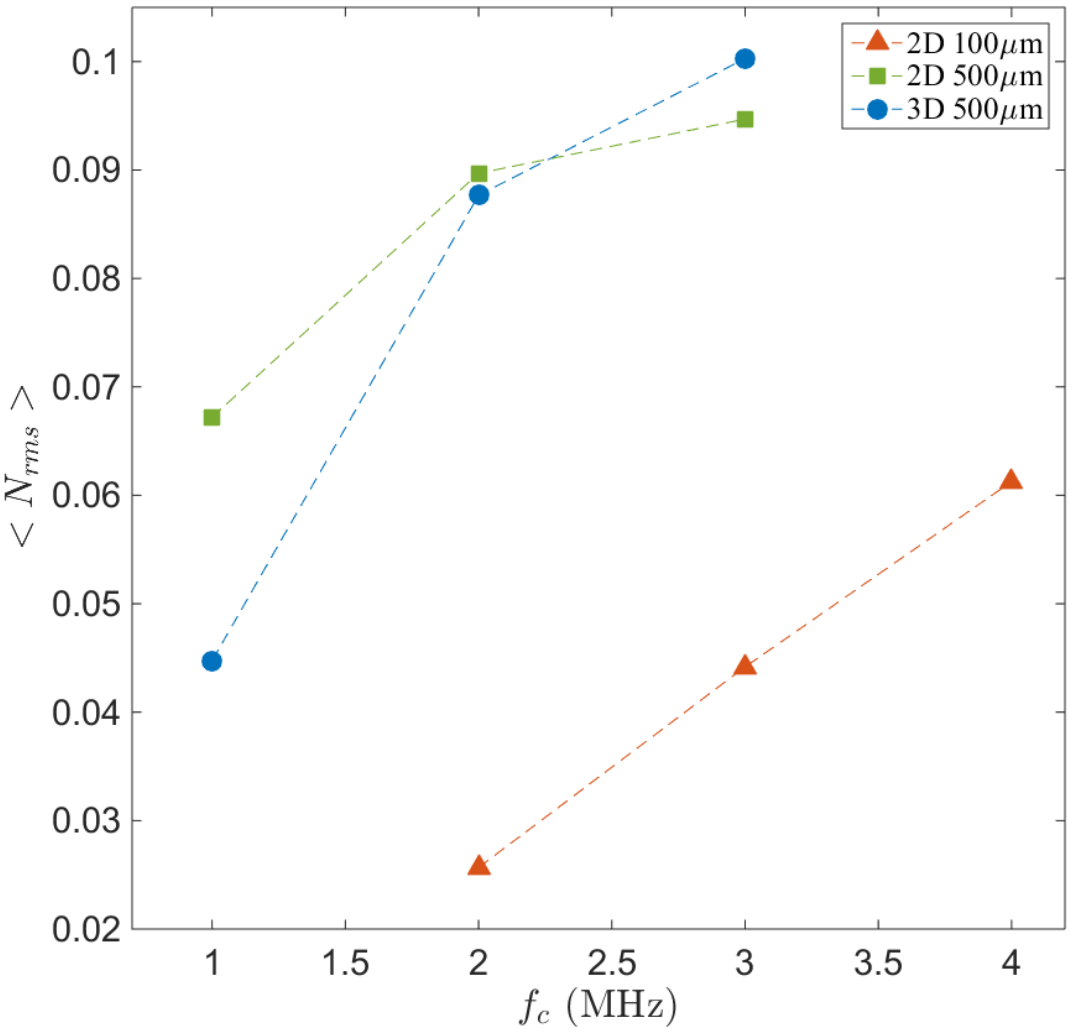

Figure 5.15: Normalised mean RMS backscatter noise (calculated from each nodal displacement) for different centre-frequency f_c simulations for 100µm 2D (triangular marker), 500µm 2D (rectangular marker), and 500µm 3D (circular marker). The backscatter noise can be seen to increase with frequency and grain size for all simulations.

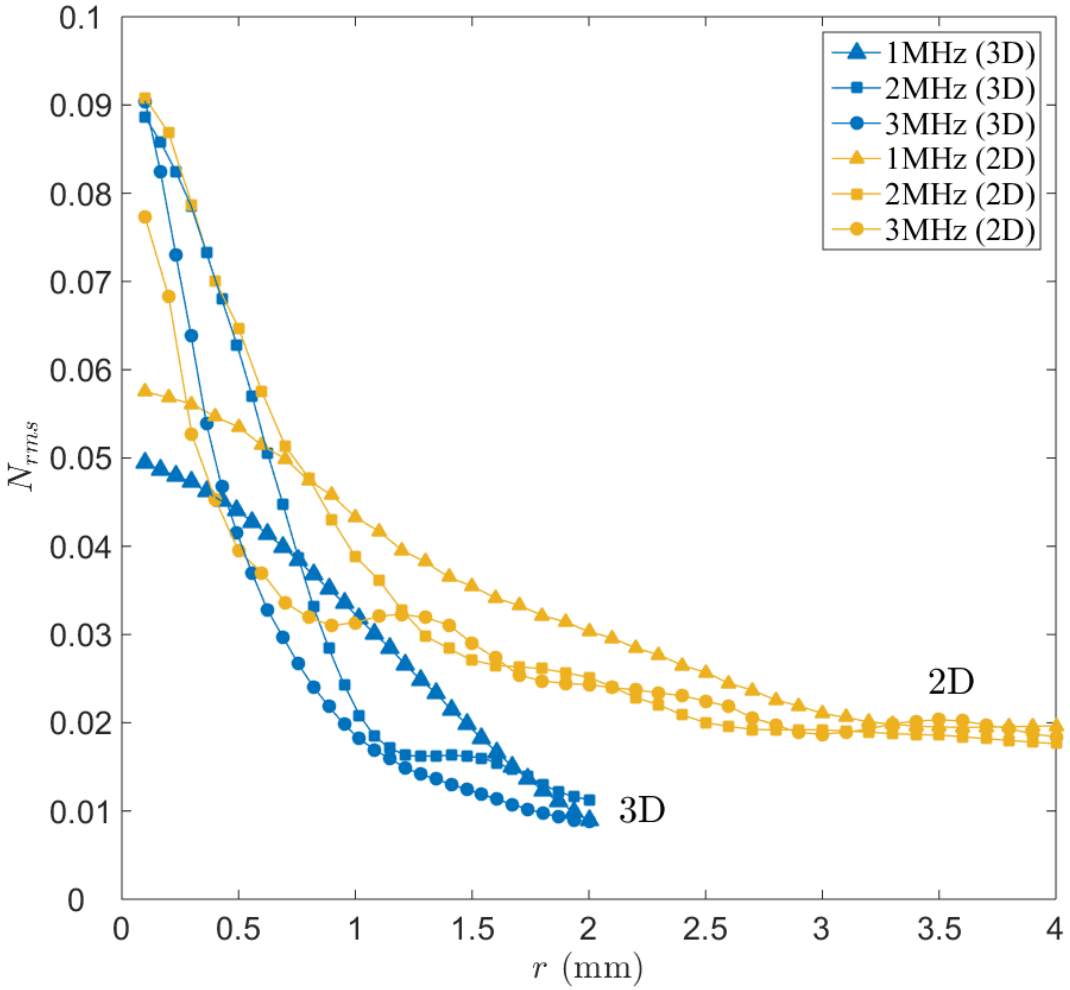

Figure 5.16: Normalised RMS backscatter grain noise against receiver radius r, for three 500µm models in 2D and 3D, each with varying centre-frequeny of 1, 2 and 3MHz. The backscatter noise can be seen to spatially avarage as the transducer radius increases.

5.4.4 Dimensionality of Scattering

In support of the previously observed dimensional behaviour of Rayleigh scattering, some additional simulations are run, this time however with repeat runs. This provides some indication of the variability of our calculation of attenuation and its power coefficient. In order to remain well within the Rayleigh scattering regime, materials are generated with a mean grain size of 100μm and insonified by a plane wave excited at a centre-frequency of 2MHz. Due to the necessity to run a statistically significant number of models - 10 are used here - instead of reconfirming the fourth order frequency dependence for 3D scattering, which would be extremely computationally expensive, it is far more efficient and equally instructive to instead confirm the relationship in 1D. The previous observations predict a second order frequency dependence for 1D scattering. The implementation of the 1D grain model is shown in Figure 5.17a, which is essentially reduced to a multi-layered material, with randomly varying thickness and stiffness for each layer. The 2D models remain unchanged from the previous results. The attenuation versus frequency plots in Figure 5.17b-c confirm the expected behaviour. It can be seen that the Rayleigh scattering mechanism has indeed reduced by an additional power in Figure 5.17b when scattering occurs in one dimension. Figure 5.17c reconfirms the third order frequency dependence for 2 dimensional equivalent. The variation in the power law obtained is sufficiently small to provide good confidence in the results.

(a)

(b)

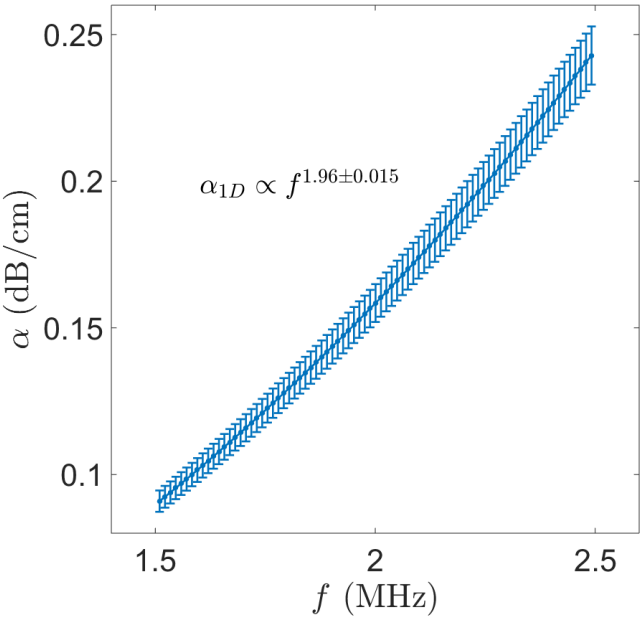

(c)

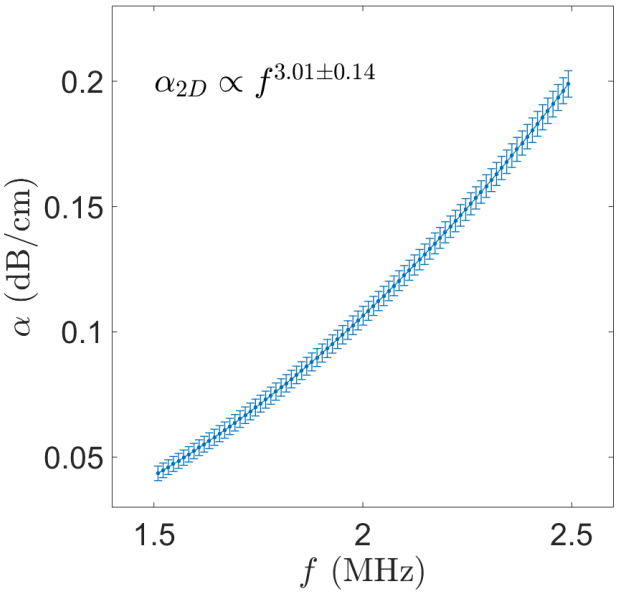

Figure 5.17: (a) 1D model layout where indicating layers of random thickness and random anisotropic orientation, attenuation results in (b) 1D and (c) 2D for 10 realisations of a 100µm grain size material with a 2MHz centre-frequency.

5.5 Conclusions

This chapter has set out to present and assess new progress in capabilities of Finite Element (FE) modelling to simulate ultrasonic scattering of longitudinal waves in an equiaxed and untextured polycrystalline material, for both 2D and 3D. The modelling adopts an established Voronoi approach to randomly generate a representative grain layout. Relying on a recently developed GPU FE solver, Pogo, large parametric studies in 2D and a single 3D model became feasible. The 2D parametric studies illustrated the mesh sampling requirements for two different types of mesh and different levels of mesh refinement, to ensure modelling accuracy and present useful guidelines for future modelling of these materials. During comparison to established theory, for both 2D and 3D, the numerically calculated attenuation and phase velocity showed good agreement across a range of scattering regimes. This is the first achievement of useful results at sensible scale in 3D and also the first time that FE has been shown to match to the Unified Theory. Hence this suggests, with relatively simple descriptions of these materials, that this type of numerical modelling has the ability to capture the key physics. Modelling limitations were also found. It was shown that 2D models fundamentally reduce the power of the frequency dependence of the scattering mechanism in the Rayleigh regime. Furthermore, the grain noise differs significantly between 2D and 3D simulations which is an important characteristic to consider for real applications. Overall, it is proposed that the progress and understanding presented in this chapter will aid the ongoing improvement of FE simulations of ultrasonic NDE of polycrystalline materials.

Chapter VI

ARRAY OPTIMISATION

A Finite Element modelling framework is outlined to simulate the ultrasonic array response of elastic waves propagating within highly scattering, polycrystalline materials. Its utility is demonstrated by investigating the performance of arrays, within both single and multiple scattering media. By comparison to well-established single scattering models, it is demonstrated that FE modelling can provide new insights to study the stronger scattering regimes. In contrast to established single scattering results, it is found that phase aberrations significantly hamper array performance, and that the noise becomes spatially incoherent. Consequently, Signal-to-Noise Ratio (SNR) no longer increases monotonically with respect to increasing aperture, which suggests that maximum apertures are not necessarily optimal. Furthermore, by measuring the SNR of the individual transmit receive combinations of the array, it is found that through separating the emitter and receiving source, it is possible to reduce the received backscatter.

This work has been submitted to NDT&E 2015 [p8].

6.1 Introduction

This chapter outlines a FE modelling framework, developed as an extension to that in Chapter 5, which enables the investigation of ultrasonic array imaging of highly scattering, polycrystalline materials. It details modelling devices which allow the isolation of different physical phenomena (e.g. element directivity, beam spreading, attenuation, backscatter) and therefore enables new and useful insights into the effects of scattering – particularly without relying on a single-scattering assumption. The methodology is applied to a relatively simple but also general case such that it both illustrates and investigates the fundamentals of array performance. The approach is now also ready for a wide variety of simulations where it can be useful in future evaluations of performance: for instance to determine the optimal configuration for a more practical inspection, quantify the smallest detectable defect, or assess new data processing algorithms such as new candidate array imaging algorithms.

The subsequent sections are organised as follows. Section 6.2 outlines the FE methodology to investigate an array probing a polycrystalline medium. Section 6.3 builds on established theory by studying array performance within a single scattering environment. Section 6.4 repeats this same procedure but extends this to stronger scattering media by introducing polycrystalline material properties. Section 6.5 compares the single scattering theory to those obtained from the highly scattering material simulations. But first, before setting out with these studies, we explore the currently established theory for determining detection performance of an array imaging a noisy but single scattering medium.

6.1.1 Established Single Scattering Theory

Single scattering models, such as the aforementioned ISM (Margetan et al. 1993; Margetan et al. 1997), determined that SNR is inversely proportional to the ultrasonic pulse volume for monolithic transducers. This led to the adoption of focused transducers to improve sensitivity of industrial inspections of scattering materials. More recently, Wilcox (Wilcox 2011) (and others e.g. (Felice et al. 2015)) found similar results for arrays

by showing SNR to depend on the Point-Spread Function (PSF) of an array (see Equation (6.1).

$$SNR(\boldsymbol{r}) = \frac{q}{\mu}\frac{|P(\boldsymbol{r},\boldsymbol{r})|}{\sqrt{\int |P(\boldsymbol{r},\boldsymbol{r}')|^2 dr'}} \tag{6.1}$$

Here μ is the backscatter coefficient, derived for polycrystalline materials by Rose (Rose 1992) and q is the scattering potential of the imaging target. Outside these two parameters, the remainder of Equation 6.1 is defined by two Point Spread Functions, e.g. $P(\boldsymbol{r},\boldsymbol{r}')$ is the image response at $\boldsymbol{r}$ of an idealised single point scatterer located at $\boldsymbol{r}'$. Thus the remaining fraction is solely determined by the imaging system and is equivalent (Wilcox 2011) to the reciprocal of the square root of the normalized PSF area, σ. For our purposes of finding an optimum, only relative SNR is of interest, and hence in the studies presented here we can disregard (Felice et al. 2015) the two parameters μ and q and redefine a relative SNR, denoted by $\overline{SNR}$ and related to σ in Equation 6.2.

$$\overline{SNR} \propto \frac{1}{\sqrt{\sigma}} \tag{6.2}$$

The relation between $\overline{SNR}$ and the PSF has several interesting connotations. It firstly implies the monotonic increase of SNR which thus maximises when using the largest possible aperture (Li et al. 2013). Aside from providing a definition of the focusing ability of the array, as shown in Equation 6.2, the PSF also allows the prediction of SNR under single scattering assumptions. The PSF area, σ, is a widely used metric and can be quantified in various ways (see e.g. (Holmes et al. 2005)), the approach adopted here is to calculate the area of the PSF which encompasses half its peak.

The PSF comprises the imaging system and can thus be controlled by optimising the array and the imaging algorithm; we will focus on the former using a specific choice of imaging algorithm.

6.1.2 Array Imaging Algorithms

Within the field of ultrasonic array imaging, there has been a recent surge of advanced imaging algorithms such as those initially outlined in Section 2.4 (see e.g.(Aubry & Derode 2009a; Simonetti 2006)) which have shown impressive progress. As highlighted in Chapter 4 however, it has proven challenging to supress coherent noise and thereby increase imaging performance beyond that provided by standard sum-and-delay beamforming. Consequently, the currently most popular algorithm, the Total Focusing Method (TFM)(Holmes et al. 2005), for the time being, remains the benchmark. It offers high performance as well as relative simplicity. Thus this chapter will rely on TFM for its investigations and illustrations; it is expected that the findings will be equally relevant for other imaging algorithms.

6.2 Method: Simulation of Highly Scattering Materials

Here we discuss the adopted methodology, firstly the incorporation of polycrystalline material properties into FE, followed by the array model, which consists of a noise and a signal model.

6.2.1 Modelling Polycrystalline Materials

The methodology from the previous Chapter 5, to incorporate polycrystalline material properties within an FE mode, is adopted here. The mesh used here comprises a structured grid of triangular elements, sampled such that the length of the element edge is finer than at least one tenth of the average grain size d to meet the criteria for convergence.

Given the computational cost, the relatively large dimensions necessary for our studies, and the interest here in performing multiple analyses in order to pursue a range of studies, the models discussed here are limited to a 2D domain. As mentioned previously, this simplification introduces certain model limitations: the scattering mechanism is reduced to a third order frequency dependence in the Rayleigh regime, and it is not obvious how to relate the spatially incoherent fields, namely the grain noise, perceived by a 2D transducer to that of a 3D one. It was previously shown however that 2D models overestimate the

absolute level of noise as there is less spatial averaging which occurs across the length of the transducer, as opposed to an equivalent area in 3D. Despite this lack of absolute accuracy, the relative accuracy is expected to be good, as the overall frequency dependent scattering behaviour has been shown to correlate well to established theory. This is deemed sufficient as we are primarily interested in examining trends and principles rather than absolute performance metrics. Moreover, the principles discussed in this chapter will apply equally well in 3D, and since 3D representation has been shown to be possible, it will only be a matter of deploying these methods in 3D once this becomes computationally feasible.

6.2.2 Array Model

The layout for the general ultrasonic array model used hereon is depicted in Figure 6.1a. The model typically simulates N=128 element arrays which are fully sampled, such that the array pitch and width both measure half a centre-wavelength. The array data acquisition adopts a Full-Matrix-Capture (FMC) (Holmes et al. 2005) approach which involves sequentially exciting all N array elements, and for each excitation, calculating the response also on all N array elements. The excitation at the i^{th} element of the array is simulated by applying a piston-like force which is perpendicular to the contact surface (producing a longitudinal wave amongst others), to all the nodes which correspond to the footprint of the i^{th} array element. In reception, the nodal displacements of all the nodes belonging to the j^{th} array element are averaged, again taking the component of displacement in the direction normal to the surface. Varying both j and i from 1 to N, this populates an FMC matrix, H, of dimensions $N \times N \times t$ where t corresponds to the number of time samples.

Noise Model

Before the introduction of any imaging target within the model, this procedure yields an array response matrix, H_N, which pertains solely to the grain noise (and reflections from the structural boundaries). This can be thought of as an artificial baseline measurement, as is commonly referred to in Structural Health Monitoring (Croxford et al. 2007).

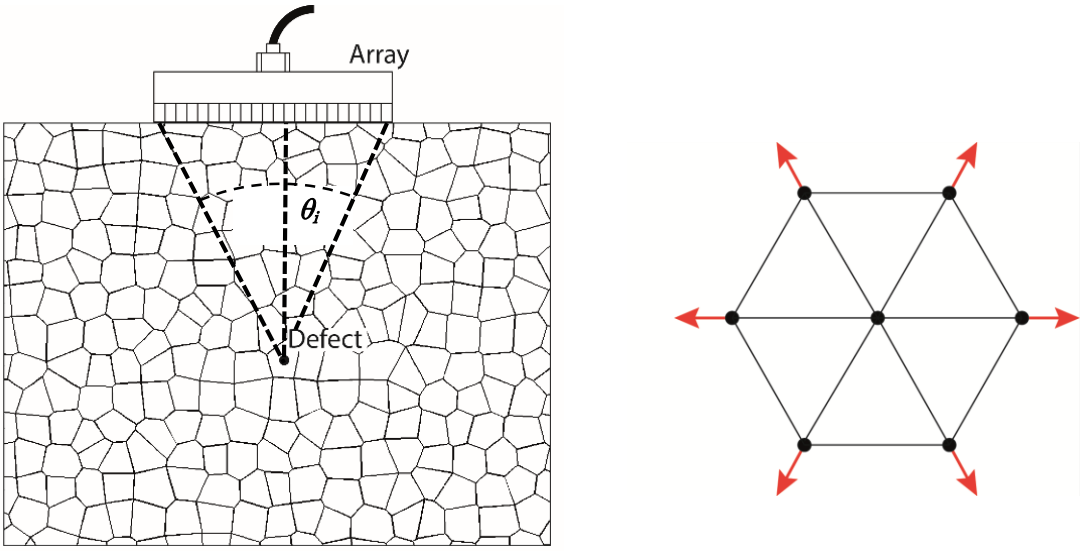

(a) (b)

Figure 6.1: (a) Schematic of the Finite Element array model layout (not drawn to scale) and (b) schematic of the true point scatterer implementation into FE.

This is useful for separately analysing the signal and noise data which enables the modelling of the true SNR. This is a valuable tool in general, as investigations are often limited to measure signals which contain noise, thereby constrained to solely measuring positive SNRs (as discussed in Chapter 4) which offer a limited utility as a performance metric. There are several ways to circumvent this, one being subtraction (Van Pamel et al. 2015), and another which is outlined in the proceeding sub-section.

Signal Model: True Point Scatterer

To obtain the previously defined PSF, we desire an ideal, omnidirectional scatterer - here referred to as a true point scatterer (TPS). A widely accepted practise is to use voids or disconnections within an FE mesh to simulate defects and scatterers. The obvious procedure to create a single point scatterer then would be to constrain or disconnect a single node. However, this does not produce a true omnidirectional scatterer; instead of the desired isotropic scattering, the scattering amplitude of the longitudinal wave varies with angle, dropping to null as the difference between the incident and scattered wave approaches 90°.

This is circumvented here by exploiting the properties of reciprocity which allows us to reverse the sender and receiver. Instead of insonifying the domain from the array, and looking for scattering back from the defect, the defect is used to insonify the domain, and the projected field is received at the array. The principle of reciprocity is then used to complete the send-receive signals and calculate the array FMC. This approach allows any desired scattering characteristic to be implemented, in this case uniform omnidirectional scattering.

The implementation of a circular wavefront outgoing from the point scatterer is achieved here by radially exciting six neighbouring nodes of a structured mesh in a hexagonal arrangement (see Figure 6.1b). To finally obtain the FMC corresponding to the signal model, requires the convolution of the 1xN vector of received signals with its transpose, to obtain an NxN data matrix, H_s. The resultant wave field from such a point source is illustrated in Figure 6.2, within (a) an isotropic and (b) polycrystalline medium.

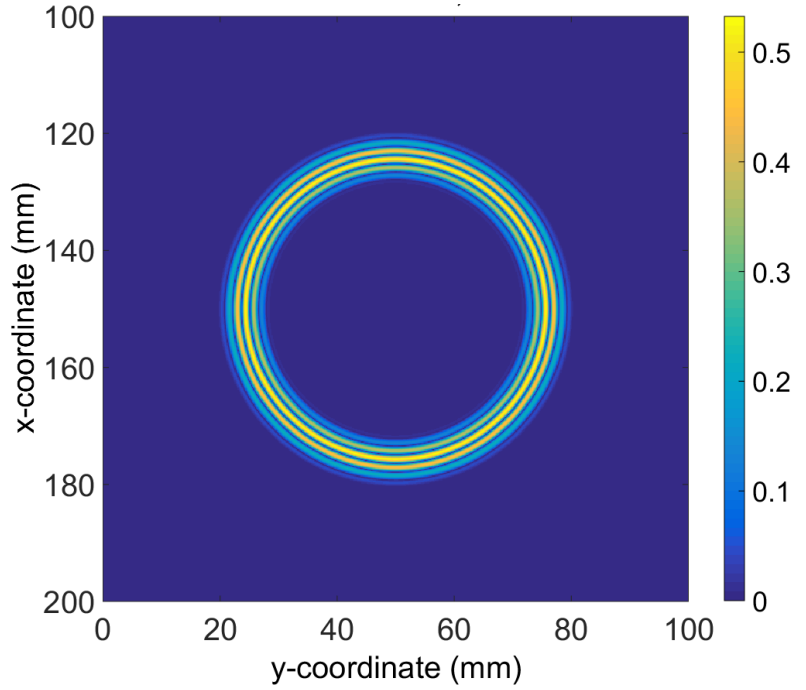

(a)

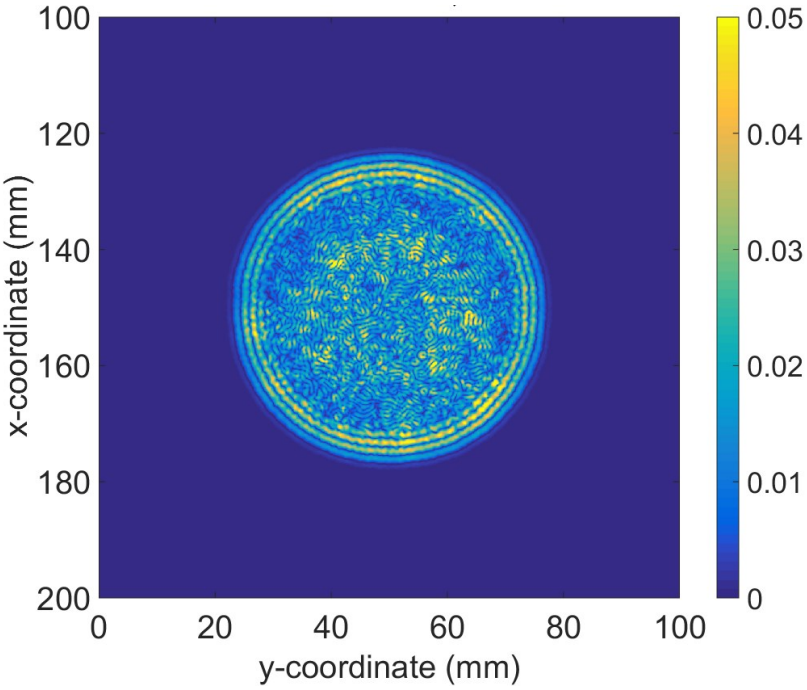

(b)

Figure 6.2: FE illustration of the wave field emanating from a true point source (TPS) propagating after 5μs within (a) isotropic homogenous and (b) polycrystalline material. The colour scale denotes the displacement amplitude at the selected time.

One noteworthy consequence of this approach is that the scattering potential of our imaging target (denoted by q in Equation (6.1) is arbitrarily defined by the excitation amplitude defined in the FE simulation. Moreover, we cannot calculate an effective incident amplitude for the circular wave, as a singularity exists at the centre where the radius equals zero and the theoretical incident amplitude tends to infinity. In our case however, as discussed in Section 6.1.1, we are only interested in a relative SNR to find an optimum, and hence to clarify this, we shall distinguish from the SNR by the term $\overline{SNR}$ to denote a relative quantity.

6.3 Results I: Simulation of Single Scattering Media

The single scattering theory (Wilcox 2011) outlined in Section 6.1 is now used to validate a Finite Element model of an array operating within a single scattering medium, here modelled by a random distribution of point scatterers within an isotropic material. By adopting a single scattering assumption, we can solve the PSF for each scatterer independently (Wilcox 2011; Jie et al. 2013). Furthermore, when considering both the noise scatterers and the imaging target as omnidirectional scatterers, the solution of one PSF provides that for all others – be it grain noise or target (Wilcox 2011; Jie et al. 2013). This approach purposefully neglects any multiple scattering effects, which serves as a benchmark for comparisons when the polycrystalline microstructure is introduced later on (see Section 6.4).

The model defines a fully sampled N=128 element array, generating a 3-cycle tone-burst longitudinal wave with a 2MHz centre-frequency in contact with an isotropic elastic material. The medium is arbitrarily defined by a longitudinal wave speed of 6123m/s (E=230GPa, ν=0.3, ρ= 8200kg/m^3). Three defect scenarios are simulated to calculate the signal FMC, H_s, for a TPS defined by the procedure in Section 6.2.2, and introduced respectively at a 25mm, 50mm, and 75mm depth within the material.

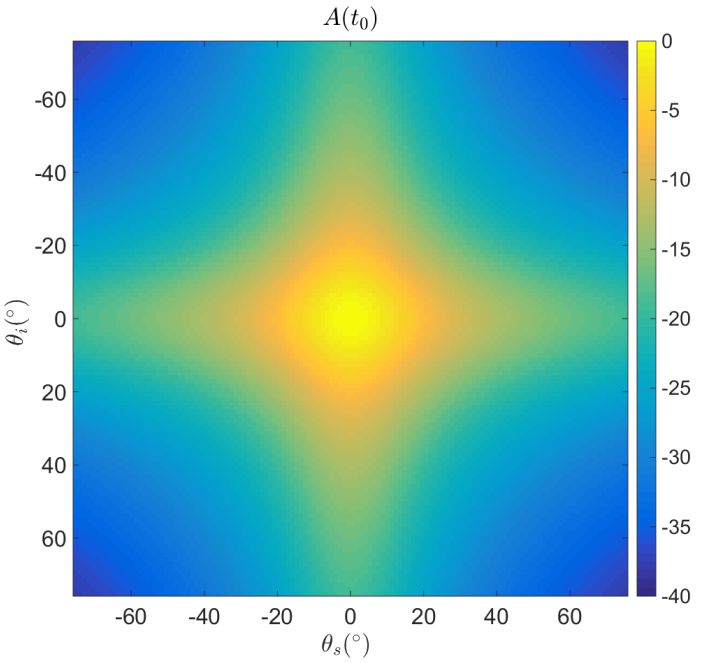

(a)

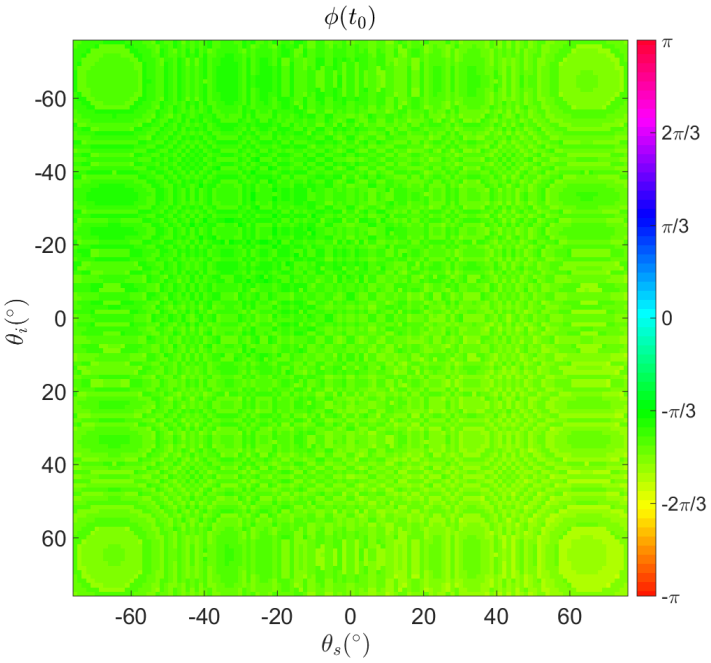

(b)

Figure 6.3: Scattering matrix for a TPS embedded in isotropic material at a 50mm depth showing (a) amplitude normalised by the peak, and (b) wrapped instantaneous phase shown in radians

6.3.1 True Point Scatterer

The behaviour of our model TPS within an isotropic material is validated by calculating apparent scattering matrices (Wilcox & Velichko 2010) from H_s. An amplitude scattering matrix procedure follows (Wilcox & Velichko 2010), however unlike classical scattering matrices our setup comprises a linear array with a limited view to sample the wave field, rather than a circumferential full view configuration. Adopting the notation of an analytical signal (Gabor 1946) the amplitude matrix used here plots the instantaneous amplitude $A(t_0)$, where t_0 corresponds to the arrival time of the signal, as a function of incident θ_i and scattered θ_s angle (defined in Figure 6.1), to produce a 2D matrix $A(\theta_i,\theta_s)$. A phase matrix is also calculated which follows the same syntax but instead of amplitude, calculates the instantaneous phase $\varphi(t_0)$ to obtain $\varphi(\theta_i,\theta_s)$. An example of how to calculate both instantaneous amplitude and instantaneous phase can be found in (Feldman 2011).

As can be seen from the amplitude scattering matrix in Figure 6.3a, the TPS defect exhibits omnidirectional scattering behaviour, as intended. Due to the absence of noise and attenuation in this case, the only drop in scattering amplitude occurs due to longer propagation distances and large receiver angles. The longer propagation distances will cause the wave amplitude to decrease due to beam spreading effects, and a loss in element sensitivity occurs at large angles as the array elements exhibit a directional sensitivity, which reduces as the incident wave moves away from the normal. The rhomboidal features manifest at larger angles are due to the linear array configuration where the propagation distance is not constant with total aperture angle (propagation distance increases non-linearly with angle). Figure 6.3b confirms the absence of aberrations in the phase matrix, as is expected for the isotropic case.

6.3.2 Point Spread Function

The PSF for a TPS within an isotropic medium is shown in Figure 6.4a, where low-intensity side lobes can also be identified. The *PSF* area (σ) is quantified as a function of the half aperture angle, denoted by θ_p and calculated by the halved sum of θ_i *and* θs (angle labelled in Figure 6.1a). As previously defined, the PSF calculates the area which encloses the PSF within -6dB from its peak.

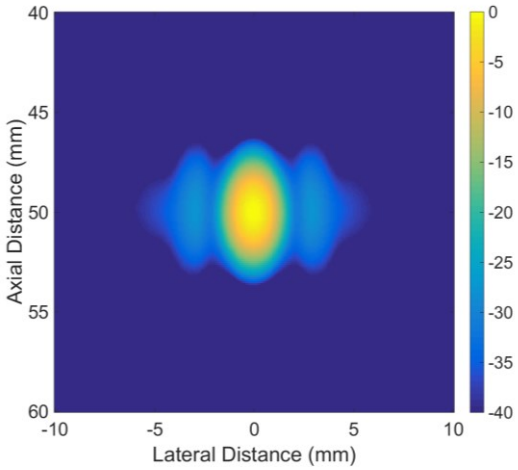

(a)

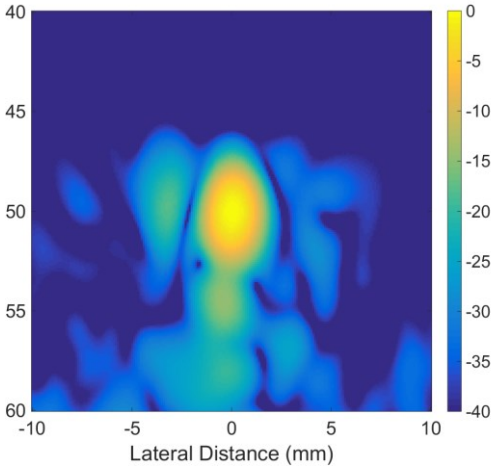

(b)

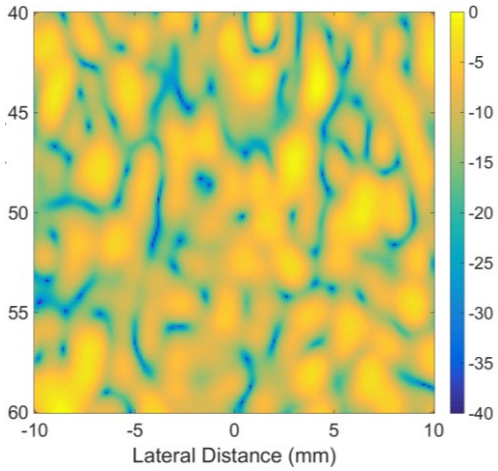

(c)

Figure 6.4: (a) Point Spread Function for an isotropic material (b) aberrated Point Spread Function for a polycrystalline material (c) noise baseline for a polycrystalline material.

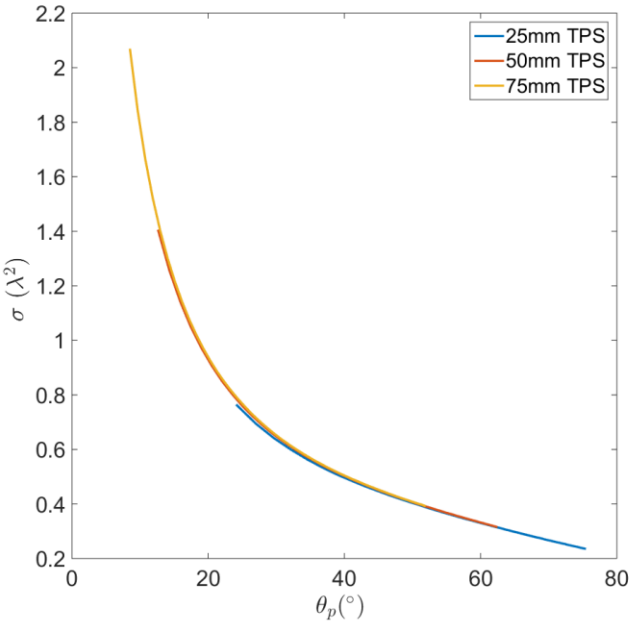

(a)

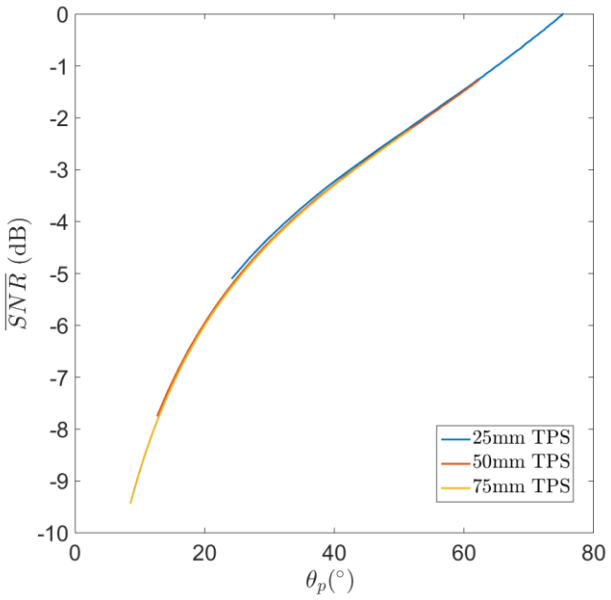

(b)

Figure 6.5: Simulation results for a true point scatterer embedded within a single scattering medium at 25mm, 50mm, 75mm depth. (a) PSF area versus half aperture angle (b) relative SNR versus half aperture angle.

Figure 6.5a plots the resultant *PSF* area (σ) as a function of half aperture angle θ_p, for three TPS at various depths. As can be seen the PSF decreases monotonically, where focusing benefits progressively lessen at high aperture angle as is dictated by the asymptotical diffraction limit (Wilcox 2011).

6.3.3 Predicted Signal-to-Noise Ratio

The previously obtained PSF area now allows the prediction of a relative SNR in a single scattering environment, $\overline{SNR}$, as defined in Equation (6.2).

Figure 6.5b shows $\overline{SNR}$ versus half aperture angle θ_p and predicts a monotonically increasing SNR, independent of defect depth. These results agree with the experimental and model findings of Wilcox (Wilcox 2011), thereby validating our single scattering model. Now we investigate the effects of multiple scattering by repeating the same simulation but with the introduction of polycrystalline material properties.

6.4 Results II: Simulation of Multiple Scattering Media

The procedure outlined in Section 6.3 is now repeated for a polycrystalline medium which introduces intrinsic scattering and thus no longer relies on a single scattering assumption. Using the same layout depicted in Figure 6.1a, with a 2MHz 3-cycle tone-burst, exciting longitudinal waves from a N=128 element array, images are acquired of targets buried at depths of 25mm, 50mm, and 75mm. The medium is non-textured cubic Inconel 600, defined by the single elastic stiffness constants taken from (Shahjahan, Rupin, et al. 2014). The grain morphology consists of equiaxed grains with their mean size set at 500μm, which places the scattering behaviour at centre-wavelength, in between the Rayleigh and stochastic scattering regimes. Although the author is not aware of a formal definition for its onset, it is assumed that multiple scattering occurs within the stochastic regime which begins for kd values of unity, where k is the wavenumber and d is the mean grain size, and given the high anisotropy value of Inconel 600 (2.8). Eight independent 2D models are run, each with the same mean grain properties but different realisations of a

random polycrystalline morphology. This provides us with a basis, albeit with a modest number of samples, to consider statistical variations.

In contrast to the earlier simulations where the PSF provided solutions for both the noise and signal model, separate simulations are required to obtain the noise data H_N and H_s: the data from a TPS.

6.4.1 Aberrated True Point Scatterer

Similarly to Section 6.3.1, we establish the behaviour of a TPS, in this instance however by considering propagation within a polycrystalline medium. In comparison to the isotropic medium considered, which incorporated beam spreading and element directivity, our signal model, H_s, now includes additional physics such as the scattering induced attenuation, dispersion, and phase aberrations.

The scattering amplitude and phase matrices (see Section 6.3.1 for methodology) for a TPS at 25mm depth, are shown in Figure 6.6 for one random realisation of a polycrystalline material. When compared to the isotropic case in Figure 6.3 it can be seen that the amplitude fluctuates, but in general depicts a similar picture to that of the isotropic case where the highest amplitudes occur at (0°,0°) angles. In terms of phase however, whereas Figure 6.3b showed no variations for the isotropic case, significant phase aberrations can be seen of up to 1π radians. This observation is further illustrated in the earlier Figure 6.2b, where aberrations can be seen to occur along the circular wavefront. However, it is interesting to note that the aberrations also seem to exhibit an underlying regularity. This could be explained by the inherent dispersion within these materials, where at long propagation distances, small errors in the assumed wave velocity cause phase shifts. This is expected to significantly hinder focusing ability of the array, which is quantified next by observing the PSF.

6.4.2 Aberrated Point Spread Function

Figure 6.4b shows an aberrated PSF for a TPS within a polycrystalline material where perturbations have now arisen when comparing to Figure 6.4a. The PSF area is calculated

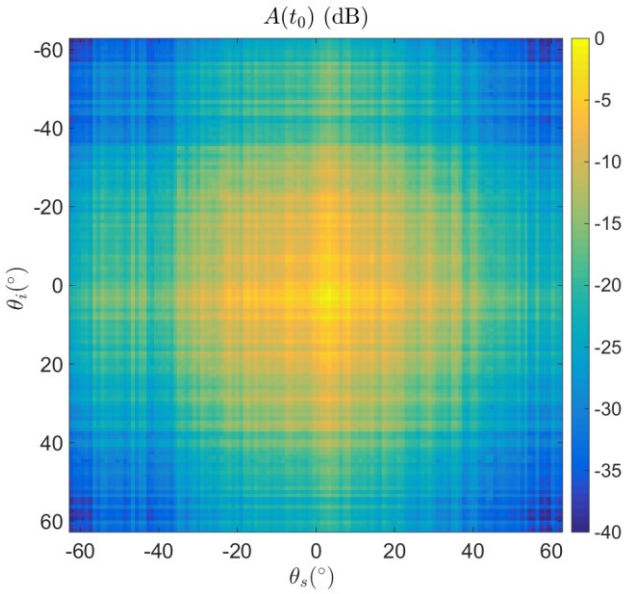

(a)

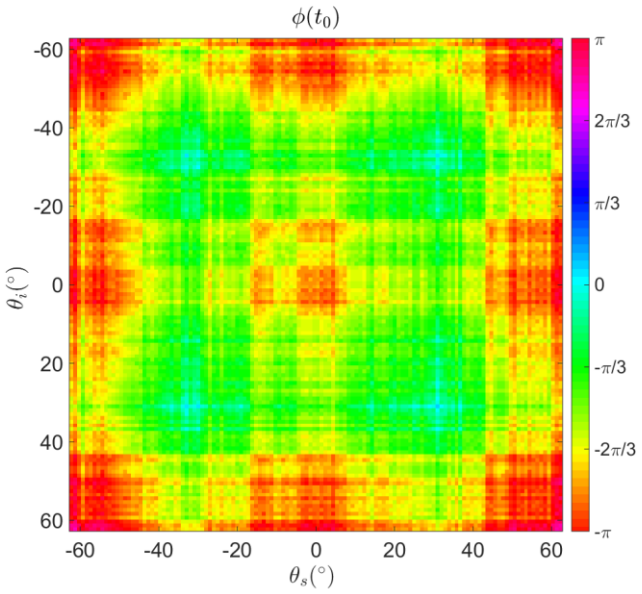

(b)

Figure 6.6: Scattering matrices for a true point scatterer at 50mm depth in a single realisation of a polycrystalline material, indicating (a) instantaneous amplitude and (b) instantaneous phase against incident and scattered angle.

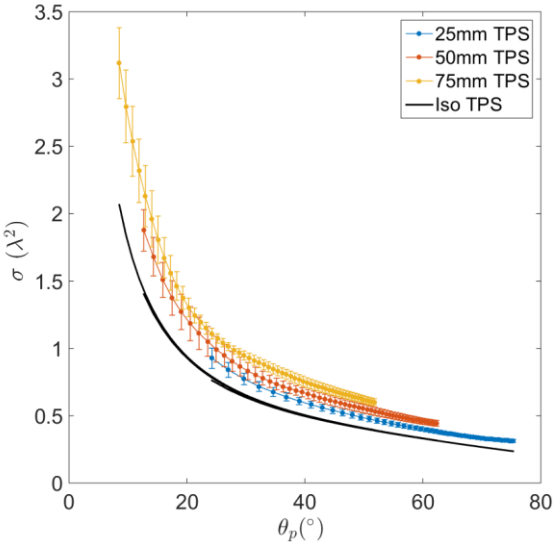

(a)

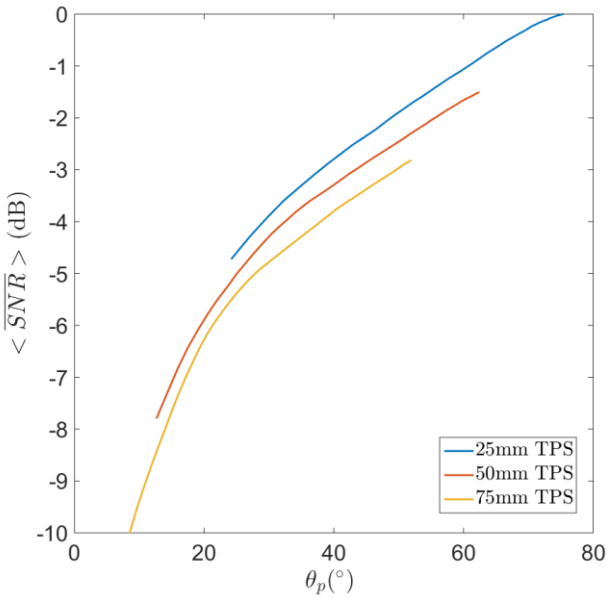

(b)

Figure 6.7: (a) PSF against half aperture angle for a true point scatterer embedded within a polycrystalline material. Results from previous section 6.3.2 shown in black. (b) mean SNR against half aperture angle calculating under single scattering assumptions. Results obtained by averaging from eight realisations of material with the same grain statistics. Error bars indicate the standard deviation variation.

similarly to the procedure adopted in Section 6.3.2, but repeated for eight realisations of a random polycrystalline material to consider statistical variation. The eight PSF areas are then averaged and their standard deviation is also recorded.

Figure 6.7a plots the mean PSF area (σ) and its standard deviation bars, versus half aperture angle, θ_p, for the three TPS cases. Comparison with Figure 6.5a reveals that the polycrystalline material has induced several changes. Firstly, it can be seen that the absolute focus has worsened, indicating a PSF which is approximately 50% relative to the previous case. This indicates that even before considering the effects of coherent noise, which probably presents further hindrance to image quality, the focus (which is related to SNR) has already been harmed. Furthermore, although PSF area remains a monotonic function with respect to aperture angle, it has become also a function of depth. Several physical effects can contribute to this effect, such as the scattering induced attenuation which removes more high frequency information for the longer propagation paths and thereby reduces focusing.

6.4.3 Signal-to-Noise Ratio

Following the procedure for a noise model outlined in Section 6.2.2, H_N is calculated by using the same eight material realisations mentioned earlier, however removing the defect and sequentially exciting the array to obtain a baseline FMC. This produces an array image as shown in Figure 6.4c. From such an array image, similar to (Jie et al. 2013), we can calculate the Root-Mean-Square (RMS) of the pixel intensities at an image area of interest, to obtain a measure of the noise.

This noise data becomes useful when combined with the previously obtained H_s, as it enables a calculation of the $\overline{SNR}$, in this case performed as a function of aperture θ_p. The signal intensity is calculated from the peak pixel intensity (pk) within the array image I_s (the PSF in this case shown in Figure 6.4b). The noise area considered forms a box around the hypothetical defect which extends 10mm beyond the defect in the negative and positive, lateral and axial directions of the noise image I_N (see Figure 6.4c). SNR is subsequently calculated as shown in Equation 6.3 where $< >_{xy}$ denotes the mean across both x and y. To distinguish from the previous SNR, we shall label this the $\overline{SNR_2}$.

$$\overline{SNR_2}\,(\theta p) = \frac{|I_s(x,y,\theta p)|_{pk}}{\sqrt{\langle I_N(x,y,\theta p)^2\rangle_{xy}}} \tag{6.3}$$

Aperture

Figure 6.8 plots the mean image $\overline{SNR_2}$ and the standard deviation bars for 8 different random realisations of materials, for the three TPS cases as a function of half aperture angle. As was previously predicted by the aberrated PSF results, SNR becomes a function of depth. Moreover, it can be seen that SNR no longer behaves monotonically with respect to aperture angle; beyond an initial increase with aperture, it decreases for the widest aperture angles. Unlike the findings from Section 6.3 this suggests, for the strong scattering regime considered here, that the largest array aperture does not always optimise image SNR.

SNR Matrix

Taking the analysis one step further, and similar to the previously mentioned scattering matrix calculations (see Section 6.3.1), equivalent noise and SNR matrices are calculated. The noise matrix is calculated using the amplitude matrix procedure but instead of calculating the instantaneous amplitude at t_0, an average noise level is calculated, defined by the RMS value of a 1μs time-window surrounding t_0. The SNR matrix is then calculated by the division of the signal and noise scattering matrices, similar to Equation 6.3.

Figure 6.9 plots the (a) mean noise and (b) mean $\overline{SNR_2}$, this time as a function of receiver and transmitted array element, i and j, instead of the equivalent angle as it bears less significance when considering noise. The results are averaged for eight realisations of the 25mm depth TPS. Several observations can be made. Firstly, high intensity noise and low SNR can be found along the leading diagonal of Figure 6.9a and Figure 6.9b respectively, which corresponds to the pulse-echo elements exhibiting the worst SNR. This is a coherent multiple scattering effect (see e.g. (Ishimaru 1978)) previously known to manifest also in FE simulations of elastic wave scattering (Shahjahan, Rupin, et al. 2014).

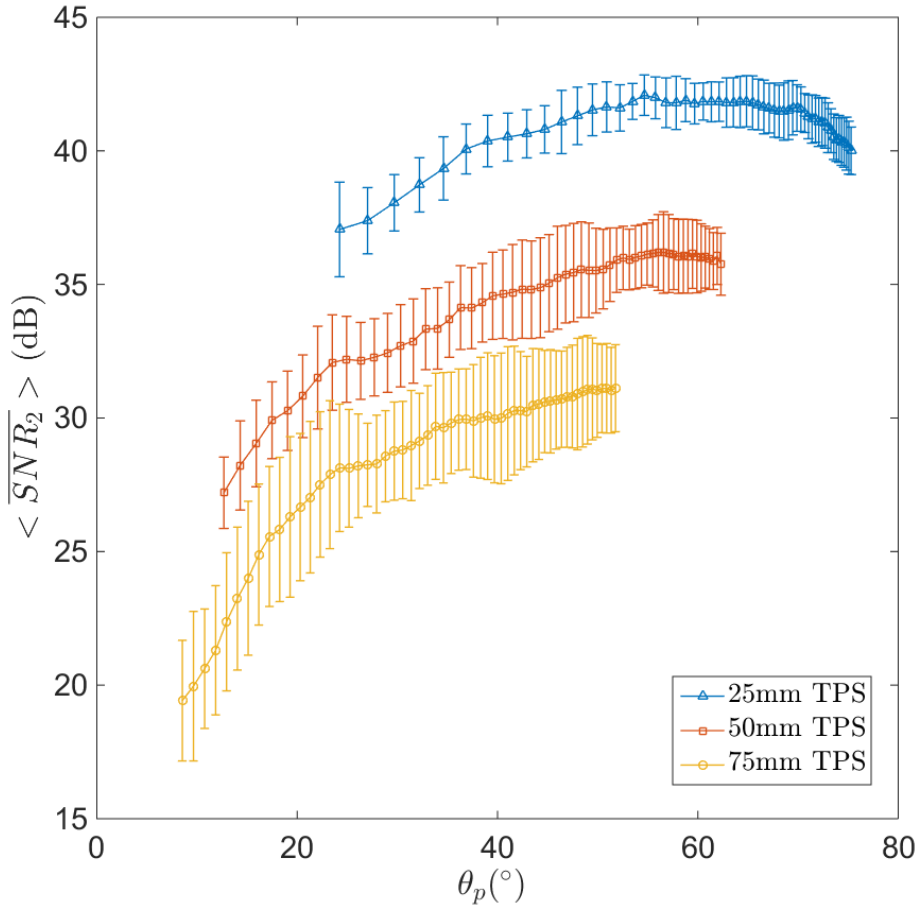

Figure 6.8: The mean SNR versus half aperture angle for 8 different realisations of a pollycrystalline material and for three true point scatterers at 25,50, and 75mm depth. Error bars indicate the standard deviation variation.

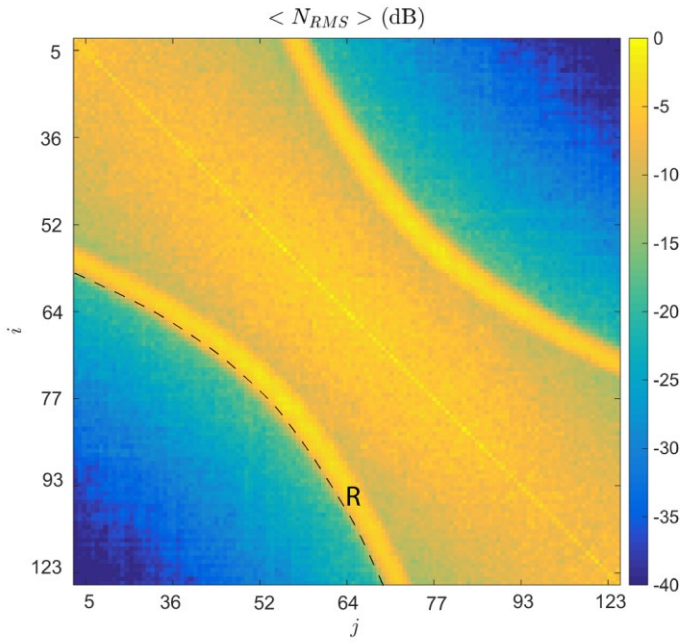

(a)

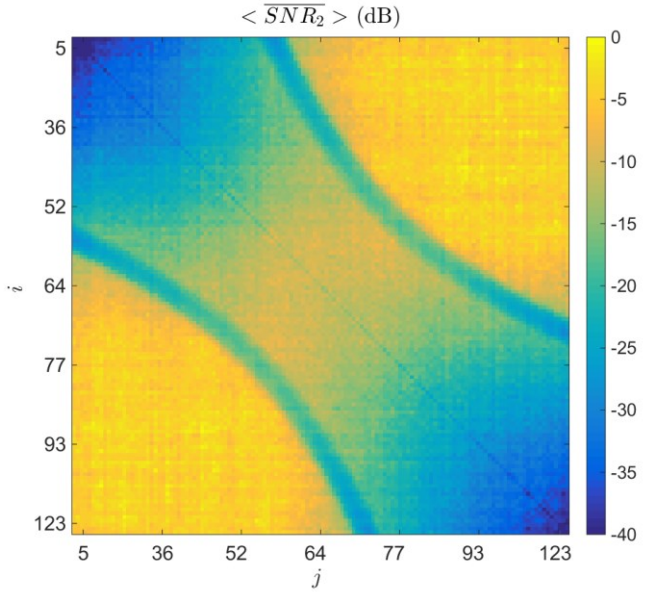

(b)

Figure 6.9: The mean (a) noise and (b) SNR averaged from eight realisations of a polycrystalline material, and calculated as a function of receiver and transmitter array element index (1 to 128). Both figures are normalised to the peak value in the image.

Due to reciprocity, a multiple scattering source-receiver path illuminates in both the direct and reciprocal directions, hence doubling the intensity of (only) the multiple scattering received noise. This explains why in practise twin-crystal probes have been observed to perform better than pulse-echo inspections relying on single probes to inspect highly scattering materials.

Another interesting feature is the presence of a (yellow) noisy region which comprises a band adjacent and roughly parallel to the leading diagonal in Figure 6.9a. Within the band we can make several observations. Firstly by assessing the change in noise along the leading diagonal and those adjacent but also parallel to it, we can see that the measured noise amplitude decays very slowly if at all, which shows that backscatter is a weak function of depth in this case. Contrarily, the noise seems to decay much quicker in the direction which is perpendicular to the leading diagonal (i.e. the anti-diagonal) which suggests a much stronger dependence on angle. Eventually, at large distances between receiver and transmitter element, there is a low-noise region (blue). However, towards the edge of the high noise region there is a significant rise in backscatter which shown by the quarter-circle bands (labelled R in Figure 6.9a) and corresponds to the Rayleigh wave (i.e. the Rayleigh wave in these cases arrived at previously defined t_0). Thus the high noise region contained within this band, hereby termed the backscatter envelope, can be broadly defined by the time-window corresponding to the arrival from transmitter to receiver of the Rayleigh wave.

The implication of the backscatter envelope is that it presents an opportunity to operate outside it. Namely, using large pitch-catch angles allows a longitudinal wave to arrive at the receiver before the majority of backscatter has arrived. This implies that pitch-catch configurations, using for example two arrays to separate the emitter and receiver, can be advantageous. At large aperture angles, it is possible that electrical noise sources become more significant as the received signal amplitudes, depending on the defect, can be significantly reduced. Such incoherent noise problems are much easier to overcome than coherent ones however, and hence in certain scenarios, it is possible that operating outside the backscatter envelope may provide benefits.

The next results section will undertake an experimental demonstration in order to illustrate this effect. First however, the next sub-section seizes the opportunity of having calculated a true SNR matrix, which offers a unique opportunity to optimise an image for the different transmitting and receiving combinations of the array.

Image Optimisation

In the previous section we have already measured the imaging SNR as the aperture is varied. Having obtained an SNR matrix however, which holds the SNR for the individual send-receive combinations of the array, we can also optimise our TFM image by calculating which individual combinations are accepted into the image to maximise SNR. Whereas this is practically impossible, as it requires a priori knowledge, it can be useful here to indicate which combinations of the array hold the best information, which should be proportional to SNR.

The optimisation routine adopted here is relatively simple. It starts by compiling a TFM image using a small aperture of 16 array elements, as the SNR should be relatively good for the elements closest to the array, and subsequently increases the aperture by compiling individual time-traces into the image one at a time, but only accepting each new candidate if they increase the image SNR. Once the full 128x128 send-receive combinations are considered, and by repeating this procedure eight times for the different realisations, we can plot an acceptance matrix which indicates how many times a particular send-receive combination of the array was accepted into the TFM image.

The resulting acceptance matrix is shown in Figure 6.10a, and one of the eight TFM images is shown in Figure 6.10b. As can be seen, the preferred combinations for maximising image SNR, are related to the SNR matrix, as expected. The comparison of the two TFM images, in in Figure 6.10b and Figure 6.10c, shows the improvement in SNR, albeit unachievable in practise.

Now we move to the next results section which aims to illustrate the backscatter envelope, experimentally.

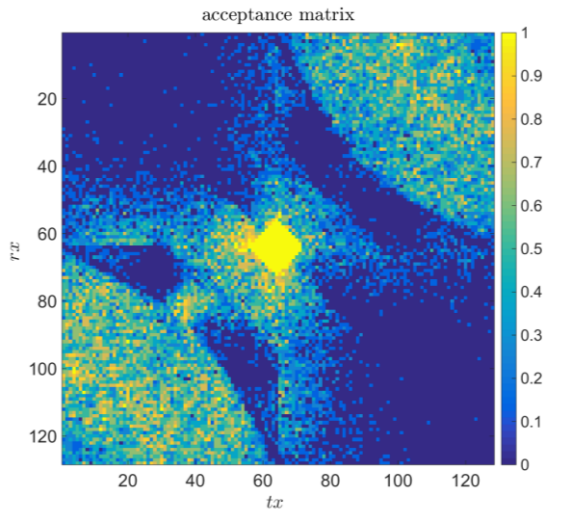

(a)

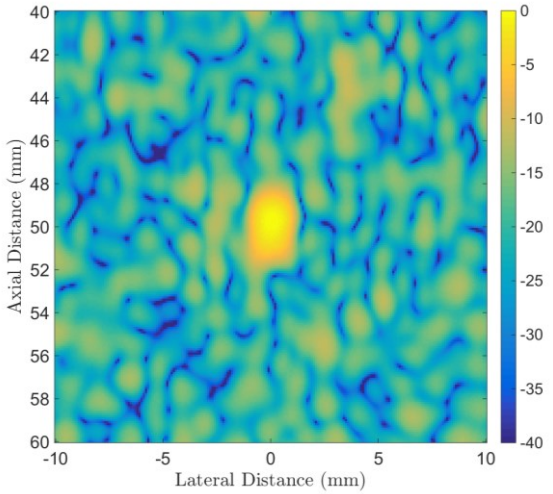

(b)

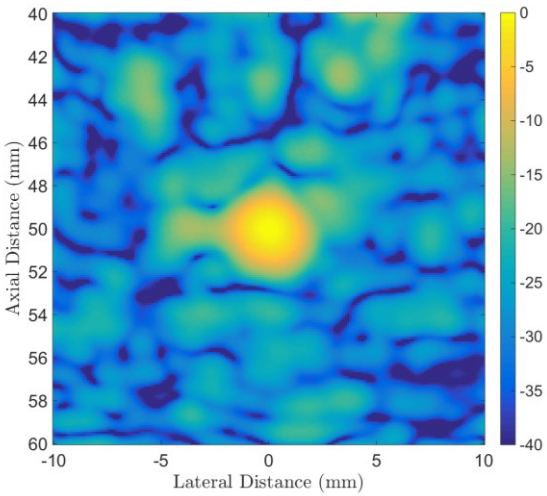

(c)

Figure 6.10. (a) Acceptance matrix (b) conventional TFM image for a 128 element aperture, and (c) optimised TFM image for a 128 element aperture.

6.5 Results III: Experimental Illustration of the Backscatter Envelope

The basis of this experiment is to acquire an ultrasonic array image with a single array, which operates well within the backscatter envelope, and compare it to the equivalent image acquired with a dual array setup (see Figure 6.11a). The dual array physically separates the arrays (in this case by 25mm) such that an opportunity exists to receive backwall signals before the majority of backscatter has arrived. The inspected material consists of a copper block and the array setup (LECOUER, France) uses two identical arrays with 64 elements and a 5MHz centre frequency (IMASONIC, France). The copper block contains no known defects, the target for the image is simply the backwall. In order to ensure electrical noise is suppressed, sufficient temporal averaging is applied to each FMC acquisition such that the noise floor is defined by coherent noise only.

Figure 6.11b plots the resultant TFM image for a single array and reveals that the backwall is detected with SNR of 50dB. The TFM image of a dual array (Figure 6.11c) conversely shows that the SNR has improved by 8dB due to a reduction in the received backscatter. This result is as was expected from the previous simulations and demonstrates that separating source and emitter can create an advantage to outrun the backscatter envelope with a longitudinal wave which corresponds to a feature of interest. The author would like to emphasize that such an opportunity will not exist for all practical scenarios, the aim is to illustrate and explain this effect which was previously only observed empirically. This demonstrates the utility of the FE modelling methodology to study such phenomena.

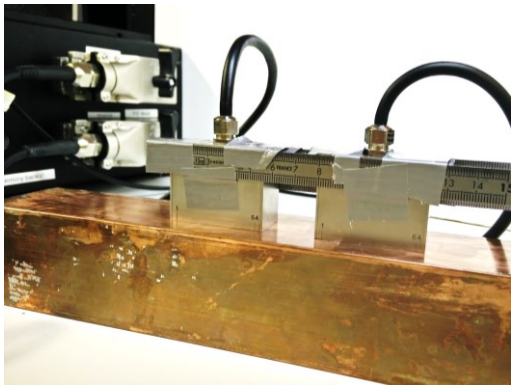

(a)

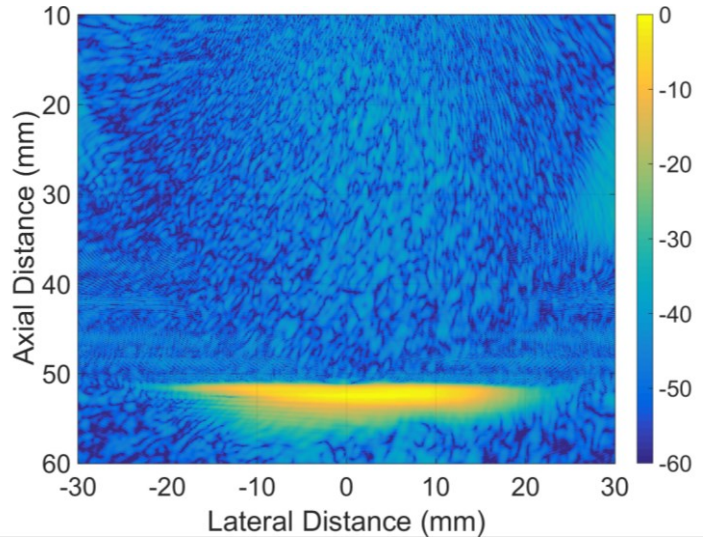

(b)

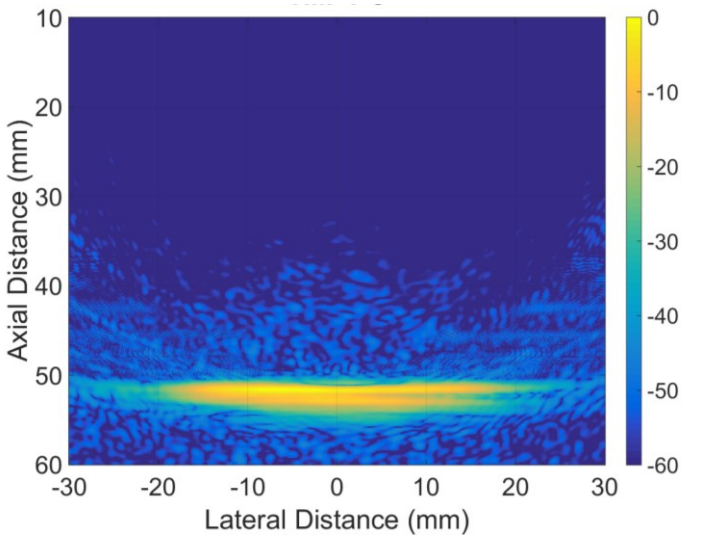

(c)

Figure 6.11. (a) Experimental setup of two arrays separated by 25mm in contact with copper. TFM images of the backwall for the (b) single array and (c) dual array setup.

6.6 Discussion: Spatial Averaging Theory

Combining the findings from Section 6.3 and 6.4, we can compare the prediction of SNR using the aberrated PSF (Section 6.4.2, Equation (6.2) to that obtained from the full signal and noise calculations (Section 6.4.3). In addition, a simplifying argument is proposed here, namely one which assumes that within highly scattering environments, the noise - albeit temporally coherent - is entirely uncorrelated between the different array elements (i.e. spatially incoherent). This would enable modelling of the noise by an averaging law which is simply the reciprocal of square-root n, where n represents the number of spatially independent time-traces used. Although the FMC holds N^2 time-traces, $n=N(1+N)/2$ due to reciprocity. Namely, in the absence of temporally random noise (e.g. electrical) only the half of the FMC (known as Half Matrix Capture (HMC)) holds unique information, as reciprocity dictates that a sender and receiver combination can be reversed and the received wave signal remains identical.

As an initial verification, Figure 6.12 plots the RMS image noise amplitude (denominator of Equation 6.3) against the averaging law. It shows good overall matching; the closest match occurring for the noise which stems deepest within the material (i.e. 75mm). Interestingly, this suggests that the grain noise is largely uncorrelated between array elements in the highly scattering regime simulated here.

6.6.1 Results versus Single Scattering Theory

The SNR comparisons are shown in Figure 6.13 a-c for the three TPS. The SNR_2 is shown with standard deviation bars, and both theoretical predictions, using the PSF and the averaging law are drawn on top. The correlations for both theories are for good at the 50mm and 75mm depths. As can be seen, the limitation of the single scattering theory, as outlined in Section 6.1.1, is that it only predicts monotonic SNR functions. Within the materials simulated, this is shown not to be the case, and hence the degradation of signal information through attenuation and phase aberration, and possibly dispersion effects is detrimental to the performance of the array.

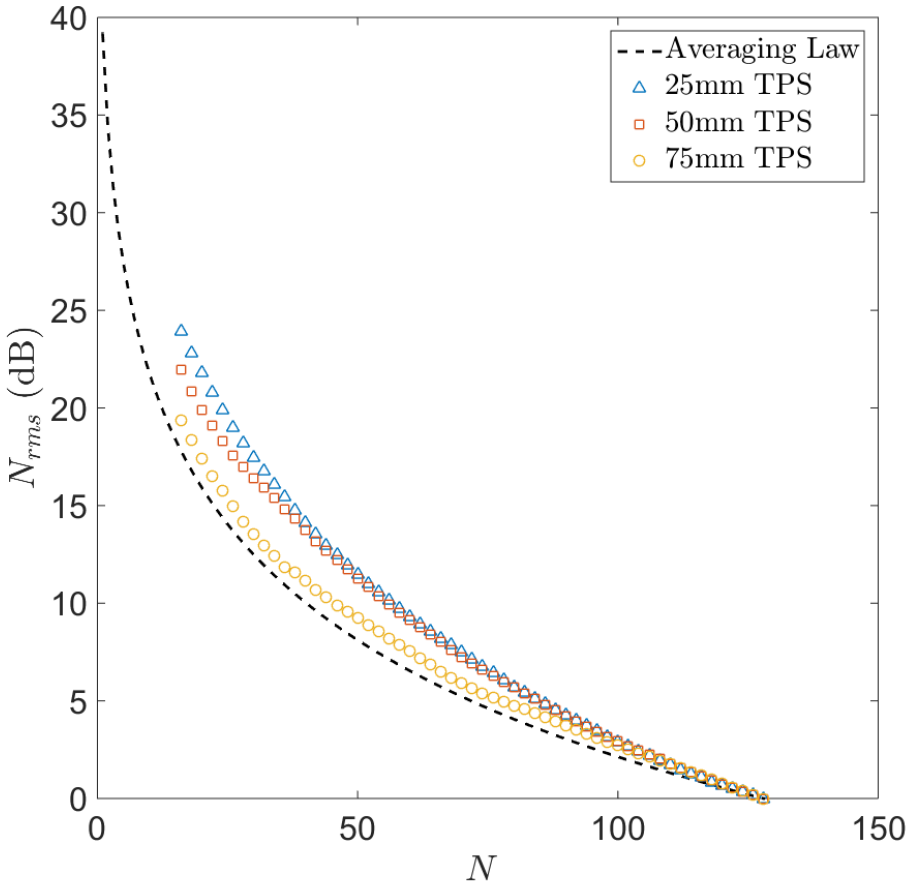

Figure 6.12: Comparison of the RMS noise calculated for a polycrystalline material as a function of the number of array elements N for three TPS cases at 25mm, 50mm, and 75mm versus the averaging law calculated from n independent time traces.

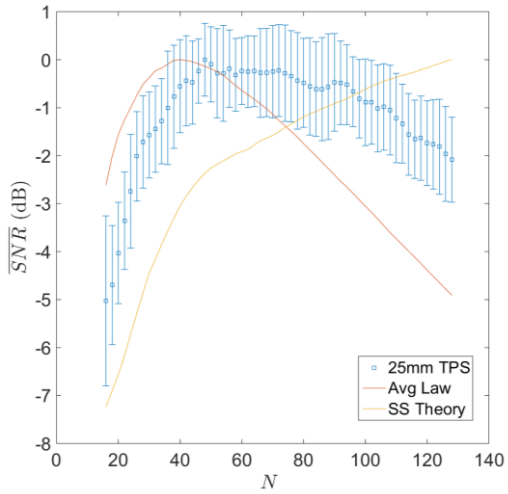

(a)

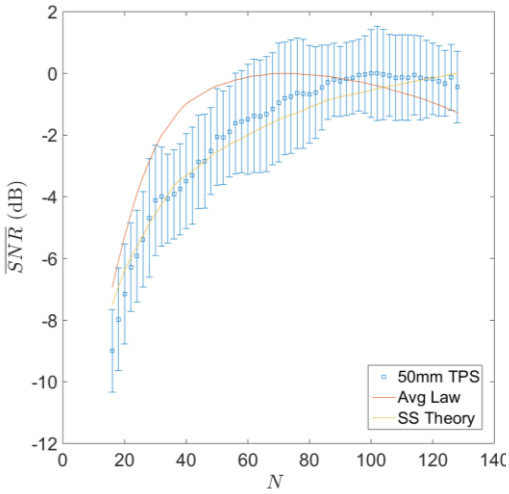

(b)

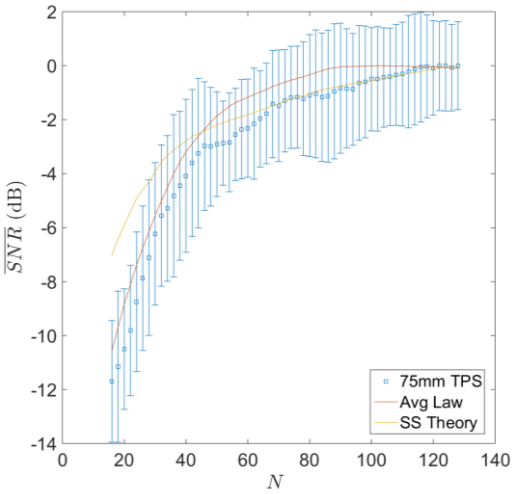

(c)

Figure 6.13: Comparison of the measured SNR, and two predicted SNRs. One uses the single scattering theory, and the second models the noise by an averaging law. The results are shown for a true point scatterer embedded at (a) 25mm (b) 50mm and (c) 75mm depth within a polycrystalline medium. Error bars indicate the standard deviation variation.

When using a single array, this notion suggests that for a given array element budget, it is preferable to spend it on a 2D configuration, i.e. rather than arrange the elements along a line, to extend them into a grid. A similar amount of spatial averaging will occur but more importantly, the signal information captured by the array will maintain a higher quality (amplitude).

6.7 Conclusions

This chapter has described a FE modelling framework to simulate ultrasonic arrays imaging within highly scattering, polycrystalline materials. Its utility is demonstrated by investigating the performance of array imaging, which is fundamentally limited by the onset of scattering noise. By comparison of multiple scattering simulations results to those of well-established single scattering models, it is found that FE modelling can provide interesting and new insights to find improved array configurations to inspect these materials.

Within highly scattering environments, the numerical simulations found that the maximum aperture does not necessarily maximise the SNR, which suggests that 2D arrays should offer improved performance over linear arrays. Furthermore, by demonstrating the existence of a backscatter envelope, it is also shown that in certain inspection scenarios, significant advantages can be derived from separating the emitting and receiving transducer. Lastly, it was found that treating the noise as spatially incoherent between the different array elements makes a good approximation in this strong scattering case.

Chapter VII

PSUEDO COLOURING

Ultrasonic imaging for NDE is limited by the challenge of detection, which relies on discriminating between objects based on their intensity. Whilst this works well in ultrasonically transparent media, in polycrystalline materials however, a host where scatterers are abundant, this is no longer the case. In such media, intensity information, as a means of interpreting an image, is compromised by the background of coherent microstructural noise. In a bid to improve this, it is suggested here to use pseudo-colouring to consider frequency information and distinguish objects based on their emitted frequency spectra. This approach exploits the frequency diversity; namely the difference in frequency dependence of the noise stemming from the material's microstructure, or backscatter, and that exhibited by the targets of interest: defects. Whereas established frequency diversity techniques exploit this additional information to reconvert it into amplitude data, colour enables the encoding of frequency and intensity information independently, yet still displayed in the same image. This chapter serves as an initial exploration of pseudo-colouring ultrasonic images for ultrasonic NDE of polycrystalline materials.

This work is published in QNDE proceedings 2014 [P4].

7.1 Introduction

Despite all the inspection difficulties which arise from a coarse grained microstructure, it also presents a unique opportunity. Namely, the microstructural noise typically exhibits a differing frequency dependence from that of the signals of interest: the defects. This phenomenon is termed frequency diversity and enables, instead of solely relying on intensity to discern between objects, to also consider differences in frequency spectra. Certain established signal processing techniques, such as SSP (Newhouse et al. 1982) (see Chapter 2) exploit this additional information but combine with the conventional intensity metric, to increase the ratio of signal to that of noise (SNR).

An alternative approach is proposed here, adding a new dimension to ultrasonic images and conjunctly displaying both types of information, by employing pseudo colouring. The objective is merely to facilitate interpretation of ultrasonic images and avoid parameter fine-tuning such as filter optimisation which would otherwise reveal the same defects. The latter, optimising filter settings, can present difficulties for practical scenarios where (1) there is no a priori knowledge of the defects, (2) there is insufficient knowledge of the grain morphology and (3) in extreme cases, the centre-frequency of defects changes with propagation distance, due to a highly frequency dependent attenuation, and hence also the optimal filter settings.

Although previously adopted by the medical community to indicate the direction of blood flow using the Doppler effect (Afruz et al. 2010; Jinxiu et al. 2007), this chapter serves as an initial exploration of pseudo-colouring for ultrasonic NDE. The concept is illustrated, first on simulation data which synthesizes frequency diversity, subsequently on Finite Element (FE) data of an array inspecting a polycrystalline material, and finally for experimental B-scans.

7.2 Concept

The concept of pseudo-colouring is illustrated here using simulated data which creates an idealized scenario. The model inspection layout consists of a 32 element array set in a

contact configuration to image three defects (see filled circles Figure 7.1a) within a polycrystalline material. The grain noise is modelled using the same model outlined in Chapter 4 (see Section 4.4.2) which generates coherent noise by superposing the separate ultrasonic responses of a random distribution of omnidirectional point scatterers (shown as small dots in Figure 7.1a). Frequency diversity, which may otherwise occur intrinsically between defect and grain, is synthesized by insonifing each with three-cycle tonebursts of different centre-frequencies, i.e. 2MHz and 5MHz respectively. In order to simulate a challenging inspection, the target point-scatterers (which are severely outnumbered) are assigned to exhibit only twice the scattering potential of a noise point-scatterer.

The ensuing computations involve calculating all the time-traces using Green's functions according to the procedure described in (Jie et al. 2013) to populate the Full Matrix Capture (FMC) array dataset. The FMC data are then used to process an image. In this example, the Total Focusing Method (Holmes et al. 2005) (TFM) is used as it provides a useful benchmark array imaging algorithm (Wilcox 2013).

The results convey a scenario where, when viewing the TFM image in Figure 7.1b, the three imaging targets cannot be detected when solely relying on the intensity of their reflection, as it fails to emerge from the background speckle. Merging frequency information (methodology described later) into the image however, through colouring as shown in Figure 7.1c, can reveal the three defects (exhibiting red and yellow colour values) which contrast in terms of their frequency content against the background noise (exhibiting blue and green colour values). Considering that the defects simulated here are relatively weak scatterers, the results indicate that differences in frequency content are a potentially sensitive metric to distinguish defects from grains.

7.3 Signal Processing Methodologies

Two methodologies were implemented as examples to explore pseudo-colouring, here named RGB and HSV for the colour models they rely on. Both methodologies are compatible with existing NDE imaging techniques including B-scans and array images such as TFM.

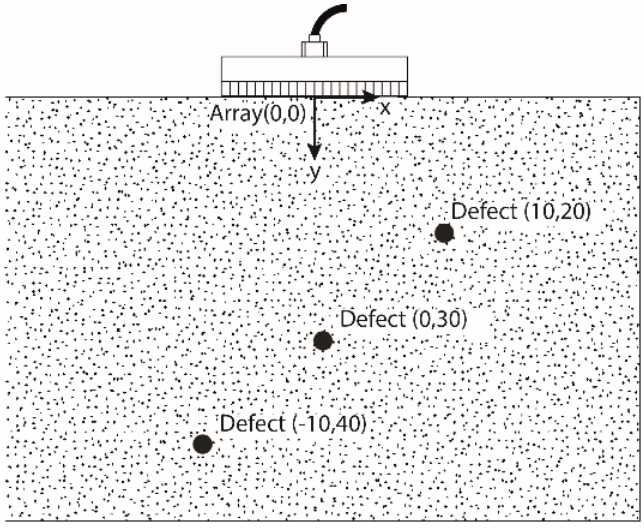

(a)

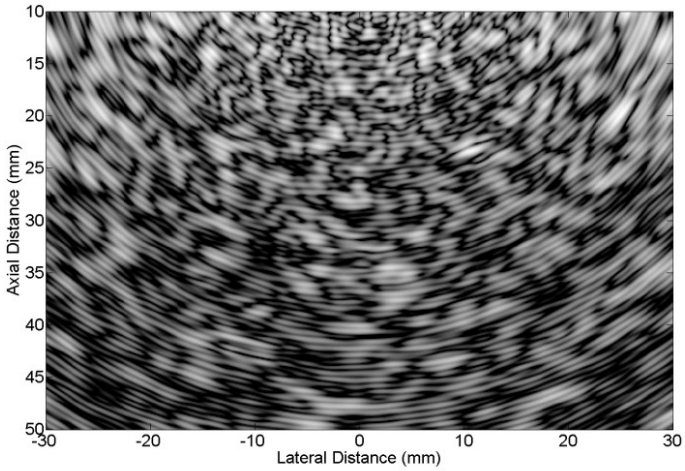

(b)

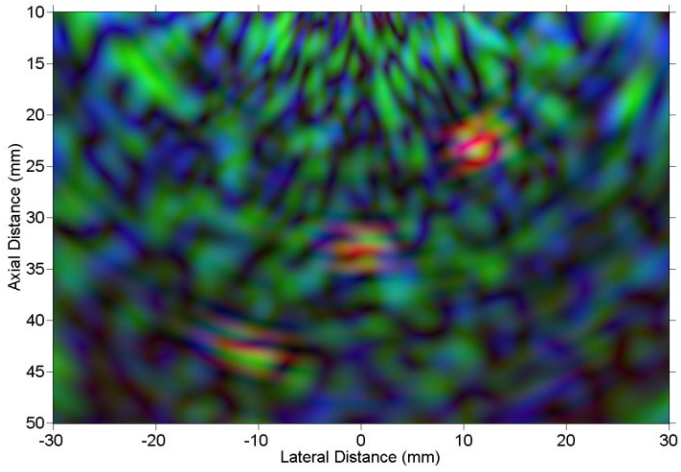

(c)

Figure 7.1: Concept of psuedo-coloring. (a) Representation of the model layout to synthesise frequency diversity for 3 imaging targets within a coherently noisy medium. (b) and (c) show the images obtained where (b) is the conventional TFM image (dynamic range of -20dB) and its TFM color equivalent (c) according to the RGB scheme.

7.3.1 RGB Colouring

The RGB colouring scheme is a simple approach to pseudo-colouring which aims to divide ultrasonic time-amplitude data into three separate bins of a low (LF), medium (MF), and high frequency (HF). Digital filtering is used to obtain a desired spectrum of the original data. The filters are implemented using Gaussian windows with -40dB at the cut-off frequencies.

When applying the RGB colouring scheme to an image, the LF, MF, and HF data are each used to process a separate image. The three resulting sub-images are subsequently combined according to the 3D colouring matrix (shown in Figure 7.2a). For example, every pixel in the image now has three values, representing an intensity at LF, MF, and HF, which corresponds to red, green, and blue respectively.

Although RGB colouring is useful to reveal differences in frequency content, it is not practical as it requires parameter fine-tuning, similar to its substitute, filter optimisation. Another algorithm, HSV, is explored to avoid this step and automatically find the appropriate colour settings.

7.3.2 HSV Colouring

HSV colouring starts by converting a conventional ultrasonic time-domain signals (A-scan) into a new signal which quantifies a centre-frequency for each time sample *(t)* of the input signal. This involves first transforming the time-domain trace into a 2D time-frequency domain signal. Although Short-time Fourier transforms are typically used for this purpose, they require fine tuning of the window time length to balance frequency resolution and time step sensitivity, and hence to circumvent this, wavelets were employed which produce a comparable end-result. Equation 7.1 shows the general wavelet transform where ψ is the wavelet shape function, n is scale parameter, b is position, $f(t)$ is the original function being transformed, and W is the resulting wavelet coefficient.

$$W(n,b) = \int_{-\infty}^{\infty} f(t) \frac{1}{\sqrt{n}} \psi^* \left(\frac{t-b}{n} \right) dt \tag{7.1}$$

The resulting 2D signal $W(n,b)$ is then converted back into a 1D time-domain signal by calculating the spectral centroid, which represents a median of the perceived spectrum, for each time sample *(t)* according to Equation 7.2 where F_c is the centre-frequency, *s(n)* is the scale number. Intensity and frequency are then mapped onto a 2D colouring matrix as shown in Figure 7.2b.

$$F_c(b) = \frac{\sum_{n=0}^{N-1} W(n,b)s(n)}{\sum_{n=0}^{N-1} s(n)} \tag{7.2}$$

7.4 Simulation Results

In contrast to the previously simulated data in Section 7.2, now a FE model is used to investigate frequency diversity and its potential merits for pseudo-colouring. As shown in Chapter 5 and Chapter 6, FE modelling benefits from a high fidelity and incorporates the frequency dependent increases in scattering activity, such that it is well suited to investigate the merits of frequency diversity. The complete account of the modelling procedure is given in Chapter 5.

Similarly to Chapter 6, a 2D FE model (shown in Figure 7.3a) is used to simulate a 5MHz array with a 32 element aperture and 0.6mm element pitch. A rigid point scatterer is introduced (not a TPS) as an imaging target by constraining displacement of six nodes in a hexagonal configuration with a diameter one twentieth of the wavelength, located centrally below the array at a depth of 25mm. The material is Inconel 600 with an average grain diameter of 500μm.

A conventional TFM image (see Figure 7.3b, -20dB dynamic range) shows that the defect is difficult to detect based solely on intensity. The pseudo-colour version of the TFM image (see Figure 7.3c) is compiled employing the RGB scheme with values of 3, 5, and 8MHz for the LF, MF, and HF centre-frequencies, each with a bandwidth of 4MHz. Figure 7.3c indicates that the defect maintains its high intensity alongside the speckle, but has now become more apparent as the other high intensity pixels which were previously obscuring the defect are coloured differently due to the differences in their reflected frequency spectra.

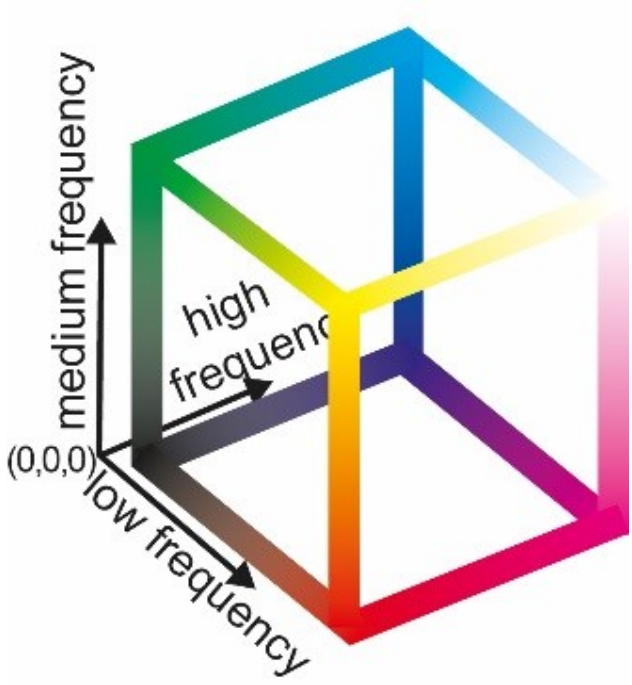

(a)

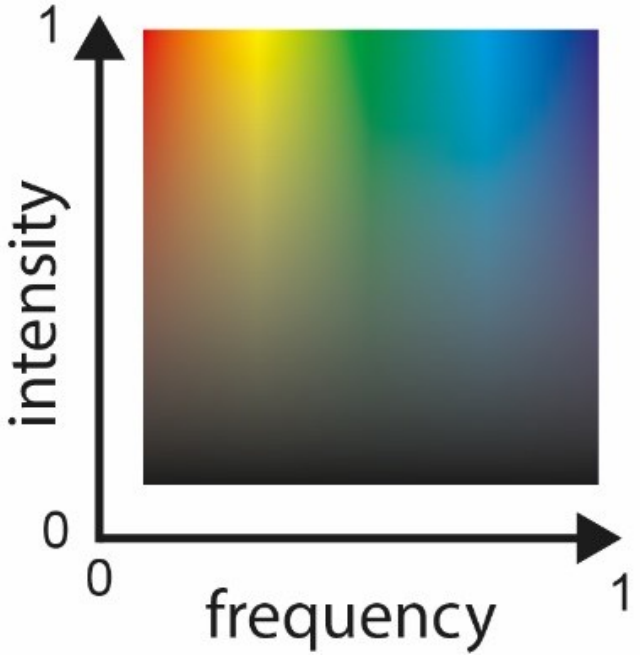

(b)

Figure 7.2: Schematics of the RGB (a) and HSV (b) colouring scheme.

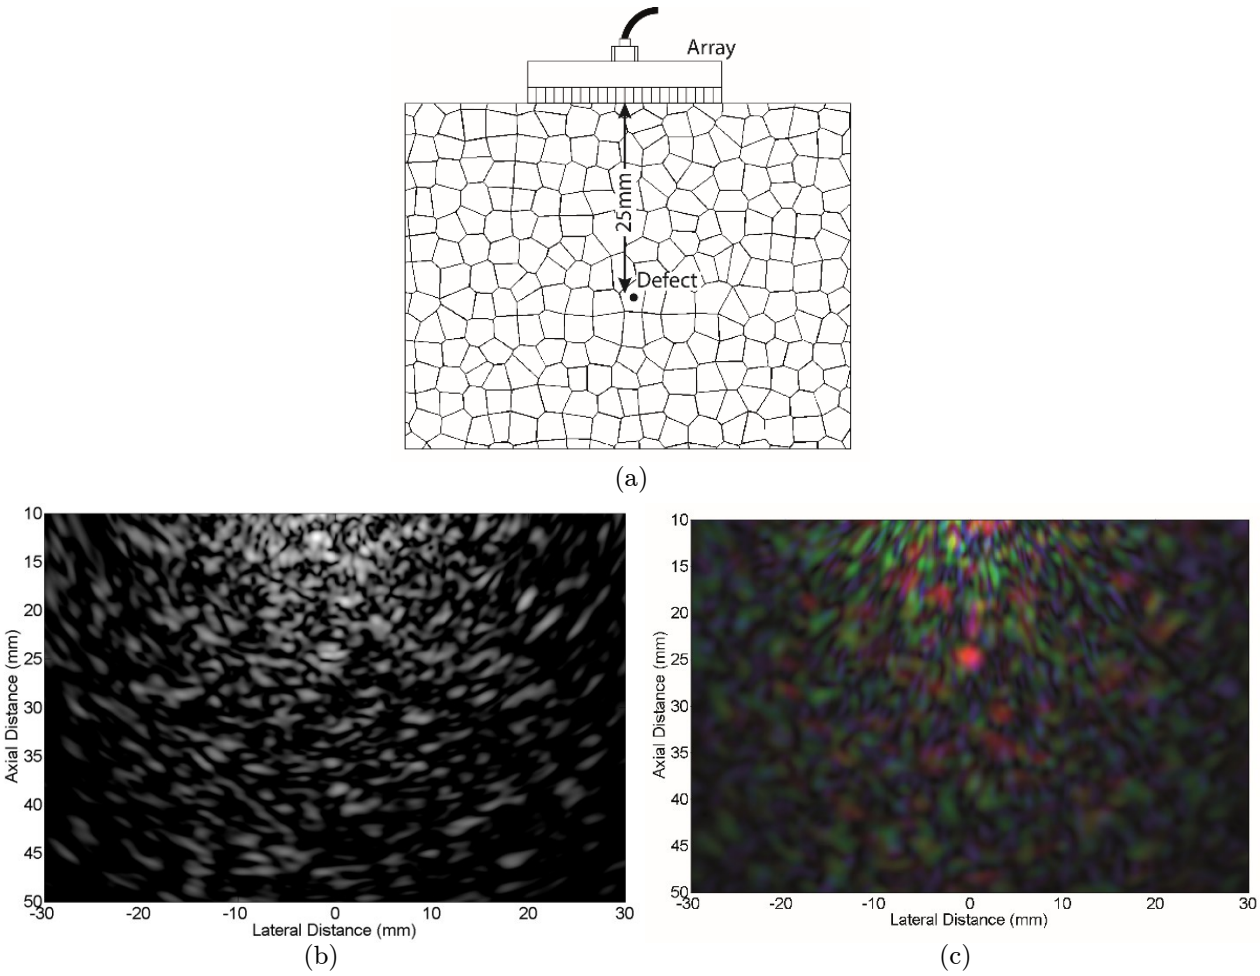

Figure 7.3: (a) Representation of the Finite Element model layout for a 32 element array inspecting a polycrystalline material with an average grain diameter of 500µm (grains not drawn to scale) and containing a point defect centrally below the array. (b) and (c) are the resulting array images, where (b) is the conventional TFM image and (c) is the RGB psuedo coloured version according to the RGB scheme.

Figure 7.4a aligns the resulting time-traces to form a B-scan, where a time-span which is in between the front and backwall of the sample is shown. It can be observed that the SDH, which is located at 12.5mm, is difficult to detect due to the backscatter. An RGB pseudo-colour version is compiled through setting the LF, MF, HF centre-frequencies at 5, 11, and 17MHz respectively with a bandwidth of 6MHz. Figure 7.4b shows that pseudo-colouring facilitates detection by considering frequency changes within the received signals. The experiment demonstrates that the grain noise has a relatively high centre-frequency when compared to the SDH's signal which has a lower centre-frequency as seen from the difference in colour (blue versus green).

The HSV pseudo-coloured version is shown in Figure 7.4c. Similarly to the RGB version, it manages to find the differences in frequency between the noise and signals, albeit less convincingly. Unlike RGB colouring however, the HSV is a more promising colouring scheme for future adaptations as it does not rely extensively on parameter fine-tuning and therefore does not require prior knowledge of the defect.

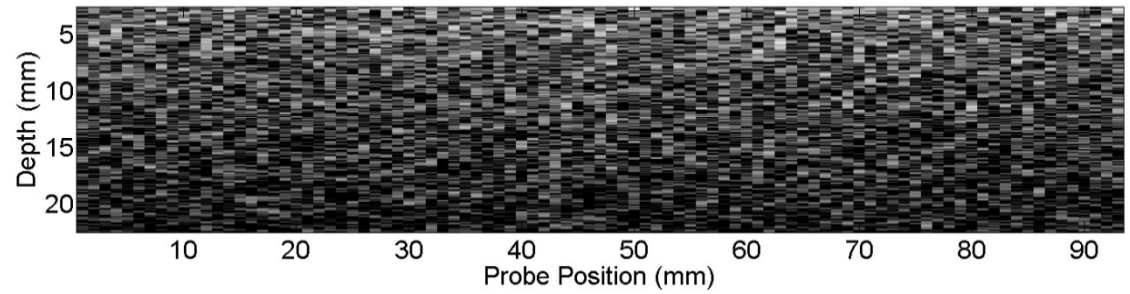

(a)

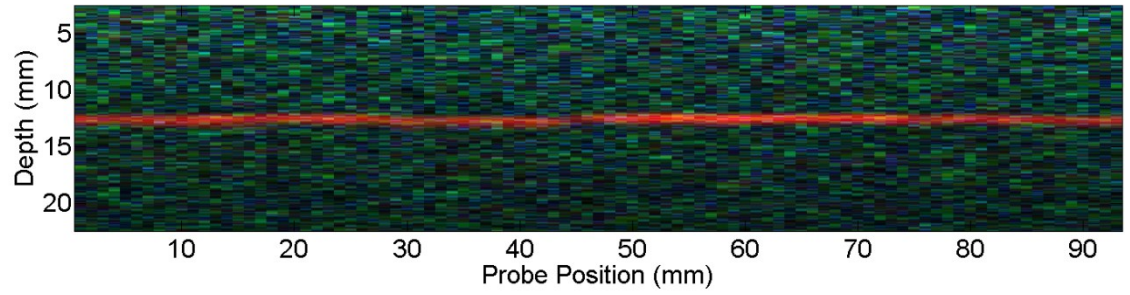

(b)

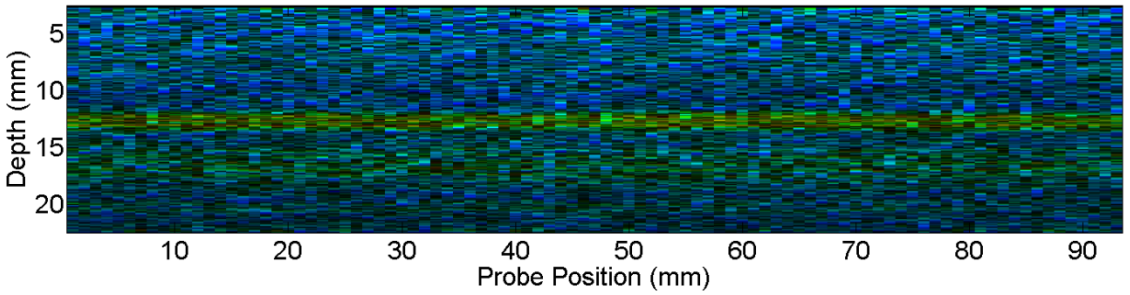

(c)

Figure 7.4: Experimentally obtained images comparing a conventional B-scan (a) to the RGB (b) and HSV (c) psuedo-colored version, all stemming from the same input data. A slight shadow can be seen in (c) which could be an artefact or a creeping wave from the SDH.

7.5 Discussion and Conclusions

Two example methodologies to produce pseudo-colour images for NDE were explored. Evidence of frequency diversity between grains and idealised defects was found by experimental results and FE simulations.

The results suggest that frequency diversity, incorporated alongside intensity through pseudo-colouring, shows promise to aid detection of defects in ultrasonic images of polycrystalline materials. It is however not necessarily impossible to detect these defects using conventional intensity based images. In fact, with the right filter settings conventional images would reveal similar imaging targets. When inspecting polycrystalline materials however, where information about the defects and materials parameters such as grain size are often unknown, it becomes difficult to know what these optimal parameters are. Therefore pseudo-colouring offers a potentially useful tool for improving interpretation of ultrasonic images.

Chapter VIII

SPATIAL FILTERING

Spatial filtering is briefly introduced as a possible approach to reduce backscatter. A numerical methodology is developed, and demonstrated on some illustrative cases. Preliminary FE tests confirm the contribution of shear waves within the scattered wave field, but when considering practical inspection scenarios, its potential benefit seems to be limited.

8.1 Introduction

Having recognised the presence of shear waves within the scattered wave field, it seems opportune to exploit this contrast with the coherent wave signal, which comprises solely a longitudinal wave. Decreasing the influence of shear waves could thus potentially reduce the contribution of backscatter, and thereby improve SNR. Achieving this in practise is not straight-forward, as both shear and longitudinal waves appear identical when viewed in the time-domain, whereto we are usually confined when observing ultrasonic waves. However, both modes exhibit a distinctive difference in wavelength, and thus as an initial investigation, we can take advantage of a simulation environment, such as the aforementioned FE, where spatial observations are possible.

This chapter proposes a methodology to filter a particular wave mode, shear or longitudinal, in attempt to provide insights into the potential gains of such an approach for ultrasonic NDE. The technique is initially demonstrated on simple yet illustrative cases, before considering scattering scenarios. The ultimate objective is to quantify the potential improvement to SNR by filtering longitudinal waves, which is demonstrated for the inspection scenario of an ultrasonic array. This preliminary investigation explores whether in principle a benefit exists, a means to achieve this practically has not been investigated and would have to be pursued in future work.

8.2 Spatial Filtering

Spatial filtering, in the sense of separating shear and longitudinal waves, can be numerically achieved in multiple ways. Perhaps the most obvious method is to rely on spatial Fourier Transforms. For the FE models considered here, this would be extremely computationally expensive, and perhaps excessive as we merely require - albeit spatially filtered - time-domain information (e.g. A-scans). Instead, the methodology proposed here relies on the previously mentioned True Point Scatterer (TPS) concept (see Chapter 6).

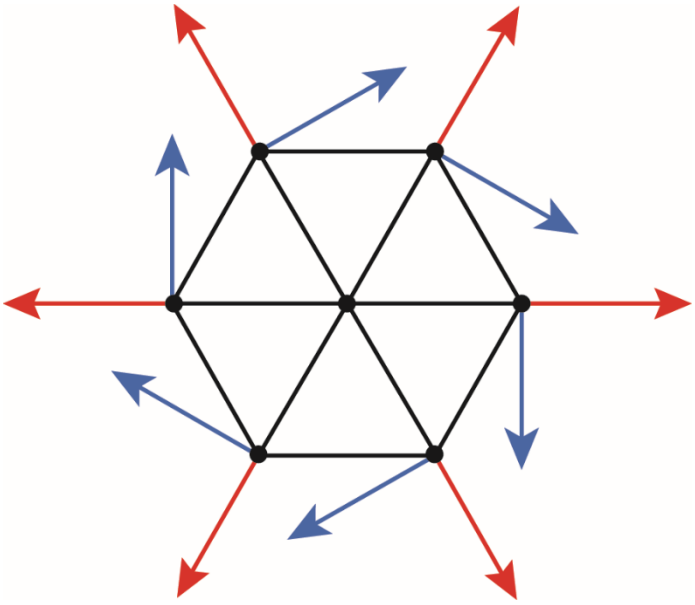

Figure 8.1: Concept of Longitudinal (red) and Shear (blue) wave Spatial Filter showing the unit vectors for each node.

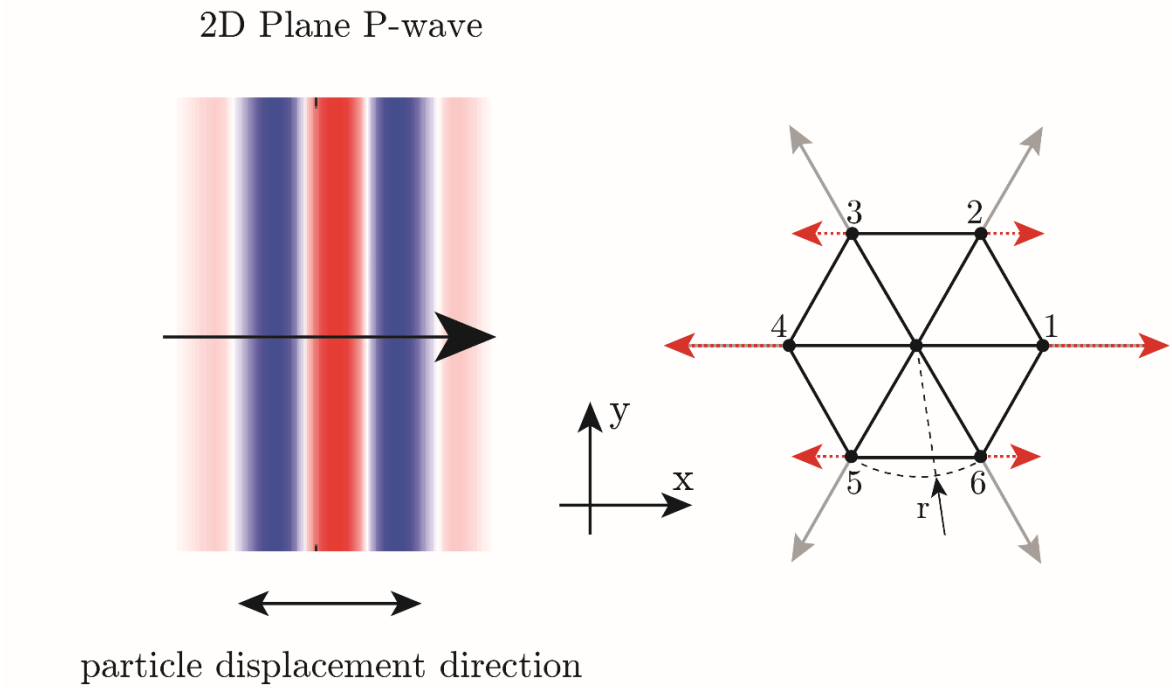

Figure 8.2: Schematic of an longitudinal plane wave incident onto a longitudinal spatial filter, indicating the x-displacement components (red arrows) from their original vector (gray). The radius, r, is labelled.

Whereas the TPS procedure was initially proposed to excite a pure omnidirectional longitudinal wave; the principle of reciprocity suggests the same calculation can be applied in reception, to purely receive longitudinal waves. This concept is termed a Longitudinal wave Spatial Filter (LSF), and can also be extended to a Shear wave Spatial Filter (SSF) as illustrated in Figure 8.1.

The following section provides the analytical background of the proposed spatial filter.

8.2.1 Analytical Validation

Before undertaking any FE simulations, let us consider an analytical case to undertake the spatial filtering calculation in known circumstances. For example, assume a longitudinal plane wave propagating in the positive x-direction (as shown in Figure 8.2). As the propagation and polarisation direction is known in this case, the x-displacement is equal to the longitudinal wave displacement, U_{L0}. The time dependent x-displacement, U_{x_0}, at a location, x_0, is given by Equation 8.1. As stated earlier, by relying on our a priori knowledge of the wave, U_{x_0} is equivalent to U_{L0} in this case.

$$U_{L0}(t) = U_{x_0}(t) = A_0 e^{i(kx_0 - \omega t)} \tag{8.1}$$

The purpose of the spatial filter is to obtain U_{L0} or U_{S0}, without relying on a priori knowledge as this will not be available once we consider a scattered wave field, where a multitude of waves propagate in random directions with unknown polarisations. The proposed spatial filter probes the two dimensional wave field, in this case at six locations surrounding x_0 in a hexagonal arrangement (an artefact of using a structured triangular FE mesh). The ensuing calculation involves projecting the displacement of each node, 1 to 6, in both the x and y direction, according to the respective unit vector shown in Figure 8.2. The resulting displacement is denoted by U_L.

In the case of the aforementioned longitudinal wave, displacement is restricted to the x-axis, and thus the vectorial calculation reduces to the summation of the nodal displacements, $U_{x,1}$ to $U_{x,6}$, with the appropriate weighting function. For our purposes at this stage, it is sufficient and equally illustrative to perform the calculation for U_L (shown

in Equation 8.2) by simply considering the wave equations at locations x_1 and x_4, neglecting the node pairs 2, 3 and 5, 6.

$$U_L(t) = U_{x,1}(t) - U_{x,4}(t) = A_0 e^{i(kx_1 - \omega t)} - A_0 e^{i(kx_4 - \omega t)} \tag{8.2}$$

To compare how the spatially filtered result, U_L, relates the known longitudinal displacement, U_{L0}, we can redefine $x_1 = x_0 + r$ and $x_4 = x_0 - r$, where r is defined in Figure 8.2. After rearranging, this yields Equation 8.3:

$$U_L(t) = A_0 e^{i(kx_0 - \omega t)}[2i\sin(kr)] = U_{L0}(t)[2i\sin(kr)] \tag{8.3}$$

This reveals that the spatial filtering procedure induces a change in amplitude, $\Delta|U_L(t)|$, and a change in phase, $\Delta\angle U_L(t)$, in comparison to U_{L0}, as described in Equation 8.4.

$$\Delta|U_L(t)| = 2\sin kr \approx 2kr\,, \quad \Delta\angle U_L(t) = \pi/2 \tag{8.4}$$

This provides several observations. Firstly, and perhaps intuitively, the calculated amplitude for U_L decreases as r becomes smaller, as the phase difference between $U_{x,1}$ and $U_{x,4}$ reduces. The dependence on k indicates that lower frequencies are amplified as they comprise larger wavenumbers. Lastly, the phase angle indicates that there is a constant phase shift which is independent of the variables.

The same analysis can be undertaken to demonstrate the parametric dependence of the SSF, which produces the identical results, shown in Equation 8.5.

$$\Delta|U_S(t)| \approx 2kr\,, \quad \Delta\angle U_S(t) = \pi/2 \tag{8.5}$$

Now that the behaviour of the spatial filter is established, some simple numerical simulations are computed to demonstrate its spatial filtering capabilities.

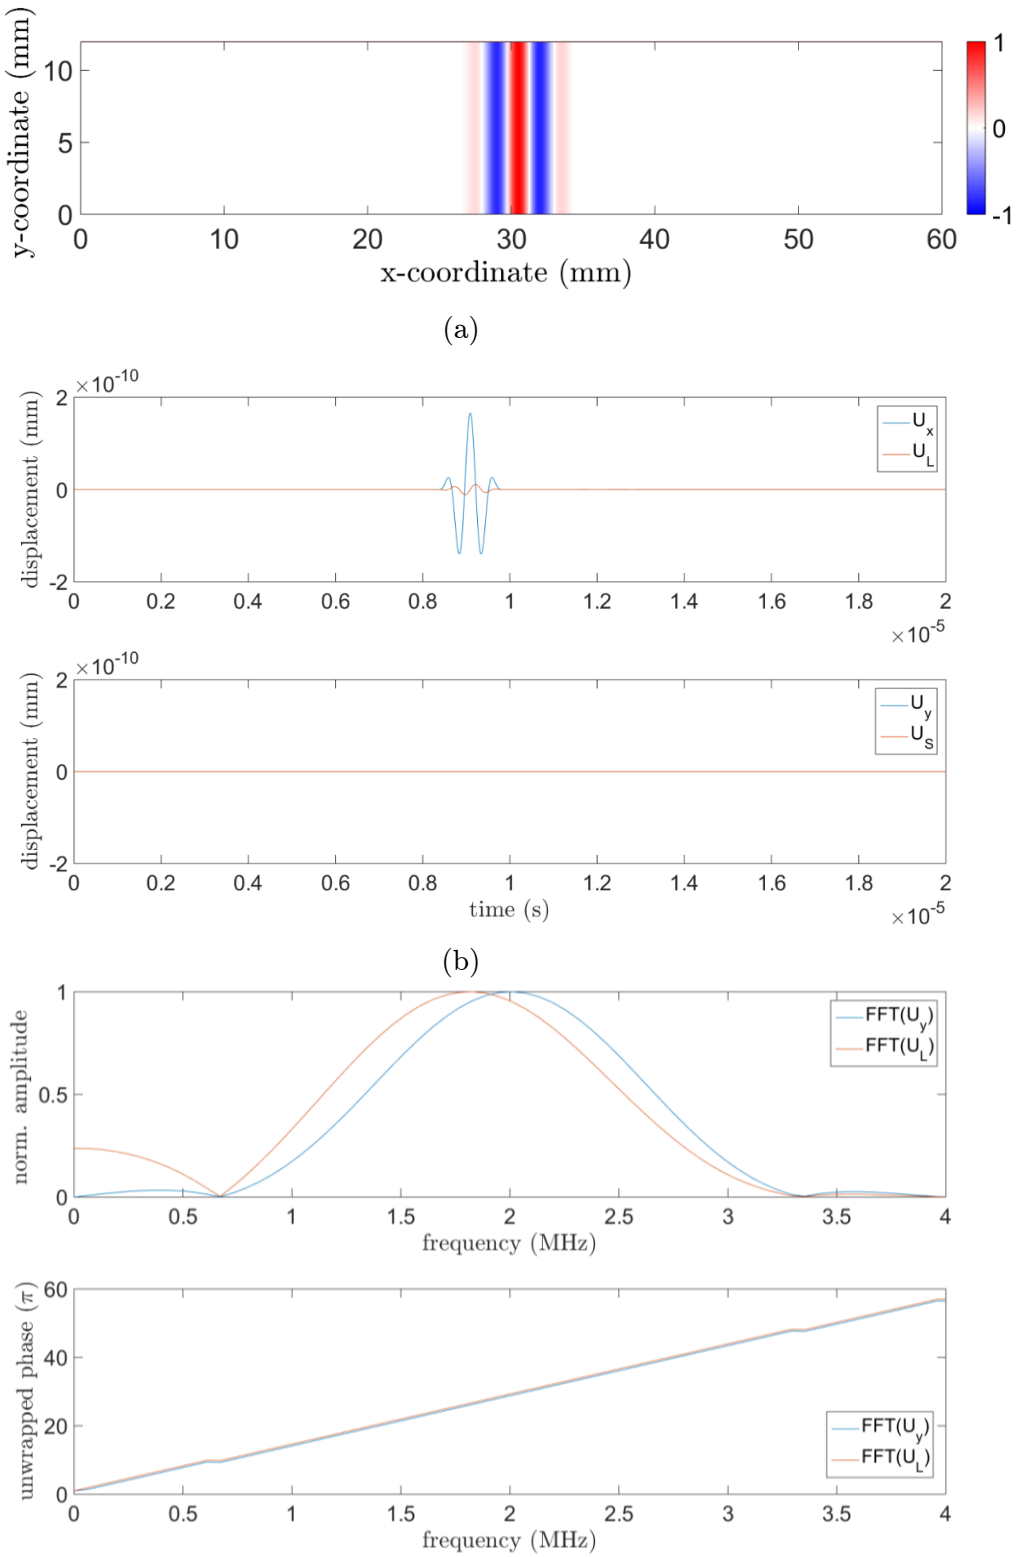

Figure 8.3: Results for longitudinal wave excitation (a) x-displacement field showing a plane wave (b) time-trace comparisons for U_{x0} vs U_L and U_{y0} vs U_s (c) Fourier analysis of previous time-traces.

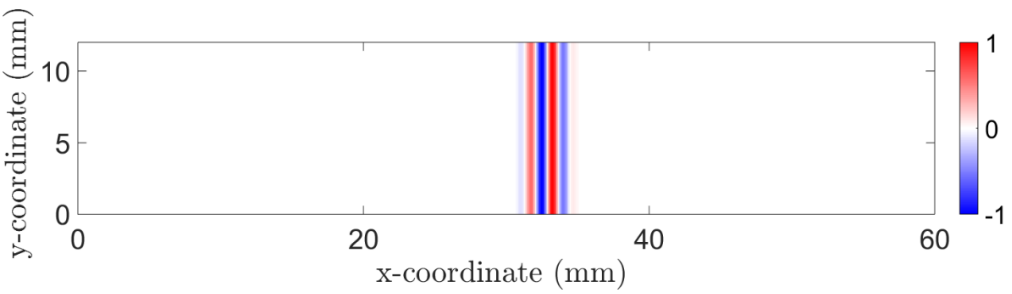

(a)

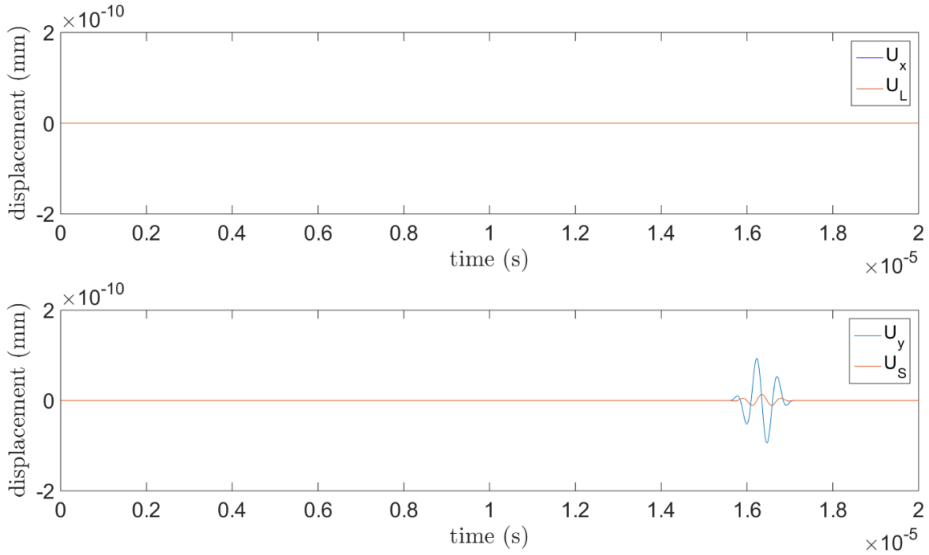

(b)

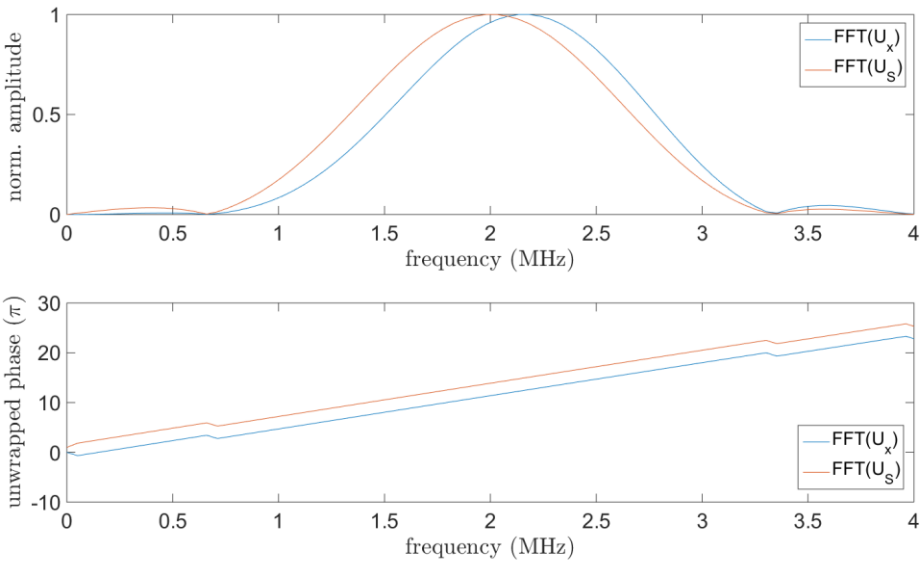

(c)

Figure 8.4: Results for shear wave excitation (a) y-displacement field showing a plane wave (b) time-trace comparisons for U_{x0} vs U_L and U_{y0} vs U_s (c) Fourier analysis of previous time-traces.

8.3 Numerical Results

The aforementioned analytical scenario is now simulated numerically; plane waves are generated within isotropic media using the FE procedure outlined in Chapter 5. A spatial filter is embedded centrally within the material (at x=30mm, and y=6mm) to calculate a resultant longitudinal and shear wave displacement, U_L and U_S, using the respective LSF and SSF procedures. Similarly to the previous section, within this controlled environment, the spatially filtered result can be verified against the known result, U_{L0} and U_{S0}. In this case, when either a longitudinal or shear wave propagates in the positive x-direction (see Figure 8.2), U_{L0} and U_{S0} are given by the x and y displacement respectively, U_{x0} and U_{y0}, of the central node within the hexagonal configuration.

8.3.1 Numerical Validation

The medium is defined to represent a typical elastic material of interest, with E=216GPa, ν=0.3 and density ρ=8000kgm^{-3}. The plane wave consists of either a longitudinal or shear wave, excited by a three cycle toneburst with a 2MHz centre frequency. For ease of implementing the hexagonal spatial filter within the FE model, a structured mesh with triangular elements is used.

Figure 8.3a displays the x-displacement field for a longitudinal plane wave propagating within a slab of isotropic material. The time-domain spatial filtering results are plotted in Figure 8.3b and reveal that the LSF has correctly identified the longitudinal wave, whilst the SSF has filtered it out (a time-trace of nulls is produced) as intended.

Figure 8.4 repeats the procedure in the previous figure but for the shear plane wave case, as seen in Figure 8.4a. This verifies that the LSF correctly filters out an incident shear wave, whilst the SSF identifies it, as shown in Figure 8.4b.

The spatial filter also seems to behave as predicted by Equation 8.4. The time-domain signals in both Figure 8.3b and Figure 8.4b reveal a reduction in amplitude, as predicted by the dependence on r. Furthermore, transformation into the frequency domain, shown in Figure 8.3c and Figure 8.4c for both cases, similarly confirm a slight shift in centre-

frequency due to the dependence on k. Lastly, the unwrapped phase shows the phase shift to be constant with frequency, and thus independent of k as expected.

The results so far suggests the spatial filtering technique is operating as intended. A question remains however concerning the spatial sampling requirements. Until now, we have adopted what was believed to be a sensible value for r, which determines the spatial sampling frequency. Similarly to time-domain measurements, it seems clear that the spatial sampling must at the very least satisfy the Nyquist criterion. A convergence study is undertaken next to determine the value for r which provides an accurate but also a computationally efficient solution.

8.4 Convergence Study

In contrast to the previous studies, scattering is now introduced to create a wave field containing both wave modes. The polycrystalline material considered here consists of Inconel 600 with an average grain size of 500μm. The resulting scattered wave field for a 2MHz longitudinal plane wave is shown in Figure 8.5a. An LSF is embedded within the material at x=20mm and y=6mm, such that a U_L time-trace is obtained as shown Figure 8.5b.

The metric to measure convergence is chosen in function of our eventual objective, namely to measure SNR. As shown in in Equation 8.6, SNR is defined as the ratio of the peak (pk) of the signal (shown in Figure 8.5b at 2.5μs) and the RMS value of the backscatter noise (time-window shown in Figure 8.5b). The error E_{SNR} is the SNR normalised against a converged value, SNR_c, which is obtained from the finest spatial sampling considered, which in this case was approximately r=λ/500.

$$SNR = \frac{S_{pk}}{N_{rms}}, \qquad E_{SNR} = \frac{SNR}{SNR_c} \tag{8.6}$$

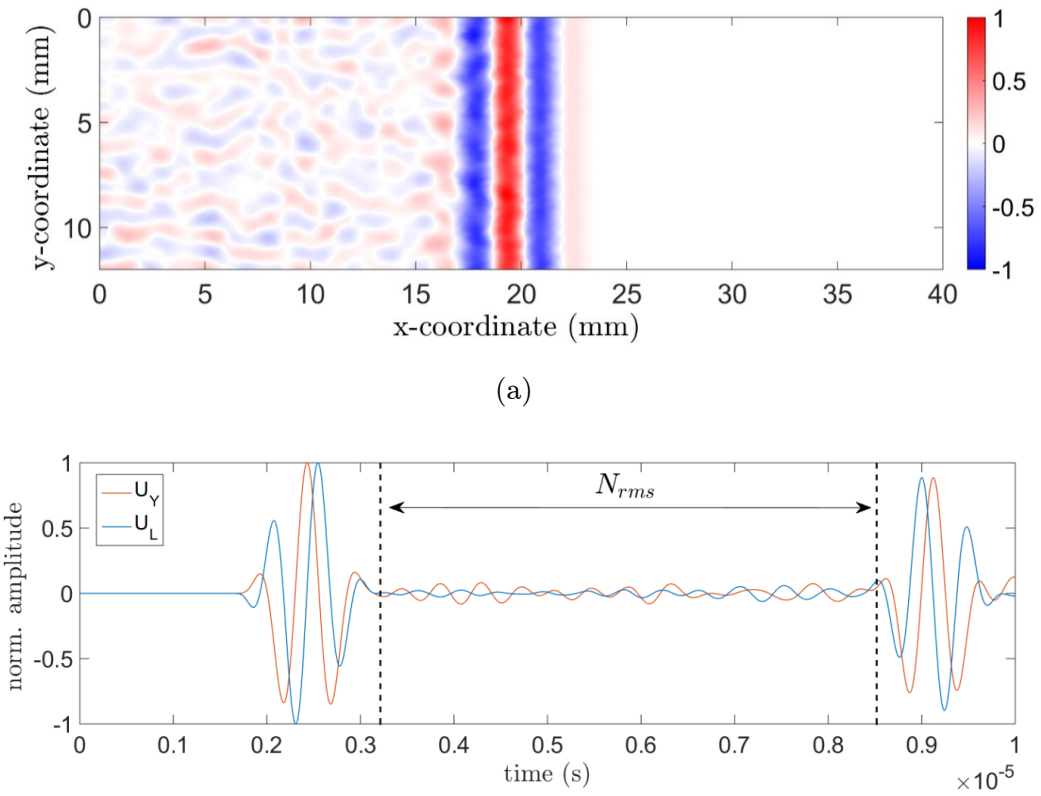

Figure 8.5: (a) longitudinal plane wave propagating in a slab of polycrystalline material. (b) The time trace obtained from the y-displacement and the filtered longitudinal displacement where shear waves are supressed.

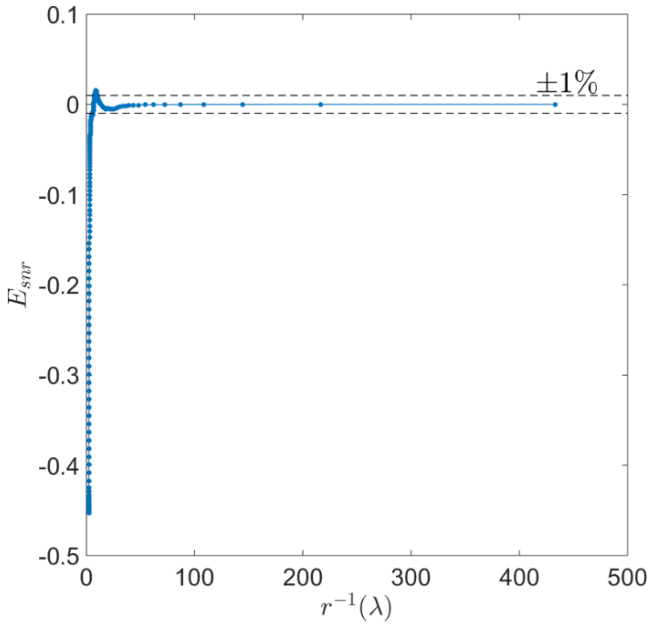

Figure 8.6: Convergence of error in SNR versus radius of point receiver quantified by portion of a wavelength.

Figure 8.6 shows the convergence of the normalised SNR error as the spatial filter radius is decreased. Convergence can be seen to be achieved to within 1% at a radius of approximately one twentieth of a longitudinal wavelength, or one tenth of a shear wavelength.

8.4.1 Longitudinal Wave Ratio

It is worthwhile noting that the potential improvement of the considered approach is fundamentally limited by the shear-to-longitudinal wave balance within the scattered field. For example, if 50% of the scattered field consists of shear waves, in a best case scenario, the filtered signal would exhibit a 6dB improvement in SNR. Although this has not been discussed at length within existing literature, Papadakis (Papadakis 1981) suggests 80% of the scattered energy is carried off by shear waves within the Rayleigh regime.

This section aims to investigate the shear-to-longitudinal wave balance, thereby verifying the premise of the proposed approach. The chosen metric measures the average intensity of the longitudinal waves within the scattered field, through spatial filtering, and compares it to the total average level of backscatter, comprising both wave modes. Due to a lack of absolute accuracy, the measurements are normalised by assuming the spatially coherent wave (i.e. the signal) comprises solely a longitudinal wave. The complete definition for measuring the ratio of shear waves (SR), is given in Equation 8.6.

$$SR = 1 - \frac{N_{rms,L}}{S_{pk,L}} \frac{S_{pk,x_0}}{N_{rms,x_0}} \tag{8.7}$$

$N_{rms,L}$ and N_{rms,x_0} are the measure of backscatter (as shown in Figure 8.5b), of the longitudinal filtered U_L and U_{x0} signal respectively. As previously mentioned, the SR is directly related to the potential improvement in SNR, and thus using the same data, an equivalent, albeit theoretical, SNR improvement can be calculated.

Both these metrics are investigated for three different grain sizes, 100, 200, and 500μm. The simulations are repeated for nine random realisations of a scattering medium (Inconel 600), which enables the consideration of statistical variations. An example simulation is shown in Figure 8.5a, the LSF is located at x=20mm and y=6mm. Similarly to the previous simulations, a longitudinal plane wave with a 2MHz centre frequency is excited.

	d=100μm	d=200μm	d=500μm
SR (%±σ) SNR	54%±9% +6.8dB	51%±11% +6.3dB	38%±18% +4.5dB

Table 8.1: Simulation results indicating the improvement in SNR when comparing a conventional time-trace to a longitudinal filtered equivalent.

The results in Table 8.1 indicate that a substantial portion of the scattered field comprises shear waves, which could enable an SNR improvement of approximately 6dB. The decrease in SR with grain size seems to suggest that the balance of shear waves decreases as larger grains are considered.

The above results have investigated the wave field using point measurements (single nodes). As established in Chapter 5, when considering receivers of finite size, such as ultrasonic transducers, spatial averaging plays an important role in determining the perceived level of noise. Since shear waves have a smaller spatial correlation length, it is expected that backscatter due to shear waves will spatially average quicker than the longitudinal counterpart. In order to investigate this, and to apply the LSF methodology to a more (relatively) practical scenario, a spatial filtering methodology is developed for an ultrasonic array.

8.5 Array Simulation

Implementing an LSF for an array requires some additional modelling devices. Firstly, instead of fully embedding the filter, the need arises to incorporate the LSF along a boundary. This is achieved here by using symmetry boundary conditions as shown in Figure 8.7a. The wave field animation resulting from an excitation with these boundary conditions in shown in Figure 8.7b. This indicates that the boundary conditions have succeeded in creating the desired semi-circular longitudinal wave front in emission.

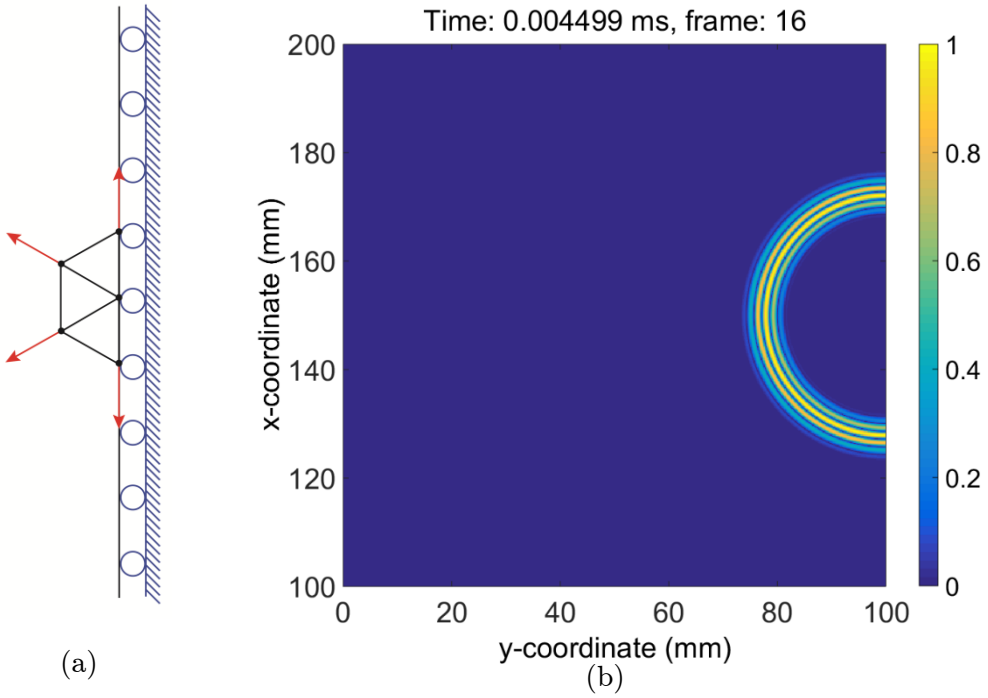

Figure 8.7: (a) Boundary conditions for an LSF along a boundary. (b) Resulting wave field when the boundary conditions in (a) are used in emission.

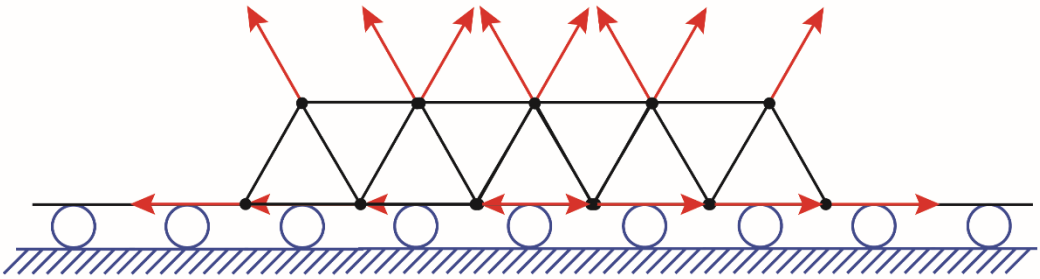

Figure 8.8: Implementation of multiple LSFs to model a line source, for clarity illustrated here for four nodes.

Again, by argument of reciprocity, this then can be applied in reception to implement a LSF along a boundary. As previously mentioned, a second change in comparison to the previous simulations is the requirement to model the footprint of the array (the array elements) which now consist of multiple nodes. Hence an LSF is placed at each node of the array element, such that a U_L is computed for each node, and is subsequently averaged for all the nodes pertaining to the array element. This is graphically equivalent to superposing the LSF, as shown in Figure 8.8.

8.5.1 Model

The FE modelling procedure to implement an array follows that outlined in Chapter 6. The layout (see Figure 8.9) comprises a 64 element linear array, each with a half centre-wavelength element pitch and element width. The array is in contact configuration with a 2D block (300mm x 100mm) (see Figure 8.9a) of Inconel 600 with an average grain size of 500µm. The array elements are excited as a piston source thus generating a longitudinal wave (but also spurious shear waves) with a 2MHz centre frequency.

A detection target is introduced into the material, centrally below the array, at a depth of 50mm as shown in Figure 8.9. The defect represents a 2mm radius SDH, which is implemented by constraining the displacement in both the x and y direction for all the nodes which reside within the circumference of the SDH (see Figure 8.9).

Whilst the same excitation procedure is used, the simulation calculates two FMCs, one pertaining to Uy displacements, which produces the conventional, non-filtered array image. The second FMC relates to the filtered time-displacements, U_L. The data are subsequently compiled into array images using the TFM algorithm (Holmes et al. 2005).

8.5.2 Results and Discussion

The conventional TFM image is shown in Figure 8.10a. The filtered TFM image, where the contribution of shear waves is filtered out, is shown in Figure 8.10b. The defect at around 50mm depth is on the verge of detection in both cases.

As can be seen, comparison of the two images reveals there is little to no improvement in image SNR in this case. Although no further analysis is performed here, this is probably

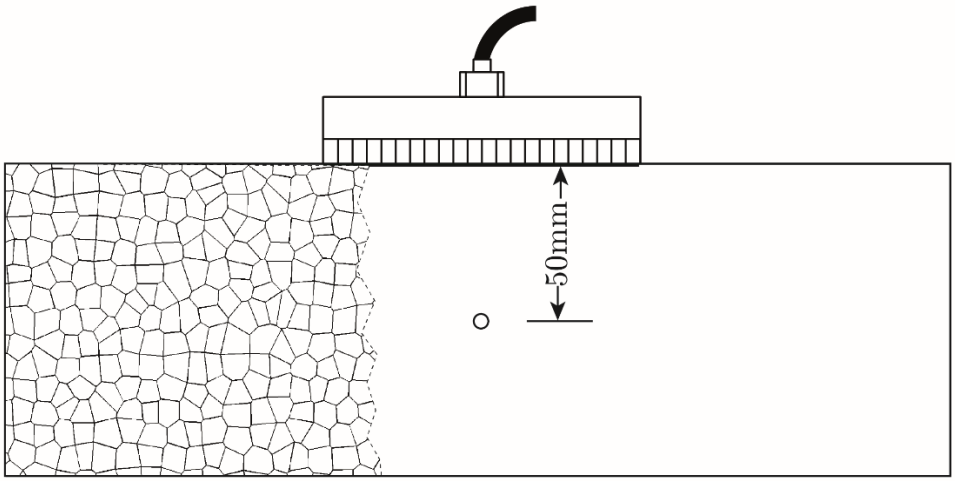

(a)

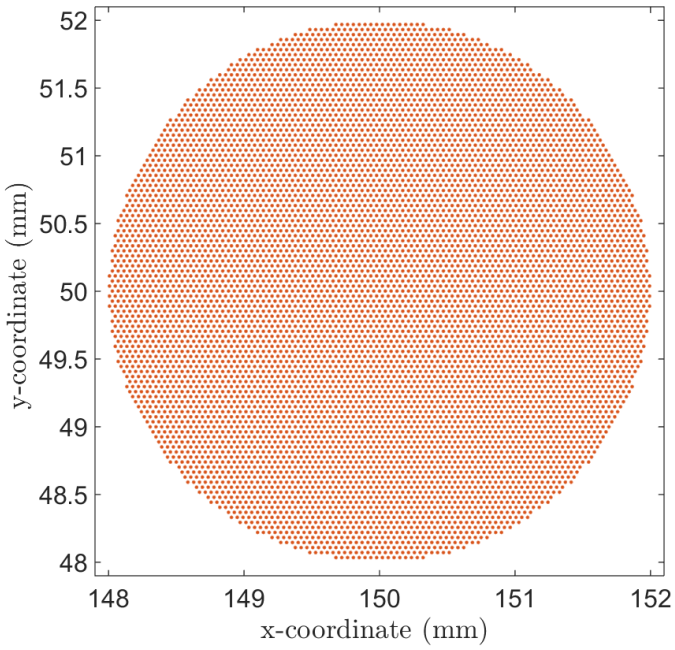

(b)

Figure 8.9: (a) Array model layout with a partly revealed polycrystalline microstructure. Model size is 300mmx100mm and a SDH defect is placed at 50mm depth. A zoom in is provided (b) of defect model, indicating the constrained nodes which conform to an imperfect shape due to the use of a structured mesh.

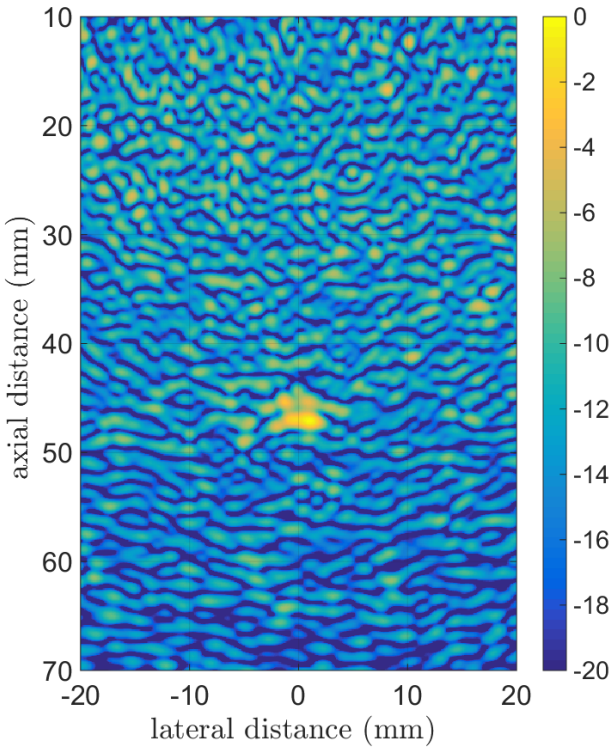

(a)

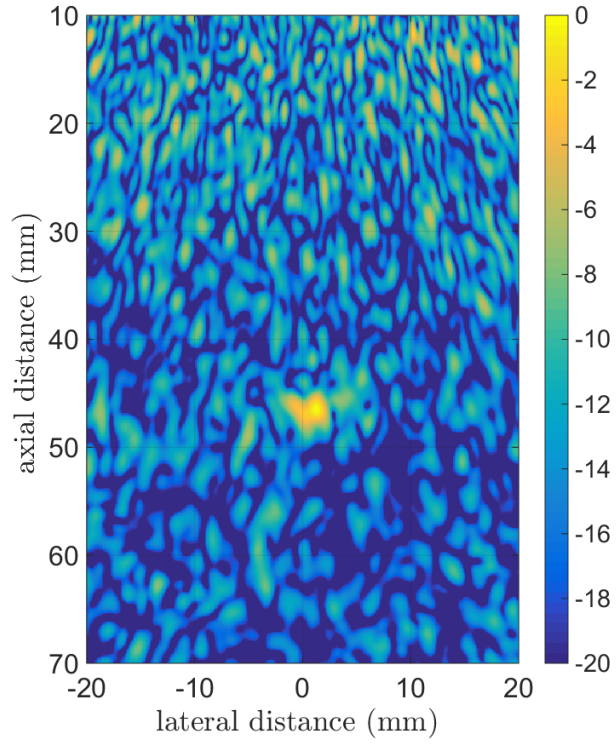

Figure 8.10: TFM array image compiled using (a) U_Y displacement and U_L displacement FMC. Imaging target at 50mm can be seen in both cases. Images are normalised to the maximum pixel intensity (signal in this case) and displayed on a dB scale.

due to the aforementioned spatial averaging effect, which inherently reduces the contribution of shear waves when considering receiving sources of finite dimensions. In addition, part of the - albeit limited - improvement of the filtered image is due to the elimination of a Rayleigh wave within the time-trace data, which is an additional noise factor for the conventional image.

Even so, the chosen 500μm grain size, because it is computationally less demanding, represents a conservative case as suggested from the previous results. It would be interesting to further quantify and understand the balance of both wave modes in additional scattering regimes. Furthermore, given that these simulations are limited to 2D, it remains to be determined how the spatial averaging effect compares in full 3D simulations, and whether the shear wave balance is different.

8.6 Conclusions

This chapter has briefly introduced a numerical method for filtering longitudinal waves, in attempt to reduce the contribution of shear waves to improve SNR. The proposed method is demonstrated to successfully separate both wave modes. However, when considering the benefits to transducers with finite size, in this case array elements with a half wavelength width, the benefit to SNR was reduced, as could be seen from comparisons of two array images.

The author believes this to be an interesting exploit, and further demonstrates the flexibility of the numerical simulations considered. As this chapter represented only an initial exploration, spatial filtering remains an idea for future work.

Chapter IX

CONCLUSIONS

Summary of thesis

9.1 Review of Thesis

Chapter II: reviewed the literature on grain scattering from a historical perspective and described the solutions considered at present to supress its adverse effect on ultrasonic NDE.

Chapter III: familiarised grain scattering through three experimental investigations. This included the inspection of a representative industrial component, the demonstration of scattering through measuring frequency dependent attenuation, and an investigation of temporally and spatially incoherent noise.

Chapter IV: outlined a benchmarking strategy for evaluating de-noising algorithms and applied it to three forerunning ultrasonic array imaging algorithms: TFM, PCI, and DORT MSF. When considering the statistics of detection, it was found that the candidate algorithms performed similarly for the assumed test case.

Chapter V: established Finite Element modelling for simulating wave propagation and elastodynamic scattering within polycrystalline materials. Several methodologies were explored to find an efficient solution, which was subsequently validated in 2D and 3D against established theory. Modelling limitations were also found.

Chapter VI: extended the Finite Element modelling approach to a framework which enables the investigation of ultrasonic array imaging within highly scattering materials. New insights were found: the results suggested 2D arrays can be advantageous and potential benefits can be derived from separating the emitter and receiver to significantly reduce backscatter.

Chapter VII: proposed a pseudo-colouring signal processing scheme to exploit frequency diversity within polycrystalline materials, thereby adding frequency information alongside intensity, to aid interpretation of ultrasonic images.

Chapter VIII: preliminarily explored spatial averaging as a means to filter-out backscatter. A numerical method was established, which showed the scattered waves to comprise shear waves, however when considering a practical scenario, the benefit to ultrasonic NDE was limited.

9.2 Key Contributions

Scattering fundamentally limits ultrasonic inspection frequency as inescapably as aerodynamic drag limits airspeed, and ultimately restricts the achievable performance of ultrasonic NDE. Whilst there seems to be no direct solution, two main philosophies can be considered to enable improvements. The first is to optimise existing ultrasound technologies, be it software (e.g. imaging algorithms) or hardware (e.g. array aperture). The second is to try and obtain new information to distinguish the targets of interest from noise. Before discussing either of these approaches, their success starts with the development of a reliable platform for understanding the physics, including the availability of the necessary tools.

Whereas current de-noising algorithms have already made significant strides, the increasing availability of computing power will continue to present future opportunities for improvements. Establishing such advances requires a rigorous tool to objectively compare competing algorithms. The approach described in Chapter IV is believed to be the first delivery of a comparison methodology for detection that provides a robust quantitative evaluation, enabling to establish future advances with objective measures of the progress that new methods achieve.

Previous studies within this field often relied on single scattering assumptions, and thus limited to relatively 'weak' scattering regimes. It is now evident that in order to advance ultrasonic NDE we must also consider multiple scattering environments. Whilst FE modelling presented a promising candidate, it remained uncertain whether it could capture the complex scattering phenomena, and furthermore, ambiguity remained regarding the most appropriate modelling methodology. These questions were answered in Chapter V where FE is revealed to capture the different scattering regimes, by validation against the widely-established Unified Theory. The approach was extended to 3D, for the first time, revealing fundamental limitations for the Rayleigh scattering mechanism simulated in reduced dimensions.

Relying on these newfound tools, the optimisation of existing technologies was initially pursued. The was demonstrated experimentally in Chapter IV which suggested the potential benefits which can be derived from considering de-noising imaging algorithms to be limited.

Given the random nature of polycrystalline materials, an experimental investigation would ideally require a statistically significant number of samples with well characterised microstructures, and even if this were practical, the fundamental difficulty of measuring signals in noisy environments still limits its efficacy. Alternatively, modelling such as the aforementioned FE technique, can randomly generate materials and offers a valuable investigative flexibility. This was realised in Chapter VI by outlining a framework to investigate ultrasonic array imaging, this time, considering the hardware aspect of the array. Significant performance enhancements were found to be possible by optimising the array configuration. Namely, 2D arrays are preferable, and the concept of a backscatter envelope was described, which provided an explanation for the improvements in SNR derived from separating emitter and receiver array. It is believed that this framework is of future value, where it can be adopted for analysing more practical and specific industrial inspection scenarios.

Beyond optimisation, attempts were also made to find new premises of detection, finding and subsequently exploiting additional information to discern signals and noise. Two such approaches were explored, one in Chapter VII, which incorporates frequency information into ultrasonic images by utilising pseudo-colouring. The approach was shown to be successful, and could potentially aid interpretation of ultrasonic images in the future.

Another such potential exploit was discovered through spatial observations which were made possible by the aforementioned numerical models. Unlike the spatially coherent wave, it was recognised that the scattered field comprises shear waves. As an initial exploration, a simulation study in Chapter 8 details a spatial filter to remove backscatter due to shear waves, and establish possible improvements. Although further work is required to fully establish the merit of this approach, initial results suggest inherent spatial averaging limits its benefit.

Aside from those previously discussed, some additional opportunities remain, some that have arisen during the present developments, which may merit future work.

9.3 Future Work

An obvious candidate for future consideration is Super Resolution techniques, as discussed in Chapter II. Recent research (Fan et al. 2014) has shown that in comparison to conventional algorithms (e.g. TFM), a higher SNR is required to operate Super Resolution algorithms, and hence it remains to be determined whether the required drop in operating frequency to improve SNR, can still outperform the currently established algorithms (which can operate at higher frequency) in terms of resolution.

Similarly, non-linear array imaging recently emerged and presents a potentially applicable imaging approach. Potter et al. (2014) have realised the possibility to exploit the non-linear behaviour of cracks for imaging purposes. Assuming grains do not exhibit such behaviour, it may be possible to utilise this additional information and derive improvements over current imaging techniques.

Although it was already briefly introduced, ultimately it seems promising to exploit the difference in spatial characteristics of the scattered field and the spatially coherent signal. A promising field for its practical implementation is that of meta-materials, where there has already been some success at for example (Shen et al. 2014) eliminating the contribution of an aberrating layer to better image an object behind it (e.g. ultrasonically accessing the brain and adding a layer in front of the skull to remove its distorting effects).

The FE modelling capability developed here is a promising tool for future research. However, given its computational cost, it is unlikely that useful sizes of 3D models will be feasible in the short term. It may be possible and also rewarding to accelerate this process by furthering our understanding of the discrepancies between 2D and 3D, in order to develop pseudo-3D models in 2D where we don't have the same limitations in computing power.

Finally, a question remains regarding the optimal inspection frequency, particularly for arrays. Within a certain regime, focusing performance (and image SNR) increases with frequency (Wilcox 2011), but eventually multiple scattering overwhelms this effect. Finding this optimum would provide far-reaching implications, and it may emerge, analogous to remaining within the subsonic drag regime, that the maximum efficiency for ultrasonic NDE lies at the limit of single scattering, before the onset of multiple scattering.

9.4 Thesis Publications

[p1] A. Van Pamel, M. J. S. Lowe, C. R. Brett, "Evaluation of Ultrasonic Array Imaging Algorithms for Inspection of a Coarse Grained Material" in Review of Progress in QNDE (AIP Conference Proceedings, Melville NY), vol. 1581, pp. 156-163 (2013).

[p2] A. Van Pamel, M. J. S. Lowe, C. R. Brett, "A Methodology for Evaluation Detection performance of Ultrasonic Array Imaging Algorithms for Coarse Grained Materials" IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 61, pp.2054-2066 (2014).

[p3] A. Van Pamel, P. Huthwaite, C. R. Brett, M. J. S. Lowe, "A Finite Element Model Investigation of Ultrasonic Array Performance for Inspection Polycrystalline Materials" in Review of Progress in QNDE (AIP Conference Proceedings, Melville NY), vol. 1650, pp.1007-1014, (2014).

[p4] A. Van Pamel, C. R. Brett, M. J. S. Lowe "An Initial Investigation into Pseudo-Colouring for Ultrasonic NDE of Polycrystalline Materials" in Review of Progress in QNDE (AIP Conference Proceedings, Melville NY), vol. 1650, pp.1031-007-1036, (2014).

[p5] P. Huthwaite, F. Shi F, A. Van Pamel, M.J.S Lowe MJS " High-Speed GPU-Based Finite Element Simulations for NDT", in Review of Progress in QNDE, (AIP Conference Proceedings, Melville NY), vol. 1650, pp. 1815-1819, (2014).

[p6] A. Van Pamel, P. Huthwaite, C. R. Brett, M. J. S. Lowe, "Finite Element Modelling of Wave Propagation in Highly Scattering Materials" in Review of Progress in QNDE, (AIP Conference Proceedings, Melville NY), in print (2015).

[p7] A. Van Pamel, P. Huthwaite, C. R. Brett, M. J. S. Lowe, "Finite element modelling of elastic wave scattering within a polycrystalline material in two and three dimensions" JASA, in print (2015).

[p8] A. Van Pamel, P. Huthwaite, C. R. Brett, M. J. S. Lowe, "Numerical Simulations of Ultrasonic Array Imaging within Highly Scattering Materials" submitted to NDT&E 2015.

REFERENCES

Abaqus 6.14, 2014, Simulia. *Dassault Systems, www.simulia.com*.

Afruz, J., Wilson, V. & Umbaugh, S.E., 2010, Frequency Domain Pseudo-color to Enhance Ultrasound Images. *Computer and Information Science*, 3(4), p.24.

Ahmed, S. & Thompson, R.B., 1996, Propagation of elastic waves in equiaxed stainless steel polycrystals with aligned axes. *Journal of the Acoustical Society America*, 99(4), pp.2086–2096.

Anxiang, L., Roberts, R. & Haldipur, P., 2003, Computational study of grain scattering effects in ultrasonic measurements. In *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 657)*. p. 117.

Asgedom, E.G.E., Gelius, L.-J.L., Austeng, A., Holm, S. & Tygel, M., 2011, Time-reversal multiple signal classification in case of noise: A phase-coherent approach. *The Journal of the Acoustical Society of America*, 130(4), p.2024.

Aubry, A. & Derode, A., 2009a, Detection and imaging in a random medium: A matrix method to overcome multiple scattering. *Journal of Applied Physics*, 106(044903).

Aubry, A. & Derode, A., 2009b, Random matrix theory applied to acoustic backscattering and imaging in complex media. *Physical Review Letters*, 102(8), p.84301.

Auld, B.A., 1990, *Acoustic fields and waves in solids vol II*, Malabar, Florida: Krieger.

Aurenhammer, F., 1991, Voronoi diagrams - a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3), pp.345–405.

Baligand, B., Grozellier, M. & Romy, D., 1986, Improvement in ultrasonic examination of austenitic steel. *Materials Evaluation*, 44, pp.577–581.

Bathe, K.J., 1996, *Finite Element Procedures*, New Jersey: Prentice Hall.

Bhatia, A.B., 1959a, Scattering of High Frequency Sound Waves in Polycrystalline Materials. II. *The Journal of the Acoustical Society of America*, 31(8), p.1140.

Bhatia, A.B., 1959b, Scattering of High-Frequency Sound Waves in Polycrystalline Materials. *The Journal of the Acoustical Society of America*, 31(1), p.16.

Bhatia, A.B., 1967, Scattering of plane elastic waves due to elastic anisotropy of the grains in polycrystals. In *Ultrasonic Absoroption*. Oxford: Clarendon Press, pp. 279–287.

Bilgutay, N.M., Bencharit, U. & Saniie, J., 1989, Enhanced ultrasonic imaging with split-spectrum processing and polarity thresholding. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 37(10), pp.1590–1592.

Bugge, J., Kjær, S. & Blum, R., 2006, High-efficiency coal-fired power plants development and perspectives. *Energy*, 31(10-11), pp.1437–1445.

Bunge, H.J. & Morris, P.R., 1982, *Texture analysis in materials science: mathematical methods*, London: Butterworths.

Camacho, J., Brizuela, J. & Fritsch, C., 2010, Grain noise reduction by phase coherence imaging. In *Review of Progress in Quantitative NDE (AIP Conference Proceedsings vol 1211)*. AIP, pp. 855–862.

Camacho, J. & Fritsch, C., 2011, Phase coherence imaging of grained materials. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 58(5), pp.1006–1015.

Camacho, J., Parrilla, M. & Fritsch, C., 2009, Phase coherence imaging. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 56(5), pp.958–974.

Chaffaï, S., Roberjot, V., Peyrin, F., Berger, G. & Laugier, P., 2000, Frequency dependence of ultrasonic backscattering in cancellous bone: Autocorrelation model and experimental results. *The Journal of the Acoustical Society of America*, 108(5), p.2403.

Chassignole, B., Duwig, V., Ploix, M.-A., Guy, P. & El Guerjouma, R., 2009, Modelling the attenuation in the ATHENA finite elements code for the ultrasonic testing of austenitic stainless steel welds. *Ultrasonics*, 49(8), pp.653–8.

Connolly, G.D., Lowe, M.J.S., Temple, J.A.G. & Rokhlin, S.I., 2009, The application of Fermat's principle for imaging anisotropic and inhomogeneous media with application to austenitic steel weld inspection. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2111), pp.3401–3423.

Croxford, A.J., Wilcox, P.D., Drinkwater, B.W. & Konstantinidis, G., 2007, Strategies for guided-wave structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2087), pp.2961–2981.

Drinkwater, B.W. & Wilcox, P.D., 2006, Ultrasonic arrays for non-destructive evaluation: A review. *NDT & E International*, 39, pp.525–541.

Drozdz, M.B., 2008, Efficient Finite Element Modelling of Ultrasound Waves in Elastic Media. *Imperial College London Mechanical Engineering PhD Thesis*.

EPRI, 2013, 21st Century Coal Power: Recent Developments in Coal Power Generation Technology. In J. Phillips, ed. *Generation Summit*. Marcus Evans.

Fan, C., Caleap, M., Pan, M. & Drinkwater, B.W., 2014, A comparison between ultrasonic array beamforming and super resolution imaging algorithms for non-destructive evaluation. *Ultrasonics*, 54(7), pp.1842–50.

Feldman, M., 2011, Hilbert transform in vibration analysis. *Mechanical Systems and Signal Processing*, 25(3), pp.735–802.

Felice, M. V, Velichko, A., Wilcox, P., Barden, T. & Dunhill, T., 2015, Optimization of ultrasonic array inspections using an efficient hybrid model and real crack shapes. *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1650)*, pp.978–986.

Feuilly, N., 2009, Etude de l'influence de la microstructure sur la diffusion d'une onde ultrasonore en vue de l'amerioration du controle non destructif des materiaux polycrystallins. *PhD Thesis Universite de la Mediterranee*.

Feuilly, N., Dupond, O., Chassignole, B., Moysan, J. & Corneloup, G., 2009, Relation between ultrasonic backscattering and microstructure for polycrystalline materials.

Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1096), pp.1216–1223.

Firestone, F.A., 1942, Flaw detecting device and measuring instrument. *US Patent*, (US2280226).

Firestone, F.A., 1946, The Supersonic Reflectoscope, an Instrument for Inspecting the Interior of Solid Parts by Means of Sound Waves. *The Journal of the Acoustical Society of America*, 17(3), p.287.

Fleming, M., 2008, Far Field Super Resolution Imaging. *Imperial College London Mechanical Engineering PhD Thesis*.

Fleming, M., Lowe, M.J.S. & Simonetti, F., 2006, Super resolution imaging: Performance studies D. O. . C. Thompson D. E., ed. *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 820)*, pp.736–743.

Gabor, D., 1946, Theory of communication. Part 1: The analysis of information. *Journal of the Institution of Electrical Engineers*, 93(26), pp.429–441.

Gerig, A. & Varghese, T., 2004, Improved parametric imaging of scatterer size estimates using angular compounding. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 51(6), pp.708–715.

Ghoshal, G. & Turner, J. a., 2009, Numerical model of longitudinal wave scattering in polycrystals. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 56(7), pp.1419–1428.

Giacomini, A. & Bertini, A., 1939, Supersonic Method for Testing Homogeneneity of Solids. *Ricerca Scientifica*, 10, p.921.

Goebbels, K., 1980, *Structure analysis by scattered ultrasonic radiation* R. S. Sharpe, ed., New York: Academic Press.

Hall, T.J., Madsen, E.L., Zagzebski, J. a & Boote, E.J., 1989, Accurate depth-independent determination of acoustic backscatter coefficients with focused transducers. *The Journal of the Acoustical Society of America*, 85(6), pp.2410–2416.

Hanley, J.A. & McNeil, B.J., 1982, The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1), pp.29–36.

Hirsekorn, S., 1988, The scattering of ultrasonic waves by multiphase polycrystals. *The Journal of the Acoustical Society of America*, 83(4), p.1231.

Hirsekorn, S., 1982, The scattering of ultrasonic waves by polycrystals. *The Journal of the Acoustical Society of America*, 72(3), p.1021.

Hirsekorn, S., 1983, The scattering of ultrasonic waves by polycrystals. II. Shear waves. *The Journal of the Acoustical Society of America*, 73(4), p.1160.

Hirsekorn, S., 1985, The scattering of ultrasonic waves in polycrystalline materials with texture. *The Journal of the Acoustical Society of America*, 77(3), p.832.

Holmes, C., Drinkwater, B.W. & Wilcox, P.D., 2005, Post-processing of the full matrix of ultrasonic transmit – receive array data for non-destructive evaluation. *NDT & E International*, 38, pp.701–711.

Hunter, A.J., Drinkwater, B.W. & Wilcox, P.D., 2008, The wavenumber algorithm for full-matrix imaging using an ultrasonic array. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 55, pp.2450–2462.

Huntington, H.B., 1950, On Ultrasonic Scattering by Polycrystals. *The Journal of the Acoustical Society of America*, 22, p.362.

Huthwaite, P., 2014, Accelerated finite element elastodynamic simulations using the GPU. *Journal of Computational Physics*, 257, pp.687–707.

IEA, 2014, *Key World Energy Statistics 2014*,

Insana, M.F., 1986, Tests of the accuracy of a data reduction method for determination of acoustic backscatter coefficients. *The Journal of the Acoustical Society of America*, 79(5), p.1230.

Ishimaru, A., 1978, *Wave Propagation and Scattering in Random Media*, New York: Academic Press.

Izquierdo, M.A.G., Hernandez, M.G., Graullera, O. & Ullate, L.G., 2002, Time–frequency Wiener filtering for structural noise reduction. *Ultrasonics*, 40(1-8), pp.259–261.

Jie, Z., Drinkwater, B.W. & Wilcox, P.D., 2013, Comparison of ultrasonic array imaging algorithms for nondestructive evaluation. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 60(8), pp.1732–1745.

Jinxiu, L., Junli, L. & Ping, W., 2007, Pseudocolor Coding of Medical Images Based on Gradient. In *International Conference on Bioinformatics and Biomedical Engineering*. IEEE, pp. 932–935.

Kalashnikov, A. & Challis, R.E., 2005, Errors and Uncertainties in the Measurement of Ultrasonic Wave Attenuation and Phase Velocity. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 52(10).

Kamaya, M., 2009, A procedure for estimating Young's modulus of textured polycrystalline materials. *International Journal of Solids and Structures*, 46(13), pp.2642–2649.

Karal, F.C. & Keller, J.B., 1964, Elastic, electromagnetic, and other waves in a random medium. *J. Math Phys.*, 5, pp.537–547.

Karaojiuzt, M., Bilgutay, N., Akgul, T. & Popovics, S., 1998, Defect detection in concrete using split spectrum processing. In *IEEE Ultrasonics Symposium*. pp. 843–846.

Karpur, P. & Canelones, O.J., 1992, Split spectrum processing: A new filtering approach for improved signal-to-noise ratio enhancement of ultrasonic signals. *Ultrasonics*, 30, pp.351–357.

Karpur, P., Shankar, P.M., Rose, J.L. & Newhouse, V.L., 1988, Split spectrum processing: Determination of the available bandwidth for spectral splitting. *Ultrasonics*, 26, pp.204–209.

Karpur, P., Shankar, P.M., Rose, J.L. & Newhouse, V.L., 1987, Split spectrum processing: optimizing the processing parameters using minimization. *Ultrasonics*, 25(4), pp.204–208.

Keller, M., Patton, T., Degtyar, A., Umbach, J., Hassan, W., Kinney, A., Roberts, R., Margetan, F. & Brasche, L., 2005, *Inspection Development for Nickel Billet - Engine Titanium Consortium Phase II*,

Kim, J., Udpa, L. & Udpa, S., 2001, Multi-stage adaptive noise cancellation for ultrasonic NDE. *NDT & E International*, 34(5), pp.319–328.

Kocks, U.F., Tomé, C.N. & Wenk, H.-R., 2000, *Texture and Anisotropy: Preferred Orientations in Polycrystals and Their Effect on Materials Properties*, Cambridge University Press.

Lan, B., Lowe, M. & Dunne, F.P.E., 2014, Experimental and computational studies of ultrasound wave propagation in hexagonal close-packed polycrystals for texture detection. *Acta Materialia*, 63, pp.107–122.

Lardner, T., Li, M. & Gachagan, A., 2014, Using phase information to enhance speckle noise reduction in the ultrasonic NDE of coarse grain materials. In *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1581)*. AIP Publishing, p. 1061.

Lehman, S.S.K. & Devaney, A.J.A., 2003, Transmission mode time-reversal super-resolution imaging. *The Journal of the Acoustical Society of America*, 113(5), pp.2742–2753.

Li, C., Pain, D., Wilcox, P.D. & Drinkwater, B.W., 2013, Imaging composite material using ultrasonic arrays. *NDT & E International*, 53, pp.8–17.

Lifshits, I.M. & Parkhomovski, G.D., 1950, On the theory of ultrasonic wave propagation in polycrystals. *Zh. Eksp. Teor. Fiz.*, 20(1), pp.175–182.

Lind, G., 1970, Measurement of sea clutter correlation with frequency agility and fixed frequency radar. *Philips Telecommunication Review*, 29, pp.32–38.

Lobkis, O. & Rokhlin, S., 2010, Characterization of polycrystals with elongated duplex microstructure by inversion of ultrasonic backscattering data. *Applied Physics Letters*, 96(161905).

Lommel, E., 1885, Die Beugungserscheinungen einer kreisrunden Oeffnung und eines kresirunden Schirmschens theoretisch und experimentell Bearbeitet. *Abh. Bayer. Akad. Wiss. Math -Naturwiss. Kl*, 15, pp.233–328.

Madsen, E., 1984, Method of data reduction for accurate determination of acoustic backscatter coefficients. *The Journal of the Acoustical Society of America*.

Marengo, E., 2007, Time-reversal MUSIC imaging of extended targets. *IEEE Transactions on Image Processing*, 16(8), pp.1967–1984.

Margetan, F. & Thompson, R.B., 1992, Microstructural noise in titanium alloys and its influence on the detectability of hard-alpha inclusions. In D. O. Thompson & D. E. Chementi, eds. *Review of Progress in Quantitative NDE vol 11*. Plenum Press, pp. 1717–1724.

Margetan, F.J., 2012, Bruce Thompson: Adventures and advances in ultrasonic backscatter. In *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1430)*. pp. 54–82.

Margetan, F.J., Gigliotti, M., Brashe, L. & Leach, W., 2002, Fundamental studies: inspection properties for engine titanium alloys. In *FAA Report DOT/FAA/AR-02/114*. Washington, DC.

Margetan, F.J., Gray, T. & Thompson, R., 1991, A technique for quantitatively measuring microstructurally induced ultrasonic noise. In D. O. Thompson & D. E. Cheminti, eds. *Review of Progress in Quantitative NDE vol 10B*. Plenum Press, p. 1721.

Margetan, F.J., Thompson, R.B. & Yalda-Mooshabad, I., 1994, Backscattered microstructural noise in ultrasonic toneburst inspections. *Journal of Nondestructive Evaluation*, 13(3), pp.111–136.

Margetan, F.J., Thompson, R.B. & Yalda-Mooshabad, I., 1993, Modeling Ultrasonic Microstructural Noise in Titanium Alloys. In D. Thompson & D. Chimenti, eds. *Review of Progress in Quantitative NDE vol 12*. Plenum Press, pp. 1735–1742.

Margetan, F.J., Umback, J., Robert, R., Friedl, J., Degtyar, A., Keller, M., Hassa, W., Brasche, L., Klassen, A., Wasan, H. & Kinney, A., 2007, Inspection Development for Titanium Forgings. In F. A. A. W. J. H. T. Center, ed. *DOT/FAA/AR-05/46*. Atlantic City, N. J.

Margetan, F.J., Yalda, I., Thompson, R.B. & Umbach, J., 1997, Ultrasonic grain noise modeling: recent applications to engine titanium inspections. In D. O. . C. Thompson D. E., ed. *Review of Progress in Quantitative NDE vol 16*. New York: Plenum Press, pp. 1555–1562.

Mason, W.P., 1947, Attenuation and Scattering of High Frequency Sound Waves in Metals and Glasses. *The Journal of the Acoustical Society of America*, 19(3), p.725.

Mason, W.P. & McSkimin, H.J., 1948, Energy losses of sound waves in metals due to scattering and diffusion. *Journal of Applied Physics*, 19, pp.940–946.

Maurel, A., Pagneux, V., Boyer, D. & Lund, F., 2006, Propagation of elastic waves through polycrystals: the effects of scattering from dislocation arrays. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 462(2073), pp.2607–2623.

Merkulov, L.G., 1956, Investigation of the scattering of ultrasound in metals. *Zh. Tekhn. Fiz.*, 26, p.64.

Metz, C.E., 1978, Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8(4), pp.283–98.

Miller, G.F. & Pursey, H., 1954, The Field and Radiation Impedance of Mechanical Radiators on the Free Surface of a Semi-Infinite Isotropic Solid. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 223(1155), pp.521–541.

Nagy, P.B. & Adler, L., 1988, Scattering induced attenuation of ultrasonic backscattering. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 7*. Plenum Press, pp. 1263–1272.

Nagy, P.B., Rypien, D. V. & Adler, L., 1987, Ultrasonic Attenuation measurement by backscattering analysis. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 6*. Plenum Press, pp. 1411–1418.

Newhouse, V.L., Bilgutay, V.L., Saniie, N.M. & Furgason, E.S., 1982, Flaw-to-grain echo ehancement by split spectrum processing. *Ultrasonics*, 20, p.59.

Nguyen, T.Q. & Jayasimha, S., 1994, Polarity-coincidence filter banks and nondestructive evaluation. *Proceedings of IEEE International Symposium on Circuits and Systems*, 2, pp.497–500.

Ogilvy, J.A., 1991, *Theory of Wave Scattering From Random Rough Surfaces*, New York: Adam Hilger.

Van Pamel, A., Huthwaite, P., Brett, C.R. & Lowe, M.J.S., 2015, A finite element model investigation of ultrasonic array performance for inspecting polycrystalline materials. *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1650)*, pp.1007–1014.

Panetta, P., Bland, L.G., Maureen, T. & Hassan, W., 2014, Ultrasonic Backscattering Measurements of Grain Size in Metal Alloys. In *TMS 2014 Supplemental Proceedings*. John Wiley & Sons, Inc., pp. 721–730.

Papadakis, E.P., 1963a, Diffraction of Ultrasound in Elastically Anisotropic NaCl and in Some Other Materials. *Journal of the Acoustical Society America*, 35(4), pp.490–494.

Papadakis, E.P., 1964a, From Micrograph to Grain-Size Distribution with Ultrasonic Applications. *Journal of Applied Physics*, 35(5), pp.1586–1594.

Papadakis, E.P., 1963b, Rayleigh and Stochastic Scattering of Ultrasonic Waves in Steel. *Journal of Applied Physics*, 34(2), pp.265–269.

Papadakis, E.P., 1981, Scattering in Polycrystalline Media. In P. D. Edmonds, ed. *Ultrasonics: Methods in Experimental Physics*. New York: Academic Press, pp. 237–298.

Papadakis, E.P., 1964b, Ultrasonic Attenuation and Velocity in Three Transformation Products in Steel. *Journal of Applied Physics*, 35(5), pp.1474–1482.

Papadakis, E.P., 1968, Ultrasonic Attenuation Caused by Scattering in Polycrystalline Media. In W. P. Mason, ed. *Physical Acoustics IV Part B*. New York and London: Academic Press, pp. 296–328.

Papadakis, E.P., 1965, Ultrasonic Attenuation Caused by Scattering in Polycrystalline Metals. *The Journal of the Acoustical Society of America*, 37(4), p.711.

Pardo, E., San Emeterio, J.L., Rodriguez, M.A. & Ramos, A., 2006, Noise reduction in ultrasonic NDT using undecimated wavelet transforms. *Ultrasonics*, 44, pp.e1063–e1067.

Phillips, J. & Wheeldon, M., 2010, Economic Analysis of Advanced Ultra-Supercritical Pulverized Coal Power Plants: A Cost-Effective CO2 Emission Reduction Option? In *Advances in Materials Technology for Fossil Power Plants Proceedings from the Sixth International Conference*. pp. 53–65.

Portzgen, N., Gisolf, D. & Blacquiere, G., 2007, Inverse wave field extrapolation: a different NDI approach to imaging defects. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 54(1), pp.118–127.

Potter, J.N., Croxford, A.J. & Wilcox, P.D., 2014, Nonlinear Ultrasonic Phased Array Imaging. *Physical Review Letters*, 113(14), p.144301.

Prada, C., Manneville, S., Spoliansky, D. & Fink, M., 1996, Decomposition of the time reversal operator: Detection and selective focusing on two scatterers. *The Journal of the Acoustical Society of America*, 99(4), p.2067.

Quey, R., Dawson, P.R. & Barbe, F., 2011, Large-scale 3D random polycrystals for the finite element method: Generation, meshing and remeshing. *Computer Methods in Applied Mechanics and Engineering*, 200(17-20), pp.1729–1745.

Rodriguez, M.A., Emeterio, J.L., Lázaro, J.C. & Ramos, A., 2004, Ultrasonic flaw detection in NDE of highly scattering materials using wavelet and Wigner–Ville transform processing. *Ultrasonics*, 42(9), pp.847–851.

Roe, R.J., 1965, Description of crystallite orientation in polycrystalline materials. III. General solution to pole figure inversion. *Journal of Applied Physics*, 36, pp.2024–2031.

Rogers, P.H. & Van Buren, a. L., 1974, An Exact Expression for Lommel's Diffraction Integral. *The Journal of the Acoustical Society of America*, 55(4), p.430.

Rokhlin, L.L., 1972, Scattering of Ultrasound in Polycrystalline Materials. *Soviet Physics – Acoustics*, 18(1).

Rose, J., 1993, Theory of ultrasonic backscatter from multiphase polycrystalline solids. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 12*. New York: Plenum Press, pp. 1719–1726.

Rose, J., 1991, Ultrasonic backscattering from polycrystalline aggregates using time-domain linear response theory. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 10*. New York: Plen, pp. 1715–1720.

Rose, J.H., 1992, Ultrasonic backscatter from microstructure. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 11*. New York: Plenum Press, pp. 1677–1684.

Rose, J.L., Nestleroth, J.B. & Balasubramaniam, K., 1988, Utility of Split-Spectrum Processing in Ultrasonic Nondestructive Evaluation. *Materials Evaluation*, 46(1), pp.144–122.

Roth, W., 1948, Scattering of Ultrasonic Radiation in Polycrystalline Metals. *Journal of Applied Physics*, 19(10), p.901.

Russell, M. & Neal, S., 1997, Experimental Evidence of Single and Multiple Scattering in Polycrystalline Materials. In D. O. Thompson & D. E. Chimenti, eds. *Review of Progress in Quantitative NDE vol 16*. New York: Plenum Press, pp. 1521–1528.

Saniie, J., Nagle, D.T. & Donohue, K.D., 1991, Analysis of order statistic filters applied to ultrasonic flaw detection using split-spectrum processing. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 38, pp.133–140.

Seo, C.H. & Yen, J.T., 2008, Sidelobe Suppression in Ultrasound Imaging Using Dual Apodization with CrossCorrelation. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 55(10), p.2198.

Shahjahan, S., 2013, *Improvement of flaw detection with ultrasonic array probed in multiple scattering polycrystalline materials by means of a random matrix approach: experimental and simulation studies*. These PhD, Universite Paris Diderot.

Shahjahan, S., Aubry, A., Rupin, F., Chassingole, B. & Derode, A., 2014, A random matrix approach to detect defects in a strongly scattering polycrystal: How the memory effect can help overcome multiple scattering. *Applied Physics Letters*, 104(23), p.234105.

Shahjahan, S., Rupin, F., Aubry, A., Chassignole, B., Fouquet, T. & Derode, A., 2014, Comparison between experimental and 2-D numerical studies of multiple scattering in Inconel600 by means of array probes. *Ultrasonics*, 54(1), pp.358–67.

Shankar, P.M., 1986, Speckle Reduction in Ultrasound B-Scans Using Weighted Averaging in Spatial Compounding. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 33(6), pp.754–758.

Shankar, P.M., Karpur, P., Newhouse, V.L. & Rose, J.L., 1989, Split-spectrum processing: Analysis of polarity thresholding algorithm for improvement of signal-to-noise ratio and detectability in ultrasonic signals. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 36(1), pp.101–108.

Shen, C., Xu, J., Fang, N.X. & Jing, Y., 2014, Anisotropic Complementary Acoustic Metamaterial for Canceling out Aberrating Layers. *Physical Review X*, 4(4), p.041033.

Shewchuk, J.R., 2002, Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry*, 22(1-3), pp.21–74.

Simonetti, F., 2006, Multiple scattering: The key to unravel the subwavelength world from the far- field pattern of a scattered wave. *Physical Review E*, 73(3), pp.1–13.

Simonetti, F., Fleming, M. & Marengo, E. a E., 2008, Illustration of the role of multiple scattering in subwavelength imaging from far-field measurements. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 25(2), pp.292–303.

Special Metals, 2013, INCONEL alloy 625. *Special Metals Material Data Sheet*, pp.1–18.

Stanke, F.E., 1985, Inversion of attenuation measurements in terms of parameterized autocorrelation function. In H. N. G. Wadley, ed. *NDE for Microstructure Process Control*. American Society for Metals, pp. 55–60.

Stanke, F.E. & Kino, G.S., 1984, A unified theory for elastic wave propagation in polycrystalline materials. *The Journal of the Acoustical Society of America*, 75(3), pp.665–681.

Strutt, J.W., 1871a, On the light from the sky its polarization and colour. *Philosophical Magazine*, 41, pp.107–120,274–279.

Strutt, J.W., 1871b, On the Scattering of Light by Small Particles. *Philosophical Magazine*, 41, pp.447–454.

Thompson, B., 2002, Elastic-wave propagation in random polycrystals: fundamentals and application to nondestructive evaluation. In M. Fink et al., eds. *Imaging of complex media with acoustic and seismic waves*. Springer Berlin Heidelberg, pp. 233–257.

Thompson, R.B. & Gray, T., 1983, A model relating ultrasonic scattering measurements through liquid–solid interfaces to unbounded medium scattering amplitudes. *The Journal of the Acoustical Society of America*, 74, pp.1279–1290.

Thompson, R.B., Margetan, F.J., Haldipur, P., Yu, L., Li, A., Panetta, P. & Wasan, H., 2008, Scattering of elastic waves in simple and complex polycrystals. *Wave Motion*, 45(5), pp.655–674.

Trahey, G.E., Smith, S.W. & Von Ramm, O.T., 1986, "Speckle pattern correlation with lateral aperture translation: Experimental results and implications for spatial compounding. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 33(3).

Turner, J. a., 1999, Elastic wave propagation and scattering in heterogeneous, anisotropic media: Textured polycrystalline materials. *The Journal of the Acoustical Society of America*, 106(2), p.541.

Turner, J. & Weaver, R., 1995, Time dependence of multiply scattered diffuse ultrasound in polycrystalline media. *The Journal of the Acoustical Society of America*, 97(5), p.2639.

Underwood, E.E., 1970, *Quantitative Stereology*, Reading: Addison-Wesley.

Viswanathan, R., 1989, *Damage Mechanisms and Life Assessment of High Temperature Components*, Metals Park, Ohio: ASM International.

Voigt, W., 1910, *Lehrbuch der kristallphysik:(mit ausschluss der kristalloptik)*, Berlin: B. G. Teubner.

Weaver, R.L., 1990, Diffusivity of ultrasound in polycrystals. *Journal of the Mechanics and Physics of Solids*, 38(1), pp.55–86.

Webster, G. & Ainsworth, R.A., 1994, *High Temperature Component Life Assessment*, Springer Science & Business Media.

Wilcox, P.D., 2011, Array imaging of noisy materials. *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1335)*, pp.890–897.

Wilcox, P.D., 2013, Ultrasonic arrays in NDE: Beyond the B-scan. *Review of Progress in Quantitative NDE (AIP Conference Proceedings vol 1511)*, pp.33–50.

Wilcox, P.D. & Velichko, A., 2010, Efficient frequency-domain finite element modeling of two-dimensional elastodynamic scattering. *The Journal of the Acoustical Society of America*, 127(1), pp.155–65.

Willems, H. & Goebbels, K., 1981, Characterization of microstructure by backscattered ultrasonic waves. *Metal Science*, 15(11-12), pp.549–553.

Yalda, I., Margetan, F.J. & Thompson, R.B., 1996, Predicting ultrasonic grain noise in polycrystals: A Monte Carlo model. *The Journal of the Acoustical Society of America*, 99(6), pp.3445–3455.

Yang, L., Lobkis, O. & Rokhlin, S., 2011, Explicit model for ultrasonic attenuation in equiaxial hexagonal polycrystalline materials. *Ultrasonics*, 51(3), pp.303–309.

Yang, L. & Rokhlin, S.I., 2011, On Comparison of Experiment and Theory for Ultrasonic Attenuation in Polycrystalline Niobium. *Journal of Nondestructive Evaluation*, 31(1), pp.77–79.

Yang, L. & Rokhlin, S.I., 2012, Ultrasonic Backscattering in Cubic Polycrystals with Ellipsoidal Grains and Texture. *Journal of Nondestructive Evaluation*, 32(2), pp.142–155.

Yu, L., Thompson, R.B. & Margetan, F.J., 2010, The spatial correlation of backscattered ultrasonic grain noise: theory and experimental validation. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 57(2), pp.363–78.

Zeng, F., Agnew, S.R., Raeisinia, B. & Myneni, G.R., 2010, Ultrasonic Attenuation Due to Grain Boundary Scattering in Pure Niobium. *Journal of Nondestructive Evaluation*, 29(2), pp.93–103.

Zhang, J., Drinkwater, B.W. & Wilcox, P.D., 2008, Defect characterization using an ultrasonic array to measure the scattering coefficient matrix. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 55(10), pp.2254–65.

Zhang, J., Hunter, A., Drinkwater, B.W. & Wilcox, P.D., 2012, Monte Carlo inversion of ultrasonic array data to map anisotropic weld properties. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 59(11), pp.2487–97.

Zhang, P., Balint, D. & Lin, J., 2011, An integrated scheme for crystal plasticity analysis: Virtual grain structure generation. *Computational Materials Science*, 50(10), pp.2854–2864.

Zhang, X.-G., Simpson, W.A., Vitek, J.M., Barnard, D.J., Tweed, L.J. & Foley, J., 2004, Ultrasonic attenuation due to grain boundary scattering in copper and copper-aluminum. *The Journal of the Acoustical Society of America*, 116(1), p.109.

Zhu, Y. & Weight, J., 1994, Ultrasonic nondestructive evaluation of highly scattering materials using adaptive filtering and detection. *IEEE Trans. Ultrason. Ferroelect. Freq. Control*, 41, pp.26–33.

Zweig, M. & Campbell, G., 1993, Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine [published erratum appears in Clin Chem 1993 Aug;39(8):1589]. *Clin. Chem.*, 39(4), pp.561–577.