Spatially Resolved Acoustic Spectroscopy (SRAS) Technology Development

Jethro Coulson, MPhys

Thesis submitted to The University of Nottingham

for the degree of Doctor of Engineering, September 2015



Contents

Abstract **2**

Acknowledgements **3**

List of acronyms **4**

1 Introduction **6**

1.1 Spatially resolved acoustic spectroscopy technology development 6

1.2 Laser ultrasonics 10

1.2.1 Ultrasonic waves 10

1.2.2 Generation mechanisms 15

1.2.3 Source types 16

1.2.4 Optical detection 17

1.3 Crystallography 23

1.4 Microstructure 26

1.5 Methods for determining microstructure 28

1.5.1 Optical techniques 28

1.5.2 X-Ray diffraction (XRD)) 29

1.5.3 Electron backscatter diffraction (EBSD) 30

1.5.4 Scanning acoustic microscopy 32

1.5.5 Neutron diffraction 32

1.5.6 X-Ray computed tomography 33

2 Spatially resolved acoustic spectroscopy **34**

2.1 SRAS background 34

2.2 EMDA instrument 39

2.2.1 Optical configuration 40

2.2.2 Detection and acquisition 43

2.3 EMDA SRAS example data 45
2.4 Crystallographic determination by SRAS 51

3 Renishaw SRAS system **60**
3.1 Customer demand for a SRAS system 62
3.1.1 SRAS measurements 62
3.1.2 The use of SRAS measurements 62
3.1.3 Market for SRAS 63
3.2 Technical requirements for the Renishaw SRAS system 64
3.3 SRAS instrument mechanical design 66
3.3.1 Constraint mechanism 67
3.3.2 Chassis design 67
3.3.3 Optical housing design 69
3.4 SRAS optical design 80
3.4.1 Detection path 80
3.4.2 Generation path 82
3.5 Manufacture and assembly 87
3.6 Lasers, electronics and data acquisition 88
3.6.1 Laser sources 90
3.6.2 KED circuit design 91
3.6.3 Signal chain 92
3.6.4 Data acquisition 93
3.6.5 Camera 94
3.7 Integration 94
3.7.1 Renishaw Equator 700 95
3.7.2 Synchronisation 96
3.7.3 Software 98
3.8 Conclusions 99
3.9 Further work 100
3.9.1 Rough surface detection 100
3.9.2 Optical efficiency 100
3.9.3 Increased automation 101
3.9.4 Data acquisition 102

4 Signal analysis **104**

4.1 Current signal processing technique issues . . . 106
4.2 Modelling SRAS signals . . . 107
4.3 SRAS data analysis techniques . . . 110
4.3.1 Maximum frequency peak (MFP) . . . 110
4.3.2 Gaussian fitting . . . 110
4.3.3 Short time Fourier transform . . . 115
4.3.4 Wavelets . . . 119
4.4 Comparison of processing techniques applied to simulated SRAS signals 124
4.4.1 Signal to noise . . . 124
4.4.2 Signal cycle number . . . 127
4.4.3 Frequency resolution . . . 130
4.4.4 Processing experimental data . . . 130
4.5 Spatial averaging . . . 136
4.6 Conclusions . . . 140
4.7 Further work . . . 140
4.7.1 Group velocity . . . 140
4.7.2 Iterative wavelet analysis . . . 141
4.7.3 Wavelet filtering for orientation determination . . . 142
4.7.4 Dual frequency signals . . . 142

5 System performance **146**

5.1 Motion platform performance . . . 146
5.2 Generation power and signal strength . . . 148
5.3 Effect of acoustic wavelength on signal strength . . . 153
5.4 Detection beam size . . . 155
5.5 Spatial resolution . . . 156
5.5.1 Orthogonal spatial resolution . . . 157
5.5.2 Effect of acoustic wavelength on spatial resolution . . . 161
5.5.3 Effect of aperture on spatial resolution . . . 163
5.6 Repeatability of acoustic wavelength calibration . . . 167
5.7 Comparison of detectors . . . 171
5.8 Repeatability of assembly . . . 176

5.9 Conclusions . . . 179
5.10 Future work . . . 180
5.10.1 Equator 700 motion . . . 180
5.10.2 Repeatability of assembly . . . 180
5.10.3 Twin KED . . . 181
5.10.4 SRAS performance data consolidation . . . 182

6 Machine vision for autofocus 183
6.1 Tolerance of SRAS to defocus . . . 185
6.2 Methods for Determining Focus . . . 189
6.2.1 Tenengrad . . . 190
6.2.2 Laplacian . . . 190
6.2.3 Gauss derivative . . . 191
6.2.4 Grey level variance . . . 191
6.2.5 Grey local level variance . . . 192
6.2.6 Gradient energy . . . 192
6.2.7 Windowed Fourier transform . . . 193
6.3 Focus method comparison . . . 194
6.4 Determination of acoustic wavelength from camera images . . . 201
6.5 Application to full SRAS scans . . . 207
6.6 Conclusions . . . 211
6.7 Further work . . . 211
6.7.1 Application to live-scans . . . 211

7 Generation of laser ultrasound on anisotropic materials 213
7.1 Motivation for developing the generation model . . . 215
7.1.1 Issues seen with the orientation determination algorithm . . . 215
7.1.2 Low generation efficiency crystal grains . . . 219
7.2 Existing methods of modelling generation of laser ultrasound . . . 219
7.2.1 Isotropic media . . . 220
7.2.2 Anisotropic media . . . 224
7.2.3 Finite element . . . 228
7.3 The generation model . . . 228
7.3.1 Finite element model of a point-source . . . 228

7.3.2 Modelling a grating . . . 231
7.3.3 Generating SAW velocity surfaces . . . 233
7.4 Unification of generation and orientation model . . . 236
7.5 Applying the hybrid orientation model . . . 237
7.6 Energy directivity . . . 240
7.7 Conclusions and further work . . . 244

8 Final conclusions **246**
8.1 Instrumentation . . . 246
8.2 Signal processing . . . 247
8.3 Renishaw SRAS system performance . . . 248
8.4 Machine vision for auto-focus . . . 249
8.5 Generation model . . . 249
8.6 Summary . . . 250

A Optical chassis design **251**
A.1 Design . . . 251
A.2 Manufacture . . . 255

B Wavelets **257**
B.1 Form of B-spline wavelets . . . 257

C Elastic constants and tensor rotation **260**
C.1 Elastic constant tensor . . . 260
C.2 Example material elastic constants . . . 261
C.3 Rotation matrix . . . 263

Bibliography **264**

Abstract

In this thesis the progress made in developing the spatially resolved acoustic spectroscopy (SRAS) technology for industrial deployment is presented. SRAS is a non-contact laser ultrasound materials characterisation instrument which uses surface acoustic waves (SAWs) to image and determine surface microstructure on metallic samples. SRAS offers advantages over existing techniques for determining microstructure, primarily speed, large sample sizes and full sample coverage. The technique is developed in four key areas: instrumentation, signal analysis, machine vision for auto-focus and modelling generation of laser-ultrasound.

A novel optomechanical system is presented which allows a compact SRAS instrument to be designed and implemented, and which is suitable for industrial manufacture, assembly and application. The capability of this new instrument is thoroughly assessed, with determinations of the noise floor, spatial resolution, and limitations on SAW acoustic wavelength presented. A review of sophisticated signal analysis techniques, suitable for application to SRAS signals, is undertaken and the performance of each technique quantified by both Monte-Carlo simulation and fitting of experimental data.

To enable scanning of non-planar sample surfaces and the use of low cost motion platforms, a means by which the optimum distance between the optical SRAS instrument and a sample surface can be maintained is described. Based on characterising changes in the imaged grating used to generate SAWs, a method is presented capable of providing absolute position corrections from single images. In addition a range of focal measures are reviewed for application to SRAS for automation of scan configuration.

It has been shown previously that SRAS is capable of determining crystallographic orientation on cubic and hexagonal materials, however some discrepancies in this method exist due to insensitivity to the generation mechanism. In this thesis a numerical model for the laser generation of SAWs on arbitrarily anisotropic materials is presented, and used to generate a family of SAW velocity surfaces for cubic nickel. The existing model is combined with the generation model to create a hybrid for use in SRAS orientation determination.

Acknowledgements

I would like to express my utmost gratitude to the Applied Optics group at Nottingham, in particular Steve, Mike, Wenqi, Rob, and Richard, without whom I wouldn't have made it past the first 3 months of the project. To Steve, my academic supervisor, I am truly indebted for the enthusiasm, wisdom and patience he has shown, but most of all for the boundless encouragement. Thank you also to Dr. Xuesheng Chen for his assistance with setting up the generation model, and to Steve and Wenqi again for allowing me to use their SRAS scan data, orientation algorithms, and orientation data.

Many thanks go to my colleagues at Renishaw for allowing me to play engineer for a few years, and the unwavering supply of biscuits, coffee and good advice. Without the advice of the experienced design and software engineers, the Renishaw SRAS prototype would never have got of the ground and I would like to express my gratitude for their support. From Renishaw, particular thanks to Colin who has been an exemplary mentor guiding me smoothly from the academic to industrial world.

My cohort at the RCNDE CDT - Adam, Ailidh, Anton, Oli, and Rob - have been invaluable for making our various courses and conferences far more bearable than they would have been otherwise, and for inspiring me, with their own achievements, to aim higher.

My thanks also go to Chris, Keith and Peter at RCNDE for making the whole project possible.

And of course thank you to my friends and family for the support shown at all stages of the last few years.

List of acronyms

- **AOG** Applied Optics Group
- **CAMPI** Confocal Fabry-Pérot interferometer
- **CT** Computed tomography
- **DFM** Design for manufacture
- **EBSD** Electron back-scatter diffraction
- **EFL** Effective focal length
- **EMDA** East Midlands development agency
- **FEM** Finite element methods
- **GC** Gaussian convolution
- **GDER** Gaussian derivative
- **GLLV** Grey level local variance
- **GLV** Grey level variance
- **GrET** Gradient of energy threshold
- **GT-PBS** Glan-Taylor polarising beam splitter
- **IG** Iterative Gaussian
- **IR** Infrared
- **KED** Knife-edge detector
- **LAPV** Laplace variance
- **LFAM** Line focus acoustic microscope

- **LMS** Least-mean square
- **MFP** Maximum frequency peak
- **MOPA** Master oscillator power amplifier
- **NIR** Near infrared
- **ND** Neutral density
- **OBD** Optical beam deflection
- **PBS** Polarising beam splitter
- **PSAW** Pseudo-surface acoustic wave
- **QWP** Quarter wave-plate
- **RMS** Root mean square
- **SAM** Scanning acoustic microscopy
- **SAW** Surface acoustic wave
- **SEM** Scanning electron microscope
- **SKED** Speckle knife-edge detector
- **SLM** Spatial light modulator
- **SRAS** Spatially resolved acoustic spectroscopy
- **STFT** Short-time Fourier transform
- **SCOE** Surface centre of expansion
- **WFT** Windowed Fourier transform
- **XRD** X-ray diffraction

Chapter 1

Introduction

1.1 Spatially resolved acoustic spectroscopy technology development

This EngD thesis concerns the general development of the spatially resolved acoustic spectroscopy (SRAS) technology and outlines the process of taking the technique from a laboratory demonstrator to an industrial prototype. Multiple aspects of the SRAS technology have been developed and these improvements form the general theme of the thesis. Consequently there is no 'main chapter' discussing one facet of this process, instead a series of pieces of work are presented, all of which have the overarching aim of improving the robustness, capability, manufacturability, and usability of SRAS. The underlying intent of the SRAS technology development project, and the industrial motivation, is to construct a functioning industrial prototype. This is achieved through new understandings supported by systematic scientific research, and technologies realised through a systematic design process.

The progress made in the project is based around four key areas, and an improved characterisation of the technology:

- **Chapter 3** - development and implementation of instrumentation. This chapter examines the optical, and optomechanical, design of the industrial instrument paying close attention to manufacturability and ease of assembly. After a process of experimentation, redesign and refinement a near production ready optome-

chanical prototype, based on a novel optical mounting system and integrated with a motion platform, is presented.

- **Chapter 4** - applying various signal analysis techniques to both simulated and experimental SRAS data. This chapter focusses on assessing the performance of various techniques in order to improve the robustness of determination of phase velocity for noisy signals. Determination of the limits on phase velocity measurement resolution is presented with optimisation of signal analysis techniques to allow real-time 'on-the-fly' capture and processing of data.

- **Chapter 5** - a thorough review of the performance of the Renishaw SRAS instrument. This includes: spatial resolution, effects of changing laser generation power and repetition rates, comparisons of different acoustic wavelengths and apertures, and an assessment of the mechanical assembly. This chapter aims to critically probe the capability and limitations of the Renishaw SRAS system.

- **Chapter 6** - investigating the use of image analysis techniques to facilitate autofocus. This chapter examines the feedback from camera images of the scanned surface in order to apply automated tracking of non-planar surfaces, and to correct motion platform instability. In addition determination of the sensitivity of SRAS measurements to sub-optimal motion platforms is assessed.

- **Chapter 7** - modelling the laser ultrasonic generation process on anisotropic materials. This chapter aims to improve understanding of the physical process occurring at the point of generation on the sample surface. Results of the generation model are applied to inform and improve the orientation fitting algorithms.

These areas were identified as priorities partially because they are under represented in the research literature and also because they are critical areas in improving the robustness of the SRAS technique. In general each of these chapters can be considered a self-contained piece of work with its own conclusions and discussion of further work, contributing to the overarching aims of the project. In addition to these key areas the following supporting chapters are also presented:

Chapter 1 - a general overview of the background science required for understanding of this project. This includes a brief foray into competitive and complementary techniques of SRAS and a discussion of the material microstructure to which SRAS is

sensitive.

Chapter 2 - a review of the state of the SRAS technique at the commencement of this thesis project. This includes the essential theory of SRAS, a description of the East Midlands Development Alliance (EMDA) SRAS system at the University of Nottingham, and progress made in crystallographic orientation determination by SRAS.

Chapter 8 - the final conclusions. This chapter summarises the work done during the project and details in which ways the SRAS technique has been improved.

The remaining portion of this chapter will cover: a brief review of laser ultrasonic techniques including detection methods, generation methods and ultrasonic waves, and a description of crystallography, microstructure and existing techniques for the determination of microstructure.

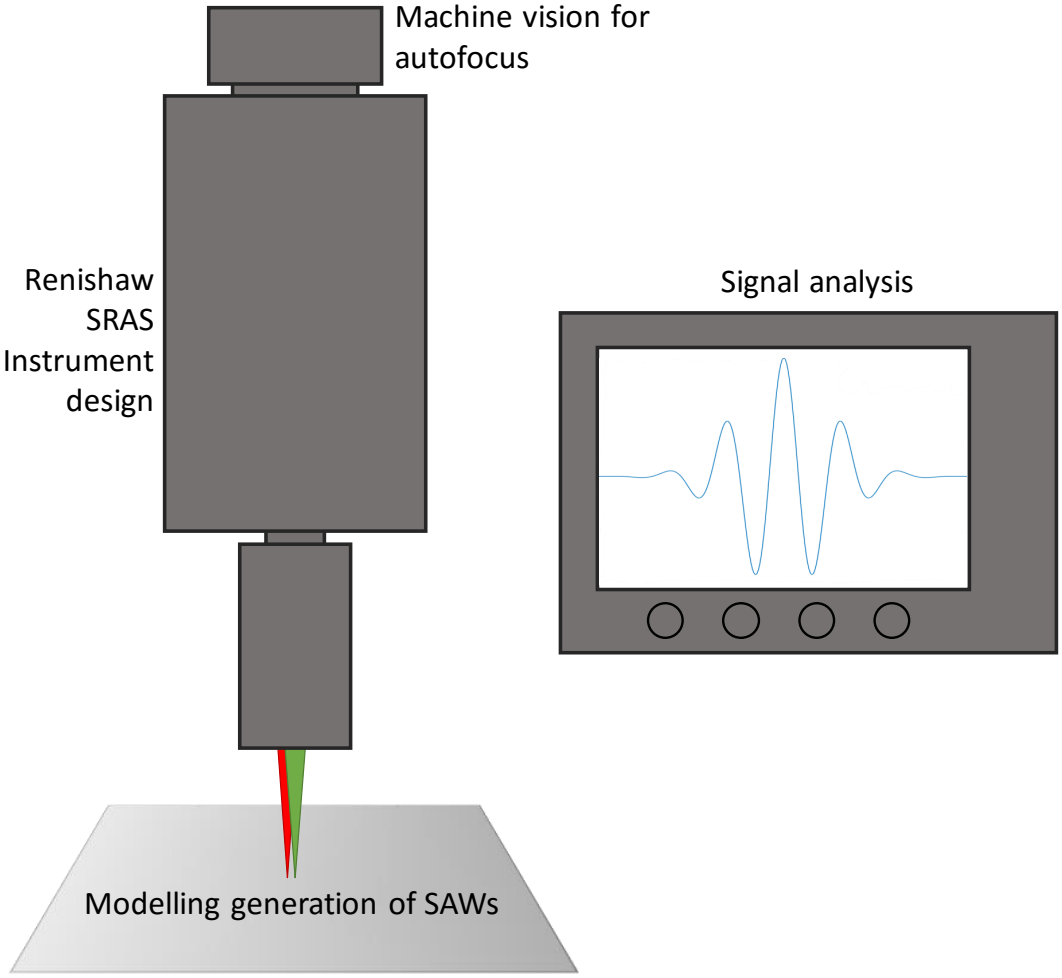


Figure 1.1: An illustration of the primary areas of research and engineering contributing to the overarching aims of the project.

1.2 Laser ultrasonics

1960 saw the first report of a functioning laser which was initially described as 'a solution looking for a problem' but rapidly spawned a vast array of applications in science and technology [1]. Lasers are a unique source of light with characteristics which make them extremely useful and versatile. Chiefly, this is their ability to generate an intense, highly collimated beam of coherent light at a well-defined wavelength. Modern engineering is capable of building lasers at wavelengths from far infra-red to mid ultra-violet, at powers from milliwatts to petawatts and with pulses widths down to femtoseconds [2, 3]. It was not until the late 1970s that lasers became a viable option for the generation and detection of ultrasound. From an industrial perspective, laser ultrasound is an extremely useful technique; it combines the probing depths of ultrasound with the non-contact, and potentially non-destructive, characteristics of optical systems. Since early development there has been much progress in the field, but it was not until the late 1990s that commercial laser ultrasound systems were readily available [4].

The following section offers a brief overview of the principle behind laser ultrasound. For a much more detailed description see reference [5].

1.2.1 Ultrasonic waves

In the same way that the tone of a struck wine glass can tell you its quality, the application of acoustics in materials can divulge a plethora of properties. Whilst tapping a wall to locate a stud is a useful if crude technique, the ultrasonic range allows significantly more finesse. The spectrum of human hearing is from approximately 20 Hz to 20 kHz, and it is sound waves at frequencies higher than this that are referred to as ultrasonic.

Acoustic waves can be generated in any form of matter by particle vibration. These waves then propagate in four modes distinguishable by the particle's oscillation direction: longitudinal, transverse, plate and surface [6]. All of these acoustic wave types can be derived from the elastodynamic wave equation for isotropic media, also known as the Navier-Green equation [7],

$$(\lambda + \mu)\nabla(\nabla \cdot \mathbf{u}) + \mu\nabla^2\mathbf{u} + \mathbf{F} = \rho\frac{\partial^2\mathbf{u}}{\partial t^2}, \tag{1.1}$$

where λ is Lamé's first parameter, μ is the shear modulus, F is the external force acting on the solid, ρ is the material density, $\mathbf{u}$ is the particle motion, and t is time.

Assuming an infinite media, neglecting external body forces, and applying the divergence operator to equation 1.1, along with the vector identity $\nabla^2\mathbf{u} = \nabla \cdot \nabla\mathbf{u}$, gives [8]

$$\nabla^2\vartheta = \frac{1}{c_l^2}\frac{\partial^2\vartheta}{\partial t^2}, \tag{1.2}$$

where $\vartheta = \nabla \cdot \mathbf{u}$. This is the scalar wave equation describing dilatational, or longitudinal, waves where the wave speed, c_l, is given by

$$c_l = \sqrt{\frac{\lambda + 2\mu}{\rho}} = \sqrt{\frac{E(1-\nu)}{\rho(1+\nu)(1-2\nu)}}, \tag{1.3}$$

where E is Young's modulus and ν is Poisson's ratio. The particle motion of longitudinal waves and the propagation direction are parallel. This motion results in sinusoidally varying areas of high and low particle density, and for this reason longitudinal waves are also known as compression waves. Examples of longitudinal waves include sound travelling through air and seismic P-waves.

Assuming an infinite media, neglecting external body forces and applying the curl operator to 1.1 gives [8]

$$\nabla^2\Omega = \frac{1}{c_t^2}\frac{\partial^2\Omega}{\partial t^2}, \tag{1.4}$$

where $\Omega = \nabla \times \mathbf{u}$. This is the scalar wave equation describing transverse, or shear, waves where the wave speed, c_t, is given by

$$c_t = \sqrt{\frac{\mu}{\rho}} = \sqrt{\frac{E}{2\rho(1+\nu)}}. \tag{1.5}$$

The propagation direction of transverse waves is perpendicular to the particle motion. In solids energy is transferred from particle to particle by a shear force, hence the

name shear waves. Examples of transverse waves include a plucked guitar string and seismic S-waves. This wave mode cannot exist in liquids and gasses; the shear force is not transmitted due to the much greater inter-atomic spacing, although combined longitudinal-transverse wave modes are supported in liquids [9].

If instead of an infinite media an elastic half-space, shown in Fig. 1.2, is assumed a new family of solutions to the elastodynamic wave equation are available which propagate along the free surface. One of these wave types are called Rayleigh waves and are a combination of both dilatational and transverse components. By assuming an exponential wave form constrained to the surface and applying the stress-free surface condition, the secular equation describing Rayleigh wave phase velocity, c_r can be derived [10],

$$\left(2 - \left(\frac{c_r}{c_t}\right)^2\right)^2 = 4\sqrt{1 - \left(\frac{c_r}{c_l}\right)^2}\sqrt{1 - \left(\frac{c_r}{c_t}\right)^2}. \tag{1.6}$$

By eliminating the trivial case of $c_r = 0$ equation 1.6 can be rearranged into a cubic form [10],

$$p^3 - 8p^2 + (24 - 16q)p - 16(1 - q) = 0, \tag{1.7}$$

where

$$p = \left(\frac{c_r}{c_t}\right)^2 \text{ and } q = \left(\frac{c_t}{c_l}\right)^2. \tag{1.8}$$

For the trivial case of an isotropic incompressible solid, where $\nu = 0.5$ and $c_l = \infty$, equation 1.7 has only one real, positive root

$$\left(\frac{c_r}{c_t}\right)^2 \approx 0.9126. \tag{1.9}$$

However, for the more common non-trivial case of compressible solids the Rayleigh wave velocity is more difficult to determine. For these cases equation 1.7 is solved by numerical methods which aim to minimise equation 1.7 and determine the real roots. Examples of this procedure can be found in references [11] and [12].

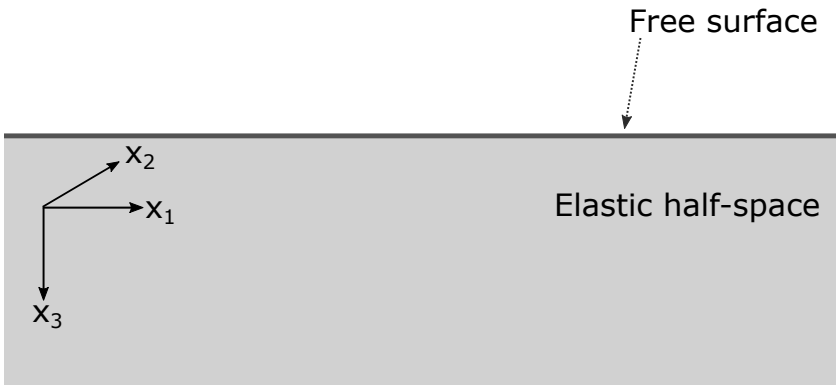


Figure 1.2: An elastic half-space with a free surface. x_1 and x_2 extend infinitely.

The Rayleigh wave is composed of dilatational and transverse components which decay exponentially with depth at different rates. Because of this the particle motion traces an ellipse with eccentricity and direction a function of depth, illustrated in Fig. 1.3. Typical velocities in engineering metals are $2000\,\mathrm{m\,s^{-1}}$ to $4000\,\mathrm{m\,s^{-1}}$ and are always smaller than the transverse and longitudinal wave velocities. True Rayleigh waves do not propagate any energy away from the surface of the half-space, however there exist a set of similar waves which do lose energy into the bulk media; these are known as pseudo surface acoustic waves (PSAWs) or 'leaky' waves [13]. These share the same solutions to the wave equation as Rayleigh waves but with a very small non-zero term perpendicular to the surface. PSAW are also non-dispersive but, unlike pure SAWs, can sometimes exceed the shear wave velocity. Since the leakage term is small over small distances it can typically be ignored, and these modes can be functionally treated as Rayleigh waves [14].

Surface waves exist in two distinct families: Rayleigh and Love [13]. Love waves travel along the surface of solids with transverse particle motion in the plane of the surface, they will not be discussed here in any detail but more information can be found in [15].

The fourth family of ultrasonic waves are plate waves, or Lamb waves, which only propagate through solids of finite thickness. By assuming a solid of thickness $2d$ and infinite extent in x_1 and x_2, as shown in Fig. 1.2, and applying the correct boundary conditions to the elastodynamic wave equation the following characteristic equations can be found [16]

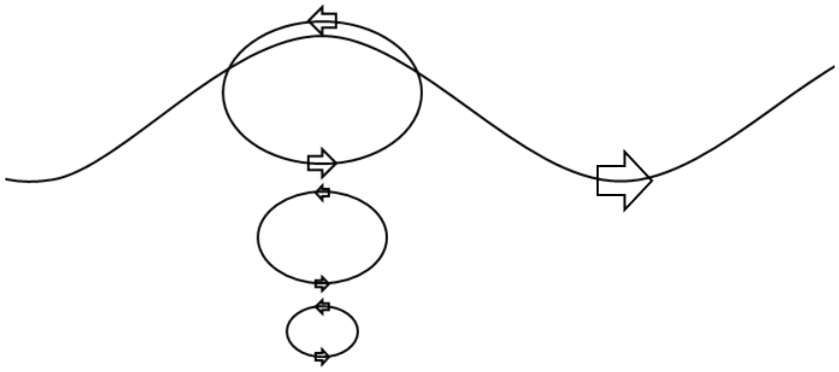


Figure 1.3: A right travelling Rayleigh wave. At the surface particle motion follows a retrograde ellipse with amplitude decreasing exponentially with depth. At a depth of approximately one fifth the wavelength, the particle motion inverts to a prograde ellipse.

$$\frac{\tanh(\beta d)}{\tanh(\alpha d)} = \frac{4\alpha\beta k^2}{(k^2+\beta^2)^2}, \tag{1.10}$$

and

$$\frac{\tanh(\beta d)}{\tanh(\alpha d)} = \frac{(k^2+\beta^2)^2}{4\alpha\beta k^2}, \tag{1.11}$$

where k is the wave-number,

$$\beta^2 = k^2 - \frac{\omega^2}{c_l^2},\ \alpha^2 = k^2 - \frac{\omega^2}{c_t^2}, \tag{1.12}$$

and ω is the angular frequency.

Lamb wave particle motion is in the plane defined by the plate normal and the propagation direction. Wave modes are either symmetrical or antisymmetrical, relative to the centre of the plate, with equation 1.11 describing antisymmetrical modes and equation 1.10 describing symmetrical modes. Although the form of equations 1.10 and 1.11 are fairly simple, there are an infinite number of solutions which can only be found through numerical methods [16]. Lamb waves velocities are a function of acoustic wavelength and plate thickness. In the regime where the acoustic wavelength is much smaller than the plate thickness the Lamb wave is converted into a true Rayleigh wave [17]. This can be seen in Fig. 1.4 where the A0 and S0 mode velocities converge to the Rayleigh wave velocity as frequency increases or thickness decreases.

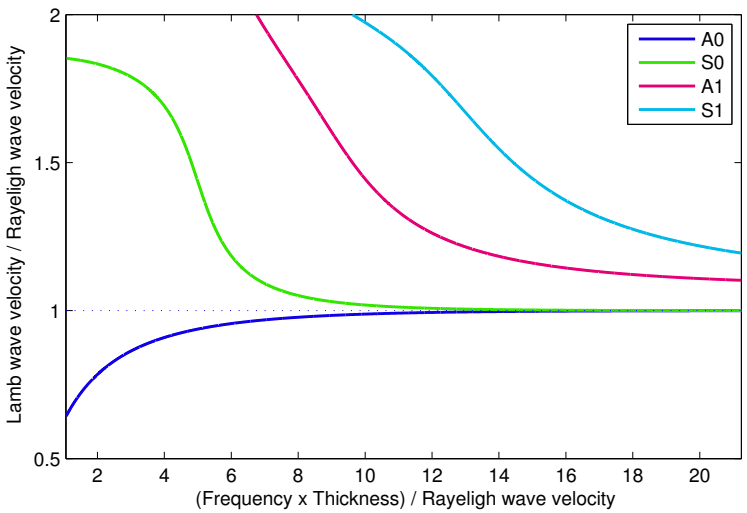


Figure 1.4: Dispersion curves for the first two symmetrical and antisymmetrical Lamb wave modes. Both axes have been normalised to the Rayleigh wave velocity.

1.2.2 Generation mechanisms

Lasers are capable of remotely generating multiple types of acoustic waves on solid surfaces. There are primarily two regimes, in addition to one special case, of laser-generated ultrasound dependent on the laser power, pulse rate and sample absorptivity [5].

Ablation

Laser energy is focused on the sample with sufficient energy density to form a plasma. The plasma is made up of a mixture of the evaporated sample and air breakdown, with a ratio dependent on the material and experimental conditions. Plasma ejection from the surface generates a reactive stress force which produces bulk acoustic waves with a large component directed perpendicular to the surface [18]. The ablative regime is by far the most useful for bulk material testing such as thickness measurements. Dependent on the application this regime can be considered non-destructive as the damaged region is highly localised, however this is essentially a destructive process.

Thermoelastic

Laser energy is focused on the sample but with energy densities lower than those capable of melting the sample. Within the material skin depth, this energy is absorbed as heat causing a local thermal expansion. This causes stresses and strains between the illuminated and unilluminated parts of the sample which in turn generates acoustic waves. These waves have large components parallel to the surface and very little energy is directed into the sample [18]. Thermoelastic generation is a completely non-destructive process.

There is not an abrupt threshold between the ablative and thermoelastic regimes. When laser energy is increased beyond the thermoelastic regime the illuminated region begins to melt and the efficiency of generation of SAWs decreases. As the input energy rises the penetration depth of the melt front increases, and bulk waves are generated with increasing amplitude [19].

Constrained

This special case generation technique employs a transparent coating which acts to constrain the surface, increasing the amplitude of longitudinal, shear and Rayleigh waves [20]. Typically energy densities in the thermoelastic range are used making this process non-destructive.

1.2.3 Source types

Generation of laser ultrasound is typically achieved using temporally short and high power laser pulses. There are numerous methods by which the required pulses can be generated and regulated. Passive mode-locking is one method to achieve this and is capable of producing ultra-short pulses, however repetition rates are fixed based on the design of the laser and typical pulse energies are of order nJ. Active Q-switching allows for much lower pulse rates which are readily adjustable. A Q-switch is a device placed in the laser cavity which acts to prevent reflection of light back to the gain medium. This prevents feedback and, as pumping continues, the energy density approaches a maximum, set by spontaneous emission and other losses. The Q-switch is then

activated to allow feedback to the gain material and a high intensity laser pulse is produced. Q-switches are normally either acousto-optic or electro-optic devices [3, 5]. An alternative to Q-switch devices is modulation of continuous-wave sources. This is achieved by applying a pseudorandom sequence to the bias current of a laser diode. The long signals and custom modulations can improve signal to noise, but the efficiency of generation compared to Q-switch generated pulses is poor [21, 22]. Active Q-switching is the most commonly used of the source types.

1.2.4 Optical detection

There are many optical detection methods available and only a thoroughly condensed account will be given here. Chapter 3 of [5] provides an excellent overview as does [23]. Although not entirely up to date, [24] offers good coverage of the scientific theory behind many of the technologies.

All optical SAW detectors require a laser beam incident on a sample surface, and some method to collect and analyse the reflected beam. To avoid signal averaging, the beam diameter is typically focused to a point smaller than half the SAW wavelength and the power density is kept low enough to avoid melting or ablating the sample. The induced temporal, spatial and phase modulations of the reflected beam are then analysed. There are several different approaches to optically detecting surface waves and these rely on one or more of the following modulations.

1. Focusing/defocusing due to the curvature of the SAW.
2. (for large beam diameters) temporally/spatially periodic phase corrugation, resulting from the pattern of acoustic displacements, this is the basis of the surface grating technique [25].
3. Beam deflection due to the SAW gradient, this is the basis of optical beam deflection (OBD) techniques.
4. Phase modulation dependent on the magnitude of the out of plane displacement of the SAW, this is the basis of interferometry.

For the purposes of SRAS, and with the ultimate aim of rough surface sample applications, the fourth modulation, i.e. the temporal/spatial periodic phase corrugation

technique, is not particularly viable as the diffraction pattern is heavily dependent on the surface finish. In addition the smaller probe beam diameters available in the other techniques is beneficial for minimising spatial resolution. For this reason interferometry and optical beam deflection are the primary choices of technique.

1. SAW focusing/defocusing

The SAW focusing/defocusing technique relies on the local curvature of the SAW acting as a concave or convex mirror, dependent on the phase [26]. A basic schematic of this detection technique is shown in Fig. 1.5. The probe beam is brought to a focus a short distance above the sample surface. As a SAW propagates beneath the probe beam the reflected beam is either focused, or defocused, by the SAW curvature. The effect of this is a modulation in the beam intensity which can be detected by a photodetector. This technique requires excellent surface preparation and excellent isolation from external sources of vibration [27].

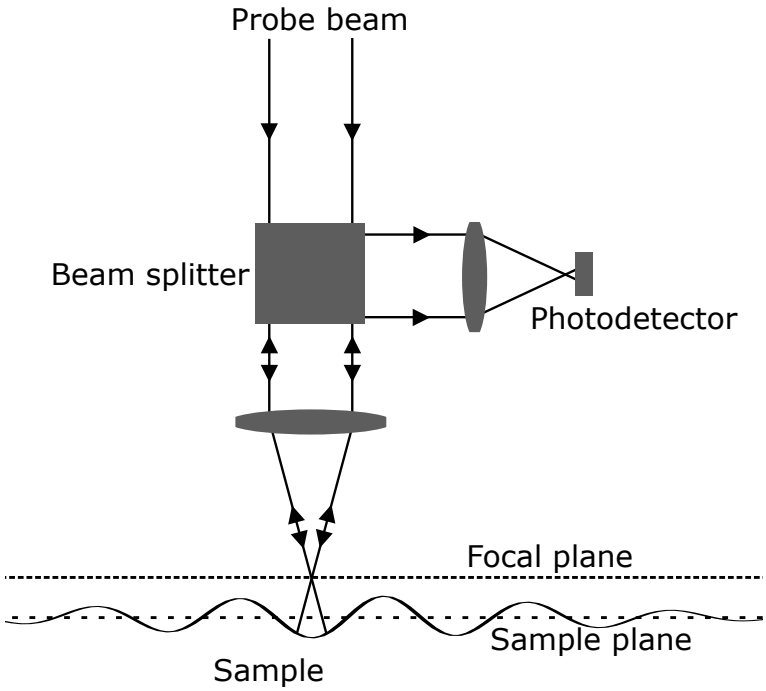


Figure 1.5: Schematic of a focusing/defocusing based SAW detector.

2. Surface grating technique

The surface grating technique deploys a probe beam diameter larger than the SAW acoustic wavelength. As the SAW passes under the detection spot it acts as a grating and the reflected beam is diffracted such that it is split into the zeroth order beam and two first order beams which are diffracted by an angle $\pm\theta_d$. This angle is dependent on the SAW acoustic wavelength and laser wavelength [24]. The interference of these three beams produces an intensity modulation which can be detected through a suitable periodic grating. A schematic of this technique is shown in Fig. 1.6. This technique requires a polished sample and is typically used for continuous SAW detection [23].

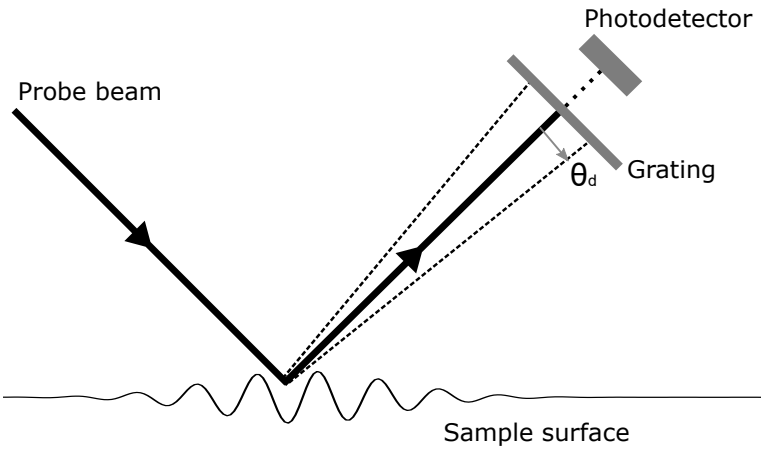


Figure 1.6: Schematic of a surface grating SAW detector. Note that the probe beam diameter must be larger than the SAW acoustic wavelength.

3. Optical beam deflection

As a SAW wave propagates under the detection point, the reflected beam is deflected by an angle twice the local wave gradient. A lens then converts this angular deflection into a spatial deflection, which is detected as an intensity variation by placing a knife-edge in the beam path, shown in Fig. 1.7. This technique is known as optical beam deflection (OBD). Since the gradient is dependent only upon the out of plane motion, the knife-edge detector (KED) is only sensitive to this modulation. With the use of a physical knife-edge blocking the detection beam the etendue, or light-gathering ability, of the detector is immediately halved. By using two photodiodes, with the beam split evenly between them, the same measurement can be made with no wasted

light. Modern photodiodes have excellent sensitivity and the limiting factor for the KED is the RMS shot noise. KEDs are cheap, adapt very quickly to surface changes, are tolerant to mechanical vibration and have increased sensitivity with increasing SAW frequency. However they rapidly lose sensitivity for rough surfaces due to the distorted reflected beam and require a highly stable single mode probe laser.

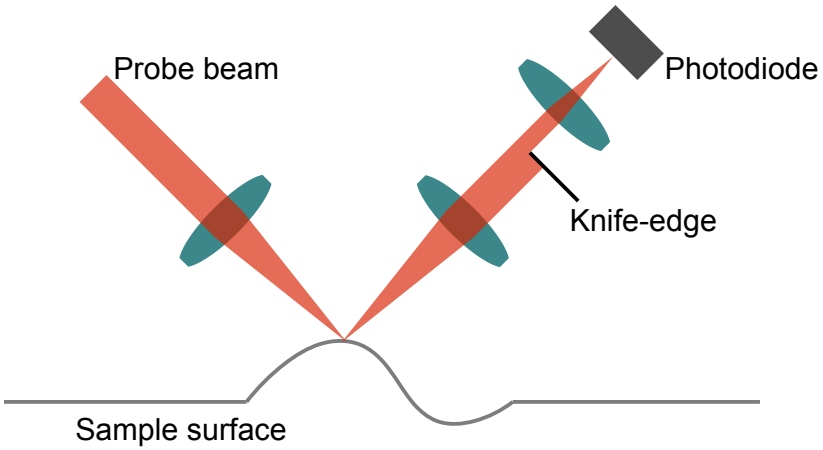


Figure 1.7: Schematic of a KED based on a physical knife-edge including the required optical components.

4. Interferometry

All interferometric techniques exploit the distortion of the wavefront in a few distinct ways illustrated in Fig. 1.8. These techniques involve interfering a wave distorted by reflection from a moving surface with: a clean reference beam from the same source, a second reflected beam or a time-delayed version of the reflected beam. The sensitivity of most interferometer types is fairly similar [5], and the choice of device is typically dependent on the application and surface condition. Whilst a vast array of exotic interferometric techniques exist, only Michelson and velocity methods will be described here as they are the most suitable for application to laser ultrasound techniques.

Michelson interferometer

The configuration of a Michelson interferometer is shown in Fig. 1.9. A beam splitter sends the input laser beam along two orthogonal paths, one directed at a mirror and the other at the sample, via a focussing lens. The two reflected beams are then

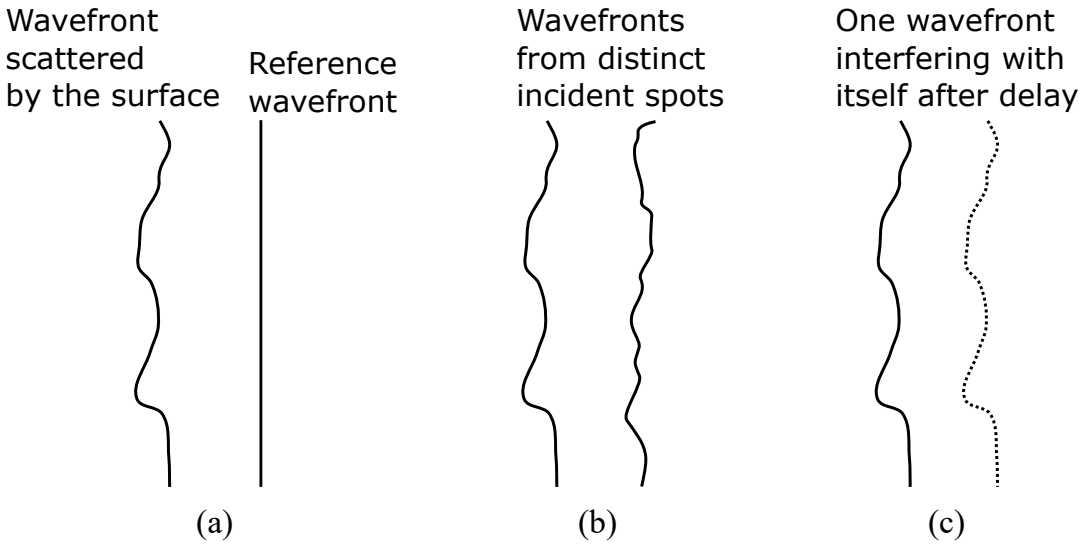


Figure 1.8: The comparison beam wavefronts used in the three main methods of interferometry. (a) Optical heterodyning where a distorted beam is interfered with a reference beam. (b) Differential interferometry where two distorted beams taken from different points on a sample surface are interfered. (c) Velocity or time-delay interferometry where a distorted beam is interfered with a time delayed version of itself. Adapted from [23].

recombined by the beam splitter and steered into a photodetector. Since the SAWs distort the sample surface, the path length of the sample beam will vary periodically. By interfering the sample beam with the stable reference beam, the slight changes in amplitude and phase can be detected and the surface motion deduced. This technique is sensitive to a wide frequency range, approximately 50 kHz to hundreds of MHz, and can detect surface motion at the order of Ångstroms.

Velocity Interferometry

This technique interferes a distorted beam with a time-delayed version of itself. The principle is based on the Doppler shift of the reflected beam caused by the modulating surface wave. One of the most promising devices to do this with is the Confocal Fabry-Pérot Interferometer (CFPI), shown in Fig. 1.10 [28, 29]. The CFPI cavity is comprised of two partially reflective concave mirrors. As light enters the cavity it is reflected off of each mirror multiple times. These mirrors allow for interference between the multiple reflections of the wavefronts within the cavity, with maximum

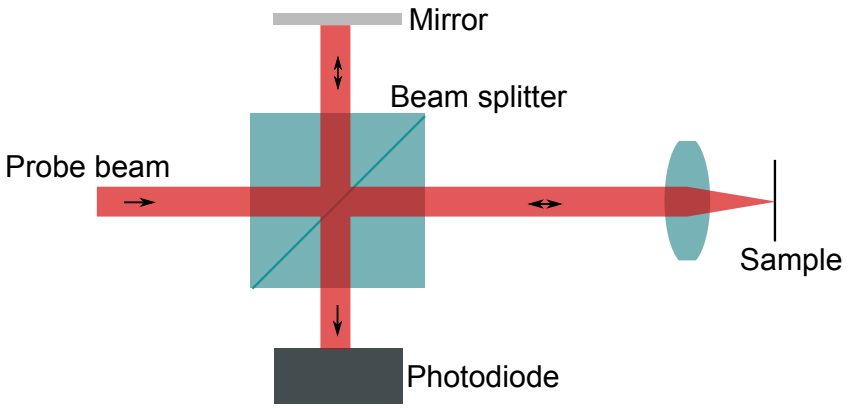


Figure 1.9: A schematic of a simple Michelson interferometer.

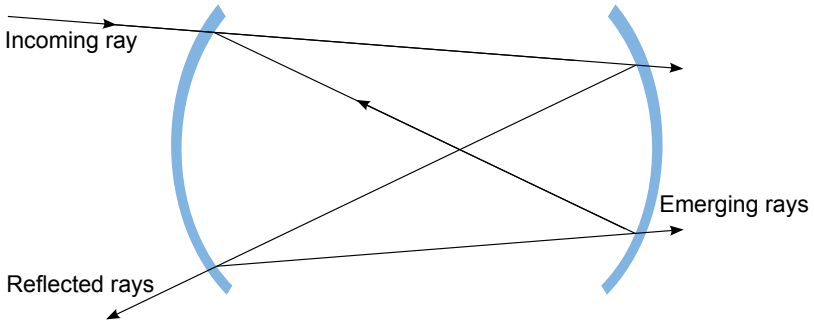


Figure 1.10: A Confocal Fabry-Pérot Interferometer cavity. Only a few rays have been shown for clarity.

transmission of light occurring when the cavity length is equal to an integer multiple of the input beam wavelength. As the SAW induces a modulation in the frequency of the detection beam, the transmission is modulated and this can be detected by a photodiode. The cavity length can be tuned to match the unmodulated wavelength of the detection beam enabling high detection sensitivity.

There are devices that do a similar task, such as the planar Fabry-Pérot interferometer, but the curved mirrors of the CFPI offer a much greater etendue due to its relative insensitivity to the incoming ray angle. Evidently, since the reflected beam is interfered with itself, distortion is unimportant and the device is capable of detecting SAWs from surfaces which generate speckle patterns in the reflected probe beam. The CFPI has the disadvantage of being very difficult to build, requires high stability in the cavity length and needs to be designed for optimisation at specific frequency ranges. As a result the CFPI is a very expensive device with typical costs of ~£80,000.

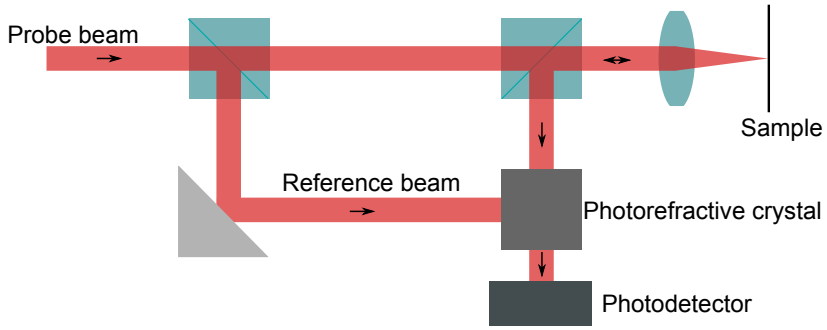


Figure 1.11: A two-wave mixer schematic using a reference beam and photorefractive crystal to generate an interference pattern.

Two-Wave Mixing

A two-wave mixer is built around a photorefractive crystal, the general configuration is shown in Fig. 1.11. A laser probe beam is directed at a beam splitter which splits it in two: the sample beam, and the reference beam which is directed into the photorefractive crystal. The sample beam is reflected from the sample and directed into the crystal, carrying with it the phase information encoded by the sample surface motion. In the crystal interference between the reference and sample beam cause a pattern of dark and bright fringes. In the bright fringes electrons absorb photons and are excited into the conduction band. Currents are thus created flowing from the bright to dark fringes, and the induced electric field alters the local refractive index of the crystal. The result is a diffraction grating which directly matches the pattern of the incident sample beam, and imposes the same pattern on the reference beam [30]. The two spatially identical beams can then be detected and the phase information extracted. This technique is excellent at coping with distorted reflected beams and has a sensitivity similar to the other interferometric techniques. Due to the slow adaptation rate of the photorefractive crystal this technique has limited scanning speed over rough surface samples.

1.3 Crystallography

Equations 1.3 and 1.5, describing the longitudinal and shear wave velocities, are valid for isotropic materials. However, a significant proportion of engineering metals are,

to varying degrees, anisotropic and these equations will return results only valid on scales much larger than the typical crystal grain size. Anisotropy arises from the basic crystal structure of the material, and the variations of the elastic constants on different crystallographic planes. The elastic constants, also called the stiffness matrix, of a material are linked to the stress, σ, and strain, ϵ, 3X3 tensors by Hooke's Law, $\sigma_{ij} = C_{ijkl}\epsilon_{kl}$ [31], where C_{ijkl} is a fourth rank-tensor with 81 independent elements. By applying symmetry arguments and employing a new index notation, a procedure described in detail in Appendix C, the elastic constant tensor can be reduced to a 6X6 matrix, C_{ij}, and Hooke's law can be written as

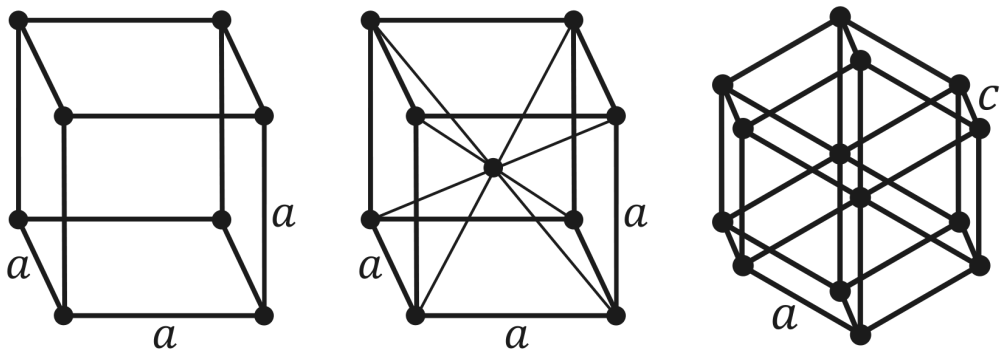
$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{21} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{31} & C_{32} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{41} & C_{42} & C_{43} & C_{44} & C_{45} & C_{46} \\ C_{51} & C_{52} & C_{53} & C_{54} & C_{55} & C_{56} \\ C_{61} & C_{62} & C_{63} & C_{64} & C_{65} & C_{66} \end{bmatrix} \times \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \epsilon_{yz} \\ \epsilon_{zx} \\ \epsilon_{xy} \end{bmatrix}. \tag{1.13}$$

Since $C_{ij} = C_{ji}$ this matrix contains only 21 independent elastic constants.

A material's crystal structure strongly affects its macroscopic properties. Chemically identical materials can often exist with different crystal structures known as phases, these materials are referred to as polymorphs. For example, tin can exist in both white (β) and grey (α) phases at low room temperatures. While white tin is malleable, grey tin is highly brittle and undesirable; the conversion of tin from the β to α phase was known as the 'tin blight' and is responsible for the decomposition of tin at low temperatures [32].

Crystalline materials form a regular grid, or lattice, of atoms and there are seven possible lattice systems and fourteen possible Bravais lattices, a full description of which can be found in reference [31]. The lattice system describes the shape of the unit cell whilst the Bravais lattice describes the discrete translation of the unit cell which makes up the periodic structure of the crystal. Fig. 1.12 shows the cubic (a, b) and hexagonal (c) lattice systems. For the cubic lattice the primitive (a) and body-centred (b) Bravais lattice types are shown.

To identify the particular plane of the crystal structure under examination, the Miller



(a) (b) (c)

Figure 1.12: a) A primitive cubic unit cell, (b) a body-centred cubic unit cell and c) a hexagonal unit cell.

indices convention is adopted. The vector [hkl] describes a direction and is derived from integers of the lattice vectors, i.e. $ha_1 + ka_2 + la_3$, where a_i are the vectors defining the unit cell. The plane, (hkl), is then a plane orthogonal to the vector and {hkl} denotes all planes equivalent to (hkl).

Dependent on the symmetry of the crystal structure of a material, the elastic constant matrix can be further simplified. In the most extreme case of a triclinic crystal structure, which has minimal symmetry, 21 independent elastic constants are required to describe the material. For the opposite case of an isotropic material only 2 independent elastic constants are required. Cubic and hexagonal crystal structures can be described by 3 and 4 independent elastic constants respectively, and the full form of the elastic constant matrices for these structures is given in Appendix C. Along with the symmetry, a useful measure of the degree of anisotropy of a material is described by the anisotropy ratio, η, such that

$$\eta = \frac{2C_{44}}{C_{11} - C_{12}}. \tag{1.14}$$

Isotropic materials have an anisotropic ratio of $\eta = 1$, anisotropic materials have $\eta < 1$ or $\eta > 1$.

Measurement of elastic constants is most commonly achieved by generating and detecting ultrasound through a material sample, and relating the ultrasonic wave velocity

to the elastic constants. For the case of an isotropic material this process is trivial. Equations 1.3 and 1.5, describing the longitudinal and shear wave velocities, can be combined to determine the Lamé moduli through measurement of the transverse and longitudinal wave velocity

$$\mu = c_t^2 \rho \text{ and } \lambda = \rho(c_l^2 - 2c_t^2). \tag{1.15}$$

The Lamé moduli can then be related back to the elastic constants since $\mu = \frac{C_{11}-C_{12}}{2}$ and $\lambda = C_{12}$ [31].

For anisotropic samples this process is more difficult as the ultrasonic wave velocities are dependent on the crystal orientation and the direction of propagation. Instead of the simple elastodynamic wave equation, the anisotropic elastodynamic wave equation must be used [33]

$$\rho \frac{\partial^2 u_i}{\mathrm{d}t^2} - C_{ijkl} \frac{\partial^2 u_k}{\partial x_j \partial x_l} = 0. \tag{1.16}$$

For high symmetry materials which can be grown into single crystal samples of size ~10 mm, time-of-flight measurements of ultrasonic wave velocity at multiple propagation directions can be used to extract the elastic constants [34]. For low symmetry, or multiple grain, samples a technique called resonant ultrasound spectroscopy is used [35]. This technique measures the response of a sample to ultrasound over a wide range of frequencies to generate a frequency spectrum with distinct resonant peaks. Numerical methods can then be used to iterate through sets of elastic constants to find the best fit to the experimental spectral data [36].

1.4 Microstructure

Crystallography manifests in macroscopic materials in a variety of ways. A notable example is the hexagonal basaltic columns at Giant's Causeway which are a result of the underlying hexagonal crystallography [37]. More typically, and especially for the case of polycrystalline engineering materials, crystallography results in a material with a granular structure at microscopic scales; this is generally referred to as

the microstructure [38]. There are two primary crystallographic contributions to microstructure: phase and orientation. Phase refers to the particular crystal structure, or lattice, which forms in a material, i.e. cubic, hexagonal, etc. Throughout a sample of uniform chemical composition, the phase can vary with a combination of lattice types forming. The number of phases and lattice types is highly material dependent with some materials existing in only a single phase. Orientation refers to the local plane and rotation of a particular crystal grain, for example the regions of uniform colour in Fig. 1.13 (a) represent regions of constant orientation.

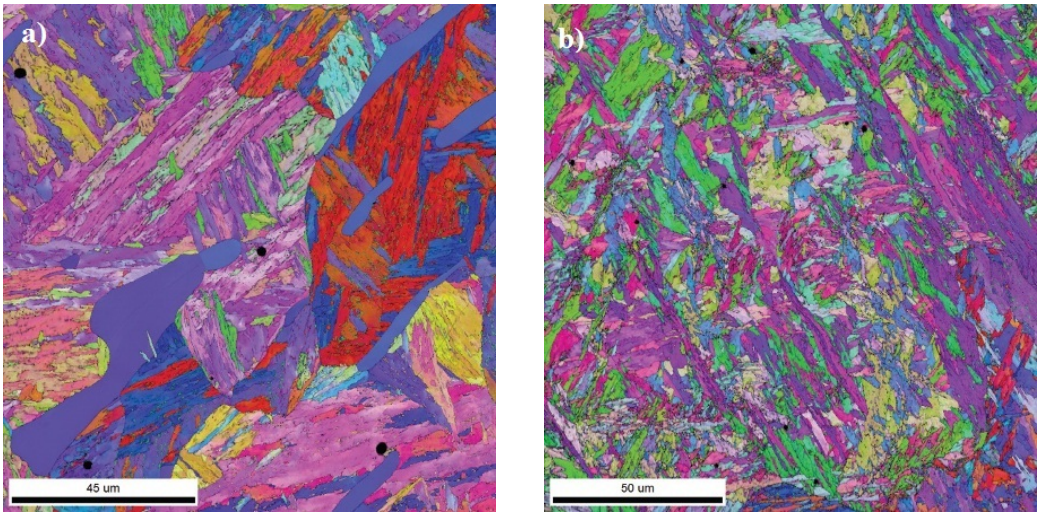


Figure 1.13: Electron backscatter diffraction (EBSD) orientation map of a section of austenitic steel gas turbine inlet guide vane. (a) leading edge (b) trailing edge [39].

There is a vast array of variables which can affect the microstructure of a component or material sample; these include, but are not limited to, [38]:

- chemical composition,
- manufacturing processes (i.e casting, forging, rolling extrusion),
- rate of cooling post-manufacture,
- physical working,
- contaminants in raw material,
- heat treatments.

For example Fig. 1.13 shows images of the leading edge and trailing edge of an inlet guide vane from a gas turbine engine. At the time of manufacture both regions had similar distributions of microstructure. During operation of the engine the different temperatures and pressures experienced at each region has caused divergence in their microstructure.

1.5 Methods for determining microstructure

With the ability of the microstructure of materials to affect the macroscopic properties it is often important to characterise the microstructure. This characterisation can be achieved, at varying resolutions and speeds, through a number of methods. Not all of these methods can provide full crystallography but instead provide a contrast image from which microstructure can be inferred. Very few techniques are capable of determining the subsurface microstructure of a part.

1.5.1 Optical techniques

Typically a polished metal surface will not exhibit its microstructure under brightfield illumination since incident light is reflected uniformly. In order to image the microstructure the differences in reflectivity of crystal grains must be enhanced and there are numerous ways to achieve this. One such method is through novel illumination techniques such as polarised light, phase contrast and differential interference contrast [40]. These techniques are non-destructive but among the optical methods, with the exception of brightfield illumination, offer the least contrast.

Chemical etching involves attacking the surface with an acid oxidising agent, typically HF, which acts to dissolve atoms at the surface. Since the rate of dissolution is dependent on the local microstructure, optical contrast is enhanced [40]. Physical etching involves either sputtering or vapour depositing a thin transparent film on the surface. As with chemical etching, since the rate of deposition of this film is dependent on the local microstructure, optical contrast is enhanced. Physical etching is not as effective as chemical etching and is typically combined with optical illumination methods for best results [40]. All of the etching methods described allow an optical contrast image

of the microstructure to be acquired, but they do not provide quantitative data nor can they provide full crystallographic orientation. Etching can be performed over large regions of a sample and full-field images can be collected very quickly.

An extension to etching is the use of microtomy: a cutting process enabling very thin sections of material to be sectioned. These thin slices are then etched and images captured for each slice. The images are then reconstructed into full 3D models of the sample microstructure [41, 42]. For extremely fine slices, $\leq$50 nm, transmission electron microscopy can be employed enabling microstructure to be resolved down to $\sim$10 nm [43].

1.5.2 X-Ray diffraction (XRD))

X-rays occupy a wavelength range of 0.1 nm to 10 nm in the electromagnetic spectrum. Since the average inter-atomic spacing is a few Ångstroms, X-rays are excellent candidates for diffraction studies. A collimated, continuous spectrum, beam of X-rays is generated by accelerating electrons into a metal target, typically copper. This beam is incident on the surface of the crystal sample where it penetrates a few atomic layers and is diffracted by the atomic electron cloud. For all planes, the spacing, d, and Bragg angle, θ are fixed, shown in Fig. 1.14. Bragg's law states,

$$n\lambda = 2dsin(\theta), \tag{1.17}$$

hence if the continuous spectrum contains the correct wavelength to satisfy Bragg's law, a reflected beam is produced [44]. The reflected beams are then imaged to produce a diffraction pattern which is then manually processed using a Greninger chart; due to multiple scattering effects, these patterns cannot generally be processed automatically [44]. This process is known as Laué beam diffraction. X-ray beam diameters are generally around 2 mm rendering this technique insensitive to crystal grains smaller than this. For larger beam diameters the diffraction pattern is an average over multiple grains. This can be used to determine the crystal structure, as with powder XRD, but not the local orientation [44]. For standard laboratory Laué machines, smaller beam sizes are impossible due to the insufficient radiation flux through the focusing pinhole required [45]. XRD is not used for imaging microstructure due to its poor resolution

and relatively slow data collection and analysis, instead it is typically combined with optical techniques to determine the crystallography of a sample at a single point or over a few points. Crystallographic orientation can be determined within an accuracy of 1-2° [44]. Beyond removing paint-layers and surface contaminants, sample preparation is non-critical and rough samples can be readily measured.

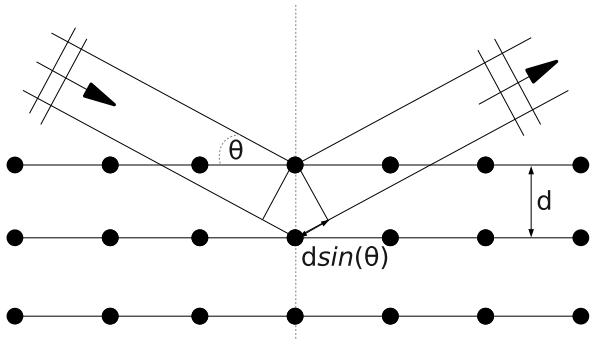


Figure 1.14: A graphic illustration of Bragg's law showing a coherent beam being diffracted.

1.5.3 Electron backscatter diffraction (EBSD)

Electron backscatter diffraction (EBSD) is a technique which exploits the particle-wave duality of electrons by diffracting them in the electronic structure of the crystal lattice. A scanning electron microscope (SEM) is used to generate and focus a narrow electron beam, with energies of 10's of keV, at a sample mounted at ~70° to the beam direction [46]. A schematic of a sample and EBSD camera mounted in a SEM is shown in Fig. 1.15. Since the electrons have a characteristic de Broglie wavelength, they are diffracted by the lattice under the conditions of Bragg's law. The diffracted beam is then directed onto a fluorescent, typically phosphor, screen which is then imaged by a CCD. The diffraction pattern obtained is called a Kickuchi pattern, after one of the pioneers in the field. Crystallographic data is then extracted automatically from the pattern using Hough transforms to identify bands, and indexing algorithms to determine orientation [47]. By scanning point-wise across a sample a map of orientation can be built which shows the fine grain structure. This technique is capable of accuracy to 0.5° angular resolution and, at its best, is able to resolve 10 nm grains at data capture rates of up to 1600 Hz [48]. Since electrons are readily absorbed by air and contaminants this technique must be undertaken in a vacuum. The surface finish of samples must be pristine, with all oxidation layers removed; typically this requires electro-polishing.

EBSD is limited to small samples, typically a few square centimetres, due to the vacuum chamber size which makes this, by necessity, a destructive technique. With the addition of an ion beam to the SEM, thin layers of the sample can be etched away and further EBSD images taken. By etching and scanning multiple times full 3D crystallography of samples can be obtained [49]. EBSD is generally considered the gold-standard of microstructure measurement and is used for comparison to, and validation of, SRAS data particularly for orientation determination in §2.4.

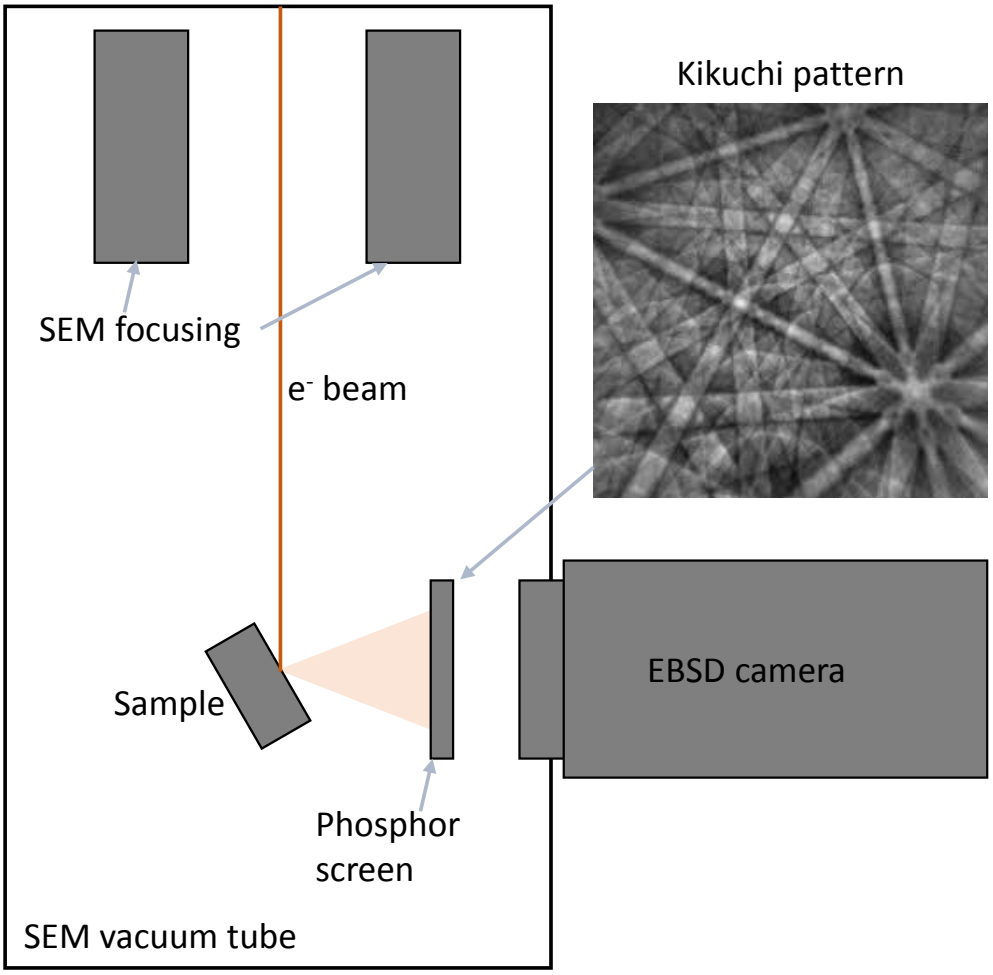


Figure 1.15: A schematic of a typical EBSD setup. A specialised EBSD camera captures images of electrons scattered from a sample, tilted at 70°, incident on a phosphor screen. EBSD is performed in a vacuum chamber with an SEM as the electron beam source.

1.5.4 Scanning acoustic microscopy

While an optical microscope is sensitive to changes in refractive index, using light as a probe, a scanning acoustic microscope (SAM) is a device which is sensitive to changes in material elastic properties, using acoustic waves as a probe [50]. A piezoelectric transducer generates acoustic waves which are focussed by an acoustic lens, through a liquid couplant into a sample. Fig. 1.16 shows a schematic of a SAM configured for a $V(z)$ type method. The $V(z)$ method is a technique whereby the focus depth is modulated in order to produce interference between bulk and surface waves [51]. The transducer generates a tone-burst which is focused at a depth Z into the sample. At the sample surface both bulk waves and 'leaky' SAWs are generated with both of these modes returning energy which is detected by the transducer. Through analysis of the interference between the bulk wave and the leaky SAW it is possible to determine the velocity of the SAW, and through this mechanism material microstructure contrast is achieved. Through use of a line focus acoustic microscope (LFAM) this contrast can be further enhanced [52].

SAM can achieve resolution down to 1 µm operating at acoustic frequencies of GHz, with data capture rates 500 Hz although requiring averaging of ~20 traces [53]. Although the technique is sensitive to sub-surface defects, the coupling of the bulk and SAWs limits the imaging of microstructure to the sample surface. Limited sample preparation is required beyond the removal of paint and coatings and no damage is inflicted upon the sample during scans.

1.5.5 Neutron diffraction

Unlike Laué and EBSD, neutron diffraction relies on the interaction of neutrons with atomic nuclei, rather than electrons or X-rays with the electron cloud. This makes it a useful technique due to the extra penetration it allows as well as the ability to detect low mass atoms such as carbon and oxygen. It works on principles very similar to Laué and is capable of producing results with excellent resolution. However this technique requires a high energy neutron source, either a synchrotron based spallation source or a lab based nuclear reactor, which are uncommon and expensive [54]. For synchrotron sources the high energy neutrons have good sample penetration such that

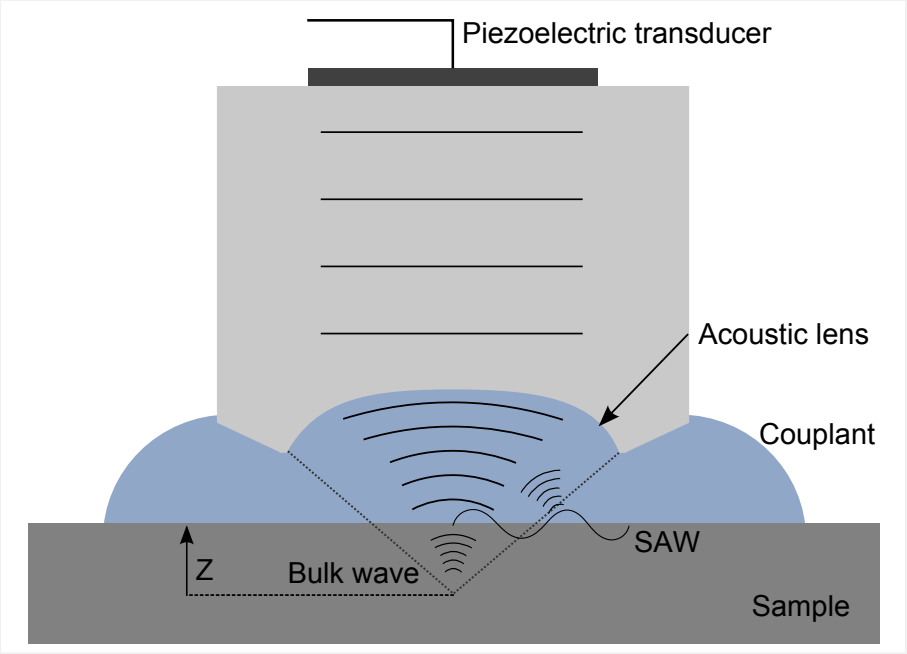


Figure 1.16: A schematic of a scanning acoustic microscope in the $V(z)$ configuration. An acoustic lens focusses ultrasonic waves through a couplant into a sample.

sample preparation is not required although lower energy sources do require the sample to be thinly sectioned. Typical spallation sources emit neutron bursts at a rate of 50 Hz and can achieve resolutions of nm [55].

1.5.6 X-Ray computed tomography

X-ray computed tomography (CT) uses a small focus X-ray source and a rotating sample stage to capture radiographs of a sample from 720 angles around a single axis, these images are then reconstructed into a 3D model of the sample [56]. Using X-ray micro-CT over repeated scans of multiple X-ray spectra it is possible to reconstruct the microstructure distribution of materials [57]. So far this technique has only been proven in this research area for imaging material phase variations, however work is progressing in identifying crystallographic phases of a uniform material. X-ray CT is a slow process requiring several seconds of exposure for each image, and more than 720 images to form a 3D model. No sample preparation is required and no damage is inflicted upon the sample.

Chapter 2

Spatially resolved acoustic spectroscopy

This chapter gives a broad overview of the state of the SRAS technology at the commencement of the EngD research period. The technique was developed by Prof Matt Clarke and Dr Steve Sharples in 2005 at the University of Nottingham. A brief explanation of the underlying science of the technique is given along with some of the limitations. The lab based EMDA SRAS system at the University of Nottingham is described with the optics, instrumentation and integration detailed. Some example SRAS velocity contrast maps are presented and a description of the process used to determine absolute crystallographic orientation given. For a more exhaustive review of the history of the SRAS technique up to the development of the EMDA SRAS system see [14].

2.1 SRAS background

SRAS is a non-contact laser ultrasonic method for the imaging and measurement of microstructure that can be deployed on most engineering metals. Two lasers are used; one for the generation and the other for the detection of SAWs. Generation is achieved by imaging a set of lines, or grating, with a given spacing, λ_a, onto a sample, this is shown in Fig. 2.1, an aperture is used to limit the diameter of the patch. In the laser illuminated regions the sample surface is rapidly heated causing it to expand. At

the boundaries with the unilluminated regions, local stress-strain fields are induced which in turn generate SAWs. Due to the grating the generated SAW packet has a dominant acoustic wavelength, λ_a, and a number of cycles, or oscillations, which is approximately equal to the number of lines in the grating. The generated SAW packet propagates primarily orthogonally to the grating lines. Fig. 2.2 shows a pair of typical SAW packets, and their frequency spectra, at two acoustic wavelengths. The Gaussian profile over the sinusoid is a result of two effects: the first is that more lines contribute to the centre of the wave packet than the extremities, the second is that the laser beam illuminating the grating has a Gaussian spatial profile.

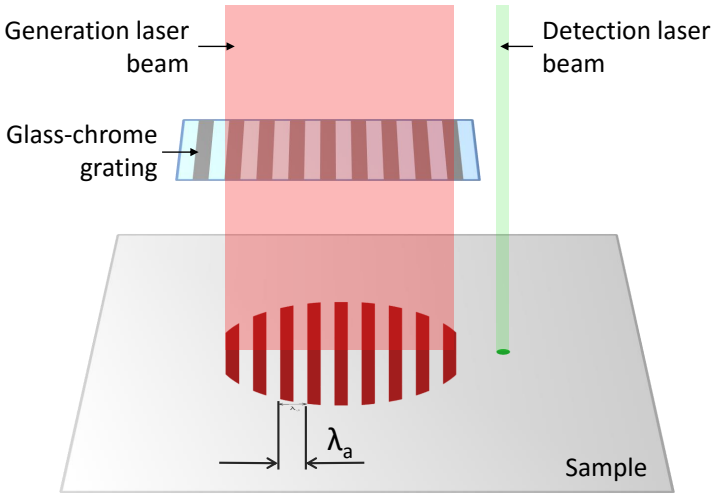


Figure 2.1: The generation patch (red) and detection spot (green) imaged onto a sample surface for generation and detection of SAWs. Typical grating spacings, and thus acoustic wavelengths, are ~20 µm.

In order to efficiently generate SAWs a pulsed laser source is essential: the temporal width of this pulse must be sufficiently short to contain frequency content within the spectrum required. Fig. 2.3 shows a selection of nano-second laser pulses and the frequency content available for generating SAWs. To efficiently generate SAWs of a particularly frequency, as determined by the grating, there must be overlap between the frequency content of the laser pulse and the chosen frequency range of the SAW.

The second, continuous wave, laser is tightly focused into a detection spot approximately 1-2 acoustic wavelengths from the edge of the patch, as shown in Fig. 2.1. As the SAW passes under the detection spot, the return detection beam is deflected by an angle dependent on the local SAW gradient. Detection is based on an optical beam

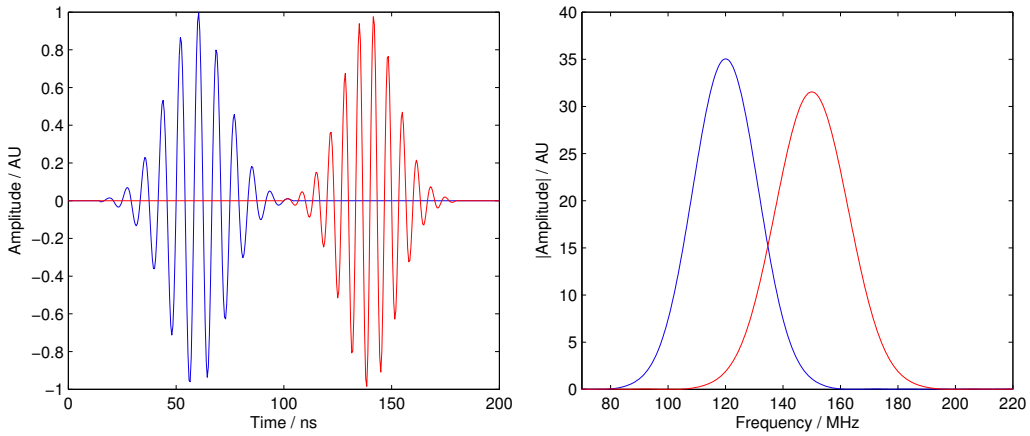


Figure 2.2: (a) A pair of ideal SRAS signals from a 9-line patch and (b) the Fourier transforms of these signals. Spacings are (blue) 25 µm and (red) 21 µm as generated on a typical aluminium sample. The traces have been shifted in time for clarity.

deflection technique, and so this beam is then collected and imaged onto a split photodiode. The detection spot diameter must be less than $\lambda_a/2$ in order to avoid spatial averaging of the SAW gradient. The output from the KED is then amplified, filtered and digitised, a Fourier transform is taken giving the frequency spectrum of the SAW, these are shown in Fig. 2.2 for an idealised SAW. Applying the wave equation, $v = f\lambda_a$, then yields the phase velocity spectrum of the SAW in the local region of the generation patch and detection spot. This is the basis of the *frequency* or *f*-SRAS technique. Previous generations of the technique implemented a Q-switched and mode-locked single frequency laser as the excitation source. A spatial light modulator (SLM) images a series of gratings each with a different acoustic wavelength, and a SAW is generated and detected for each change. This yields a peak in the wavelength spectrum which, with the known frequency of the laser and wave equation, gives a peak phase velocity. This technique is referred to as *k*-SRAS. The difference between the two techniques is in how a narrow-band SAW is created. *f*-SRAS applies a narrow-band grating filter to a broadband laser source to produce a narrow-band SAW packet. *k*-SRAS applies the same narrow-band grating filter to an already narrow-band laser source to produce a narrow-band SAW packet, with an amplitude dependent on how well the two filters overlap. In both techniques the broadest filter is the optical grating and so the generated SAWs are very similar [58]. Due to the requirement to generate multiple traces to determine phase velocity, *k*-SRAS is considerably slower than *f*-SRAS.

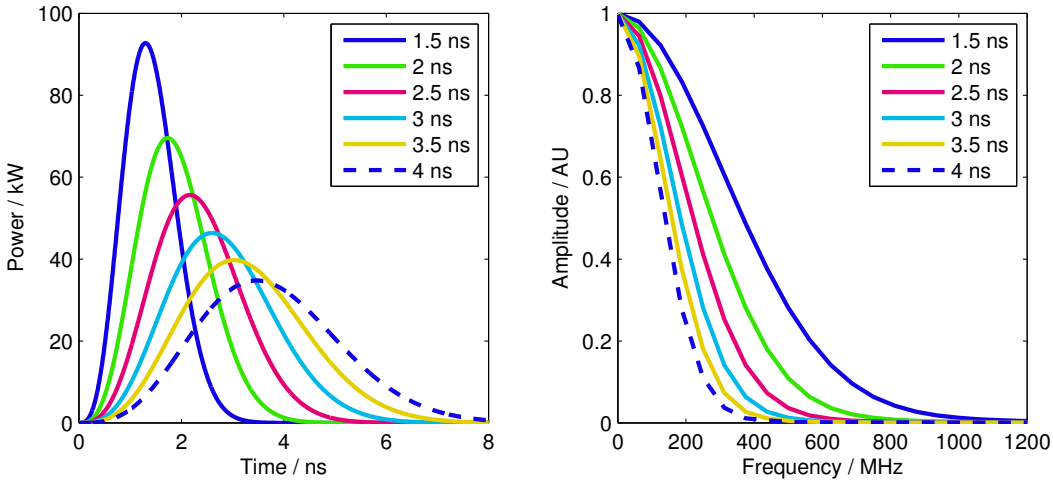


Figure 2.3: A set of laser pulses (left) each with total energy 120 µJ at various temporal widths. The Fourier transforms (right) of the laser pulses showing the frequency content available for generation of SAWs at various frequencies.

In order to generate a map of phase velocity over a sample, the generation patch and detection spot are raster scanned over a regular grid. Currently the sample is moved at a constant velocity underneath the SRAS instrument, collecting a trace every time the generation laser fires. Typical SAW velocities range between $1500\,\mathrm{m\,s^{-1}}$ to $3500\,\mathrm{m\,s^{-1}}$ whilst the stage scanning speeds are $20\,\mathrm{mm\,s^{-1}}$ to $500\,\mathrm{mm\,s^{-1}}$. As such the scanning speed does not contribute significantly to errors in the velocity measurement. A phase velocity image is then built by finding the peak in each velocity spectrum as described above. Fig. 2.4(a) shows a typical velocity contrast map captured using the EMDA SRAS instrument, the spacing between velocity measurements is 5 µm in the x-axis and 10 µm in the y-axis. Note that for all SRAS scan velocity maps shown in this thesis the units of the colourbars are always $\mathrm{m\,s^{-1}}$, whilst the amplitude maps and optical power maps are indications of relative amplitude and are not calibrated. Fig. 2.4(b) show the peak amplitude of the velocity spectrum, which gives a very approximate indication of the SAW amplitude. Fig. 2.4(c) shows the total optical power measured at the detector. These values can both be used to determine the quality of the velocity determination and to generate data selection masks where scanning passes over sample edges.

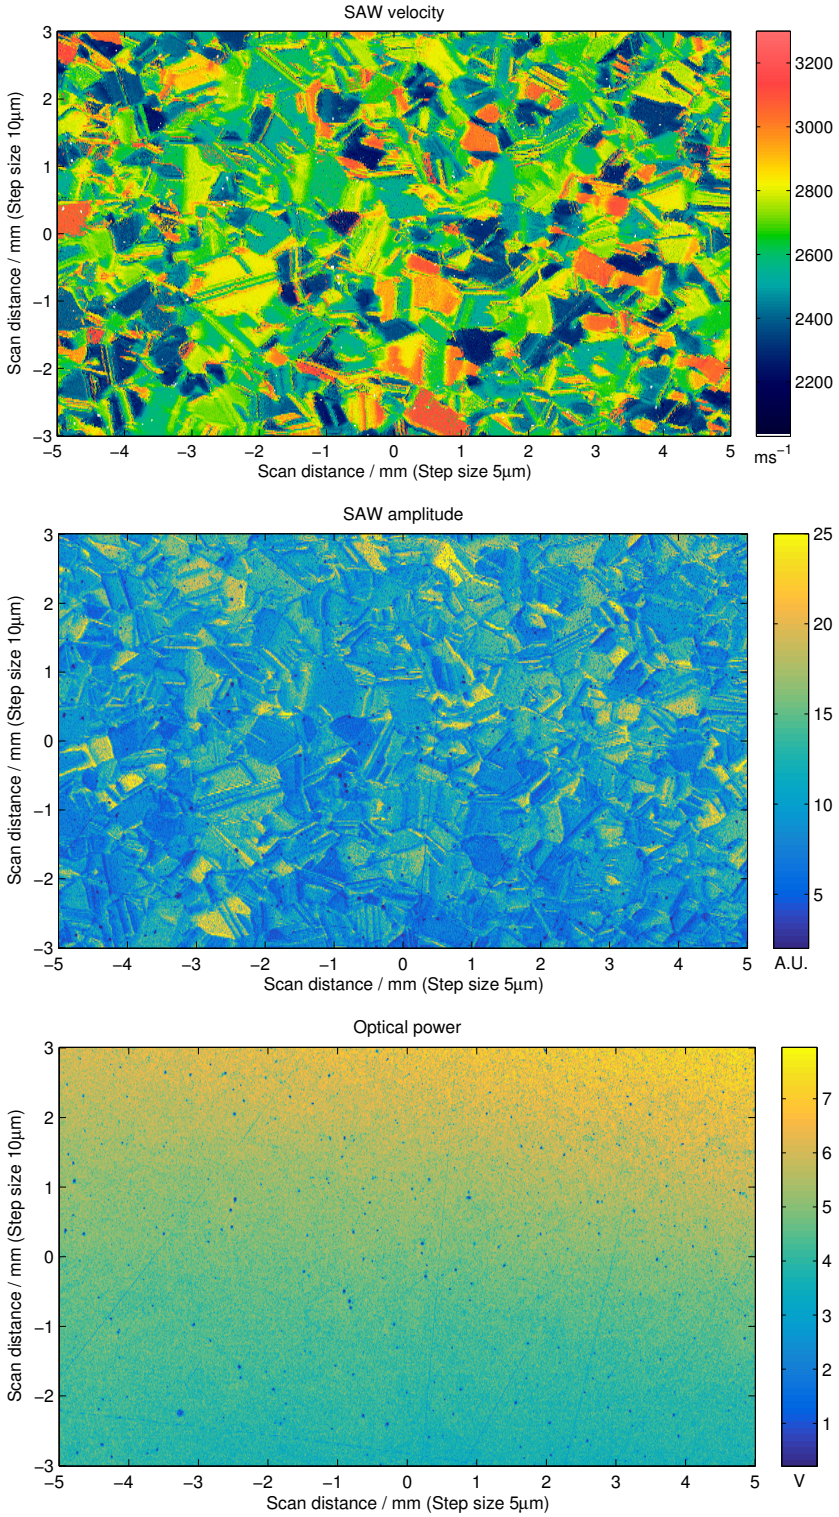


Figure 2.4: A typical SRAS scan showing a 10 mm by 6 mm area of an Inconel 617 sample. An acoustic wavelength of 22 µm was used with an aperture of 1 mm. In addition to the velocity contrast map, the signal amplitude and optical power at the detector are also captured. This Inconel 617 sample was provided by E.ON and the University of Strathclyde.

2.2 EMDA instrument

A few SRAS instruments have been built and implemented to date and these are discussed in reference [14]. As of the start of the EngD project, the latest instrument is the EMDA SRAS system. Based on the *f*-SRAS technique the EMDA system, shown in Fig. 2.5, utilises a Q-switched, broadband 1064 nm laser for excitation and a continuous wave 532 nm single-longitudinal mode laser for detection. The EMDA system was designed and built by Dr. Richard Smith in 2011 and, to date, has seen little development or change apart from a switch to a fibre coupled detection laser from the originally bread-board mounted configuration.

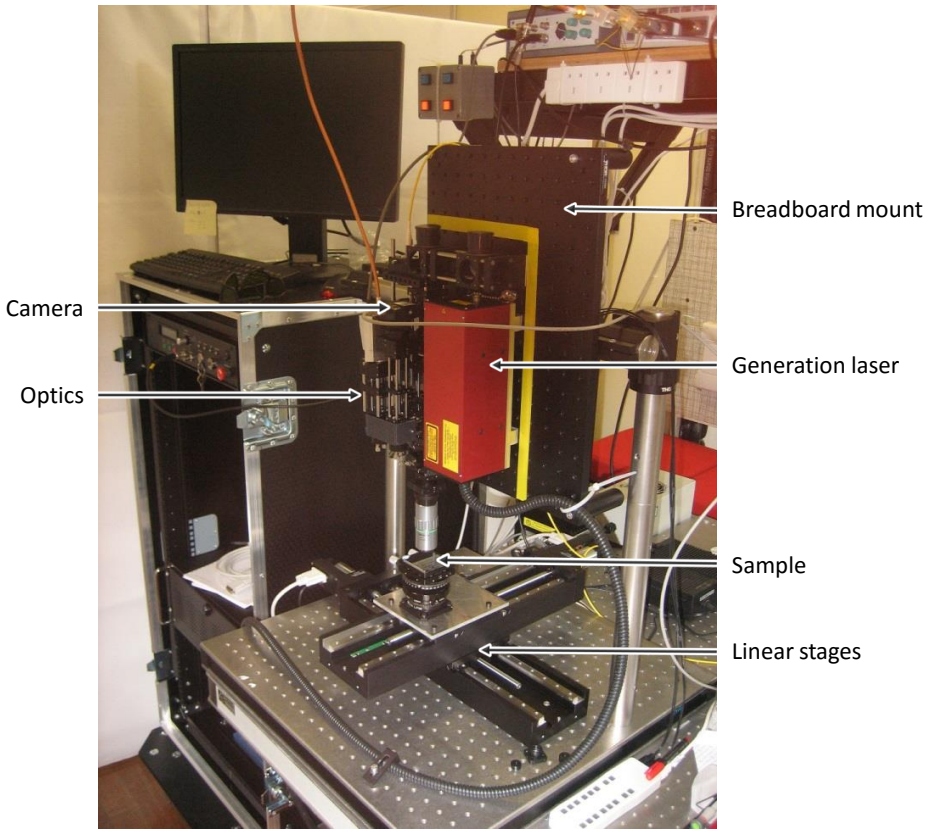


Figure 2.5: A labelled photograph of the EMDA system at the University of Nottingham.

2.2.1 Optical configuration

Fig. 2.6 shows the optical configuration of the EMDA SRAS system. All of the optical components, with the exception of the masks, are standard off-the-shelf parts mounted in a optomechanical cage system. There are essentially two beam paths: the generation path and the detection path.

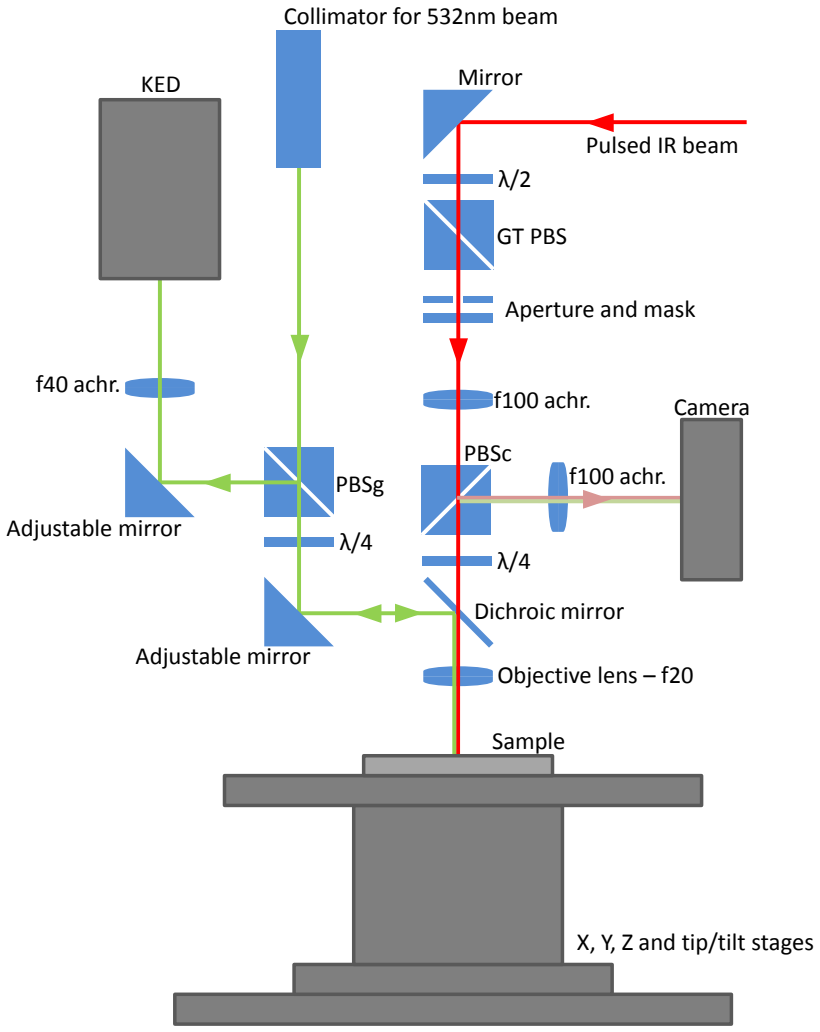


Figure 2.6: An optical schematic of the EMDA instrument. All optics apart from the masks are standard and mounted in off-the-shelf optomechanics. The IR laser is mounted directly to the breadboard along with the optics. The visible laser is fibre-coupled. The abbreviations used are: achr. - achromatic lens, PBSc - polarising beam splitter combine, PBSg - PBS green and GT PBS - Glan-Taylor PBS.

Generation path

The generation laser is an Nd:YAG Q-switched and polarised source with a centre wavelength of 1064 nm, which is mounted directly to the breadboard. Pulse lengths are ~1 ns to 3 ns with energies of 100 µJ to 200 µJ. Adjusting the laser pump current increases and decreases the pulse energy, but also affects the pulse temporal width. Therefore a rotatable quarter-wave plate (QWP) and a Glan-Taylor polarising beam splitter (GT PBS) are used to split and siphon off excess intensity from the beam. The pump current is kept at its maximum where pulse length is at a minimum. Power control is necessary to ensure generation is kept in the thermoelastic regime, thereby avoiding damage to samples. The beam is turned through two mirrors and directed towards the QWP.

After the power control is a mask and foil aperture assembly. The mask is a fused quartz substrate, onto which chrome lines are deposited, and the aperture is a foil pinhole centrally aligned to the beam. A small rotation stage allows rotation of the mask to propagate SAWs in any chosen direction. Masks and pinholes are chosen based on the inspection requirement, but typical values are 1 mm for the pinhole diameter and 100 µm line spacing for the mask.

An f100 achromatic tube lens is placed 100 mm along the path from the chrome-side of the mask, such that the mask is at the back focal plane. Combined with the final f20 objective lens this achieves a factor 5 demagnification of the aperture and mask, i.e. a 1 mm diameter aperture forms a 200 µm image on the sample with the 100 µm line spacing mask generating a 20 µm acoustic wavelength SAW. Next, a PBS (PBSc) is mounted at the same orientation as the GT-PBS, achieving maximum transmission. A quarter wave plate then alters the polarisation from linear to circular. Following this is a dichroic mirror coated for transmission of 1064 nm and reflection of 532 nm light. Finally the beam passes through the objective lens and is focused onto the sample, the objective is a long-working distance, infinity corrected 10X lens. This lens produces a very well-focused image and provides a working distance of 34 mm.

The reflected IR beam comes back through the objective and dichroic mirror, at the wave-plate the polarisation is modified back to linear, but is orthogonal to the input beam. This beam is then reflected by the PBSc into a f100 achromat which images

the beam onto the camera. By matching the camera lens and the tube lens, a 1:1 scale image of the mask is produced on the camera sensor.

Detection path

The source of the detection path is a fibre coupled, frequency doubled, single longitudinal mode, 532 nm continuous wave Nd:YAG laser with a maximum output power of 400 mW. Single-mode, polarising maintaining fibre is used, as a result the collimated beam has near diffraction limited divergence but is reduced in power by ~40%.

The collimated beam passes through a polarising beam splitter (PBSg), which is aligned for maximum transmission, before passing through a quarter wave plate which modifies the polarisation from linear to circular. A piezo-enabled right angle mirror then turns the beam toward the dichroic mirror where it is reflected towards the final objective lens. The collimated beam is then focused to a point on the sample surface, in the same image plane as the IR grating image. Lateral movement of the focused spot to the correct position relative to the grating is achieved by the piezo adjustable mirror. Each time the mask is rotated the detection spot must be reconfigured.

The reflected beam from the sample is then captured by the objective lens and recollimated. The dichroic mirror then turns the beam back towards the piezo adjustable mirror. Although the reflection of the dichroic mirror is high at 532 nm, a small proportion, ~1%, of the light is transmitted through the dichroic. This is then converted back to linear polarisation and reflected towards the camera by the PBSc. The camera imaging lens then forms an image of the detection spot on the camera sensor.

The majority of the beam is turned by the piezo mirror towards the quarter-wave plate where the polarisation is modified back to linear, although orthogonal to the original beam. This beam is then reflected by the PBSg towards another, manually adjustable mirror, where the beam is turned towards an f40 achromat which focuses the beam towards the detector. A segmented photodiode is the optically sensitive part of the detector. The detector is located such that the focusing beam occupies an equal area of each photodiode segment, and occupies as much total area as possible across both segments. The manually adjustable mirror allows the beam to be translated across the photodiode. Each time the detection spot location is altered, or the propagation

direction changed, the location of the beam on the photodiode must be re-optimised.

2.2.2 Detection and acquisition

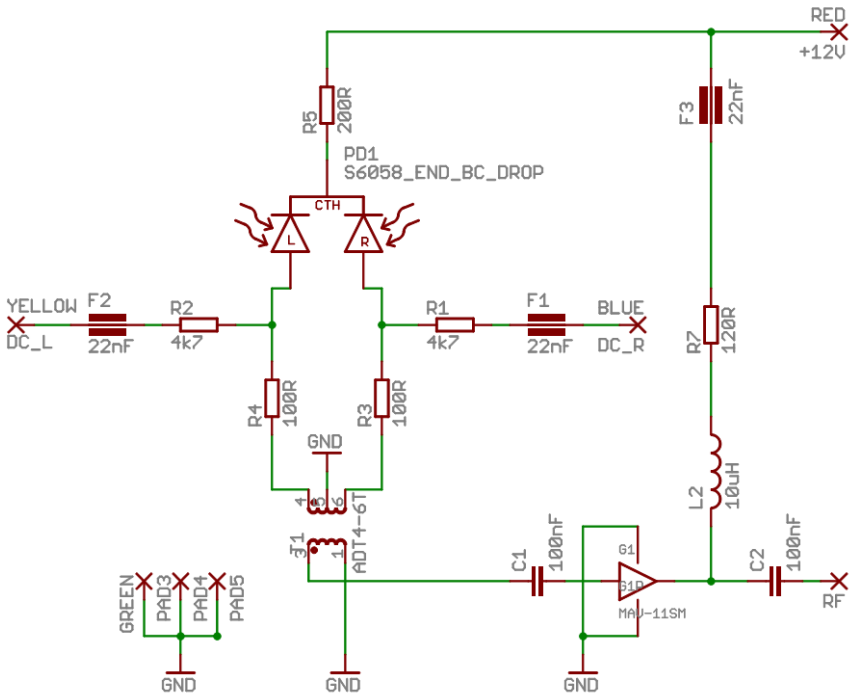


Figure 2.7: A circuit diagram of the KED used on the EMDA system.

Fig. 2.7 shows the circuit diagram for the KED originally designed by Dr. Steve Sharples. The photodiodes are reverse biased to 12 V which helps to improve the bandwidth of the detector. Current is limited to ~40 mA in order to protect the photodiodes from very bright laser sources. A 100 Ω load resistor is used to siphon off the DC from each segment of the photodiode, which is then low-pass filtered. The sum and the difference of these load voltages is used to optimise the beam location on the photo sensitive region of the detector. The sum is also digitised by an oscilloscope and is used to form the optical images shown in Fig. 2.4(c). As a SAW deflects the detection beam from side-to-side across the detector the intensity on each photodiode segment changes slightly. This causes a change in the direction and amplitude of current through the transformer. The coupled coil of the detector picks up this current at a voltage ratio of 1:2, due to the coil winding. This RF current is then amplified and output via a coaxial cable into a set of discrete filter and amplifiers shown in Fig. 2.8.

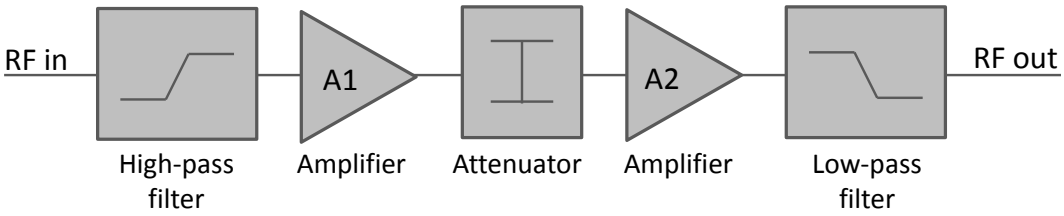


Figure 2.8: A circuit schematic of the signal chain between the KED and oscilloscope.

Dependent on the scan parameters the passive components in the signal chain typically have values: high pass filter 50 MHz or 100 MHz, RF attenuator 3 dB or 6 dB, and the low pass filter 150 MHz to 250 MHz. The amplifiers are: A1 - 10 MHz to 500 MHz at 19 dB and A2 - 0.1 MHz to 500 MHz at 24 dB. The result is a signal with a maximum amplitude of 0.05 V to 1.5 V which is detected on a digital oscilloscope.

A typical SAW signal length is in the range 50 ns to 200 ns and this signal is acquired by the oscilloscope at a bit depth of 8 and a sample rate of 1.5 GHz to 3 GHz. The scope is triggered by the sync pulse from the generation laser, which is output each time the Q-switch opens releasing a pulse. Although there is some jitter in the laser Q-switch, this is of order nano-seconds and is significantly shorter than the temporal length of the signal. Since only the frequency content of the signal is of interest this effect is not an issue.

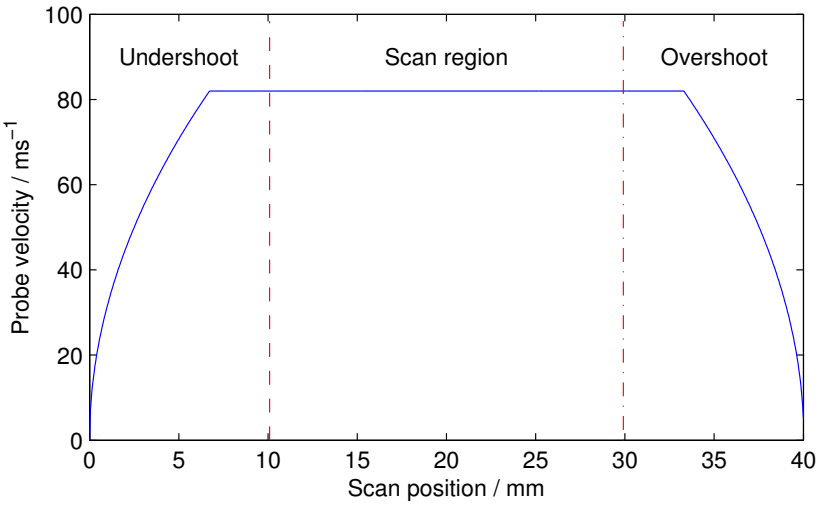


Figure 2.9: A figure illustrating the overshoot and undershoot regions used to ensure constant scanning velocity during data collection.

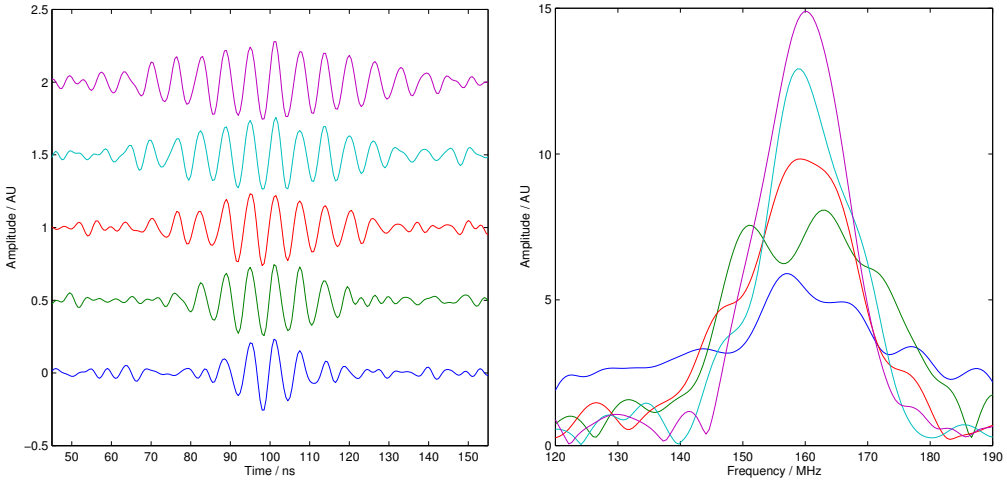
Scanning is performed ‘on-the-fly’ with linear stages initialising and terminating ac-

quisition based on passing certain points in a line-scan. An undershoot and overshoot are used to ensure a constant velocity during acquisition, Fig. 2.9 shows these regions for each line-scan and the scan region where data is collected. Traces are buffered onto the fast memory of the oscilloscope during each line-scan, this limits the length of each line-scan to the maximum capacity of the memory: 8,000 traces. The buffered traces are then downloaded via ethernet to a computer and stored on the hard drive. The next line-scan cannot start until the memory has been fully emptied. Typically less than 4000 are traces captured per line-scan, above this value data transfer rates decrease non-linearly.

2.3 EMDA SRAS example data

The EMDA SRAS system can be configured with acoustic wavelengths between 10 µm to 50 µm and has a minimum step between point measurements of order 1 µm. Based on experimental evidence the maximum achievable spatial resolution achieved with the instrument is ~30 µm using a patch size of 60 µm. The trade off for high spatial resolution is poorer velocity resolution due to the reduction of cycles in the SAW signal and the resultant frequency peak broadening [59], Fig. 2.10 shows this effect. Currently, the absolute limit on acoustic wavelength is the detector which has a bandwidth of 300 MHz. A more typical EMDA SRAS configuration is an acoustic wavelength of 24 µm and patch diameter of 200 µm which generates a SAW of approximately 8 cycles. This is the default setting used for most scans.

The noise floor of the instrument is limited by a few factors including temperature, sample flatness and motion platform smoothness. Of most significance is the error introduced in signal processing, shown in Fig. 2.11, where despite having a strong signal, the noise distribution is still sufficient to distort the frequency peak such that the error in the measurement of peak frequency is 5.8 MHz. For the standard acoustic wavelength of 24 µm this corresponds to an error in velocity of $145\,\mathrm{m\,s^{-1}}$. More advanced signal processing techniques have been used on SRAS data previously, however with average data collection rates now ~1200 points per second over the length of a scan, these proved to be infeasibly slow [60]. That said, Chapter 4 examines the application of a range of fitting techniques to SRAS data to reduce errors in velocity



(a) (b)

Figure 2.10: A set of simulated SRAS signals over a range of cycles. (a) shows the raw traces and (b) shows the frequency spectrum of each trace.

determination.

Fig. 2.12 shows a SRAS scan of an isotropic aluminium sample, over which the SAW velocity should be constant. The first image shows the single-shot results of the scan where every trace has been analysed, for the second image 20 averages have been taken before analysis. For the first image the average velocity is $2934\,\mathrm{m\,s^{-1}}$ with a standard deviation of $13\,\mathrm{m\,s^{-1}}$; this equates to a measurement error of 0.44%. For the second image the average velocity is $2932\,\mathrm{m\,s^{-1}}$ with a standard deviation of $7\,\mathrm{m\,s^{-1}}$; this equates to a measurement error of 0.24%. Evidently averaging can be applied to reduce the velocity measurement error, however this is at the cost of spatial resolution.

Fig. 2.13 shows an example SRAS scan taken with a 24 µm acoustic wavelength and 200 µm aperture on a large-grained sample of Inconel 617. Immediately clear is that the typical grain size of the material is ~500 µm, which is much larger than the SRAS resolution of ~50 µm. Operating in this regime, grains are well resolved and the microstructure can be trivially interpreted. The right edge shows the effect of non-flat surfaces on the SRAS signal. As the generation patch is defocused two things occur: a change in acoustic wavelength due to the change in line spacing, and a decrease in signal amplitude due to decreasing power density. In addition the balance of light on the segmented photodiode is disturbed, decreasing the sensitivity of the detector. A secondary effect is seen at the bottom edge of the sample; here only part of the patch

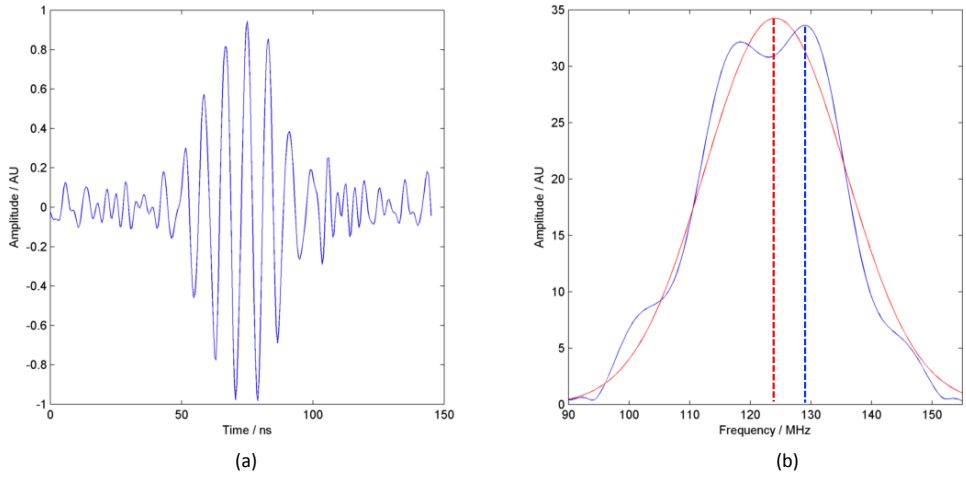


Figure 2.11: a) A simulated SRAS signal of SNR 8 dB (left). (b) The frequency spectrum of the signal (blue) and the frequency spectrum of the signal without noise (red).

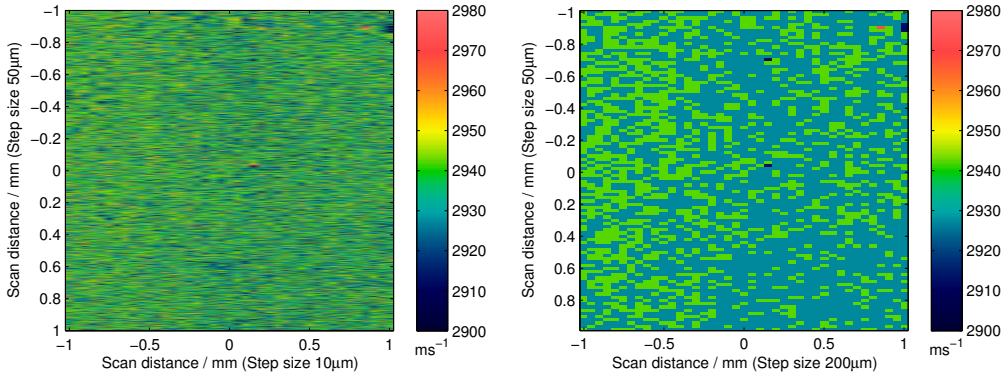


Figure 2.12: A SRAS scan of an isotropic aluminium sample using the default mask configuration. (a) The raw fitted data, (b) 20 averages taken before fitting the data.

is on the sample, generating signals with fewer, or no cycles and, therefore, yielding a noisy measurement. In order to determine the quality of the signal the total light on the photodiode is recorded, thresholded, and used as a mask for the velocity map. For this reason the effect of a partial patch on the sample is not seen at the top edge since the detection beam is the first part of the patch to pass over the sample edge.

Fig. 2.14 (a) and (b) show velocity contrast maps on a sample of titanium 834, in (a) the SAW is propagating from left to right, with respect to the image, and (b) is propagating from top to bottom. Unlike the previous sample the typical grain size of this sample is less than the resolution of the measurement. Instead of a well defined

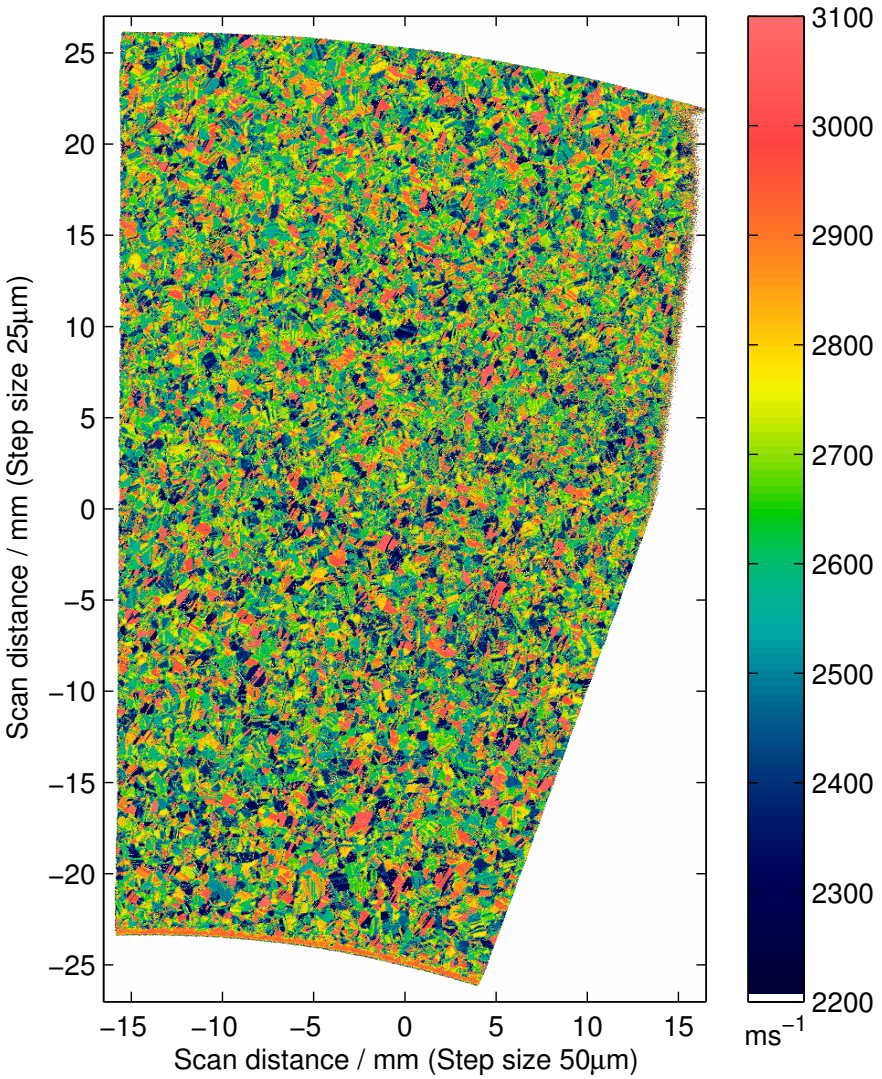


Figure 2.13: An example SRAS scan taken with a 24 µm acoustic wavelength and 200 µm aperture on a large-grained sample of Inconel 617. The SAW is propagating from bottom to top. A zoomed region of this sample was shown in Fig. 2.4.

Colour	Scan 1 SAW speed	Scan 2 SAW speed
White	Slow	Slow
Green	Fast	Slow
Red	Slow	Fast
Blue	Fast	Fast

Table 2.1: The 2D colour map colour specification based on the SAW speed from two orthogonal scans.

grain structure instead a more statistical measurement of the average SAW velocity is returned. In order to enhance this microstructure, two orthogonal SRAS scans can be combined This is shown in Fig. 2.14c and is called a velocity vector map. The 2D colour map used is defined in Table 2.1.

The velocity vector map has a secondary function which is to better show unique grains crystallographies. A single propagation direction scan gives only one SAW velocity, which is insufficient to highlight grains uniquely, therefore two grains may have exactly the same SAW velocity but have entirely different crystallographic orientations. Thus by combining two measurements the likelihood of two planes having the same SAW velocities in two orthogonal directions is greatly diminished, and unique planes are more easily identified.

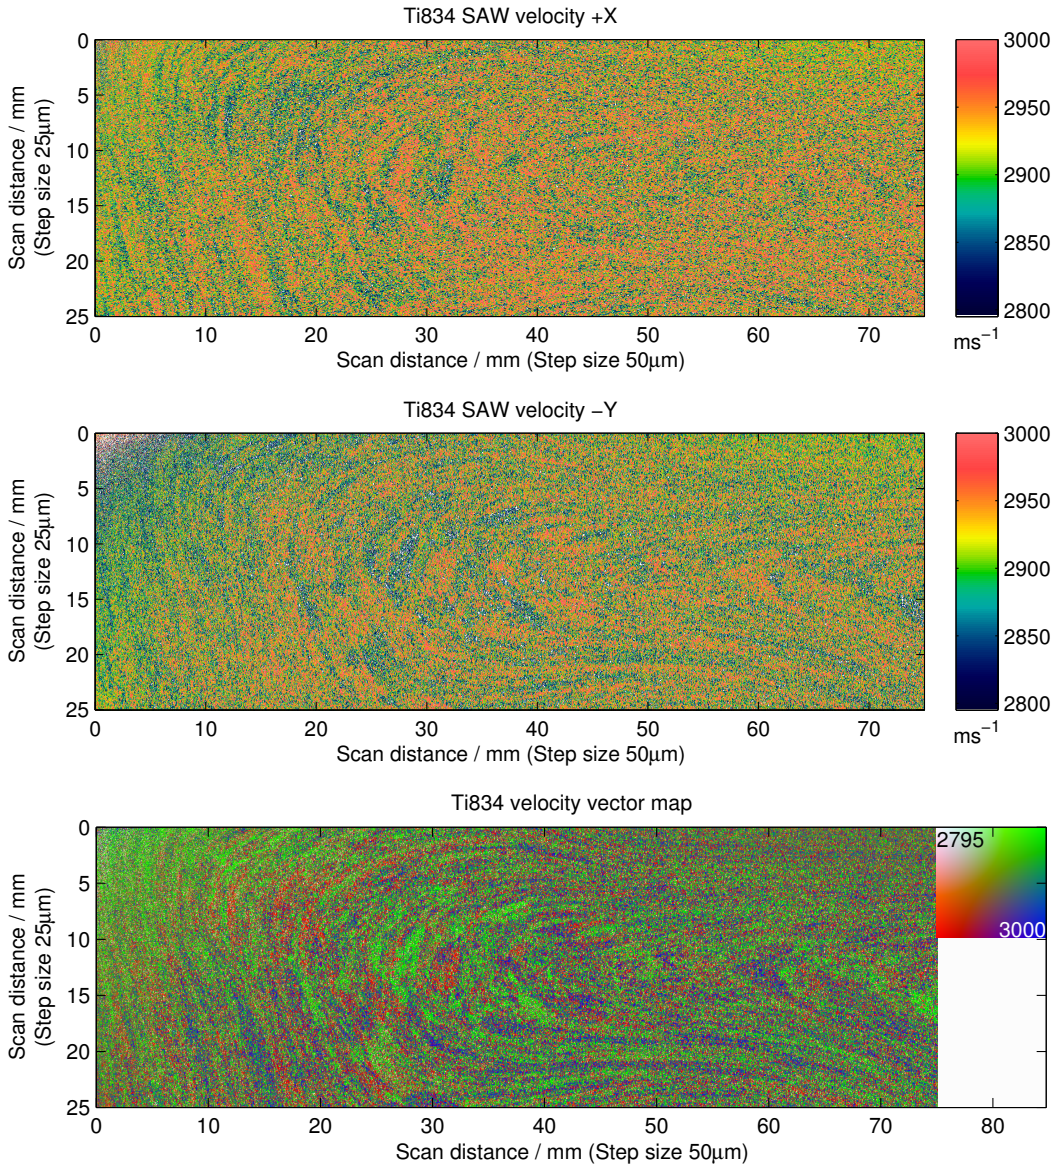


Figure 2.14: Two orthogonal SRAS scans taken on a sample of Ti-834 and a combined vector map which improves visibility of features in the scans.

2.4 Crystallographic determination by SRAS

One of the more powerful applications of SRAS is its application to the quantitative determination of material crystallographic orientation. This work was done primarily by Dr. Wenqi Li and is described in detail here [14]. The orientation determination work remains a very active area of research and the information presented here represents very recent work.

As the direction of propagation of a SAW is rotated on an anisotropic sample the measured SAW velocity will change, with the change dependent on the local crystallography. This can be seen in Fig. 2.15 which shows a set of scans of a subregion of the Inconel 617 samples shown in Fig. 2.13. The propagation direction has been rotated by 30° between each scan. Of primary interest in this figure is that the grain structure remains consistent while the SAW velocity changes. A particularly clear example of this is highlighted by the pink circles in Fig. 2.15.

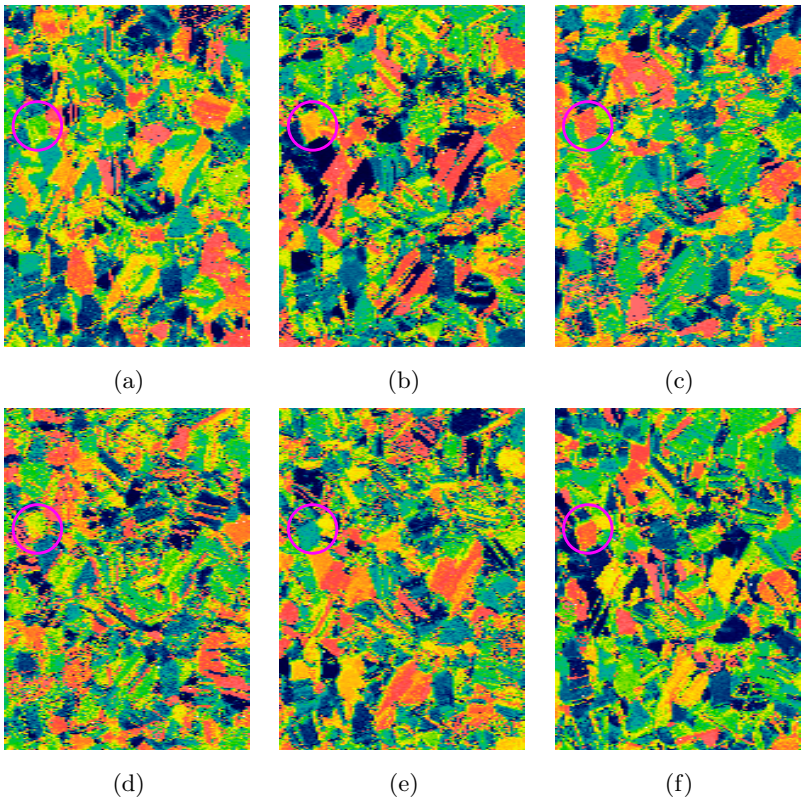
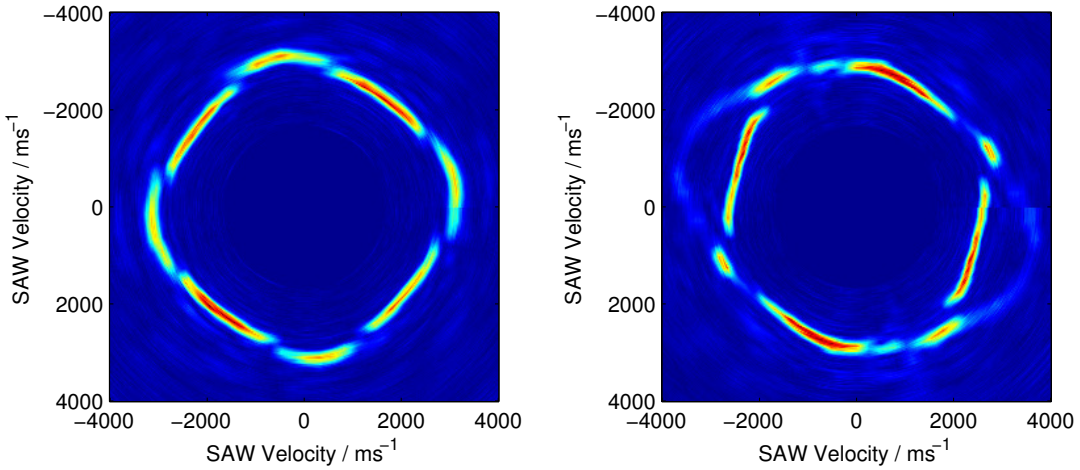


Figure 2.15: SRAS scans of the same region of an Inconel 617 sample with SAW propagation direction stepping from (a) 0° to f) 150° with steps of 30°. The full dataset contains 18 propagation directions stepping every 10° from 0° to 180°.

A velocity surface for a particular pixel is generated by choosing a single pixel and, at this point, collecting a set of scans including propagation directions from 0° to 180° with steps of 10°. Instead of plotting the dominant SAW velocity the full velocity spectrum is plotted as a polar plot with the SAW velocity along the radial axis. A set of velocity surfaces are shown in Fig. 2.16.



(a) (b)

Figure 2.16: Two experimentally derived nickel velocity surfaces at different orientations and rotations. Each surface shows differences in signal strength and velocity with propagation direction. Some surfaces exhibit high symmetry, such as 4-fold in (a) while (b) exhibits the minimum 2-fold symmetry. In certain propagations two SAW modes are present, as seen in (b). These scans were taken by Steve Sharples and Wenqi Li on nickel samples provided by Rolls-Royce.

There exists a relationship between the SAW velocity, crystallographic orientation and elastic constants, shown in Fig. 2.17. The 'forward model', that is determining SAW velocity for a given set of elastic constants and crystallography, is tractable and a complete analytical solution is available [61, 14]. Solving the 'inverse model', that is determining crystallography from a set of SAW velocities and elastic constants, is analytically intractable apart from a few special high symmetry cases. A thorough description of the analytical forward model is given in reference [14], but the process is described briefly here.

- Choose a crystal type and crystallographic plane.
- Combine the anisotropic elastodynamic wave equation (equation 1.16) and the

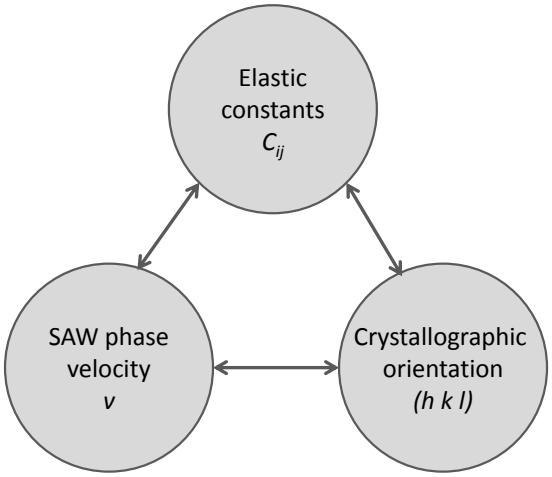


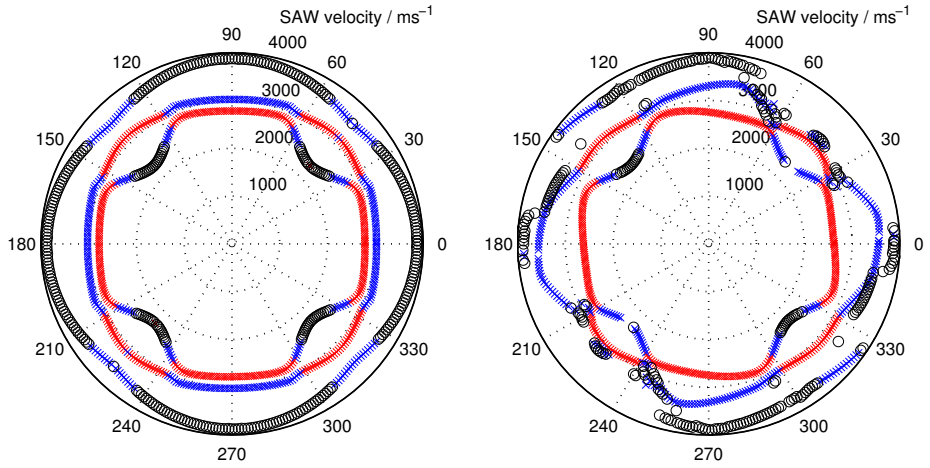
Figure 2.17: A diagram illustrating the interrelation between the crystallography, SAW velocity and material elastic constants. With any two of these the third is theoretically obtainable.

stress free boundary condition to give the determinant of the wave equation.

- Iterate through possible SAW velocities to find minima in the determinant of the wave equation. These minima are taken to be valid wave modes.
- Calculate the particle motion for each valid mode, and use the out-of-plane displacement to determine the 'dominant' mode.
- Repeat procedure for all possible propagation directions and crystal planes.

The result of this process is a database for a particular material of the dominant SAW modes for the velocity surface on each crystallographic plane, two such planes are shown in Fig. 2.18.

Since solving the inverse problem is not a possible route for determining crystallography from SAW velocity surfaces, a procedure based on the forward model is used. Orientation is determined by finding the optimum matching between the experimental and modelled velocity surfaces. Each experimental velocity surface is compared to every possible modelled plane and rotated through every possible propagation direction. A figure of merit is generated for each attempt by correlating the experimental and modelled surfaces. The final orientation is then taken to be the plane of maximum figure of merit and the plane rotation is taken to be the rotation of the propagation direction yielding the maximum figure of merit. Fig. 2.19 shows a selection of experimental velocity surfaces with the optimum model velocity surface overlaid.



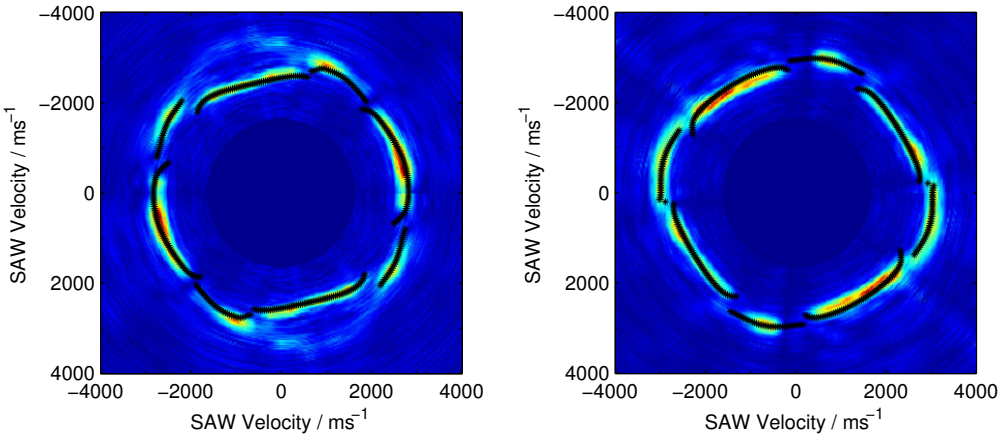
(a) (b)

Figure 2.18: A pair of example modelled SAW velocity surfaces for two planes on nickel. The dominance of each mode is coloured from red (for most dominant) to black (for least dominant).

Fitting orientation is a fairly slow process for two reasons: the large volume of data collected and the data processing rate. A full set of 18 SRAS scans is currently used for determining microstructure, it is possible to use fewer directions although this reduces the accuracy of the orientation determined [62]. Due to the requirement to search the model databases exhaustively, fitting data takes ~0.1 s per point. Fig. 2.20 and Fig. 2.21 show the output of the fitting algorithm for the sample of Inconel 617 shown in Fig. 2.13. In order to minimise the fitting time, four averages were taken along the y-axis before fitting. Fig. 2.20 is an inverse pole figure with the orientation plotted by Miller index, Fig. 2.21 shows the rotation of the local crystal on the plane.

EBSD data for this Inconel sample were also taken over a small region, and these are shown with the SRAS data in Fig. 2.22. These images show clearly the differences in the spatial resolutions of the technique, while EBSD has a resolution of a few nm, SRAS is ~3 orders of magnitude poorer. However, within the large grains on the Inconel shown, SRAS compares well to EBSD.

In order to further qualify the fitting algorithm a comparison has been made on a nickel superalloy between orientation from SRAS and orientation from EBSD, this is shown in Fig. 2.23. For this plot the SRAS step-size between points is 25 µm by 7 µm whilst the EBSD step-size is 5 µm by 5 µm. The SRAS measurement resolution is of order 3-4 steps whilst the EBSD resolution is significantly less than the step-size. Despite



(a) (b)

Figure 2.19: A set of fitted SAW velocity surfaces for a pair of planes near (0 0 1) on a nickel sample. 18 propagation directions have been used to fit the data. Only the dominant wave modes have been overlaid. These scans were taken by Steve Sharples and Wenqi Li on nickel samples provided by Rolls-Royce.

this the SRAS data again compare favourably with EBSD data in terms of imaging the microstructure, and in the absolute measurement of orientation.

Currently, fitting of cubic materials is well understood and well qualified, and full orientation has been demonstrated and validated against EBSD data. Significant progress has been made in determining orientation on hexagonal close packed materials, such as titanium, although the scope for this is limited by the acoustic symmetry of the hexagonal lattice, as such only partial determination of orientation can be made. Fig. 2.24 shows a comparison between orientation determination of a hexagonal titanium, taken by both SRAS and EBSD. Since only partial determination of orientation can be made Fig. 2.24 shows the first Euler angle instead of the Miller indices. Euler angles are an alternative representation of orientation which describe a crystallographic plane by three rotation transforms away from the (0 0 1) plane [64]. For the large grained sample, SRAS orientation determination performs similarly to EBSD. The reduced sensitivity to surface condition of SRAS compared to EBSD has allowed more data to be captured. This an active area of research. More information about the forward model for determination of crystallographic orientation can be found in Chapter 7 where enhancements to the model, based on the generation mechanism, are discussed.

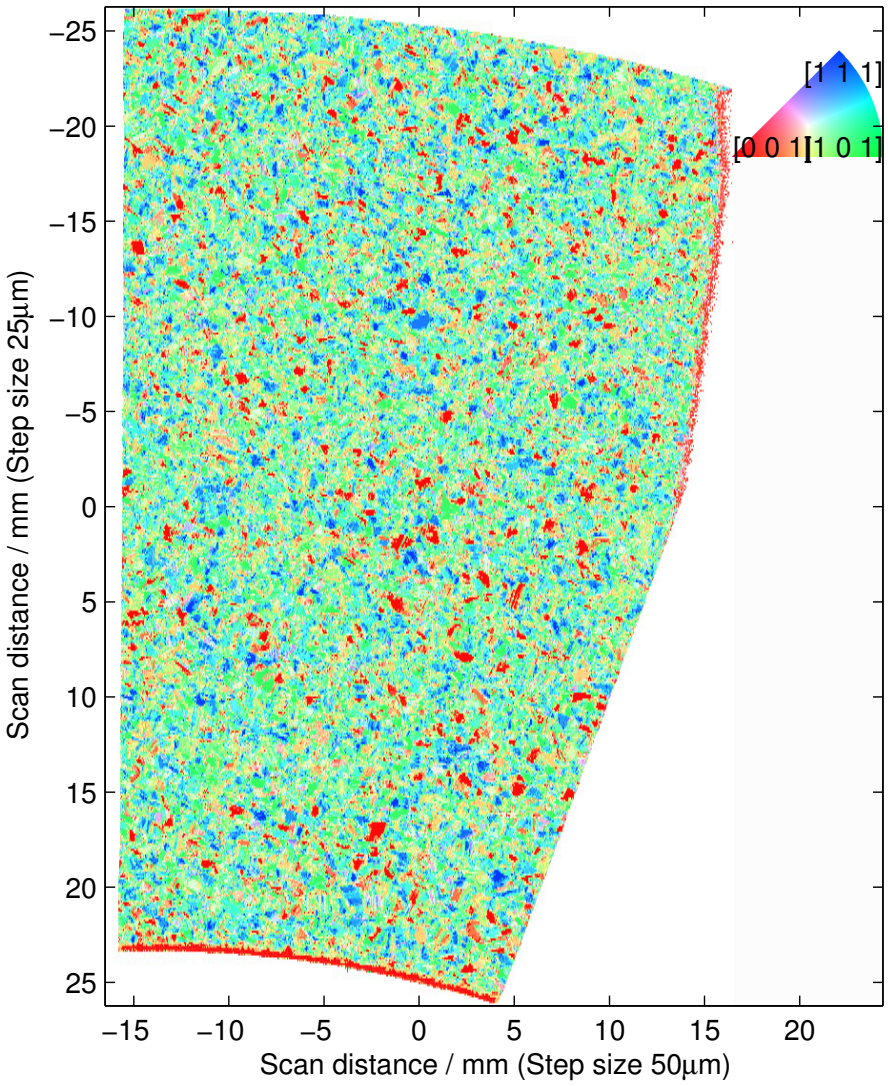


Figure 2.20: An inverse pole figure showing the orientation of a multi-grain cubic Inconel 617 sample. Data have been averaged 4 times before fitting giving pixel sizes of 100 µm in the y direction and 50 µm in the x direction.

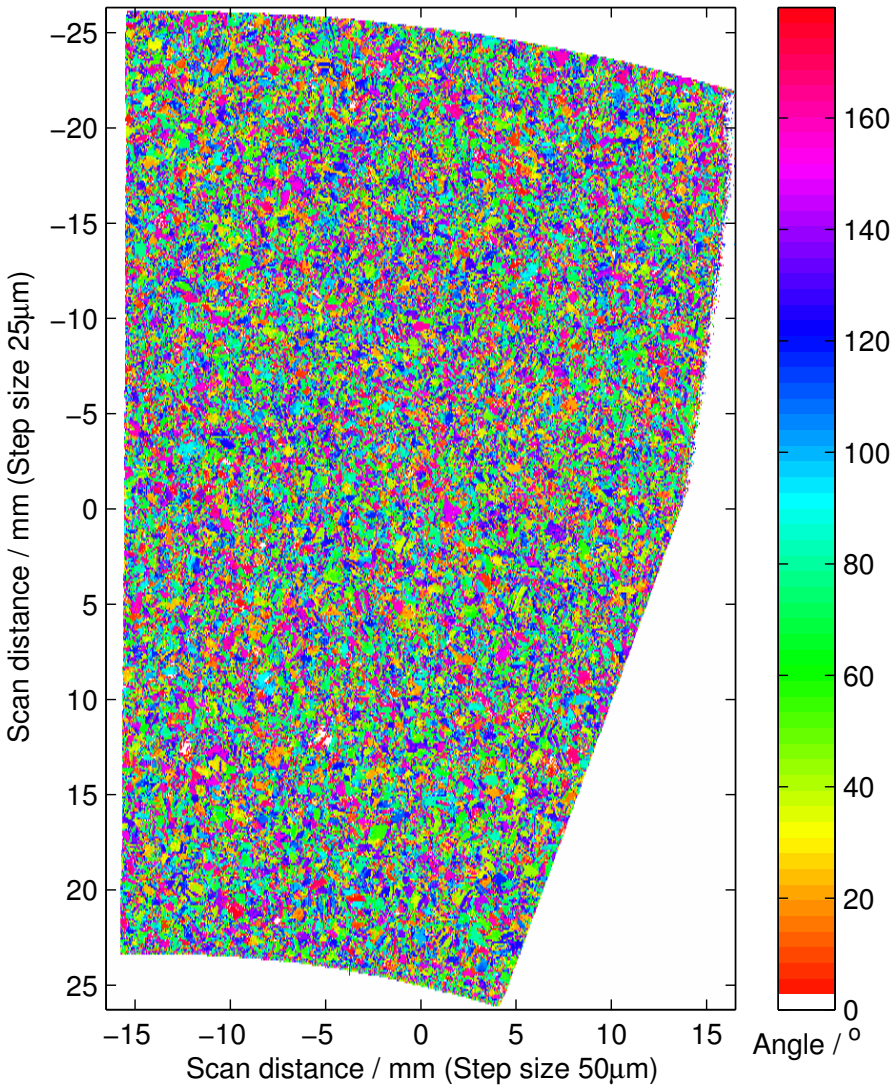


Figure 2.21: A figure showing the rotation on the plane for the inverse pole figure of a multi-grain cubic Inconel 617 sample. Data have been averaged 4 times before fitting giving pixel sizes of 100 µm in the y direction and 50 µm in the x direction.

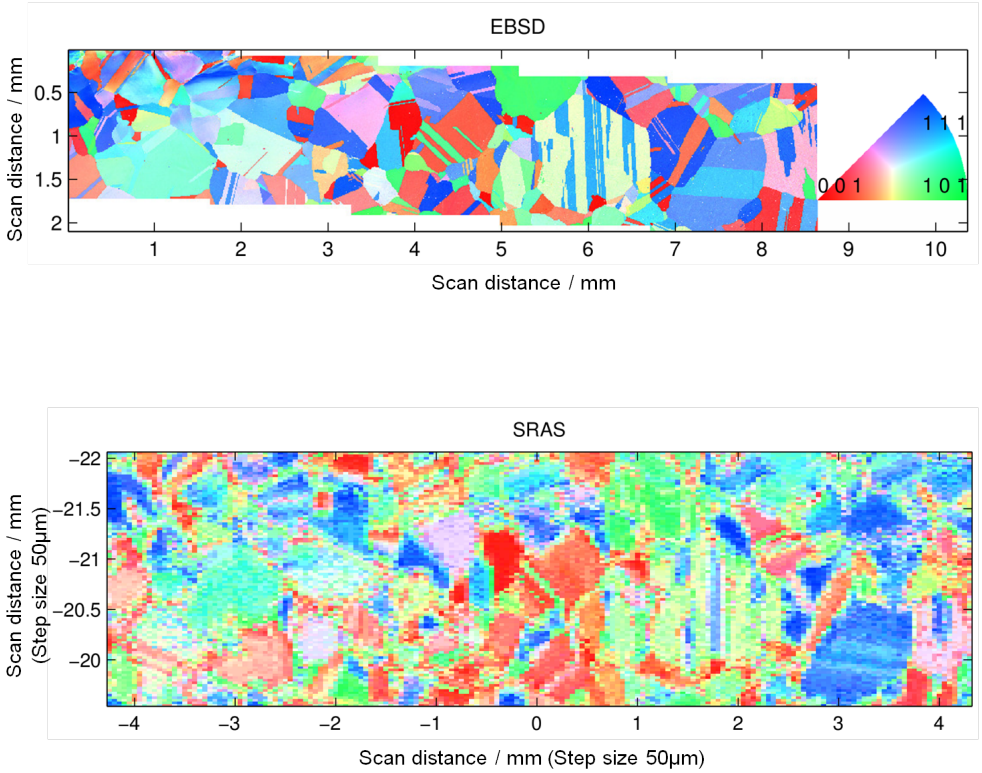


Figure 2.22: A comparison of EBSD data, top, and SRAS data, bottom for a small section of the Inconel sample in Fig. 2.20. No averages were taken during fitting. EBSD data were provided by University of Manchester.

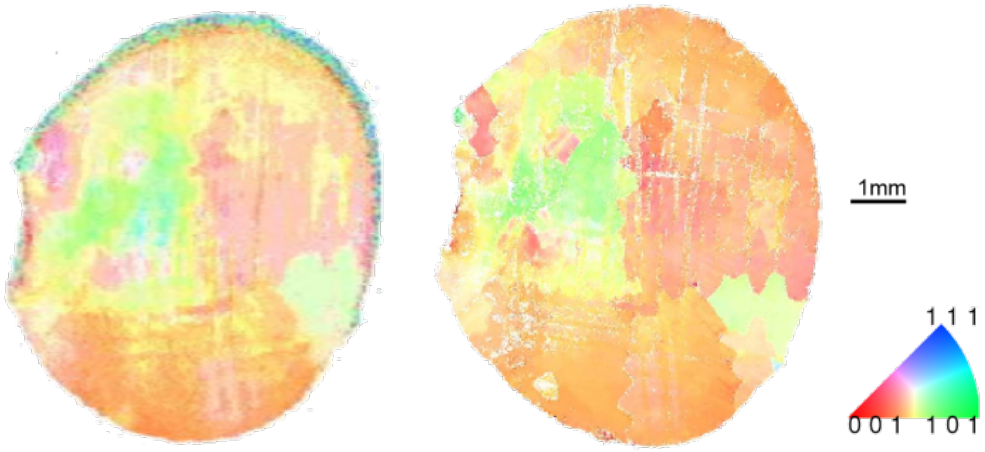


Figure 2.23: An inverse pole figure of a cubic nickel superalloy sample. On the left is the SRAS data and the right is EBSD. The SRAS resolution is 7 times less than the EBSD resolution. 18 propagation directions were used to fit the data [63]. The sample and the associated EBSD data were provided by Rolls-Royce and the University of Cambridge.

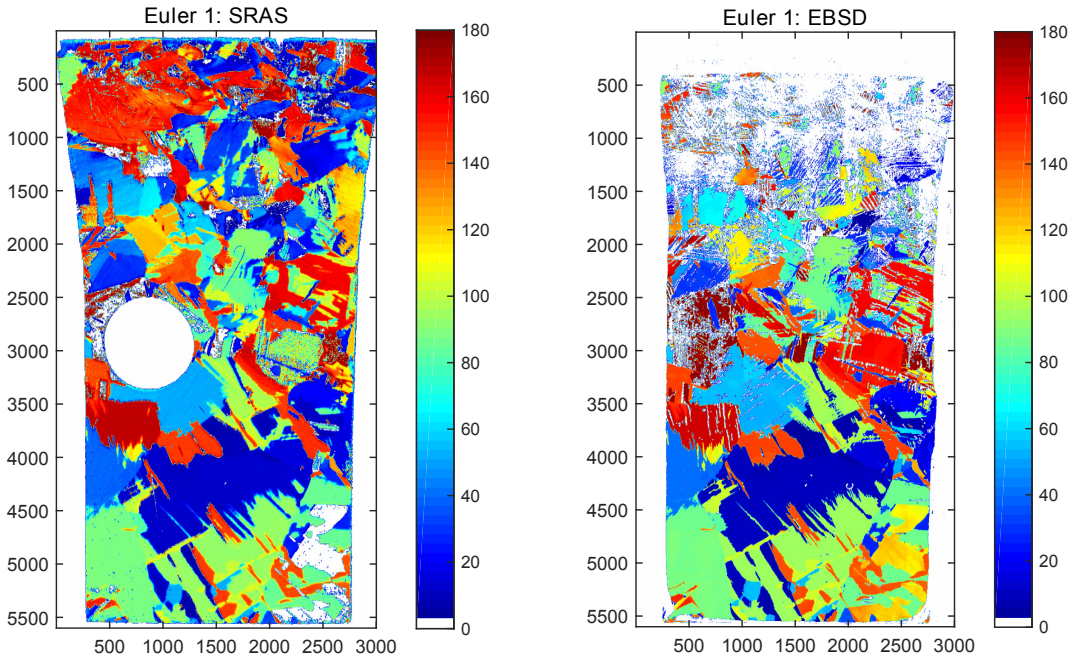


Figure 2.24: A comparison of the first Euler angle, extracted from the orientation of a sample of hexagonal titanium, for EBSD and SRAS. The EBSD data were provided by the University of Cambridge, and the SRAS data were collected and processed by Dr. Wenqi Li.

Chapter 3

Renishaw SRAS system

The EMDA SRAS system at Nottingham works very well in its predefined range: on small, flat and well prepared samples. However it was designed as a test-bed for the SRAS technique and as a laboratory demonstrator and therefore has multiple limitations and drawbacks when considered as a potential industrial prototype. To ameliorate some of the shortcomings of the EMDA system the majority of work done during this EngD project has been on designing, assembling and refining a fully bespoke SRAS instrument better suited to industrial implementation. This chapter focuses and reports on that body of work. The principle areas covered are: optical design, optomechanical and mechanical design, instrument manufacture, electronics and integration with motion platforms, with the final outcome being an operational SRAS instrument far more suited to industrial production and commercial applications. The final instrument is referred to as the Renishaw SRAS system. This work has been done in-situ at Renishaw with full access to the substantial design, prototyping and manufacturing resources available.

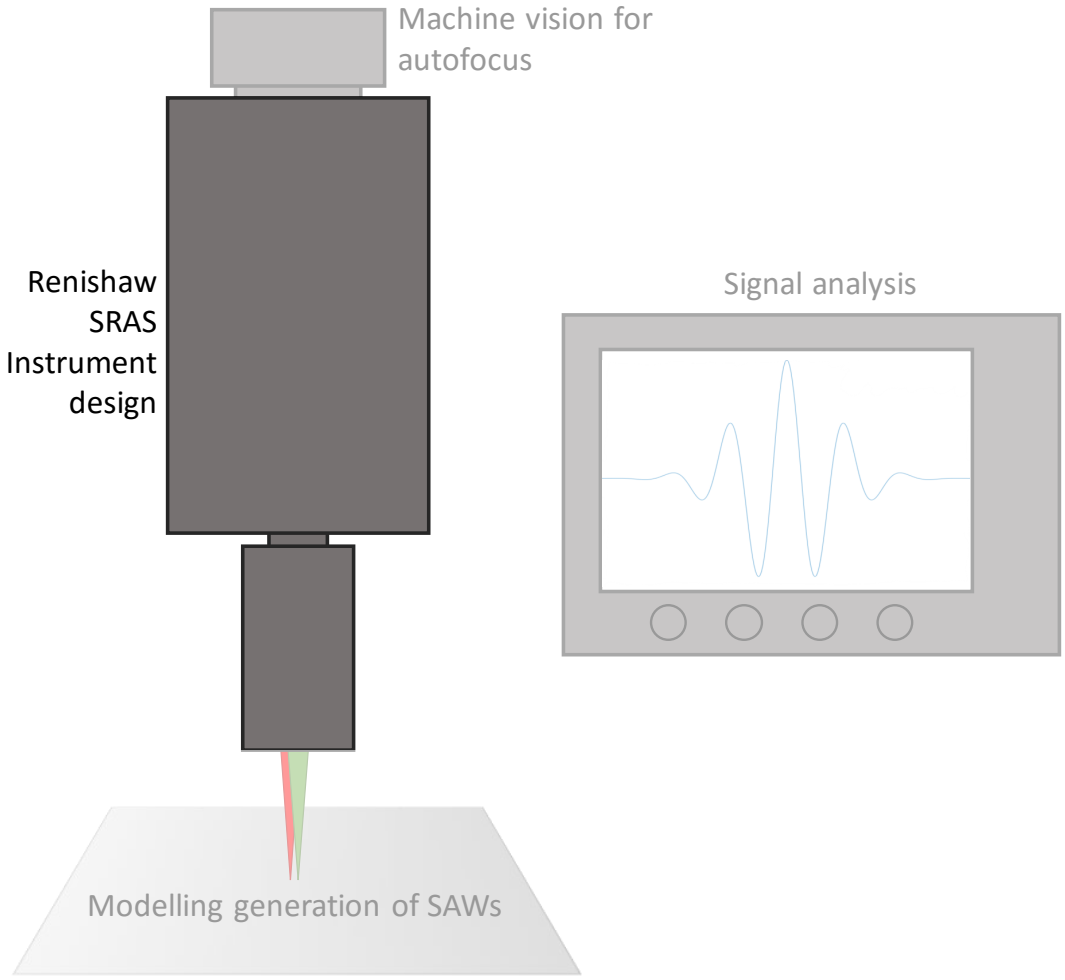


Figure 3.1: This chapter focuses on the optical, mechanical and electronic design of the Renishaw SRAS instrument.

3.1 Customer demand for a SRAS system

3.1.1 SRAS measurements

There are two modes of operation for the SRAS instrument: imaging of grain structure through velocity contrast maps, and full orientation determination. For many applications, some of which are listed in the next section, a simple contrast image of the grain structure is sufficient to identify potential issues with microstructure. From these images the average grain size, typical grain shape, distribution of grains, and changes in microstructure due to processing can be determined.

Full orientation determination of a part requires knowledge of the elastic constants of the material from which the forward model can generate a velocity surface database. A set of 18 SRAS scans must then be taken over a range of SAW propagation directions and each pixel from the scan can then be fitted to the database to determine orientation. This is useful when the absolute orientation of a component is required.

3.1.2 The use of SRAS measurements

As was discussed briefly in §1.4, the majority of contemporary engineering materials are anisotropic and exhibit microstructure. The potential sources and effects of microstructure are numerous and a few examples which are relevant problems in modern engineering are listed here.

Changing the process, opening a new production line or altering an alloy for material manufacture can alter the microstructure of the final product. This can occur as altered material composition and material processing, affect both grain growth and distribution, as well as material crystallography. In this example SRAS would allow rapid feedback for the development of a particular grain structure for new alloys, and validation of expected grain structure for process changes.

Machining of very-high precision parts is dependent on the microstructure, where the level of precision and machined surface form is determined by the interaction of the local grain structure with machine tools [65]. Therefore for high-precision engineering of parts, the material microstructure must first be validated in order to achieve the required tolerances.

Some materials, such as the nickel-chromium alloy Inconel, have inherently large grain structure which makes traditional ultrasonic non-destructive evaluation difficult due to scattering [66]. For the next generation of power plants this is a primary candidate for the high pressures and temperatures desired. In order to perform NDT, the material must have a microstructure which allows proper inspection.

Welding of material results in a much larger grain structure in the heat-affected zone than in the rest of the material. This is particularly problematic in materials such as stainless steel which exhibit typical grain size in bulk material of ~1 µm [67]. The effect of this added large grain structure is reduced resistance to corrosion and reduced yield strength.

Additive manufacturing yields parts with an array of microstructure distributions, particularly in welding type technologies where grain evolution can be aggressive and undesired. Within single components large varieties of grain structure have been observed and the exact effect this has on a part is not yet fully understood [68]. SRAS is capable of providing rapid validation of the microstructure of printed components.

Modern aerospace engines include nearly 200 nickel high-performance alloy single-crystal turbine blades. Compared to multi-grain blades, the single-crystal blades offer up to 9 times the life in terms of creep strength and resistance to thermal fatigue [69]. On a production line these blades need to be inspected for microstructure which is currently achieved by etching and optical inspection. This is a subjective and highly manual process requiring several stages of preparation and post-inspection reconditioning [70]. SRAS would allow objective measurements of microstructure to be taken along with absolute measurement of crystal orientation as a single inspection process.

3.1.3 Market for SRAS

The initial target market for SRAS is the aerospace industry and research community. Significant interest and funding has been directed towards SRAS already by the aerospace industry. Within this sector the interest in component microstructure is very high, as indicated by the example applications, especially where gains on component lifetime can be made.

The research community is the second target for two reasons. The first is exposure of

the technology and validation of the measurements SRAS can make. This is critical for the long-term adoption of SRAS and for building confidence in SRAS measurement accuracy and validity. The second reason is that many future applications for SRAS will start with early stage research. Having an instrument available for this research will help identify these applications and help promote SRAS as a useful tool.

SRAS is still an early stage technology which does not have widespread exposure to industry, and has not yet seen adoption by industry. For Renishaw, SRAS as a product is a significant diversification from existing products and this process involves taking a new technology into an unfamiliar market. While individual applications have been identified it has proved difficult to translate these into sales forecasts. For these reasons the market size and potential for a SRAS product are not well known.

3.2 Technical requirements for the Renishaw SRAS system

There are a great number of ways to discuss and list the requirements of the Renishaw SRAS instrument which depend largely on the interested party. To the potential customer: capability, reliability, total cost and the capacity to improve profits are priorities. To Renishaw of greatest importance is offering customers an innovative solution to a problem and reducing customer costs. In practice, as a business, profit margin, market size, marketability and maintaining the company's reputation as a market leader in innovation and metrology are nearly as important. At the start of the EngD project the actual requirements were not yet defined and a general brief given to simply transfer the technology from the University of Nottingham to Renishaw. Consequently the primary target has been to create an instrument at least as capable as the EMDA SRAS system - in terms of sensitivity, resolution and measurement robustness - and in a form suitable for production manufacture. With this brief the priority throughout the project has been on the mechanical design and optical integration of the instrument with less attention being given to overhauling the established optical implementation of the SRAS technique. Given this focus, a set of requirements can be identified which fall into a few key areas.

Motion

The EMDA system is limited in scanning volume size due to the linear stages used to move samples beneath the fixed instrument. In addition it is unable to cope with non-planar sample surfaces due to the lack of automated Z-axis or tilt manipulation. This architecture is not suitable for industrial application where parts range from millimetres to several metres in size and can be of complex geometries. The obvious alternative is to move the SRAS instrument instead, introducing a new set of challenges. For rapid scanning with accelerations and decelerations up to $10\,\mathrm{m\,s^{-2}}$ a very lightweight and rigid SRAS instrument is required. Optics need to be sufficiently well constrained to prevent dislodgement and misalignment due to motion, and to prevent unintended scatter of laser radiation. Laser sources require fibre delivery due to their bulk and heat generation. To be able to mount an instrument on as wide a variety as possible of motion platforms it must be much more compact than the EMDA system.

Operation

The EMDA system requires manual optical and sample stage set-up for every new sample inspected, and the instrument must be manually reconfigured to alter the propagation direction. All of these processes require class IV laser safety training and are skilled tasks requiring experience and skill. For an industrial instrument all of these processes must be fully automated, or require minimal skill and training to perform. The requirement for users to be laser safety trained is untenable and as such the final instrument must be a Class I laser device.

SRAS is still at a very early stage with only minor industry adoption in the aerospace sector. Early users of an industrial instrument will be University and research labs and the instrument therefore needs to be flexible enough to cope with the variety of demands research requires; the freedom to change acoustic wavelengths, aperture sizes and laser power densities is essential. With a long term outlook a SRAS system may be installed in poorly controlled and hostile environments where ruggedness against thermal gradients, electrical noise and mechanical noise is essential and so this must be considered at the design stage.

Manufacture and assembly

Assembly of the EMDA instrument is a difficult, fiddly and highly skilled task requiring patience and experience in optomechanical systems, and laser safety training. Off-the-shelf optomechanics are expensive and bulky, and the result of their use is a large and heavy instrument which is costly and difficult to make. For an industrial instrument an entirely different production philosophy and route are required, by which the following requirement are achievable: low cost and easy manufacture of parts, non-skilled assembly process with no requirement for laser safety training, and simple integration with motion platforms. A lot of these requirements are covered in the principles of design for manufacture and assembly (DFM), which is an approach to engineering which aims to include input from all stages of prototyping, design and inspection in order to ensure that parts are easy and cost effective to manufacture.

DFM principles cover many areas of design and manufacture:

- limiting material stock, material forms and tooling available for production
- designing parts which can be machined in as few stages as possible,
- designing parts to be multi-purpose so as to limit tooling requirements,
- promoting early communication between production engineers and designers to ensure parts are designed such that critical dimensions and tolerances can be inspected.

Throughout the design process of the SRAS prototype significant consideration has been given to DFM.

3.3 SRAS instrument mechanical design

In the early stages of the SRAS instrument design process many potential assembly types were considered and assessed for feasibility, a brief overview of this process is given in Appendix A. The final design selected is based on a single-piece tube-like solid chassis, into which modular optical housings can be constrained. Cut tube or extrusion is used to separate the optical housings to obtain the correct optical path. Lasers and

electrical cables are delivered from a rack mount cabinet, through an umbilical to the chassis. This section describes how the chassis and individual modular elements work.

3.3.1 Constraint mechanism

There are numerous methods to constrain optics, many of which are used in off-the-shelf optomechanical systems. Many of the optical components to be used have large aspect ratios, i.e. they have a large diameter and are thin. When constraining this shape around its edge it is difficult to ensure that the optical axis is correctly aligned. Fig. 3.2 shows the difference in axis alignment for components with large and small aspect ratios as constrained in a tube or rail system. For a custom instrument the optical constraint system must take account of this. Taking inspiration from camera lenses the idea of constraining optics within tubes was pursued. Fig. 3.3 (a) shows such a system where an optical component is placed within a tube and constrained on its top and bottom surfaces by a pair of smaller diameter tubes. This offers many advantages including: excellent parallelism, easy adjustment in height, well-secured optics, simple construction of assemblies and elimination of the issue found with low aspect ratio components. In order to constrain optics in their flat plane a three point constrain system is applied and this is shown in Fig. 3.3 (b). The circular feature C1, in combination with the square cut-out of the component, effectively constrains the part allowing only rotation around the centre of the circular feature. Applying a small load towards the pair of circular features C1 and C2, fully constrains the part giving excellent centring and repeatability, and avoiding over-constraint.

3.3.2 Chassis design

The described constraint mechanism has been deployed into a custom chassis designed to be manufactured by extrusion, a profile of this extrusion is shown in Fig. 3.4. Consisting of four identical channels the extrusion allows modular optical housings, of form shown in Fig. 3.3 (b), to be slotted into a channel and constrained by a pair of rubber O-rings. Initial prototypes used a plate, as shown in Fig. 3.3 (a), which slotted into the extrusion instead of the O-rings, this did not work well due to the difficulty of assembly and the lack of access to the components once mounted. The O-rings

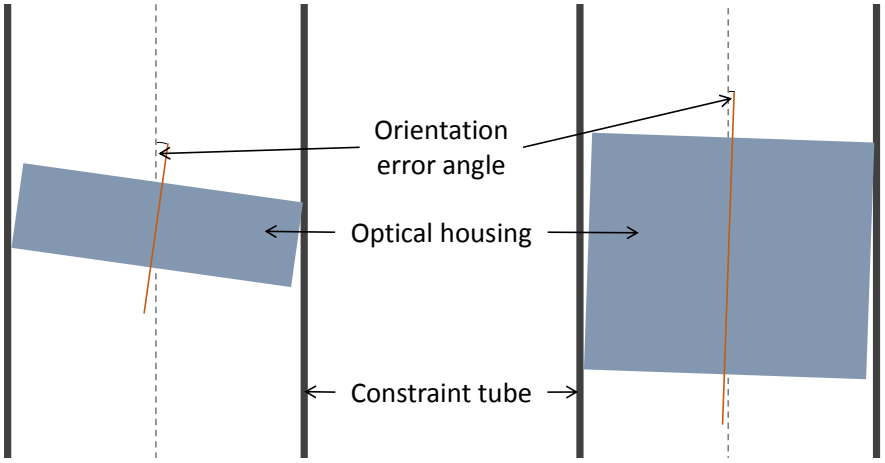
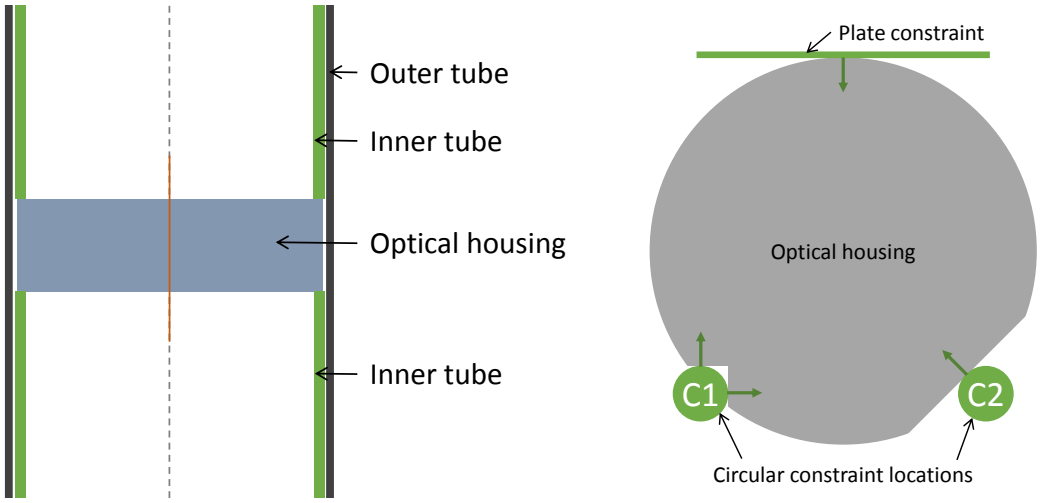


Figure 3.2: A diagram indicating the misalignment of the optical axis caused by high aspect ratio optical components mounted in a tube or on a rail system.

are 3 mm diameter nitrile rubber, other materials were trialled but did not provide a suitable load for good constraint of the optical housings while allowing easy assembly. Each channel has a second pair of openings, adjacent to the O-ring constraint features, which are tapped for assembly such that plates can be added to the top and bottom of the extrusion, as shown in Fig. 3.5. In order to constrain the optical components vertically, cut tube is slotted into the channels and the grub screws in the top plate secure all the components in place. Fig. 3.6 shows an example of the assembled chassis with a lens housing and cut tube slotted into the extruded chassis and constrained by the slightly compressed rubber O-ring.

This extrusion based design has some useful features:

- Manufacture is simplified; the extrusion is simply cut to length and a few slots milled to allow beams to pass between tubes.
- The extrusion is rigid while being very light-weight.
- Tolerances of extrusion, especially straightness and twist, are well within those required ($\sim$0.5 mm over the 182 mm length). Any errors in the beam paths under this tolerance can be corrected using adjustable mirrors.
- A common centre axis is supplied for all components to be precisely constrained with centring to $\sim$0.1 mm.
- Cost of extrusion is very low; around £3500 for tooling and 250 kg of aluminium



(a) (b)

Figure 3.3: a) A part constrained vertically in a outer tube by two sets of inner tubes, this constraint provides parallelism between the part and the outer tube. (b) A circular part with features added to allow an ideal three point constraint to be applied.

extrusion.

- An opening in the profile can be included for removable components such as filters.
- Assembly is extremely simple with pre-cut tube defining the optical path lengths, with parts simply sliding into place in a predefined order.

There are a few downsides to using an extrusion process and the largest of these is the difficulty in prototyping. Although tooling and a run of extrusion is low cost on a production scale, it is not well suited to one-off units. An additional issue is surface finish; extruded parts often show major surface defects and marks after manufacture, requiring multiple finishing processes to achieve the Renishaw standard for customer visible faces. Finally, the modular parts intended to slot into the extrusion are somewhat limited in size and form, as such standard parts such as automated mirror mounts cannot be used.

3.3.3 Optical housing design

With the optical chassis design finalised, the engineering of the optical components follows fairly organically. The outer profile of each component matches the profile

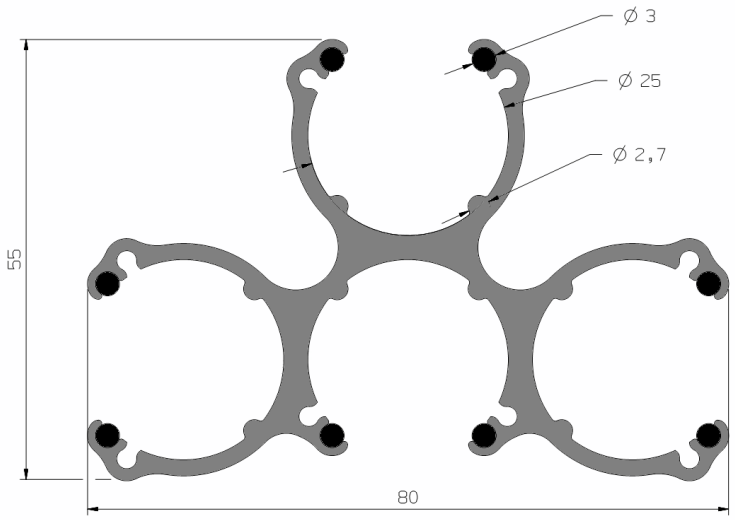


Figure 3.4: Profile of the aluminium extrusion which is used as the main chassis for the SRAS head, all measurements in mm. The black rubber features provide a pre-load to the inserted modular components, constraining them against the circular features of the chassis.

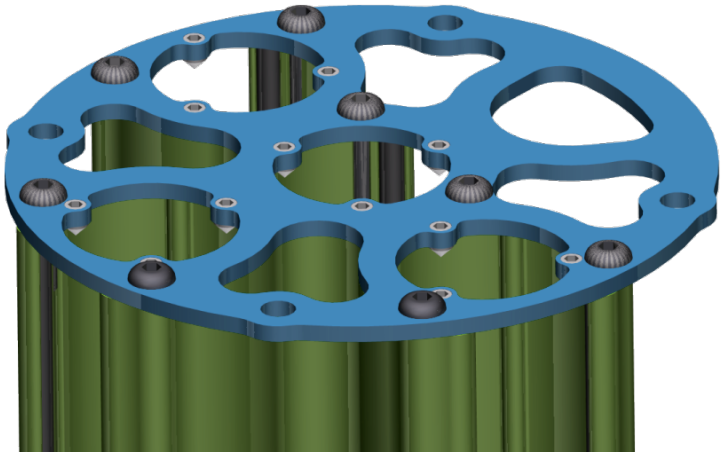


Figure 3.5: A diagram of the cap plates for the extrusion which are attached by threaded fasteners. The grub screws are used to constrain the optical housings and cut tube in the vertical axis.

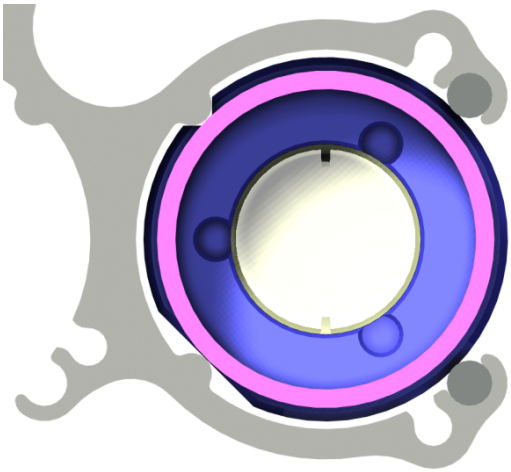


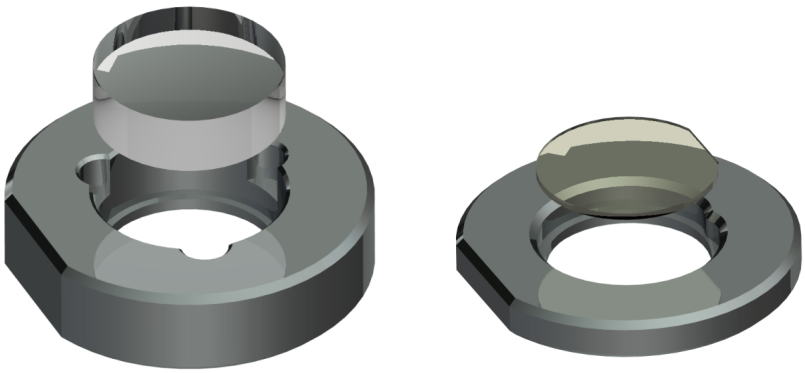
Figure 3.6: A rendered image showing an example assembly of the chassis, optical housing (blue), cut tube (pink) and rubber O-ring.

shown in Fig. 3.6 with the internal form dictated by the optic. DFM is a large factor in the engineering design of the optical housings and where possible, parts are turned rather than milled and serve multiple functions. Due to the target size of the instrument head, most of the housings are designed for the use of adhesive to secure the optics.

Lens and wave-plate mounts

Multiple lenses of varying focal length are used throughout the instrument and to simplify the assembly all lenses share an outer diameter of 12.5 mm. Fig. 3.7 (a) shows the final design for the lens holder. Lenses are secured in a counter-bored hole deep enough to protect the lens surfaces. Three circular cut-outs are used to ease the fitting of lenses and for the application of adhesive. Since lenses are rotationally invariant, no features are required to constrain rotation. Due to the differing curvatures and geometries of lenses, precisely where they locate relative to the lip of the counterbored hole varies. To remove errors in calculating lens focal lengths relative to the lens mount, a space model for each lens has been generated and used to find the lens surface positions relative to the lens mount.

Waveplate mounts are similarly designed but with the addition of a feature to constrain the quarter wave plate (QWP) at 45° to the polarisation of the detection laser. At this angle the QWP converts linearly polarised light into circular, and vice-versa, this



(a) (b)

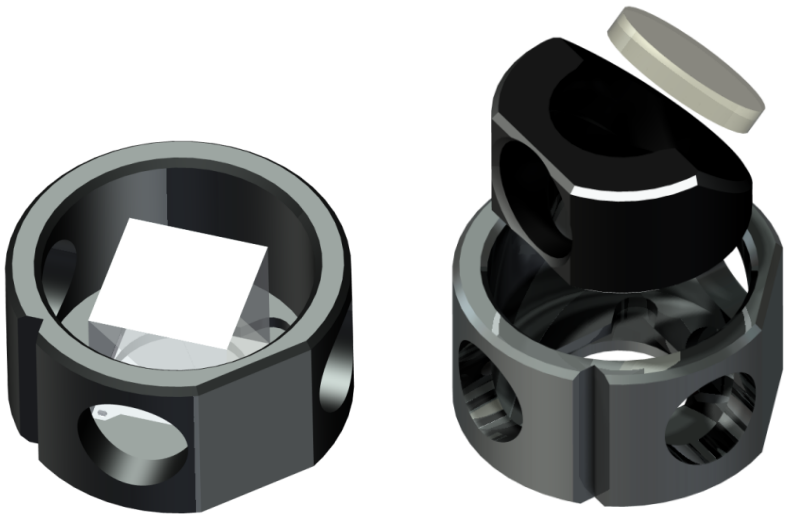
Figure 3.7: a) 12.5 mm lens holder and (b) 15 mm waveplate holder with alignment feature for conversion between linear and circular polarisation.

part is shown in Fig. 3.7 (b) and is also suitable for manufacture by mill-turn. While this part could be thinner, in the interests of DFM, the minimum depth is set by the manufacturing technique.

Mirror mounts

There are four mirror requirements in the chassis: flat mirrors - such as dichroic mirrors, right angle mirrors, beam splitter cubes and steerable mirrors. A multi-purpose part was designed to allow 10 mm square base beam splitters and right-angle mirrors to be held and is shown in Fig. 3.8 (a). This part has a milled square pocket in which adhesive is used to secure mirrors in whichever orientation is required. To enable constraint of flat mirrors a simple plastic adapter has been designed which can be slotted into the square mirror holder at the desired orientation, this is designed such that the front face of the mirror is the reflective surface and maintains beam centricity. The square mirror holder is a turned component, while the adapter is designed for injection moulding but is currently made by additive manufacture.

While the design of the chassis and optics holders should keep laser beams central and parallel to the chassis channels, there is still a requirement for beam steering in order to position the visible and infra-red beams on the sample, and to correct for small manufacturing tolerances in the chassis and fibre launch optics. Ideally motorised mirror mounts would be used but, for the compact extrusion, off-the-shelf solutions are too large to fit. A study of potential motorised custom solutions was undertaken



(a) (b)

Figure 3.8: a) A PBS cube mounted in a mirror holder. (b) A plastic adapter is used to hold flat mirrors such that the beam strikes the centre of the mirror.

and the added complexity, mass and control electronics required was deemed to be to great a penalty for the benefit motorised mirrors give. The current solution is to use a more simple manual design, and this is shown in Fig. 3.9. Fine thread, low backlash, brass inserts and ball tipped screws are bonded into the aluminium surround while two strong neodymium magnets are mounted into the insert and surround and constrain the insert against the brass screws. The fine M2 thread allows very precise and stable positioning of the mirrors for beam steering.

Grating, aperture and filter holders

As per the requirements stated in §3.2 the gratings, filters and apertures need to be changeable by the user of the instrument. In addition a kinematic mount is required so that as gratings are replaced their location relative to the generation beam is consistent. With the design of the chassis there is no access to optics from below or above, only through the open side of each channel of the extrusion. A cartridge system has been designed which can be inserted from the side of the chassis and uses spring plungers for constraint and location, this is shown in Fig. 3.10. A permanent holder is included in the chassis during assembly which allows either the filter holder or mask holder to be inserted and removed. Two spring plungers act with locating features to form a

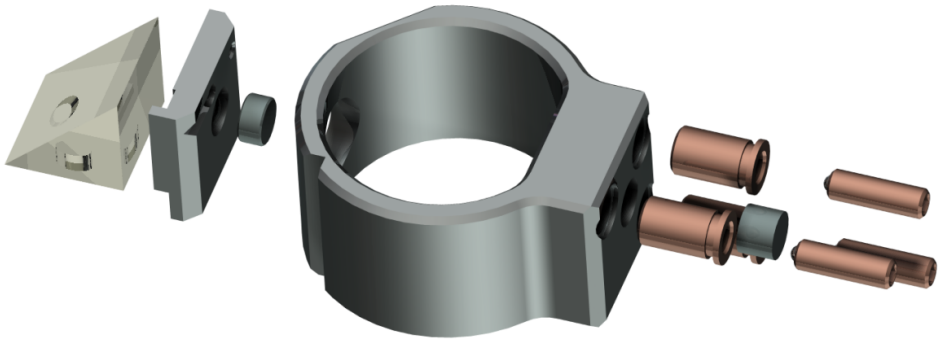


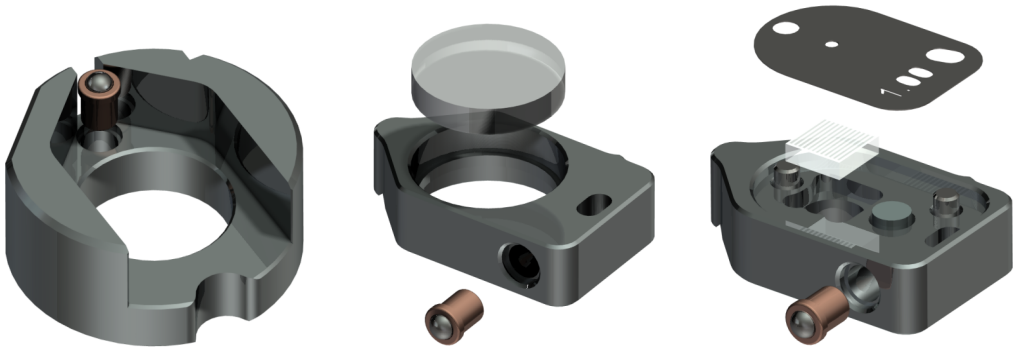
Figure 3.9: An exploded view of the adjustable mirror assembly. The mirror is mounted into an insert which is magnetically constrained against brass screws in the outer section. Finely pitched brass inserts and round tipped screws are used to manipulate the orientation of the mirror.

repeatable three-point constraint system. The filter holder has a simple counter-bore to allow 12.5 mm diameter filter plates to be held. The slightly more complex mask holder uses a square pocket into which the glass grating is bonded, and locating dowel pins and a magnet to secure the steel apertures. A small slot at the front edge of each cartridge allows them to be easily removed from the chassis.

Gratings are manufactured from chrome-coated soda lime glass which are patterned by photolithography, Fig. 3.11 shows a glass plate with various line spacings patterned into the surface. The glass plate is then sliced into 5 mm squares which fit into the custom mask holders. Although the first production run of masks were not coated for anti-reflection at 1064 nm, future batches will be coated before slicing. Apertures are laser cut from 100 µm thickness stainless steel in-house with the capability to create pinholes as small as 200 µm.

Fibre collimators

The generation laser outputs short pulses of high instantaneous power; ~10 kW over 2 ns to 3 ns. This necessitates the use of a 200 µm core, multi-mode fibre which can handle the thermal shock without damage. An unfortunate consequence of this is that the output beam from the fibre cannot be well collimated and suffers from large



(a) (b) (c)

Figure 3.10: (a) The cartridge holder permanently mounted into the chassis assembly. (b) A filter cartridge which is inserted into the cartridge holder and located by spring plungers. (c) A mask and aperture cartridge which allows a selection of different mask gratings and aperture sizes to be inserted into the chassis.

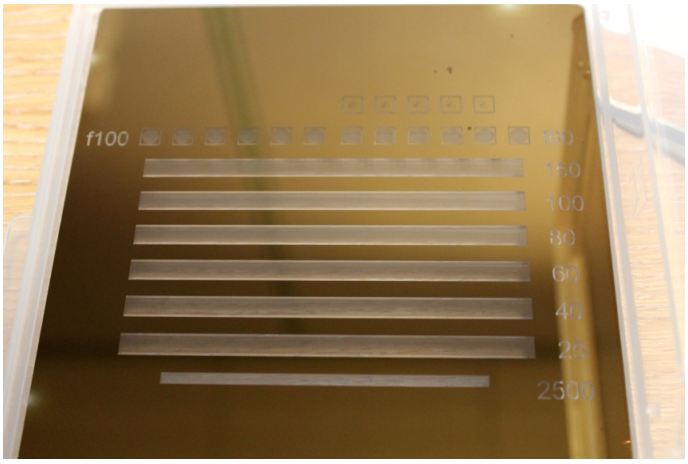
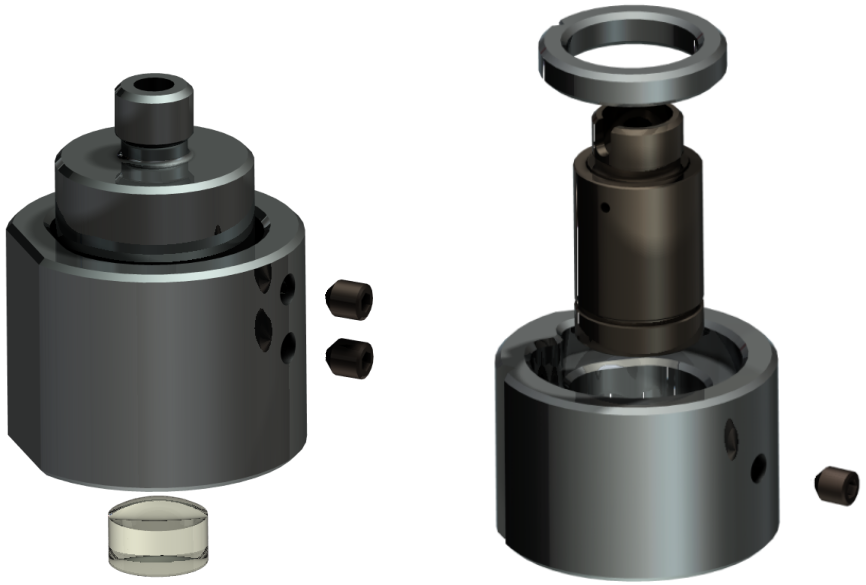


Figure 3.11: A 5" square chrome on glass mask plate manufactured by J.D. Photo. A selection of gratings have been etched in preparation for cutting into smaller masks.



(a) (b)

Figure 3.12: a) A bespoke fibre collimator for the generation beam. An achromatic lens is mounted in an aluminium holder and an SMA adaptor allows the focal distance to be adjusted. Grub screws are used as part of an internal three-point constraint. (b) An adaptor for a standard Thorlabs collimation package for the detection beam. A collar is applied before assembly to ensure correct polarisation.

divergence. Off-the-shelf collimation packages proved inadequate in limiting beam spread and a bespoke solution was designed and built, this is shown in Fig. 3.12 (a). An anti-reflection coated f10 achromatic lens acts as the collimation lens, with a sliding fibre adaptor allowing fine adjustments to be made in distance to the fixed lens. Thus by adjusting the sliding mount and monitoring laser power, optimal collimation can be achieved. Optical adhesive is used to bond the lens to the holder while nylon tipped grub screws hold the circular sliding fibre adaptor in a three point constraint configuration.

The visible detection laser beam is coupled into a single-mode, polarisation maintaining fibre with a core diameter of ~5 µm. As such collimation is trivial and a Thorlabs fixed focus, narrow-key, angled fibre collimation package, pre-aligned for a 532 nm beam, is used. With a focal length of 10.9 mm, this gives a beam diameter of 2.1 mm with diffraction limited dispersion. The detection beam path relies on fixed polarisation and so, to ensure the collimator is correctly orientated, an optical jig based on a PBS and

power meter is used to align an orientating collar to the collimator, this can be seen in the exploded view shown in Fig. 3.12 (b). With this set-up the entire collimator can be removed from the SRAS head during assembly and disassembly, which is preferable to handling the exposed delicate fibre output connector. Nylon tipped grub screws hold the inserted collimator in a three point constraint configuration providing repeatable centring of the output beam.

KED case

The knife-edge detector is composed of a small PCB and a segmented photodiode. Historically, the detection electronics of the EMDA system have been fairly sensitive to the electrical coherent noise broadcast by the Q-switch of the generation laser, which occurs in the same time frame as the SRAS signal. As a result, for the Renishaw SRAS system, significant engineering effort has gone into shielding of the detection electronics to ameliorate the noise issue. Fig. 3.13 shows an exploded view of the KED case and components. Two extruded sections and two turned end-plates are combined to create a tightly sealed box, the round extrusion can be removed for easy access to the PCB. Dowel pins and screws are used to constrain and assemble the case. Filtered feed throughs remove RF signals from the DC power and return DC light level signals, and a ground pin is tightly bonded to the case providing a low impedance path for high-frequency noise. The segmented photodiode is tightly constrained and centred in a close fitting plastic collar which also determines the direction of sensitivity of the photodiode, in most cases the detector is sensitive to beam deflections parallel to the PCB. Through these engineering efforts the coherent noise from the generation laser has been completely eliminated, leaving only incoherent noise on the measured signals.

Twin photodiode KED

Due to changes in supplier inventory the original Hamamatsu segmented photodiode used in the EMDA KED is no longer available. While there are alternative suppliers of similar devices they do not match the low capacitance and large active area of the discontinued photodiode. Therefore an alternative design was proposed which uses two independent high-speed photodiodes to achieve a knife-edge, with a cube prism used to

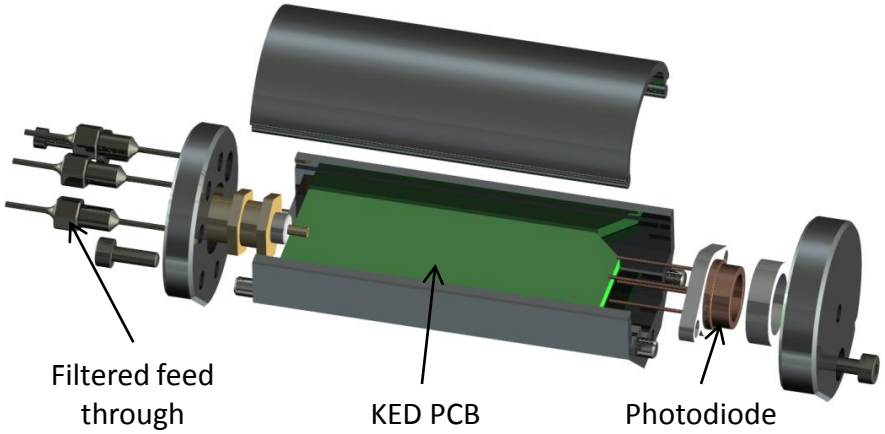


Figure 3.13: An exploded view of the KED case with PCB and photodiode. Filtered feed throughs are used to channel DC signals into and out of the case.

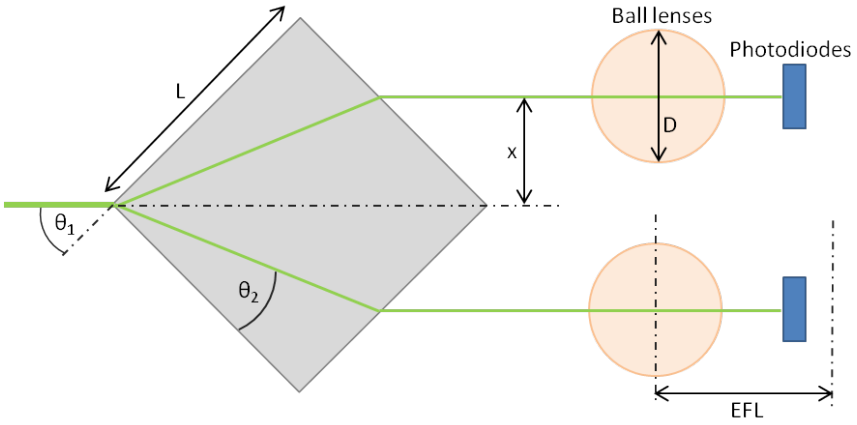


Figure 3.14: Schematic of the optically split, twin photodiode knife-edge detection scheme. The input beam is split at the corner of the cube with the split beams focussed onto two photodiodes by ball lenses.

optically split the detection beam. A schematic for this twin photodiode KED is shown in Fig. 3.14. The advantage of using a cube, over other prisms, is that the split beams are parallel to the original beam and only have a separation of a few millimetres. Thus the two single element photodiodes can be placed in close proximity on the same PCB used for the original KED, allowing a high-frequency differential signal to be detected. An additional benefit comes from the use of ball lenses to focus the segmented beam, allowing for small active area photodiodes to be used. A comparison of the twin KED and segmented KED is given in Chapter 5.

In order to calculate the split beam separation, Snell's Law is used;

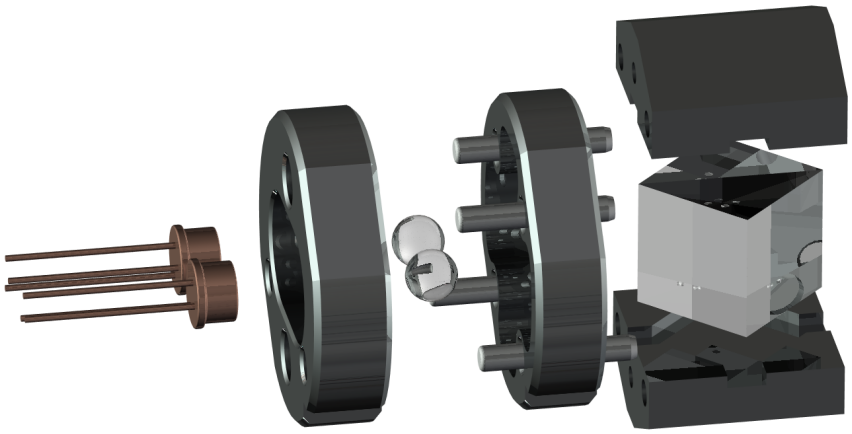


Figure 3.15: An exploded view of the twin photodiode knife-edge detector assembly. The top components holding the photodiodes and circuit board have been omitted for clarity but are shown in Fig. 3.13.

$$\frac{n_2}{n_1} = \frac{\sin(\theta_1)}{\sin(\theta_2)}, \tag{3.1}$$

where n_2 and n_1 are the refractive indices of the prism and air respectively. Based on the geometry of Fig. 3.14 clearly $\theta_1 = \pi/4$ and n_1 can be taken as unity. Then by analysing the geometry it is found that

$$x = L\frac{\sin(\frac{\pi}{4} - \theta_2)}{\cos(\theta_2)}. \tag{3.2}$$

Thus, combining the equations and assuming a 10 mm cube made from standard optical BK7 glass [71], it is found that $x = 3.29\,\text{mm}$. A pair of ball lenses are then used to focus each beam onto a photodiode, the offset between the photodiodes and lenses is determined by calculating the effective focal length (EFL) of the ball lens using [72];

$$\text{EFL} = \frac{nD}{4(n-1)}, \tag{3.3}$$

where n is the refractive index of the lens and D is the ball lens diameter. To avoid damage the photodiode is offset a small distance, ~0.5 mm, from the lens focal plane. Based on these dimensions an assembly has been designed and is shown in Fig. 3.15.

3.4 SRAS optical design

Like the EMDA SRAS instrument the Renishaw instrument has two independent optical paths which are combined to form the generation patch and detection spot on the sample. The design of the chassis lends itself very well to optical isolation of each beam path, as such each channel is dedicated to a particular optical task with beams passing between channels through milled slots.

3.4.1 Detection path

A simplified schematic of the detection path is shown in Fig. 3.16. Via optical fibre, a single longitudinal mode, polarised beam is delivered into a collimator which outputs a low divergence 2.1 mm diameter visible beam, an optional neutral density (ND) filter allows the beam power to be attenuated if the detector is being saturated. The beam is then turned by an adjustable mirror towards a polarising beam splitter (PBSg) coated and aligned for reflection of the S-polarised visible beam towards the sample. A $\frac{\lambda}{4}$ QWP aligned at 45° to the axis of polarisation converts the S-polarised beam to circularly polarised, the handedness of which depends on the assembly of the QWP in the chassis. Passing through a dichroic mirror, coated for reflection of 1064 nm and transmission of 532 nm, the beam is focused to a spot on the sample by the f20 objective lens. The first adjustable mirror allows the beam to be translated across the sample surface into the required location relative to the generation patch.

In reflection of the beam from the sample the polarisation handedness is inverted. The beam is collected and re-collimated by the objective lens before passing through the dichroic mirror. The QWP then converts the circularly polarised beam to a P-polarised beam which is transmitted by the PBSg to a cold mirror, which turns the beam towards a second adjustable mirror. This mirror directs the return beam through a 532 nm-filter to a lens which focuses the beam close behind the photodiodes of the KED. The positioning of the lens is such that the beam spot fills the segmented photodiode and the adjustable mirror facilitates balancing of the beam between the two segments. Sufficient beam intensity is transmitted through the cold mirror to allow a lens to form an image of the sample surface on the camera sensor. An optional filter is used to avoid saturation of the camera.

Fig. 3.17 shows the actual implementation of the detection beam path in the extruded chassis, the components used are labelled in Fig. 3.16. The objective lens is a 10X Mitutoyo long working distance objective, corrected for the visible to NIR range. Orientation of the segmented photodiode in the KED is such that only SAW packets propagating into and out of the plane of Fig. 3.17 are detected.

For assembly, optical paths are defined by the cut tube lengths and orientation of the optical holders. The only critical optical length in the detection path is the 14.2 mm between the detector lens and photodiodes, which has a tolerance of 100 µm. All other path lengths have a tolerance of 250 µm due to the corrections available with the adjustable mirrors.

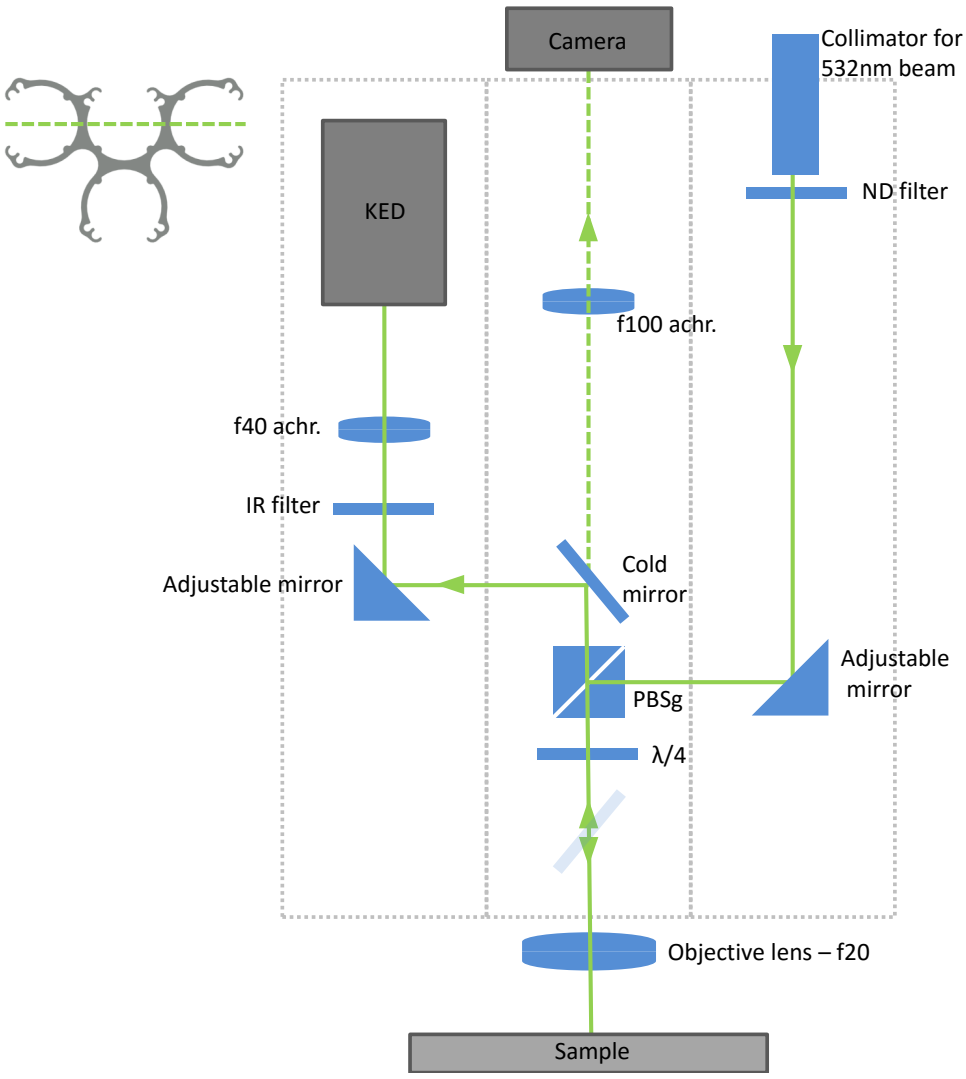


Figure 3.16: A simplified schematic of the optical detection path and components required. The path lengths between optics are shown in Fig. 3.17.

3.4.2 Generation path

A simplified schematic of the generation path is shown in Fig. 3.18. Via a multi-mode fibre a pulsed IR beam is delivered to the custom collimator shown in Fig. 3.12, which outputs a high divergence beam with starting diameter ~3 mm. This beam passes through an aperture and grating holder before reaching an optional reflective ND filter which can be used to attenuate the beam power. Attenuation by filter is preferable as the characteristics of the laser pulse are adversely altered if the laser pump current is adjusted to control power. An f100 achromatic lens, placed with the

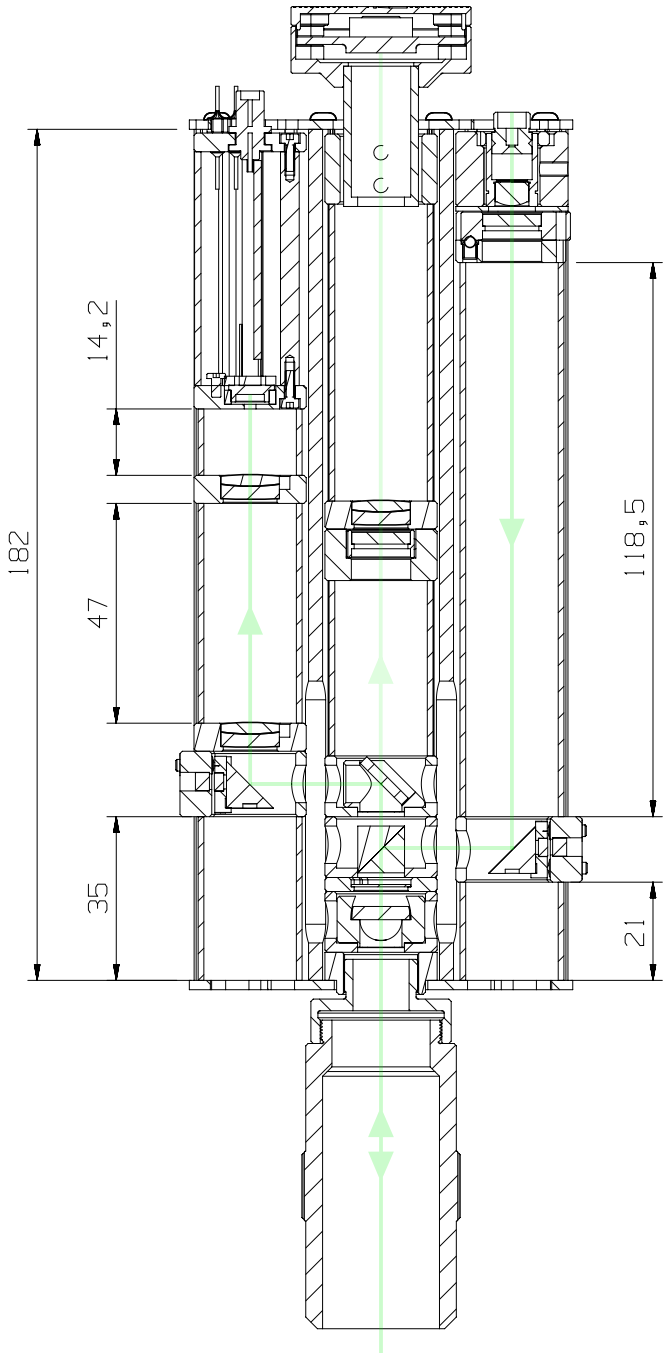


Figure 3.17: A to scale schematic of the detection path as implemented in the SRAS instrument head. The component labels can be found in the simplified schematic shown in Fig. 3.16.

focal plane at the mask surface, then collimates the grating image and focuses the laser beam. Due to the large divergence the beam is not brought to a focus by this lens, instead it acts to improve collimation of the beam. An adjustable mirror then turns the beam towards a dichroic mirror, coated for reflection of 1064 nm, which then turns the beam again towards the sample. The Mitutoyo objective lens then focuses the grating image and beam onto the sample, the high quality optic gives a very well focused image. In combination with the f100 achromat, the f20 objective lens acts to reduce the grating size by a factor 5; as such a grating with 125 µm spacing will generate SAWs of acoustic wavelength 25 µm. Reflected IR light is captured by the objective and sufficient intensity is transmitted through the optical system to form an image on the camera, the cold mirror used in the detection path is critical for this.

Fig. 3.19 shows the actual implementation of the generation beam path in the extruded chassis, the components used are labelled in Fig. 3.18. As with the detection path, the critical optical lengths are determined by the cut tube lengths for ease of assembly. Of note here is the position of the collimator in the chassis. Due to the poor collimation, the beam source is located as close as possible to the sample but still allowing the required optical path to form the grating image. This ensures the minimal spread and maximises energy density on the sample. From the perspective in Fig. 3.19, the detector is sensitive to SAWs propagating in a left-right direction, therefore masks must be mounted with lines running into the plane of the image. Only one cut tube dimension is critical; the 81.75 mm length between the filter and f100 achromat. This controls the focal plane of the grating image in the sample surface, and this must be coincident with the point of optimum energy density and the focal point of the detection spot. The distance between the camera lens and camera sensor is also important and this is configured by focusing the camera, with the camera lens, at infinity before attaching the objective.

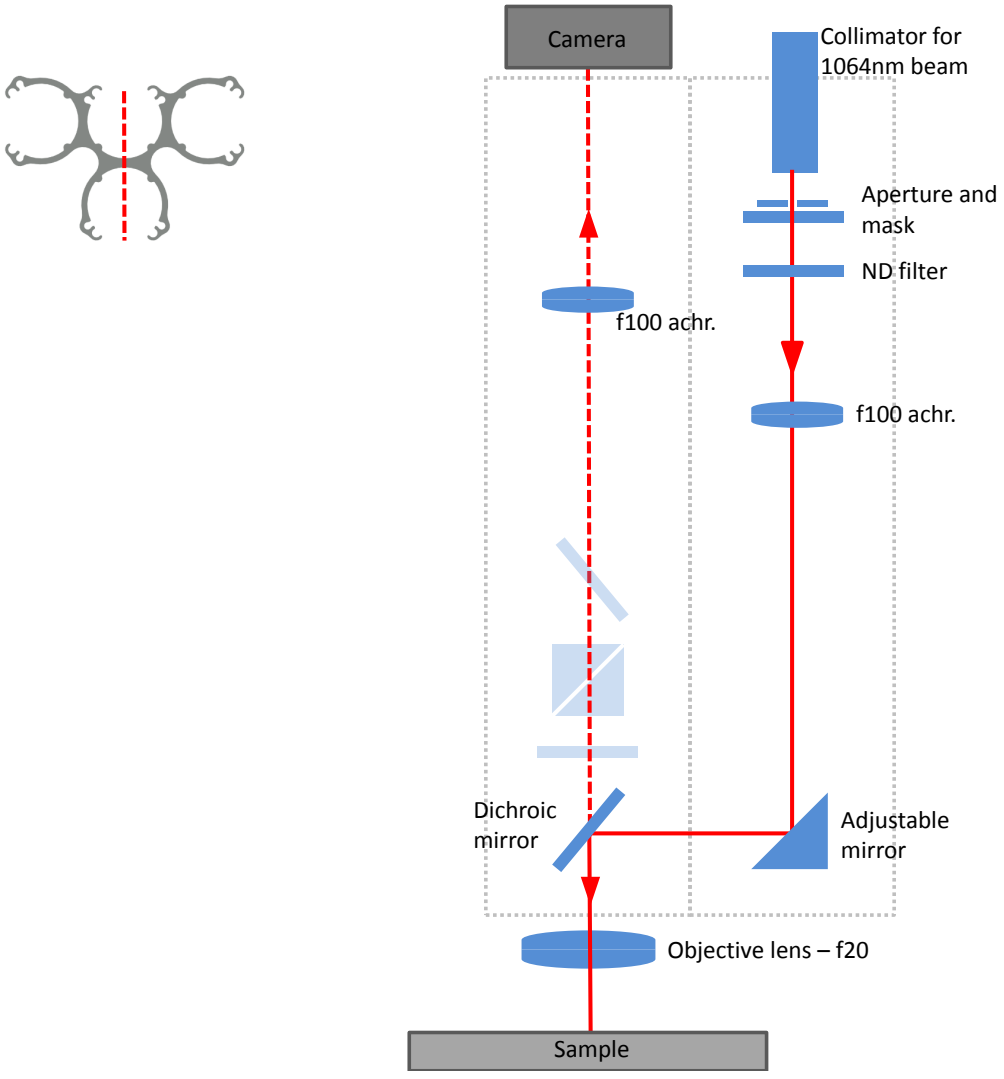


Figure 3.18: A simplified schematic of the optical generation path and components required. The path lengths between optics are shown in Fig. 3.19.

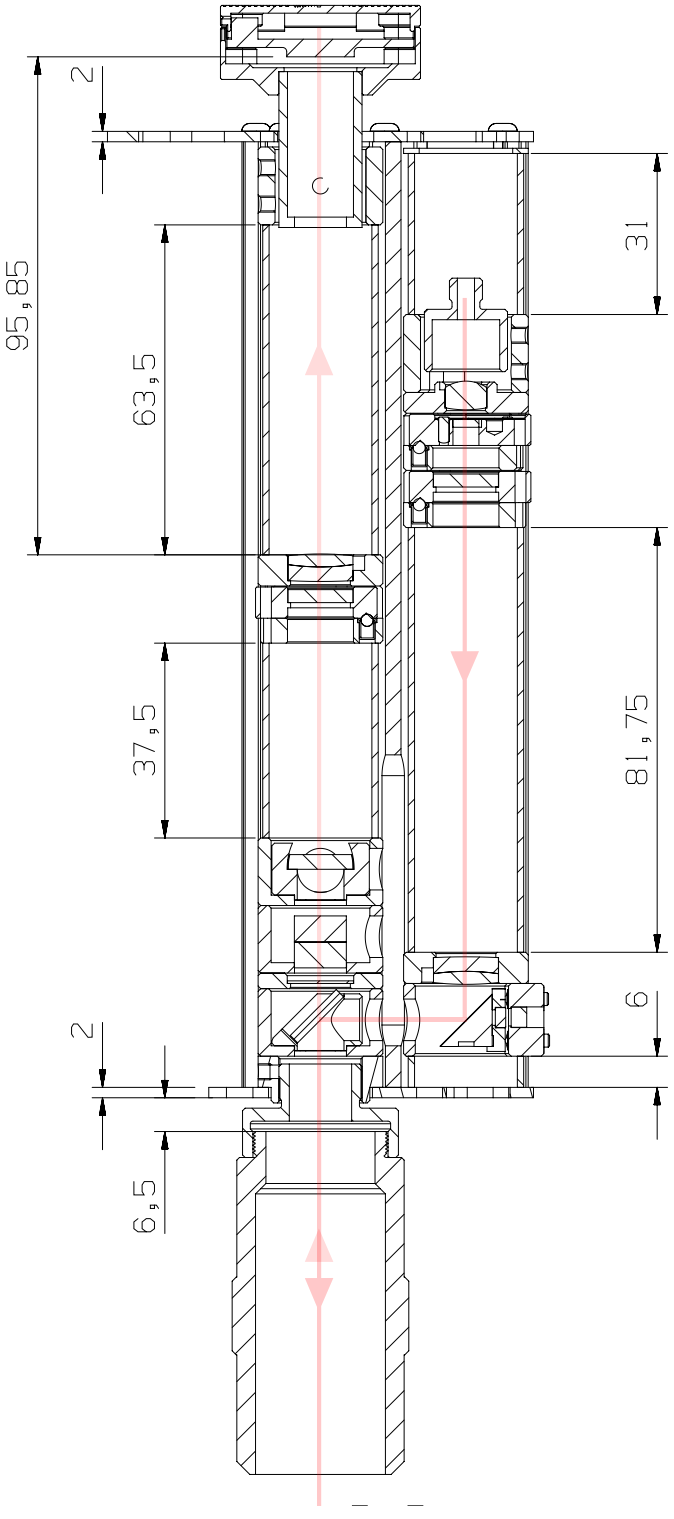


Figure 3.19: A to scale schematic of the generation path as implemented in the SRAS instrument head. The component labels can be found in the simplified schematic shown in Fig. 3.18.



Figure 3.20: A selection of internally made optical mounts with optics bonded by UV cure adhesives.

3.5 Manufacture and assembly

A process of design, prototyping, refinement and manufacture has been ongoing at Renishaw throughout the development of the SRAS instrument. Additive manufacture and rapid prototyping methods have been used extensively to expedite the design process, with the final designs of parts manufactured through the Renishaw machine shop. For assembly, optics were bonded into the machined components using a UV cure adhesive and a UV chamber, although the adhesive is matched in refractive index to the glass, care was taken to keep active optical surfaces clean. The advantage of the UV cure is the ability to precisely tweak the positioning of the optic within the semi-cured adhesive. In terms of optical constraint all of the designed housings worked as expected, after some revision, and assembled easily, Fig. 3.20 shows a small selection of assembled optics and housings. Although most of the parts shown are raw aluminium in the long term they will all be black anodised like the lens holders shown.

As discussed briefly in §3.3.2 the disadvantage of using an extrusion process for manufacture is the difficulty in prototyping and for this reason numerous methods of prototype manufacture were attempted; primarily plastic additive manufacture and metal additive manufacture, this process is described in more detail in Appendix A. For the

working prototype the chassis manufacture route was by electrical discharge machining which is well suited to the profile required, slots were then post-machined into the chassis. Aluminium tubes were cut to length on a mill-turn, and then filed and sanded to meet the tolerance requirements. For long term production these parts will be laser cut from thin walled stainless steel or aluminium tube.

With all the required parts manufactured a statically mounted prototype assembly was built, this is shown in Fig. 3.21. This assembly then went through a process of optimisation and refinement to arrive at the final design shown in Fig. 3.17 and Fig. 3.19. The processes of assembly and optical alignment are fairly simple and easy to perform. In terms of assembly parts simply slide into place, already constrained by the rubber O-ring material, before the extrusion top plate constrains them rigidly. Since the optical paths were pre-calculated and the tube cut to the correct length, the location of optics requires no adjustment. Adjustable mirrors can be pre-aligned by looking down the centre of each channel and adjusting until the final exit aperture of the instrument is well centred, thus when the lasers are introduced only small tweaks are required.

With regard to the suitability for industrial manufacture the assembly is not ideal. In general, when designing for assembly, the principle of poka-yoke is followed which is essentially mistake-proofing by design [73]. For example connectors only fitting in a single orientation or mating parts only fitting into an assembly at one location. This allows leaner, less skilled and faster assembly, in turn reducing the overall instrument cost. Since the SRAS head assembly is designed to be modular the risk of a component being placed in the wrong location or in the wrong orientation is high and, dependent on the mistake, errors are difficult to diagnose. As such assembly is a skilled task and the primary method for preventing errors is good labelling of parts and thorough assembly work instructions.

3.6 Lasers, electronics and data acquisition

The lightweight chassis is designed to be mounted on motion platforms and manipulated, and so the lasers, electronics, power supplies and data acquisition are mounted separately. A sturdy umbilical houses and protects the cables and fibres, delivering

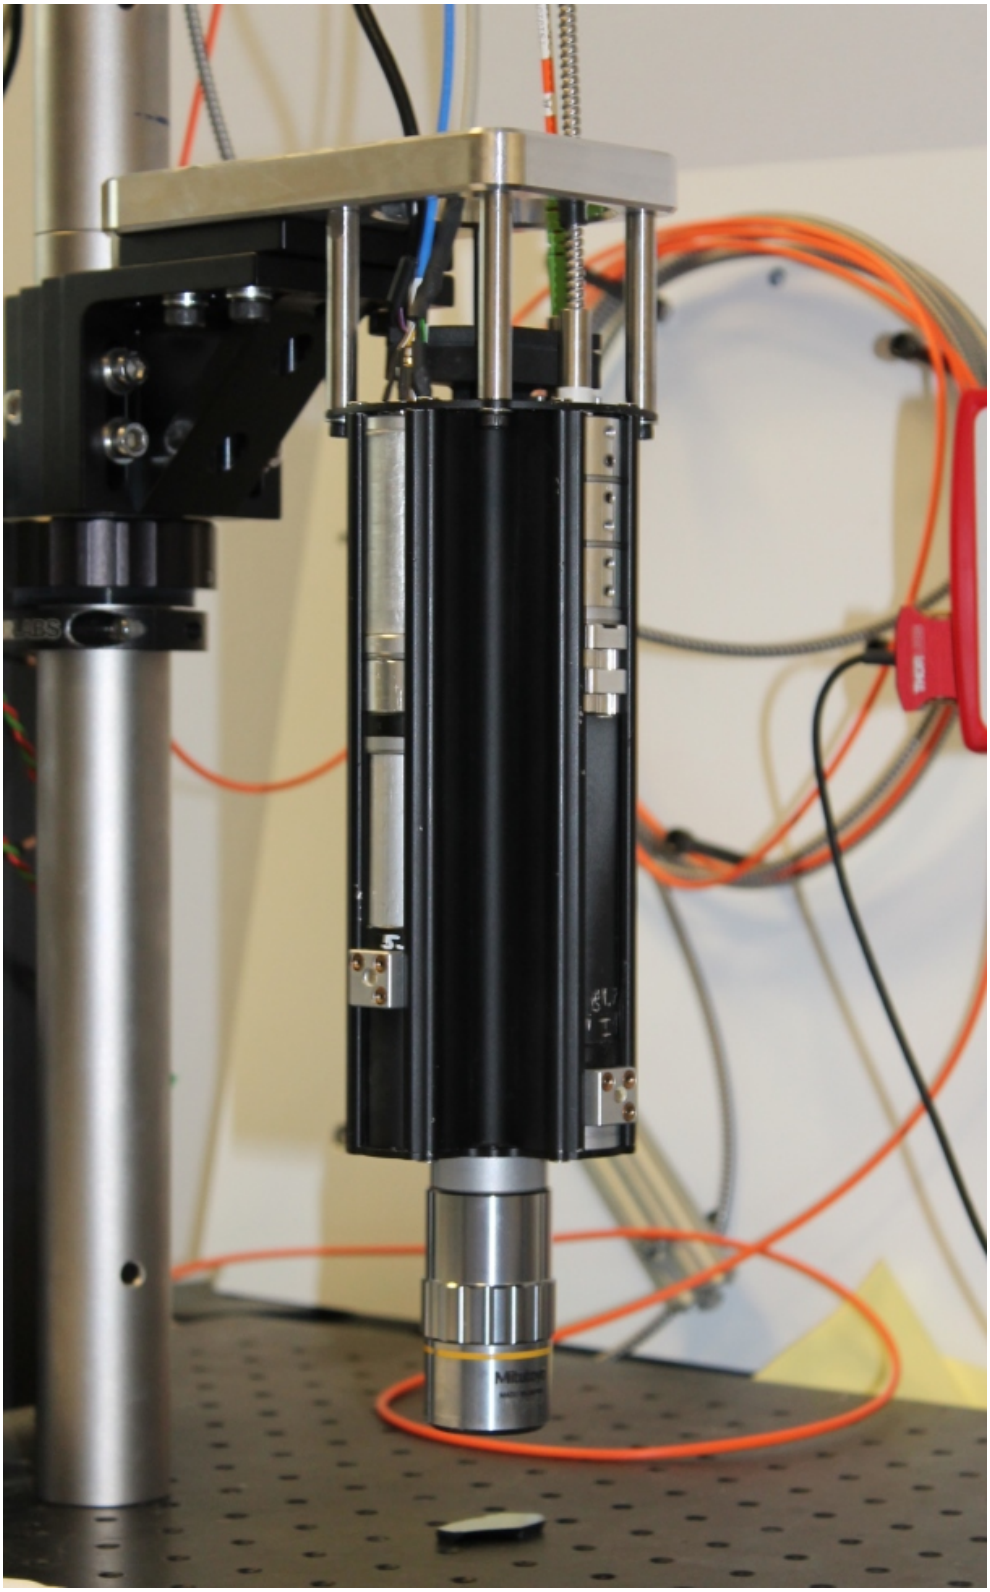


Figure 3.21: An assembled and statically mounted SRAS head. The generation fibre is the stainless steel coil entering the top of the chassis and the detection fibre is orange.

the required signals, power and light to and from the chassis. The hardware required to power the SRAS instrument are housed in a rack mount.

3.6.1 Laser sources

There are two fibre-coupled lasers required for SRAS: one for detection of SAWs and one for generation of SAWs. Both of these lasers are Class IV and require multiple safety measures such as: suitable light attenuating protective eye-wear, full enclosure of the SRAS instrument, and use of interlocks to prevent unauthorised entry to the enclosure.

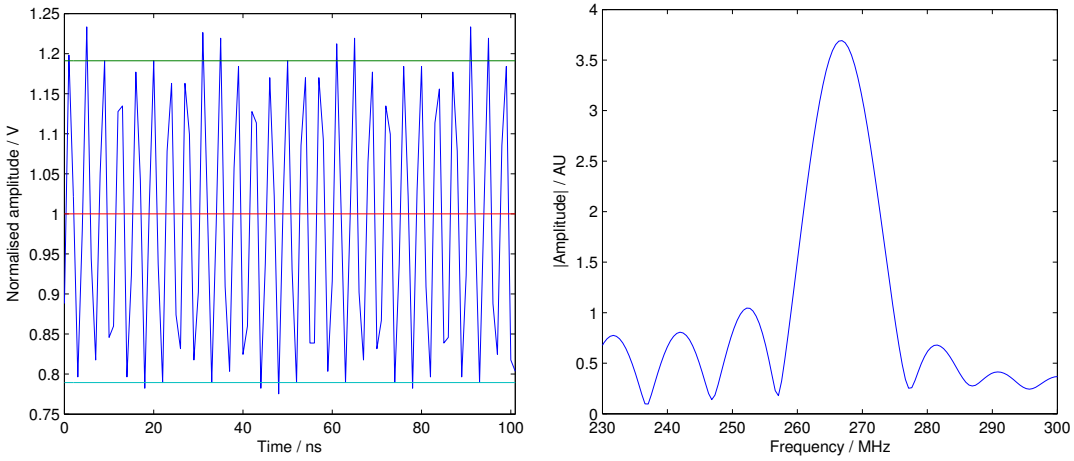
Detection

Manufactured by Laser Quantum, the detection laser is diode pumped and solid state using an Nd:YAG crystal as the lasing medium. The output beam contains a single longitudinal mode at a wavelength of 532 nm, with an output power of 500 mW. A Thorlabs fibre coupler couples the beam into a polarisation maintaining single-mode fibre at an efficiency of ~40%. After passing through the fibre and the optics detailed in Fig. 3.17, at the sample surface, with no attenuating filters, the laser power is 75 mW.

Although very stable up to 50 MHz, the detection laser has an issue with high-frequency amplitude modulation and this is shown in Fig. 3.22, as measured on an unfiltered high bandwidth amplified photodiode. This modulation issue strongly affects signal strength due to saturation of the amplifiers with the primary symptom being a ~factor 3 drop in signal amplification. A low pass filter has been applied before amplification to ameliorate this issue, however this in turn limits the maximum SRAS signal frequency to below 250 MHz.

Generation

Manufactured by Innolas, the generation laser is a Q-switched Nd:YAG with pulse lengths of 2.8 ns to 3.2 ns, a centre wavelength of 1064 nm, and repetition rates of up to 10 kHz. The beam is coupled into a 200 µm multi-mode fibre using a large diameter



(a) (b)

Figure 3.22: (a) Temporal modulation of the detection laser showing an amplitude modulation. (b) The FFT of the temporal modulation showing a strong peak at 267 MHz.

IR coated achromatic lens. To cope with the high power a fibre with an exposed core at each end is used to prevent damage from rapid thermal expansion. Pulse energies vary with repetition rate and are shown, measured after collimation from the fibre, in Fig. 3.23. An amplified version of this laser which uses a master oscillator power amplifier (MOPA) is also available and, given the issues with fibre collimation, will be used for future instruments. The MOPA is an optical amplifier which increases pulse energy but temporally distorts the pulse. The laser can be triggered either internally or externally and outputs a trigger signal each time the Q-switch is activated, releasing a pulse. The pulse temporal width is dependent on both pump-current and repetition rate and so, to maintain consistency, the pump-current is set to maximum and filters are used to control laser power.

3.6.2 KED circuit design

The KED PCB is adapted from the University of Nottingham standard design, and has been tweaked and laid out by Tim Probert, an electrical design engineer at Renishaw. Fig. 2.7 shows the circuit diagram for the KED originally designed by Prof. Matt Clarke and Dr. Steve Sharples. A built PCB and photodiode assembly is shown mounted in the custom KED case in Fig. 3.24.

In addition to the RF signal, DC signals indicating light levels for each photodiode

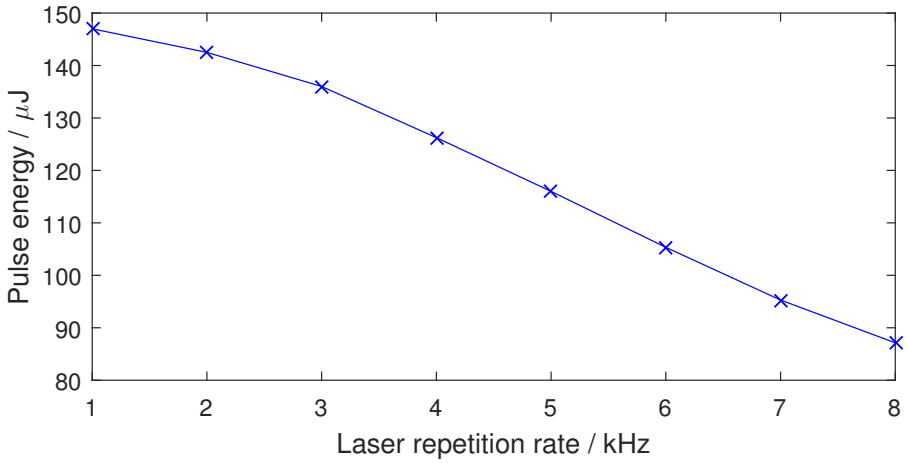


Figure 3.23: The pulse energies of the generation laser with repetition rate measured at the output of the multi-mode fibre.

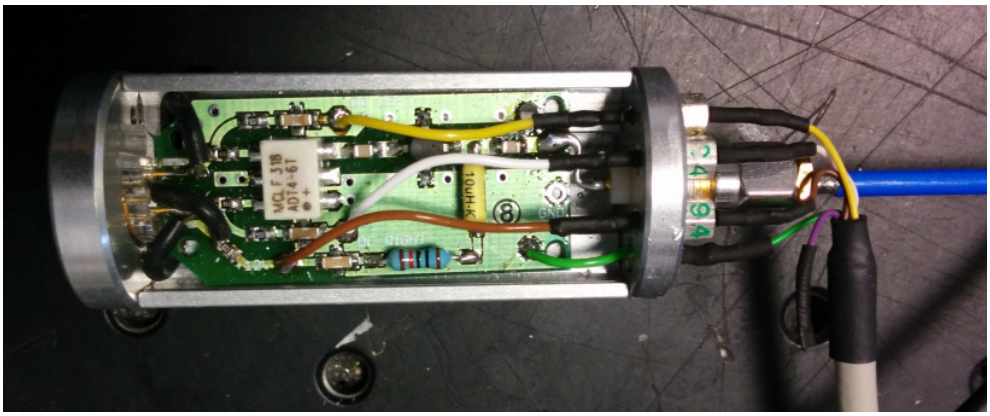


Figure 3.24: A photograph of a KED PCB mounted into the shielded case. A semi-rigid coaxial cable carries the RF signals, and standard shielded electrical cable carries the DC signals: yellow DC-left, white DC-right, brown - 12 V power and green - 0 V.

are captured, these are amplified and the sum and difference taken. A gain of ~6 dB is applied to the sum and a gain of 9 dB is applied to the difference. These are then digitised and saved during scans, they are also used during the optical alignment and instrument set-up to ensure the detector is balanced.

3.6.3 Signal chain

The instrument signal chain is shown in Fig. 3.25. A voltage of order mV is amplified internally in the KED and output into a 3 dB attenuator followed by a 250 MHz low-pass filter. The purpose of this attenuator and filter is to remove the detection laser

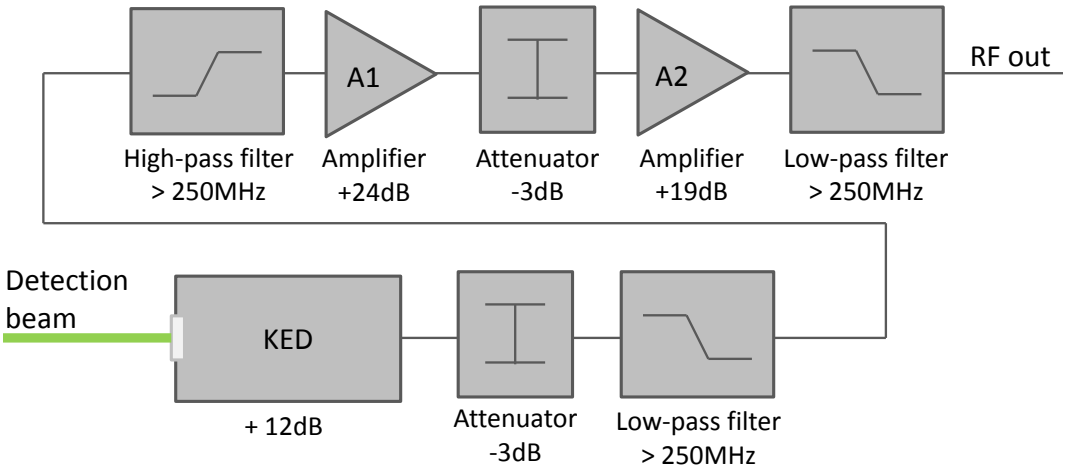


Figure 3.25: Renishaw SRAS system signal chain. An attenuator and filter remove the detection laser modulation before the signal is high-pass filtered, heavily amplified and low-pass filtered.

power modulation shown in Fig. 3.22 which would otherwise saturate the amplifiers. With the modulation suppressed a 50 MHz high-pass filter removes any DC and low frequency components from the signal before entering two low-noise, large bandwidth amplifiers. The first amplifier is a Mini Circuits ZFL-500LN+ which has a bandwidth of 0.1 MHz to 500 MHz and a noise figure of 2.9 dB. A second 3 dB attenuator removes some ringing caused by the high gain amplifier before a second amplifier is applied. The second amplifier is a Mini Circuits ZFL-500HLN+ which has a bandwidth of 10 MHz to 500 MHz and a noise figure of 3.8 dB. The final RF signal is then sent to the data acquisition system with a typical peak to peak voltage of 1.5 V. In a bid to reduce interference at all stages in the signal chain, components are either directly coupled or rigid, and semi-rigid coaxial cable used. The separation between the detector and signal chain is kept to a minimum; when mounted on the motion platform a 150 mm length of coaxial cable is used. Powering the amplifiers is a high stability, low noise 15 V dedicated power supply, DC signals are coupled over high quality Lemo connectors and shielded cable.

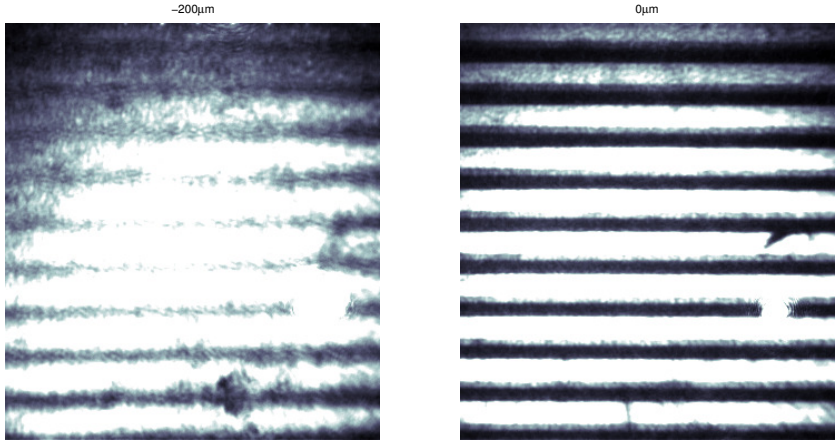
3.6.4 Data acquisition

As with the EMDA system data acquisition is still performed by a digital oscilloscope. The laser trigger output is used to trigger a trace collection on the oscilloscope at a sample rate of 2 GHz for approximately 400 samples or 200 ns. Samples are scanned in distinct lines, during each line scan the traces are buffered in the fast memory for the duration of a single line scan and are then transferred to a PC and saved via an ethernet connection. The oscilloscope can capture up to 4000 traces before the memory is filled, however the transfer times are non-linear with traces captured and so a more typical line-scan is composed of ~2500 traces. With oscilloscope buffering and PC transfer, average data acquisition rates are roughly one half as fast as the laser repetition rates.

The oscilloscope used to capture the SRAS data displayed throughout this thesis is approaching the end of its life and as a result is not entirely reliable. During some of the SRAS scans an occasional line of data is missing or slightly displaced, this is because the oscilloscope has a tendency to attempt to recalibrate itself between line captures. Since the data acquisition system is not sufficiently sophisticated to be able to automatically correct for this issue these bad data remain in the scan data.

3.6.5 Camera

A camera is used with an f100 imaging lens to view the generation and detection patches on the sample surface. The camera board is a 1.3MP large sensor monochromatic CMOS distributed by Thorlabs which has been repackaged into a low profile 3D printed custom case. A USB2.0 interface is used to download and display images on a PC and, by selecting only the region of interest in which the generation patch is displayed, frame rates of up to $100\,s^{-1}$ can be achieved. Fig. 3.26 shows captured images for a poorly focussed and well focussed generation patch. Due to the relative intensities of the lasers and spectral sensitivity of the sensor, a cold mirror filter is used to allow equal exposure of both the visible and IR beams. Even with heavy filtering, the exposure time and sensitivity of the camera are set near the minimum values.



(a) (b)

Figure 3.26: Two camera images from a sub-region of the sensor showing a poorly focussed generation patch, (a) and a well focussed patch (b). The grating spacing is 40 µm on the sample.

3.7 Integration

The statically mounted SRAS head went through a thorough optimisation and refinement process, at the end of this a finalised design was generated and the prototype head reassembled into this design. This assembly was then integrated into a motion platform to allow scanning of samples.

3.7.1 Renishaw Equator 700

Due to easy availability and the freedom of 6-axis motion, the chosen motion platform was the Renishaw Equator 700. This is an in development self-supporting hexapod with a 700 mm diameter and 700 mm tall cylindrical working volume. Full Cartesian motion in X,Y and Z is available within this volume as well as some limited rotational motion about all 3 Cartesian axes. Accelerations in excess of 1G are achievable but for the longevity of the platform, accelerations are limited to 1G. An Equator 700 is shown with the SRAS head mounted in Fig. 3.27. The Equator offers a few advantages over other motion platforms such as 3-axis gantries or stacked linear stages:

- With the ability to tilt, non-planar sample surfaces can be scanned and scan paths can be adjusted for poorly mounted samples.

- The large working volume and sturdy base allow large samples to be scanned.
- High scanning speeds and accelerations are possible.
- The central quill is hollow allowing tidy, and safe, delivery of fibre optic cables to the SRAS head.
- Since it is a Renishaw product, custom solutions for electronic and mechanical integration are readily available.
- High precision optical encoders on each strut along with sophisticated motor control allow very fine, ~1 µm, moves to be made.

For delivery of the umbilical robust cable carrier is used to protect the fibre optics and electrical cables in a chain-drag configuration up to the quill, which allows for the changing length of the hexapod struts. Control of the motion platform is realised through a software interface or a joystick. In order to mount the SRAS head a simple mechanical assembly was designed and manufactured which interfaces with the existing Equator mounting kinematics. Since the SRAS head is much larger then the typical probes mounted on the Equator a minor modification was made to shorten the length of the central quill to give as a large a working range as possible.

3.7.2 Synchronisation

For synchronisation of position and SRAS data the Equator 700 controller has been implemented as the laser control and data acquisition trigger source. There are two primary triggers available: the servo clock, which is the base clock used for all synchronisation within the controller and fires every 152 µs, and the generic controller which stays high for 152 µs every 7 servo clocks. The generic controller triggers reading of the positional encoders, such that positional data is updated at a rate of 1.064 kHz. Since the laser is capable of repetition rates up to 10 kHz flexibility was desired in setting this rate and, instead of using the generic controller pulse, the laser is triggered after a chosen integer number of servo clocks, typically 3 or 4. This synchronisation scheme can be seen in Fig. 3.28.

Constantly switching the triggers to the laser on and off is detrimental to the stability of the output power. To counter this the laser is constantly triggered and a second

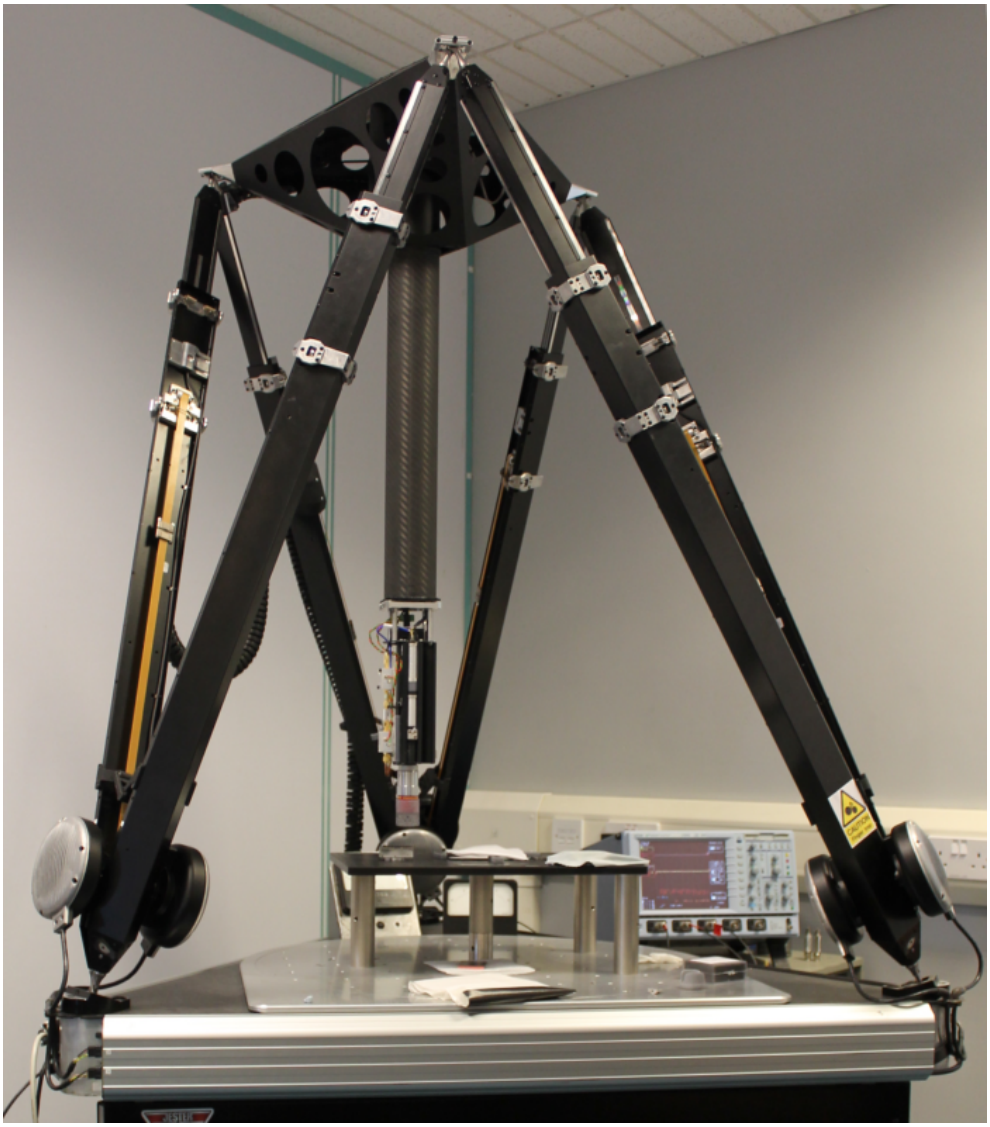


Figure 3.27: An Equator 700 motion platform with a mounted SRAS head.

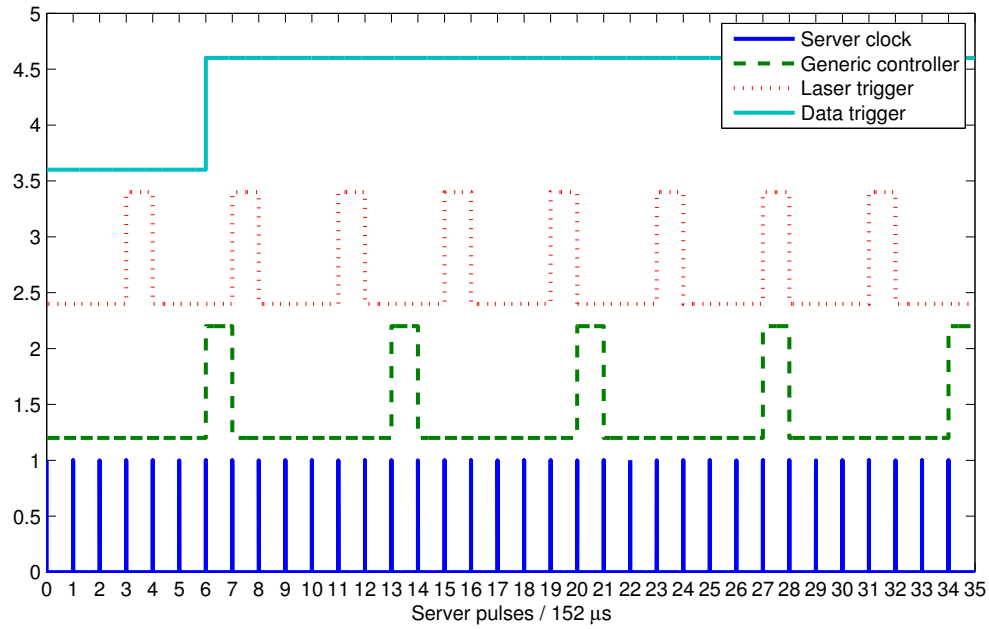


Figure 3.28: The synchronisation of the data collection and laser triggering with the Equator 700 controller.

'data trigger' signal is used to indicate when data collection is requested. This signal is raised at the start of a line scan and dropped at the end and is used with an AND gate to control the laser output trigger to the oscilloscope, as shown in Fig. 3.29. As a line scan starts the laser trigger outputs are passed through to the oscilloscope initiating the buffering of traces. The data trigger can only be raised and dropped simultaneously with the generic controller, as a result there is an uncertainty of up to 6 servo clocks in the synchronisation initialisation of data collection and the first laser trigger.

3.7.3 Software

The current software solution is unsophisticated and largely driven from MATLAB. A set of moves in X, Y and Z are defined alongside a data request variable and the set-up parameters for the oscilloscope. A small piece of software, developed by Matt Torr at Renishaw, interfaces with the Equator controller and oscilloscope, and initiates motion and data collection. The positional and SRAS data are then transferred and saved to a PC a line-scan at a time ready for processing. Further software development, which will allow more complex 3D motion and path generation, has been delegated to the

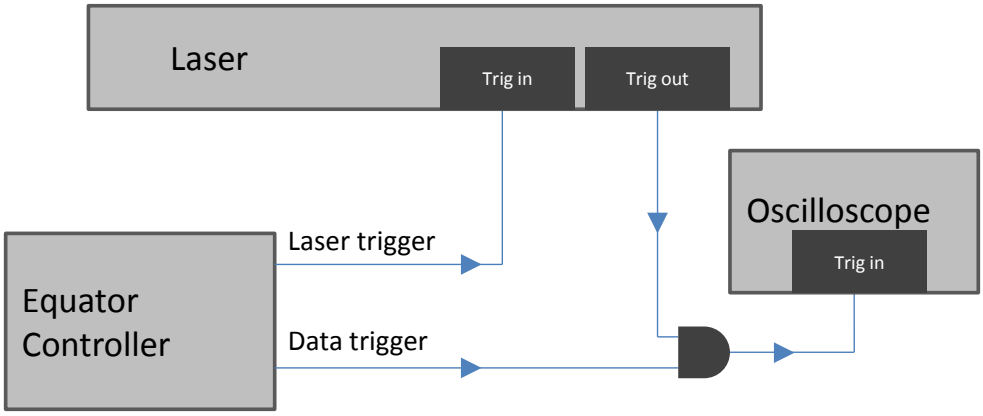


Figure 3.29: The control path for triggering the laser and oscilloscope.

Renishaw software division.

3.8 Conclusions

This chapter describes the advances made in the instrumentation used to implement the SRAS technique and developments made towards an industrial prototype. In order to meet the requirements of an easy to assemble and low cost SRAS head, a novel optical constraint system scheme has been devised and implemented. This rigid chassis based system has multiple advantages including: low weight, rigidity, customisability and compactness and compares favourably to alternative, off-the-shelf optomechanical systems.

The optical paths for implementation of SRAS have been shown and optical housings designed to realise these paths. In combination with the rigid chassis these housings allow excellent passive beam centring and alignment such that final alignment of beams is straight-forward and trivial. A SRAS head assembly has been constructed and integrated with a motion platform allowing static samples to be scanned.

Some initial work has been done in the ancillary aspects of the SRAS instrumentation, such as signal chain, data acquisition and software but a considerable body of work remains in developing these areas before they can be considered production ready. Fortunately the majority of this work does not require expertise in SRAS or optics, and can be delegated to the relevant departments at Renishaw.

3.9 Further work

3.9.1 Rough surface detection

One of the fundamental limitations of the current Renishaw SRAS system is the sensitivity to rough surfaces. Currently samples need to be prepared and polished to a surface roughness of less than 1 µm R_a, which is approaching a mirror finish. As surface roughness increases the reflected detection beam becomes less specular, is scattered over a greater range of angles, and starts to form speckle patterns on the KED. All of these effects act to reduce the sensitivity of the detector.

There are multiple alternative optical detection techniques which are less sensitive to rough surfaces which were discussed in §1.2.4. Of these two-wave mixing and the Confocal Fabry-Pérot Interferometer are the most promising. TWM is currently not commercially available as a fibre-coupled detection technique and puts severe limitations on the scanning speed. The CFPI is available in a fibre-coupled form and can maintain the scanning speeds currently used with the KED. However at ~£80,000 the cost of this technology is prohibitive for Renishaw.

There is a technology in development at the University of Nottingham called the speckle knife-edge detector (SKED). This device comprises of an array of photo-sensitive pixels which can configure themselves to a speckle pattern imaged onto the sensor. In the configured state the sensor treats each speckle as an individual laser spot with the local pixels set up as individual segmented photodiodes [74]. The SKED can be fabricated for costs similar to the existing KED and can be used with the existing detection laser. Initial trials of the SKED have shown it is capable of detecting SAWs up to frequencies of ~80 MHz. More development is required before the full-range of SAW frequencies used for SRAS can be achieved.

3.9.2 Optical efficiency

Due to the poor collimation possible on the generation beam, the final power density at the sample surface is reduced from the desired levels. The result of this is a decrease in SAW amplitude, and thus a decrease in SNR. There are two potential paths for reclaiming some of the lost energy density. The first is the addition of a MOPA to the

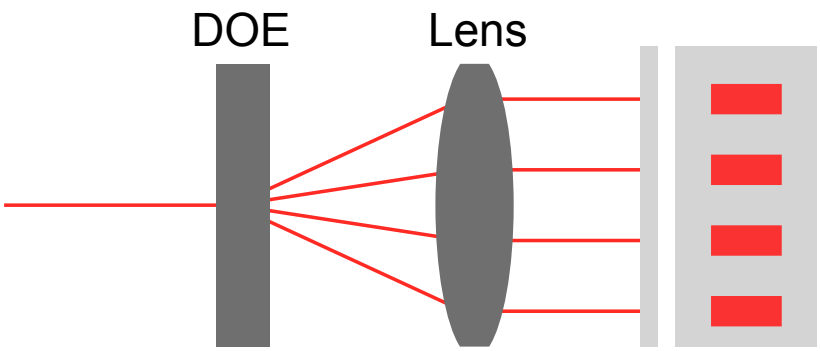


Figure 3.30: A schematic illustrating the application of a DOE to create the generation patch image. This is not to scale.

generation laser, which would double the energy per pulse. However the pulse temporal length is also increased, limiting the frequency range of SAW generation. The second option is to use more sophisticated optics for creating the generation grating image; the currently used chrome mask is effective but effectively discards half of the laser energy. One promising option is diffractive optical elements (DOE) which are fabricated from a computer generated hologram to produce nearly any desired wavefront modification [75]. For application to SRAS a fairly simple beam splitter and line generator is required, the effective schematic is shown in Fig. 3.30. In this configuration the laser energy can be concentrated into only the area which contributes to the detected SAW, improving energy density. In addition a typical DOE efficiency is ~75% meaning much less power is wasted than by reflection from the chrome mask.

For the existing system the issue of low SNR is tackled by the application of more sophisticated techniques for determining peak velocity from SRAS signals. This work is detailed in Chapter 4.

3.9.3 Increased automation

Part of the design philosophy for the SRAS head is a set-and-forget functionality, i.e. the optics need only be calibrated and aligned once. Unfortunately there are two functional requirements which both demand reconfiguration of the head: changing masks or apertures, and altering propagation direction. The cartridge system means that changing masks and apertures is fairly trivial, however when changing between very different generation patch sizes the distance between the generation and detection

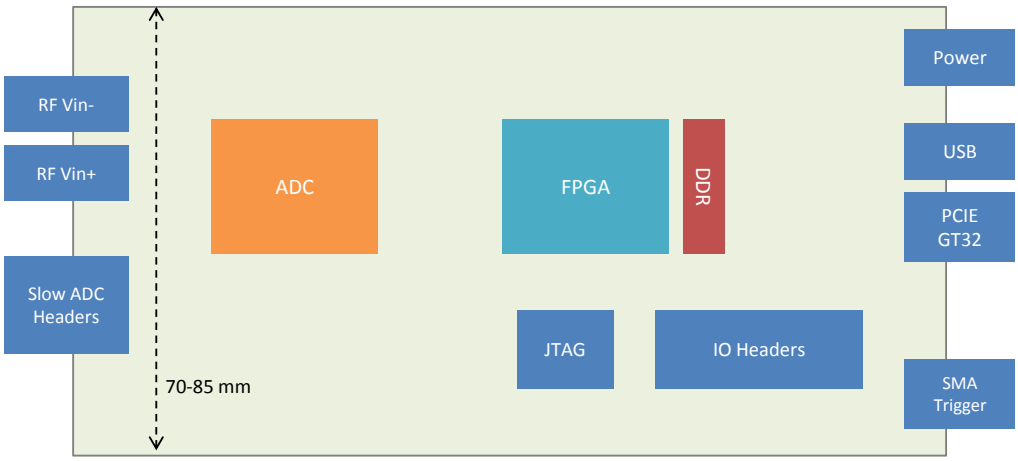


Figure 3.31: A rough layout for a custom data acquisition board to be used for the production SRAS instrument.

patches requires optimisation. Currently this is done by manually altering a screw on the adjustable mirror for the generation patch while the lasers are active. To remove this procedure a clip-on motorised actuator is in development which will allow remote manipulation of the generation patch. With this addition the SRAS unit can be fully enclosed, and laser safe, during configuration.

For the EMDA system altering propagation direction is a process which requires reconfiguration of both the generation grating and the detection beam alignment. This can only be done with the lasers active and requires manual manipulation of the optics. To avoid the requirement for any reconfiguration of the Renishaw SRAS system a motorised rotary stage is under development which will locate within the Equator 700 quill and rotate the entire SRAS chassis.

3.9.4 Data acquisition

As discussed in §3.6.4 the currently used oscilloscope is less than ideal for data acquisition and needs review. An oscilloscope is not the production solution for the Renishaw SRAS system: it is too slow and too expensive. A custom design is currently under review which will enable live, and continuous, streaming of data to a PC, a rough layout of this board is shown in Fig. 3.31.

This board is specified with a 12 bit ADC capable of sample rates up to 4 GHz. Al-

though the specifications of USB2 indicate sufficient bandwidth for this application, in practice these are rarely achieved and so a more reliable and faster PCI-e interface will be used.

Chapter 4

Signal analysis

Obtaining phase velocity from a signal time trace and a known acoustic wavelength is one of the most critical areas of the SRAS technique. In order to deploy SRAS industrially as a robust and reliable instrument this process needs to be examined and refined. This chapter is aimed at investigating and applying a range of signal analysis techniques in order to determine which is most suitable for SRAS data, and also to determine the minimum SNR from which reliable data can be obtained. The techniques examined include several varieties of wavelet analysis, short time Fourier transforms (STFT), Gaussian fitting and use of simulated signals to fit data. All of these techniques are designed to find the dominant peak in the frequency spectrum of a measured trace. Then by applying the wave equation, $v = f\lambda_a$, the dominant phase velocity is obtained.

With the Renishaw SRAS system signal analysis is critical: fibre-coupling of the generation laser through a multi-mode fibre reduces the energy density in the generation patch by ~40%. As a result the SAW amplitude is decreased and the measured signal is decreased. Some of this drop in SNR has been recovered in more effective shielding and grounding of the detector and lasers, however significant incoherent noise from the photodiodes and electronics remains.

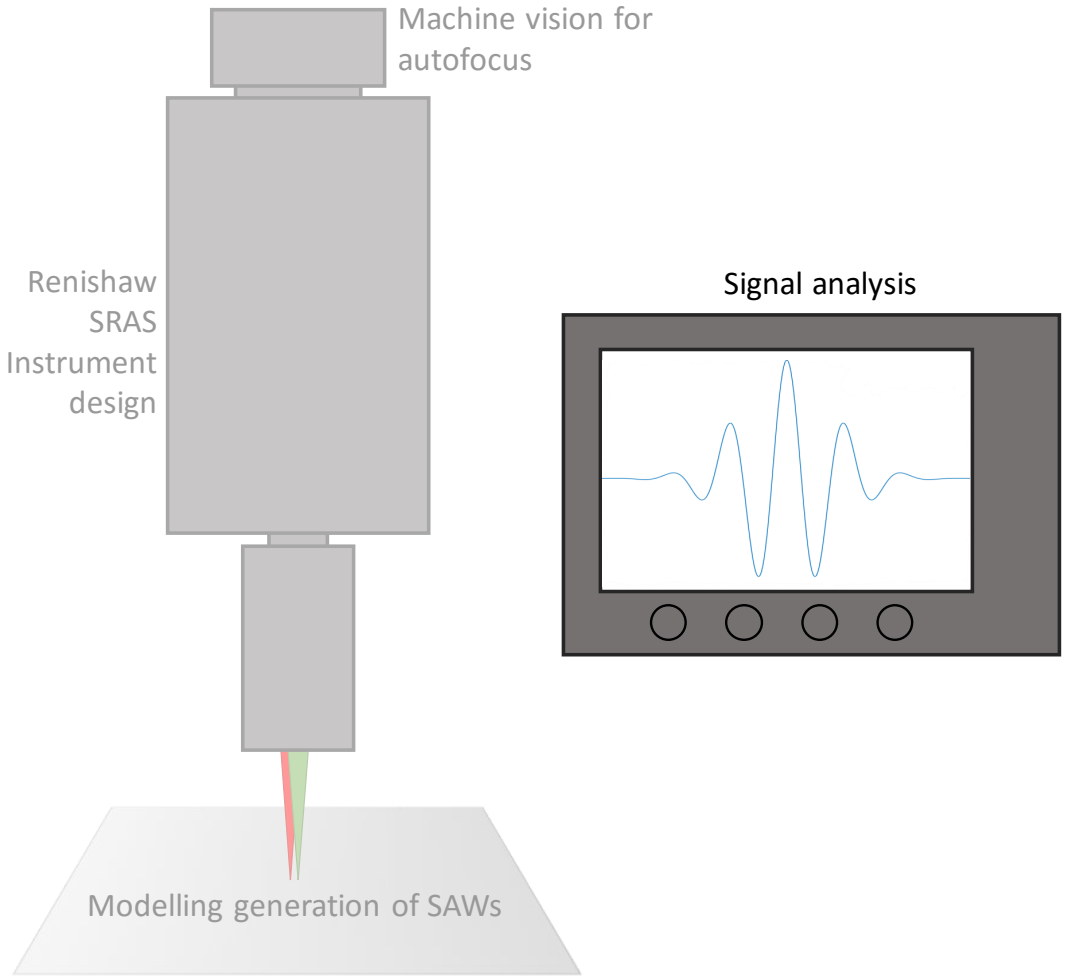


Figure 4.1: An illustration of the primary areas of research and engineering covered in this thesis. In this chapter signal processing is investigated.

4.1 Current signal processing technique issues

Early implementations of the SRAS technique collected data at ~10 points per second and used an iterative, least mean squares, Gaussian fitting technique to extract the dominant frequency peak. However with data acquisition rates increasing to ~2000 points per second with subsequent implementations, this fitting technique became too slow and instead a more computationally simple technique was adopted. Zero-padding and taking the FFT of a SRAS signal gives an interpolated frequency spectrum. Simply taking the peak of this spectrum within the frequency range of interest gives a reasonable estimate of the dominant frequency peak. This technique requires very little computation and processing is far quicker than data collection. The issue with this technique is shown in Fig. 2.11 where a simulated SRAS signal with a SNR of 8 dB has been padded to a length of 2048 and an FFT has been taken. A peak frequency of 129 MHz was found, however the actual frequency of the signal is 124 MHz, giving a measurement error of 5 MHz. This occurs due to the distribution of noise distorting the shape of the Gaussian peak in the frequency spectrum. When this technique is applied to a scan with low SNR, the result is a noisy image, an example of which is shown in Fig. 4.2 where the sample has been etched before scanning. Although the grain structure is fairly clear, previous scans on the EMDA system have shown that this sample should be fairly consistent in velocity across each grain, indicating that these variations are due to noise.

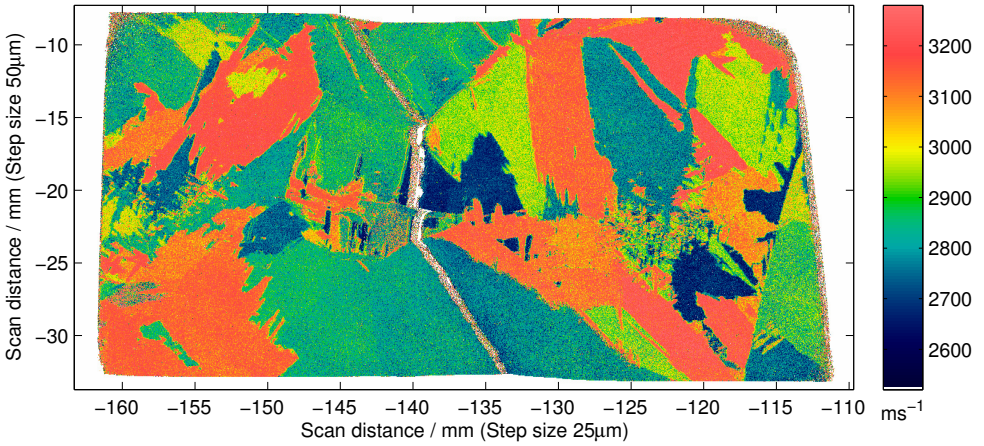


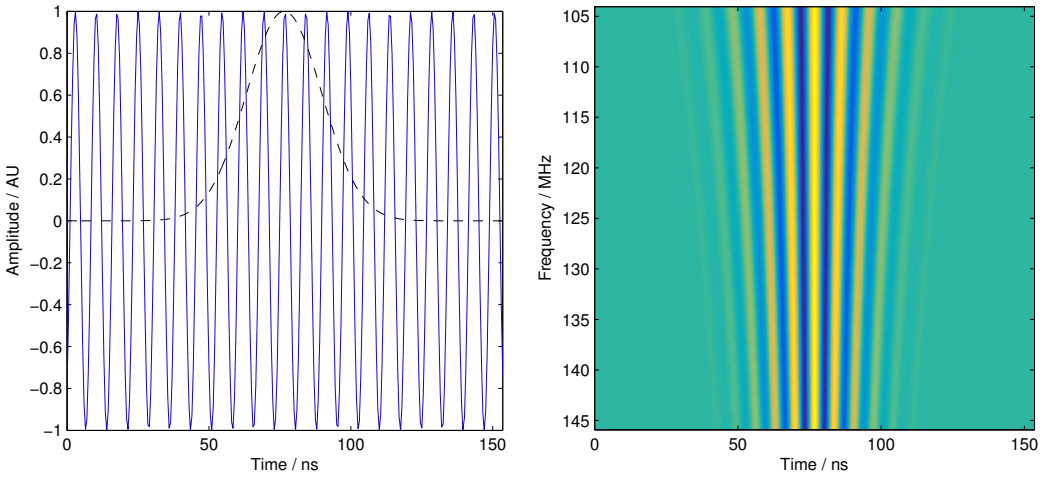
Figure 4.2: A SRAS scan taken on the Renishaw prototype of a cracked sample of Ti 685. The sample has been etched resulting in low signal strength, as a result the SRAS data shows significant noise across homogeneous grains.

4.2 Modelling SRAS signals

A SRAS signal can be simulated by the convolution of a confined Gaussian, or Tukey, window and a sinusoid, or sum of sinusoids, at a given frequency, or set of frequencies, i.e.

$$s_c(t, f) = \cos(2\pi f t) \times e^{-(\frac{\pi f t}{c})^2}, \tag{4.1}$$

where s_c is the SRAS signal, f is the fundamental frequency and c is the desired number of cycles in the signal. The factor π in the numerator gives a constrained Gaussian window of the correct width for the desired number of cycles. Fig. 4.3 (a) shows the construction of a single frequency SRAS signal while (b) shows a set of simulated signals over a frequency range of 105 MHz to 145 MHz with a cycle number of 7. To match the data acquisition a sample rate of 2 GHz has been used.

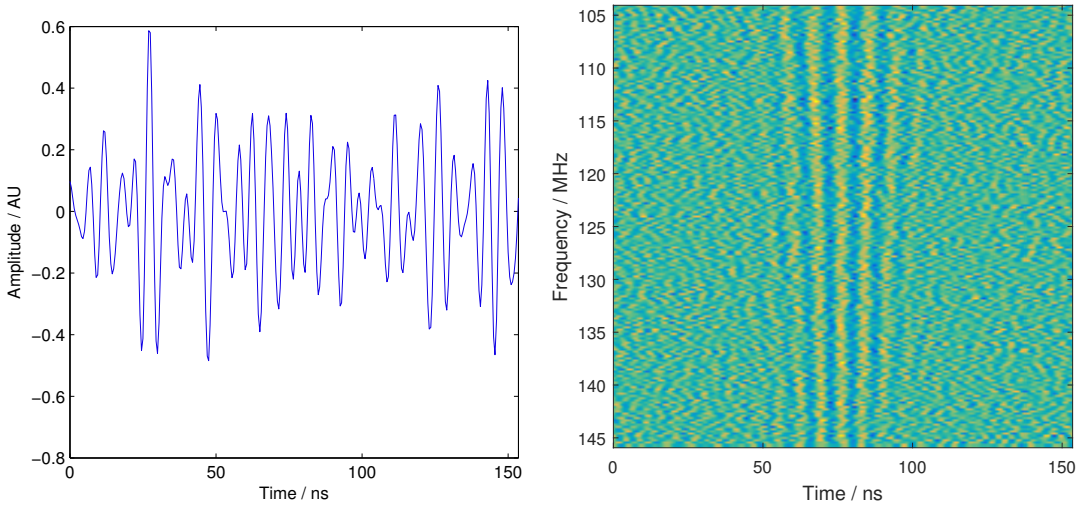


(a) (b)

Figure 4.3: a) A single frequency sinusoid (solid) and a Gaussian window (dashed). (b) A set of simulated SRAS signals over a frequency range of 105 MHz to 145 MHz.

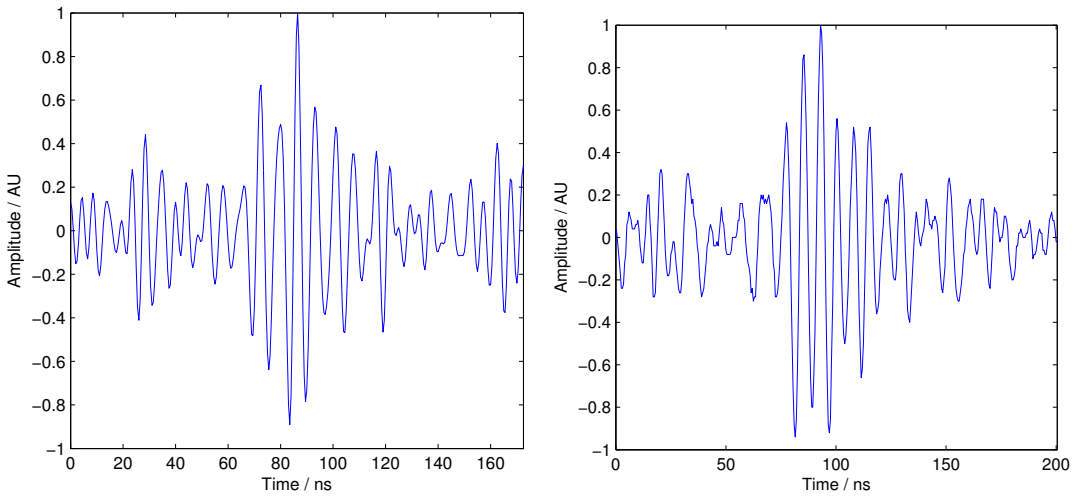
Incoherent noise is then simulated using a normally distributed set of pseudo-random numbers. This series of numbers is Fourier transformed and bandpass filtered before an inverse Fourier transform is taken. To match the filtering used in the actual instrument data collection, the bandpass is 50 MHz to 200 MHz. A typical noise trace is shown in Fig. 4.4 (a) with a maximum amplitude of 0.5 compared to the reference signal. This noise is scaled in amplitude and summed with the simulated SRAS signal, the result

is shown in Fig. 4.4 (b) where a 0 dB SNR set of signals has been simulated. For each signal in the frequency set a unique noise trace is simulated and filtered, by altering the scaling of the noise any desired SNR signal can be generated. Fig. 4.5 shows a comparison of a simulated signal with 8 cycles and a SNR of 0.6 dB with a measured SRAS signal. The two signals show good agreement and validate the signal simulation for use with modelling signal processing techniques.



(a) (b)

Figure 4.4: a) A typical noise trace which is bandpass filtered in the range 50 MHz to 200 MHz. (b) The set of simulated SRAS signals shown in Fig. 4.3b) with filtered noise added to give a SNR of 0 dB.



(a) (b)

Figure 4.5: a) Simulated SRAS signal of SNR 0.2 dB. (b) Experimentally measured SRAS signal.

4.3 SRAS data analysis techniques

The following section details the theory, and method of application, of a range of signal processing techniques. While there are a vast number of possible frequency, and time-frequency analysis techniques available, a reasonably small section of these is reviewed here. An initial pre-screening was performed of a larger range of techniques assessing their suitability for transient, typically single frequency, SRAS signals, with the best suited techniques presented herein.

4.3.1 Maximum frequency peak (MFP)

The fastest and easiest analysis method to implement is the maximum frequency peak (MFP) method. Here, an FFT of the SAW trace is taken and the maximal value of the magnitude of the frequency spectrum is found in a given frequency range of interest. This is the current standard technique used in SRAS signal processing. In practice this technique is applied to a full set of signals captured during a line scan simultaneously.

4.3.2 Gaussian fitting

Gaussian fitting is performed in the frequency domain and there are two methods used.

Iterative LMS fit

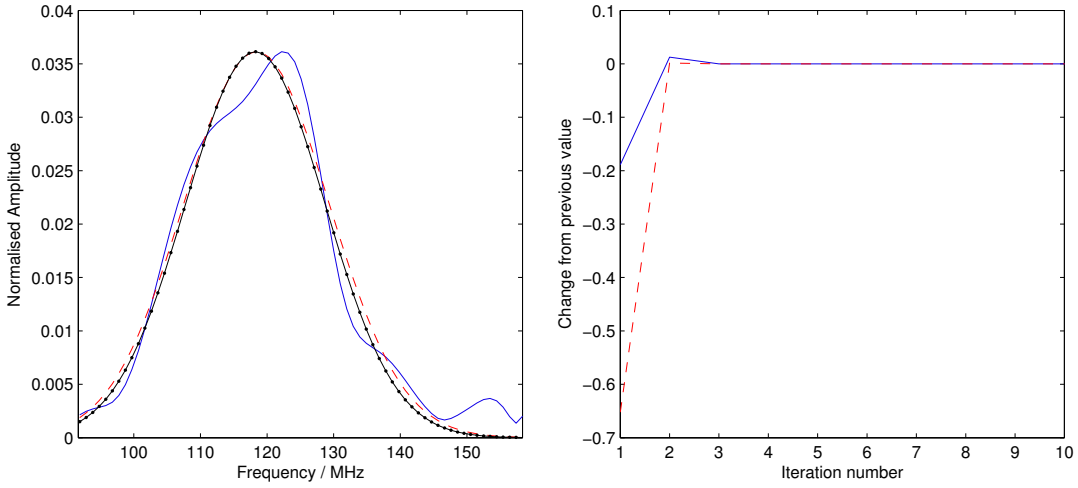
The first option is an iterative non-linear least mean square fit. Firstly the trace is padded, an FFT is taken and the frequency spectrum is sliced to the range of interest. A Gaussian probability density function, $\psi_{\mu,\sigma}$ is assumed of the form

$$\psi_{\mu,\sigma}(t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{t-\mu}{\sigma})^2}, \tag{4.2}$$

where σ is the Gaussian width and μ the mean value. An initial guess of μ_0 is taken to be the maximum frequency peak of the signal and $\sigma_0 = \frac{1}{\sqrt{2\pi}\ \max(s(t))}$. An iterative least squares fit is then applied for up to 25 iterations, or until the changes between subsequent μ and σ values are small enough to consider the fit to have converged, here

a change of less than 0.1% in either value is used. The fitting algorithm is complex and is described in detail in reference [76].

Fig. 4.6 shows a fitted simulated SRAS signal frequency trace with the ideal frequency spectrum overlaid. The ideal signal frequency is 118.8 MHz whilst the measured dominant frequency peaks are 119.3 MHz from an iterative LMS fit and 116.3 MHz from taking the maximum frequency peak.



(a) (b)

Figure 4.6: a) The frequency spectrum of a noisy simulated SRAS signal (solid), ideal SRAS signal (dash-dot) and the fitted Gaussian curve (dashed). (b) The convergence of μ (solid) and σ (dashed) over multiple iterations. With a good initial guess of μ_0 and σ_0 the fitting algorithm converges very quickly and would finish after 4 iterations.

Gaussian convolution

Unlike many other signal processing requirements, for SRAS the form of the signal is already well-known and varies over a small frequency range - dependent on the material. As such it is possible to use this information to assume a certain signal form which can be convoluted with the measured signal in order to extract the dominant phase velocity. To formalise this process mathematically, the operation for finding $S_c(f_s)$, the summed correlation curve in chosen frequency steps, f_s, is

$$S_c(f_s) = \sum_{\omega=0}^{\Omega} \left[\int_{-\infty}^{\infty} e^{-i\omega t} s_c(t, f_s)\ \mathrm{d}t * \int_{-\infty}^{\infty} e^{-i\omega t} x(t)\ \mathrm{d}t \right], \quad f_s \subseteq \Omega, \tag{4.3}$$

where ω is the frequency of the signal to be fitted, Ω is the full frequency domain, s_c is given by equation 4.1 for a particular number of cycles and $x(t)$ is the measured signal. In practice this process is applied through the following steps.

1. A set of simulated signals with a given cycle number, over the expected frequency range are generated, such a set is shown in Fig. 4.7 (a) for the range 110 MHz to 150 MHz with 100 steps in frequency. Traces are generated using equation 4.1 with a Gaussian envelope width chosen to give 9 cycles. The choice of cycle number is dependent on the desired sensitivity and resistance to noise. A long signal will yield a sharp frequency peak which will give better resolution in frequency but will be sensitive to noise. A short signal will give a broad frequency peak which is fairly insensitive to noise but gives poor frequency resolution.

2. These traces are then Fourier transformed along the time axis, the absolute value of this is shown in Fig. 4.7 (b).

3. Processing is performed on one trace at a time, a single trace is captured and an FFT taken, shown in Fig. 4.7 (c), this is windowed to the frequency range of interest. The sample rate and zero-pad for both the simulated and real signals are kept the same.

4. For each of the 100 frequency steps, the complex signal frequency trace and simulated frequency traces are multiplied, and the inverse Fourier transform taken.

5. After taking the absolute of this the result is a peak in frequency-time space and is shown in Fig. 4.7 (d). A dominant frequency peak is then found by summing across time for each frequency step, and finding the maximum value.

Fig. 4.8 shows $S_6(f_s)$ and $S_{12}(f_s)$ for a 9 cycle signal. The measured frequency peaks are: ideal 107.8 MHz, MFP 105.5 MHz, $S_6(f_s)$ 106.5 MHz, $S_{12}(f_s)$ 107.8 MHz.

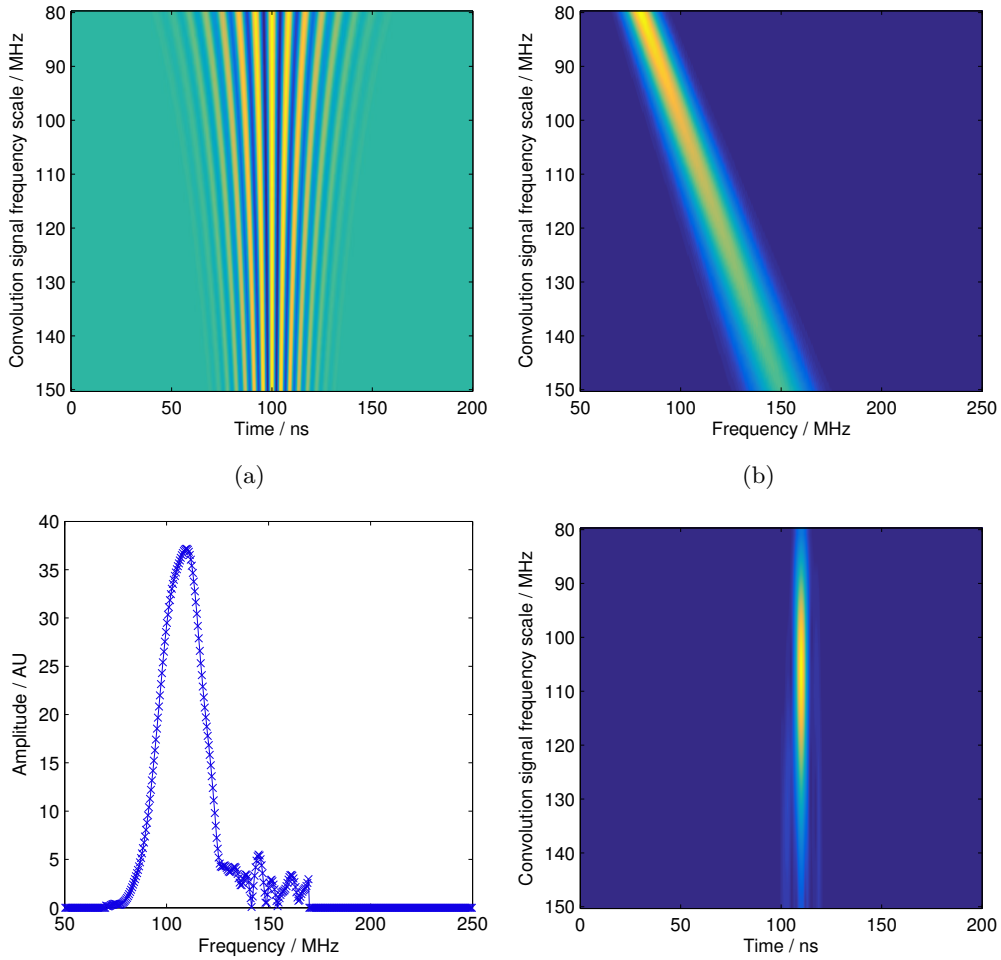


Figure 4.7: a) A set of simulated signals with 9 cycles and a frequency range of 110 MHz to 150 MHz. (b) the Fourier transform of the traces in (a). (c) The Fourier transform of the measured signal. (d) The IFFT of the multiplication of the measured and simulated frequency spectra.

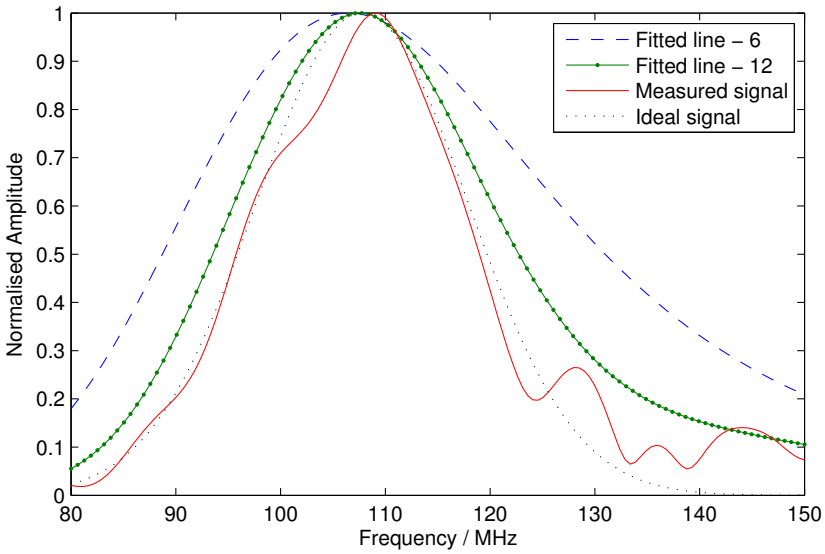


Figure 4.8: A simulated 9 cycle noisy frequency spectrum plotted with the ideal frequency spectrum (solid red and black dotted respectively). Two sets of simulated signals have been fitted to this data, one with 6 cycles (dashed blue) and the other with 12 cycles (dot-dash green).

4.3.3 Short time Fourier transform

A short time Fourier transform (STFT) applies a window function to the signal before calculating the FFT. The premise is to isolate the signal region of interest in time and reduce the influence of noise. Typically these are used for signals with continuously varying frequency, for example musical tracks, but they should also work with SRAS signals due to the transient nature of the SAW packet in time. Formally, the STFT can be applied with discrete window delays, $t - \tau$ using

$$S(\omega) = \sum_{\tau=0}^{T} \int_{-\infty}^{\infty} x(t)w(t-\tau)e^{-i\omega t}\,\mathrm{d}t, \tag{4.4}$$

to give a spectrum in frequency, $S(\omega)$. Here $w(t)$ is the chosen window and T is the length of the window in time. There is a vast array of windows to choose from and the choice for these is dependent on the required shape and properties of the window in the frequency domain [77]. A thorough discussion of the benefits and drawbacks of using a variety of window functions can be found in [77]. The properties involved in the choice of a window are highest side-lobe level, gain and side-lobe drop off. Three windows have been chosen based on compromising in these areas: Hann, Tukey and confined Gaussian.

Confined Gaussian windows can be generated using equation 4.2 setting $\mu = \frac{T}{2}$ and $\sigma = 0.1T$, the amplitude term before the exponential is neglected. This window gives a narrow frequency width with rapid drop-off and a relatively small gain [78]. The STFT with a Gaussian function is also known as a Gabor transform.

The Hann, or Hanning, window is a compromise between the three discussed factors having middling values for side-lobe level, gain and side-lobe drop off. Hann windows are based on the cosine and can be generated using [79]

$$w(t) = \frac{1}{2}\Big(1 - \cos\Big(\frac{2\pi t}{T-1}\Big)\Big). \tag{4.5}$$

Tukey windows are effectively the opposite extreme to the Gaussian, they have a slow drop-off and large side lobes but have a very high gain. Like Hann windows, Tukey windows are also based on a cosine but are convolved with a rectangular window. They can be generated using [79]

$$w(t) = \begin{cases} \frac{1}{2}\left[1 + \cos\left(\pi\left(\frac{2t}{\alpha(T-1)} - 1\right)\right)\right] & \text{if } 0 \leqslant t \leqslant \frac{\alpha(T-1)}{2} \\ 1 & \text{if } \frac{\alpha(T-1)}{2} < n \leqslant (T-1)(1-\frac{\alpha}{2}) \\ \frac{1}{2}\left[1 + \cos\left(\pi\left(\frac{2t}{\alpha(T-1)} - \frac{2}{\alpha} + 1\right)\right)\right] & \text{if } (T-1)(1-\frac{\alpha}{2}) < n \leqslant (T-1). \end{cases} \tag{4.6}$$

Here α determines how rectangular the window is, from $\alpha = 0$ for a rectangular window to $\alpha = 1$ for a Hann window. Fig. 4.9 shows each of these three windows with various delays applied.

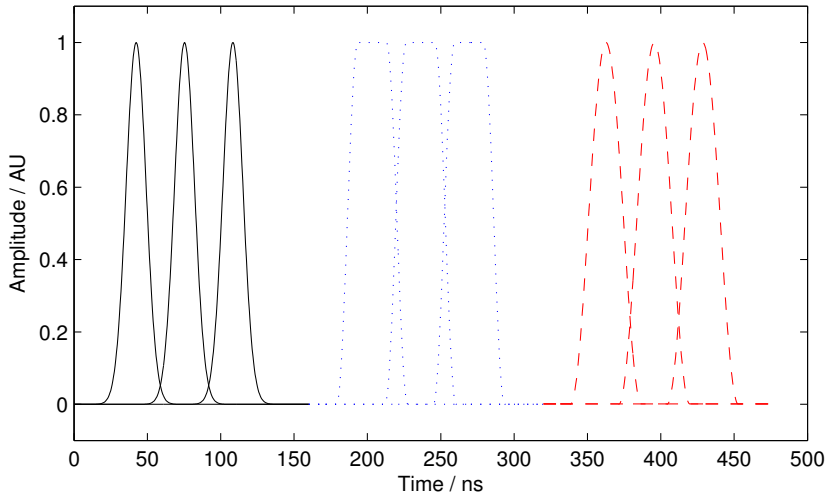


Figure 4.9: Constrained Gaussian windows of width 48 ns (solid black), Tukey windows of width 48 ns and $\alpha = 0.6$ (dotted blue), Hann windows of width 48 ns (dashed red).

In order to apply the STFT to a SRAS signal the following process is followed:

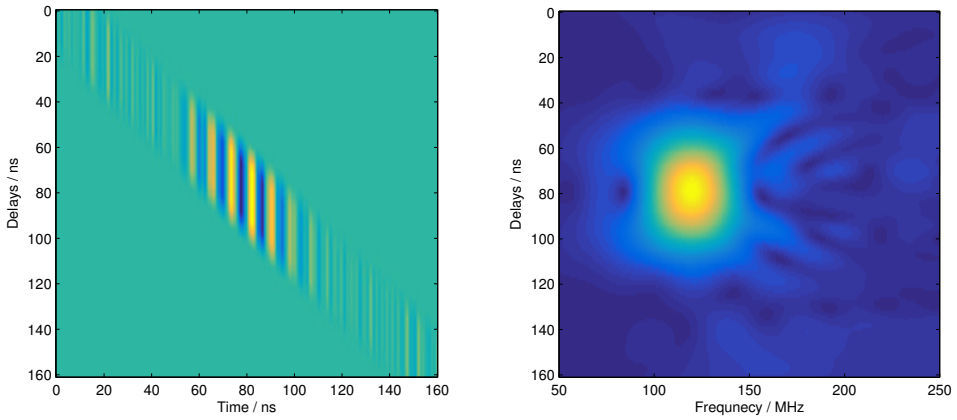
1. Choose window type, window width and number of delay steps in time. The choice of delay step size will increase time resolution but also increase the computation required. Window width should be approximately equal to the SRAS signal length, which can either be measured or simulated. The maximum SRAS signal length occurs at the lowest frequency.

2. Generate a set of windows with various delays, τ, this can be done using delay laws or, in the case of Tukey windows, by shifting a single window in time. When implementing these windows in MATLAB the $(T-1)$ terms become T due to the peculiarities of indexing.

3. Multiply the time trace of the signal by each window, the result of this is shown in Fig. 4.10 (a).

4. Take the FFT of the windowed signals to obtain a plot in delay-frequency space, shown in Fig. 4.10 (b).

5. Summing across each delay gives $S(\omega)$, shown in Fig. 4.11 for a Tukey and Hann window. The measured peak frequencies in Fig. 4.11 are: ideal - 119.8 MHz, Gaussian (not plotted) - 120.2 MHz, Tukey - 120.7 MHz, Hann - 120.7 MHz and MFP - 123.1 MHz.



(a) (b)

Figure 4.10: a) A single simulated SRAS signal with a set of Tukey windows applied over a range of delays. The window width is 48 ns and $\alpha = 0.6$. (b) The magnitude of the FFT of the windowed SRAS signals.

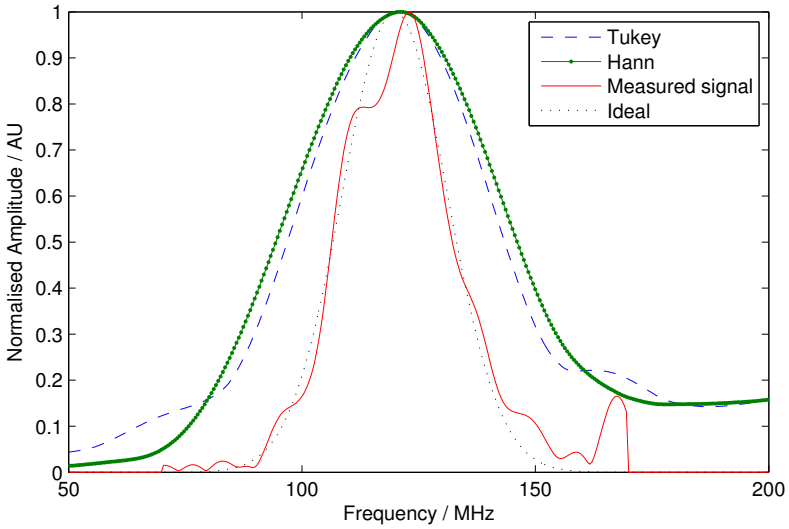


Figure 4.11: A simulated 9 cycle noisy frequency spectrum plotted with the ideal frequency spectrum (solid red and black dotted respectively). Both Tukey (dashed blue) and Hann (dot-dash green) STFTs have been applied and both show good agreement with the ideal signal. The effect of the difference in window shape can be seen in the frequency widths; the broader Tukey window has a narrower frequency peak than the Hann window.

4.3.4 Wavelets

Wavelet analysis is a relatively young field of signal processing, only gaining traction and widespread use in the 1970's [80]. Like Fourier transforms, wavelet transforms are used to split a signal into a set of single frequency components. The strength of wavelet analysis over Fourier transforms is in analysing transient and non-periodic signals [81]. There exist a great many types and families of wavelet, and a number of ways of representing wavelets. Many of these are orthogonal and are used in more sophisticated orthogonal wavelet analyses which rely on decomposition of signals into a Fourier series [82, 83]. Here wavelets are defined as orthogonal if the inner product of each mutual Fourier component is zero. Orthogonal wavelets limit the choice of frequency values to a discrete set whilst non-orthogonal wavelets can be generated for an arbitrary set of frequencies [84]. For SRAS signal analysis fine steps in frequency are required therefore only non-orthogonal wavelets are suitable. The chosen wavelets are listed below, these have been selected for their range over frequency and temporal widths and lack of discontinuities.

Morlet

The Morlet wavelet is quite similar in form to a standard SRAS signal; it is composed of a complex exponential windowed by a Gaussian and can be generated as [81]

$$\Psi_\sigma(t) = \frac{\left(e^{i\sigma t} - e^{-\frac{1}{2}\sigma^2}\right)}{\pi^{\frac{1}{4}} e^{\frac{1}{2}t^2} \sqrt{1 + e^{-\sigma^2} - 2e^{-\frac{3}{4}\sigma^2}}}, \tag{4.7}$$

where t is time and σ is the wavelet order which determines the temporal extent of the wavelength. Fig. 4.12 shows Morlet wavelets of orders 4 and 8. Compared to other wavelets of the same order, the Morlet wavelet has good frequency resolution and fairly poor temporal resolution.

Marr

The Marr wavelet, also known as the Mexican hat wavelet, is the real valued second derivative of a Gaussian function. A real valued wavelet is less successful at capturing

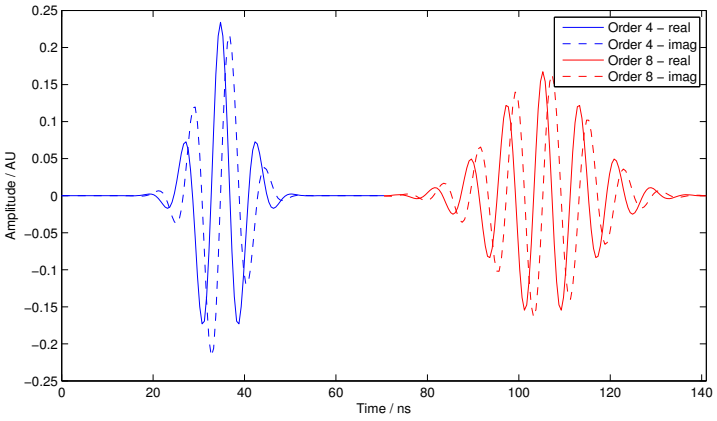


Figure 4.12: Morlet wavelets of order 4 (blue) and 8 (red). Real components (solid) and imaginary components (dashed) are overlaid.

oscillatory behaviour in the signal and is generally more useful for finding peaks in amplitude, as well as discontinuities [82]. Compared to other wavelets of the same order, the Marr wavelet has reasonable frequency and temporal resolution.

$$\Psi_\sigma(t) = \frac{1}{\sqrt{2\pi\sigma^3}}\left(1 - \frac{t^2}{\sigma^2}\right)e^{\frac{-t^2}{2\sigma^2}}. \tag{4.8}$$

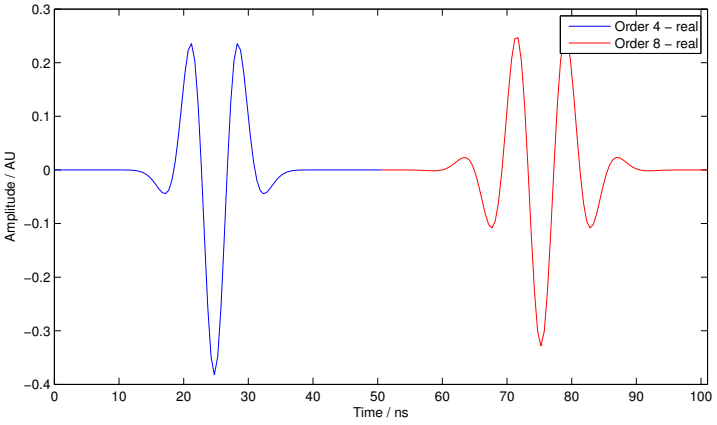


Figure 4.13: Marr wavelets of order 4 (blue) and 8 (red). Marr wavelets are purely real.

Paul

The Paul wavelet is complex valued and are typically used in areas where phase-sensitivity is required [85]. Compared to other wavelets of the same order, the Paul wavelet has poor frequency resolution and good temporal resolution. Paul wavelets can be generated as such [82]

$$\Psi_\sigma(t) = \frac{2^\sigma i^\sigma \sigma!}{\sqrt{\pi(2\sigma)!}}(1-it)^{-(\sigma+1)}. \tag{4.9}$$

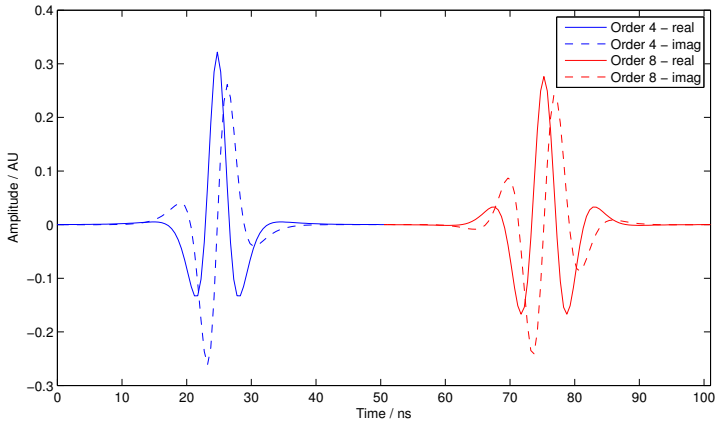


Figure 4.14: Paul wavelets of order 4 (blue) and 8 (red). Real components (solid) and imaginary components (dashed) are overlaid.

B-spline

Spline wavelets are unlike the wavelets already discussed above, and are a much more recent discovery [86]. These wavelets are constructed from spline functions which are curves described piece-wise using polynomial functions. There exist a huge variety of spline wavelets with a vast array of properties and uses. Most widely used are compactly supported B-spline wavelets since they achieve almost optimal localisation in frequency time, and can be constructed to suit the application required [87, 88]. The full closed form expression used to generate the third and fourth order B-spline wavelets can be found in Appendix B. The B-spline wavelets shown in Fig. 4.15 are complex valued, smooth and non-orthogonal.

Application of wavelet analysis

Wavelet analysis can be applied in numerous ways such as: continuous transforms, discrete transforms and wavelet decomposition. For the analysis of continuous SRAS signals, only continuous wavelet transforms are applied. The process is:

1. Choose wavelet daughter and desired frequency resolution, typically defined as the scale factor. The frequency resolution choice is a trade-off between resolution

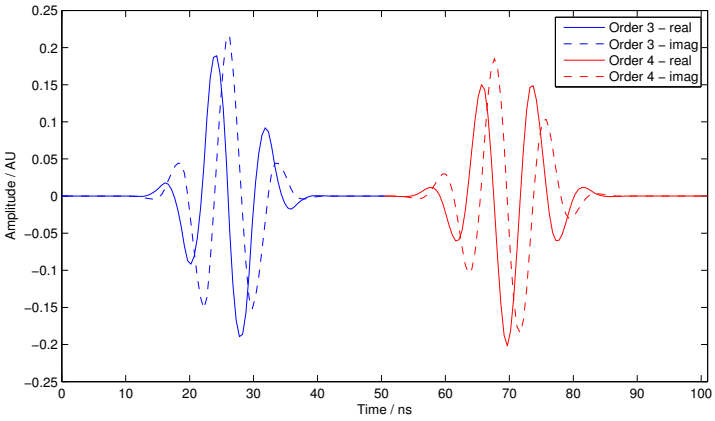


Figure 4.15: B-spline wavelets of order 4 (blue) and 8 (red). Real components (solid) and imaginary components (dashed) are overlaid.

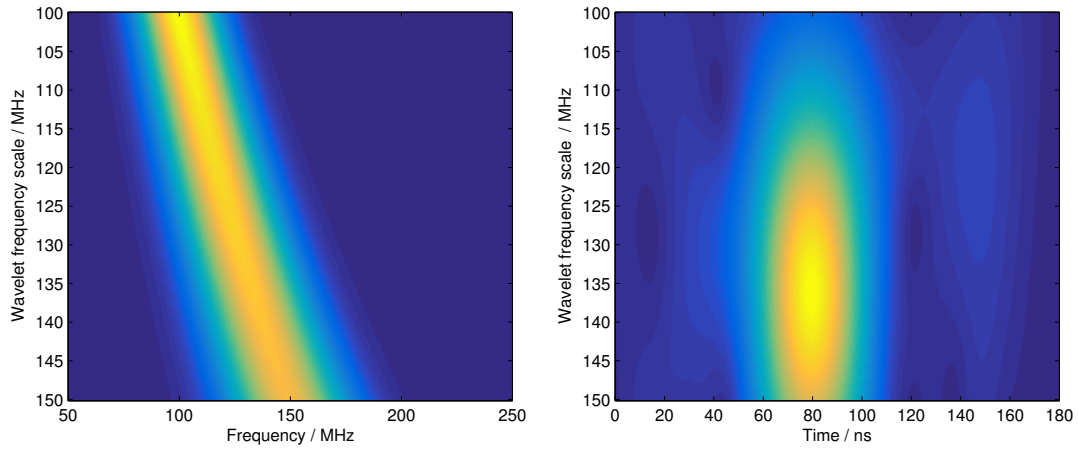
of the dominant frequency determination and computation time.

2. Generate a family of wavelet daughters, discretely spaced in frequency space, over a given frequency range, shown in Fig. 4.16 (a).

3. Take the FFT of the input signal and multiply by each wavelet, this is the continuous wavelet transform. The IFFT of this is shown in Fig. 4.16 (b).

4. Sum along the continuous frequency axis and find the peak in this curve, these curves are shown in Fig. 4.17. If required it is also possible to sum across the orthogonal axis of the plot in Fig. 4.16 (b) and find the peak in time.

The result of the application of a set of wavelet daughters to a simulated SRAS signal is shown in Fig. 4.17. Since the wavelets applied have varying temporal widths and shapes the frequency peak shapes vary greatly in width. Finding the peaks of each of the frequency traces gives the following errors relative to the ideal signal: Paul -1.23%, Morlet -0.11%, Marr -0.99%, and MFP -2.20%. Clearly, for this example, all of the wavelet transforms have returned a frequency value closer to the expected result than the MFP method.



(a) (b)

Figure 4.16: a) The magnitudes of a set of 112 scaled Morlet wavelets in the frequency domain. (b) The magnitude of the wavelet transforms of the SRAS signal. This is plotted in time-frequency space.

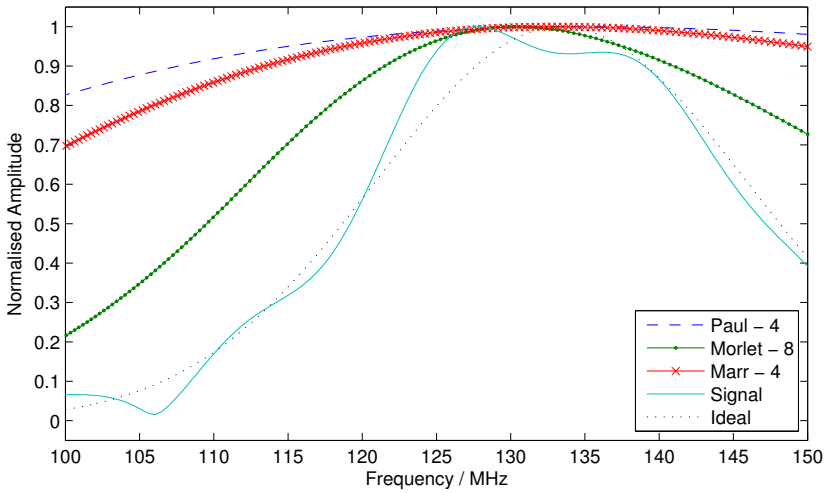


Figure 4.17: A simulated 9 cycle noisy frequency spectrum plotted with the ideal frequency spectrum (solid blue and black dotted respectively). Three wavelet transforms have been applied and summed across all times. The transforms are order 4 Paul (dashed blue), order 8 Morlet (dot dash green) and order 4 Marr (cross dash red).

4.4 Comparison of processing techniques applied to simulated SRAS signals

In order to assess the various signal processing techniques described above a Monte-Carlo type simulation was applied. A set of SRAS signals was simulated in the frequency range 110 MHz to 140 MHz with steps of 0.5 MHz and with cycles between 4 and 12 in steps of 2. For the commonly used SRAS acoustic wavelength of 23 µm, this frequency represents a SAW phase velocity range of $2530\,\mathrm{m\,s^{-1}}$ to $3220\,\mathrm{m\,s^{-1}}$. This range is consistent with the SAW velocity ranges on typical engineering metals. To each of these simulated signals, filtered normally distributed random noise was added as described in §4.2. The amplitude of the added noise was scaled such that the SNR varied between −4 dB to 20 dB with 20 steps.

These signals were then processed using the techniques described above to find the dominant frequency peak, and the disagreement with the ideal frequency calculated for each signal. For the processing techniques which require a set of daughter functions over a range of frequencies, a frequency range of 100 MHz to 150 MHz was used. A mean error and standard deviation was taken of this disagreement over all frequency steps, with the errors at each frequency step expressed as a percentage error of the expected frequency. Once complete a new noise distribution was added to the simulated signals and the traces were processed again. This cycle was repeated 30 times in total for each step in signal cycles. These simulations have been implemented in MATLAB and take approximately 7 hours on a fast, multi-core processor PC to process.

4.4.1 Signal to noise

The results of the simulated signal processing are shown in Fig. 4.18 and Fig. 4.19. For each of the processing techniques the standard deviation in frequency of each fitted point signal from the ideal signal frequency was averaged for each noise level. This is shown in Fig. 4.18 for the case of 6 cycle signals. Similarly, Fig. 4.19 shows the mean error away from the ideal signal frequency for the same fitted traces. Essentially the mean shows how close to the expected measurement you are, while the standard deviation shows how much variation there is around the ideal value. Together these values indicate how well each method determines frequency. For example, the mean of

the fitted values from the MFP method are all close to zero over the full noise range, however the standard deviation is large especially at low SNRs. As shown in Fig. 2.11, the MFP method is sensitive to distortions of the frequency peak due to noise and since the noise applied is incoherent, the probability of the distorted peak being determined as larger or smaller than the ideal frequency are the same. Therefore, over several repeats, the determined value averages out to a mean error close to zero.

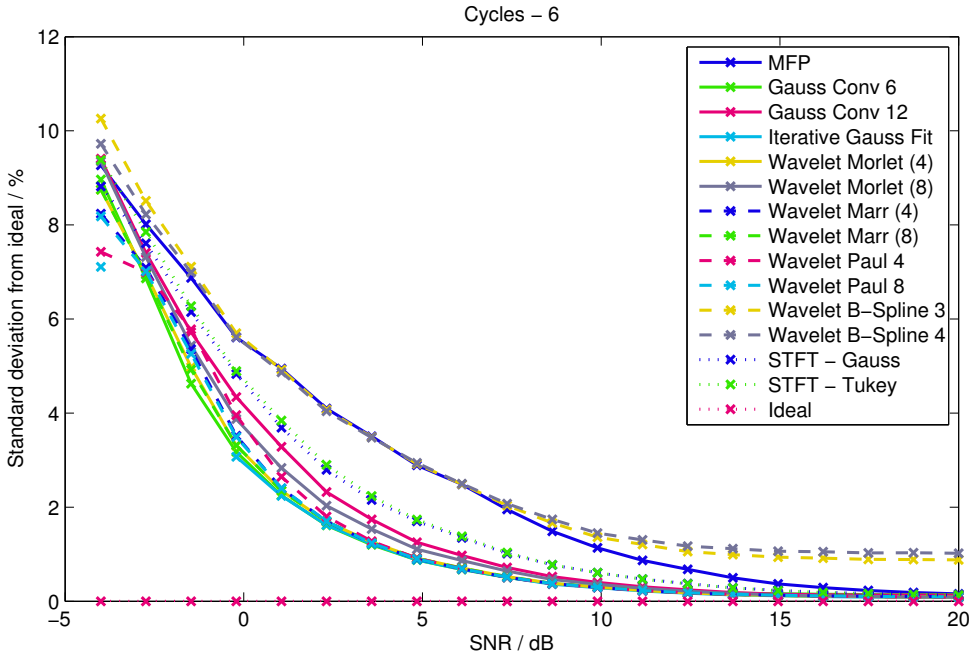


Figure 4.18: The standard deviation of the differences between the measured frequency peaks and actual frequency peaks for a set of analysis techniques. The simulated SRAS signals have 6 cycles and have had random noise added to give SNRs from -4 to 20.

Fig. 4.18 indicates that all fitting techniques perform almost equally poorly at very low SNRs; there is no stand-out technique. As SNR improves the efficacy of all techniques improves in a fairly similar manner, with most techniques converging at 20 dB with the exception of both orders of B-spline wavelets. The fact that B-splines do not match the other techniques suggests that they are poorly suited to the application of SRAS signals. Most likely this a result of the slight discretisation in centre frequencies required to construct the wavelet. This would explain why the standard deviation is large but consistent, as there is always an offset between the SRAS frequency and discretised wavelet frequency, with the low mean error indicating that the wavelets are actually doing a reasonable job of determining the correct frequency. Gaussian iterative

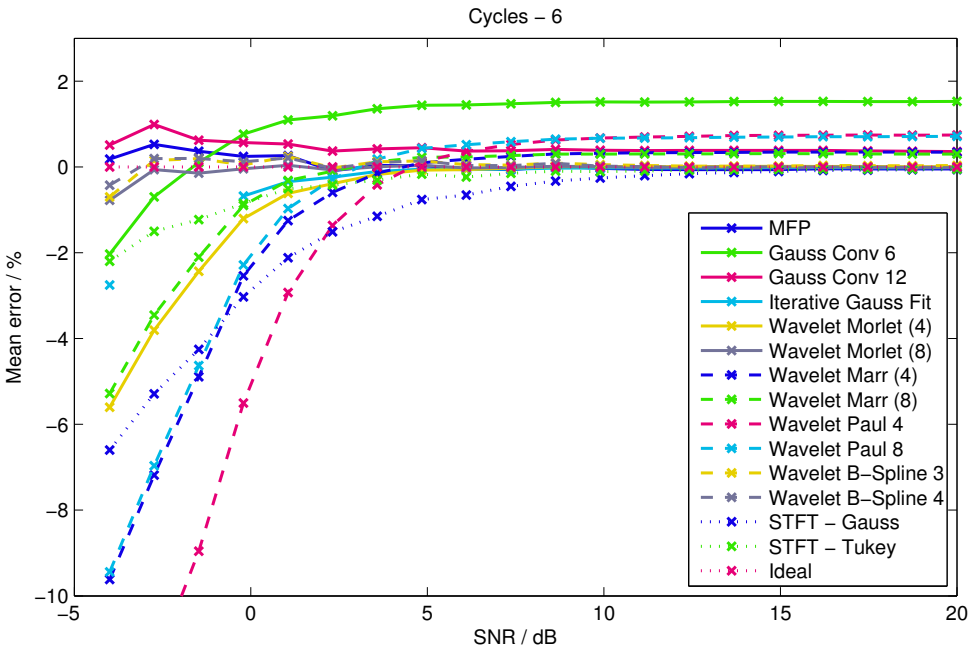


Figure 4.19: Mean values of the differences between the measured frequency peaks and actual frequency peaks for a set of analysis techniques. The simulated SRAS signals have 6 cycles and have had random noise added to give SNRs from -4 to 20.

fitting and the order 4 Morlet, Paul and Marr wavelets show the best performance in standard deviation.

Looking at Fig. 4.19 shows that Paul wavelets and order 6 Gaussian convolution do a very poor job of finding the correct frequency value at all SNRs. In fact they both seem to have a fairly consistent offset which is significantly larger than the frequency spacings between adjacent wavelet daughters. Paul wavelets are very sensitive to phase and this effect may be dominating over detection of oscillations in signal amplitude. Marr wavelets and the order 12 Gaussian convolution agree slightly better with the ideal measurement but are still fairly poor at all SNRs. At SNRs less than ~3 dB only MFP, order 8 Morlet and B-spline wavelets perform reliably, although given the results from the standard deviation figure the B-spline wavelets are less than ideal. The mean values show that across the plotted noise range, order 8 Morlet wavelets have excellent agreement with the ideal measurement and that the order 4 Morlet and Gaussian iterative fit also have good agreement. The STFT method gives a reasonable standard deviation but quite poor mean value, especially when a Gaussian window is applied. An STFT using Hann windows was also applied but the results proved to

be indistinguishable from the application of a Gaussian and so were omitted. From the data plotted so far the most suitable techniques are order 8 Morlet wavelets and Gaussian iterative fitting.

4.4.2 Signal cycle number

The same simulations were run for SRAS signals of cycle numbers 4, 6, 8, 10 and 12, and in general the same trends were seen for each of the techniques shown in Fig. 4.18 and Fig. 4.19. To avoid overloading the reader with figures for every cycle number a comparison over all techniques and cycle numbers was generated by averaging the mean and standard deviation values over all SNRs. This obfuscates some of the more subtle variations between the techniques, however these have been highlighted already in the previous figures, and ultimately the target of this analysis is to choose the most suitable technique over all conditions, therefore the subtle differences can be neglected. Fig. 4.20 shows the averaged values for the standard deviation and mean respectively plotted as a bar chart, each bar has been segmented based on the contributions at each cycle number.

Examining the data for the 6 cycle signals shows good agreement with Fig. 4.18 and Fig. 4.19. For the standard deviation all of the wavelet analyses, with the exception of B-splines, are closely grouped and B-spline, MFP and STFT techniques are again shown to be fairly poor. Looking at the mean values shows order 8 Morlet wavelets as the stand-out technique, with Paul wavelets and order 6 Gaussian convolution as the worst of the techniques.

Across all of the data in Fig. 4.20 it is evident that the efficacy of all the techniques improves with cycle number, intuitively this is an expected result as there is simply more signal to fit to and a sharper peak in frequency. The improvement with cycle number is non-linear, there is a distinct step from 4 to 6 cycles while further increases in cycle number show a more marginal improvement. This is a good result as it shows that the benefit of very long cycle signals is small and that shorter signals, such as 8 cycle, can be used with little penalty.

Looking at the results for standard deviation shows that many of the techniques are fairly comparable; Gaussian convolution, iterative Gaussian and wavelet analyses, with

the exception of B-splines, are very similar especially if the 4-cycle signal case is ignored. B-spline wavelets perform very poorly at all cycle numbers with MFP also showing a high variance around the expected measurement.

Turning to the mean values shows a much greater difference between the techniques, with a much clearer indication of the best and worst performing. By a significant margin Gaussian convolution is the worst technique for correctly finding the frequency peak, followed by Paul wavelets, STFTs and Marr wavelets. Although the standard deviations are the worst, B-spline wavelet analysis is amongst the best techniques in terms of mean values. This indicates that the analysis works poorly but is randomly equally weighted either side of the correct frequency peak. The most effective techniques are MFP, iterative Gaussian and Morlet wavelets with order 8 Morlet wavelet analysis proving to be, by far, the best.

Consistently over all of the comparisons used, order 8 Morlet wavelets have performed excellently. Given the similarity of the Morlet wavelet to SRAS signals, and that wavelet analysis is designed to pick out certain features of a signal, it is not unexpected that this technique is thoroughly suitable. As such Morlet wavelet fitting has been chosen as the default signal processing method for SRAS signals, and Gaussian iterative fitting has been identified as the second choice.

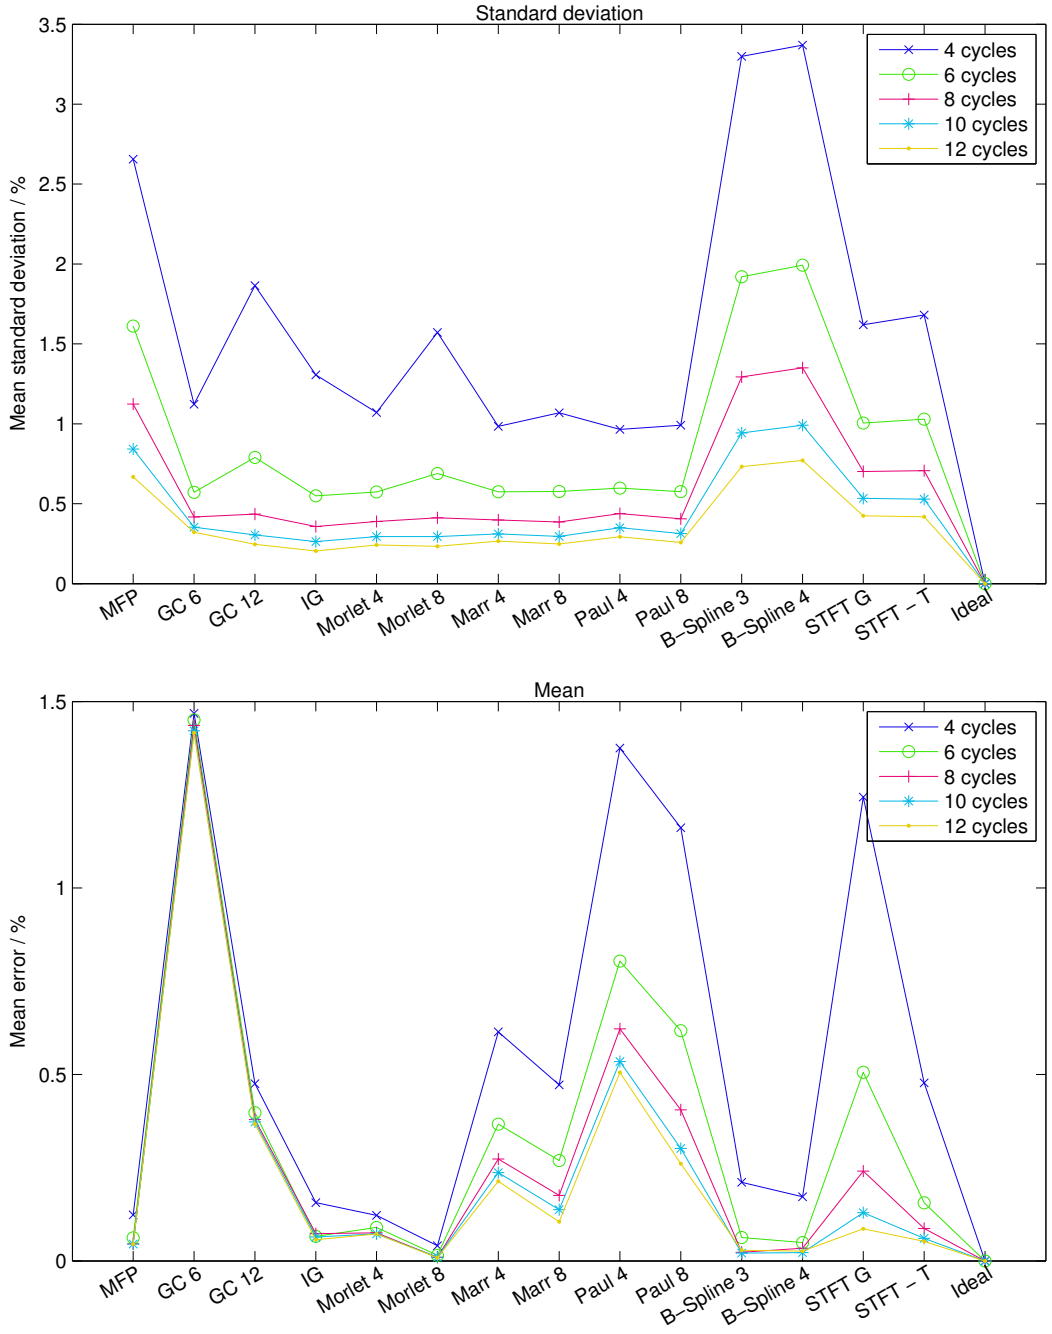


Figure 4.20: The sums of the standard deviation and mean values of the differences between the measured frequency peaks and actual frequency peaks for a set of analysis techniques and signal cycle numbers. The unused abbreviations are: Gaussian convolution order 6 and 12 - GC 6/12, iterative Gaussian - IG, and STFT Gauss and Tukey - STFT G/T.

4.4.3 Frequency resolution

Morlet wavelet fitting, like most of the techniques used, applies a set of wavelet daughters scaled in frequency to the SRAS signal. The choice of range in frequency is easily determined by the material being fitted, while the choice of step size in the frequency range is free. A very fine step-size will give better frequency resolution but at the cost of computation time. To achieve real-time data processing a compromise has to be made. Fig. 4.21 shows a comparison between multiple scale parameters of Morlet wavelet applied to simulated data. The scale parameter describes the frequency resolution where a small scale parameter indicates a fine frequency resolution. The standard deviations show the expected result: as scale parameter decreases frequency step also decreases, and a finer resolution is returned. Although there is a slight trend in the mean values towards better fitting at smaller scaler parameters, there is a lot of variation in this result and any differences are negligible.

At all scale factors the Morlet wavelet analysis significantly outperforms the MFP technique. Fig. 4.22 shows the relative processing time and minimum frequency step for all of the Morlet wavelet scale parameters. Processing time appears to be non-linear with frequency step size. Fitting an exponential curve to processing time, t, against frequency step in MHz, ν, shows that $t(\nu) \simeq \exp(-1.3\nu + 5.8)$. Due to this non-linearity a good compromise is difficult to choose and must therefore be determined by the processing time available. For the Renishaw SRAS system a Morlet parameter of 0.006 is typically used as this gives equal data capture and processing rates. Smaller scale parameters are used for low anisotropy materials or where velocity resolution is critical.

4.4.4 Processing experimental data

With the preferred fitting methods identified and an optimum frequency resolution chosen, processing can be applied to experimental SRAS data. A scan of a titanium 685 sample has been taken with an acoustic wavelength of 22.9 µm, a 200 µm patch size and with steps sizes of 25 µm in the x direction and 50 µm in the y direction. The sample scanned has been etched to enable optical measurements and this has left a low reflectivity surface causing noisy SRAS signals, typical SNRs are 2 dB to 4 dB with

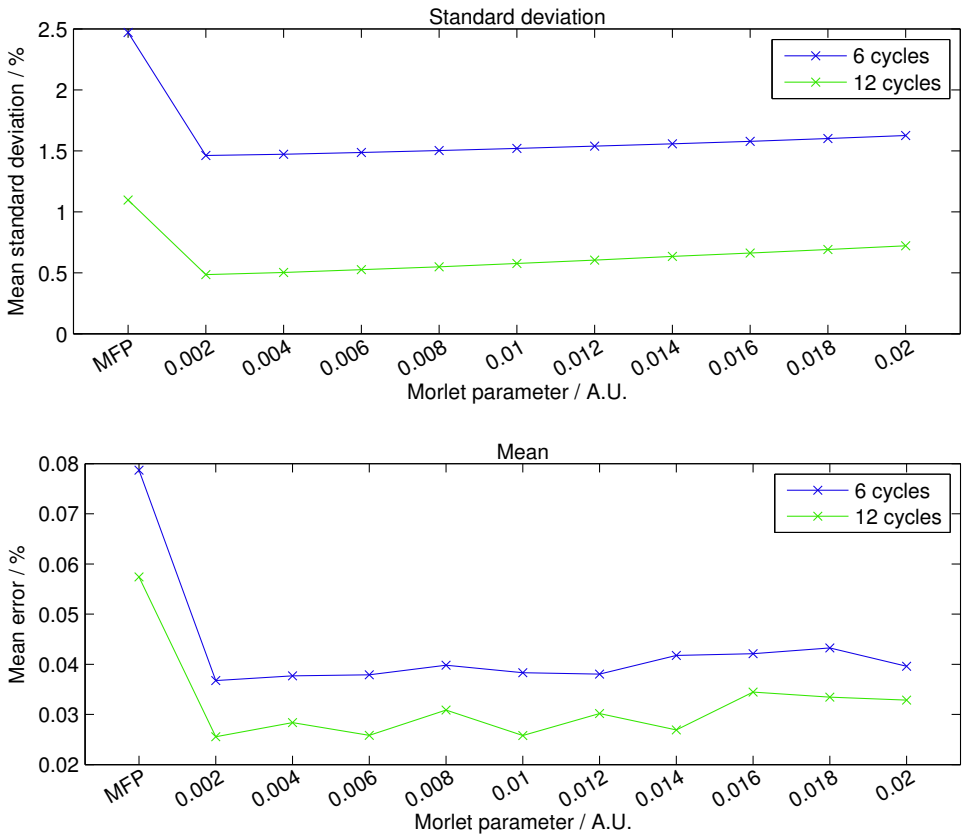


Figure 4.21: The sums of the standard deviation and mean values of the differences between the measured frequency peaks and actual frequency peaks for a set of Morlet wavelet scale parameters for 6 and 12 cycle signals. A smaller parameter indicates a finer frequency step.

signal cycle numbers between 6 and 8. Four signal processing techniques have been applied: MFP, iterative Gaussian fitting, and Morlet wavelets of order 4 and 8. The results are shown in Fig. 4.23.

When analysing SRAS velocity contrast maps there are a two key markers to look for. The first is low amplitude signals which are indicated by the white part of the colour map. Each image is masked using the optical illumination feedback and the fitted signal strength. Where these are below a conservatively set threshold, typically 10% of the median signal strength, the pixel is painted white. The second marker is good spatial correlation and consistency of velocity. SRAS is a single-shot measurement and thus each measurement is completely independent and should only depend on random incoherent noise and the local microstructure. Therefore, within a single grain, each determined velocity should be consistent with other measurements taken on the same

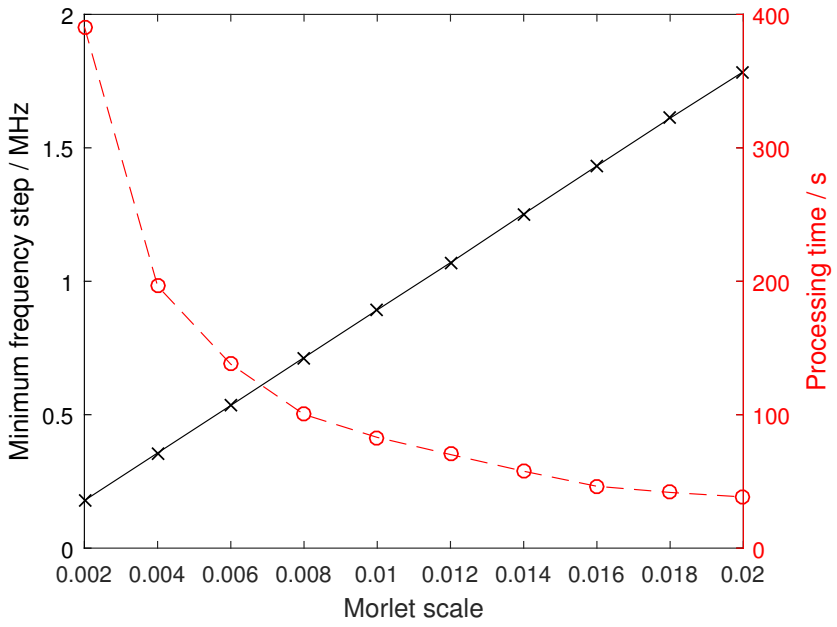


Figure 4.22: The minimum frequency (black solid) and relative processing time (red dashed) for a range of Morlet wavelet scale parameters. Processing time is summed over a large data set.

grain. Large variations of velocity within a grain are indicative of noise.

In general there is not a huge difference between each of the images; the velocity ranges are similar and the major grains have been identified in all cases. However, in terms of consistency and noise across grains the MFP method is considerably worse than the others. Ti 685 typically exhibits small scale intra-granular long and thin aligned features, called acicular grains [89]. In the MFP image these features are obscured in the noise, whereas on certain grains from the other images a streaked pattern is visible.

Between the Morlet and Gaussian methods there are some subtle differences. Most noticeable is the slight difference in SAW phase velocity on slower grains where the Gaussian method seems to overestimate velocity compared to MFP and Morlet techniques. The difference is small, $\sim 25\,\mathrm{m\,s^{-1}}$, but noticeable in the images. Based on the form of the raw data, the cause of this is most likely slight clipping of the signal due to saturation of the amplifiers in the signal chain which alters the shape of the signal frequency peak. The Gaussian iterative method is more sensitive to this effect than the Morlet wavelet analysis. This highlights another benefit of the Morlet analysis, as processing is ideally insensitive to small errors in instrument set-up. By eye the

Technique	Region 1		Region 2		Region 3	
	STD	**Mean**	**STD**	**Mean**	**STD**	**Mean**
MFP	65.6	2780	75.8	3146	53.1	2697
Gaussian iterative	23.4	2818	28.5	3139	26.54	2734
Morlet 4	26.8	2779	30.5	3135	25.3	2679
Morlet 8	27.2	2787	31.3	3141	29.5	2687

Table 4.1: The standard deviations and mean, measured in $\mathrm{m\,s^{-1}}$, for the three regions shown in Fig. 4.23 for 4 different fitting techniques. The regions are identified by colour in Fig. 4.23: 1 - white, 2 - magenta and 3 - cyan.

difference between the Morlet wavelet analyses is not that visible, in agreement with the earlier simulations they both do a good job in the SNR range of this scan.

Several sub-regions of Fig. 4.23 have been selected, and the standard deviations and means of these regions calculated. These data are shown in Table 4.1. In general these data are in agreement with the conclusions drawn from the velocity maps, with a large difference in standard deviation from MFP to the other techniques. Gaussian iterative fitting gives a very small variance which reflects the findings shown in the comparison on simulated data. The simulation also showed Gaussian iterative fitting gave a worse estimate of velocity than the Morlet techniques, and this is highlighted in the table where it shows poor agreement with the other techniques.

Fig. 4.24 shows a second example of a rolled stainless steel sample with poor surface condition and a fine grain structure. A 4 cycle SRAS signal was used in order to maximise spatial resolution such that the small scale microstructure could be resolved. Due to this, and the poor surface reflectivity, SNR was very low, $\sim$1 dB to 3 dB and the MFP processed image does not clearly show the distribution of the microstructure. However applying an order 8 Morlet wavelet produced an image with a much better average SNR and far fewer areas of poor signal amplitude. In addition it is more clear that there are regions of consistent SAW velocity which better indicate the sample microstructure.

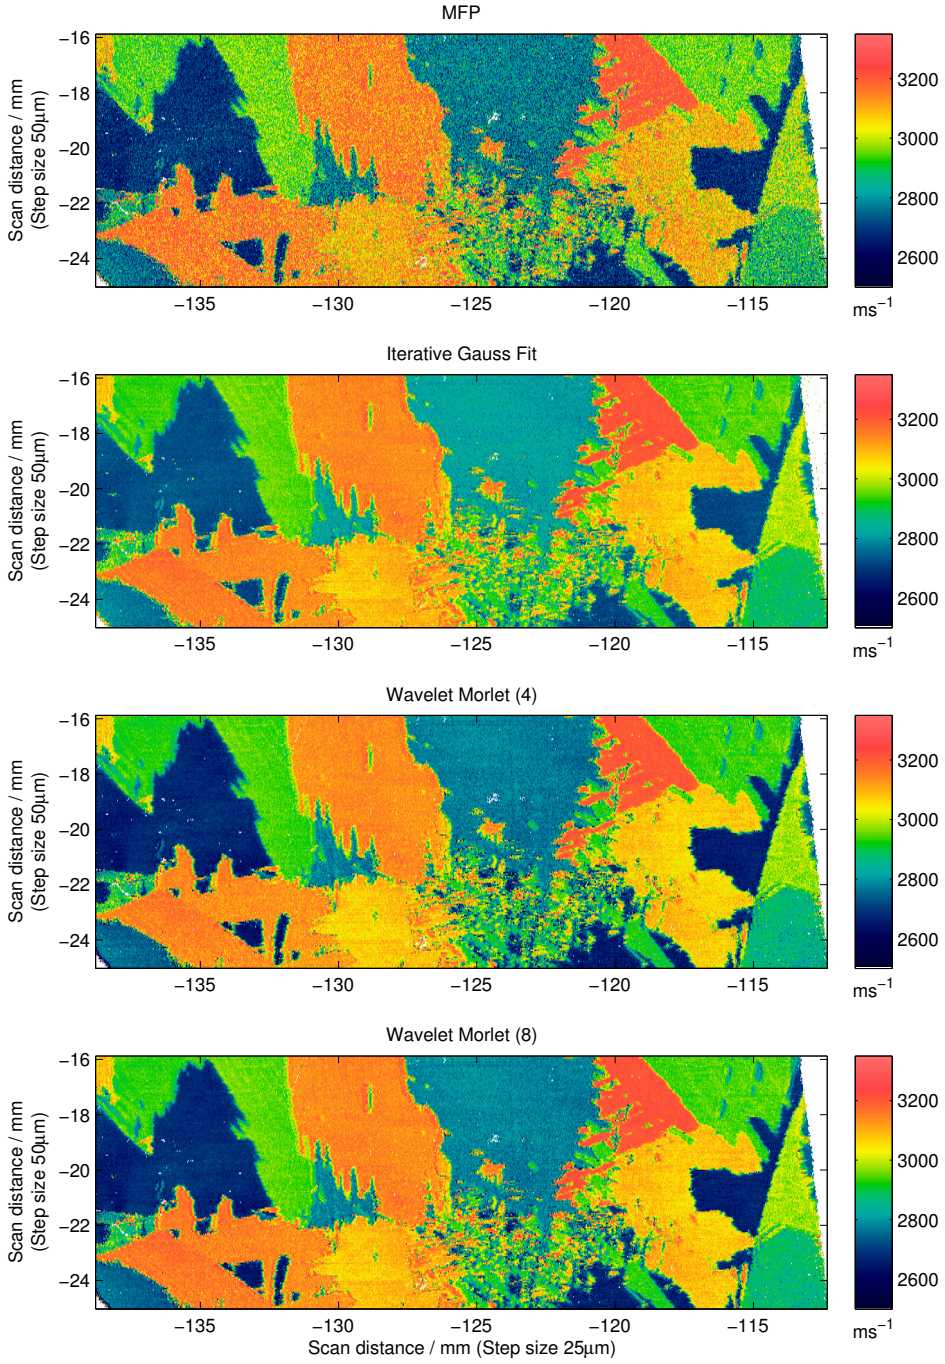


Figure 4.23: SRAS scans of a Ti 685 sample processed using four different techniques. From top to bottom: MFP, iterative Gaussian fit, Morlet wavelet of order 4, and Morlet wavelet of order 8. Three sub-regions have been taken from each data-set and are compared in Table 4.1.

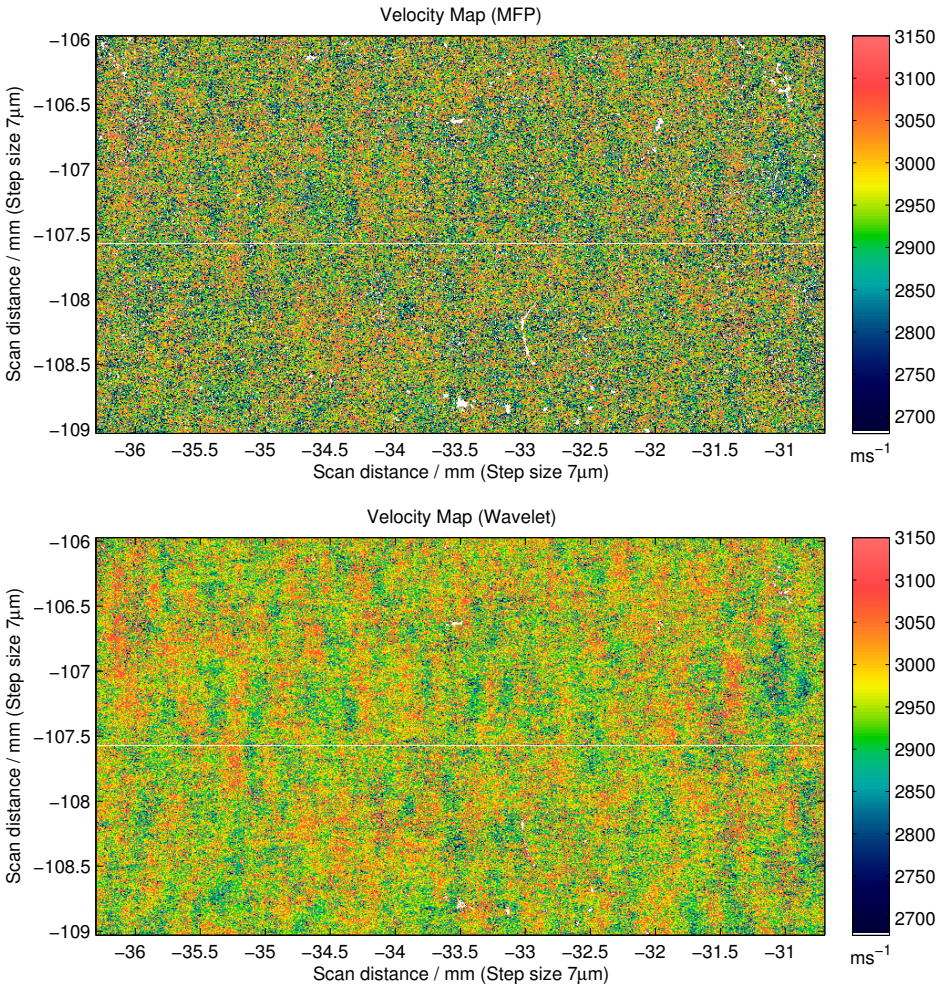


Figure 4.24: A rolled stainless steel sample with very poor surface condition. A 5 cycle SRAS signal was used in order to improve spatial resolution.

4.5 Spatial averaging

A second method for improving the SNR of the dominant frequency measurement is to average two or more traces which are spatially contiguous. Spatial resolution is decreased in doing this by approximately the same factor as the number of averages. This can be simulated using the SRAS signal model to determine the theoretical benefit of taking averages. Fig. 4.25 shows the method for applying an averaging simulation. A set of N signals, where N is the number of averages to be taken, are generated with identical cycle numbers and frequency but with a unique random noise distribution. A small temporal jitter - up to 1.5 ns - is also applied to each trace which replicates the temporal jitter of the data acquisition system. These set of signals are then averaged and the Fourier transform taken and used to find the signal amplitude. The Fourier spectrum of the idealised signal is then subtracted and the result used to determine the noise amplitude.

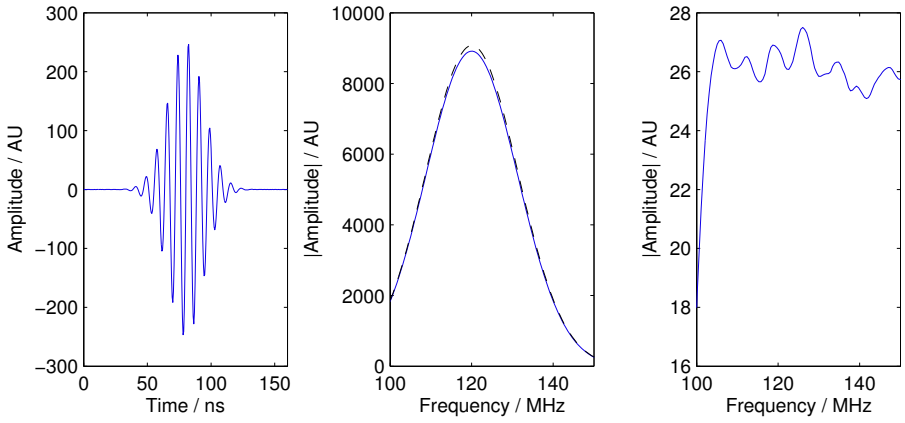


Figure 4.25: An 8 cycle SRAS signal with added noise is Fourier transformed and the ideal signal subtracted to estimate the SNR.

This process has been applied to a set of signals with a range of initial SNRs between 1 dB and 24 dB and the result is shown in Fig. 4.26. The expected improvement of SNR with averaging is proportional to $\sqrt{N}$. While a linear relationship has not been recovered, clearly SNR does increase with the number of averages. Over 256 averages the improvement to all signals is ~12 dB. The non-linear trend suggests that either the signals are not well correlated or that the noise is mildly correlated. Repeating the simulation with zero jitter, i.e. with perfect signal correlation, yields similar results in SNR improvement, thus there is some correlation in noise and this is most likely due

to the discretisation required for computation. An important point here is that the improvement in SNR is independent of the initial SNR, allowing improvements to any signal to be reliably predicted.

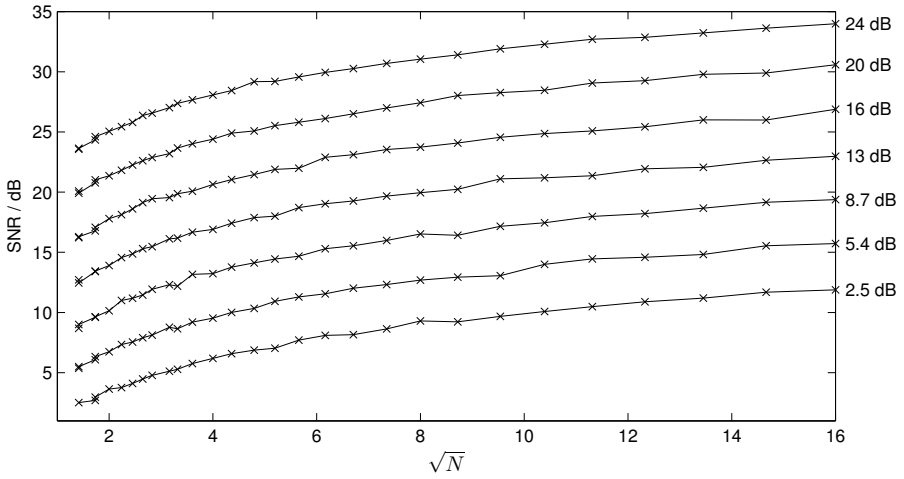


Figure 4.26: The result of applying the averaging simulation to a number of 8 signal cycles over a range of initial SNRs. The initial SNR of each trace is shown along the right-hand y axis.

SRAS is applied as a scanning technique and a constant speed is maintained throughout data capture. Therefore, in order to apply averaging, signals captured from different spatial locations have to be used and the overall spatial resolution of the scan is reduced. Fig. 4.27 shows various degrees of spatial averaging applied to measured SRAS data and fitted using the MFP method, these data are from the same scan shown in Fig. 4.23. Raw traces are extracted either as a line in the scan direction or as a square/rectangle and averaged before processing by MFP and Morlet wavelet analysis. The improvement in the image with averaging is not very good; the speckled noise within grains is not removed at any level of averaging. Spatial averaging reduces the spatial resolution and reduces the ability to resolve features and edges. One slight improvement can be seen in areas where the signals have been too noisy or small to fit, plotted as white. With averaging, there are slightly fewer white patches over the image. For comparison, Fig. 4.28 shows the same averaged SRAS data, but with a Morlet wavelet analysis applied. As before there is little improvement in the, initially very good, image.

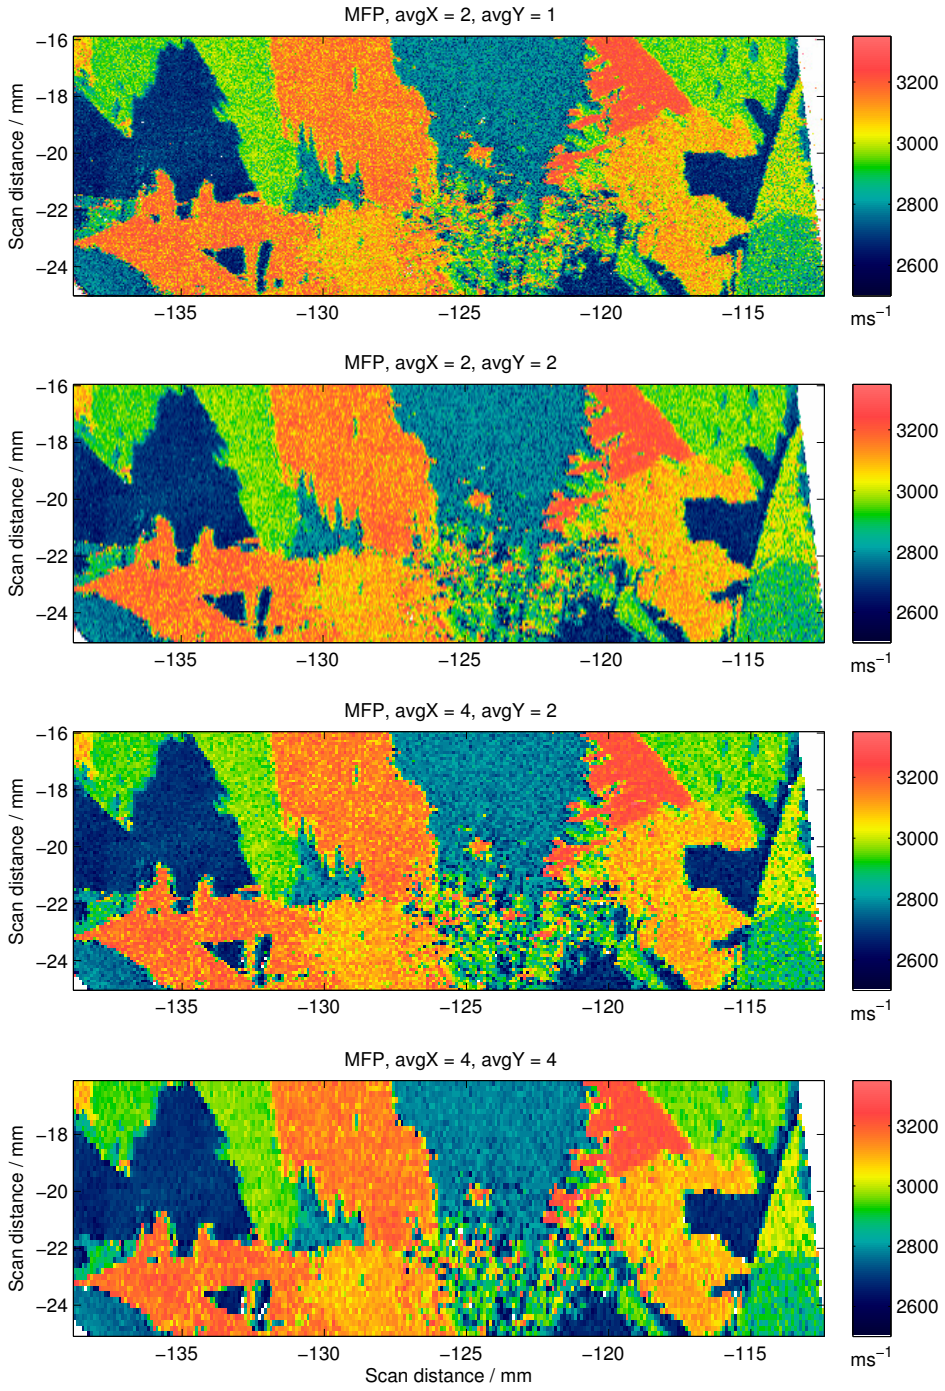


Figure 4.27: A SRAS scan which has been spatially averaged in multiple ways and processed using the MFP method. The initial pixel size is 25 µm in X and 50 µm in Y.

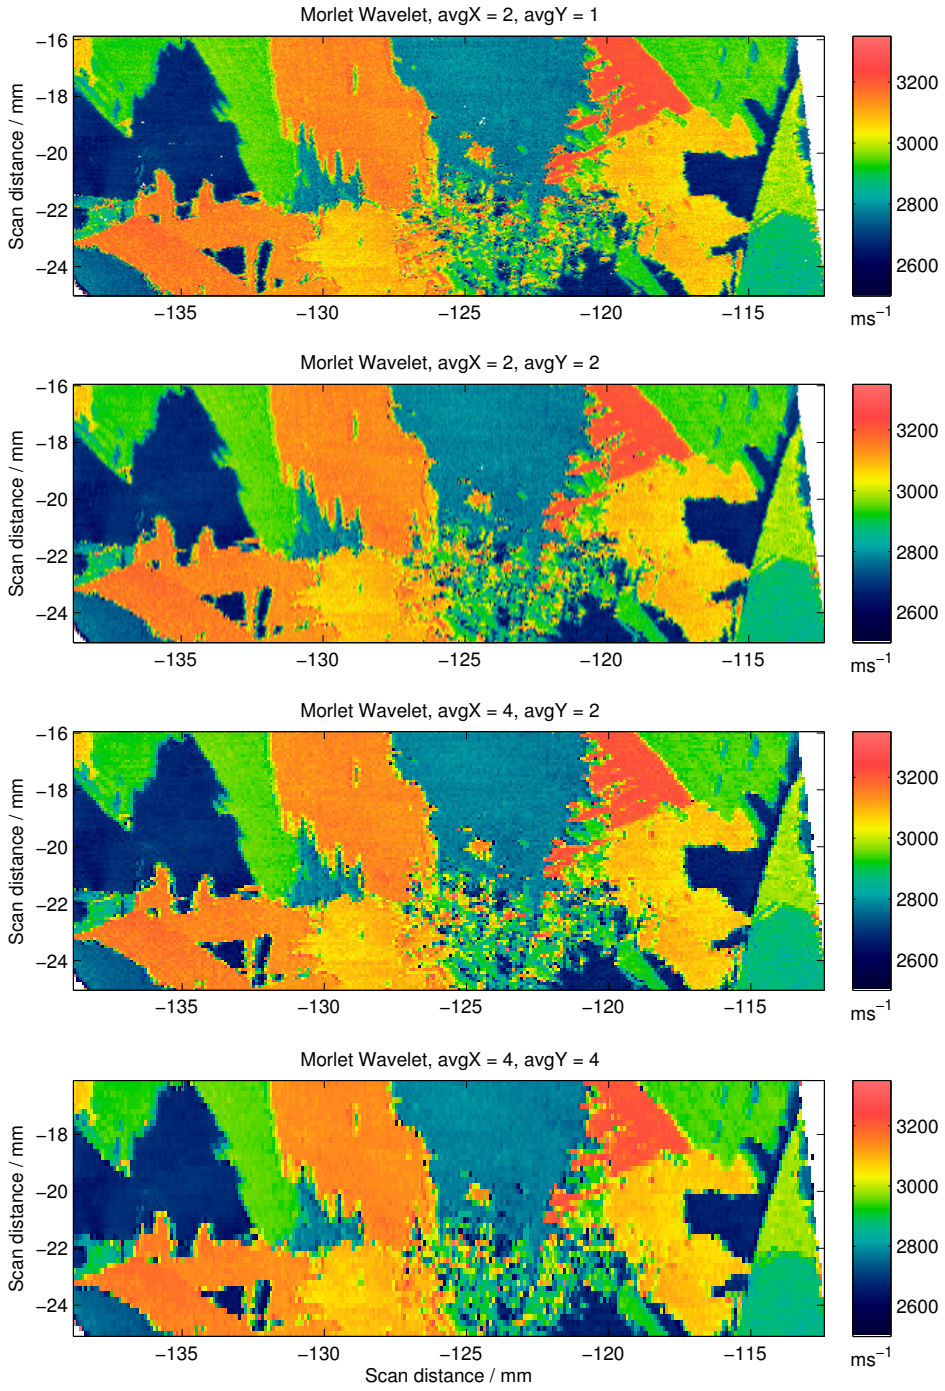


Figure 4.28: A SRAS scan which has been spatially averaged in multiple ways and processed using a order 8 Morlet wavelet. The initial pixel size is 25 µm in X and 50 µm in Y.

4.6 Conclusions

In this chapter the aim was to investigate whether a more robust means for determining SAW phase velocity from noisy signals existed, and to quantify what improvements from the MFP method could be made. The result is the identification of Morlet wavelet analysis as a suitable candidate. It has been shown quantitatively from Monte-Carlo type simulations of SRAS signals that Morlet wavelet analysis can more frequently and more reliably find the maximum frequency peak in addition to the other methods trialled. From a qualitative analysis of SRAS scan images, it has been shown that Morlet wavelet analysis can produce better images than MFP analysis. These images have less noise and improve the visibility of various sample features.

In comparing different orders and scale factors applied to Morlet wavelet analysis it has been shown that the best compromise is to use order 8 wavelets with a scale factor of 0.006. With these values the target of processing data in real-time as it is captured has been achieved. It has also been shown that this processing technique can be tweaked if greater frequency resolution or minimal processing time is required in order to achieve the desired result.

4.7 Further work

4.7.1 Group velocity

In addition to determining frequency, the Morlet wavelet technique also isolates the SRAS signal in time, allowing the measurement of arrival time to be made. From the generation and detection patch configuration the travel distance of the SAW can be estimated. The travel time has not been accurately estimated and is based entirely on the delay between the electronic trigger output from the laser, and the start of signal capture. This allows a consistent, although poor, time-of-flight estimate of SAW velocity to be made. Effectively this is group velocity, $v_g = \frac{\mathrm{d}\omega}{\mathrm{d}k}$, but for non-dispersive SAWs $\omega = v_p k$ and therefore $\frac{\mathrm{d}\omega}{\mathrm{d}k} = v_p = v_g$, therefore the SAW phase velocity measurements and group velocity measurements should be identical. Fig. 4.29 shows a comparison of the phase velocity and group velocity images. The group velocity was estimated from the measured delay time, fibre optic lengths and distance between the

centre of the generation patch and detection spot. Evidently the estimation of group velocity is fairly poor, however if considered only as a contrast mechanism, then the two images are fairly similar. Many, but not all of the grain boundaries are visible in both images. There is significant noise on the time delay measurement which is due to the temporal jitter of the laser Q-switch. Over similar grains this noise dominates over the small differences in phase velocity.

One of the most interesting features of the image are the grains which show large disagreement in phase and group velocities, for example the very slow central grain in the phase velocity image is nearly indistinguishable in the group velocity image. This could occur for a few reasons: a change in the signal shape, a thin grain layer - SAWs are dispersive if in constrained layers, or dominant PSAW mode - PSAWs are slightly dispersive [90]. The raw time traces show very little difference, and without destructive testing of the sample or methods to determine SAW wave mode shape this remains an open problem.

4.7.2 Iterative wavelet analysis

In comparison with the previously used MFP fitting method, wavelet analysis is ~2 orders of magnitude slower to apply. The reason for this is the requirement to multiply every signal trace by a fairly large set of wavelet daughters, typically 100 or more, before finding the peak value. Although, with the application of parallel processing and batch processing, wavelet analysis can keep up with data collection, future plans for the technique involve faster laser repetition rates, greater DAQ sample rate and higher bit depth, all of which increase the processing computational requirement. In addition it was shown in Fig. 4.21 and Fig. 4.22 that to achieve optimum fitting, finer step sizes in frequency are required and that this comes at the cost of processing time.

To ameliorate both of these issues an iterative process is proposed to replace the current exhaustive process. By fitting with a coarse frequency step an interim maximum can be found and a set of wavelet daughters with a smaller scale factor generated and fitted near the maximum. With repeated iteration an arbitrarily high frequency resolution can be achieved. Fig. 4.30 shows this process applied to a single trace. To cover the frequency range with the exhaustive method and a spacing of 0.34 MHz takes

147 calculations. With an iterative method, over the same frequency range, only 36 operations are required to achieve a frequency spacing of 0.14 MHz. Clearly a much better resolution is achieved with lower computational requirement. The future work here is in ensuring that this process is sufficiently robust to avoid local maxima and to determine the optimum number of iterations required. In addition this iterative technique could also be implemented as a LMS algorithm.

4.7.3 Wavelet filtering for orientation determination

Morlet wavelets are very good bandpass filters, Fig. 4.31 shows an input, noisy SRAS signal and the returned, filtered signal. Use of this, in combination with the frequency extraction work described in the previous section, could be used to clean-up noisy signals before the orientation determination process is applied.

4.7.4 Dual frequency signals

SRAS signals are not always single frequency and there are two primary mechanisms by which multiple frequency SRAS signals are generated. The first is due to the support of multiple SAW and PSAW modes on certain crystallographic grains. Typically one of these modes is dominant over the other, but in some cases the amplitudes can be fairly similar. The second is at grain boundaries where the generation patch is split over two grains causing delays, and phase changes, in the measured signal. Since Morlet wavelet analysis is a time-frequency technique, these different modes can be separated in both frequency and time. The signals affect the determination of SAW velocity; instead of a dominant peak at one of the frequencies, an average over multiple frequencies is returned. Therefore, separation and identification of each frequency in a signal will give a more true measure of SAW phase velocity. Fig. 4.32 shows a simulated SRAS signal which has dual frequency content and the result of fitting with both order 3 and order 8 Morlet wavelets. Clearly order 3 Morlet wavelets give better temporal resolution allowing the two SAW modes to be separated in both time and frequency. Here further work is required in how to extract this frequency information from these fitted data and exploring the value of data fitting at multiple orders of wavelet daughter.

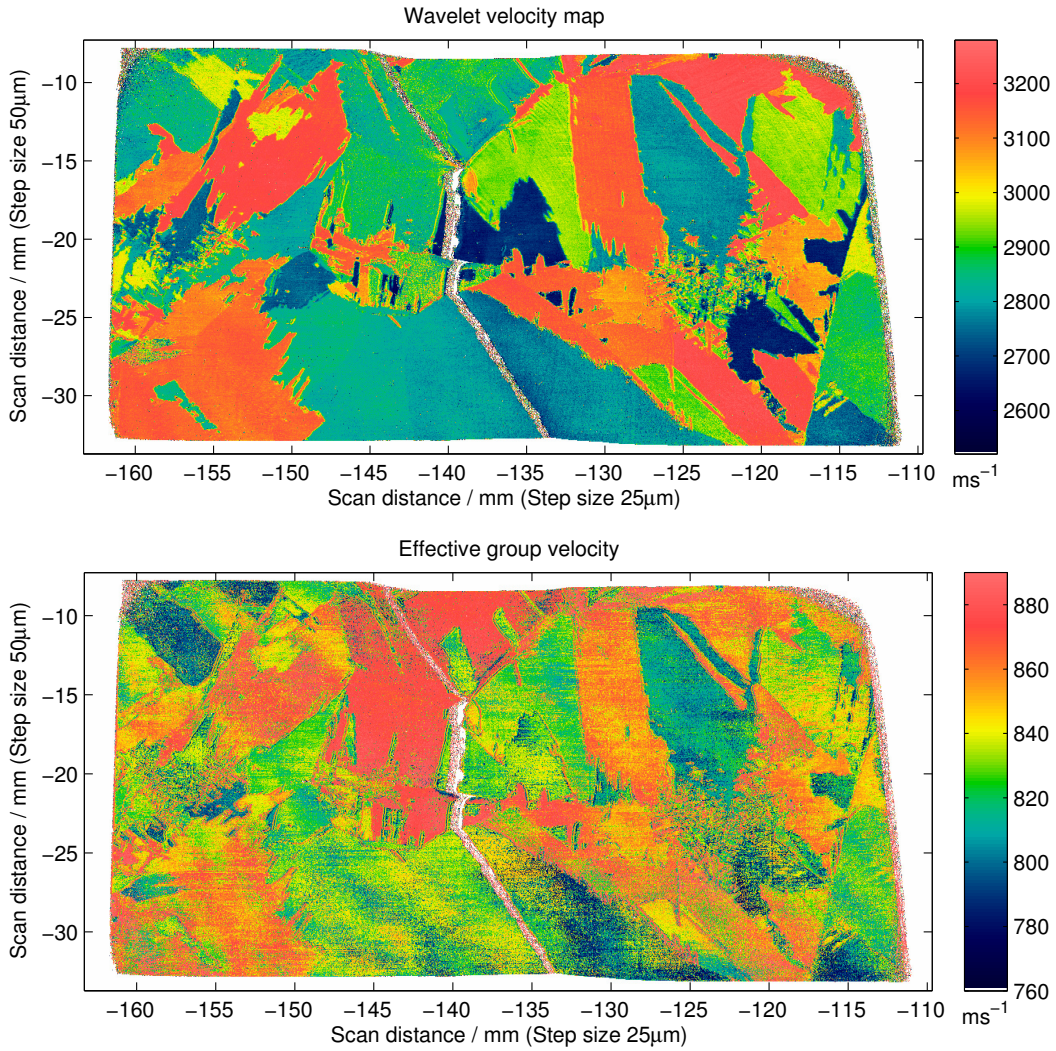


Figure 4.29: SRAS scans of a Ti 685 sample processed using the Morlet wavelet technique. Top: the phase velocity determination, bottom: the effective group velocity based on the time delay measurement.

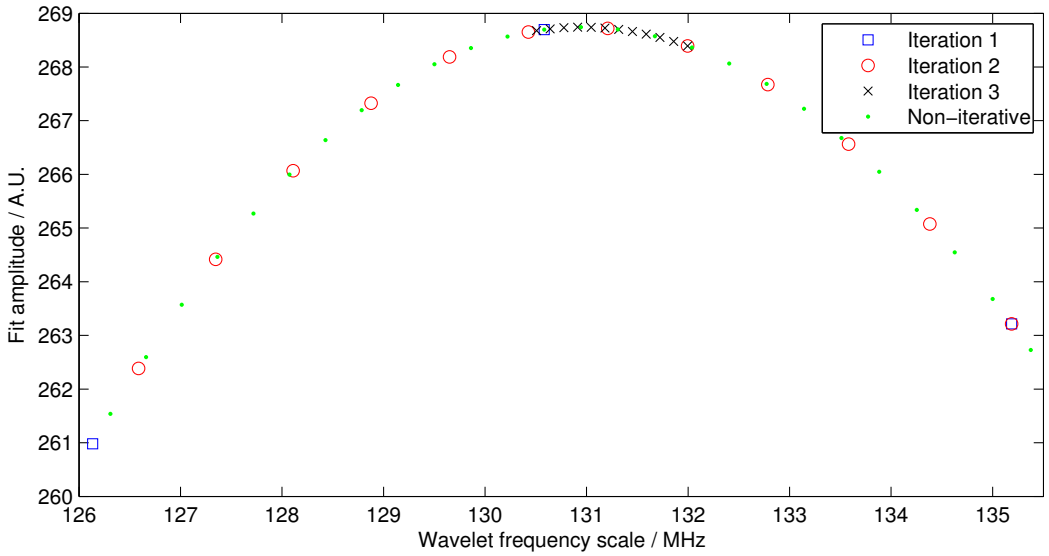


Figure 4.30: An iterative wavelet processing technique applied to SRAS trace compared to an exhaustive processing technique. The iterated region is determined by defining the new frequency range as the two neighbours to the maximum amplitude point. Only 3 iterations were required to exceed the exhaustive search frequency spacing.

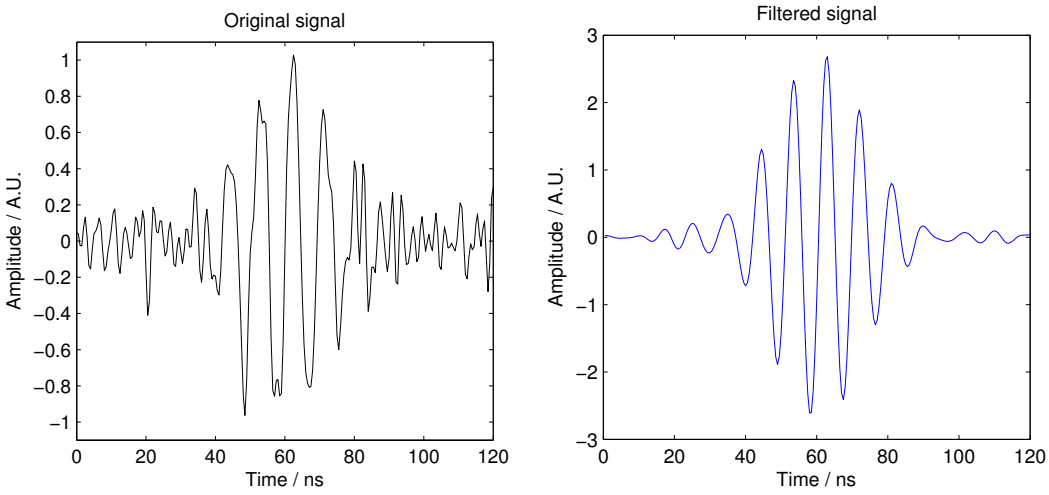


Figure 4.31: A noisy simulated SRAS signal after filtering using an order 8 Morlet wavelet.

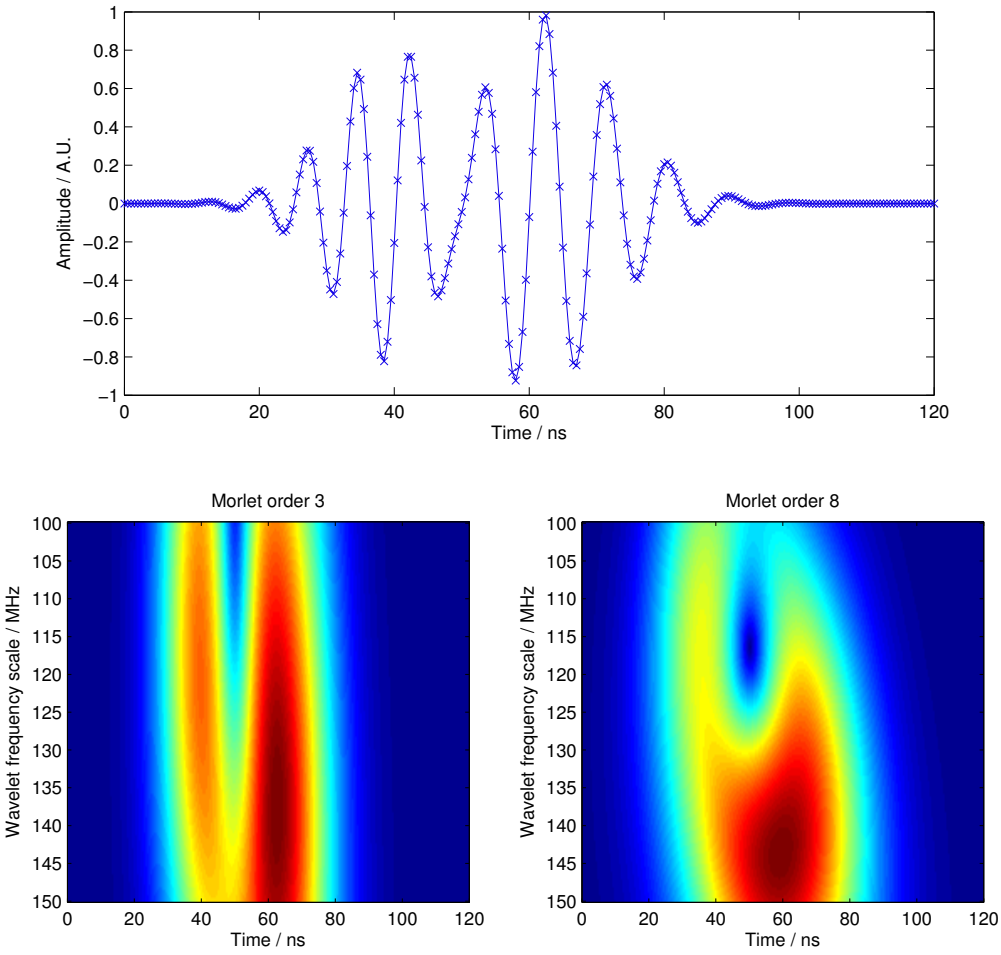


Figure 4.32: A dual frequency simulated SRAS signal fitted with both order 3 and order 8 Morlet wavelets. Order 3 wavelets give much better temporal resolution than order 8 wavelets.

Chapter 5

System performance

Chapter 5 qualitatively and quantitatively examines the overall performance of the Renishaw SRAS instrument. The instrument is composed of multiple independent systems, with each one contributing to the system performance through a diverse set of mechanisms. Each of these systems and mechanisms is isolated and assessed. The motion platform is analysed for flatness, repeatability and stability of motion. A thorough review is undertaken on the factors affecting signal strength including: the generation laser power, SAW acoustic wavelength, apertures and detector type. The spatial resolution of the SRAS technique is determined for multiple acoustic wavelengths and aperture sizes. Repeatability of instrument calibration is assessed and the effects of poor assembly, and variations in manufactured part dimensions, on signal strength are investigated.

5.1 Motion platform performance

The Equator 700 is not designed as a scanning motion platform; it is a gauging motion platform which uses probe head deflection techniques to precisely compare parts to a reference part or CAD model. The motion, and any issues with the motion, need to be characterised and, where possible, corrected. For scanning the ideal scenario is to maintain a constant speed such that the spacing of collected data is consistent while maintaining a constant distance from the sample surface. To achieve this for the line scans used for SRAS, an undershoot and overshoot region are defined to allow for

acceleration of the SRAS head to and from stationary. An acceleration region of 1.5 times the actual distance required for a given acceleration and velocity is programmed into the scan path. SRAS data is collected during these acceleration phases but is typically not plotted. The Equator 700 provides instantaneous measurements of position at a rate of 940 Hz, set by the 7 152 µs servo clocks, and these measurements are recorded with the acoustic data.

In order to determine the stability and flatness of the motion platform a set of scans were taken on a flat sample at two scan velocities. Fig. 5.1 shows the deviation from the expected scan head position in the X, Y and Z axes, deviation measurements are in µm. The scan path is along the x-axis in a positive direction with accelerations of $1500\,\mathrm{mm\,s^{-2}}$ up to a maximum velocity of $82.2\,\mathrm{mm\,s^{-1}}$. Since the mounting of the sample is not perfectly flat, a programmed change in Z height to correct for the slight tilt of the sample is applied. A slightly unwieldy value for velocity is required in order to achieve spacings of 25 µm at the laser repetition rate of 1645 Hz. This repetition rate is set by the 4 152 µs Equator 700 servo clocks used to trigger the laser. A spacing of 25 µm has been used as it is the most typically used spacing in SRAS scans. Both the undershoot and overshoot regions have been cropped to better highlight the scanning errors affecting the region of useful SRAS data.

The positional errors in X and Y are not large, nor critical as these simply alter the location of the plotted pixel in the SRAS images. Clearly the initial acceleration causes the most significant deviations in X and Z and although the scale is fairly small, unexpected displacements in Z have a strong affect on the sharpness of the grating image and the resulting reliability of the measured SRAS signal, this issue is discussed in detail in Chapter 6. The artefacts in X and Z due to acceleration are well spatially correlated but occur several mm after maximum velocity has been achieved. By visual inspection of the SRAS head during scanning shows that these artefacts relate to resonant vibration modes induced during acceleration which are not well damped. In addition to the deviations due to the initial acceleration for each line scan there is also some clustered correlated noise during scans particularly noticeable in Z.

Fig. 5.2 shows a second scan over the same region of the aluminium sample taken at twice the scanning speed, $164\,\mathrm{m\,s^{-1}}$ and with a larger acceleration region, the step size

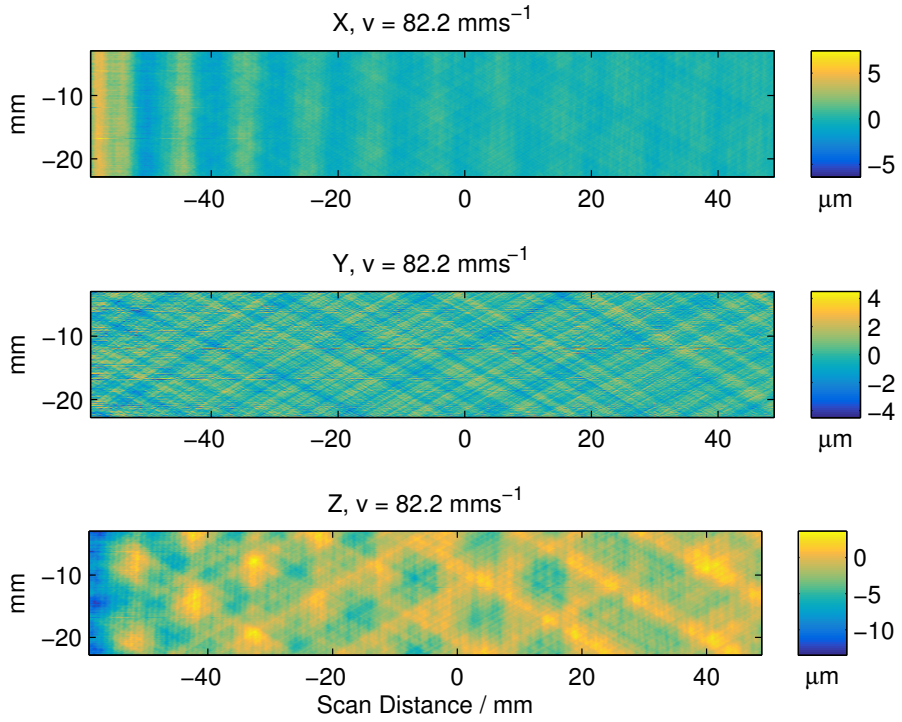


Figure 5.1: The deviation from the expected position of the Equator 700 in X, Y and Z as measured by the Equator position encoders. The colour bar units are µm and the step size for this scan is 25 µm in X and 50 µm in Y. The scan data has been cropped to remove the accelerations and decelerations from stationary. The velocity of the scanning head was set to 82.2 mm s^{-1}.

in X is 100 µm and the step size in Y is 100 µm. A larger undershoot and overshoot region was specified and, as a result, the oscillation due to acceleration has been largely eliminated, however the clustered correlated errors remain very prominent. Comparing the errors at the two velocities shows excellent spatial correlation in Y and Z which indicates that these are not random but are a result of systematic imperfections in the scanning platform. Unfortunately this is not something that can be trivially removed from the motion without delving into the mechanics and fundamental design of the Equator 700.

5.2 Generation power and signal strength

It is known that the fibre coupled Renishaw SRAS system is lacking in generation pulse energy and the result of this is poor SNR. The generation power and pulse energies

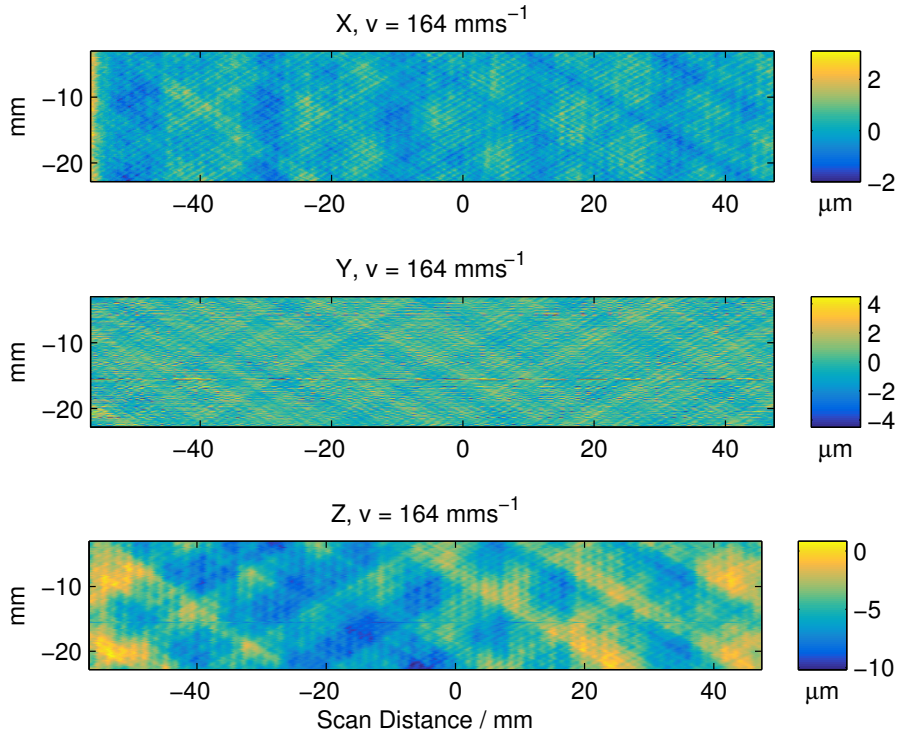
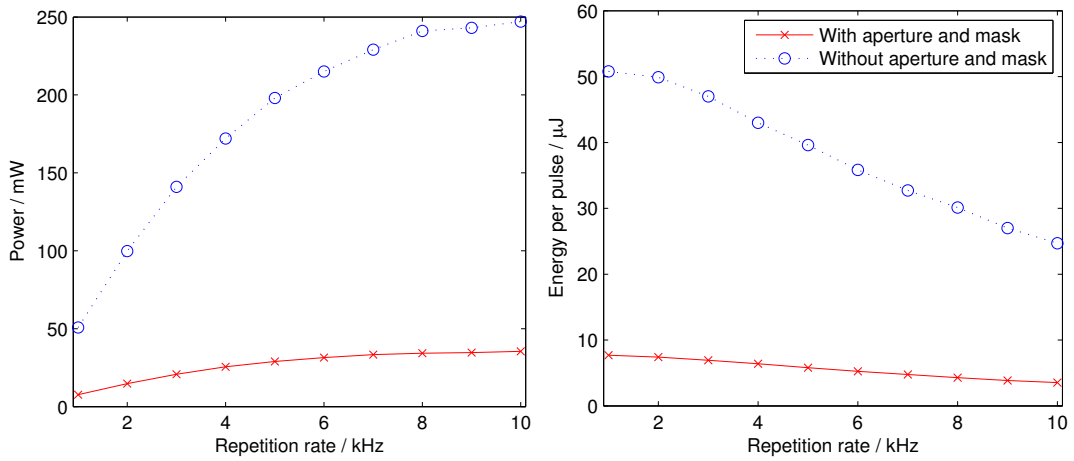


Figure 5.2: The deviation from the expected position of the Equator 700 in X, Y and Z as measured by the Equator position encoders. The colour bar units are µm and the step size for this scan is 25 µm in X and 50 µm in Z. The scan data has been cropped to remove the accelerations and decelerations from stationary. The velocity of the scanning head was set to 164 mm s^{-1}.

are shown in Fig. 5.3 as measured at the sample surface having passed through the full optical assembly although without an aperture or mask. As the repetition rate of the generation laser is increased the output power also increases but at the cost of individual pulse energy which, for SAW generation, is more critical. The decrease in available pulse energy from 1 kHz to 10 kHz is fairly linear, suggesting that this trend will continue beyond repetition rates of 10 kHz.

In order to determine how pulse energy affects signal strength a short scan was repeated on an anisotropic sample over a range of laser repetition rates. The aperture and mask used for measuring pulse energy were consistent at 1.5 mm and 23 µm, and the scan region was 2 mm by 15 µm with 5 µm steps taken in each axis. Each scan was averaged to give a single trace at each repetition rate and this is shown in Fig. 5.4 (a). From this plot there are three relevant observations. Firstly, as the laser repetition rate is

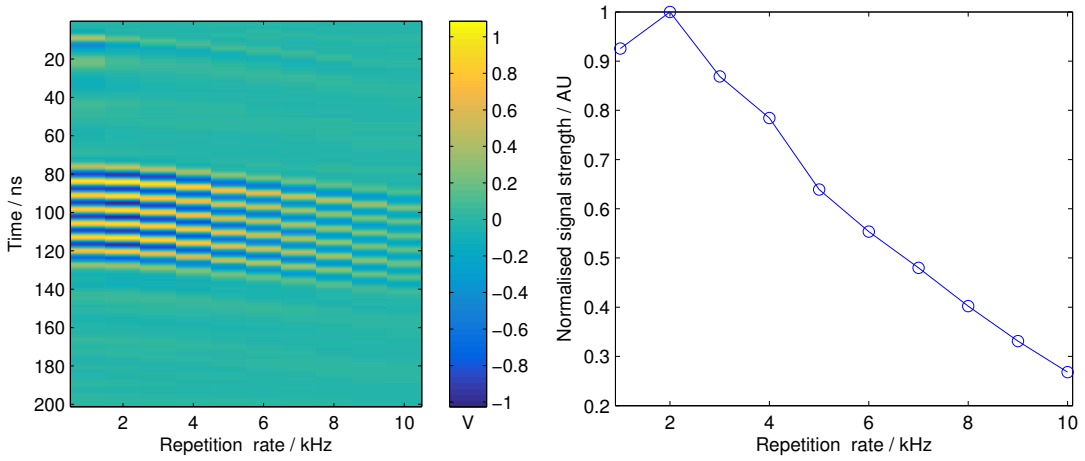


(a) (b)

Figure 5.3: (a) The generation beam power measured at the sample surface as a function of generation laser repetition rate with and without a mask and aperture cartridge. (b) The energy per pulse delivered to the sample surface as a function of repetition rate.

increased the signal amplitude is decreased; this is a result of the decreasing laser power. Secondly, the form of the signal remains the same at each repetition rate indicating that this is determined only by the aperture and grating configuration. Lastly there is a signal delay of up to 20 ns as the repetition rate increases from 1 kHz to 10 kHz. Although there is a small change in pulse width from 2.8 ns to 3.2 ns, as measured by the laser manufacturer, this is too small to explain the observed delay. Therefore the most likely cause of this delay is the electronic control of the generation laser.

To directly compare amplitude, an FFT of all of the traces at each repetition rate was taken and the median amplitude calculated and plotted in Fig. 5.4 (b), the amplitudes were normalised for clarity. This figure shows that the decrease in signal amplitude occurs at a greater rate than the decrease in laser energy. Ignoring the slightly anomalous result at 1 kHz, the decrease in energy from 2 kHz to 10 kHz is 50% whilst the detected signal amplitude decreases by 70%. The SRAS generation and detection system is fairly complex and so there are multiple potential sources of non-linearity. The KED is sensitive to the gradient of the SAW which is linear with amplitude. The measured signal amplitude is determined by an FFT of the measured signal which is also a linear process with amplitude. Since the generated SAWs have the same frequency content over all repetition rates any, non-linearity in the signal chain should be negligible. It is known that, up to the ablation threshold, optical energy is



(a) (b)

Figure 5.4: (a) A summed set of 1400 traces measured on an isotropic aluminium sample over a range of generation laser repetition rates.(b) The mean signal strength taken from the absolute of Fourier transform of the traces shown in (a) as a function of repetition rate. These data have been normalised to the maximum amplitude.

linearly proportional to generated acoustic wave amplitude [5], however amplitude can be affected by the small change in pulse width as repetition rate is increased.

Following the isotropic model for generation of SAWs by Achenbach, [91, 92], it is possible to predict the SAW displacement at a given generation laser pulse width. This model has been applied and the normalised maximum SAW displacements are plotted in Fig. 5.5. The effect of pulse temporal width on displacement is surprisingly potent and, in combination with the decrease in pulse energy, explains the drop in signal amplitude with repetition rate. Clearly this is an issue for scanning at higher repetition rates especially on noisy samples and sets an upper limit on scanning speeds that must be considered on a sample by sample basis. In addition this effect highlights the importance of the choice of generation laser for making good SRAS measurements.

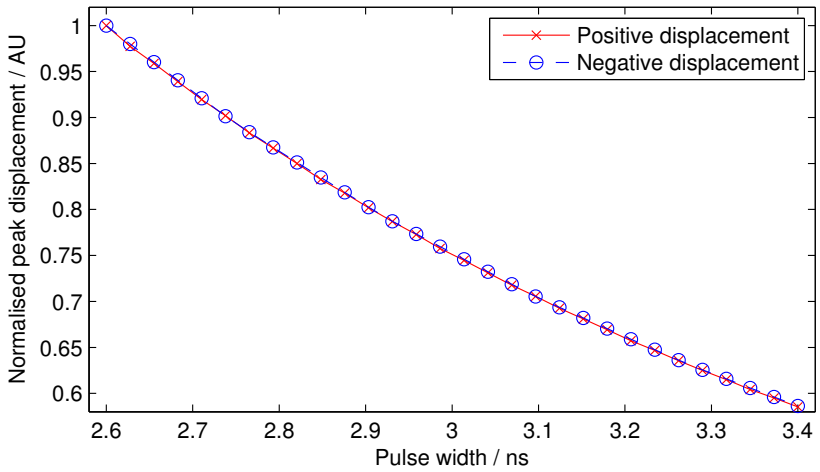


Figure 5.5: The normalised theoretical maximum positive and negative displacements of a laser generated SAW plotted against the generation laser temporal width. This plot has been generated using the method described in [91].

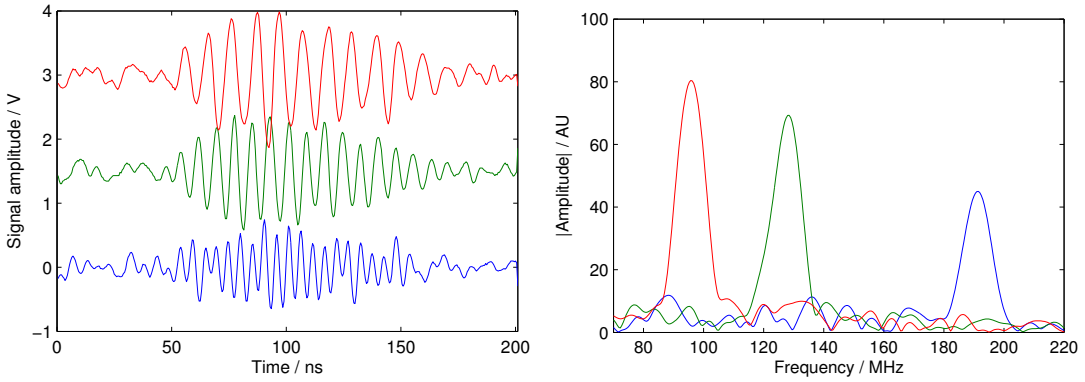
5.3 Effect of acoustic wavelength on signal strength

The generation patch size and acoustic wavelength can be set within a range dependent on the available chrome glass masks, apertures and objective lenses. For the Renishaw SRAS system gratings which give acoustic wavelengths of approximately 15 µm, 23 µm, and 30 µm are available along with apertures of 750 µm, 1000 µm, and 1500 µm.

On a standard metallic sample typically velocities are up to 3500 m s^{-1}, for the gratings used this gives a maximum frequency of ~230 MHz. Based on Fig. 2.3 and the <3 ns pulse length of the generation laser, there is only a small drop off in available power from the laser at this upper frequency and the effect on signal strength will be minimal. Similarly the signal chain has components with bandwidths of 500 MHz, except for a low-pass filter which can be altered as required, and will therefore not affect signal strength either. At higher frequencies the gradient of SAWs is steeper causing greater deflection of the detection beam and therefore increasing the sensitivity of the KED. The KED photodiodes have finite capacitance, in the range 1.5 pF to 4 pF, which limits the bandwidth of the detector to approximately 330 MHz, well above the maximum frequency at which SRAS is being applied. As a result of the above points the detected SAW amplitudes as a function of acoustic wavelength - or frequency - can be considered to be dominated by the generation process rather than the detection process.

In order to determine the effect of acoustic wavelength on signal strength an empirical approach has been taken. Fig. 5.6 (a) shows a trio of SRAS signals generated at the same point on an isotropic aluminium sample with identical configuration parameters apart from a change in generation grating. Masks with chrome spacings of 80 µm, 120 µm and 160 µm were used without apertures and with constant generation laser parameters to ensure that the generation energy was consistent. The figure shows each signal has a similar window, determined by the constant generation beam diameter, but different frequencies and maximum amplitudes with the amplitude decreasing with frequency. Fig. 5.6 (b) shows the FFT of the signals and further highlights the decreasing signal amplitude with frequency.

To further highlight the differences between acoustic wavelengths, a small scan was undertaken for each grating on the same isotropic aluminium sample. The maximum magnitude of the FFT of the signals are plotted in Fig. 5.7. These amplitude plots



(a) (b)

Figure 5.6: (a) A trio of SRAS signals generated on an isotropic aluminium sample with various generation gratings and no apertures. The grating spacings are, from top to bottom, 160 µm, 120 µm and 80 µm. (b) The FFT of each of the signals shown in (a) with a zero-pad of 2^{12}.

have been normalised to the maximum amplitudes from the longest wavelength grating. Calculating the median values of each plot gives 0.45, 0.67 and 0.72 for the 15 µm, 23 µm and 30 µm gratings respectively confirming the findings from Fig. 5.6.

The cause of the decrease in signal amplitude with signal frequency is most-likely related to the sharpness of the imaged patch. The Rayleigh criterion states that $\delta l = \frac{1.220 f \lambda}{D}$, where δl is the spatial resolution, f is focal length and D is the lens aperture. For the objective lens $f = 20\,\text{mm}$ and $D = 24.6\,\text{mm}$, thus the spatial resolution is ~1.1 µm on the sample. For a 15 µm acoustic wavelength the width of a single illuminated line is 7.5 µm and therefore approximately 30% of this line is poorly resolved, compared to ~20% and 14% for the 23 µm and 30 µm masks respectively. The end result is less efficient generation of SAWs for finer grating spacings.

At higher frequencies the greater number of cycles should yield a sharper peak in frequency. However a secondary effect of the reduction in imaged line sharpness is an increase in the bandwidth of the generated SAW packet. This effect can be seen in Fig. 5.6 (b) where there is no sharpening of frequency peaks for higher frequency signals.

For the 15 µm acoustic wavelength there may be a secondary effect due to the detection method. If the detection spot size is larger than 7.5 µm then as the 15 µm wavelength SAW passes under it the sinusoidal amplitude will be slightly averaged, reducing the sensitivity of the detector.

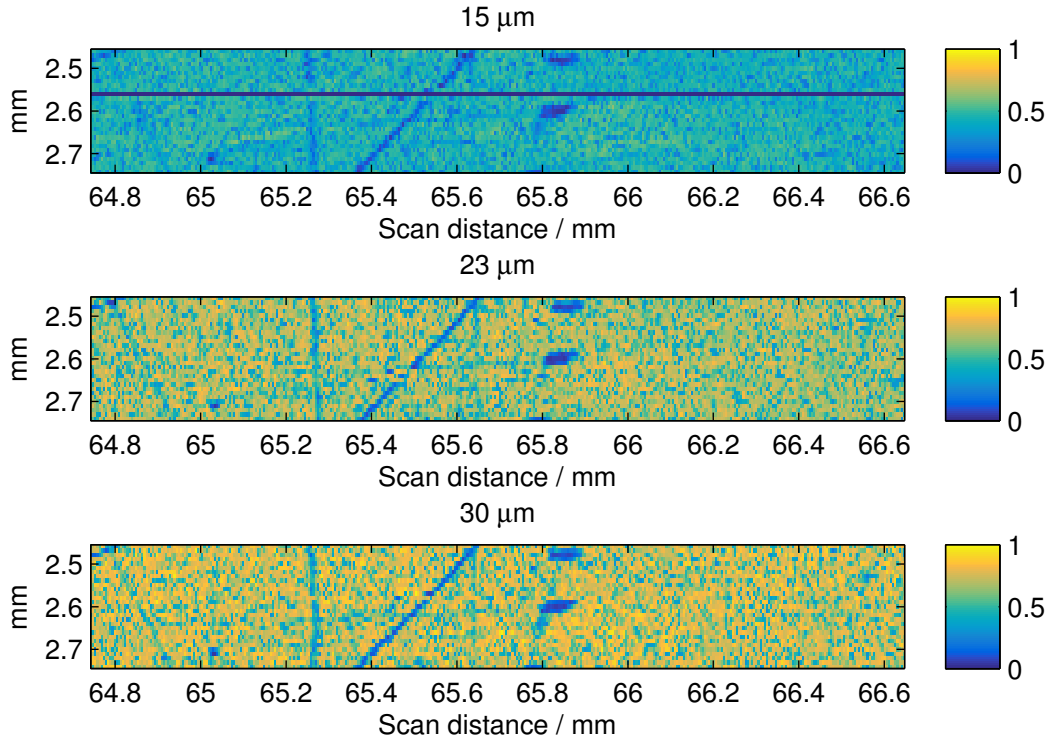
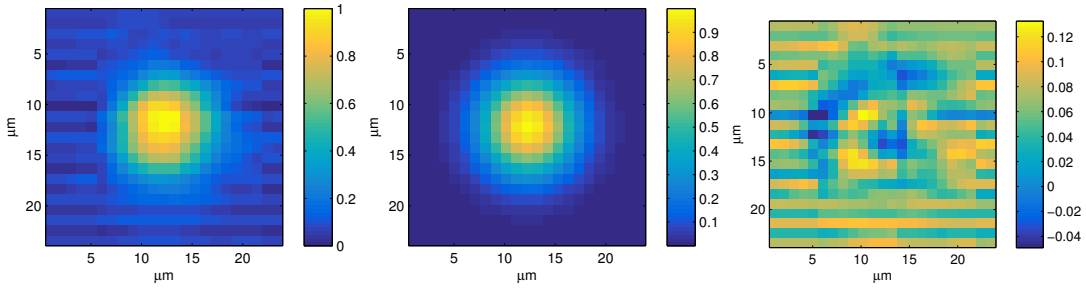


Figure 5.7: The signal amplitudes from a set of scans taken with gratings of acoustic wavelength 15 µm, 23 µm and 30 µm. Scans were taken on an isotropic aluminium sample with step sizes of 5 µm in X and 10 µm in Y. The signal amplitudes have been normalised to each other and plotted using the same colour bar scale.

5.4 Detection beam size

As noted in the previous section, it is important that the detection beam spot size is smaller then half the acoustic wavelength to avoid averaging over a SAW cycle. Due to the high focussed power and small spot size, an absolute measurement of the detection beam spot size on the sample surface is difficult to make. Fig. 5.8 (a) shows the detection beam profile imaged using the camera. Comparing the imaged beam to a fitted 2D Gaussian, shown in Fig. 5.8 (b) and (c), shows excellent agreement which indicates that the detection beam is passing through the fibre optic and optical system as expected. However the imaged beam size is heavily dependent on the arbitrary attenuation of the beam and camera exposure settings, meaning an estimate is difficult to make from the camera image. For this reason the 12 µm estimate made from the camera image can only be taken to be an upper limit on the true detection spot size.

A lower bound estimate of spot size can be found from calculation of the theoretical



(a) (b) (c)

Figure 5.8: a) A normalised image of the detection spot captured by the SRAS mounted camera. (b) A 2D Gaussian fit of the imaged beam. (c) Residual errors between the imaged and fitted Gaussian spots. Using the $1/e^2$ estimate, the beam spot size is 12 µm.

ideal case. For a focussed beam an estimate of the spot size, w_0, can be taken as [93]

$$2w_0 = \frac{4\lambda F}{\pi D}, \tag{5.1}$$

where λ is the laser wavelength, F is the objective lens focal length and D is the illuminated beam diameter on the lens. Applying the correct values gives, with $D =$ 2.1 mm, a focussed beam diameter of 3.2 µm. Beam divergence of the collimating optic is, at worse, 0.03° over the full beam path and this increases the focussed beam diameter to 3.5 µm. The optical path for delivery of the detection beam to the sample is quite simple and only the objective lens actively alters the beam size. Given that the Mitutoyo objective is optimised for the focussing of 532 nm light, it is not unreasonable to assume that the divergence from the ideal case is fairly small. Even assuming that the final beam diameter is 2 times worse than the ideal case, this is still less than half of any of the acoustic wavelengths used.

5.5 Spatial resolution

The spatial resolution of the SRAS technique is a critical factor for the overall usefulness of the technology. A poor spatial resolution means that fine-grained materials and samples do not exhibit SAW phase velocity contrast in the results, as the SRAS measurement will then yield only an average velocity over multiple grains. Determining absolute resolution of the SRAS technique is a non-trivial task; the spatial resolution is determined by a combination of the generation patch size, acoustic wavelength, propa-

gation direction and detection spot size. Fig. 5.9 indicates the contributing factors to the spatial resolution relative to the generation patch size in the propagation direction, L_x, and in the orthogonal direction, L_y.

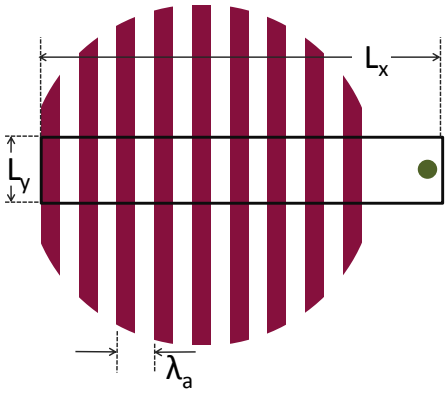


Figure 5.9: A schematic indicating the contributing factors to the SRAS spatial resolution. Resolution is different in orthogonal directions due to the geometries of the generation patch and detection spot. The rectangle L_x, L_y indicates the active generation and detection region.

5.5.1 Orthogonal spatial resolution

The method adopted for determining typical SRAS spatial resolution is by empirical observations and measurements on samples with a fine-grain microstructure, typically 10 µm. Two finely stepped scans, 5 µm in both X and Y, were taken on a sample of stainless steel in orthogonal directions using an acoustic wavelength of 23 µm and an aperture of 1 mm. Due to the method of manufacture of the sample - rolling - the microstructure consists of long and thin grains orientated parallel to the rolling direction. This structure is useful as it will highlight the anisotropic resolution of the SRAS technique in orthogonal directions. As an aside, this sample was taken from the Equator 700 struts which were failing due to the alignment of the grains relative to the sheet metal folding direction.

Fig. 5.10 shows the two orthogonal SRAS images. The sample was rotated with the SRAS orientation and propagation direction kept constant, as such the grain structure is aligned along the x-axis for the 0° scan, and along the y-axis for the 90° scan. The difference between the two scans is immediately clear; the narrow grains are well spatially resolved in the 0° scan and are not at all well resolved in the 90° scan. While useful analysis is possible from the 0° image, it would be difficult to determine in which

direction the grains are actually aligned or even that the fine-grain structure exists from the 90° image. The cause of this is down to the shape of the grain and the mechanism of generation and detection. When a SAW is generated by the illuminated patch only the SAWs generated in the central region of the patch, highlighted in Fig. 5.9, are detected. Therefore when the grain is aligned with the highlighted rectangle the SAW packet contains only the phase velocity of the grain and very good lateral resolution can be achieved. When the orthogonal grain structure is imaged the generation patch sits over multiple grains and therefore the detected SAW packet has content from each of the contributing grains and the determined phase velocity is an average of these contributions, as such a poor spatial resolution is achieved.

For a more direct comparison the two orthogonal scans were aligned using the optical image and a sub-region chosen for plotting, this is shown in Fig. 5.11 where the 90° scan has been rotated such that the grain structure in both SAW phase velocity maps is aligned. From these scans an estimate for the spatial resolution can be made by measuring the minimum resolvable width and lengths of the grains. Doing this gives $L_y \leq 10\,\mu\text{m}$ and $L_x \leq 30\,\mu\text{m}$ for this particular scan where the SRAS configuration used a 23 µm acoustic wavelength and a 1 mm aperture. The reason that L_x is significantly smaller than the patch size, 200 µm, is that the measured SAW velocity over a small-grain region is a weighted average over each of the grains under the patch. Where the grains under the patch are larger than average, or have very high or low SAW velocity, the average phase velocity is weighted towards the true value of these grains, resulting in enhanced velocity contrast.

It is important to note here that for this special case of elongated grains the lateral spatial resolution is significantly better than the spatial resolution in the propagation direction. However this only occurs because the grains are long enough to fit both the detection spot and a sufficient proportion of the generation patch to generate a SAW packet with mostly the local grain phase velocity. For the case of small circular grains this is not reproducible and the lateral spatial resolution will increase. A second point to note is that the stainless steel sample used has fairly high anisotropy; the velocity range is $\sim 480\,\text{m}\,\text{s}^{-1}$. The result is good velocity contrast between grains which effectively improves resolution. However for low anisotropy samples, such as aluminium, velocity contrast is poorer resulting in a larger spatial resolution.

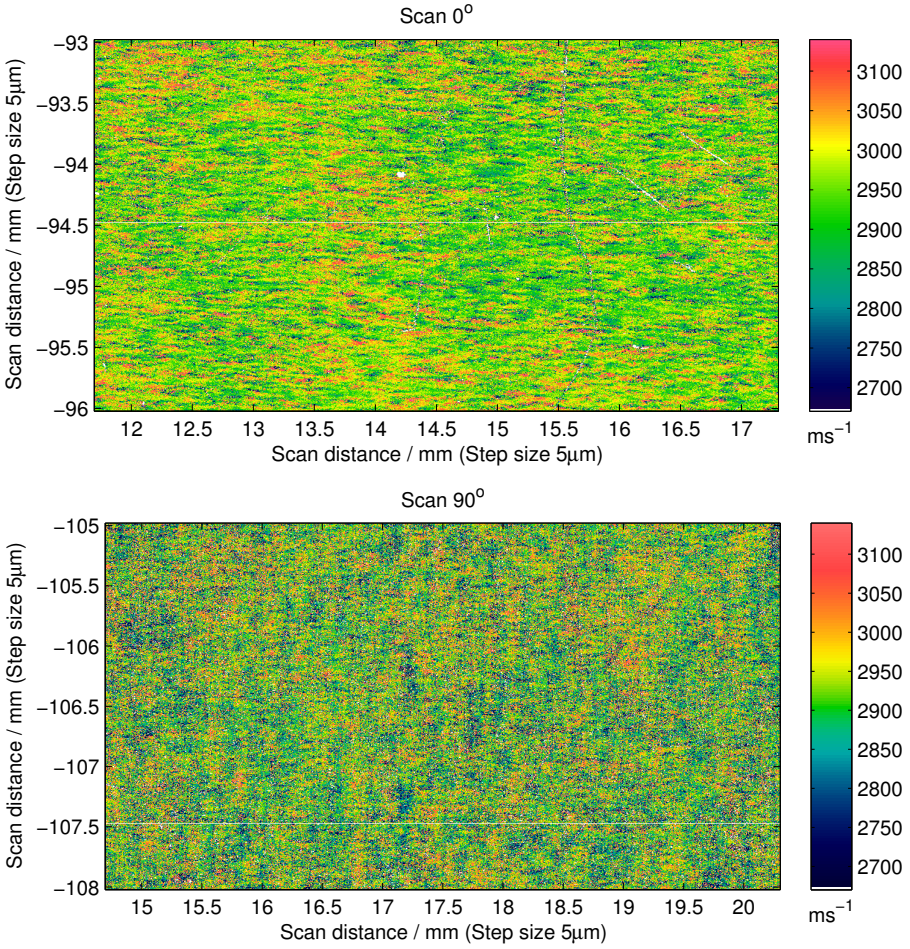


Figure 5.10: Two SRAS scans of a rolled stainless steel sample taken in orthogonal directions. 0° is taken to be parallel to the positive x-axis. The step size was 5 µm in both X and Y, and an acoustic wavelength of 23 µm and an aperture of 1 mm were used. The sample has been rotated under the SRAS instrument such that for both images the propagation direction is along the x-axis.

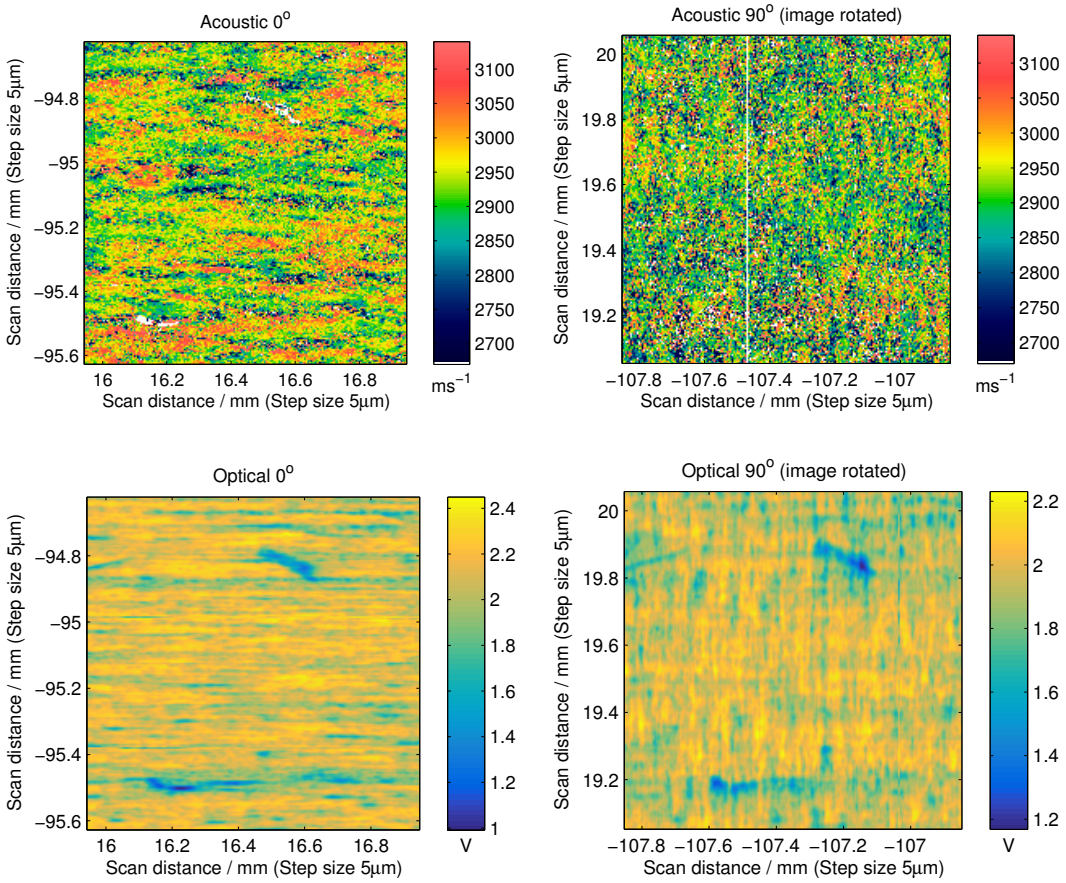


Figure 5.11: Aligned sub-regions of two SRAS scans of a rolled stainless steel sample taken in orthogonal directions. 0° is taken to be parallel to the positive x-axis. The step size was 5 µm in both X and Y, and an acoustic wavelength of 23 µm and an aperture of 1 mm were used. 90° scan has been rotated such that the grain structure in both SAW phase velocity maps is aligned. The optical images used for aligning the acoustic data are also shown for both sub-regions.

5.5.2 Effect of acoustic wavelength on spatial resolution

SRAS allows different gratings to be inserted altering the acoustic wavelength of the generated SAWs. In theory the chosen wavelength will affect the spatial resolution of the scan at length scales similar to the wavelength. For grains smaller than a wavelength the SAW velocity will be an average over the local grains, whereas for grains larger than the wavelength the SAW velocities are determined only by the local grain. The previous section indicated that the spatial resolution achievable on a fine-grained sample is similar to the acoustic wavelength. To further test the effect of acoustic wavelength on spatial resolution three SRAS scans were taken on a large grain Ti 685 sample with the three available gratings. Step size was 10 µm along the x-axis and 20 µm along the y-axis, propagation was in the positive x direction. To decouple the effect of apertures, no apertures were used, although due to the generation beam focus on the sample surface the effective generation patch size on the sample surface was ~350 µm, equivalent to a 1.8 mm aperture. The three scans are shown in Fig. 5.12.

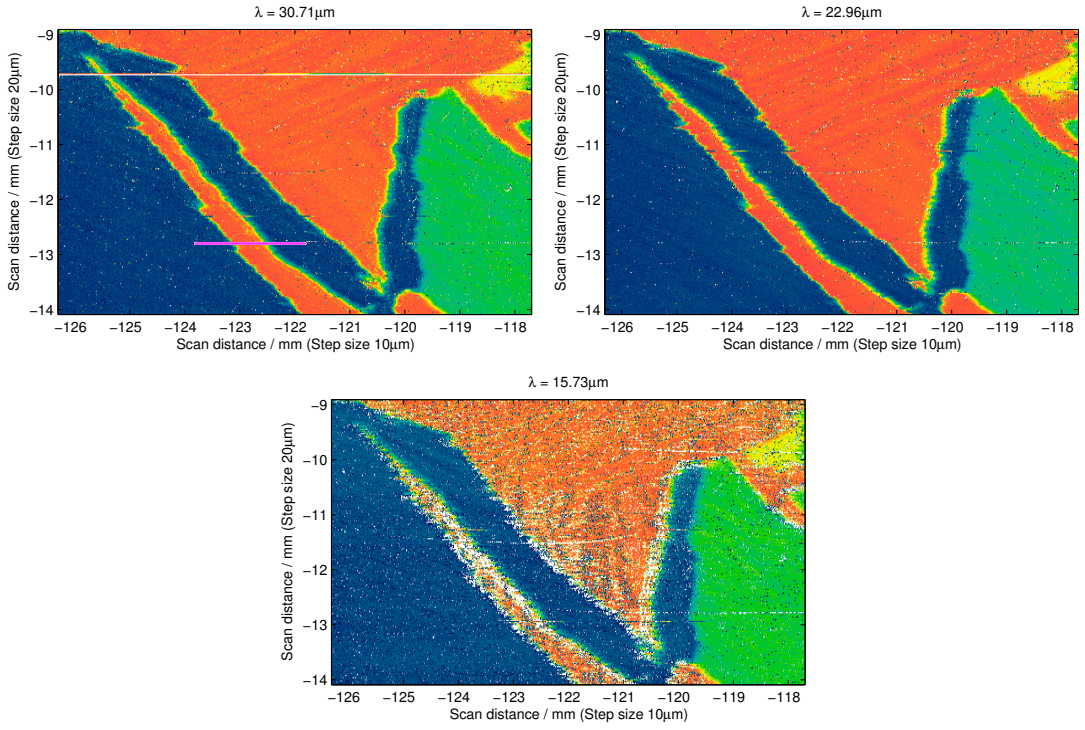


Figure 5.12: Three SRAS scans taken on a large grain Ti 685 sample at three acoustic wavelengths and without apertures. Step size was 10 µm along the x-axis and 20 µm along the y-axis, propagation was in the positive X direction.

The noise on the 15 µm acoustic wavelength scan shows poorer SNR than the two other scans which is a reflection of the findings in the previous sections of this chapter relating to efficiency of SAW generation, this issue is particularly pronounced due to the poor surface condition of the titanium sample. Looking at the edges of grains and the transitions between grain velocities shows very little difference between the three acoustic wavelengths used. This is because the velocity determined at grain boundaries is a weighted average dependent on the location of the generation patch which is the same size at each wavelength. A slice of the thin, fast grain on the left of the image is shown in Fig. 5.13 which confirms the fairly consistent size of the grain boundary regions.

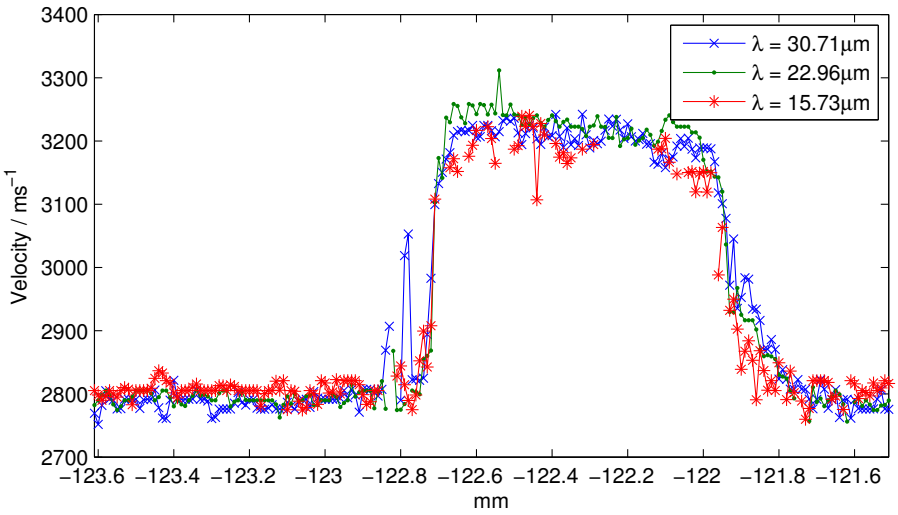


Figure 5.13: A slice through the thin red grain shown in Fig. 5.12 at y = −13.28 mm indicating the feature width at three acoustic wavelengths. Due to the noise in the 15 µm scan some points have been omitted from the plot.

There is some asymmetry in the size of the grain boundary transition regions relative to the propagation direction of the SAWs which is particularly noticeable on the transitions between the slow blue grains and fast red-orange grains. In all images the SAWs are propagating from right to left, which is also the scan direction, with the detection spot ahead of the generation patch. To explain this there are two possibilities: either the SAWs generated on the slower blue regions are not efficiently transmitted across the boundary to the detection spot, or the SAWs generated on the slow region are fairly dominant such that only a small portion of the patch has to be over the slow region to affect the velocity determination. Examining the raw data shows very low

amplitude signals at the boundary region indicating that the former explanation is correct.

Besides noisiness, the primary difference between the images is in how visible the intra-granular features are. Focussing on the large blue region on the left of each image, there are some indications of faint thin features running from top-left to bottom-right. The contrast of these features, relative to the background, increases as the acoustic wavelength decreases for two reasons. First is the increased sharpness of the SAW frequency peak due to the increased number of cycles at smaller wavelengths, improving how well the velocity can be determined; Second is the expected interaction of the SAW packet with the very fine grain structure which is of similar size to the acoustic wavelengths used.

5.5.3 Effect of aperture on spatial resolution

Adding apertures to the generation beam limits the patch size and therefore determines the length and number of cycles in the generated SAW packet. Apertures do not concentrate the generation energy but block it from reaching the sample, therefore limiting the energy available in the SAW packet and reducing SNR. Fig. 5.14 shows the decrease in pulse energy with shrinking apertures and compares the theoretical energy loss, based on a Gaussian beam, with the measured values. There is some disagreement between the expected measured energy values because the beam is not a perfect 2D Gaussian due to the large core multi-mode fibre, instead the beam is closer to a Tukey or top-hat function. In addition the apertures are not perfectly centred to the beam, for the larger apertures this is non-critical but for the smaller apertures the loss of energy becomes more pronounced.

The advantage of introducing an aperture is the improvement in spatial resolution. In order to quantify this improvement a set of scans have been taken using various aperture sizes. Four scans were taken over the same sub-region of the Ti 685 sample used in the previous section. An acoustic wavelength of 23 µm was used and step sizes of 10 µm along the x-axis and 20 µm along the y-axis were taken with SAWs propagating in the positive X direction. Fig. 5.15 shows these four scans with circular apertures of values 2 mm, 1.5 mm, 1 mm and 0.75 mm. Given the ~5X demagnification

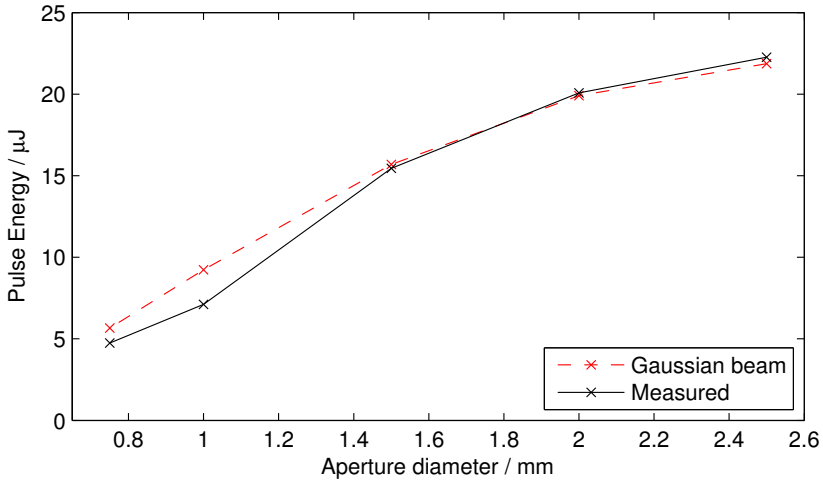


Figure 5.14: Generation beam pulse energy available at the sample surface with a variety of apertures used compared to the theoretical change in pulse energy based on an aperture applied to a perfect Gaussian beam.

of the system this corresponds to patch sizes of 400 µm, 300 µm, 200 µm and 150 µm.

As expected the use of apertures has adversely impacted the SNR, and the image captured with a 0.75 mm patch size is the most noisy. The differences between the scans with a 1.5 mm and 2 mm aperture are quite small, this is because the natural aperture of the generation beam is, at 1.8 mm, smaller than the largest aperture used and not much larger than the 1.5 mm aperture. Comparing these two scans with the scan at 23 µm with no aperture in Fig. 5.12 also shows minimal differences. By visual inspection of the images it is clear that the spatial resolution has improved with smaller apertures. This can be seen in the grain boundary transition regions which increase in sharpness, and also in the small feature at (−120.5 mm, −13.5 mm) which becomes much better separated from the nearby larger grain.

To quantify this improvement a sub-region was chosen where feature widths could be measured for each scan, this region is indicated by the black rectangle in Fig. 5.15. For this region a horizontal slice was taken from each of the four scans and a threshold was set and used to measure the total feature width. This was repeated at multiple steps along the y-axis, and the resultant widths are plotted for each aperture in Fig. 5.16.

In agreement with the initial observations there is very little difference between use of a 2 mm and 1.5 mm aperture; the median difference between the feature widths measured

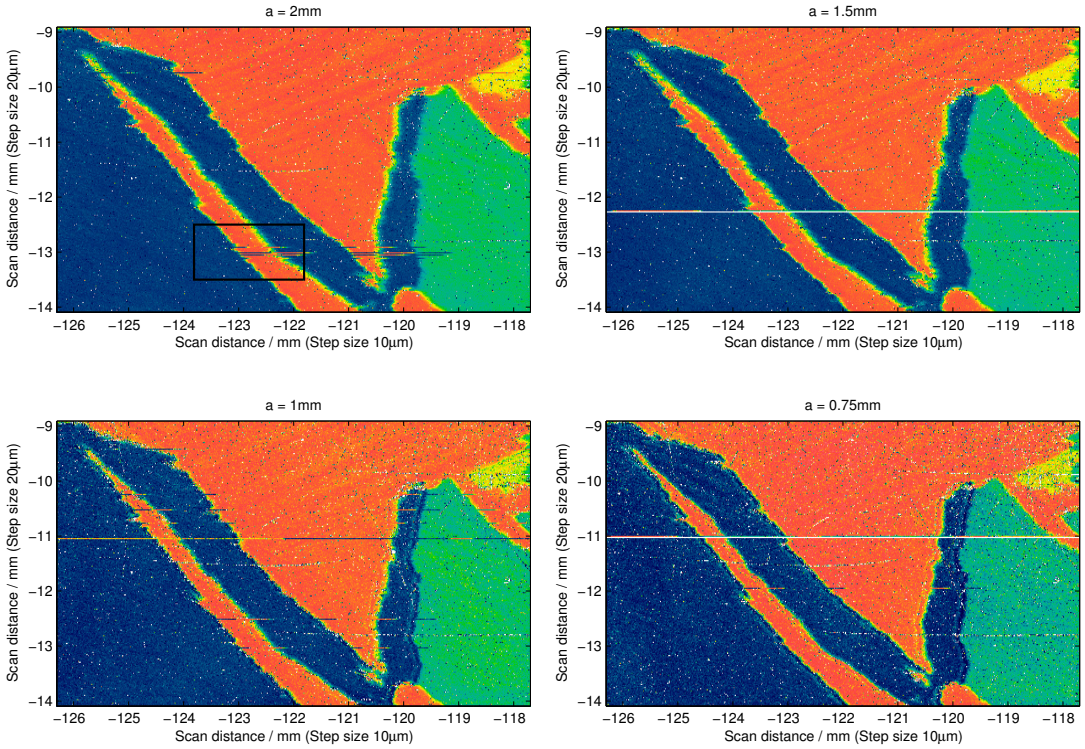


Figure 5.15: Four SRAS scans taken over the same sub-region of the Ti 685 sample with various apertures used. An acoustic wavelength of 23 µm was used and step sizes of 10 µm along the x-axis and 20 µm along the y-axis were taken with SAWs propagating in the positive x direction. The colour map is the same for each image.

for these two scans is 0. For smaller apertures the measured feature width decreases as the smaller patch straddles the grain boundaries for fewer pixels, resulting in a sharper transition between grain velocity measurements. Comparing the curves in Fig. 5.16 shows a decrease in feature width from the 1.5 mm aperture to the 1 mm aperture of 70 µm. From the 1 mm aperture to the 0.75 mm there is a 30 µm decrease in feature width. Since the feature width measurement is affected by both sides of the feature the effective improvement in spatial resolution can be estimated to be half of these values. This improvement scales almost linearly with the decrease in generation patch size, 100 µm and 150 µm respectively, indicating that the aperture is the dominant factor affecting spatial resolution in SRAS measurements.

To explore the influence of samples with lower anisotropy, the aperture experiment was repeated on a large-grained sample of aluminium using the same scan parameters and selection of apertures. Slightly finer step sizes of 5 µm in x and 10 µm in y were used

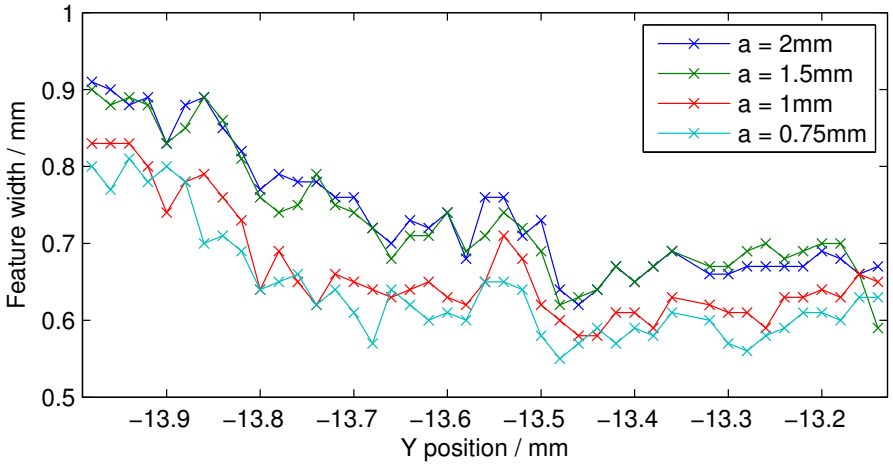


Figure 5.16: Measurements of the width of the feature highlighted by the black rectangle in Fig. 5.15 at several y-values and for each aperture. Due to the shrinking of the grain boundary transition region the smallest apertures give the smallest measurement of feature width.

for these scans which are shown in Fig. 5.17. As with the titanium sample the decrease in SNR with smaller apertures is fairly clear although, due to the low velocity contrast on the low anisotropy sample, the improvement in sharpness at grain boundaries is less evident. One marker of improved spatial resolution is the feature at (−121.5 mm, −13.5 mm) which only becomes evident at apertures of less than 1.5 mm diameter.

The same feature width analysis used on the titanium sample was applied to the blue feature at the centre of the scans along the horizontal axis and the result of this is shown in Fig. 5.18. Due to the low anisotropy, grain velocity differences are relatively small. Because of this feature width measurements were more difficult to make than on the titanium sample, especially for the smallest aperture. Large spikes in the feature width curves caused by noise have been omitted from the plot to improve clarity. As seen with the titanium sample, smaller apertures give a sharper image and a smaller feature size. The difference between the 2 mm and 1.5 mm aperture is again close to 0. Comparing the 1.5 mm aperture to the 1 mm aperture shows a decrease in feature width of 70 µm, and comparing the 1 mm aperture to the 0.75 mm aperture shows a decrease in feature width of 20 µm. These numbers are broadly similar to the values found for titanium and any discrepancy can be explained by the noisiness of the data used for width measurement.

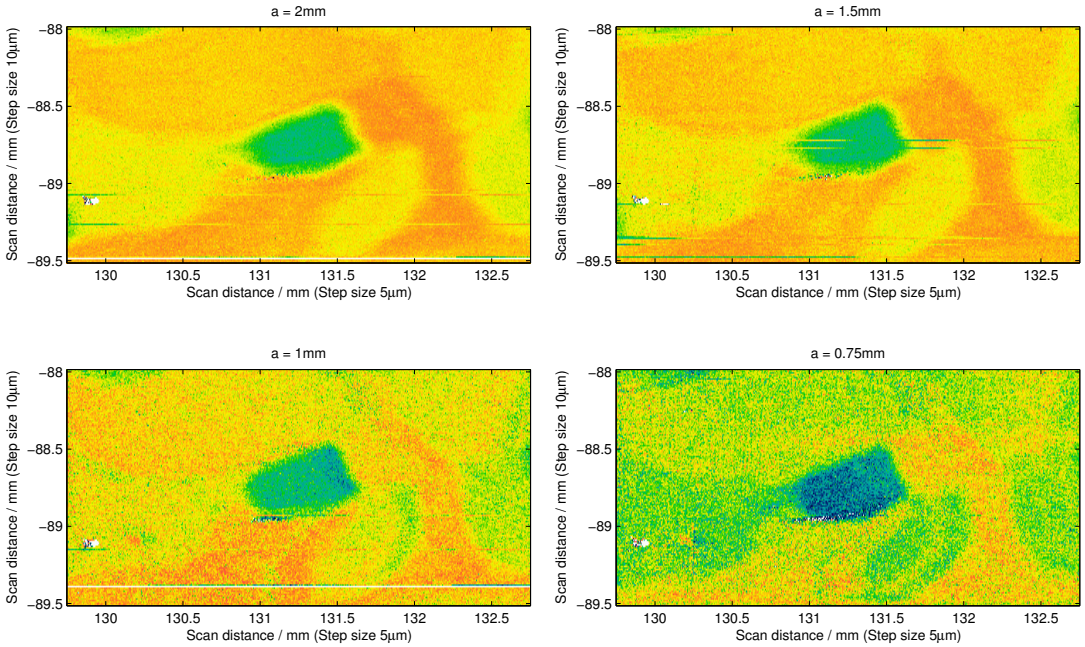


Figure 5.17: Four SRAS scans taken over the same sub-region of a large grained aluminium sample with various apertures used. An acoustic wavelength of 23 µm was used and step sizes of 5 µm along the x-axis and 10 µm along the y-axis were taken with SAWs propagating in the positive x direction. The colour map is the same for each image.

5.6 Repeatability of acoustic wavelength calibration

To calibrate the acoustic wavelength of the inserted grating, a reference sample of isotropic aluminium is used for which the SAW phase velocity is known to be 2933 $\mathrm{m\,s^{-1}}$. By sampling over multiple points on the aluminium surface and averaging the measured frequencies, the instrument acoustic wavelength can be determined. For flexibility the gratings are mounted into cartridges which are slotted into a kinematic mount in the head and this raises two potential issues: is the acoustic wavelength the same after removing and replacing a grating, and is the acoustic wavelength variant over different gratings from the same batch due to variations in mounting?

In order to answer these questions a simple trial was devised whereby multiple gratings were swapped in and out and the calibration procedure repeated. The same mounted grating was scanned 5 times, in between each scan the grating was removed and reinserted. Two alternate gratings of the same spacing were then each scanned twice, with the gratings reinserted between each scan. No aperture was used in order to optimise

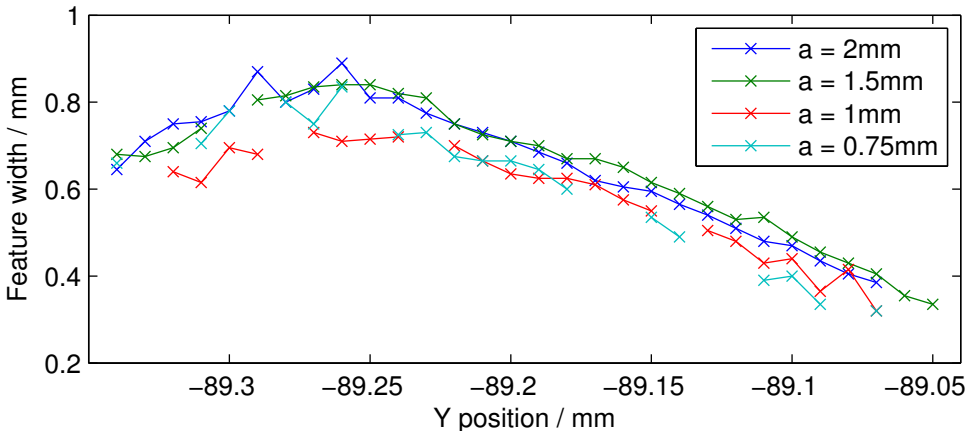


Figure 5.18: Measurements of the width of the blue/green feature at the centre of each of the scans shown in Fig. 5.17 at several y-values and for each aperture. Due to the shrinking of the grain boundary transition region the smallest apertures give the smallest measurement of feature width. For the small apertures the velocity data is quite noisy and so anomalous points have been omitted from the figure.

frequency resolution, and the scan parameters were unchanged for each grating such that this experiment tests only the repeatability of the kinematic mount and similarity of masks. A low speed scan was used over a very small region, 2 mm by 0.5 mm, in order to avoid the issues seen with the motion platform in §5.1, steps were 5 µm in the X direction and 10 µm in the Y direction. To determine acoustic wavelength the frequency data were thresholded, based on the optical and acoustic amplitude maps, and the median value taken along with the standard deviation. This was repeated for three gratings with chrome spacings of: 80 µm, 120 µm and 160 µm. Based on the optical configuration the expected acoustic wavelength is one fifth of these values, and these data are shown in Fig. 5.19.

One of the striking features of the figure is the consistent difference between the MFP and wavelet determinations of frequency. Based on the simulations performed in Chapter 4 the expectation is for the wavelet analysis to be more reliable and this is shown in the standard deviation measurement where the wavelet analysed acoustic wavelength values are in all cases more consistent than the MFP values. In reality, as long as the choice of method is consistent, it is not critical whether wavelets or MFP are used. For all of the grating spacings the first mask tested, samples 1 to 5, the acoustic wavelength is fairly consistent across all repeats. The range of calibrated values is less than 0.15%

in all cases. This suggests that re-calibration when the same mounted grating has been removed and reinserted is unnecessary as the kinematic mount provides very repeatable positioning. For the second and third sets of gratings agreement in measured acoustic wavelength is less consistent. These sets of gratings were not fully-bonded into the mask holder and this is the most-likely cause of the variance. Therefore, during manufacture it is essential that glass gratings are well bonded into their holders, and that the system is re-calibrated for each new grating used.

Comparing the measured acoustic wavelengths and physical grating sizes indicates that the imaging system has a fairly consistent magnification factor of 5.2X. This is slightly larger than the expected 5X and is simply result of slight variances in lens focal lengths away from the stated f100 and f20 values and the asymmetric back and forward focal lengths.

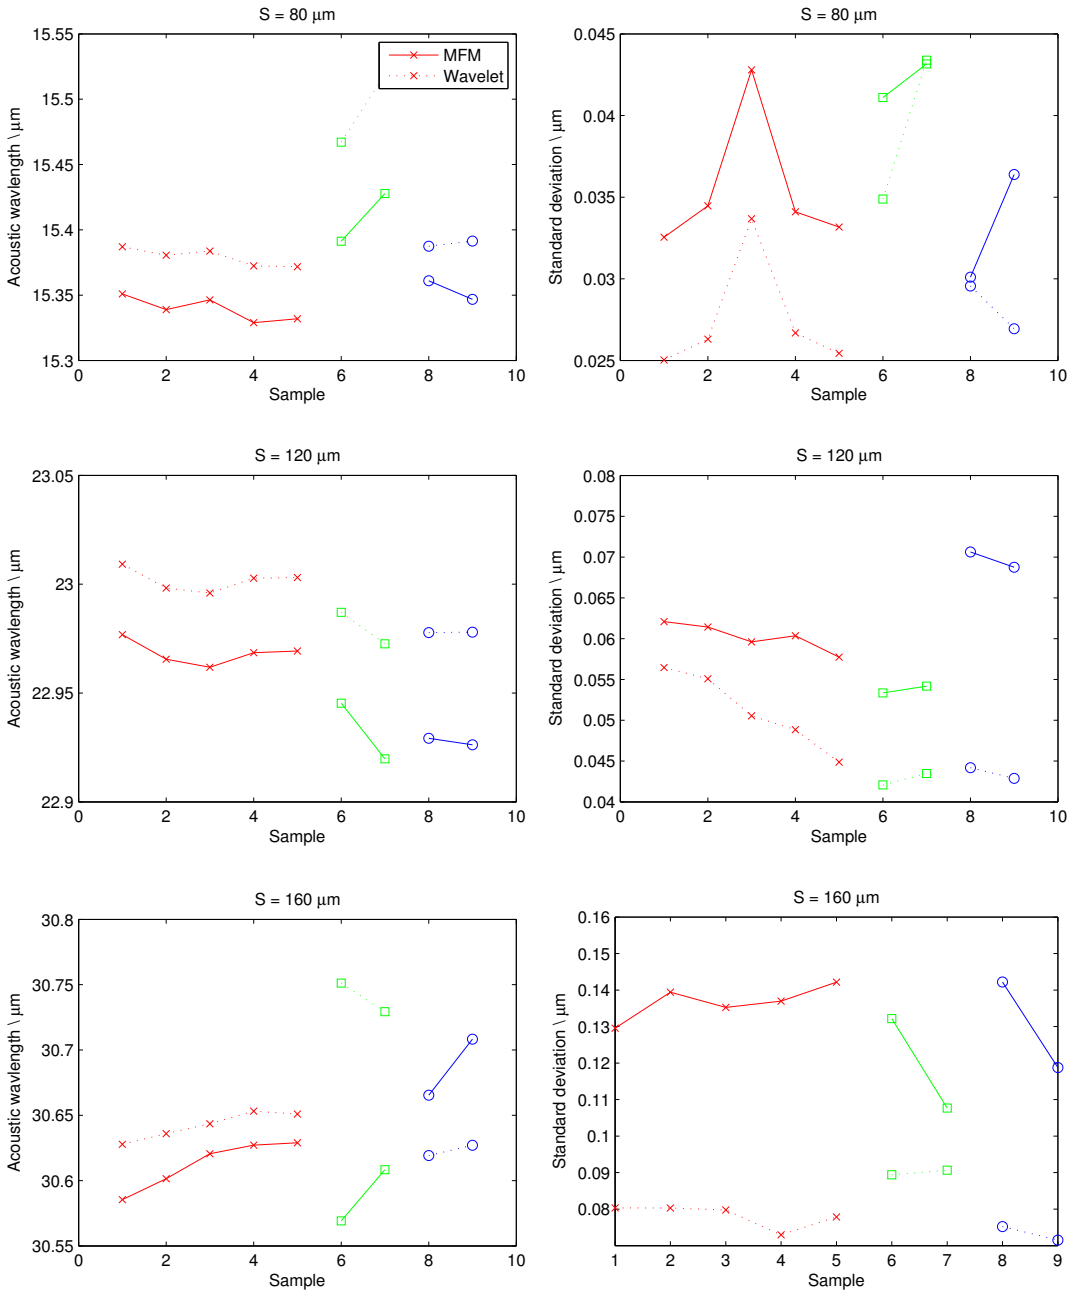


Figure 5.19: A set of outputs from repeated acoustic wavelength calibration procedures from 3 grating spacings. For each set of spacings 3 near identical gratings were calibrated using both the MFP method (solid) and a Morlet wavelet (dashed) for determining peak frequency, each grating is identified by a different style of line plot. The measured median acoustic wavelength is shown on the left and the standard deviation of the measured frequencies is shown on the right.

5.7 Comparison of detectors

An alternative to the segmented photodiode KED was discussed in §3.3.3 which uses a prism to split the detection beam before it is focussed onto a pair of single element photodiodes, this is referred to as the twin KED. The segmented KED, and the associated focussing lens, were removed from the SRAS head assembly and the twin KED inserted. Compared to the segmented KED, the twin KED is much more difficult to align and balance the detection beam equally over the two photodiodes. This is in part due to the small target of the prism edge but is primarily due to the lack of a lens after the steering mirror. In this configuration the beam can only be pointed at the detector rather than moved laterally over the detector, as with the segmented photodiode. A comparison scan was taken on an isotropic aluminium sample for each detector and the amplitudes are shown in Fig. 5.20. Step sizes of 5 µm and 10 µm were taken along the x and y axes respectively, propagation is along the x-axis with an acoustic wavelength of 23 µm and no aperture. Due to a slight shift in the location of the head following reassembly, the position of the generation patch relative to the sample was not the same between scans. Over the two scans the only change is in the detection path, the generation remains constant and therefore SAW amplitude and wavelength can be considered to be identical.

The signal amplitude as measured by the twin KED is worse than that measured by the segmented KED by a factor of 2.03 over the scan. This drop is primarily a result of the lack of detection laser power at the photodiodes, the optical sum is down by a factor 1.8. Partially this is due to reflection at the additional glass surfaces in the optical path added by the AR coated prism and uncoated ball lenses. In addition the prism has a 0.25 mm chamfer on the splitting edge which means that the central portion of the detection beam, where the most power is contained, is not split between the photodiodes but instead travels undeviated through the prism. With less optical power the signal amplitude is reduced, however the drop in measured signal amplitude is less than the drop in optical power indicating that the twin KED is more sensitive than the segmented KED.

To test the performance of the twin KED on real samples a pair of scans were taken on a large-grain aluminium sample using both detectors, these are shown in Fig. 5.21.

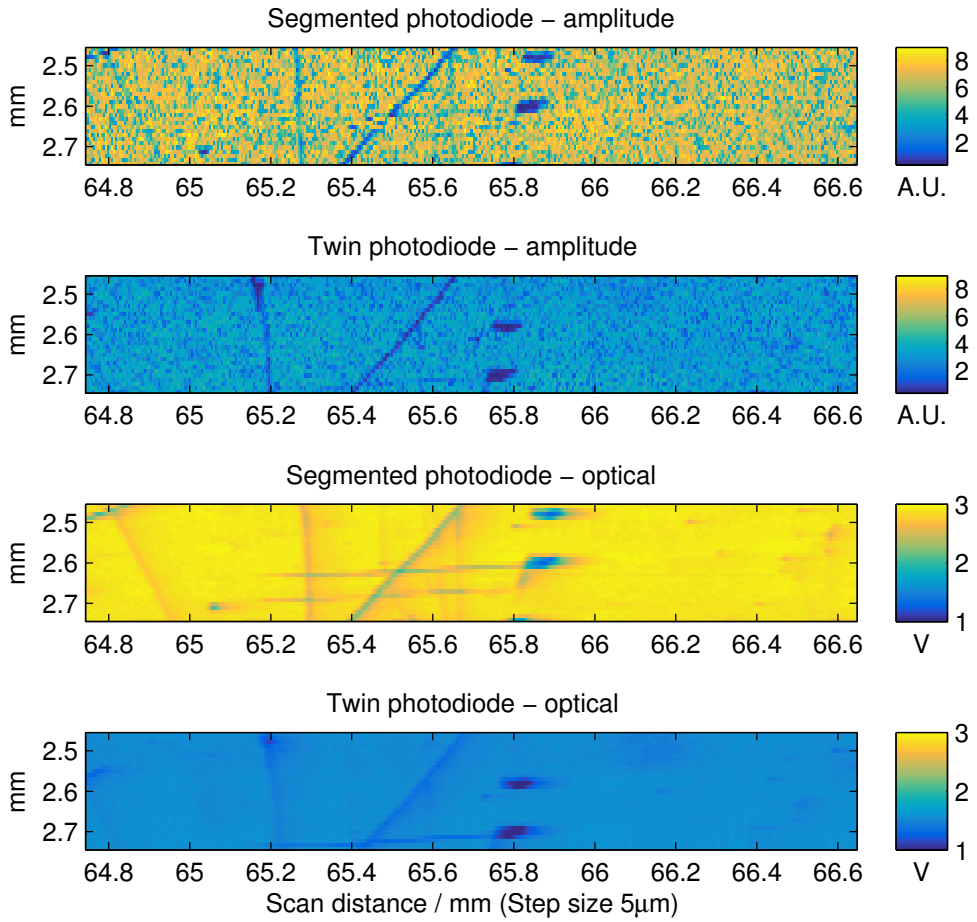


Figure 5.20: Two comparison scans taken on an isotropic aluminium sample with the segmented KED and twin KED, only the signal amplitude and optical sum are plotted and the same colour scale is used. Step sizes of 5 µm and 10 µm were taken along the x and y axes respectively, propagation is along the x-axis with an acoustic wavelength of 23 µm and no aperture.

Both scans were taken with an acoustic wavelength of 23 µm, an aperture of 1 mm, propagation in the positive x direction and step sizes of 10 µm in x and 20 µm in y. While the scan taken with the segmented KED has better SNR there is not a great disparity between the two. In terms of generating an image from which the sample grain structure can be determined both are successful. However in some regions the twin KED signal is too small, compared to the noise floor, to make a correct determination of SAW phase velocity, whilst the segmented KED can make the measurement. The signal amplitude ratio, which is the segmented photodiode signal amplitude divided by the twin photodiode signal amplitude, is 1.28 while the optical ratio is 1.24. This indicates a fairly linear drop in signal amplitude with optical power.

For highly reflective samples, where there is ample detection power, the twin KED works quite well, however many samples have poor surface finish where the optical power budget is critical. To test the performance on such samples two scans were taken on a Ti 685 sample using both detectors and the results are shown in Fig. 5.22. Both scans were taken with an acoustic wavelength of 23 µm, an aperture of 1.5 mm, propagation in the positive x direction and step sizes of 10 µm in x and 20 µm in y. The effect on the scan quality due to the loss of signal amplitude is much more pronounced than in the scans on aluminium. Quite large regions of the twin KED scan are dominated by noise, making identification of grain structure impossible. Between the segmented and twin KED signal amplitude ratio and optical ratios which were found to be 2.06 and 1.58 respectively. This is a significantly larger drop in amplitude relative to optical power compared to the scans on aluminium. The reason for this is that the titanium sample scans are in the regime where the already low detection power drops to the point at which the signal is indistinguishable from the noise.

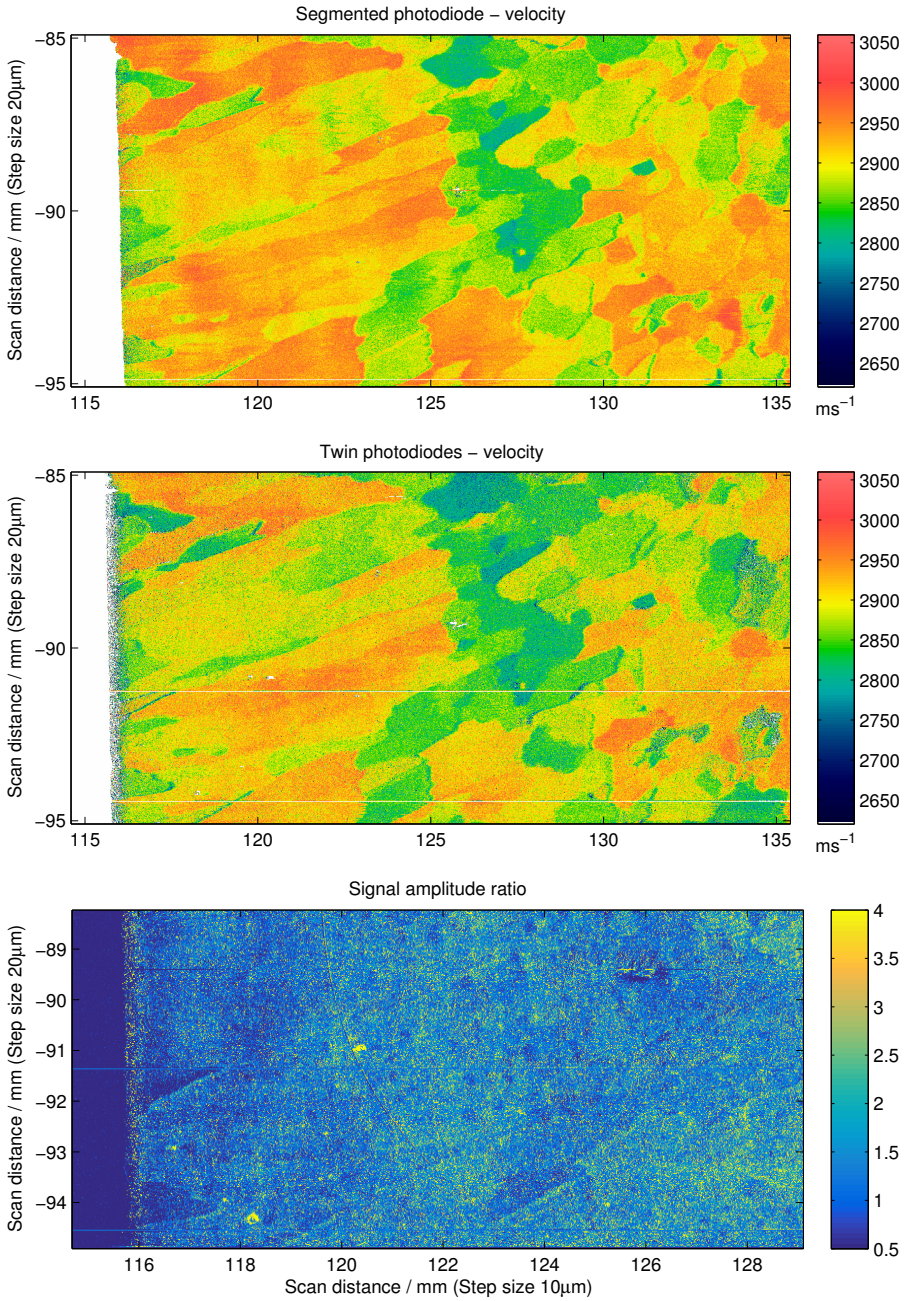


Figure 5.21: Two comparison scans taken on a large grain aluminium sample with the segmented KED and twin KED. Step sizes of 10 µm and 20 µm were taken along the x and y axes respectively, propagation is along the x-axis with an acoustic wavelength of 23 µm and a 1 mm aperture. The velocity image colourbars are in units of $\mathrm{m\,s^{-1}}$.

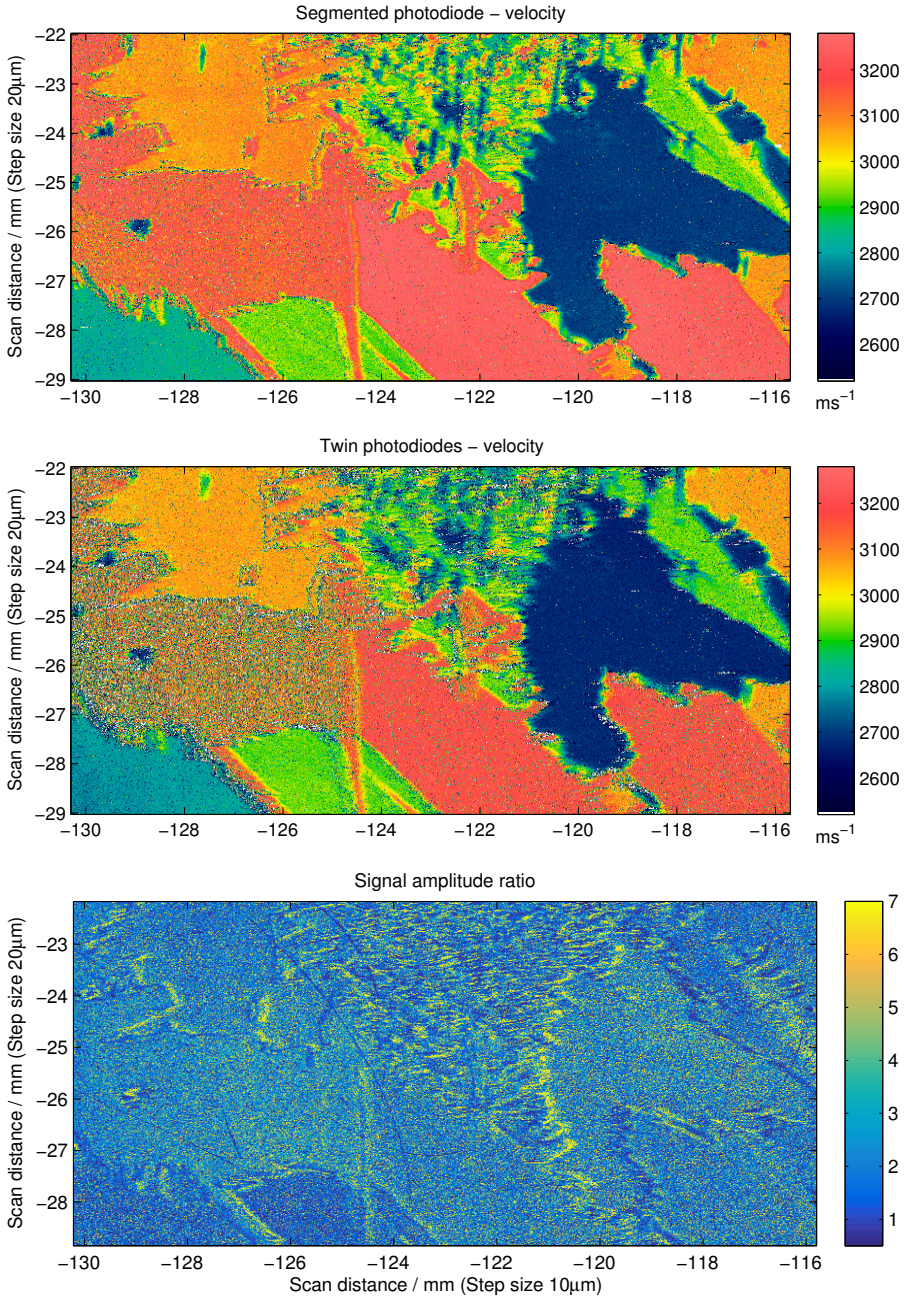


Figure 5.22: Two comparison scans taken on a low reflectivity Ti 685 sample with the segmented KED and twin KED. Step sizes of 10 µm and 20 µm were taken along the x and y axes respectively, propagation is along the x-axis with an acoustic wavelength of 23 µm and a 1.5 mm aperture.

5.8 Repeatability of assembly

The criteria for a good repeatability of assembly is threefold:

- after a modification or instrument head re-build, only a small realignment is required to return the system to working condition,
- after a modification or instrument head re-build, the performance of the instrument is negligibly affected,
- the performance differential between two identically constructed assemblies is negligible.

The latter of these is not addressed here simply because, for the prototype instrument, only one assembly exists. In addition to the above criteria there is a commutability requirement for the optics and optical housings used in the assembly. This requirement is that instrument performance is minimally affected by the small deviations inherent in optical and mechanical component manufacture; primarily deviations in part size, focal lengths of lenses, and flatness of mirrors.

Based on the regular disassembly and reassembly of the instrument, the repeatability of assembly is quite good. Between many of the scans shown in all sections of this thesis the SRAS head has been dismounted, disassembled, tweaked and reassembled with no evidence of any loss in performance. The main reason for this is the solid tubes used to set the optical path lengths. As discussed previously the only critical path lengths are between the mask and tube lens, and between the KED and the detector focussing lens. Both of these paths employ a single section of tube which ensures excellent repeatability. Optical alignment following a reassembly is also a quick and simple procedure due to the kinematic mounts used for the fibre optics, and the chassis kinematics constraining the mirrors.

With regard to the requirement of commutability, there are two aspects: the in-house manufactured optical housings, and the bought in optics. For in-house manufactured parts the design process and drawings include thorough tolerancing and dimensioning of parts, which are then inspected post-manufacture and discarded if out of tolerance. As such the equivalence of these parts to each other, and to the original design, is both very good and well known. Typically critical features are toleranced to less than

100 µm at which point the effect of any variations on the system performance has proved to be negligible.

The quoted tolerance on focal length for the critical f100 lens in the generation patch imaging path is ± 2%, this gives a focal length tolerance of ± 2 mm, a value far larger than the typical manufacturing tolerances of the mechanical components in the assembly. This level of variance in the optics is high and therefore, for the critical generation path tube lens, the focal length must be measured using a diopter gauge before assembly and corrected for with shims or custom cut tube. A low-cost gauge can typically determine focal length to approximately ± 0.5% leaving a potential focal length error of ± 0.5 mm. To determine whether this length error is a critical issue, a set of scans comparing an optimal set-up with a compromised set-up were performed. On a sample of Ti 685 two scans with apertures of 1 mm and acoustic wavelength of 23 µm were taken at step sizes of 10 µm along the x-axis and 20 µm along the y-axis. One scan was taken with an optimum optical path, while the distance between the mask and f100 tube lens was increased by 0.5 mm for the other, the results are shown in Fig. 5.23.

When the generation path is sub-optimal the focal plane of the generation patch relative to the focal plane of the detection spot is changed. Therefore either the generation or detection efficiency is compromised dependent on the configuration of the scan. The end result of either is a loss in measured signal strength. For the defocused scan shown in Fig. 5.23 there is a slight shift in velocities which indicates that the acoustic wavelength has changed and so the generation patch is not optimally focussed. The effect on the image is a general increase in noise due to the fairly significant factor 1.36 drop in signal strength. This is a critical issue and to diagnose this from a SRAS image is very difficult, making good measurements of optical component parameters essential.

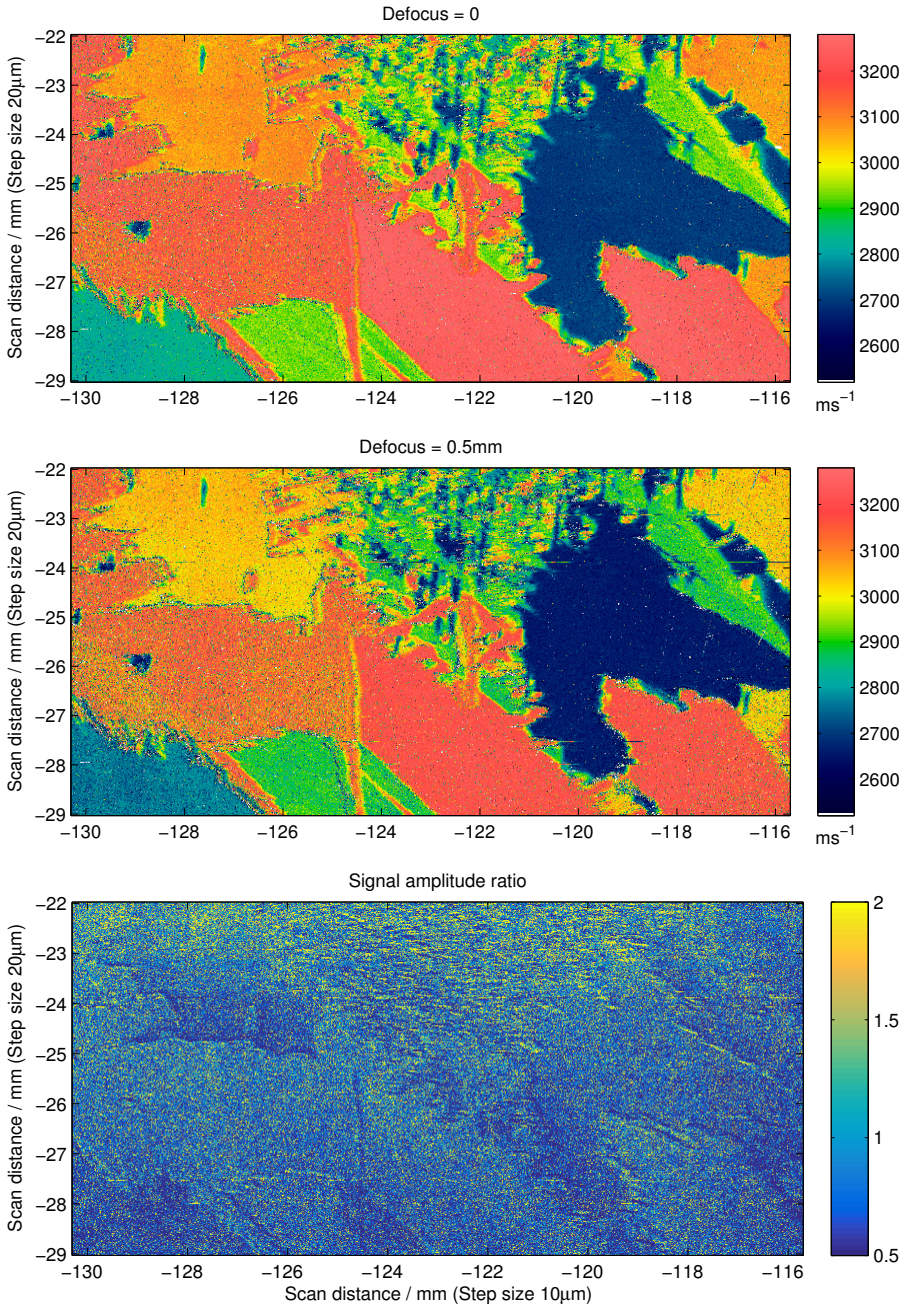


Figure 5.23: Two SRAS scans taken on a large grain Ti 685 sample with a optimal optical set-up and a set-up with the generation path defocused by 0.5 mm. Step size was 10 µm along the x-axis and 20 µm along the y-axis, propagation was in the positive x direction. An acoustic wavelength of 23 µm and an aperture of 1 mm were used.

5.9 Conclusions

In this chapter the general performance of the Renishaw SRAS instrument has been tested fairly exhaustively, and where possible quantified.

Although generally well suited to the task of scanning, some issues have been noted concerning the flatness of the Equator 700 motion. This issue is exhibited as a regular pattern of ~5 µm deviations in the Z axis which is repeatable over multiple scans and at multiple scan speeds. How these Z modulations affect the quality of SRAS data is explored further in Chapter 6.

The energy available for generation supplied by the Innolas laser has been explored, and an issue of decreasing energy with laser repetition rate, due to both a lack of power and temporal beam widening, has been noted. Although this issue has not affected the SRAS data used so far, the long term plan is to collect data at greater repetition rates where a lack of generation energy critically limits the range of rates available. The possibility of applying a MOPA to the generation laser was raised in discussion of instrumentation, however, in view of the dependence of SAW amplitude on temporal pulse length, this will not provide the gains expected. Consequently the optical path efficiency is clearly highlighted as an area of high importance for further work.

SRAS spatial resolution has been shown to be highly dependent on the sample grain shape, SAW acoustic wavelength and sample anisotropy, but is primarily determined by the aperture size, scaling with a factor approximately 1/3 the aperture size. Acoustic wavelength affects spatial resolution but in a more subtle manner than the aperture, and is generally exhibited as greater velocity contrast of grains equal to or smaller than the SAW acoustic wavelength. For a 1 mm aperture, with a 23 µm acoustic wavelength, a resolution of ~30 µm has been achieved on a finely-grained, high anisotropy, stainless steel sample.

Smaller acoustic wavelengths allow more cycles in a SAW packet for a given aperture diameter, thereby increasing the reliability of the SAW phase velocity measurement. However the generation efficiency rapidly drops with smaller acoustic wavelengths and it has been postulated that this is primarily due to vignetting of the imaged grating lines. As such, the generation efficiency is reduced by ~35% for an acoustic wavelength reduction of 50%.

A new OBD detector has been tested and shown to successfully measure beam deflection. Due to a lack of suitable optics, the twin photodiode suffers from high detection beam power wastage and, as a result, poor sensitivity. However where detection power is abundant, the twin KED performs similarly to the segmented KED.

Finally, the optical assembly and repeatability of assembly has been shown to be quite good. It has also been shown that the instrument performance is critically affected by small, 0.5 mm, changes in the optical path lengths.

5.10 Future work

5.10.1 Equator 700 motion

The Equator 700 motion is not ideal and, given the repeatability, appears to be an inherent feature of the mechanics of the platform and is an issue which cannot be eliminated. There are two possible, non-mechanical, solutions to reduce this issue. The first is through software calibration whereby the entire volume of the Equator 700 is mapped, and corrections are applied to the motion paths based on the known calibration. A second potential solution is to issue on-the-fly corrections. Equator 700 controllers include a scanning probe feedback mechanism, which is a feature designed to enable scanning probes to report required motion corrections to maintain contact with a surface. If the deviation from the sample can be determined then this probe feedback can be used to apply corrections. Chapter 6 looks at the feasibility of using the SRAS camera to determine the required corrections.

5.10.2 Repeatability of assembly

One of the requirements for determining the repeatability of an assembly is that multiple assemblies of the same design have the same performance, and this requirement has not yet been tested. As such a clear piece of future work is to build multiple instruments and determine whether their performance is equivalent.

One of the issues raised in the assembly was the critical influence of small deviations in optical focal lengths to the system performance, particularly for the generation

path tube lens. A typical diopter scale does not offer the required accuracy for the small diameter lenses used in the SRAS instrument and so a custom solution for determining lens focal lengths is required. Fig. 5.24 shows a basic schematic for a system to determine the lens focal lengths. By adjusting the lens position relative to a grating mask, the mask image can be optimally collimated and the camera used at multiple positions to determine the collimation. Through this method the critical distance between lens and mask can be measured and a custom tube cut for each lens.

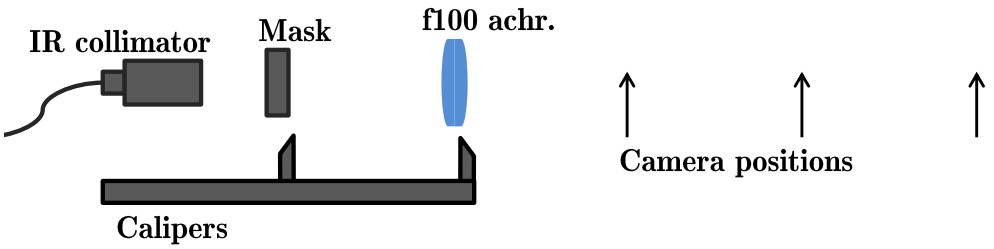


Figure 5.24: An optical jig for determining optimum separation between the grating and the f100 achromatic lens to be used to measure lens focal lengths before assembly.

5.10.3 Twin KED

The twin KED showed some promise as a detector but was badly affected by poor optical efficiency and was difficult to align. As such a few tweaks need to be made to the detector. By replacing the chamfered prism with a sharp-edged prism the light throughput can be easily increased. Using a beam expander before the prism will make alignment easier as the detector will be less sensitive to small changes in beam position on the prism. However, this has the disadvantage of decreasing the change in light intensity on each photodiode due to a SAW, therefore decreasing detector sensitivity. A suitable compromise needs to be found between sensitivity and ease of alignment.

The optically split beams offer the opportunity of fibre coupling the beams such that detection can be performed remotely from the instrument head. With this system the photodiodes can be housed in a very well sealed low noise enclosure, reducing shot noise and electrical interference, and cooled to help reduce thermal noise and increase consistency. A secondary advantage is that this removes a heat source from the SRAS head, which is good for metrology in terms of thermal expansion. In addition fewer cables need to be routed through the motion platform, reducing assembly complexity.

5.10.4 SRAS performance data consolidation

In this chapter each of the systems and mechanisms contributing to the performance of the SRAS instrument have been independently assessed. While some of these results, such as wavelength calibration and assembly repeatability, are useful for informing assembly and operational procedures, many of the results are not that useful in isolation. For an operator of the instrument one of the primary challenges is deciding which SRAS configuration parameters are suitable for a given inspection requirement. Currently this question can only be answered with experience and expertise in the SRAS technology. Thus, a significant body of work remains in consolidating the characterisations performed in this chapter into a form in which suggestions for SRAS configurations can be intelligently made.

Chapter 6

Machine vision for autofocus

The Renishaw SRAS system is designed to be tolerant to the choice of motion platform, and complexity of samples over which it is scanned. Between motion platforms there is a vast range of capabilities, from 1 nm precision piezo-electric stages to 100 µm low cost 6-axis articulated robotic arms. The Renishaw Equator 700 used for the SRAS prototype has ~1 µm repeatability, but is not designed or calibrated for absolute positioning within the measurement volume, or for perfectly smooth motion. In addition samples are not necessarily manufactured to reference CAD models, or drawings, and can be of complex, unpredictable, surface geometry due to the manufacture method itself, for example additive manufacture. If CAD models are used for generating scanning paths, or smooth surfaces falsely assumed, then this will cause issues with the quality and reliability of SRAS data.

Under such conditions the SRAS instrument is placed in an incorrect position and orientation relative to the sample surface under inspection. While there is some feedback from the optical sum and balance measurements, this feedback is not sufficient for correcting both location and orientation. During SRAS scans a camera is used to capture images of the generation patch. This chapter will describe the automated use of the camera generated images to determine how well focussed the patch is, and in which direction the instrument needs to move to improve focus.

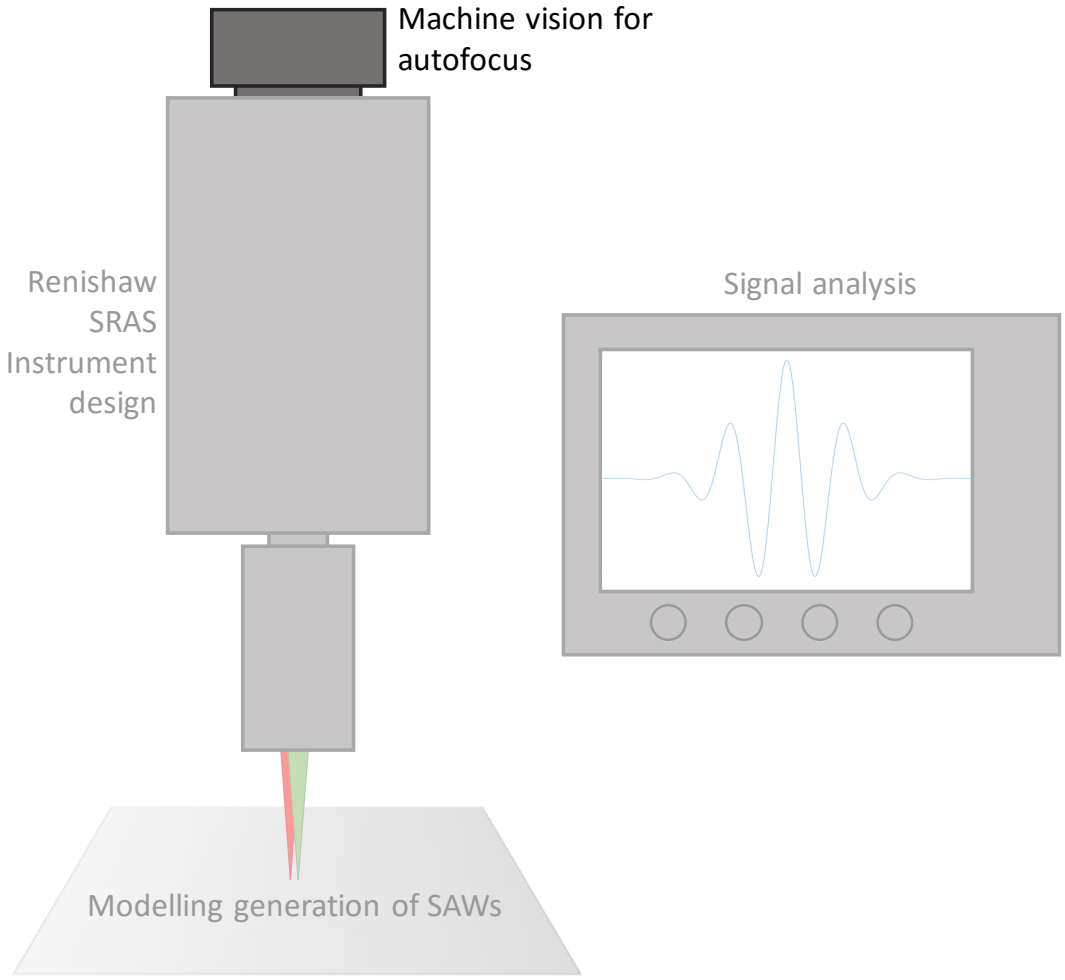


Figure 6.1: This chapter concerns image and focus analysis, and will determine whether, by analysing images of the generation patch, a focus correction feedback can be determined

6.1 Tolerance of SRAS to defocus

The SRAS instrument has three main factors relating to orientation and placement which can affect the reliability of measurements: the focus of the generation patch, the focus of the detection spot on the sample and detector, and the balance of the detection laser over the photodiodes of the KED.

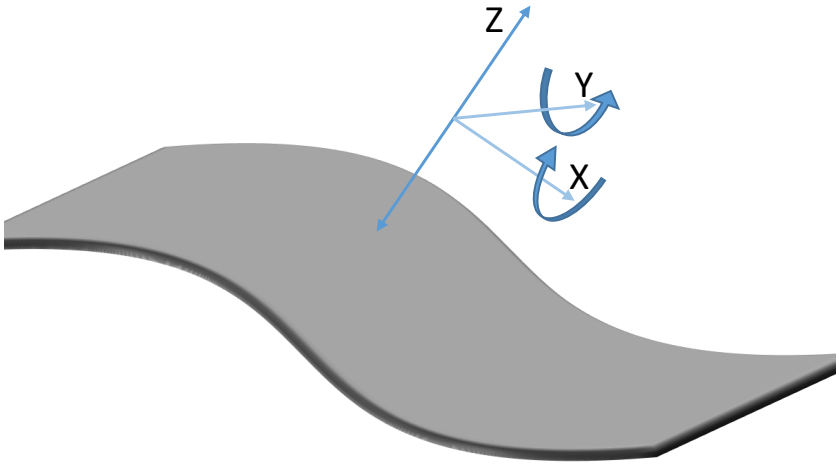


Figure 6.2: The axes around which misalignment of the SRAS instrument can occur relative to the sample surface.

Fig. 6.2 shows the potential misalignment axes for the SRAS instrument inspecting a curved sample surface. The first main misalignment is positional errors in X, Y and Z over the surface. Other than their effect on Z, and the local surface geometry, errors in X and Y do not contribute to a decrease in signal quality but do affect the correct registration of the SRAS measurement with respect to the sample. Changes in Z can affect signal quality in multiple ways: decreasing the generation laser power density, increasing the focal spot size of the detection laser, defocusing of the generation patch and changing the grating line spacing of the generation patch. The second misalignment is rotation around the X and Y axes such that the instrument is not collinear with the local sample surface normal vector. While this can have a small effect on the generation patch focus, the primary result is poor balance of the detection beam over the segmented photodiode, reducing detector sensitivity. Of these misalignments, primarily the linear change in Z is investigated.

The tolerance to changes in Z has been tested experimentally and is shown in Fig.

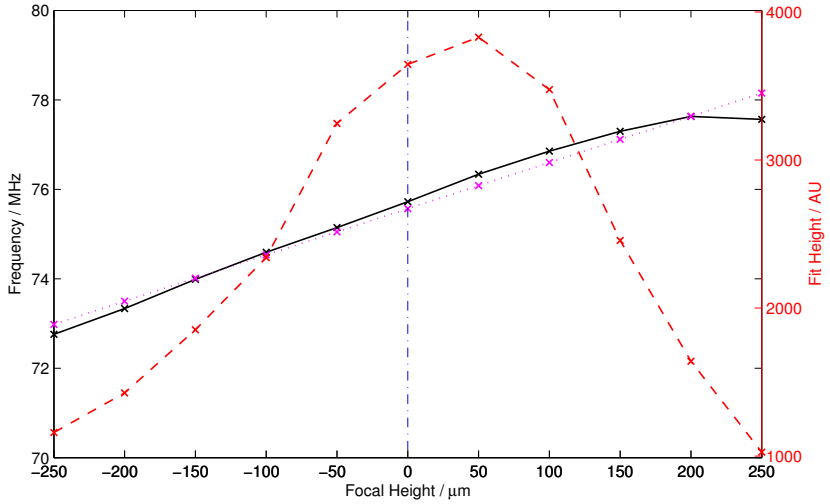


Figure 6.3: The measured average frequency (solid) and signal strength (dashed), or fit height, at a single point on an isotropic sample over a range of Z displacements from optimum focus. A best fit line has been applied to the average frequency (dotted red).

6.3. These data were collected on a static and rigid set-up of the SRAS instrument with a 50 µm acoustic wavelength and no aperture. An isotropic aluminium sample was placed on a manual Z-stage and 500 samples were collected and processed, the median peak frequency and signal amplitude were then determined, before the sample was moved by 50 µm in Z. The zero-height was taken to be the Z displacement at which the camera image of the generation patch was visually at sharpest focus. There are three striking features of Fig. 6.3. The first is that there is a very strong peak in signal amplitude at an optimum focal distance, and the second is that this does not occur at the zero focal height. This is due to a slightly mis-calibrated camera position which occurs during the assembly of the SRAS head. An initial estimate of the correct camera position is set by placing the camera sensor at the theoretical focal plane of the camera focus lens. However, due to an ~2% tolerance in the lens focal length this camera position is not always correct. The third feature is the very clear linear relationship between Z and the measured frequency, caused by changes in the acoustic wavelength. Fig. 6.4 shows the cause of this relationship. Acoustic wavelength should be symmetric about optimum focus and this result has not been returned, indicating that the static SRAS optical configuration is poorly optimised.

Due to the imaging system implemented in the generation optical path, the focal plane of the grating is independent of the focal point of the laser beam. As such, moving in

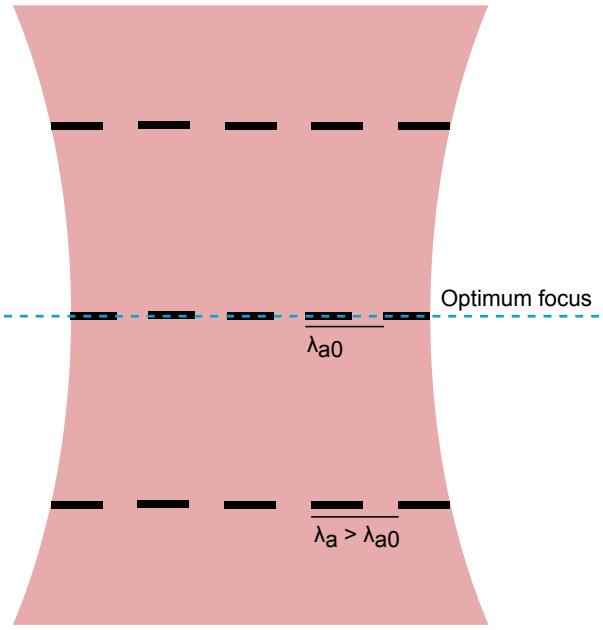


Figure 6.4: A drawing showing the change in line spacing of the grating at different heights in the focusing generation laser beam from the final objective lens. The optimal focal plane, where the sample surface is ideally placed, is shown by the blue dashed line.

positive Z, i.e. away from the sample, should give a larger acoustic wavelength and, since $v = f\lambda$, a smaller frequency. In Fig. 6.3 the opposite effect is seen and this is because of the poor optical configuration. Regardless it is clear that changes in focal height can affect a change in the grating separation and thus acoustic wavelength, highlighting this as a potential contrast mechanism for auto-focus. Taking the known SAW phase velocity of the isotropic sample as $2933\,\mathrm{m\,s^{-1}}$ and fitting a straight line to the frequency data it is found that the change of frequency with Z is $10\,\mathrm{kHz\,\mu m^{-1}}$, and the change of acoustic wavelength with Z is $5.2\,\mathrm{nm\,\mu m^{-1}}$.

Using the SRAS signal frequency is not a feasible route for auto-focus as it is very sensitive to the local microstructure, however images of the generation patch may provide useful data. In addition to measuring the signal strength and mean measured frequency, images of the generation patch were captured for each step in Z; these are plotted in Fig. 6.5. There is a circular feature visible in the middle right of most of these images, this is caused by a reflection of the detection laser within the SRAS instrument. Visibly the $0\,\mu\mathrm{m}$ image has the sharpest focus, with $-50\,\mu\mathrm{m}$ having the second sharpest. Conversely, the measured amplitude data show a peak at $\sim 40\,\mu\mathrm{m}$, further highlighting the poor configuration of the camera. By eye there is no clear

change in grating spacing.

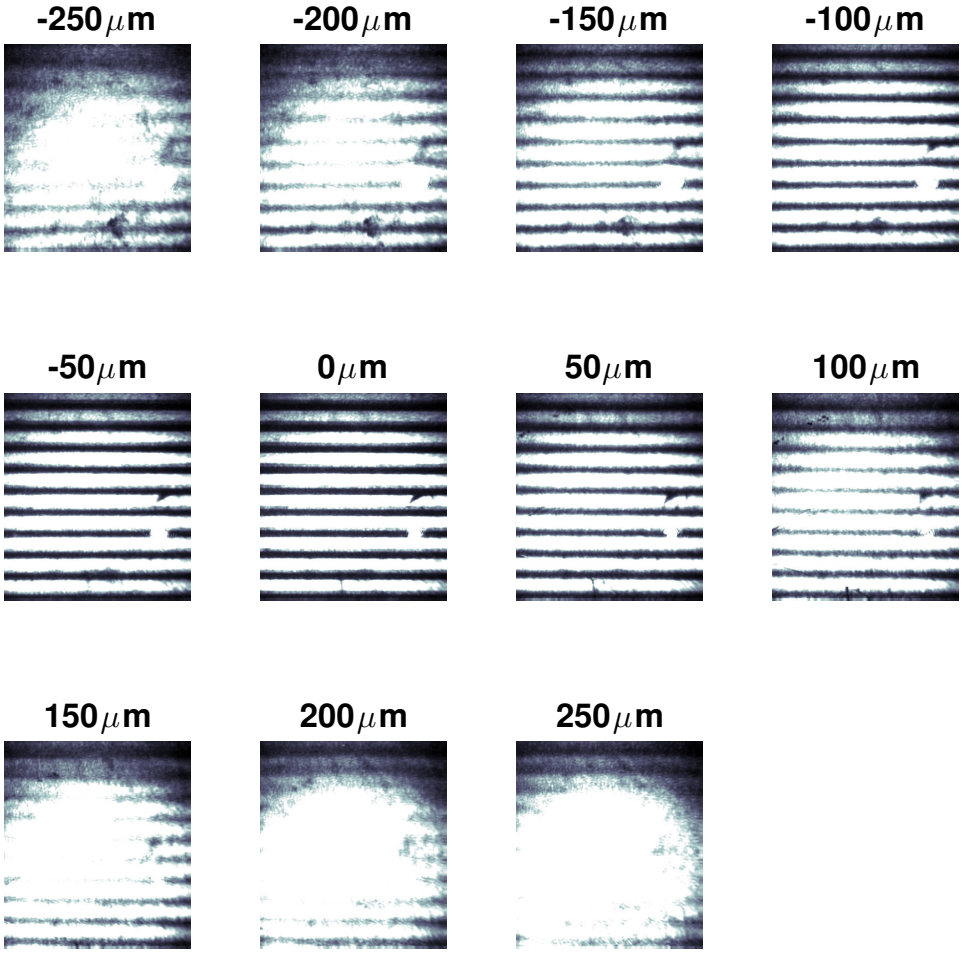


Figure 6.5: Camera images of the generation patch captured over a range of Z displacements. The images correspond to the data shown in Fig. 6.3. The patch only occupies a sub-region of the camera image which have been cropped to a 350 pixel square.

6.2 Methods for Determining Focus

The study of techniques for focusing is a well-trodden subject and excellent and thorough reviews are available which tackle the problem in a variety of fields. These include:

- Cameras, where the aim is to provide rapid auto-focusing of an image before a photograph can be taken. An overview of techniques is given in references [93] and [94], while reference [95] describes the efficient implementation of these techniques.
- Microscopy, where auto-focus algorithms are used to find the optimal objective-image distance [96, 97], and can also be used to enhance contrast and identify features in fluorescent image analysis [98]. A review of these techniques is given in reference [99].
- Machine vision, where focus can be used to determine the range to a part in a production environment [100], and can improve the determination of component shape [101]. A review of these techniques is given in reference [102].

As challenges in image focus analysis go, the generation patch is a very simple example case with high symmetry, good contrast and little noise, therefore sophisticated techniques are simply not required. As such, finding the sharpness of the images shown in Fig. 6.5 is a fairly trivial task and so the purpose of this review is to identify which technique provides optimum dynamic range, continuity over the full focal range and optimum processing speed. The ultimate aim of this is to enable measurement of sharpness on-the-fly during scans, and to feed this metric back to the motion platform and enable rapid corrections to be made to the scanning path and head position.

Images are captured at rates of up to 100 Hz dependent on the size of the generation patch on the camera sensor. Any processing technique used must be capable of exceeding this rate, as such certain techniques have been rejected due to their computational requirements, primarily discrete cosine transforms and wavelet transforms: these take several orders of magnitude longer to compute than many other techniques.

6.2.1 Tenengrad

The Tenengrad measure is based on finding the magnitude of the image gradient through application of a Sobel kernel, a discrete differentiation operator which efficiently estimates the gradient of an image [100]. Sobel filters are applied in each axis of the image independently, the two kernels, K_x and K_y are given by

$$K_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} \text{ and } K_y = K_x^{\mathsf{T}}. \tag{6.1}$$

The filtered image, $G_{x,y}$, is generated by convolving each of the kernels with the original image, $I_{x,y}$, i.e.

$$G_{x,y} = (K_x * I_{x,y})^2 + (K_y * I_{x,y})^2. \tag{6.2}$$

The Tenengrad measure, ϕ is then determined by taking the mean of the filtered image

$$\phi = \frac{1}{N_x N_y} \sum_{x=1,y=1}^{N_x N_y} G_{x,y}, \tag{6.3}$$

or alternatively by calculating the variance of the image,

$$\phi = \frac{1}{N_x N_y} \sum_{x=1,y=1}^{N_x N_y} (G_{x,y} - \overline{G})^2. \tag{6.4}$$

6.2.2 Laplacian

The Laplace operator, $\triangle$ is a differential operator which determines the divergence of the gradient of an image, this can be thought of as the second derivative of the image intensity,

$$G_{x,y} = (\bigtriangledown \cdot \bigtriangledown (I_{x,y}))^2 = (\triangle I_{x,y})^2. \tag{6.5}$$

In practice this is a computationally demanding operation and so, as with the Tenengrad measure, a estimate of the Laplace operator is applied as a discrete kernel [103],

$$G_{x,y} = K * I_{x,y} \text{ with } K = \frac{1}{3}\begin{bmatrix} 0.5 & 2 & 0.5 \\ 2 & -10 & 2 \\ 0.5 & 2 & 0.5 \end{bmatrix}. \tag{6.6}$$

The Laplacian kernel operates in two dimensions simultaneously, as such only a single convolution with the original image is required. A final focus measure is determined by taking either the mean or variance of the Laplace filtered image.

6.2.3 Gauss derivative

It has been shown that the first order derivatives of a 2D Gaussian function provide very good enhancement of edges and gradients in images [97]. The Gaussian function, $\gamma_{x,y}$, is derived from,

$$\gamma_{x,y} = \exp\left(-\frac{1}{\pi(\sigma_x + \sigma_y)}\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)\right), \tag{6.7}$$

where σ_x and σ_y determine the width of the Gaussian in each direction, these are set such that the Gaussian is confined within the chosen filter size, N, i.e. $\sigma = 0.1N$. Note that the Gaussian function is not normalised; since the same kernel is applied to each image, normalisation is unnecessary. The filtered image is generated by taking two orthogonal convolutions with the first derivatives of this 2D Gaussian function,

$$G_{x,y} = \frac{\mathrm{d}}{\mathrm{d}x}\gamma_{x,y} * I_{x,y} + \frac{\mathrm{d}}{\mathrm{d}y}\gamma_{x,y} * I_{x,y}. \tag{6.8}$$

As with the other techniques, a focus measure is then determined by taking either the mean or variance of the filtered image.

6.2.4 Grey level variance

There are some alternative methods which do not involve the application of filters, one such method is the grey level variance measure. Essentially this method determines focus by taking the standard deviation of the original image on the basis that a focussed image will have greater variance than an unfocussed image [93].

$$\phi = \frac{1}{N_x N_y} \sum_{x=1,y=1}^{N_x N_y} (I_{x,y} - \overline{I})^2. \tag{6.9}$$

6.2.5 Grey local level variance

An adaptation of the grey level variance technique is the grey level local variance [104]. Rather than taking the standard deviation of the entire image, the standard deviation of a sub-region around each pixel in the image is calculated

$$G_{x,y} = (L_{v,(x,y)} - \overline{L_v})^2, \tag{6.10}$$

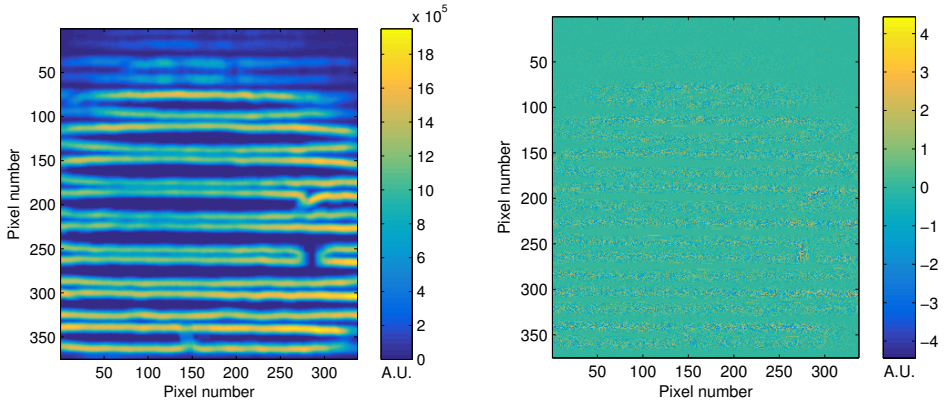
where $L_{v,(x,y)}$ is a sub-region of $I_{x,y}$ centred at (x, y) and with a width (w_x, w_y). A focus measure is then determined by taking the global variance of the local variances. This technique requires application of a sliding neighbourhood algorithm and is therefore computationally demanding. A modified version of the grey level local variance is calculated by [104]

$$G_{x,y} = K * I_{x,y}^2 - \frac{1}{w_x w_y}(K * I_{x,y})^2, \tag{6.11}$$

where K is simply a matrix of ones of size (w_x, w_y) . This simplification works by ignoring the small cross-terms in the expansion of equation 6.10 and is significantly less computationally demanding to calculate. Fig. 6.6 (a) shows $G_{x,y}$, as determined by the computationally efficient modified grey level local variance technique. In order to determine the effect of the simplification, the grey level local variance has also been computed by the sliding neighbourhood method, and the result of subtracting this from Fig. 6.6 (a) is shown in Fig. 6.6 (b). The computation errors between the two techniques is less than 0.0025%, therefore the modified grey level local variance can be used.

6.2.6 Gradient energy

The energy of the gradient technique simply takes the gradient of an image directly in orthogonal directions and sums the result to form an image with enhanced edges [103].



(a) (b)

Figure 6.6: (a) $G_{x,y}$ as determined by the method given in equation 6.10 for a grating image. (b) The difference between the two methods for determining $G_{x,y}$ given in equations 6.10 and 6.11. A kernel of size [7 7] was used.

Conceptually this is a very similar process to the Laplace and Tenegrad methods, and is calculated by

$$G_{x,y} = \left(\frac{\mathrm{d}}{\mathrm{d}x} I_{x,y}\right)^2 + \left(\frac{\mathrm{d}}{\mathrm{d}y} I_{x,y}\right)^2. \tag{6.12}$$

A focus measure is then determined by taking the standard deviation or mean of this gradient image. To nullify the influence of slow changes over an image a slight modification of the technique introduces a threshold value, τ, below which values in the gradient image are ignored, with the focal measure then given by the mean value of the thresholded gradient image,

$$\phi = \frac{1}{N_x N_y} \sum_{x=1,y=1}^{N_x N_y} G_{x,y},\ G_{x,y} \geq \tau. \tag{6.13}$$

6.2.7 Windowed Fourier transform

Since the features of interest are entirely parallel to one axis - the y-axis as plotted in Fig. 6.5 - the focal information of interest, i.e. the sharpness of the grating line edges, will be contained only in the perpendicular axis. As such focus measure techniques which apply only to one axis should be very applicable to this problem, and a custom technique has been designed around this principle.

- The median value of each row of pixels, i.e. parallel to the grating lines, is taken.
- A padded FFT is taken along the column of the average row values.
- An integral sum is computed over a given range in frequency space which is defined by the expected feature width. Since this width is just the grating width on the camera it can be estimated from the grating size and camera pixel size.

This technique is called a windowed Fourier transform (WFT) and formally this is formulated as

$$\phi = \int_{\nu(1)}^{\nu(2)} \left| \int_{-\infty}^{\infty} \left(\sum_{y=1}^{N_y} I_{x,y} \right) e^{-2\pi i \nu} dx \right| d\nu, \tag{6.14}$$

where $\nu(1)$ and $\nu(2)$ are the frequency range determined by the feature width.

6.3 Focus method comparison

With a set of focus methods determined and implemented in MATLAB, a set of sample SRAS data were collected. To avoid the influence of the motion platform on the stability of focus height, the SRAS head was mounted statically, and a micrometer driven Z-stage was used to set the focal height. A 40 µm acoustic wavelength was used with no aperture, and for each focal height 250 samples were taken on an isotropic aluminium sample. The measured signal amplitude and frequencies are shown in Fig. 6.3 and the captured images of the generation patch are shown in Fig. 6.5. To take advantage of the known directionality of the grating lines, where applicable kernels are rectangular with size [11, 7] to help enhance sensitivity to focus in the direction perpendicular to the grating lines. Where kernels are pre-set for the technique, i.e. Laplace and Sobel, the kernels given in the descriptions are applied. Fig. 6.7 shows the returned focus measure for each technique over the range of focal heights; both the normalised and absolute values are shown.

In order to choose the most suitable technique there are a few primary considerations: a smooth curve over all focal heights, no anomalous values, good dynamic range and low computational requirements. Based on the first three requirements the Laplace

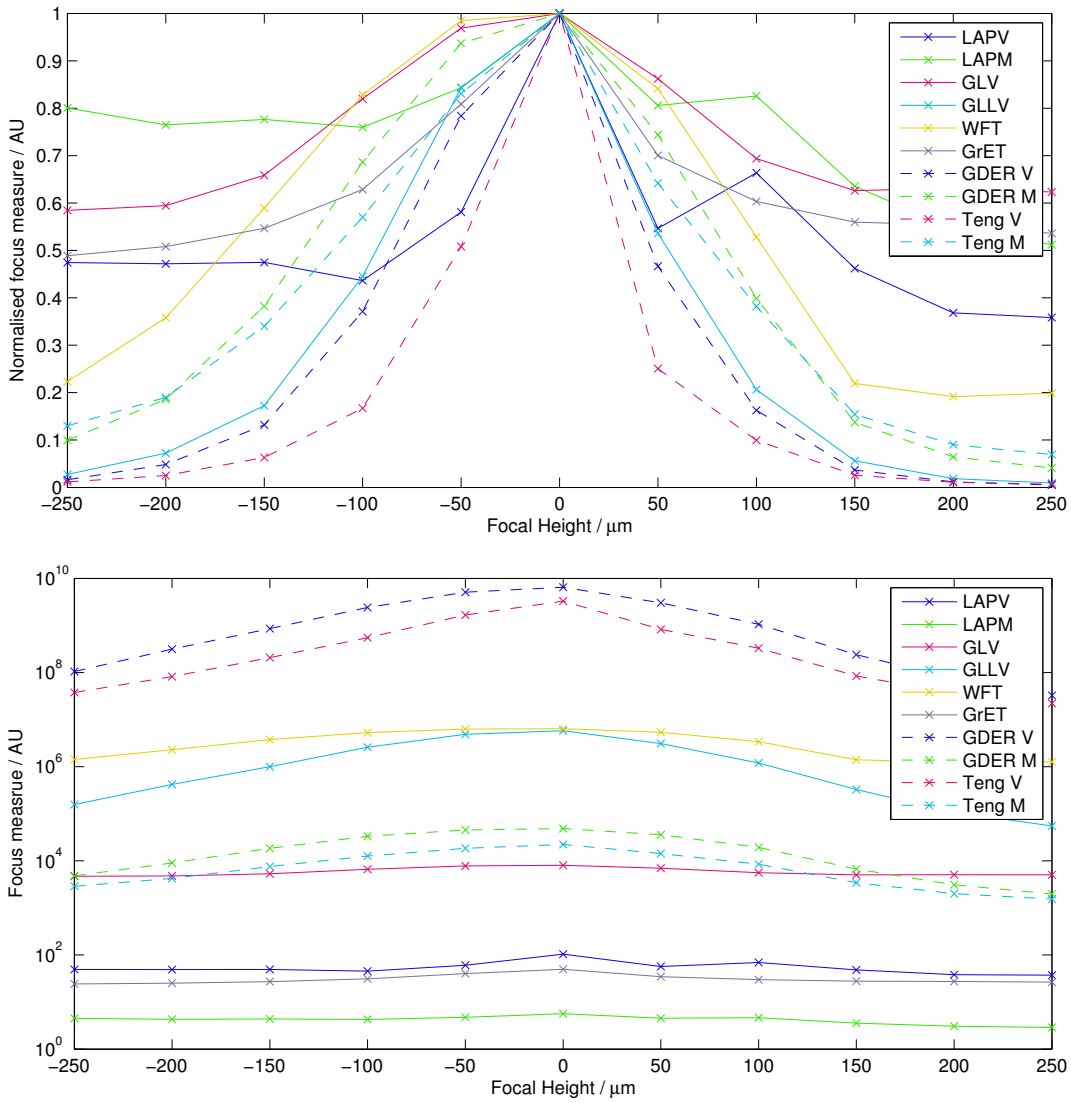


Figure 6.7: The results of applying the focal techniques to each of the images shown in Fig. 6.5. The data are shown normalised to their maximum value (top) and as the absolute figure returned by the processing technique (bottom). The techniques applied are: Laplacian variance and mean (LAPV/M), grey level variance (GLV), grey level local variance (GLLV), windowed Fourier transform (WFT), gradient of energy thresholded (GrET), Gaussian derivative variance and mean (GDER V/M) and Tenengrad variance and mean (Teng V/M).

methods can be immediately discarded since these are quite sharp, have values in disagreement with the other techniques and have poor dynamic range. Although smooth, the grey level variance has poor dynamic range, as does the gradient energy threshold technique. The Tenengrad methods are reasonably good, however they are both quite sharp and either drop to zero too quickly, in the case of the variance, or do not offer the dynamic range of some of the other techniques, in the case of the mean.

One of the main considerations is computation time, and this is shown in Fig. 6.8. Based on this, the windowed Fourier transform can be immediately discarded, otherwise all of the other techniques are quite similar in computation time apart from the very quick grey level variance. All of the techniques, other than WFT, have computation less than the available 10 ms between camera frames, given the 100 Hz camera repetition rates. Note that none of these techniques have been optimised yet for rapid computation and so the given values represent an upper limit.

In general the trend for the techniques used shows that most of the techniques with set, small, kernel sizes applied in 2D have not been very efficient. This is because the contributions to focus come from small scale noise rather than the longer scale grating line features which can be enhanced with larger kernels. With all others eliminated only the grey level local variance and the Gaussian derivative methods remain. Of the two Gaussian derivative methods the variance seems more suitable due to its greater range, and as such, the grey level local variance and Gaussian derivative variance have been chosen as the primary focus techniques for further evaluation.

A comparison of the signal amplitudes and focus measure for the data taken on the static SRAS configuration is shown in Fig. 6.9. This highlights again that the configuration of the camera is incorrect with the peaks in amplitude and focus offset by 40 µm. It is also of note that the signal amplitude drop-off is more gentle than that of the patch focus measurement. This is a good characteristic because it means that the response from the focal measurements will allow quick adjustments with a comparatively small decrease in signal strength.

In order to test the focal methods in a realistic scanning application, further data were taken but with a more representative SRAS experimental configuration. The SRAS head was mounted on the Equator 700 motion platform and a set of four acoustic

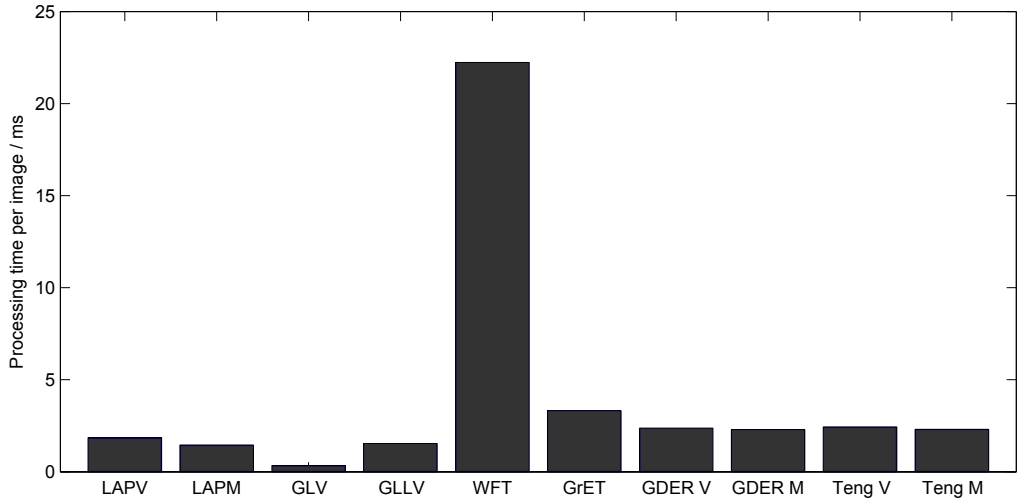


Figure 6.8: The computational time taken per image for each of the applied techniques. Computation time was calculated by averaging over 110 applications of each technique. The technique shortened names are described in Fig. 6.7.

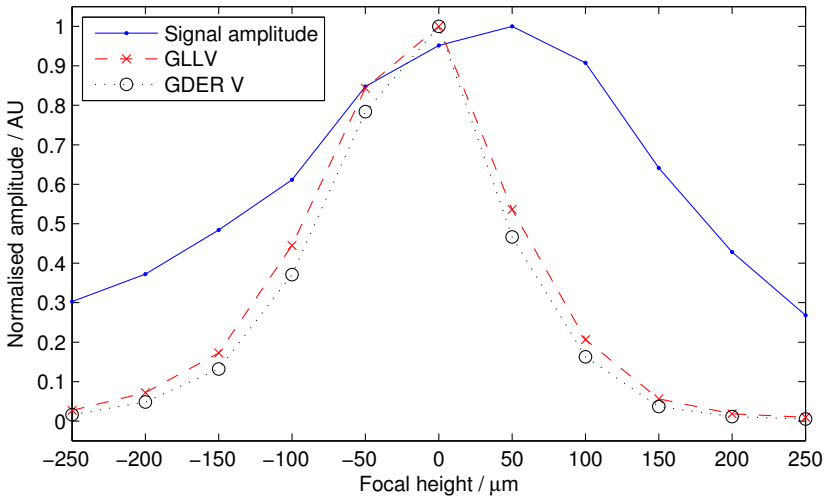


Figure 6.9: A comparison of the measured focus, determined using grey level local variance and Gaussian differential variance, with the measured signal amplitude for each of the patches in Fig. 6.5. An average of 128 traces at each focus height were used to determine signal amplitude. Scans were taken with a 40 µm acoustic wavelength and no aperture while the SRAS head was statically mounted.

wavelengths were tested, each with an aperture of 1.5 mm. Scans were taken on an isotropic aluminium sample with step sizes of 1 µm in X and 25 µm in Z, for each Z height 1000 traces were taken and analysed with the determined velocities averaged. Images were captured of the generation patch at the maximum rate the camera could run, ~100 Hz, and analysed using the grey level local variance and Gaussian derivative variance techniques with kernel size of [11 7]. The results for each acoustic wavelength are shown in Fig. 6.10.

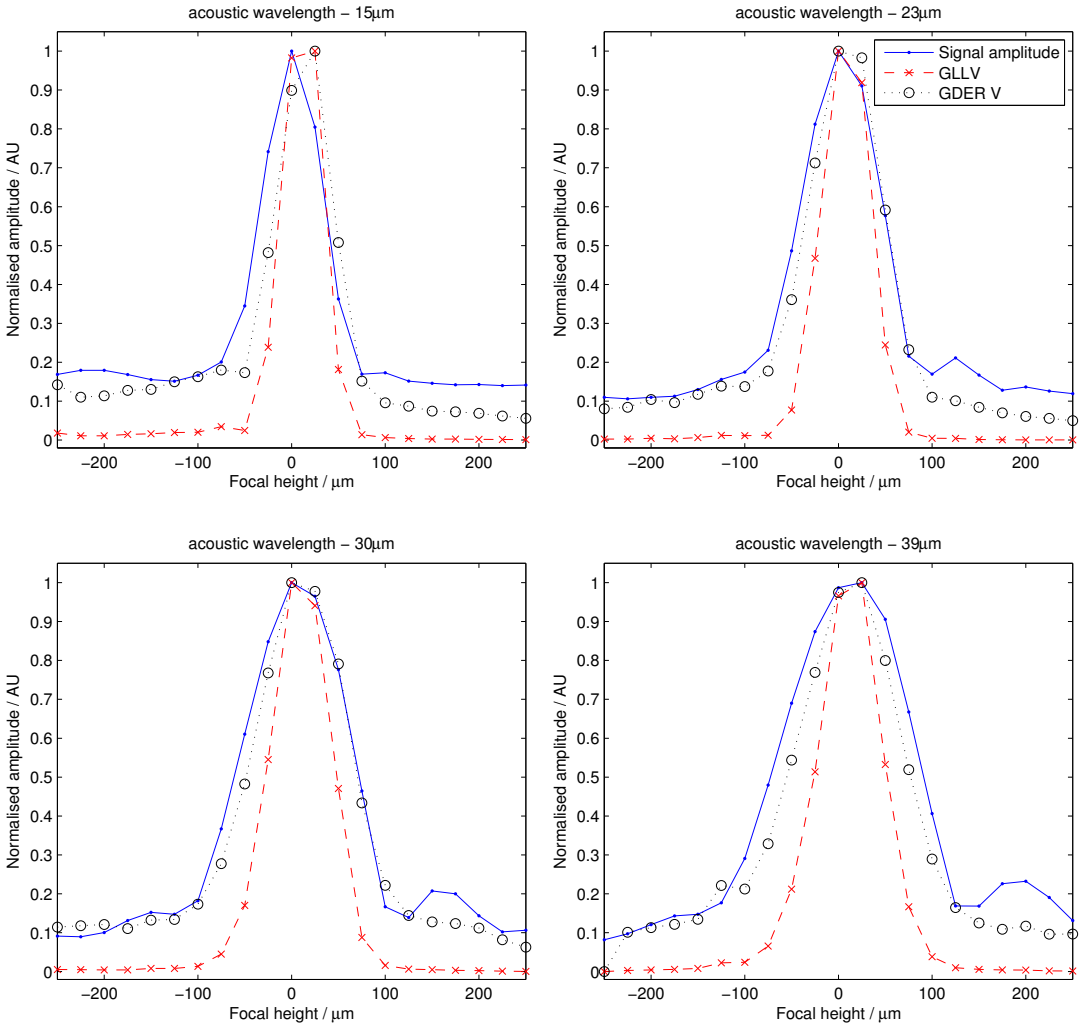


Figure 6.10: A set of scans taken at a range of focal heights with four different acoustic wavelengths. Images of the generation patch were also captured and two focal measurement techniques applied: grey level local variance and Gaussian differential variance. Average signal strength was determined from 1000 averages taken on an isotropic aluminium sample. For all acoustic wavelengths a 1.5 mm aperture was used.

With the addition of an aperture there is less energy available for generation, and as such the width of the signal amplitude curve over focal length has been reduced compared to the static SRAS head case. There is also strong correlation between the robustness of the SRAS technique to defocus and acoustic wavelength with the 15 µm acoustic wavelength signals dropping off much more rapidly than the 40 µm acoustic wavelength signals. This relates to the effect described in §5.8 whereby for smaller acoustic wavelengths a greater proportion of the generation patch is out-of-focus and therefore not generating SAWs efficiently. As the generation patch is defocused this effect is exacerbated and also seen in the measures of focus which also decay more rapidly for smaller wavelengths.

In contrast to the static SRAS head results the grey level local variance yields a significantly sharper peak than the Gaussian derivative variance. Since the only change in the image is the addition of the aperture this is most likely due to the edge effects introduced by the sharp aperture. Both focus methods however are in good general agreement with identical peak positions at all acoustic wavelengths. In terms of which focal measure is more useful there are two conflicting points. For a rapid response to defocus the grey level local variance is most useful as it has the largest gradient near the peak value, however it also decays to zero quite quickly meaning that if focus is more than ~60 µm out in one direction the focal measure becomes useless. Conversely the Gaussian derivative variance shows better contrast over the full range of focal heights but has a slower response with defocus. A compromise solution would be to continue using both techniques, as Fig. 6.8 indicates these techniques are two of the fastest available and there is little penalty in applying them both to each image.

From these signal amplitude data the focus range in which useful SRAS data can be taken can be quantified by fitting a Gaussian and taking the full width at half maximum for each amplitude curve. The centre of each amplitude curve varies only by ~2 µm over all acoustic wavelengths, while the full width at half maxima are: 50.0 µm, 59.7 µm, 73.8 µm and 79.7 µm, going from smallest to largest acoustic wavelength respectively. These values are quite promising and indicate that Renishaw SRAS instrument can operate over a fairly large focal range.

A secondary use of the focal measurement techniques is to provide rapid feedback for the correct positioning of the camera. The optimum objective to sample distance can

be found by positioning the SRAS head at a focal height which gives maximum SAW signal amplitude. By tweaking the position of the camera while monitoring the focal measure, the optimum camera position can also be found. This procedure has been adopted for all of the data shown in Fig. 6.10 and the result can be seen in the good alignment of the peaks in focal measure and signal strength. The experiment shown in Fig. 6.3 has been repeated after a correction in the camera position and the results are shown in Fig. 6.11. The improved agreement between the focal measure peak and signal amplitude peak indicates that the camera alignment procedure has been successful.

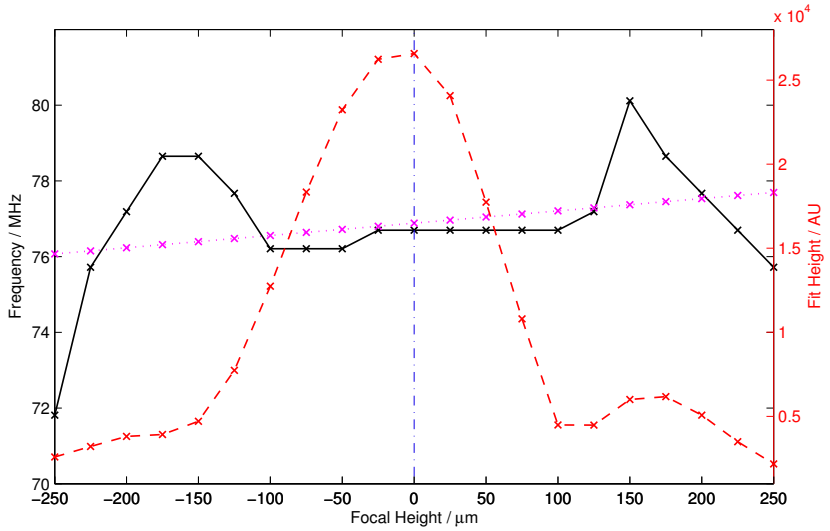


Figure 6.11: The measured average frequency (solid) and signal strength (dashed), or fit height, at a single point on an isotropic sample over a range of Z displacements from optimum focus. A best fit line has been applied to the average frequency (dotted red).

6.4 Determination of acoustic wavelength from camera images

Focus curves are symmetric which means that when the generation patch goes out of focus it is impossible to know whether the sample is too close or too far from the objective lens, the same is true of the signal amplitudes. It would be therefore useful to have a metric which determines the direction of defocus and in this section the feasibility of using the imaged grating spacings as this metric is investigated. As Fig. 6.4 shows there is an expected change in mask spacing with defocus. If this can be detected on the camera image then two useful properties may be ascertained: the change in acoustic wavelength and thus, from this, the direction and magnitude of defocus.

In order to determine the image grating spacings a broadly similar technique to the windowed Fourier transform is applied, but instead of summing over frequency space a peak is found. The process is:

- From the known grating size define a range of expected feature spacing, typically $\pm$ 25% of the expected value. Use the known camera sensor pixel pitch, 5.5 µm, to describe this in terms of pixel number.
- Collect an image of the generation and define the direction of interest as perpendicular to grating lines.
- Use a 2D Fourier transform to filter the low frequency information with the minimum frequency threshold set as 3 times the acoustic wavelength. The original image and filtered image are shown in Fig. 6.12 (a) and (b) respectively. The filter process has removed the background beam profile and improved the contrast of the grating lines.
- Take the mean across the rows, shown in Fig. 6.12 (c), this helps reduce the influence of high frequency noise, such as scratches on the sample surface, on the peak measurement and reduces the computational requirements. It also reduces the influence of slight rotation of the camera and mask relative to each other.
- Calculate the zero-padded FFT of the mean and cut it to the predefined feature spacing range, shown in Fig. 6.12 (d), from this curve the peak frequency is

taken.

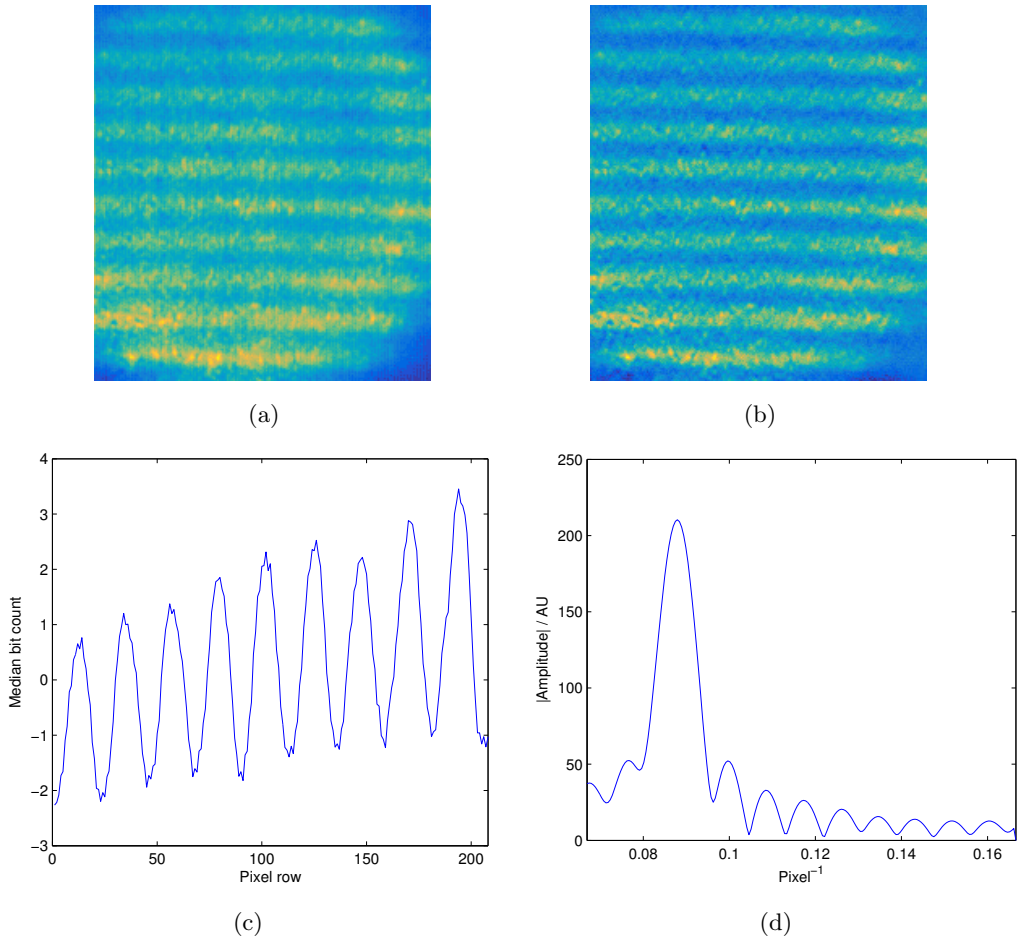


Figure 6.12: The steps taken to calculate the grating line spacings. (a) the input image of the generation patch. (b) the high-pass filtered version of (a). (c) the mean of the rows of (b). (d) the FFT of (c) cut to the region of interest.

This analysis technique was applied to the images captured on the static SRAS instrument, shown in Fig. 6.5, and is plotted along with the acoustic wavelength determined from the SRAS signals in Fig. 6.13. For easy comparison the measured image grating width has been set equal to the SRAS measured acoustic wavelength at optimal focus, without this the measurement of grating spacing has an offset of ~1.5 µm which is a result of slight uncertainty in the magnification of the SRAS system, and precise camera pixel pitch. When the image is poorly focussed determination of the grating spacing becomes unreliable however, where there is reasonable focus, a value is returned which is in good agreement with the SRAS signal derived SAW acoustic wavelength.

As with the focus measure techniques the grating width measurement was also applied

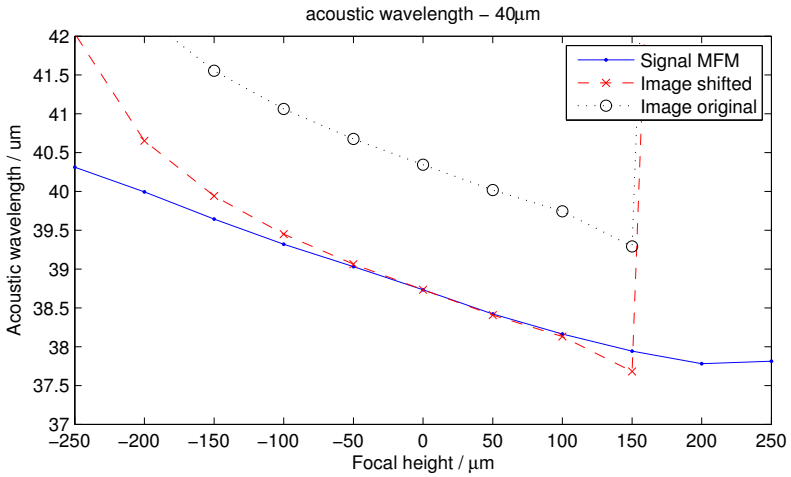


Figure 6.13: A comparison of the measured grating spacing, determined using the described method, with the measured signal acoustic wavelength for each of the patches in Fig. 6.5. An average of 128 traces at each focus height were used to determine acoustic wavelength. Scans were taken with a 40 µm acoustic wavelength and no aperture while the SRAS head was statically mounted.

to the data collected and shown in Fig. 6.10. These data were collected using the Equator 700, with nominal acoustic wavelengths of 15 µm, 23 µm, 30 µm and 39 µm. The grating width measurement has been applied to the camera images from each scan, and the determination of each acoustic wavelength at each focus height is shown as hollow circles on Fig. 6.14. In addition, the SRAS data were analysed using both Morlet wavelet and MFP analysis to determine the acoustic wavelength. These data are shown by the solid blue line and dashed red lines respectively on Fig. 6.14.

These results are in disagreement with the initial results measured on the static SRAS configuration, this is due to the optimisation and tweaking performed before mounting the SRAS head on the motion platform. Each of the SRAS derived acoustic wavelengths in Fig. 6.14 exhibit quite small changes over the focal height ± 50 µm with some forming a bowl shape around the 0 µm value of focus. This indicates that the optical paths are correctly configured such that the effect indicated in Fig. 6.4 is being observed.

Where there is poor image focus and signal amplitude, measurements of grating spacing become unreliable and this effect is especially prevalent on the shorter acoustic wavelength scans. In the regions where the focus is good, it can be seen that grating width measurement increases as the distance to the sample increases and this rela-

tionship does not correlate with the change in acoustic wavelength from SRAS data. The cause of this is the method by which images of the sample surface are captured and imaged. As the objective lens moves further from the sample the reflected image spreads out more before being captured resulting in a larger measured grating spacing, the converse is true as the objective lens moves closer to the sample. Consequently this measure of grating spacing is not useful for estimating changes in actual acoustic wavelength. However this trend is asymmetric about the point of optimal focus, and is therefore a useful metric for correcting the instrument motion as it moves too close and too far from samples.

To quantify how the measured grating spacing changes with defocus a set of scans were taken. These were of the same sample region as those in Fig. 6.14, but with a smaller and finer range of focal heights. Rather than steps of 25 µm between ± 250 µm, 5 µm steps were taken between ± 50 µm. The results of these scans are shown in Fig. 6.15 and display largely the same results as the previous scans. However, the finer step size allows straight lines to be fitted to the measure of grating separation and from these gradients can be determined. The measured gradients were, for each grating in ascending order: $4.2\,\mathrm{nm\,\mu m^{-1}}$, $9.1\,\mathrm{nm\,\mu m^{-1}}$, $9.9\,\mathrm{nm\,\mu m^{-1}}$ and $8.2\,\mathrm{nm\,\mu m^{-1}}$. The slightly unusual units represent the change in grating separation in nm due to a given change in focal height measured in µm. Determining these values gives a calibration for any feedback mechanism which adjusts the focal measure for each grating, i.e. a calculation can be made for a required correction based purely on the change in image grating spacings.

There is a slight discrepancy with the 40 µm grating gradient compared to the others, which tend towards greater steepness with acoustic wavelength. This is likely due to the relatively small number of grating lines in the captured image. When the FFT is taken in the algorithm for determining grating spacing, the small number of lines will give a broader frequency peak and a less accurate determination of spacing.

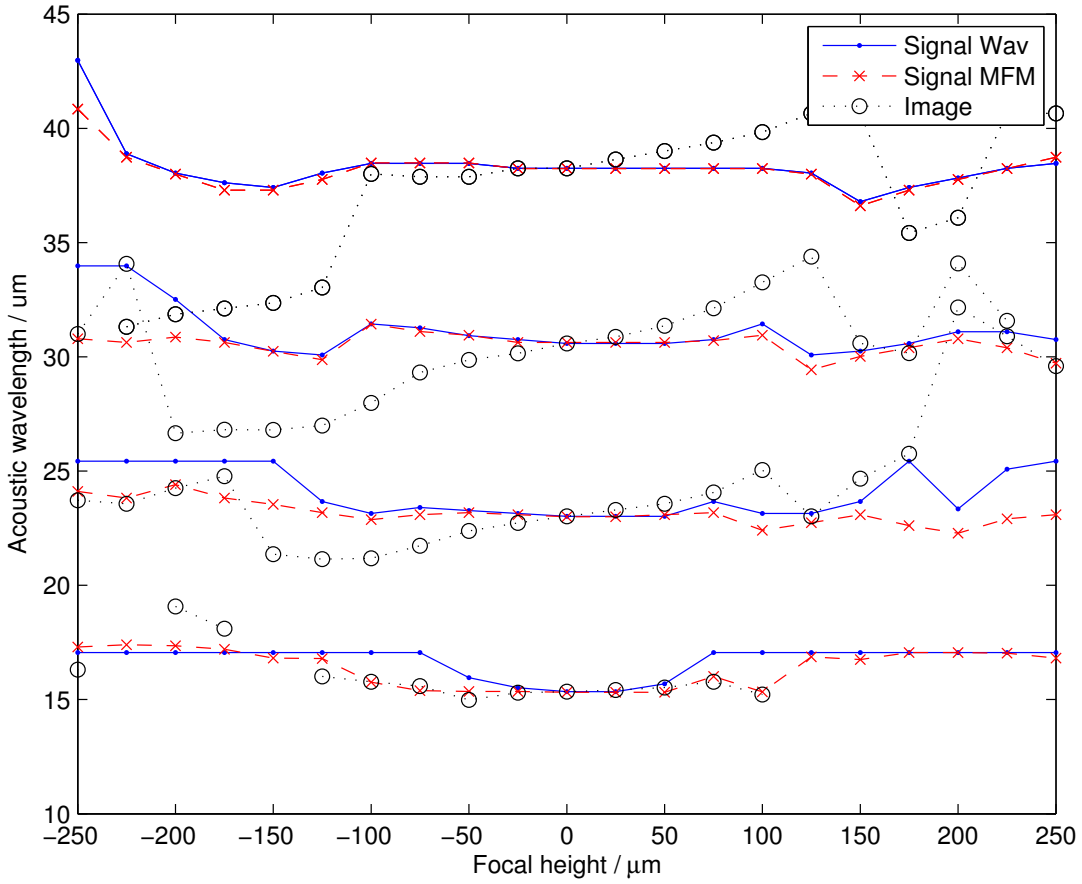


Figure 6.14: A set of scans taken at a range of focal heights with four different acoustic wavelengths. Images of the generation patch were also captured and the described method for determining grating spacing was applied. Average acoustic wavelengths were determined from 250 average acoustic wavelengths with a 4X average applied before signal processing with both Morlet wavelet and MFP methods. Scans were taken on an isotropic aluminium sample and for all acoustic wavelengths a 1.5 mm aperture was used. Anomalous measurements of acoustic wavelength due to poor focus have been omitted to improve clarity.

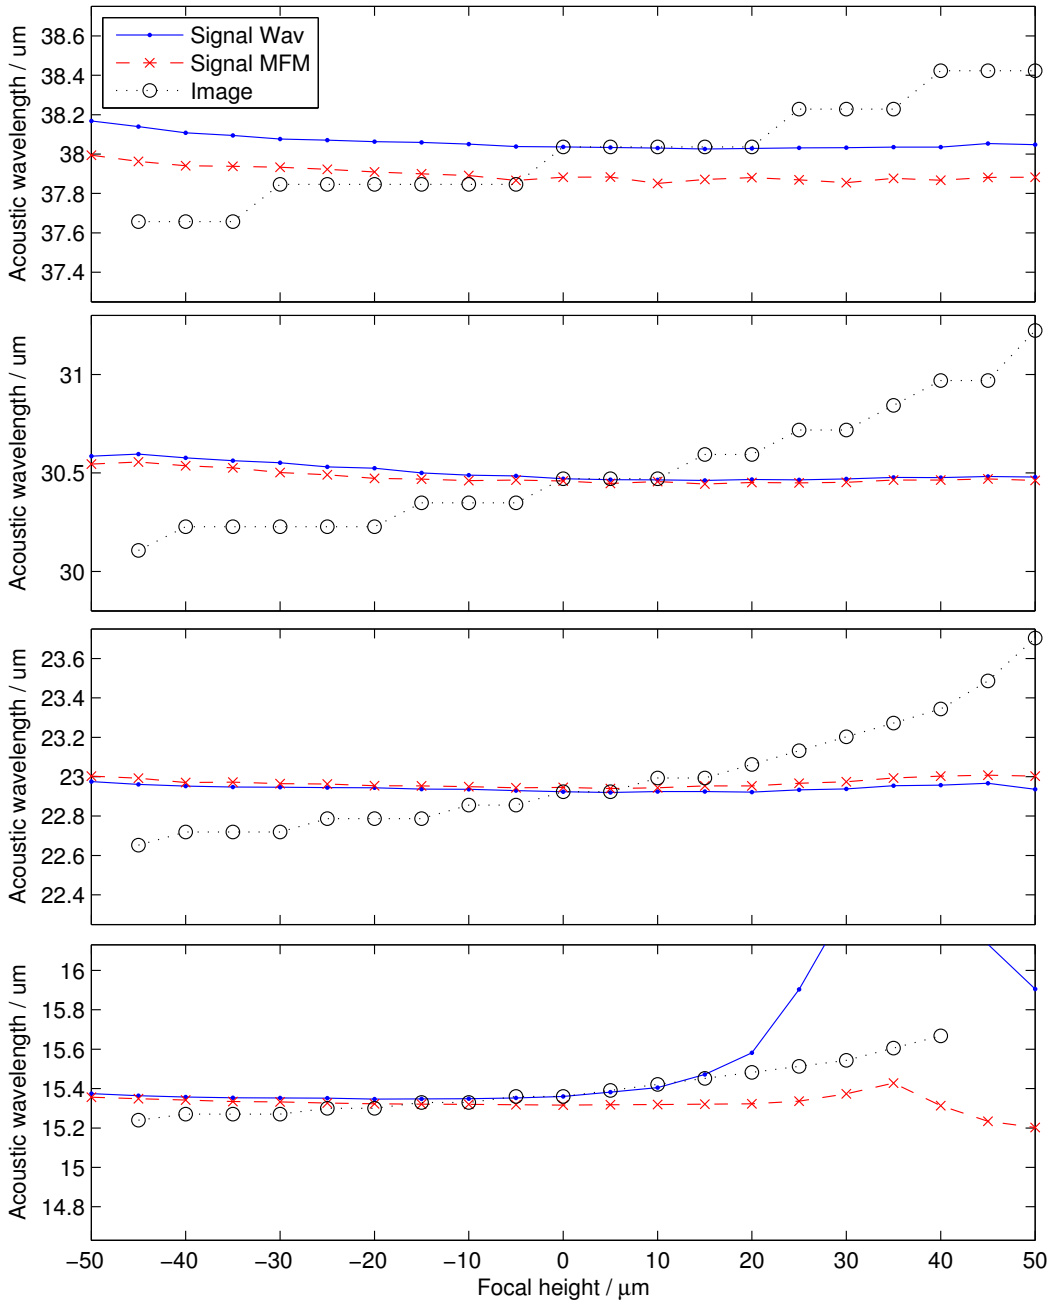


Figure 6.15: A set of scans taken at a range of finer range of focal heights with four different acoustic wavelengths. Images of the generation patch were also captured and the described method for determining grating spacing was applied. Average acoustic wavelengths were determined from 250 average acoustic wavelengths with a 4X average applied before signal processing with both Morlet wavelet and MFP methods. Scans were taken on an isotropic aluminium sample and for all acoustic wavelengths a 1.5 mm aperture was used. Anomalous measurements of acoustic wavelength due to poor focus have been omitted to improve clarity.

6.5 Application to full SRAS scans

The next step in determining the feasibility of focal height correction from the camera images is to test the analyses developed so far on a real SRAS scan. A scan was taken on an isotropic sample of aluminium, step sizes were 20 µm by 50 µm in X and Y respectively, and a 23 µm acoustic wavelength was used with a 1.5 mm aperture. A Z offset was set diagonally over the scan, such that the corner at (80 mm, −8 mm) had an offset of 50 µm and the corner at (48 mm, 8 mm) had an offset of −50 µm. During the scans the camera collected images at a frame rate of ∼90 Hz. Since the camera was not synchronised to the data collection the image locations were reconstructed using a time-stamp on the captured image, and the known velocity the of the SRAS head. With the scanning speed set to $25\,\mathrm{mm\,s^{-1}}$, the average step size in focus data is ∼250 µm. All captured images were processed for focus and grating spacing off-line, and the results of processing are shown in Fig. 6.16. The figure shows the mean image intensity, which is used to threshold data for plotting, the focal measure as determined by grey level local variance, and finally the determined focal height generated from the measured grating spacing.

Image intensity is fairly consistent over the entire sample this is because, although the image is defocused, the same light flux is returned to the camera and thus the same intensity measured. The cause of the horizontal banding seen in both the image intensity and focus measure images is currently unknown. Between bands there is no change in Z offset or sample surface condition and so the cause is most likely in the data collection system, perhaps an overheating camera issue. Although this banding does affect the focus measure, it only affects the intensity of the grating image and thus does not contribute to the measure of grating spacing. Looking at the focus measure shows the expected result; at the centre of the sample, where the Z offset is minimal, optimum focus is achieved.

To create the focal height image the measured grating spacings were normalised to a spacing taken from the centre of the sample, where Z offset is zero. The gradient for the 23 µm mask, found in the previous section, was then used to determine the focal height offset. As is clear from the image this process has worked very well, especially given that the only data used were a set of images of the generation patch. Feasibility

of focal height correction looks promising and, with the camera returning a simple correction value, complex PID type control mechanisms are not required.

With feasibility proven on a contrived case of defocus, a second scan was taken without any added deviations in focal height, on the sample used to test the motion platform in §5.1. A scan was performed with a 50 µm step size in X and a 100 µm step size in Y, an acoustic wavelength of 23 µm and an aperture of 1.5 mm. Images of the generation patch were collected continuously, although the higher scan speed used on this sample, compared to the previous isotropic aluminium sample, resulted in an average step size in the generation patch image of 900 µm. As with the previous sample the generation patch images were analysed for focus measure, intensity and grating separation off-line, and the focal height offset was calculated using the defocus gradient obtained from Fig. 6.15. The focal height offset is shown in Fig. 6.17 along with the Z position error reported by the Equator 700.

The figure shows the same pattern in height seen in previous scans from §5.1, which are a result of the mechanics of the Equator 700. Although there is a dominant gradient, due to a slightly tilted sample, the same pattern can be seen in the focal height offset. This shows that the technique for determining focal height offsets from generation patch images has good sensitivity to small values of defocus, capable of sensing offsets as small as ~3 µm. Previously, the SRAS signal tolerance to focal offset was measured at a minimum of 50.0 µm, for a factor 2 decrease in signal strength. Given this, for application to scanning this 3 µm sensitivity is more than sufficient.

In §6.4 a conclusion was drawn that this technique could not be used to correct for changes in acoustic wavelengths in post-processing of SRAS data, due to the non-linear change of acoustic wavelength with focal height offset. To further highlight the rationale for this decision a histogram was taken of SAW velocity against focus offset for the SRAS scan of the large mirror, this is shown in Fig. 6.18. While there is a slight trend between the two variables it is quite weak, and the wide spread in SAW phase velocity data means that post-scan correction is infeasible.

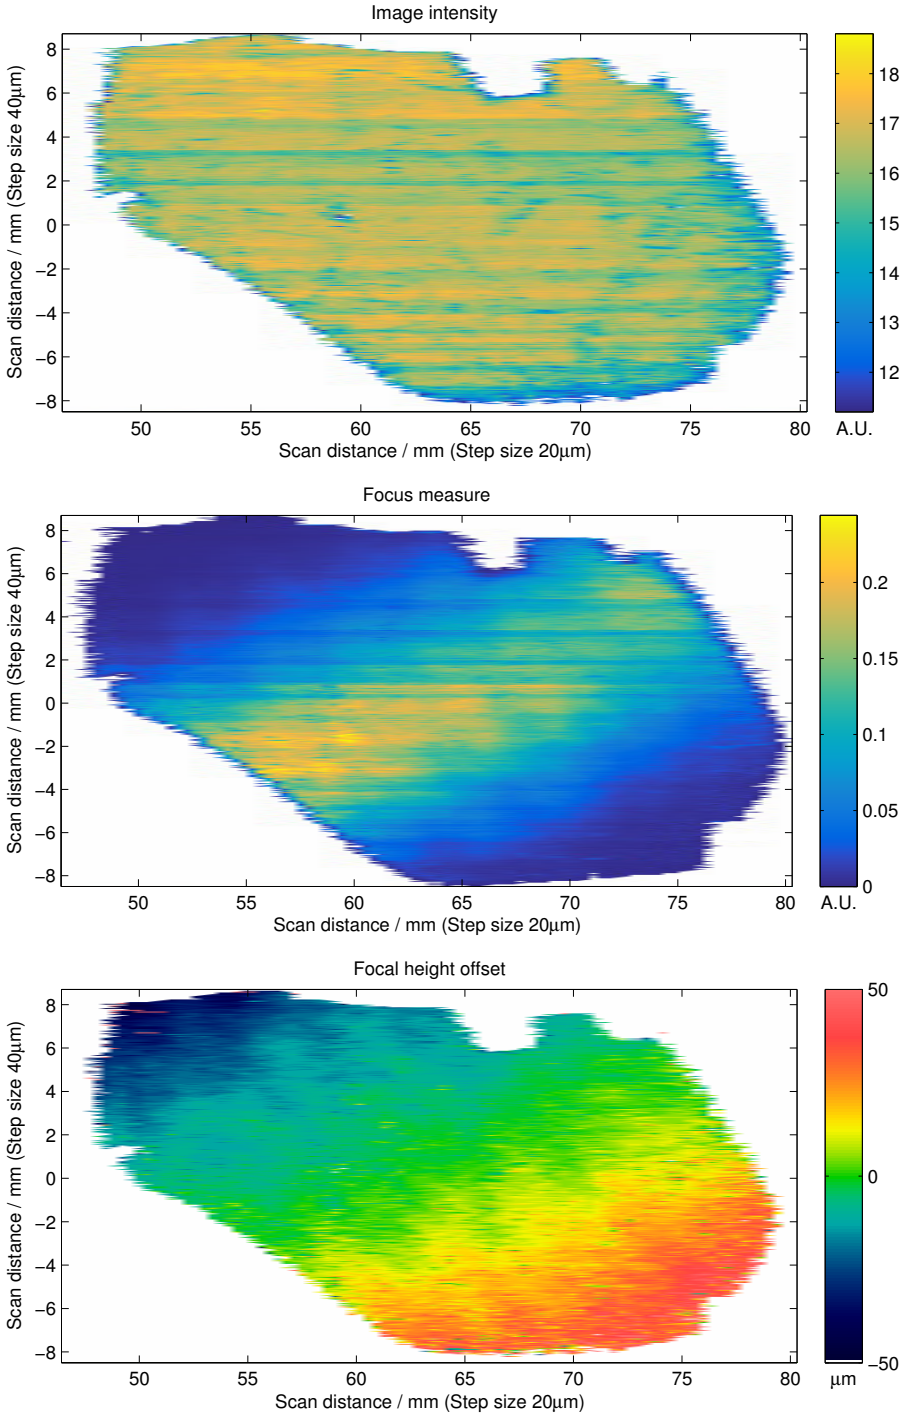


Figure 6.16: Processed generation patch image data for a scan on an isotopic sample of aluminium. The step size of images in X is ~250 µm and in Y is 50 µm. The acoustic wavelength is 23 µm and a 1.5 mm aperture has been used. The three images are: the average intensity of each generation patch image, the focus measure for each image as measured by the grey level local variance, and the focal height offset, in µm, determined from grating spacing changes.

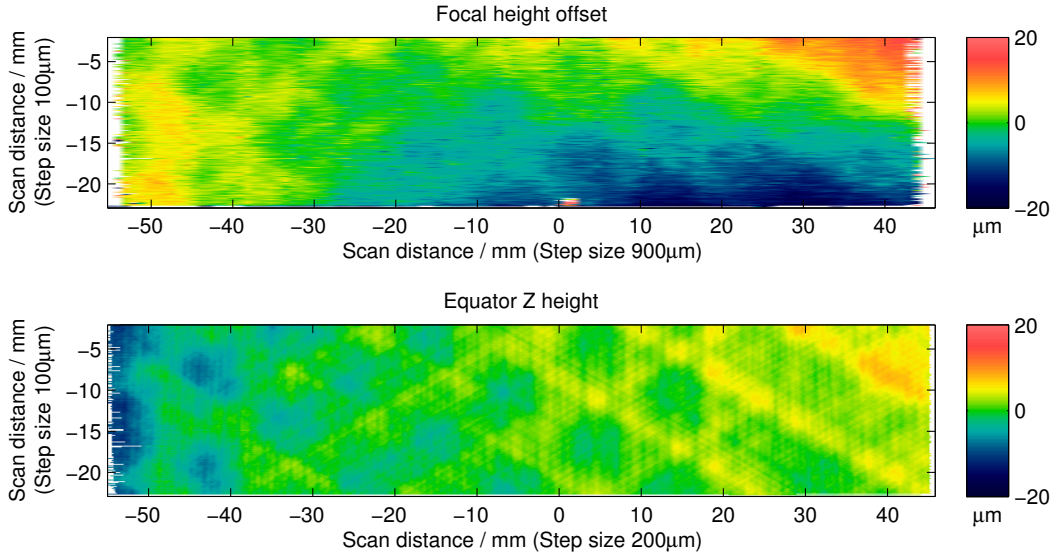


Figure 6.17: The results of a SRAS scan on a large, flat, aluminium mirror taken with an acoustic wavelength of 23 µm and an aperture of 1.5 mm. The focal height offset, measured in µm has a step size of 900 µm in X, based on the camera image frame rate, and 100 µm in Y. The values have been derived from the measured change in grating spacing. The Equator Z height, measured in µm, has a step size of 50 µm in X and 100 µm in Y and shows the reported deviation of the motion platform from the expected position.

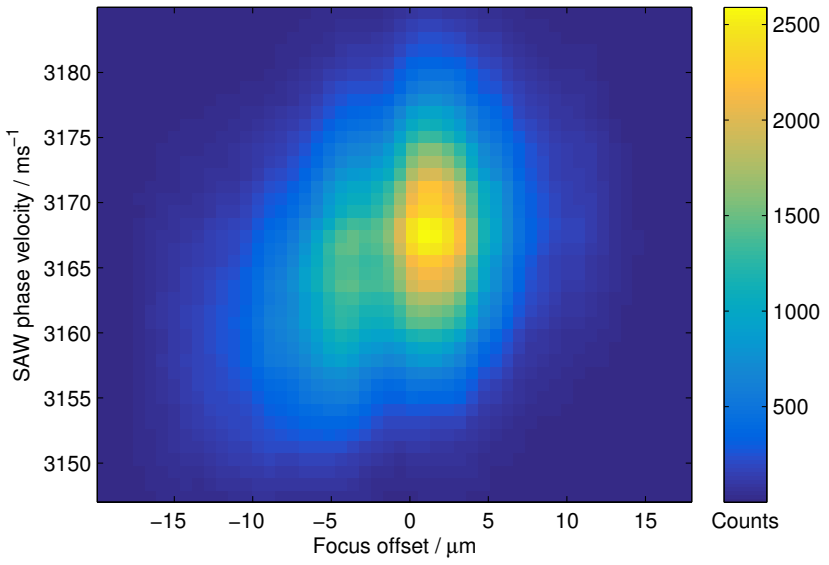


Figure 6.18: A histogram showing focus height offset against measured SAW phase velocity for the scan shown in Fig. 6.17 of a large aluminium mirror. There are 60 bins in X and 30 bins in Y and a filter has been applied to improve the clarity of the underlying trend.

6.6 Conclusions

In this chapter the initial aim was to determine the feasibility of applying image analysis techniques to the generation patch, in order to perform auto-focus. From a comparison of various focus measures, the optimum techniques were found to be gray level local variance and the variance of Gaussian derivative. Across all grating sizes tested, both of these yielded a measure which reliably indicated the defocusing of the generation patch due to focal offsets in the Z direction.

Comparisons of these focal measures to signal amplitude allowed measurements for the tolerance of SRAS to be determined over multiple grating size configurations. It was found that, based on a 50 % decrease in amplitude, the working range of SRAS was $\pm$ 25 µm for the smallest 15 µm acoustic wavelength, rising to $\pm$ 37 µm for a 30 µm acoustic wavelength.

Although it is possible to apply focussing algorithms to auto-correct for Z offsets, these can be complex to implement and require several camera frames to operate. Instead the possibility of using the changes in the geometry of the generation patch with defocus was investigated and a solution based on filtered FFTs was shown to be very effective at measuring changes in grating line spacing. By generating calibration curves for changes in grating spacing with defocus it was shown that an absolute focus offset measure could be made from a single image frame, and that a sensitivity of ~3 µm in Z was achievable.

6.7 Further work

6.7.1 Application to live-scans

So far focal offset correction measurements have only been determined off-line and have not yet been applied to a live scan. Before this can be done a few changes need to be made to the implementation of the algorithms, primarily: rewriting them in C++ from their current MATLAB implementation, and optimisation of the code to ensure minimum processing time. From this point there are two routes to implementation. The first is to use the focal data from the previous line scan to correct for the next

line scan. Since the typical SRAS patch is 2-3 times larger than the typical step size between line scans, the focus measure varies slowly between line scans, this is shown in Fig. 6.19. Therefore the data from the previous scan line is a reasonable estimate for the next line scan, but only when operating on planar surfaces.

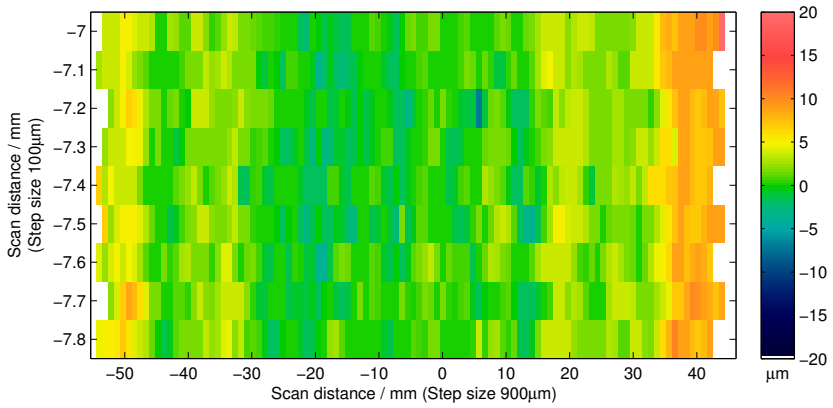


Figure 6.19: A zoomed in section of Fig. 6.17 showing the similarity of adjacent line scans.

The second route involves true live correction to the motion paths. The Equator 700 has a probe deflection feedback mechanism allowing tracking of surfaces by issuing motion corrections. Instead of a probe feedback, the focal offset measure could be used.

Chapter 7

Generation of laser ultrasound on anisotropic materials

In this chapter a model for the generation of laser ultrasound on anisotropic materials is presented and is applied to some typical engineering metals to predict SAW velocity surfaces for an arbitrary crystallographic plane. This model is a hybrid technique which uses a finite element model to compute the SAWs generated by a point source, before applying a Huygen's principle method to simulate the SRAS grating source. The advantage of this method is the vast reduction in computational time required to compute velocity surfaces where gratings at multiple propagation directions are necessary.

The motivation for developing this model is two-fold: to attempt to explain why certain propagation directions on certain crystallographic planes generate very small SRAS signals, and to ameliorate some sources of systematic errors observed in the output from the orientation determination algorithm, the outline of which was presented in §2.4. Both of these points will be expanded on in the following section.

Before the model is described a literature review is presented of existing research in modelling laser ultrasonic generation. This covers the work on point-sources on isotropic materials which then expands to finite line and point sources on samples of high crystallographic symmetry. In addition, the related field of SAW propagation on anisotropic media is discussed before recent progress in FEM is given.

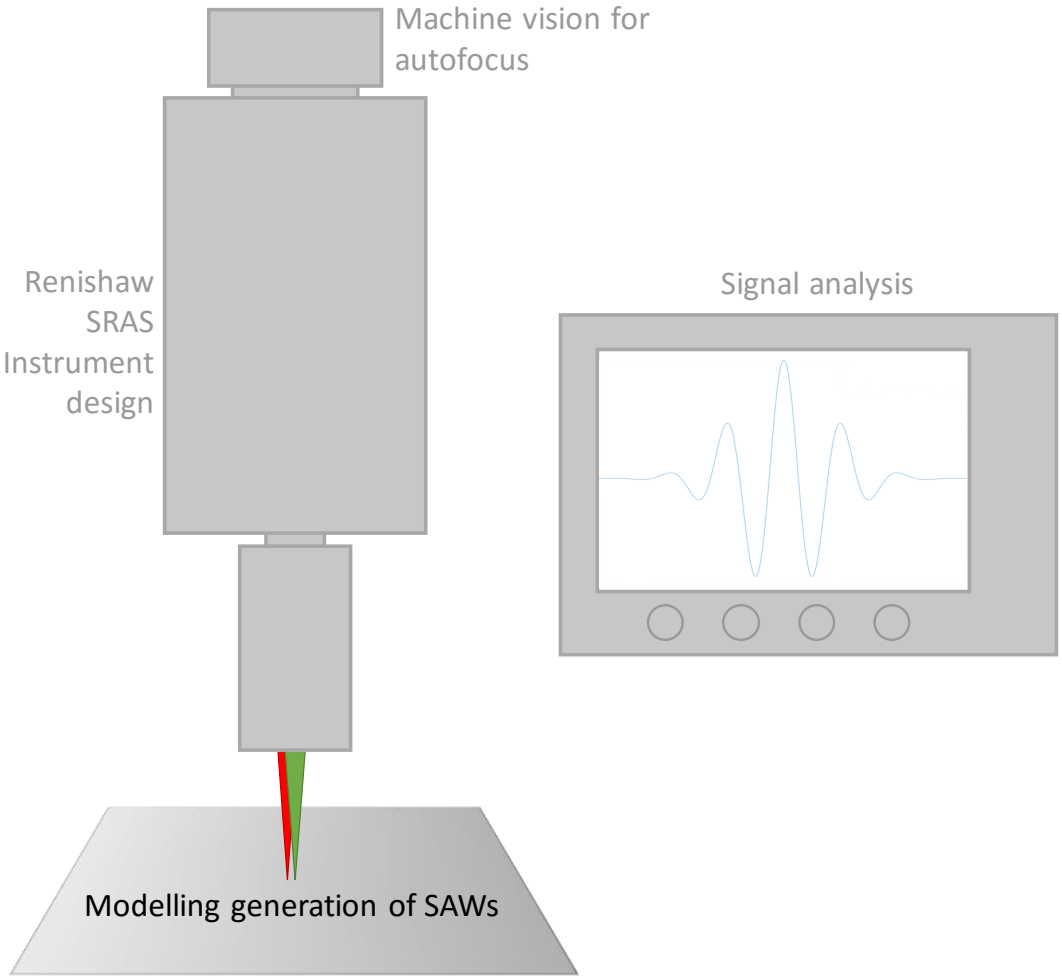


Figure 7.1: This chapter focuses on the modelling of laser generation of SAWs on anisotropic materials and developing algorithms for predicting SAW velocity surfaces.

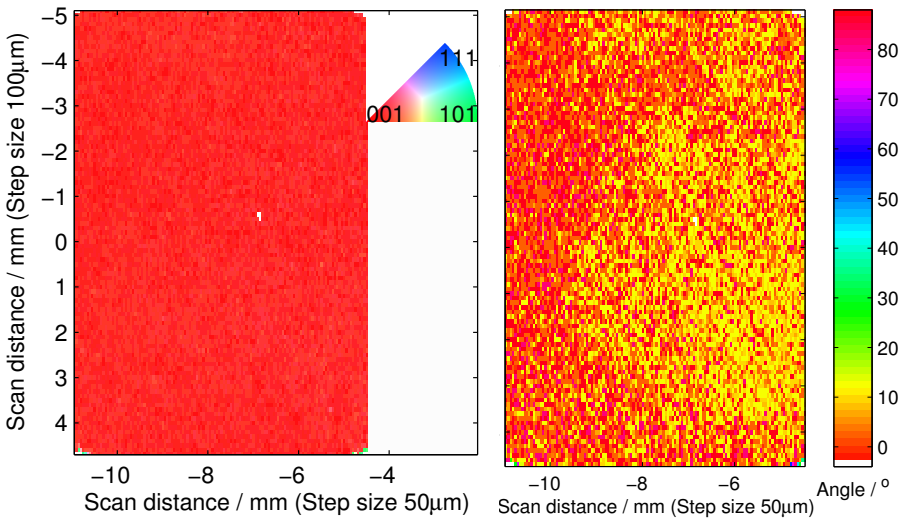
7.1 Motivation for developing the generation model

7.1.1 Issues seen with the orientation determination algorithm

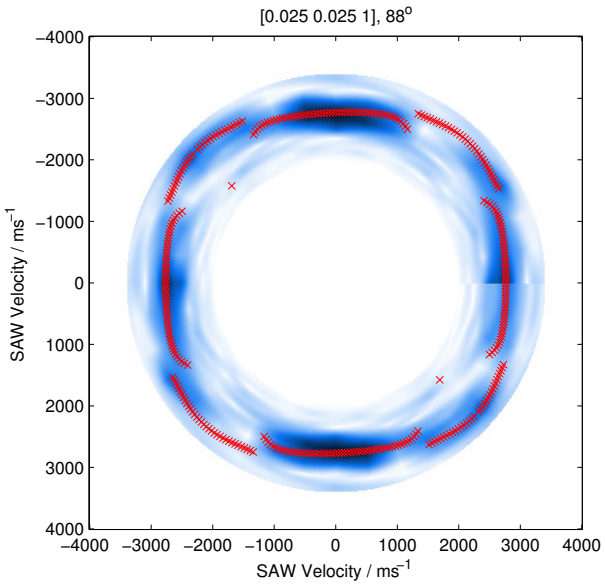
SRAS can be applied in two distinct modes. The first, and by far the quickest, is the imaging of samples for velocity contrast with many such images shown throughout this work. These images can indicate the location, size, distribution and shape of the granular structure but do not provide the actual crystallographic orientation of the sample. The second mode of operation is the use of SRAS to determine absolute crystallography over an entire sample. In general this process has been shown to work very successfully [14, 59, 63], however there are some uncommon cases where the process does not work as well as expected.

A pair of well polished, single crystal cubic nickel samples were scanned in 18 propagation directions on the EMDA SRAS system. With the known elastic constants a database of modelled velocity surfaces were generated and used to determine the crystallographic orientations of the samples. Fig. 7.2 and Fig. 7.3 (a) and (b) show the inverse pole figures and plane rotation for the two single crystal samples indicating orientation close to the (0 0 1) plane. Over the inverse pole figures the determined orientation looks very consistent while the plane rotation is slightly noisier. Together, and by first impression, these images indicate a fairly good determination process. However, when the raw data is compared to the fitted orientation it is clear that the fit is not ideal. Fig. 7.2 and Fig. 7.3 (c) show the raw data velocity spectra with the fitted forward model optimum match overlaid. Although the shape, and thus the plane, looks fairly close to correct, the rotation of the plane is systematically wrong by a few degrees in a clockwise direction. For Fig. 7.2 the rotational error is ~20° while for Fig. 7.3 the error is ~10°.

Looking at the overlaid velocity surface data gives some indications of why this error has occurred. The raw data velocity spectrum is quite broad, due to the limited number of cycles in the SRAS signal, and so there is overlap between the forward model predicted mode velocities, and the edges of the broad velocity peaks where there should be no match. As a result, false positives are added to the figure of merit used to determine optimum matching. Therefore a possible strategy for reducing the rotational errors is to reduce the chance of false positives.

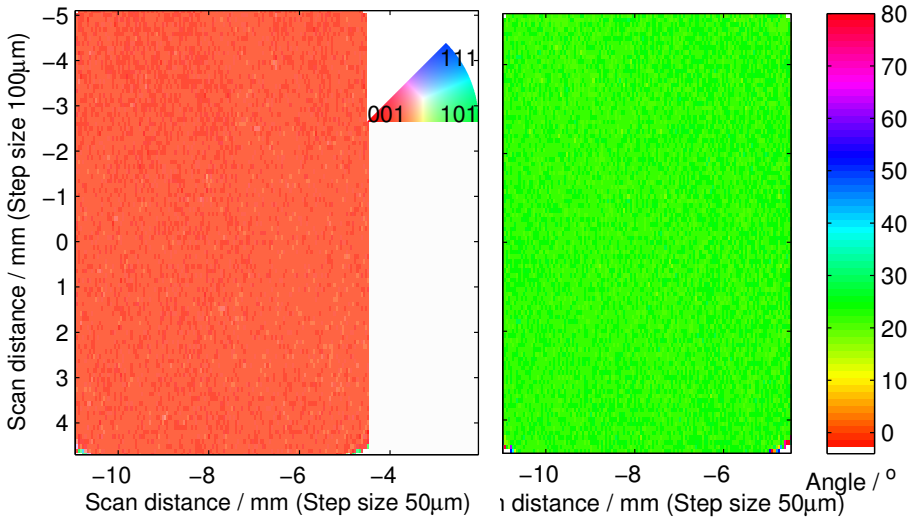


(a) (b)

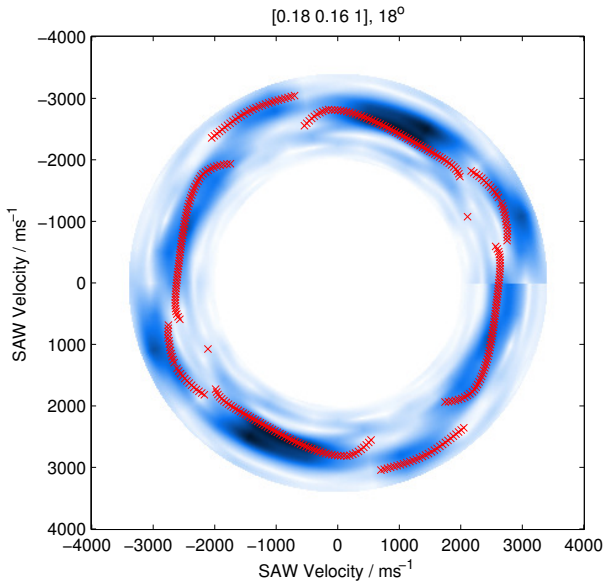


(c)

Figure 7.2: A single grain nickel sample of orientation close to (0 0 1). (a) the inverse pole figure, (b) the plane rotation, (c) the fitted plane shown against the raw data.



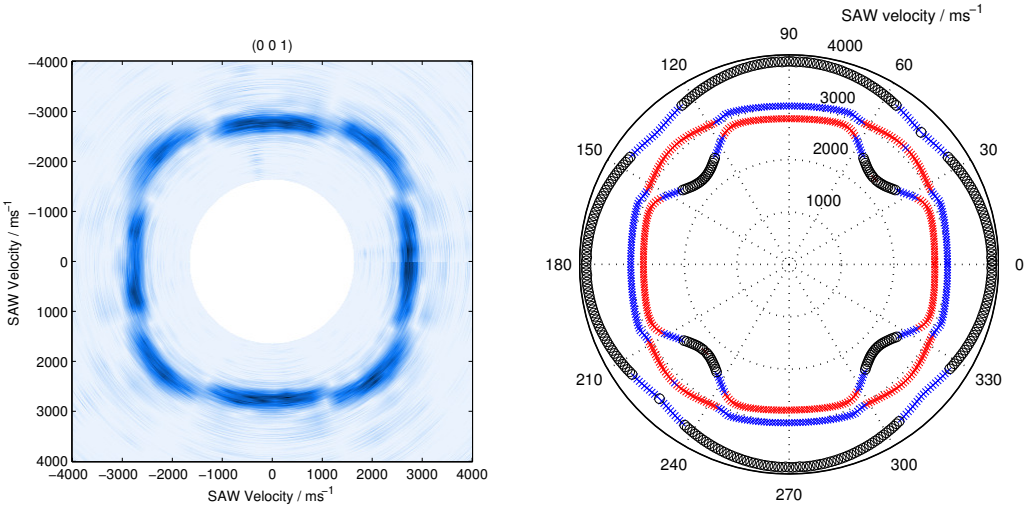
(a) (b)



(c)

Figure 7.3: A single grain nickel sample of orientation close to (0 0 1). (a) the inverse pole figure, (b) the plane rotation, (c) the fitted plane shown against the raw data.

The forward model is a purely propagational one and is capable of precisely predicting the velocities and slowness surfaces of all of the SAW modes which can exist on a particular crystallographic plane. However, when the model predictions are compared to the experimental data it is clear that not all of the possible modes are either generated or detected. Fig. 7.4 shows this discrepancy for the (0 0 1) plane of nickel. Clearly there are multiple modes predicted by the forward model while, in each propagation direction, only one of these is experimentally observed, and in some propagation directions the SAW amplitude is nearly zero.



(a) (b)

Figure 7.4: a) Experimentally obtained slowness surface for the (0 0 1) plane of nickel. (b) The forward model prediction of the (0 0 1) plane of nickel. Relative mode strength is indicated, in decreasing order, by red, blue black.

The cause of the discrepancy is the fact the forward model takes no account of how the SAWs are generated. An estimate of the ranking of each possible mode in a particular propagation direction can be made, however this does not indicate whether this mode is generated or its amplitude relative to other propagation directions. Thus, a further model is required to consider the generation mechanism involved which can then be used to ignore modes which are not generated. This will in turn reduce the probability of overlap false positives and hopefully reduce the rotational error.

7.1.2 Low generation efficiency crystal grains

As already discussed the Renishaw SRAS system has limited generation power and efforts have been made in optical efficiency and signal processing to minimise the influence this has on collected data. Another factor affecting signal strength is the local grain of propagation; on certain grains SRAS signals can have small amplitudes compared to other grains on the same sample. Fig. 7.5 shows the measured signal amplitudes on a Ti 685 sample and highlights strong changes in signal strength with the local microstructure. For scans taken with already low signal strength, i.e. with small acoustic wavelengths or on dull samples, this can yield scan images which differ greatly in quality dependent on the area of the sample, this effect is shown clearly in Fig. 5.22, particularly for the twin-photodiode scan.

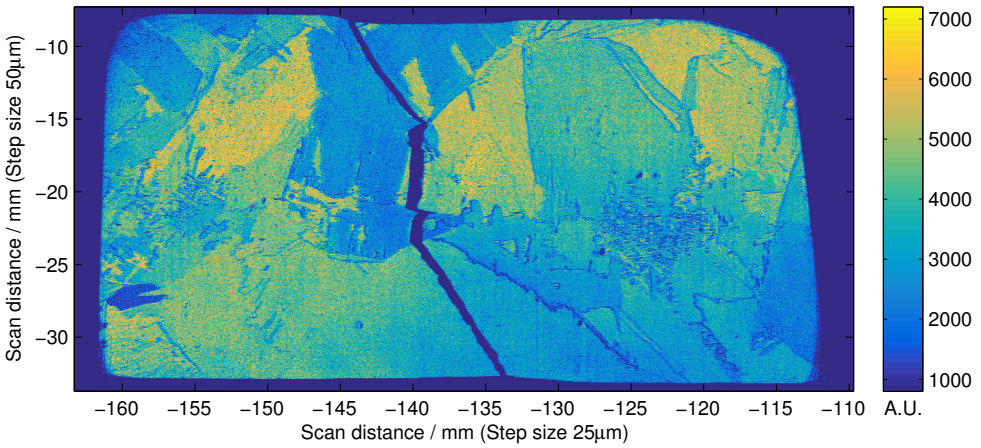


Figure 7.5: An image of a Ti 685 sample showing the high variation in SRAS measured trace amplitude across grains.

Here the generation model is of use in understanding why these areas of low signal strength occur, and to explore the possibility of altering the exact nature of generation and detection to minimise the issue.

7.2 Existing methods of modelling generation of laser ultrasound

The thermoelastic generation of ultrasound by a laser point source is conceptually a simple process. A laser beam pulse deposits heat onto a sample surface, causing rapid

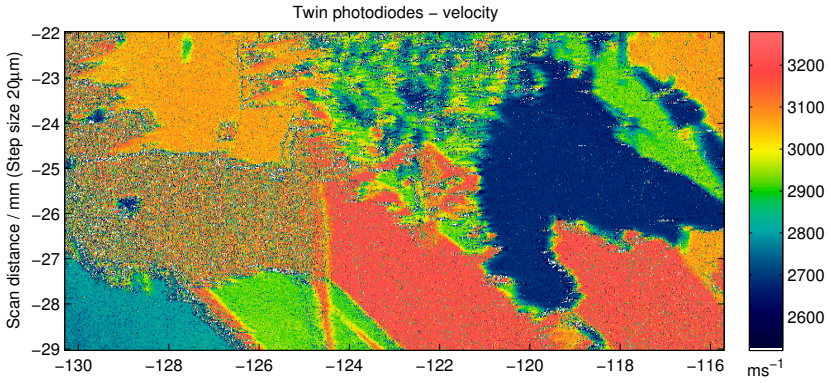


Figure 7.6: A SRAS scan taken on a low reflectivity Ti-685 sample with the twin KED. Step sizes of 10 µm and 20 µm were taken along the x and y axes respectively, propagation is along the x-axis with an acoustic wavelength of 23 µm and a 1.5 mm aperture.

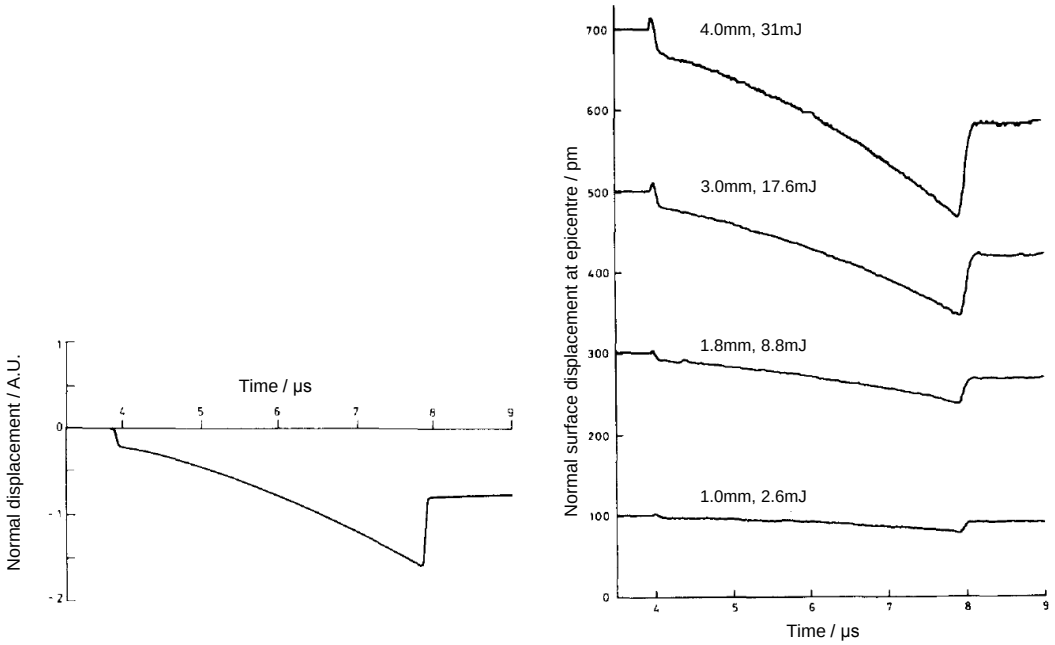
thermal expansion of the illuminated region. At the boundary between the expanding and static regions a stress field is induced which in turn generates SAWs.

Significant work has been undertaken in the field of modelling acoustic waveforms, amplitudes and directivity patterns generated by lasers on solid media. In this section a review of this work is presented with a particular focus for the generation of SAWs, and generation on anisotropic media. Due to the intended application of a model to the SRAS technique only nanosecond temporal width generation models are discussed, there is also a significant volume of work concerning picosecond laser ultrasonic models, see for example [105, 106, 107].

7.2.1 Isotropic media

Similar problems to that of modelling the generation of laser-ultrasound have been studied since before lasers were in common use. White, in 1963, built a thermal stress-strain model that described the waves produced by the harmonic heating of the entire surface of a sample [108]. He demonstrated that the amplitude of the stress-wave varied proportionally with the absorbed energy density, and was heavily dependent on the thermal and elastic material properties. His experimental work was conducted with an electron ray rather than laser source, but the underlying principle is compatible with a laser source. While experimental work on laser generated acoustic waves continued through the 60s and 70s [109, 110, 111], it was not until the early

80s that further models were developed. In 1980 Scruby et. al. developed a model for a pulsed laser source on an isotropic material [112]. Instantaneous and uniform heating of a small surface volume element were assumed, and the temperature rise, and subsequent expansion, calculated. The expansion was shown to be equivalent to the insertion of a small volume element and was modelled as a set of edge dislocation loops. Using the material elastic constants, the distribution of dislocation loops was related to a distribution of force dipoles. Then, using a suitable Green's function, the epicentral wave amplitude at the far side of the sample was predicted, shown with experimental data in Fig. 7.7. In agreement with White they found that the acoustic wave amplitude varied proportionally with the incident energy density, and found that their predicted waves agreed reasonably well with experiment although failed to predict the fast initial increase in amplitude before the slower decrease, known in the literature as the precursor.

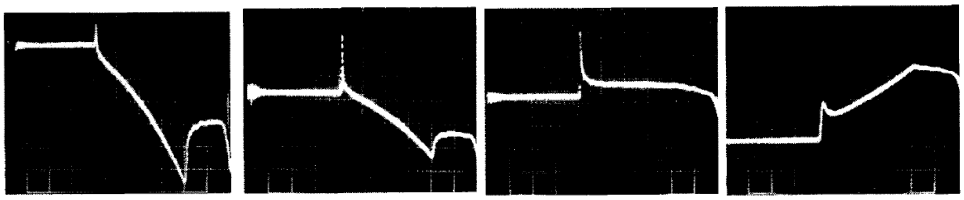


(a) (b)

Figure 7.7: a) Prediction of the measured epicentral displacement from a thermoelastic laser source. (b) Experimentally measured epicentral displacements for a range of thermoelastic laser power densities. These bulk waves were propagated through a sample and detected on the opposite side to generation. Adapted from [112].

This work was continued and it was found that the wave-form produced in the ablative

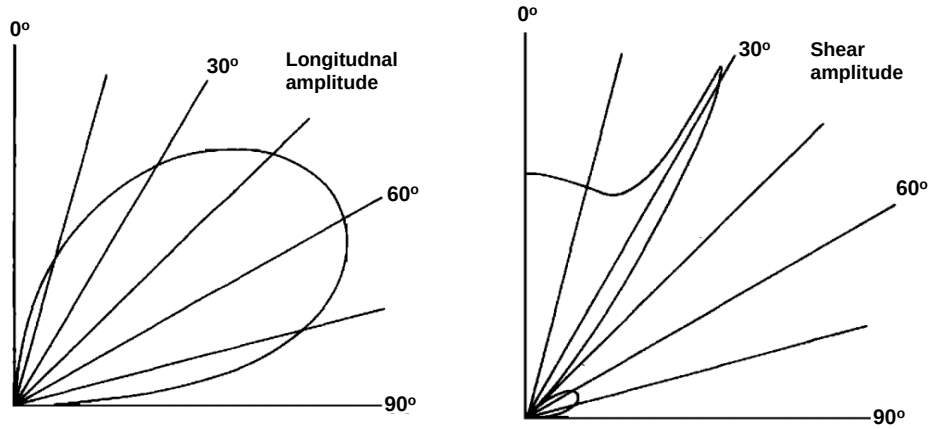
regime is significantly different to that in the thermoelastic [113]. As the power density is increased to the point of generating a plasma, the contribution of the thermal expansion becomes negligible, and the amount of energy directed into the sample increases. This process was modelled as the combination of a thermoelastic expansion, and an ablative force, and the relative amplitude of each was estimated by fitting the model to the experimental data, shown in Fig. 7.8. It was found that as the energy density increases into the ablative regime, the contribution of the thermoelastic expansion to the measured wave-form becomes negligible. This is most likely caused by the proportion of optical energy absorbed to generate the plasma increasing. Hutchins et. al. expanded the thermoelastic point-source model to predict the acoustic wave amplitude at a range of angles into the sample, produced by disk and line-sources [18]. This model assumed a disk source on the surface of a sample with a uniform and harmonic normal force. Directivity patterns for both the shear and longitudinal waves were predicted, shown in Fig. 7.9 and these showed good agreement with the experimental results. Further results and predictions were also presented for the directivity patterns of a line source. The dramatic change in the longitudinal and shear wave forms generated in the thermoelastic and ablative regimes was again noted and explained by the stress free surface condition.



(a) (b) (c) (d)

Figure 7.8: Experimental results for the epicentral displacement measured for increasing incident energy densities. a) $9 \times 10^7\,\mathrm{W\,cm^{-2}}$ (b)$1.3 \times 10^8\,\mathrm{W\,cm^{-2}}$ (c) $3.2 \times 10^8\,\mathrm{W\,cm^{-2}}$ (d) $4.5 \times 10^8\,\mathrm{W\,cm^{-2}}$. c) shows the cross-over point between the thermoelastic and ablative regimes. Note that the scale of (c) and (d) is 5 times that of (a) and 3 times that of (b). Taken from [113].

Rose, in 1984, took a slightly different approach to the model. Assuming a point-source in the thermoelastic regime, for an isotropic material, the resultant wave forms were predicted using a volume source theorem [114]. By way of an extremely thorough mathematical analysis, Rose showed that the point-source assumption is valid under



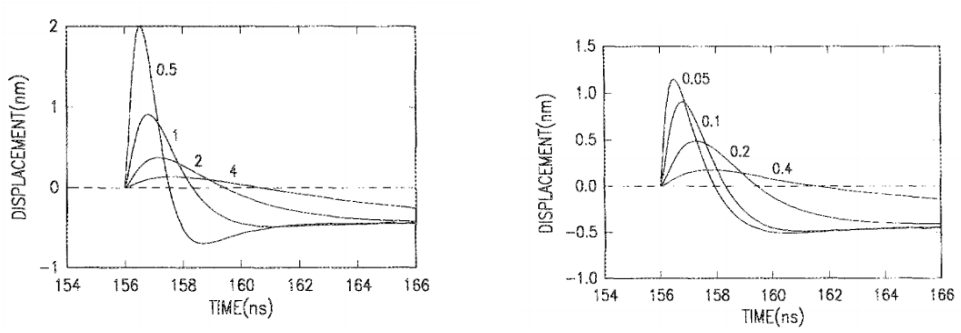
(a) (b)

Figure 7.9: Theoretically derived directivity patterns produced by a infinite line source on aluminium in the thermoelastic regime. Longitudinal (a) and shear (b) are shown respectively. Adapted from [18]

limits dependent on the wavelength, source size, pulse length and elastic properties. He named this model the surface centre of expansion (SCOE) and demonstrated excellent agreement between the predicted wave forms and experiment. However, this excellent agreement was limited to well focused and short pulses, and, as with the model by Scruby et al, it failed to predict the initial small peak in the epicentral longitudinal wave-form, known as the precursor [113]. This precursor had previously been seen in models by Dewhurst et. al. but only under the assumption of an ablative force, and not as a result of a purely thermoelastic model.

Up to this point the models developed, while fairly successful, made assumptions which limited their applicability. In particular, many models assumed no thermal diffusion and typically relied on a point-source assumption [20]. However, for long pulse lengths, in the thermoelastic regime, these assumptions are no longer valid. In the early 90s another set of models were developed, by several authors, which took full account of the finite temporal and spatial dimensions of the laser pulse as well as thermal diffusion [115, 116, 117]. These models successfully predicted the precursor in the longitudinal wave-form, shown in Fig. 7.10 [116], with the presence of the precursor attributed to the effects of thermal diffusion at the point of generation.

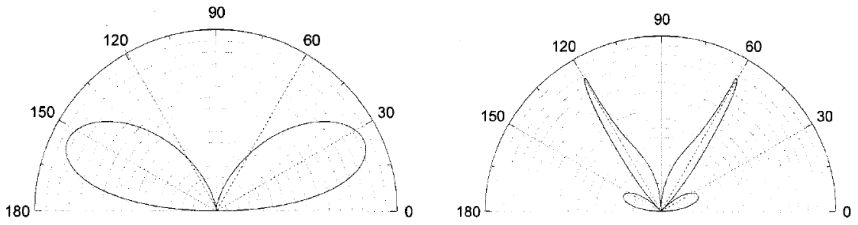
Future models generally focused on refinement of previous models and their appli-



(a) (b)

Figure 7.10: Modeled prediction of the precursor seen in the longitudinal wave form. (a) various laser pulse times in nm. (b) various beam radii in mm [116].

cations [118, 119, 120], and on developing models for laser line sources [121]. The advantages of a line source, or set of line sources, are primarily the increase in signal to noise, and the generation of planar SAWs to probe surface defects [122]. Although Hutchins had developed a model for a line source previously [18], it did not allow for finite width or thermal diffusion. By adopting an infinite length line source, Bernstein and Spicer were able to correctly predict the bulk wave directivity patterns [123]. Their method was to formulate the displacements, and solve them using Laplace transforms, the results are shown in Fig. 7.11. They found that in two dimensions the directivity pattern form and amplitude was identical to those of a point source, but that the temporal wave-field was different. Of note here is the significant difference in the epicentral longitudinal amplitude to that predicted by Hutchins, shown in Fig. 7.9. This is a result of the difference in the models; Bernstein and Spicer use a full volume of thermal heating while Hutchins employed a surface only model. This line model, however, did not take into account thermal diffusion and the finite temporal dimensions of the source. Arias and Achenbach derived a two dimensional model capable of accounting for thermal diffusion, optical penetration and source pulse shape [91, 124]. By analytically solving the full thermoelastic problem, they were able to predict the wave-forms and directivity patterns and showed excellent experimental agreement, predicted SAW profiles are shown for a line focus and point focus in Fig. 7.12.



(a) (b)

Figure 7.11: Directivity patterns of thermoelastically generated acoustic waves predicted by Bernstein, (a) longitudinal, (b) shear [123].

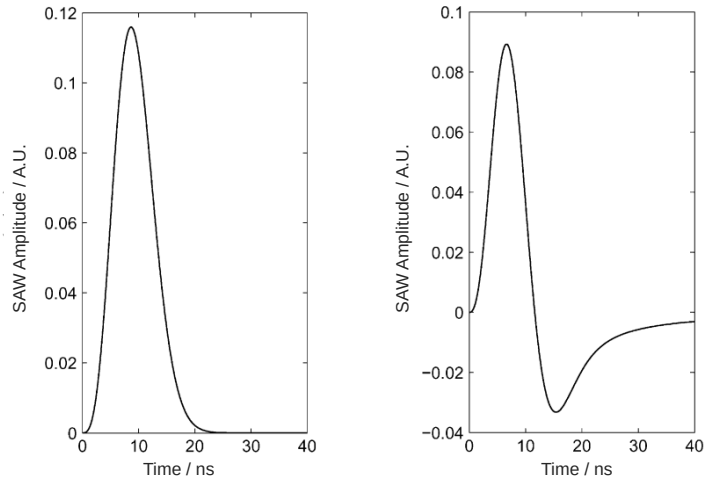


Figure 7.12: Predicted SAW amplitude from an (a) line and (b) point laser source [91].

7.2.2 Anisotropic media

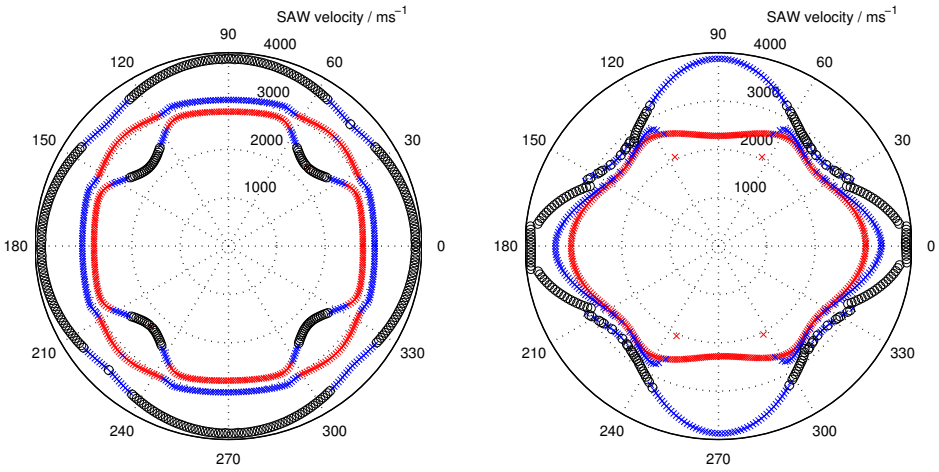
Isotropy has been assumed for nearly all works cited so far. For isotropic material the elastic properties are identical in every direction of propagation and hence, only one solution for the stress-strain boundary condition is required. Anisotropic materials, conversely, require unique solutions for each propagation direction, on every possible plane, and are a much more complicated problem.

Some of the earliest work in predicting surface wave velocities and modes was done by Stoneley in 1949 in the field of seismology [125]. Stoneley was interested in the differences between seismic waves propagating through isotropic and transversely isotropic media. He showed that Rayleigh waves can exist on the basal plane of transversely isotropic materials, with the same velocity and amplitude in all propagation directions. He later extended this treatment of Rayleigh and Love waves to cubic materials,

specifically the (1 0 0) plane [126], on which he demonstrated that Rayleigh waves propagating in the [0 1 1] and [0 0 1] directions follow elliptical particle motion, with amplitudes heavily dependent on the elastic constants. Whilst possible on some materials, he showed that for certain materials, such as copper, no Rayleigh wave exists in the (0 1 1) direction. This treatment extended only to a few directions on a single plane of a cubic material, and involved significant labour for each new propagation direction calculation. In addition, the neglect of an entire class of solutions, leading to the prediction of no Rayleigh waves, was shown to be erroneous, and this set of solutions was later shown to be valid [127].

With the advent of greater computing power, numerical solutions for more general directions became practical, and this led to a flurry of work through the 1960s. For the basal plane of cubic materials Rayleigh waves were shown to exist for a wide range of elastic constants and directions [128]. Further work on a variety of materials, including iron, copper and aluminium, showed that for the planes (1 0 0) and (1 1 0) there existed regions in which the surface wave degenerates into a shear wave [129, 130]. Continuing work began to focus on lower symmetry conditions and it was soon shown, contrary to earlier beliefs, that the existence of Rayleigh waves on arbitrary planes is extremely likely [131, 132]. In 1970 Farnell produced a full treatment of the Rayleigh wave problem capable of predicting all possible modes, for any propagation direction, on any plane of any material [133]. His method was to iteratively seek minima in the determinant of the elastic wave equation given an appropriate set of boundary conditions, and as well as predicting the possible Rayleigh modes, it also predicted the pseudo-Rayleigh waves. A full description of this method can be found in references [133] or [14]. Fig. 7.13 shows the highest amplitude three modes on the (0 0 1) plane of nickel predicted using Farnell's method.

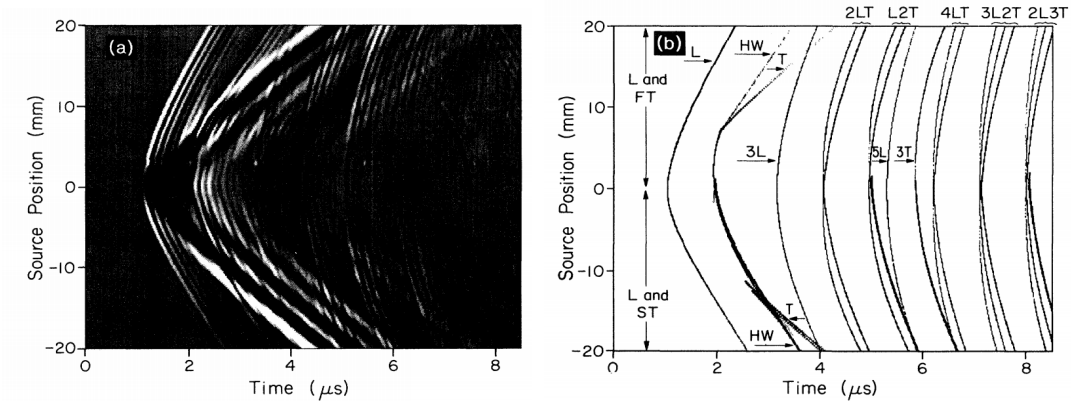
Whilst Farnell's method, and the work leading to it, provide predictions in excellent agreement with experiment, they do not take account of the generation mechanism and thus cannot yield the amplitude of each mode. In order to predict the mechanical displacement of each wave mode the wave equation must simultaneously satisfy three coupled equations of propagation, and three boundary conditions [134]. Due to this complexity the focus on predictions of amplitude largely followed that of the progress in the first Rayleigh wave work, and early progress was made on only high symme-



(a) (b)

Figure 7.13: Numerically derived slowness surfaces for first three modes on the (a) (0 0 1) and (b) (0 1 1) planes of nickel. The modes are ordered as red, blue, black which indicates decreasing strength.

try planes, with propagation typically along crystal axes. Every et al. adopted an approach, for point sources on high symmetry planes of silicon, which used a Monte Carlo ray simulation to predict an image, shown in Fig. 7.14(a). For the planes (1 1 1), (1 0 0) and (1 1 0) they showed good agreement with experiment but were not developed beyond these high symmetry cases [135].



(a) (b)

Figure 7.14: SAW propagation on a nickel [1 1 1] plane from (a) experiment and (b) a Monte-Carlo ray simulation model [135].

The majority of other research tended to use the Cagniard-de Hoop method [136, 137]

which was developed in the 1960s for applications in seismology, ostensibly for layered media based on Laplace transforms. This method was successfully applied to some simple examples of anisotropic materials; either on very high symmetry planes or conditions of transverse and orthotropic isotropy [138, 139, 140, 141]. However one thing all of these models assumed was the laser source acting as a stress dipole, a condition which does not hold for low symmetry planes. Since the temporal evolution of the SAW waveform is heavily dependent on the source characteristics, an analytical model valid for the case of arbitrary anisotropy cannot be built using these methods [142].

7.2.3 Finite element

With the constantly growing availability of computer power, finite element methods (FEM) for prediction of acoustic waveforms emerged in the mid 2000s, initially for a constrained layer problem [143] and also for the thermoelastic case [144]. While still isotropic, these models took full account of finite laser temporal and spatial profiles and provided a framework for future development. This led to simple transverse isotropy cases being studied for applications to fibre composites [145, 146, 147]. Through these works the applicability of FEM to the problem of anisotropic laser-ultrasonic generation was shown to be good over a wide range of source parameters.

7.3 The generation model

Based on the findings of the literature review, FEMs were identified as the only avenue for developing a model of generation of SAWs on an arbitrarily anisotropic media. This thermoelastic laser generation model is composed of two distinct steps; modelling the SAWs produced by a point source on an anisotropic medium, and applying Huygen's principle to predict the SAWs produced by various geometries of source, in various propagation directions.

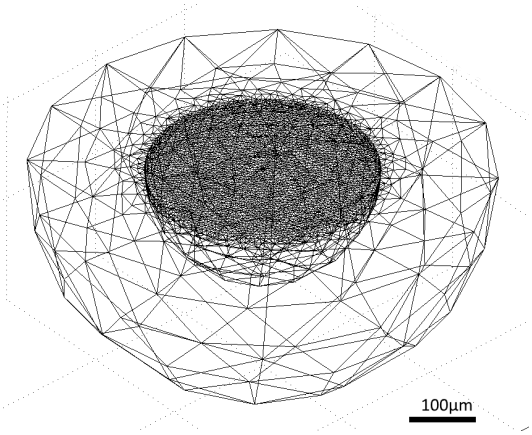


Figure 7.15: A 45000 element mesh of two hemispheres.

7.3.1 Finite element model of a point-source

The first step in the model is performed in the finite element software Comsol and was initially developed in collaboration with Dr. Xuesheng Chen of the Applied Optics Group (AOG) at the University of Nottingham. A half space is defined by a hemisphere which is meshed finely near the centre and coarsely toward the edges, the size of the hemisphere is much larger than the region of interest for generation. The mesh is triangular and in total contains approximately 45000 elements concentrated at the flat surface, the meshed hemisphere is shown in Fig. 7.15. Onto this mesh the material constants are applied with the relevant elastic constants being rotated to the crystallographic plane being modelled, the rotation matrix, elastic constants and thermal constants used are given in Appendix C. The laser pulse is simulated as a heat source applied at the centre of the hemisphere with a spatial profile described by a 2D Gaussian, $G(x, y) = exp(-(x^2+y^2)/r_0^2)$, where r_0 describes the width. Temporally the heating is described by a sharp asymmetric Gaussian pulse, $I(t) = (8t^3/v^4)exp((-2t^2)/v^2)$ [148], where v is a characteristic time.

To correctly simulate the generation laser used in SRAS, and a typical grating configuration, the temporal pulse width is set to 2.5 ns and the spatial beam width to 12 µm. Since it is small compared to the mesh size and lateral beam width, the beam extinction value is considered to be infinitely small, i.e. all absorption occurs at the surface. With the material and source defined, the FEM is then applied in two stages. Firstly, the 'heat transfer by conduction' module of Comsol is used to predict the thermal diffusion and expansion of the illuminated region. Since the materials mod-

elled are anisotropic, thermal expansion is anisotropic [149] and so this step cannot be performed analytically. The second stage applies the 'solid stress-strain' module to calculate the stress fields resulting from the thermal expansion, and in turn the particle displacements. The generated SAWs are then propagated across the surface of the hemisphere, with the surface displacement in three orthogonal axes determined for each mesh point and time step. Time steps are of length 0.1 ns with a total of 400 time steps taken.

From the arrays of surface particle displacement the out-of-plane displacements were extracted, as this is the only motion to which KED is sensitive to. For the first run of the model the (0 0 1) plane of hexagonal titanium was used as it is known to be isotropic. Fig. 7.16 shows the out-of plane displacement for this simulation at three points in time which broadly represent three phases in the model: initial thermal expansion, the near-field, and the far-field. In all cases the generated SAWs are axisymmetric although the noise introduced by the fairly coarse mesh is randomly distributed. Comparing the form of the SAW to those predicted by Achenbach in Fig. 7.12 shows that, for the far-field, there is agreement, although the initial positive displacement from the FEM simulation is much smaller than that predicted by Achenbach, and the random noise exists only in the FEM model. Between the near and far field surface displacements there are significant differences. The multiple wave modes which exist in the near-field are rapidly attenuated leaving only a single cycle pulse at larger propagation distances. Consistent with the directivity patterns predicted in Fig. 7.9 and Fig. 7.11, there is no indication of the shear wave mode in the simulation.

With validation of the model for the isotropic case indicated, anisotropic cases were tested. Fig. 7.17 and Fig. 7.18 show point source simulations at three times for the (1 1 1) and (0 0 1) planes of cubic nickel respectively. As with the isotropic case, the simulations are shown at three phases. For the (1 1 1) plane the results are promising; the expected 6-fold symmetry has been returned. Again, there are significant differences between the near and far field with the initially complex waveforms decaying into a single wave-mode. Interestingly, the symmetry of the heated region is only 3-fold which suggests that the initial stress field is 3-fold symmetric, while the dominant SAW propagation mode is 6-fold symmetric. For the (0 0 1) nickel plane the results are equally promising; the expected 4-fold symmetry has been returned. The point sources

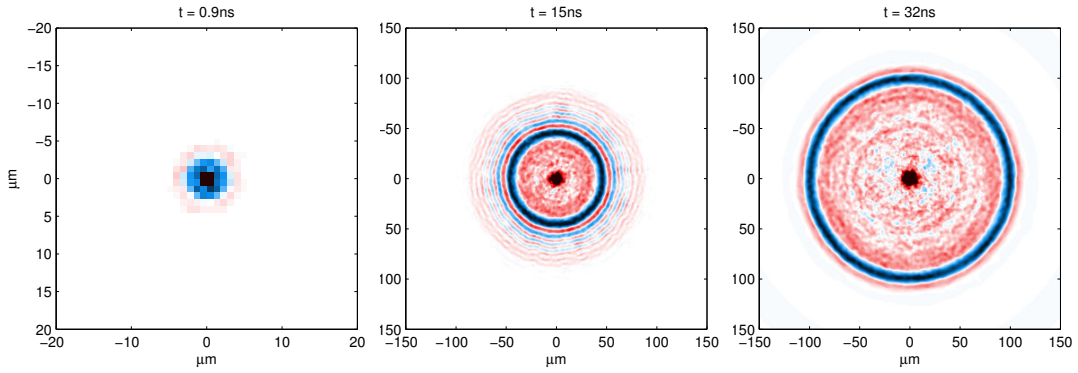


Figure 7.16: Out-of-plane surface displacement generated by a laser point source at times 0.9 ns, 15 ns, 32 ns for the (0 0 1), basal, plane of hexagonal titanium.

all show variations in SAW velocity with propagation direction as well as the variation in amplitude, indicating that the model is generating the required information. For example, on the (0 0 1) plane, the SAW propagating along the x-axis appears to have greater amplitude than the SAWs propagating at 45° to the x-axis.

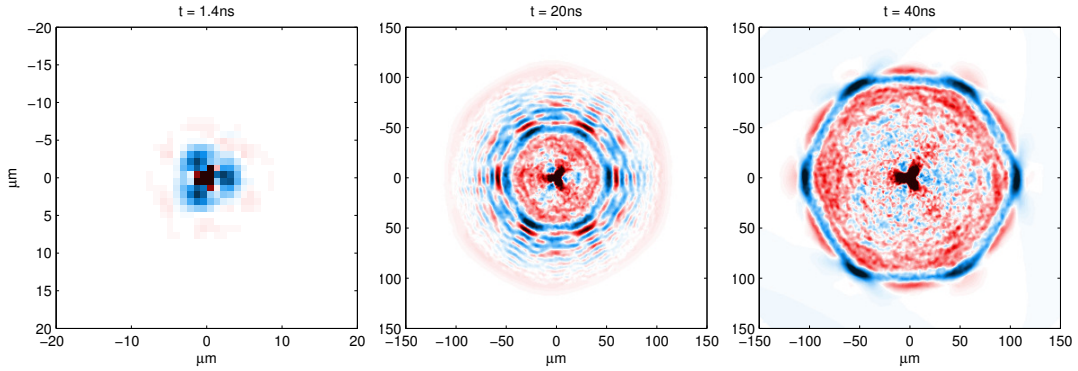


Figure 7.17: Out-of-plane surface displacement generated by a laser point source at times 1.4 ns, 20 ns, 40 ns for the (1 1 1) plane of cubic nickel.

7.3.2 Modelling a grating

With the ability to simulate point sources models on any given plane with only a set of elastic constants required, the next stage of the model can be addressed. The second part of the model, coded in MATLAB, takes point sources and uses them to build any desired geometry source by superposition; for application to SRAS the desired

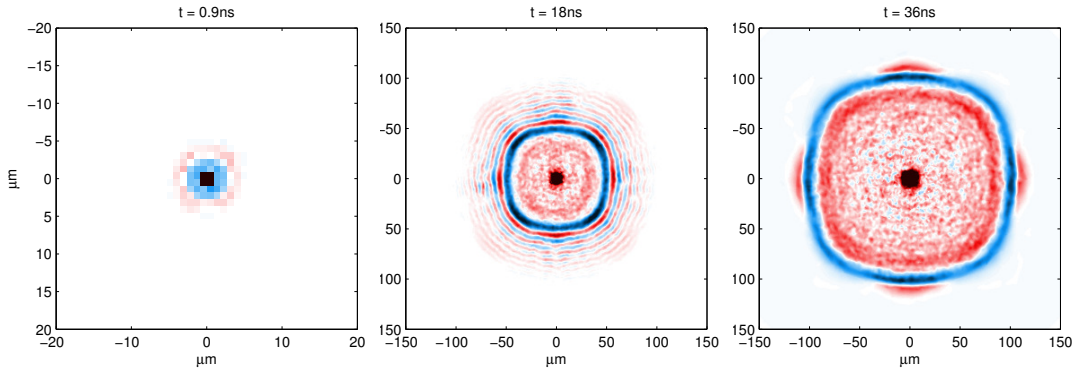


Figure 7.18: Out-of-plane surface displacement generated by a laser point source at times 0.9 ns, 18 ns, 36 ns for the (0 0 1) plane of cubic nickel.

geometry is a diffraction grating with equally spaced line sources. Although it would be possible to do this by FEM, it would require much greater computational resources than a point source, and would have to be repeated for every propagation direction. Instead this process takes advantage of the Huygen's principle, on the basis that it has been shown to work for analytical models [114].

Constructing the grating geometry is fairly trivial, the meshed point source is interpolated on to a 2D regular grid with spacings similar in size to the mesh spacings, mesh points not on the hemisphere surface are discarded. This grid is then rotated to the desired propagation direction and multiple sources are tiled into a line, typically 100-200 point sources are used to generate a single line. Rotation is defined from the central point of the point source with the positive Y axis set as 0°, and rotations defined in a clock-wise direction. Fig. 7.19 shows a tiled line source at two times on the (0 0 1) plane of nickel, with a 0° rotation. As with the point sources there is a clear difference between the near and far-field wave modes. For this highly symmetric plane, the line-source generates a symmetrical SAW propagating orthogonally from the source.

To generate the final grating a set of the lines are tiled with a given spacing which represents the effective acoustic wavelength. Emulating the typical SRAS configuration, a trace is then measured at a point approximately 2-3 acoustic wavelengths in front of the grating. This is shown in Fig. 7.20 for a grating on the (0 0 1) plane of nickel for both the 0° and 45° propagation directions. In the SRAS experiment signals are high and

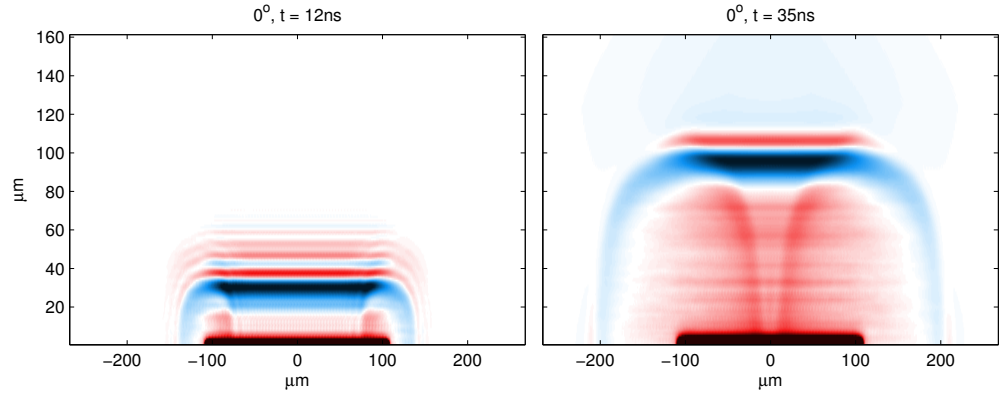


Figure 7.19: Out-of-plane surface displacement generated by a simulated laser line source at times 12 ns and 35 ns for the (0 0 1) plane of cubic nickel.

low-pass filtered and the same process has been applied to the simulated signals. Note that although the y-axis shows amplitude in nm this has not been calibrated and is not a reliable measurement. The simulated signals strongly resemble a typical SRAS signal; they both show a sinusoidal behaviour modified by a Gaussian-like window. Comparing the two signals shows that the respective amplitudes and velocities shown by the point source have been carried through to the final simulated signal with the 45° signal of lower amplitude and greater velocity - indicated by the higher frequency.

7.3.3 Generating SAW velocity surfaces

To generate a SAW velocity surface for a given plane, the procedure for measuring a grating trace is repeated with the point source systematically rotated between 0° and 180°, the two-fold symmetry of cubic and hexagonal crystal structures is exploited to generate the full 360° surface. An FFT is taken of each trace and the acoustic wavelength, set by the grating spacing, applied to give a velocity spectrum. Fig. 7.21 shows an experimentally derived velocity surface and a simulated velocity surface for the (0 0 1) plane of nickel. The experimental data for this, and the other velocity surfaces shown, were collected by Dr. Wenqi Li and Dr. Steve Sharples. On the whole the model results seem quite promising; both the velocities and amplitudes match well to the experimental data. In addition the separation between modes, where the (0 0 1) plane goes from the slower pure Rayleigh wave mode to an alternate faster SAW mode, is reproduced. The peak widths of the simulated velocity surface are wider than for

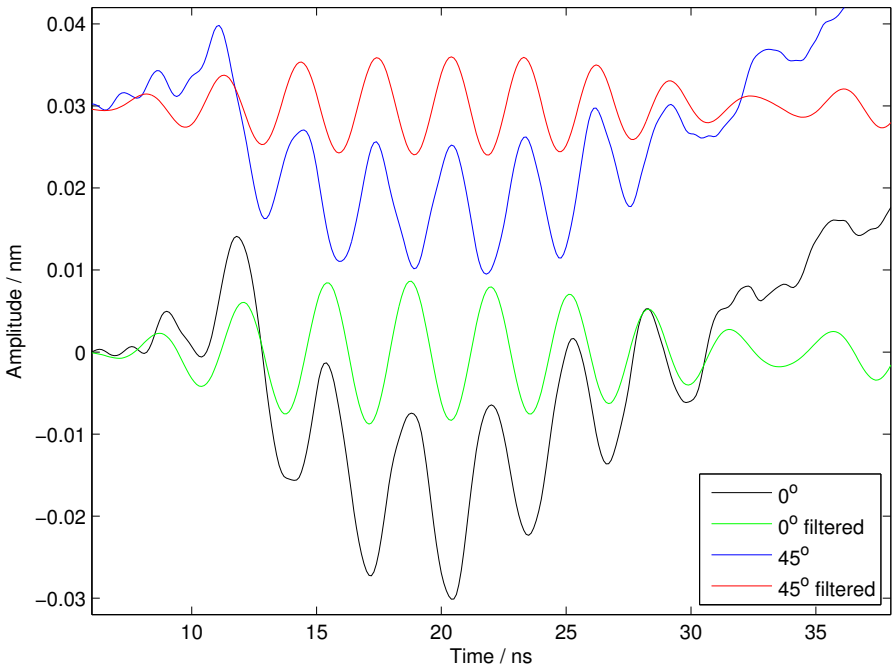
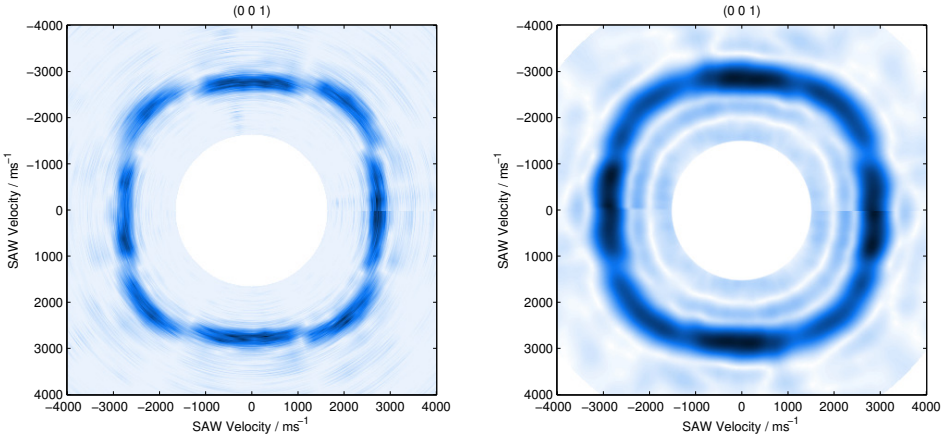


Figure 7.20: Simulated detected trace from a grating on the (0 0 1) plane of cubic nickel in the 0° and 45° propagation directions. The raw signals are shown along with the filtered version of the signal. An offset has been artificially added to the 45° signal.

the experimental data due to the fewer cycles in the simulated signal. As a result of this extra width, the transition between modes is slightly blended giving poorer contrast. Like the point source the 4-fold symmetry expected on the (0 0 1) plane has been reproduced.

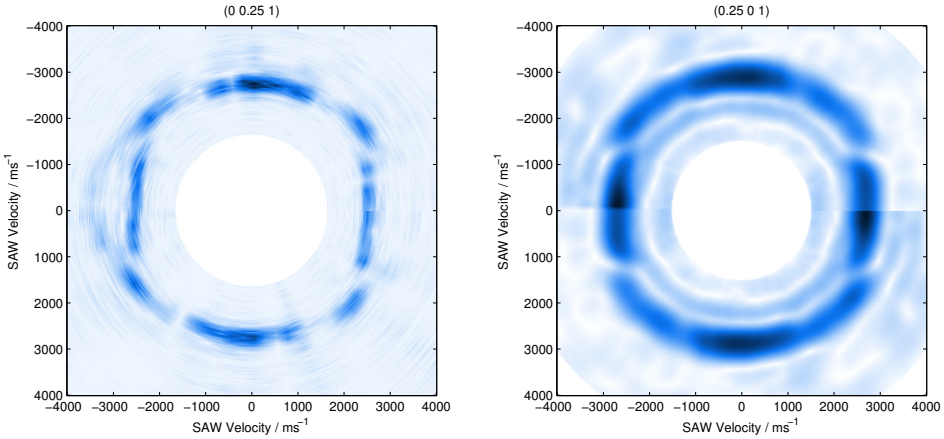
For the fairly simple case of high symmetry planes the generation model works quite well, and so some less symmetrical planes were simulated. Fig. 7.22 and Fig. 7.23 show experimental and simulated velocity surfaces for the (4 0 1) and (3 10 1) planes of cubic nickel respectively. Both of these planes are highly anisotropic with only 2-fold symmetry. For the (4 0 1) plane the results are fairly good; the SAW phase velocities are approximately correct and the relative amplitudes also closely matched. Where the dominant mode changes is visible, but again slightly obscured by the width of the velocity peaks. This issue is particularly evident in the 135° propagation direction - defined by a clock-wise rotation from 0° at positive y-axis - where the large separation between modes in the experimental data is not replicated.

Finally looking at the (3 10 1) plane shows the worst agreement of the three planes seen



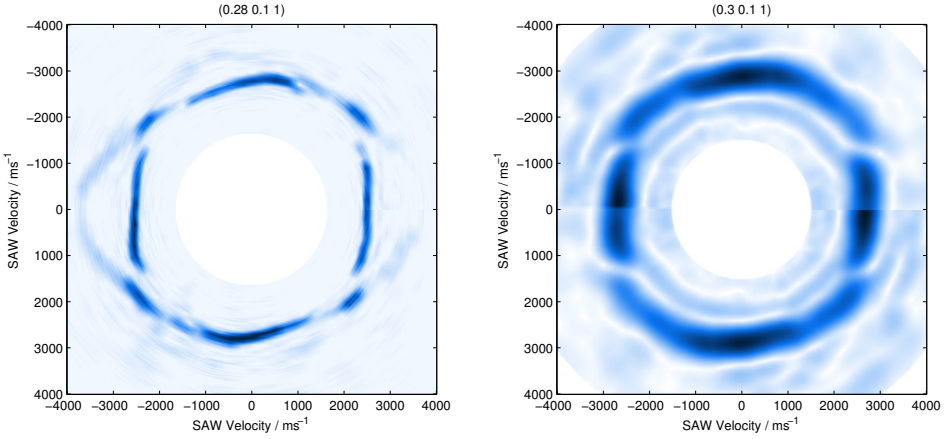
(a) (b)

Figure 7.21: a) experimentally measured and (b) modelled slowness surfaces for the (0 0 1) plane of nickel.



(a) (b)

Figure 7.22: a) experimentally measured and (b) modelled slowness surfaces for the (4 0 1) plane of nickel.



(a) (b)

Figure 7.23: a) experimentally measured and (b) modelled slowness surfaces for the (3 10 1) plane of nickel.

so far. Although the general shape, values of velocity and amplitude are correct the wide velocity peaks mean that the subtle features in the experimental data are entirely missed. Additionally the generation model fails to replicate the propagation directions in which multiple modes exist. A quick proviso is required here as the experimental and simulated data were not generated under identical conditions. The experimental data for the (3 10 1) were measured for use in the initial testing of the orientation algorithms and so, to maximise velocity resolution, a SRAS configuration with an atypically high number of grating lines was used. In contrast the generation model uses the more standard SRAS generation configuration of ~10 cycles and therefore inherently gives poorer velocity resolution. The experimental velocity surfaces in Fig. 7.21 and Fig. 7.22 are far more representative of typical data. This limit of 10 on the number or cycles which can be used by the generation model is set by the length of time the point source was simulated for.

While precise predictions of SAW phase velocity are not available in the generation model, the predictions of amplitude can be combined with the precise velocity predictions of the forward model. This produces a SRAS-specific slowness surface for a given generation patch configuration, and it is only within these limits that the generation model is useful to orientation. As such, the generation model was run on a full 21x21 set of planes for cubic nickel, matching the planes generated using the forward model. Computation time for this set was approximately 17 days, with each surface taking

1-2 hours to generate on a desktop computer. The modelled velocity surfaces shown throughout this section were taken from this large set.

7.4 Unification of generation and orientation model

Given the broadness of the generation model velocity surface peaks, and the imperfect agreement with the experimental results, using the generation model to directly fit experimental data is infeasible. Instead the generation model can be combined with the forward model to produce a hybrid amplitude-weighted forward model, i.e. the generation model is used to decrease the amplitude weighting of unobserved SAW modes. Thus when experimental data are fitted the influence of false positives can be reduced. Rather than outright discarding modes, the combined model simply re-ranks the modes, keeping only the two most dominant, and records the amplitude. When fitting experimental data, the thresholds used for rejecting modes from fitting can be set to suitable values, or not used at all. A secondary benefit of the hybrid model is that fewer modes need to be fitted to the experimental data, typically 5 modes are used from the forward model, but by eliminating modes based on amplitude this can be reduced to 2 and fitting speed increased.

Fig. 7.24 shows the forward and generation model predictions being superimposed for three planes of cubic nickel. The first two ranked modes from the forward model are shown as red and green crosses respectively, and the amplitude-ranked first mode is shown by yellow circles. Since the generation model was simulated for the same planes as the forward model, this process can be repeated for every plane. In terms of agreement between the models the results are a little varied. While the (5 9 20) and (16 17 20) planes show good agreement across all propagation angles, the (0 0 1) mode appears to have a slight offset in velocity. Given that the elastic constants used for each model are identical, the exact source of this discrepancy is not fully understood. Almost certainly the source of the error is in the generation model, and most likely stems from compounded computation and rounding errors generated during the interpolation of both the mesh grid and of the rotated point source grid.

For the plane (0 0 1) the amplitude-weighted dominant mode is in good agreement with the predictions from the forward model at nearly all angles, with the exception

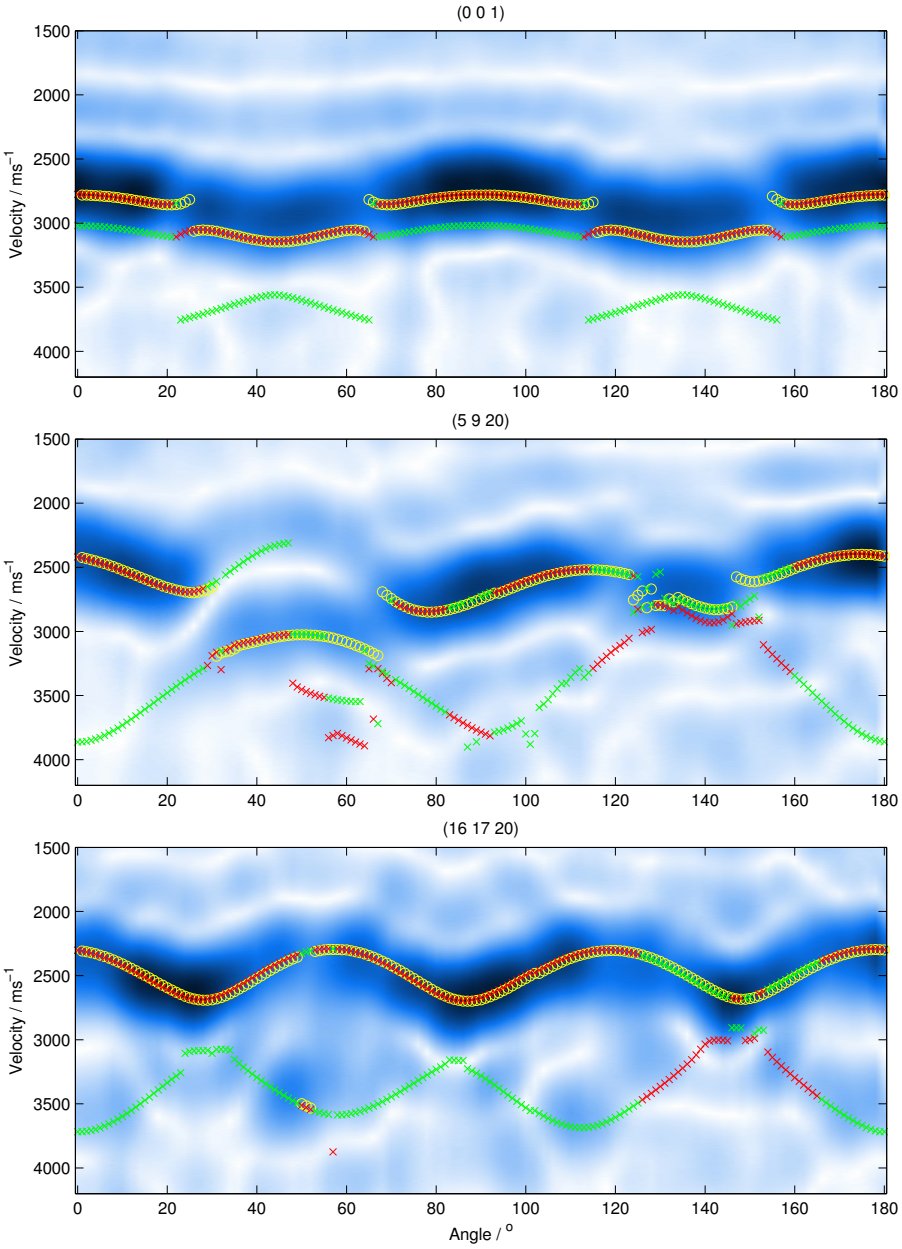


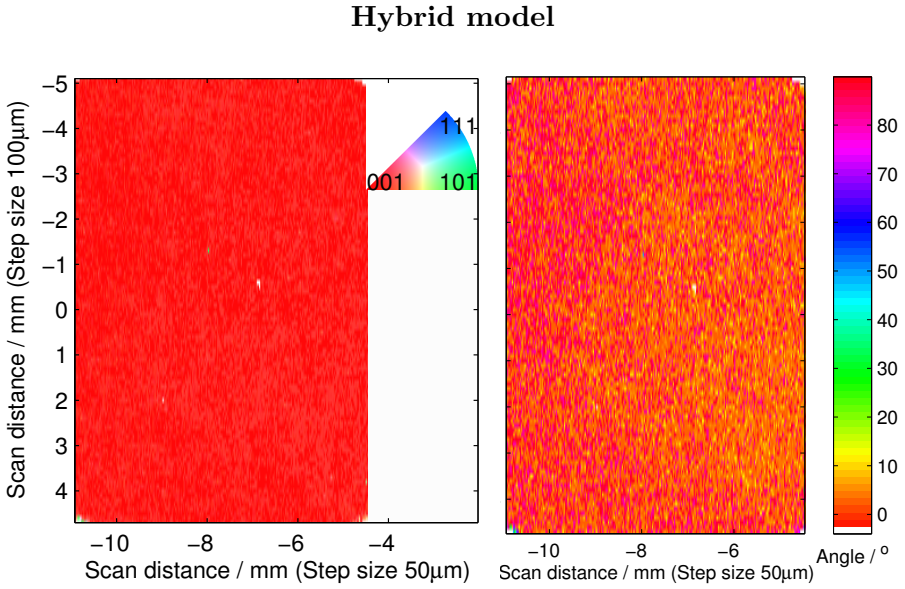
Figure 7.24: A comparison of the forward and generation model predictions for three planes of cubic nickel. The value of the generation model at the location of each of the points in the forward model is extracted and used to build the combined model. The forward model predicted modes are shown as crosses and ranked with red as the dominant mode. The hybrid weighted mode is shown as yellow circles.

occurring at transitions between wave modes, a similar result is seen for the plane (16 17 20). However for the plane (5 9 20) there is significant disagreement, with amplitude-weighted mode frequently not found within the top 2 forward model modes, this indicates that for low symmetry planes with high discontinuity between SAW modes the forward model does a poor job of ranking the SAW modes.

7.5 Applying the hybrid orientation model

Application of the hybrid orientation model is nearly identical to the application of the forward model. For each point at which data was collected the raw traces for 18 propagation angles are loaded and processed into the SAW velocity surface. The two amplitude weighted SAW modes are then compared against the velocity surface for all possible planes and a figure of merit generated at each plane and rotation angle. Although the code was written to threshold the amplitude weighted SAW modes, to reduce the number of variables in the fitting procedure this was not used, and so only the effect of amplitude ranking was tested. A slight alteration was made to the fitting model to improve speed. Instead of comparing every plane in the database, a coarse grid of planes is first applied and then refined within the range of maximum figure of merit. These changes were tested on the forward model first and showed identical orientation determination with a speed increase of ~4 times. From the final figure of merit map, the best fit plane and plane rotation were then extracted.

The experimental data used for fitting were the same data from the single crystal nickel samples used to plot Fig. 7.2 and Fig. 7.3. Fig. 7.25 shows the results of fitting using both the forward model and hybrid model. Initial results look quite promising, in terms of base functionality the hybrid model clearly works and is able to determine orientation to good agreement with the forward model. Comparing the inverse pole figures shows very similar results, with the hybrid model selecting grains closer to (0 0 1) with slightly more regularity. Fig. 7.26 shows a histogram of the angles of each fitted point to the (0 0 1) plane for both the hybrid and forward model. For plane rotation the results for the hybrid are again quite similar to those for the forward model, apart from a slightly more compact range of rotations. This is a good result as it indicates that the hybrid model is producing a more consistent result over the



(a) (b)

Figure 7.25: The orientation determination results on a single grain nickel sample of orientation close to (0 0 1) using the hybrid model. (a) The inverse pole figure and (b) the plane rotation.

known single crystal grain.

Fig. 7.27 and Fig. 7.28 show the experimental SAW velocity surfaces for the same two points on the nickel sample as shown in Fig. 7.2 and Fig. 7.3, with the fitted plane overlaid. For the forward model fitted data, only the first dominant mode is shown, while for the hybrid model the first amplitude-weighted mode is shown. Results from the first figure show the orientation determinations to be identical, the same plane has been fitted with the same rotation. This is not necessarily a terrible result as it shows that the hybrid model is performing similarly to the, already well tweaked and validated, forward model on its first application. As before, the plane appears well chosen for both models whilst the rotation has an error of ~10°.

Moving to the Fig. 7.3 shows a more promising result. Although the rotation is correct for the forward model determination, the chosen plane does not look quite right. However the hybrid model appears to have returned a better match for the plane and has returned a fairly close plane rotation. Note that the rotations stated on the figure are quite different for the two models, this is simply due to the near 4-fold symmetry producing an offset of 90°.

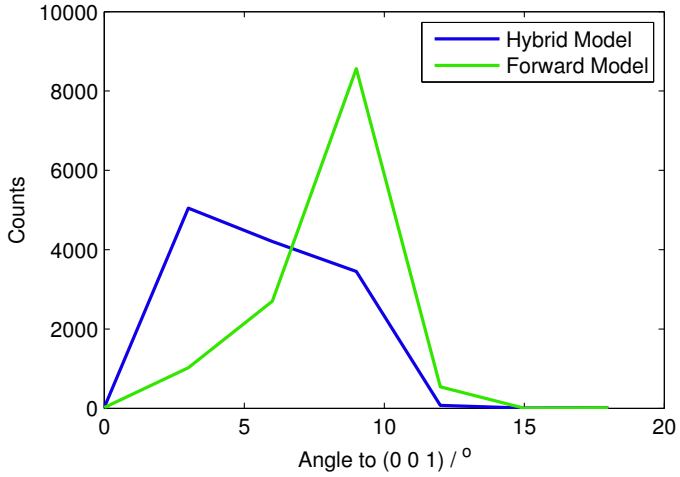
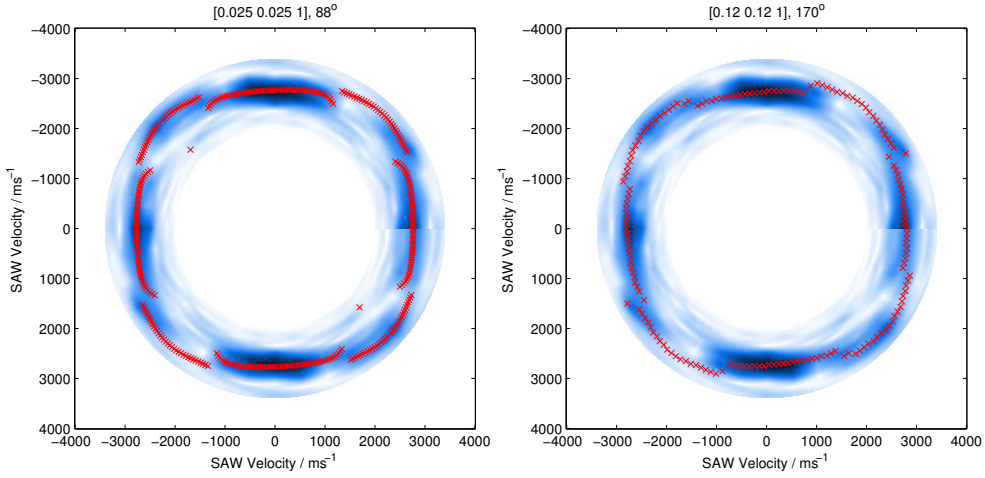
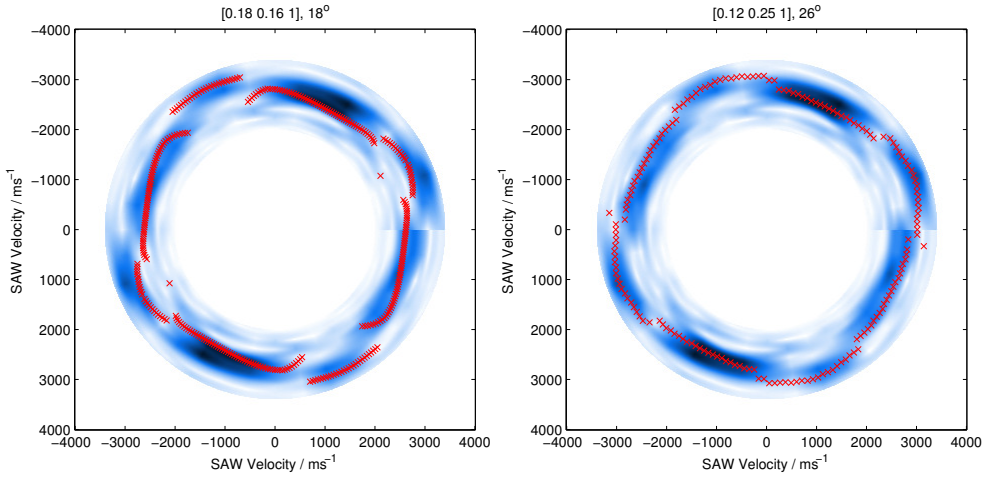


Figure 7.26: A histogram showing the distribution of the angles to the (0 0 1) plane of the fitted orientations for both the hybrid and forward models.



(a) (b)

Figure 7.27: The fitted planes shown against the raw SAW velocity surfaces for (a) the forward model and (b) the hybrid model for a single point on a sample of single crystal nickel.



(a) (b)

Figure 7.28: The fitted planes shown against the raw SAW velocity surfaces for (a) the forward model and (b) the hybrid model.

7.6 Energy directivity

Aside from the orientation determination, a secondary motivation for the generation model was to better understand why certain grains during a scan yield very low signal amplitudes in certain propagation directions. This effect has been shown clearly in some of the simulated slowness surfaces, where in the region between SAW modes the amplitudes are low, this is particularly well observed in Fig. 7.22 for both the experimental and simulated surfaces. Using the point sources this can also be shown in terms of available energy in each propagation direction. Fig. 7.29 shows the measured traces for each propagation direction on the planes (0 0 1), (1 0 1) and (1 1 1) of cubic nickel. For the (0 0 1) plane and especially the (1 0 1) plane it is clear that the transition between SAW velocity modes are fairly smooth and continuous. This is particularly evident on the mode between 12 ns to 15 ns on the (1 0 1) plane where there are no discontinuities and very consistent SAW amplitudes. Conversely, for the (1 1 1) plane, the transitions between modes are quite discontinuous leading to high variations in amplitude.

Summing the square of the amplitudes in each propagation direction gives a reasonable estimate of the energy travelling in that propagation direction, and this is shown in Fig. 7.30. For the (1 0 1) plane the variation in energy is fairly small compared to the

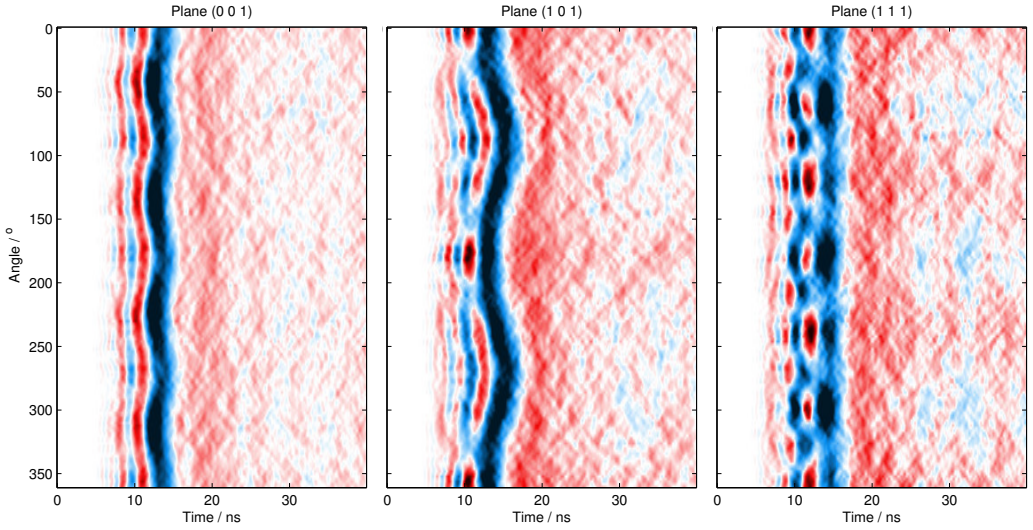


Figure 7.29: Measured traces for each propagation direction on the planes (0 0 1), (1 0 1) and (1 1 1) of cubic nickel taken from the simulated point sources.

two other planes, except for two distinct peaks. The main peaks in the energy occur at 0° and 180° which corresponds to the high large SAW modes which exists only between 170° and 190°. Either side of these peaks are slight troughs, which suggest that energy is being preferentially directed into the mode at 180°. A similar effect is seen on the (1 1 1) plane where the high amplitude modes occurring at 60°, and repeating every 60°, show peaks in the energy directivity with troughs either side. The available energy in each propagation direction is not the only factor affecting the final SRAS signal amplitude, however it does set the upper limit on SAW packet amplitude.

The conclusion here is that for certain planes, energy is preferentially directed into modes in certain propagation directions, minimising the energy available for less preferable SAW modes. Ultimately, the process which determines preferable modes is simply down to how well the SAW mode particle motion and thermally generated stress field are matched. In terms of reducing the impact of this on SRAS scans, one possible solution is to use slightly arc shaped generation lines. This would avoid the narrow directions in which SAW generation efficiency is low by averaging over a range of directions, however the SAW phase velocities would also be averaged over and the velocity contrast in the SRAS image degraded. Given that the primary purpose of SRAS is to provide good contrast between crystal grains, this issue must be solved by improving

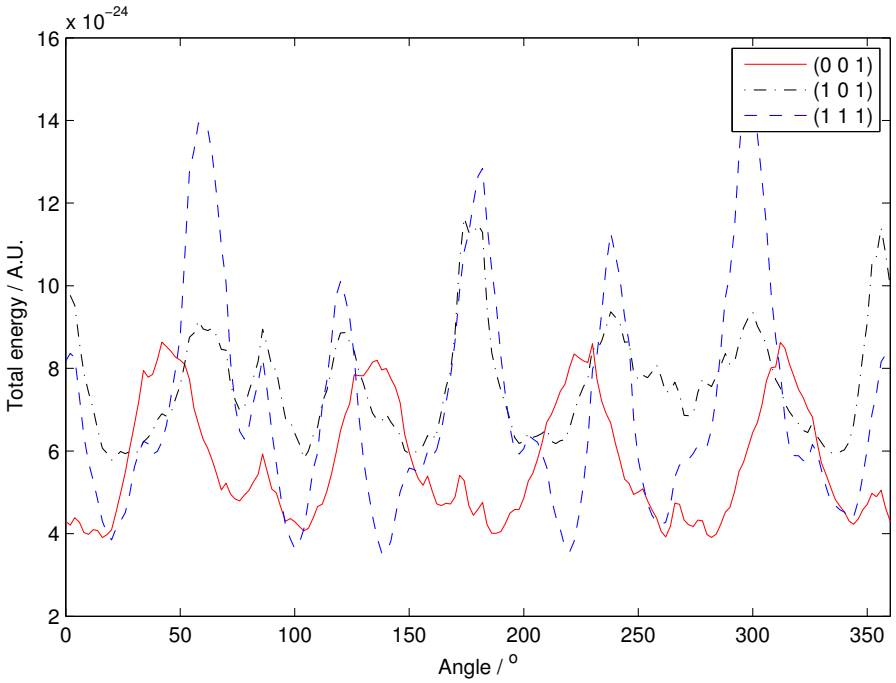


Figure 7.30: An estimate of the energy available in each propagation direction for the (0 0 1), (1 0 1) and (1 1 1) of cubic nickel taken from the simulated point sources.

signal amplitude in other ways.

7.7 Conclusions and further work

The generation model comprises two steps: a FEM to produce a point source, and manipulation of the point source to produce SAW velocity surfaces. For the first stage of this process the simulated point sources on anisotropic materials appear to show consistency with the expected symmetry and, when used to produce velocity surfaces, show agreement with experiment. However these point sources have not been independently validated against experimental data and so their accuracy is simply unknown. Thus a clear piece of future work is to apply SRAS as a focussed point source to a known orientation sample and use the result to validate the model.

The point sources are slow to simulate, typically taking 30 minutes per plane on a fairly standard desktop PC. Thus, building a full database of planes is extremely time consuming and, although this process only needs to be performed once per material, the issue of whether this is truly feasible given the range of materials and alloys SRAS

can scan, remains. A secondary motivation for building the generation model was to explore why certain grains have such large variations in generation efficiency. For this purpose the point sources have proved quite useful, showing that for propagation angles which sit in the transition between two SAW modes, less energy is available.

The second stage of the model takes the simulated point sources and superimposes them to generate line and grating sources and there are a few issues with this process. Tiling the same point source repeatedly results in compounding the computational errors from the simulation, ideally multiple point sources would used for tiling each with a slightly different initial mesh. However, the required computation time for this is infeasible and initial trials of applying dithering were unsuccessful. In addition interpolation has to be applied for both generating the regular grid from the irregular mesh, and for rotation of the point source before tiling, with this process introducing error to the predicted surface displacements. Although possible to generate finer meshes and finer grids to help avoid the interpolation requirement, computation time makes this difficult.

Generating a line source from a infinite point source was initially shown to be a valid operation by Rose [114] and was successfully applied to isotropic materials. However more recent work based on anisotropic materials have shown that, for bulk acoustic waves, this assumption is invalid due to focussing of energy [150]. Saying that, a comparison was made between a FEM simulation of a line source and the tiled point source and shown to have good agreement, and the SAW velocity surfaces show fairly good agreement with experiment, Therefore, while the assumption may not be perfectly valid, it appears valid enough for the purposes of the model.

The final stage of the model uses simulated velocity surfaces to weight the forward model SAW modes based on generation amplitude, to produce the hybrid model. This was applied to experimental data to determine orientation and was compared to the forward model. While the hybrid was no worse than the forward model, it also did not show any significant improvement, other than decreasing processing time due to the reduction in the number of SAW modes compared to the data. However, because there are so many parameters to tweak in the generation model which affect the SAW velocity surfaces, it is difficult to determine whether the generation will successfully improve the orientation determination procedure without significant further work.

The generation model was explored as a means by which some observed issues with the orientation determination could be ameliorated. To be clear this chapter is presented as a somewhat open-ended piece of work with only a description of the process and the initial results shown. As such, the question of whether this initial aim has been met is unanswered, but from the initial findings the usefulness of the model in general appears to be good. It is also important to note that, throughout the body of work undertaken in modelling the laser ultrasonic generation process, modelling of point sources and velocity surfaces on arbitrarily anisotropic materials has never been done before. Furthermore, although this thesis focuses largely on the engineering aspects of SRAS and making quantifiable improvements to capability, the opportunity to understand the SRAS process better, even if only for the sake of it initially, should not be missed.

Chapter 8

Final conclusions

As stated in the introduction, the primary aim of this EngD thesis project was to take a SRAS laboratory demonstrator and develop it into a functioning industrial prototype. This aim has been achieved with the Renishaw SRAS system fully-operational and collecting data. To reach this point progress has been made in three areas: instrumentation, signal processing and determination of the Renishaw SRAS system performance. In addition, further progress has been made in two areas which have not yet been applied to the industrial prototype: machine vision for auto-focus and the modelling the generation of laser ultrasound.

8.1 Instrumentation

At the start of this project the state-of-the-art was the EMDA system at the University of Nottingham. This is a very capable technology demonstrator but is not well suited to: ease of operation, large sized samples, assembly and manufacture, and low cost construction. Through the design and implementation of a novel optomechanical chassis, nearly all of the EMDA system limitations have been ameliorated for the Renishaw SRAS system. Based on a rigid chassis and modular optical housings, the full instrumentation details are given in Chapter 3. This industrial prototype is lightweight and rigid, allowing it to be fibre-coupled and mounted to a scanning motion platform. Manufacture and assembly of the optical housings and head is far simpler than for the EMDA system, and is well suited to industrial manufacture.

The new system is not without its own problems. Fibre-coupling of the multi-mode fibre has reduced the energy density available for generation resulting in reduced SNR. Although improvements have been made to SNR with more effective electrical shielding and the elimination of coherent noise, the final noise floor of the system is higher than the EMDA system. Some additional research into more sophisticated optical configurations have been suggested to improve energy density, while in the interim more sophisticated signal processing is applied. Although the optomechanics are well developed, significant work is still required in other areas of integration: primarily data acquisition and mechanics for head rotation.

8.2 Signal processing

The Renishaw SRAS system has reduced SNR compared to the EMDA system although, unlike the EMDA system, this noise is completely incoherent. As a result advanced signal processing techniques can be readily applied to the measured signals, and this was the subject of Chapter 4. A range of processes have been compared for both simulated and experimental SRAS data, and it was found that Morlet wavelet analysis is very effective at determining dominant phase velocity. This wavelet analysis assumes a signal form and generates a set of wavelet daughters over a given frequency range. These daughters are then compared to the SRAS signal and the peak frequency is taken where highest correlation occurs. Morlet wavelets analysis requires two main parameters: the order of the wavelet, and the scale factor. The order dictates the temporal length of the wavelet and the optimum was found to be 8. The scale factor represents how closely spaced in frequency the wavelet daughters are, and is a compromise between computation time and frequency resolution. For typical analysis of SRAS signals the optimum was taken to be 0.006. Comparing Morlet wavelet analysis to the previously used MFP technique on experimental data shows an approximately factor two improvement in determined velocity standard deviation on single crystal grains. With minor optimisation of code the Morlet technique has been applied in real-time to SRAS scans although, with increasing data capture rates the fitting techniques will require further optimisation. Wavelet analysis gives time-frequency information and full analysis of dual-frequency signals remains a promising area for further work.

8.3 Renishaw SRAS system performance

In Chapter 5 a thorough assessment of the capabilities and performance of the Renishaw SRAS system was undertaken. The Equator motion platform was found to have periodic and repeatable deviations from the expected scan path, while this problem could not be addressed mechanically a solution based on auto-focus was sought in Chapter 6. Determining generation laser output energy showed that at higher repetition rates the energy per pulse decreases and the pulse length increases. As a result the generated SAW signal strength decreases rapidly between 1 kHz and 10 kHz, which limits the useful range of repetition rates. Comparing the generation efficiency of different gratings showed that, for shorter acoustic wavelengths, generation efficiency decreases due to the optical resolution of the optics. This affects SNR and is typically evidenced by noise in the velocity contrast maps. If very short acoustic wavelengths are required a higher magnification objective lens will improve efficiency.

From a stainless steel sample with long thin grains, the spatial resolution of SRAS in orthogonal directions was estimated as 30 µm in the propagation direction, and 10 µm in the orthogonal direction, for a 23 µm acoustic wavelength and 1 mm aperture configuration. For more typical samples with small, round grains the spatial resolution is limited to 30 µm. The primary factor affecting spatial resolution was found to be the aperture size, with which spatial resolution scales linearly. Although acoustic wavelength does influence spatial resolution, it is in a more subtle manner than the aperture.

The alternate twin KED was compared to the segmented KED and found to perform similarly where detection power is ample, but was more sensitive to loss of optical power. A path for improvements in the detector was detailed along with a plan for full fibre-coupling. Finally, the system was purposefully detuned from the optimum optical path in order to test the tolerance to optical component manufacturing variances. It was found that for the generation beam the optical path lengths are critical and a means for measuring each lens prior to assembly was suggested.

8.4 Machine vision for auto-focus

The Equator 700 motion platform was shown to give imperfect motion with small deviations in all Cartesian axes, especially during acceleration phases. To counter this Chapter 6 examined the feasibility of using autonomous focus techniques based on images of the generation patch. A set of focus measure techniques were compared with the gray-level local variance and Gaussian derivative methods identified as the most suitable. These focal measures are useful for repeatable configuration of scans, which previously relied on a 'by-eye' estimate of focus.

Analysis of the generation patch was extended and a Fourier transform technique devised and applied to measure changes in grating spacing. Then, by calibrating the change in grating spacing with the focal height offset, it was shown that the focal height offset could be determined from a single camera image. Although this was tested, and shown to work well, for real experimental data this process has only been performed off-line. The next logical step is to apply this process to live scans for autonomous, and real-time, focal height correction.

8.5 Generation model

Two effects have been seen which motivated the research into an anisotropic model of laser ultrasound. The first was systematic errors in the observed orientation determination for nickel and the second was the variation in SRAS signal strength with local crystallography. Chapter 7 details the progress made in developing a model for laser ultrasonic generation of SAWs. The existing research into laser ultrasonic generation is reviewed and from this the conclusion is drawn that FEM is the most suitable technique for further research. A combined generation model was devised and implemented which uses FEM to simulate a point-source on an arbitrary crystallographic plane. These point sources showed the expected symmetry and variations in velocity and amplitude, but were not validated against experiment. A superposition method was then applied to tile the point sources into a grating representative of a typical SRAS configuration. From these gratings velocity surfaces were simulated which were shown to have reasonable agreement with experiment.

Taking the existing forward model, the two models were unified to create a hybrid model. This used the high precision velocity prediction from the forward model, and the amplitude weighting of the generation model. Performing orientation determination with the hybrid model showed good agreement with the forward model fitting, with some suggestion of improved plane determination. Due to the large number of parameters available for tweaking the generation and hybrid model, there remains a large body of work to determine whether the hybrid model can remove the issues seen with the forward model.

The generation model was also used to analyse the directivity of energy of the point sources. This showed that for planes with a large number of steps between SAW modes there was poor SAW amplitude in the transition between modes. Consequently there is no improvement in SNR to be made in these transition propagation directions.

8.6 Summary

In combination, the completed areas of research contribute quantitatively to an improvement in the robustness and capability of the SRAS technique. For the research which has not yet been fully implemented, there are sufficient indications of promise to suggest that they can be applied usefully to the technique. While this progress is encouraging in its own right, the more important aspect for the future of the technology is how this affects long-term industrial deployment. As an indicator of progress from an industrial perspective, during this EngD the SRAS technology has cleared three major project milestones of the Renishaw new product development process. This indicates that the technique is past the early technology feasibility stages and is currently undergoing assessment for product feasibility. Given that relatively few projects reach this stage without becoming a released product, there is good reason to be optimistic that the SRAS technology will become available to industry as a mature materials characterisation instrument.

Appendix A

Optical chassis design

In this appendix the process of design and refinement of the SRAS head optomechanical chassis is detailed. While this is by no means an essential read as part of the thesis, it is hoped that this review gives some insight into the rationale for the decision to use the design presented in Chapter 3.

A.1 Design

A variety of chassis designs were trialled and studied for feasibility throughout the SRAS project. The first design was based upon optics held by V-shaped clamps in a milled aluminium chassis with sufficient space to allow for tweaking of the optical path and is shown in Fig. A.1. This design is the most similar to the way in which typical Renishaw parts are made. Unfortunately this design contained numerous flaws some of which were fairly fundamental. For example the aluminium chassis, while advantageous in its simplicity of manufacture and thermal stability, does not offer much adjustment, making prototyping very difficult. In addition the V-clamp system does not work well for constraining high aspect ratio optics.

Taking inspiration from the constraint of lenses within DSLR camera lens assemblies, the next prototype design employed a system of concentric tubes, used to hold circular optics both vertically and around the perimeter, this is shown in Fig. A.2. From this mock-up a few advantages can be seen. With the ability to use a more three-dimensional design, rather than the two-dimensional limitation of the milled chassis,

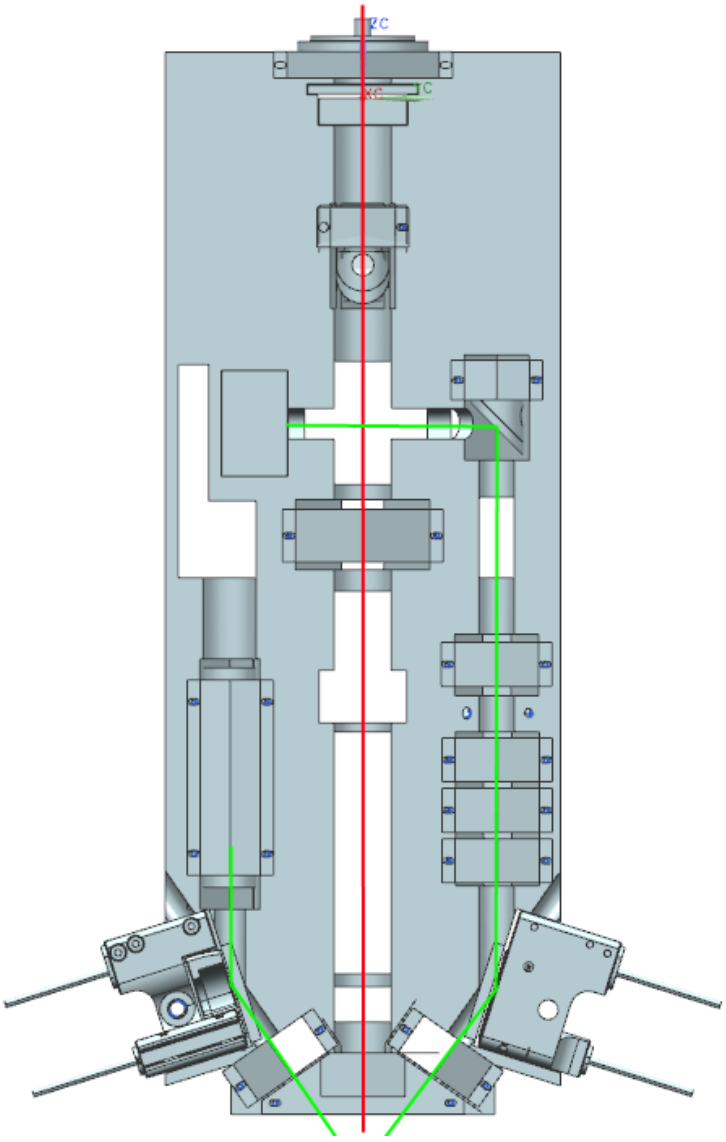


Figure A.1: A working design of the single-block chassis for containing and aligning the optics. The approximate optical paths are shown in red (for the IR) and green (for the visible).

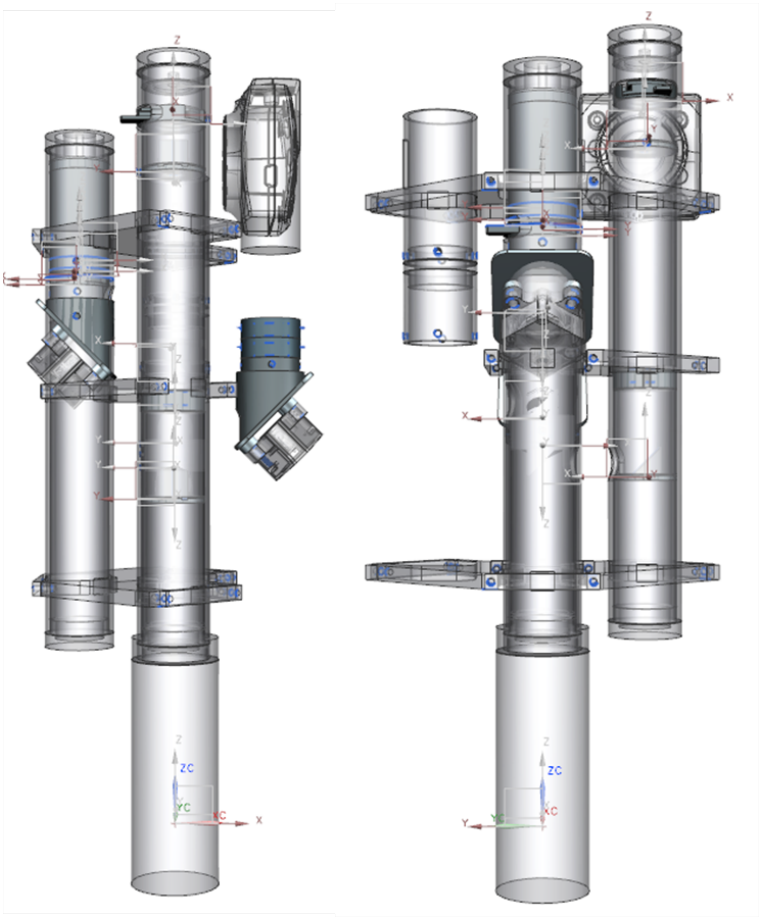


Figure A.2: A mock up of the tube constraint system applied to a full instrument design. Four tubes are used for different optical paths: IR beam delivery, visible beam delivery, visible beam return and camera image return.

optical paths can be completely separated into their own tubes, increasing the optical path simplicity and the ease of alignment. In addition optics are well constrained and prototyping can be achieved by simply cutting new tube lengths. However, assembly of this instrument is not very simple; there are multiple different lengths of tube to hold together which must be kept highly parallel. Also changing components, such as apertures, masks and filters, is difficult without disassembly. Although circular optics are very well constrained, for optics which are rotationally variant, such as wave plates, constrain is difficult.

From this tube design emerged a concept which aimed to combine the benefits of the milled chassis, i.e. rigidity and ease of manufacture, and the tube system, i.e optical constraint and easy prototyping. This to the rigid extruded chassis which,

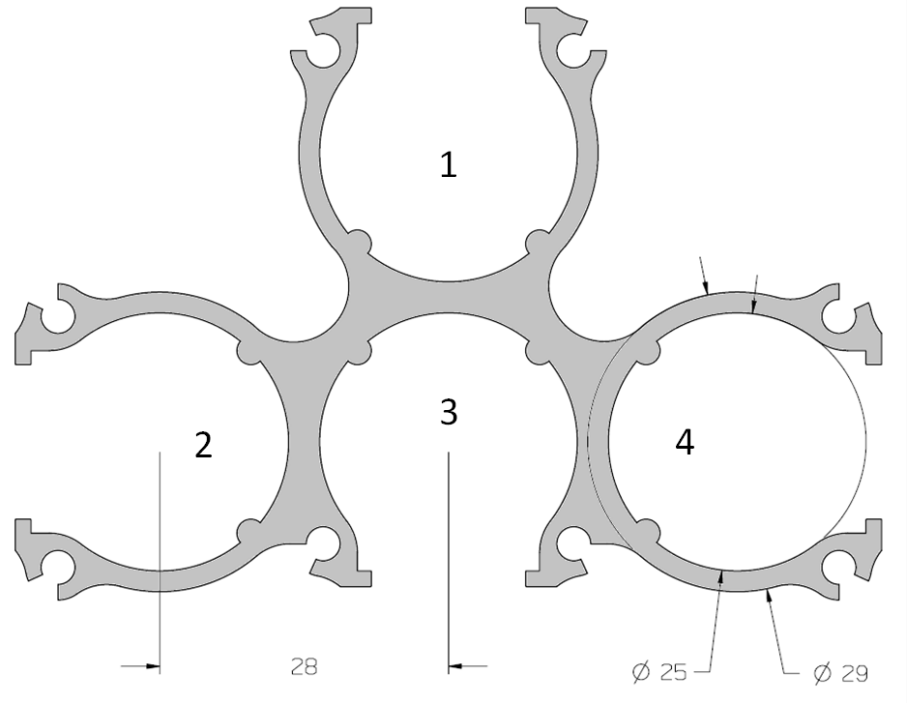


Figure A.3: A profile view of an extruded chassis design which uses flat plates to constrain optics.

in combination with cut tube, constrains optical housings accurately and with good repeatability. Fig. A.3 shows one of the early designs applying this concept which uses a flat plate at the mouth of each channel to constrain the optics. A cap plate is placed at the top and bottom of the extrusion and is used to apply a load to and constrain the cut tube and optical housings vertically. Unfortunately this prototype did not work as well as expected; the plates were simply to difficult to slide into place. In addition having a solid plate meant that access to the optical housings was not possible and so apertures and gratings could not be easily changed.

Thus, the final design was arrived at by a fairly simple change of the plate constraint, this is shown in Fig. A.4. Instead of a plate inserted after assembly, a permanent pair of rubber strips are used to provide a constant pre-load against the horizontal constraint. For assembly this works very well as the pre-load is gentle enough to allow easy sliding of optics, for long-term stability the optical housings and tube are clamped vertically.

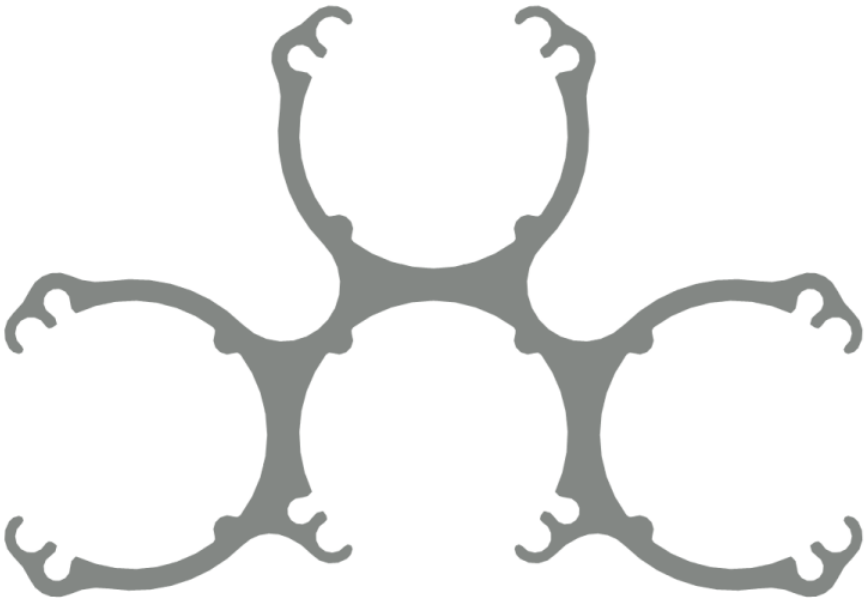


Figure A.4: A mock up of the tube constraint system applied to a full instrument design. Four tubes are used for different optical paths: IR beam delivery, visible beam delivery, visible beam return and camera image return.

A.2 Manufacture

Both the designs shown in Fig. A.3 and Fig. A.4 were manufacture and tweaked at various stages in the design process. Initial prototyping was by selective laser sintering, however this method resulted in a prototype which was too flexible to be able to assess how well the rubber constraints were working. To create a more rigid prototype, selective laser melting (SLM), an additive manufacture process, was used and a titanium version of the chassis was printed, both methods are shown in Fig. A.5. SLM was also being explored as potential route for manufacture. Unfortunately the titanium chassis did not work entirely as expected as the print process results in a very rough surface. A smooth surface is required for sliding optical mounts into place and for sliding in the constraining plates, neither of these things were possible with the titanium chassis. The final prototype, used throughout the thesis, was manufactured by wire erosion.



(a) (b)

Figure A.5: a) Prototype of the SRAS extrusion built using 'Selective layer sintering', a process in which a profile is melted into a thin layer of powdered nylon which is re-laid for each new layer. This prototype has a 3 mm wall thickness which was found to be too thin. b) Selective laser melted (SLM) titanium chassis.

Appendix B

Wavelets

B.1 Form of B-spline wavelets

The closed form expression for the B-spline wavelets used in Chapter 4 can be found below [151]. Families of these wavelets are constructed as described in [152].

$$\Psi_n(x) = x^n\psi_0 + x^{n-1}\psi_1 ... + x^0\psi_n, \tag{B.1}$$

where Ψ is the final B-spline wavelet form, n is the wavelet order and ψ are the B-spline wavelet coefficients used to construct Ψ. The third order B-spline wavelet has coefficients:

$$a_3\psi_{C,3}(x) = \begin{cases} 1, 0, 0 & \text{for } 0 \leqslant x < \frac{1}{2} \\ -31, 32, -8 & \text{for } \frac{1}{2} \leqslant x < 1 \\ 206, -442, 229 & \text{for } 1 \leqslant x < \frac{3}{2} \\ -626, 2054, -1643 & \text{for } \frac{3}{2} \leqslant x < 2 \\ 1056, -4674, 5085 & \text{for } 2 \leqslant x < \frac{5}{2} \\ -1056, 5886, -8115 & \text{for } \frac{5}{2} \leqslant x < 3 \\ 626, -4206, 7023 & \text{for } 3 \leqslant x < \frac{7}{2} \\ -206, 1618, -3169 & \text{for } \frac{7}{2} \leqslant x < 4 \\ 31, -278, 623 & \text{for } 4 \leqslant x < \frac{9}{2} \\ -1, 10, -25 & \text{for } \frac{9}{2} \leqslant x < 5 \end{cases} \tag{B.2}$$

$$a_3 = 240.$$

The fourth order B-spline wavelet has coefficients:

$$a_4\psi_{C,4}(x) = \begin{cases} 1,0,0,0 & \text{for } 0 \leqslant x < \frac{1}{2} \\ -127,192,-96,16 & \text{for } \frac{1}{2} \leqslant x < 1 \\ 2052,-6345,6441,-2163 & \text{for } 1 \leqslant x < \frac{3}{2} \\ -13308,62775,-97239,49677 & \text{for } \frac{3}{2} \leqslant x < 2 \\ 47349,-301167,630645,-435579 & \text{for } 2 \leqslant x < \frac{5}{2} \\ -105099,842193,-2227755,1946421 & \text{for } \frac{5}{2} \leqslant x < 3 \\ 154624,-1495314,4784766,-5066100 & \text{for } 3 \leqslant x < \frac{7}{2} \\ -154624,1751790,-6580098,8192908 & \text{for } \frac{7}{2} \leqslant x < 4 \\ 105099,-1364886,5886606,-8429364 & \text{for } 4 \leqslant x < \frac{9}{2} \\ -47349,693162,-3374610,5462460 & \text{for } \frac{9}{2} \leqslant x < 5 \\ 13308,-216693,1174665,-2119665 & \text{for } 5 \leqslant x < \frac{11}{2} \\ -2052,36747,-219255435855 & \text{for } \frac{11}{2} \leqslant x < 6 \\ 127,-2475,16077,-34809 & \text{for } 6 \leqslant x < \frac{13}{2} \\ -1,21,-147,343 & \text{for } \frac{13}{2} \leqslant x < 7. \end{cases} \tag{B.3}$$

$$a_4 = 30240.$$

Appendix C

Elastic constants and tensor rotation

C.1 Elastic constant tensor

For the case of an anisotropic elastic solid Hooke's law can be expressed as [153]

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl}. \tag{C.1}$$

Here σ_{ij} and ϵ_{kl} are 3X3 matrices representing the components of stress and strain respectively, and C_{ijkl} is an 81 element fourth-rank tensor representing the elastic constants. Both σ_{ij} and ϵ_{kl} are symmetric tensors and so

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl} = \sigma ji = C_{jikl}\epsilon_{kl}, \tag{C.2}$$

and

$$\sigma_{ij} = C_{ijkl}\epsilon_{kl} = C_{ijlk}\epsilon_{lk}. \tag{C.3}$$

Therefore

$$C_{ijkl} = C_{ijlk} = C_{jikl}. \tag{C.4}$$

Tensor Index	Matrix Index
11	1
22	2
33	3
23	4
13	5
12	6

Table C.1: The index transformations used to convert the elastic constant fourth-rank tensor into a 6X6 matrix.

Applying these symmetries means that the elastic constant tensor can be reduced from 81 independent elements to 36. With this simplification it become convenient to represent the elastic constant tensor as a 6X6 matrix and the stress and strain tensors as 6X1 matrices. To do this the index transforms given in Table C.1 are used.

Applying this procedure reduces Hooke's law to $\sigma_i = C_{ij}\epsilon_j$, or in full

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{21} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{31} & C_{32} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{41} & C_{42} & C_{43} & C_{44} & C_{45} & C_{46} \\ C_{51} & C_{52} & C_{53} & C_{54} & C_{55} & C_{56} \\ C_{61} & C_{62} & C_{63} & C_{64} & C_{65} & C_{66} \end{bmatrix} \times \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \epsilon_{yz} \\ \epsilon_{zx} \\ \epsilon_{xy} \end{bmatrix}. \tag{C.5}$$

Note that in this notation $C_{ij} = C_{ji}$.

C.2 Example material elastic constants

For the generation model both cubic nickel and hexagonal titanium have been simulated. The elastic and thermal properties are presented herein.

For isotopic material the elastic constant matrix is.

$$C_{ij} = \begin{pmatrix} c_{11} & c_{12} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{12} & c_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{c_{11}-c_{12}}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{c_{11}-c_{12}}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{c_{11}-c_{12}}{2} \end{pmatrix}. \tag{C.6}$$

For cubic materials the elastic constant matrix is.

$$C_{ij} = \begin{pmatrix} c_{11} & c_{12} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{12} & 0 & 0 & 0 \\ c_{12} & c_{12} & c_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & c_{44} \end{pmatrix}. \tag{C.7}$$

For hexagonal materials the elastic constant matrix is.

$$C_{ij} = \begin{pmatrix} c_{11} & c_{12} & c_{13} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{13} & 0 & 0 & 0 \\ c_{13} & c_{13} & c_{31} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & 2(c_{11}-c_{12}) \end{pmatrix} \tag{C.8}$$

Material	Moduli				Density / $\mathrm{kg\,m^{-3}}$	Ref
	c_{11}	c_{12}	c_{13}	c_{44}		
Nickel	247	153	-	122	8912	[14]
Titanium α	163.6	92.3	67.92	47.05	4429	[14]

Table C.2: Table giving the mechanical properties of various media.

Material	Heat capacity / $\mathrm{J\,K^{-1}}$	Thermal conductivity / $\mathrm{W\,m^{-1}\,K}$	Thermal expansion $\mathrm{\mu m\,K^{-1}}$	Ref.
Nickel	445	90.7	13.4	[154]
Titanium	522	21.9	13.4	[154]

Table C.3: Table giving the thermal properties of various media.

C.3 Rotation matrix

For rotation the 6x6 elastic constant matrix, C_{ij}, is converted into the 4D stiffness tensor, C_{ijkl} and rotated as follows [153].

$$C_{ijkl} = \gamma_{im}\gamma_{jn}\gamma_{ko}\gamma_{lp}C_{mnop}, \tag{C.9}$$

where

$$\gamma_{ij} = \begin{pmatrix} \mathbf{m_1}\cdot\mathbf{e_1} & \mathbf{m_1}\cdot\mathbf{e_2} & \mathbf{m_1}\cdot\mathbf{e_3} \\ \mathbf{m_2}\cdot\mathbf{e_1} & \mathbf{m_2}\cdot\mathbf{e_2} & \mathbf{m_2}\cdot\mathbf{e_3} \\ \mathbf{m_3}\cdot\mathbf{e_1} & \mathbf{m_3}\cdot\mathbf{e_2} & \mathbf{m_3}\cdot\mathbf{e_3} \end{pmatrix}. \tag{C.10}$$

Here e_i are the initial orthonormal basis vectors and m_i are the orthonormal basis vectors to be rotated to.

Bibliography

[1] L. Garwin and T. Lincoln, *A Century of Nature: Twenty-One Discoveries That Changed Science and the World*. University of Chicago Press, 2010.

[2] G. Mourou, C. Barty, and M. Perry, "Ultrahigh-intensity lasers: Physics of the extreme on a tabletop," *Physics Today*, vol. 51, no. 1, pp. 22–8, 1998.

[3] C. Blain, *Lasers: a guide to the book literature*. Nova Science Publishers, 2002.

[4] J.-P. Monchalin, C. Néron, J. F. Bussière, P. Bouchard, C. Padioleau, R. Héon, M. Choquet, J.-D. Aussel, G. Durou, and J. A. Nilson, "Laser ultrasonics: from the laboratory to the shop floor," *Advanced Performance Materials*, vol. 5, pp. 7–23, 1998.

[5] C. Scruby and L. Drain, *Laser Ultrasonics: Techniques and Applications*. London: Taylor and Francis, 1990.

[6] J. Cheeke and N. David, *Fundamentals and Applications of Ultrasonic Waves*. CRC Press, 2002.

[7] L. Landau, E. Lifshits, A. Kosevich, and L. Pitaevskiĭ, *Theory of Elasticity*. Course of theoretical physics, Butterworth-Heinemann, 1986.

[8] K. Graff, *Wave Motion in Elastic Solids*. Oxford University Press, 1975.

[9] L. Holthuijsen, *Waves in Oceanic and Coastal Waters*. Cambridge University Press, 2007.

[10] L. Rayleigh, "On waves propagated along the plane surface of an elastic solid," *Proceedings of the London Mathematical Society*, vol. 1, no. 1, p. 4, 1885.

[11] M. Rahman and J. Barber, "Exact expressions for the roots of the secular equation for rayleigh waves," *Journal of Applied Mechanics*, vol. 62, no. 1, pp. 250–252, 1995.

[12] D. Nkemzi, "A new formula for the velocity of rayleigh waves," *Wave Motion*, vol. 26, no. 2, pp. 199–205, 1997.

[13] A. Oliner and G. Farnell, "Types and properties of surface waves," in *Acoustic Surface Waves*, vol. 24 of *Topics in Applied Physics*, pp. 13–60, Springer Berlin / Heidelberg, 1978.

[14] W. Li, *Laser ultrasonic method for determination of crystallographic orientation of large grain metals by spatially resolved acoustic spectroscopy.* PhD thesis, University of Nottingham, 2012.

[15] A. Love, *Some Problems of Geodynamics.* CUP Archive, 1967.

[16] H. Lamb, "On waves in an elastic plate," *Proceedings of the Royal Society of London. Series A*, vol. 93, no. 648, p. 114, 1917.

[17] W. Gao, C. Glorieux, and J. Thoen, "Laser ultrasonic study of lamb waves: determination of the thickness and velocities of a thin plate," *International journal of engineering science*, vol. 41, no. 2, pp. 219–228, 2003.

[18] D. Hutchins, R. Dewhurst, and S. Palmer, "Directivity patterns of laser-generated ultrasound in aluminum," *The Journal of the Acoustical Society of America*, vol. 70, p. 1362, 1981.

[19] A. Every, Z. Utegulov, and I. Veres, "Laser thermoelastic generation in metals above the melt threshold," *Journal of Applied Physics*, vol. 114, no. 20, p. 203508, 2013.

[20] S. Davies, C. Edwards, G. Taylor, and S. Palmer, "Laser-generated ultrasound: its properties, mechanisms and multifarious applications," *Journal of Physics D: Applied Physics*, vol. 26, p. 329, 1993.

[21] S. Pierce, B. Culshaw, and Q. Shan, "Laser generation of ultrasound using a modulated continuous wave laser diode," *Applied Physics Letters*, vol. 72, p. 1030, 1998.

[22] T. Murray and O. Balogun, "High-sensitivity laser-based acoustic microscopy using a modulated excitation source," *Applied Physics Letters*, vol. 85, p. 2974, 2004.

[23] J. Monchalin, "Optical detection of ultrasound," *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, vol. 33, pp. 485–499, 1986.

[24] R. Whitman and A. Korpel, "Probing of acoustic surface perturbations by coherent light," *Applied Optics*, vol. 8, no. 8, pp. 1567–1576, 1969.

[25] A. Harata, H. Nishimura, and T. Sawada, "Laser-induced surface acoustic waves and photothermal surface gratings generated by crossing two pulsed laser beams," *Applied Physics Letters*, vol. 57, no. 2, pp. 132–134, 1990.

[26] P. V. Santos *et al.*, "Modulation of photonic structures by surface acoustic waves," *Reports on Progress in Physics*, vol. 68, no. 7, p. 1639, 2005.

[27] P. V. Santos, "Acoustic field mapping on gaas using microscopic reflectance and reflectance anisotropy," *Applied Physics Letters*, vol. 74, no. 26, pp. 4002–4004, 1999.

[28] J. Monchalin, "Optical detection of ultrasound at a distance using a confocal fabry–perot interferometer," *Applied Physics Letters*, vol. 47, no. 1, pp. 14–16, 1985.

[29] J. Monchalin and R. Heon, "Laser ultrasonic generation and optical detection with a confocal fabry-perot interferometer," *Materials Evaluation*, vol. 44, pp. 1231–1237, 1986.

[30] R. Ing and J.-P. Monchalin, "Broadband optical detection of ultrasound by two-wave mixing in a photorefractive crystal," *Applied Physics Letters*, vol. 59, no. 25, pp. 3233–3235, 1991.

[31] H. Rosenberg, *The Solid State*. Oxford University Press, 2000.

[32] L. Gray, *Tin*. Benchmark Books, 2004.

[33] I. Tsvankin, K. Helbig, and S. Treitel, "Seismic signatures and analysis of reflection data in anisotropic media," 2001.

[34] H. McSkimin, "Measurement of elastic constants at low temperatures by means of ultrasonic waves–data for silicon and germanium single crystals, and for fused silica," *Journal of Applied Physics*, vol. 24, no. 8, pp. 988–997, 1953.

[35] J. Maynard, "Resonant ultrasound spectroscopy," *Physics Today*, vol. 49, no. 1, pp. 26–31, 2008.

[36] M. Deschamps and C. Bescond, "Numerical method to recover the elastic constants from ultrasound group velocities," *Ultrasonics*, vol. 33, no. 3, pp. 205–211, 1995.

[37] S. I. Tomkeieff, "The basalt lavas of the giants causeway district of northern ireland," *Bulletin Volcanologique*, vol. 6, no. 1, pp. 89–143, 1940.

[38] M. Durand-Charre, *The Microstructure of Superalloys.* Taylor & Francis, 1998.

[39] EDAX Inc, *Steel Property Optimization: Using EBSD To Control Retained Austenite Content*, 8 2012.

[40] G. Petzow, *Metallographic Etching, 2nd Edition: Techniques for Metallography, Ceramography, Plastography.* Techniques for metallography, ceramography, plastography, ASM International, 1999.

[41] J. Alkemper and P. Voorhees, "Quantitative serial sectioning analysis," *Journal of Microscopy*, vol. 201, no. 3, pp. 388–394, 2001.

[42] G. McMahon and T. Malis, "Ultramicrotomy of nanocrystalline materials," *Microscopy Research and Technique*, vol. 31, pp. 267–274, 1995.

[43] B. Tanem, J. Mol, H. Terryn, J. dewit, L. Katgerman, *et al.*, "A combined tem and skpfm investigation of the surface layers on rolled aa5050 aluminium alloy using ultra-microtomy," *Surface and Interface Analysis*, vol. 40, no. 8, pp. 1157–1163, 2008.

[44] B. Warren, *X-Ray Diffraction.* Addison-Wesley, 1990.

[45] D. Bilderback, S. Hoffman, and D. Thiel, "Nanometer spatial resolution achieved in hard x-ray imaging and laue diffraction experiments," *Science*, vol. 263, no. 5144, p. 201, 1994.

[46] A. Schwartz, M. Kumar, and B. Adams, *Electron Backscatter Diffraction in Materials Science*. Kluwer Academic, 2000.

[47] V. Randle, *Microtexture Determination and its Applications*. Maney Publishing, 2003.

[48] F. Humphreys, "Review grain and sub-grain characterisation by electron backscatter diffraction," *Journal of Materials Science*, vol. 36, no. 16, pp. 3833–3854, 2001.

[49] R. Wirth, "Focused ion beam (fib) combined with sem and tem: Advanced analytical tools for studies of chemical composition, microstructure and crystal structure in geomaterials on a nanometre scale," *Chemical Geology*, vol. 261, no. 3, pp. 217–229, 2009.

[50] R. Lemons and C. Quate, "Acoustic microscopescanning version," *Applied Physics Letters*, vol. 24, no. 4, pp. 163–165, 1974.

[51] J.-I. Kushibiki and N. Chubachi, "Material characterization by line-focus-beam acoustic microscope," *IEEE transactions on sonics and ultrasonics*, vol. 32, no. 2, pp. 189–212, 1985.

[52] J. Kushibiki, A. Ohkubo, and N. Chubachi, "Linearly focused acoustic beams for acoustic microscopy," *Electronics Letters*, vol. 17, no. 15, pp. 520–522, 1981.

[53] A. Briggs, *Advances in Acoustic Microscopy*, vol. 1. Springer Science & Business Media, 2013.

[54] M. Graef and M. McHenry, *Structure of Materials: An Introduction to Crystallography, Diffraction and Symmetry*. Cambridge University Press, 2007.

[55] I. Gardner *et al.*, "A review of spallation neutron source accelerators," in *Proc. EPAC*, vol. 98, 1998.

[56] F. Mees and G. S. of London, *Applications of X-ray Computed Tomography in the Geosciences*. No. v. 1 in Applications of X-ray Computed Tomography in the Geosciences, Geological Society, 2003.

[57] S. C. Mayo, A. M. Tulloh, A. Trinchi, and S. Yang, "Data-constrained microstructure characterization with multispectrum x-ray micro-ct," *Microscopy and Microanalysis*, vol. 18, no. 03, pp. 524–530, 2012.

[58] W. Li, S. D. Sharples, M. Clark, and M. G. Somekh, "Frequency spectrum spatially resolved acoustic spectroscopy for microstructure imaging," in *Ultrasonics Symposium (IUS), 2009 IEEE International*, pp. 791–794, IEEE, 2009.

[59] R. J. Smith, W. Li, J. Coulson, M. Clark, M. G. Somekh, and S. D. Sharples, "Spatially resolved acoustic spectroscopy for rapid imaging of material microstructure and grain orientation," *Measurement Science and Technology*, vol. 25, no. 5, p. 055902, 2014.

[60] S. D. Sharples, M. Clark, W. Li, and M. G. Somekh, "Rapid imaging of microstructure using spatially resolved acoustic spectroscopy," in *Proc. of the First Internat. Symp. on Laser Ultrasonics: Science, Technology and Applications (July 16–18, 2008 Montreal)*, 2008.

[61] G. Farnell, *Properties of Elastic Surface Waves, in: W.P. Mason, Physical Acoustics, vol.6.* New York: Academic Press, 1970.

[62] W. Li, S. Sharples, R. Smith, M. Clark, and M. Somekh, "Determination of crystallographic orientation of large grain metals with surface acoustic waves," *The Journal of the Acoustical Society of America*, vol. 132, p. 738, 2012.

[63] W. Li, J. Coulson, J. W. Aveson, R. J. Smith, M. Clark, M. G. Somekh, and S. D. Sharples, "Orientation characterisation of aerospace materials by spatially resolved acoustic spectroscopy," *Journal of Physics: Conference Series*, vol. 520, no. 1, p. 012017, 2014.

[64] J. Diebel, "Representing attitude: Euler angles, unit quaternions, and rotation vectors," *Matrix*, vol. 58, pp. 15–16, 2006.

[65] H. H. Xu, N. P. Padture, and S. Jahanmir, "Effect of microstructure on material-removal mechanisms and damage tolerance in abrasive machining of silicon carbide," *Journal of the American Ceramic Society*, vol. 78, no. 9, pp. 2443–2448, 1995.

[66] A. Van Pamel, C. R. Brett, and M. J. Lowe, "A methodology for evaluating detection performance of ultrasonic array imaging algorithms for coarse-grained materials," *Ultrasonics, Ferroelectrics, and Frequency Control, IEEE Transactions on*, vol. 61, no. 12, pp. 2042–2053, 2014.

[67] O. Grong and D. K. Matlock, "Microstructural development in mild and low-alloy steel weld metals," *International Metals Reviews*, vol. 31, no. 1, pp. 27–48, 1986.

[68] B. Baufeld, O. Van der Biest, and R. Gault, "Additive manufacturing of ti–6al–4v components by shaped metal deposition: microstructure and mechanical properties," *Materials & Design*, vol. 31, pp. S106–S111, 2010.

[69] T. Pollock and A. Argon, "Creep resistance of cmsx-3 nickel base superalloy single crystals," *Acta Metallurgica et Materialia*, vol. 40, no. 1, pp. 1–30, 1992.

[70] R. C. Reed, *The superalloys: fundamentals and applications*. Cambridge university press, 2006.

[71] M. Weber, *Handbook of Optical Materials*. Laser & Optical Science & Technology, Taylor & Francis, 2002.

[72] W. J. Smith, *Modern Optical Engineering*. Tata McGraw-Hill Education, 1966.

[73] S. Shingo, *Zero quality control: Source inspection and the poka-yoke system*. Productivity Press, 1986.

[74] S. Sharples, R. Light, S. Achamfuo-Yeboah, M. Clark, and M. Somekh, "The sked: speckle knife edge detector," in *Journal of Physics: Conference Series*, vol. 520, p. 012004, IOP Publishing, 2014.

[75] D. C. O'Shea, T. J. Suleski, A. D. Kathman, and D. W. Prather, *Diffractive Optics: Design, Fabrication, and Test*, vol. 62. Spie Press Bellingham, WA, 2004.

[76] D. M. Bates and D. G. Watts, *Nonlinear regression: iterative estimation and linear approximations*. Wiley Online Library, 1988.

[77] F. J. HARRIS, "On the use of windows for harmonic analysis with the discrete fourier transform," *Proceedings of The IEEE*, vol. 66, pp. 51–83, 1978.

[78] L. B. Almeida, "The fractional fourier transform and time-frequency representations," *Signal Processing, IEEE Transactions on*, vol. 42, no. 11, pp. 3084–3091, 1994.

[79] K. Prabhu, *Window functions and their applications in Signal Processing*. CRC Press, 2013.

[80] I. Daubechies *et al.*, "Orthonormal bases of compactly supported wavelets," *Communications on Pure and Applied Mathematics*, vol. 41, no. 7, pp. 909–996, 1988.

[81] C. Chui, *An Introduction to Wavelets*. Wavelet analysis and its applications, Academic Press, 1992.

[82] C. Torrence and G. P. Compo, "A practical guide to wavelet analysis," *Bulletin of the American Meteorological society*, vol. 79, no. 1, pp. 61–78, 1998.

[83] H. Weng and K. Lau, "Wavelets, period doubling, and time-frequency localization with application to organization of convection over the tropical western pacific," *Journal of the Atmospheric Sciences*, vol. 51, no. 17, pp. 2523–2541, 1994.

[84] M. Farge, "Wavelet transforms and their applications to turbulence," *Annual review of fluid mechanics*, vol. 24, no. 1, pp. 395–458, 1992.

[85] M. Afifi, A. Fassi-Fihri, M. Marjane, K. Nassim, M. Sidki, and S. Rachafi, "Paul wavelet-based algorithm for optical phase distribution evaluation," *Optics Communications*, vol. 211, no. 16, pp. 47 – 51, 2002.

[86] C. K. Chui and J.-z. Wang, "A cardinal spline approach to wavelets," *Proceedings of the American Mathematical Society*, vol. 113, no. 3, pp. 785–793, 1991.

[87] C. K. Chui and J.-z. Wang, "On compactly supported spline wavelets and a duality principle," *Transactions of the American Mathematical Society*, vol. 330, no. 2, pp. 903–915, 1992.

[88] M. Unser, "Ten good reasons for using spline wavelets," in *Storage and Retrieval for Image and Video Databases*, 1997.

[89] M. Donachie, *Titanium: a Technical Guide, 2nd Edition.* ASM International, 2000.

[90] M. Levy, H. Bass, and R. Stern, *Handbook of Elastic Properties of Solids, Liquids, and Gases, Four-Volume Set.* Elsevier Science, 2000.

[91] J. Achenbach, "Laser excitation of surface wave motion," *Journal of the Mechanics and Physics of Solids*, vol. 51, no. 11, pp. 1885–1902, 2003.

[92] J. Achenbach, "Simplifications for the calculation of surface wave pulses generated by laser-irradiation," *The Journal of the Acoustical Society of America*, vol. 116, no. 3, pp. 1481–1487, 2004.

[93] M. Subbarao, T.-S. Choi, and A. Nikzad, "Focusing techniques," *Optical Engineering*, vol. 32, no. 11, pp. 2824–2836, 1993.

[94] F. C. Groen, I. T. Young, and G. Ligthart, "A comparison of different focus functions for use in autofocus algorithms," *Cytometry*, vol. 6, no. 2, pp. 81–91, 1985.

[95] J.-H. Lee, K.-S. Kim, B.-D. Nam, J.-C. Lee, Y.-M. Kwon, and H.-G. Kim, "Implementation of a passive automatic focusing algorithm for digital still camera," *Consumer Electronics, IEEE Transactions on*, vol. 41, no. 3, pp. 449–454, 1995.

[96] Y. Sun, S. Duthaler, and B. J. Nelson, "Autofocusing in computer microscopy: selecting the optimal focus algorithm," *Microscopy Research and Technique*, vol. 65, no. 3, pp. 139–149, 2004.

[97] J.-M. Geusebroek, F. Cornelissen, A. W. Smeulders, and H. Geerts, "Robust autofocusing in microscopy," *Cytometry*, vol. 39, no. 1, pp. 1–9, 2000.

[98] A. Santos, C. Ortiz de Solorzano, J. J. Vaquero, J. Pena, N. Malpica, and F. Del Pozo, "Evaluation of autofocus functions in molecular cytogenetic analysis," *Journal of microscopy*, vol. 188, no. 3, pp. 264–272, 1997.

[99] S. Pertuz, D. Puig, and M. A. Garcia, "Analysis of focus measure operators for shape-from-focus," *Pattern Recognition*, vol. 46, no. 5, pp. 1415–1432, 2013.

[100] E. Krotkov and J.-P. Martin, "Range from focus," in *Robotics and Automation. Proceedings. 1986 IEEE International Conference on*, vol. 3, pp. 1093–1098, IEEE, 1986.

[101] S. Pertuz, D. Puig, and M. A. Garcia, "Analysis of focus measure operators for shape-from-focus," *Pattern Recognition*, vol. 46, no. 5, pp. 1415–1432, 2013.

[102] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*. Cengage Learning, 2014.

[103] M. Subbarao and T.-C. Wei, "Depth from defocus and rapid autofocusing: a practical approach," in *Computer Vision and Pattern Recognition, 1992. Proceedings CVPR'92., 1992 IEEE Computer Society Conference on*, pp. 773–776, IEEE, 1992.

[104] J. L. Pech-Pacheco, G. Cristóbal, J. Chamorro-Martinez, and J. Fernández-Valdivia, "Diatom autofocusing in brightfield microscopy: a comparative study," in *Pattern Recognition, 2000. Proceedings. 15th International Conference on*, vol. 3, pp. 314–317, IEEE, 2000.

[105] T. Saito, O. Matsuda, and O. Wright, "Picosecond acoustic phonon pulse generation in nickel and chromium," *Physical Review B*, vol. 67, no. 20, p. 205421, 2003.

[106] D. Hurley, O. Wright, O. Matsuda, V. Gusev, and O. Kolosov, "Laser picosecond acoustics in isotropic and anisotropic materials," *Ultrasonics*, vol. 38, no. 1, pp. 470–474, 2000.

[107] V. Gusev and O. Wright, "Ultrafast nonequilibrium dynamics of electrons in metals," *Physical Review B*, vol. 57, no. 5, p. 2878, 1998.

[108] R. White, "Generation of elastic waves by transient surface heating," *Journal of Applied Physics*, vol. 34, no. 12, pp. 3559–3567, 1963.

[109] G. Cachier, "Laser excitation of microwave sound in solids," *The Journal of the Acoustical Society of America*, vol. 49, p. 974, 1971.

[110] R. Von Gutfeld and R. Melcher, "20-MHz acoustic waves from pulsed thermoelastic expansions of constrained surfaces," *Applied Physics Letters*, vol. 30, no. 6,

pp. 257–259, 1977.

[111] E. Ash, E. Dieulesaint, and H. Rakouth, "Generation of surface acoustic waves by means of a cw laser," *Electronics Letters*, vol. 16, no. 12, pp. 470–472, 1980.

[112] C. Scruby, R. Dewhurst, D. Hutchins, and S. Palmer, "Quantitative studies of thermally generated elastic waves in laser-irradiated metals," *Journal of Applied Physics*, vol. 51, no. 12, pp. 6210–6216, 1980.

[113] R. Dewhurst, D. Hutchins, S. Palmer, and C. Scruby, "Quantitative measurements of laser-generated acoustic waveforms," *Journal of Applied Physics*, vol. 53, no. 6, pp. 4064–4071, 1982.

[114] L. Rose, "Point-source representation for laser-generated ultrasound," *The Journal of the Acoustical Society of America*, vol. 75, p. 723, 1984.

[115] J. Spicer, *Laser ultrasonics in finite structures: comprehensive modelling with supporting experiment.* PhD thesis, The John Hopkins University, 1991.

[116] F. McDonald, "On the precursor in laser-generated ultrasound waveforms in metals," *Applied Physics Letters*, vol. 56, no. 3, pp. 230–232, 1990.

[117] U. Schleichert, K. Langenberg, W. Arnold, and S. Fassbender, "A quantitative theory of laser-generated ultrasound," *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 8, pp. 489–496, 1989.

[118] T. Sanderson, C. Ume, and J. Jarzynski, "Hyperbolic heat equations in laser generated ultrasound models," *Ultrasonics*, vol. 33, no. 6, pp. 415–421, 1995.

[119] T. Murray, S. Krishnaswamy, and J. Achenbach, "Laser generation of ultrasound in films and coatings," *Applied Physics Letters*, vol. 74, p. 3561, 1999.

[120] M. Dubois, F. Enguehard, L. Bertrand, M. Choquet, and J. Monchalin, "Modeling of laser thermoelastic generation of ultrasound in an orthotropic medium," *Applied Physics Letters*, vol. 64, no. 5, pp. 554–556, 1994.

[121] A. McKie, J. Wagner, J. Spicer, and C. Penney, "Laser generation of narrowband and directed ultrasound," *Ultrasonics*, vol. 27, no. 6, pp. 323–330, 1989.

[122] P. Doyle and C. Scala, "Near-field ultrasonic rayleigh waves from a laser line source," *Ultrasonics*, vol. 34, no. 1, pp. 1–8, 1996.

[123] J. Bernstein and J. Spicer, "Line source representation for laser-generated ultrasound in aluminum," *The Journal of the Acoustical Society of America*, vol. 107, p. 1352, 2000.

[124] I. Arias and J. Achenbach, "Thermoelastic generation of ultrasound by laser line-source irradiation," *Int. J. Solids and Struct*, vol. 40, pp. 6917–6935, 2003.

[125] R. Stoneley, "The seismological implications of aeolotropy in continental structure," *Royal Astronomical Soc. Monthly Notices, Geophysical Supplement*, pp. 343–353, 1949.

[126] R. Stoneley, "The propagation of surface elastic waves in a cubic crystal," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 232, no. 1191, pp. 447–458, 1955.

[127] V. Buchwald, "Rayleigh waves in anisotropic media," *The Quarterly Journal of Mechanics and Applied Mathematics*, vol. 14, no. 4, pp. 461–469, 1961.

[128] D. Gazis, R. Herman, and R. Wallis, "Surface elastic waves in cubic crystals," *Physical Review*, vol. 119, no. 2, p. 533, 1960.

[129] V. Buchwald, "Rayleigh waves in transversely isotropic media," *The Quarterly Journal of Mechanics and Applied Mathematics*, vol. 14, no. 3, pp. 293–318, 1961.

[130] V. Buchwald and A. DAVIS, "Surface waves in elastic media with cubic symmetry," *The Quarterly Journal of Mechanics and Applied Mathematics*, vol. 16, no. 3, pp. 283–294, 1963.

[131] K. Ingebrigtsen and A. Tonning, "Elastic surface waves in crystals," *Physical Review*, vol. 184, no. 3, p. 942, 1969.

[132] T. Lim and G. Farnell, "Search for forbidden directions of elastic surface-wave propagation in anisotropic crystals," *Journal of Applied Physics*, vol. 39, no. 9, pp. 4319–4325, 1968.

[133] G. Farnell, *Properties of Elastic Surface Waves, in: W.P. Mason, Physical Acoustics, vol.6.* New York: Academic Press, 1970.

[134] D. Royer and E. Dieulesaint, "Rayleigh wave velocity and displacement in orthorhombic, tetragonal, hexagonal, and cubic crystals," *The Journal of the Acoustical Society of America*, vol. 76, no. 5, pp. 1438–1444, 1984.

[135] A. Every and W. Sachse, "Imaging of laser-generated ultrasonic waves in silicon," *Physical Review B*, vol. 44, no. 13, p. 6689, 1991.

[136] A. De Hoop, "A modification of cagniard's method for solving seismic pulse problems," *Applied Scientific Research, Section B*, vol. 8, no. 1, pp. 349–356, 1960.

[137] L. Cagniard, E. Flinn, and C. Dix, *Reflection and Refraction of Progressive Seismic Waves*. McGraw-Hill Book Company, 1962.

[138] A. Mourad, M. Deschamps, and B. Castagnede, "Acoustic waves generated by a transient line source in an anisotropic half-space.," *Acta Acustica united with Acustica*, vol. 82, no. 6, pp. 839–851, 1996.

[139] D. Hurley and J. Spicer, "Point-source representation for laser-generated ultrasound in an elastic, transversely isotropic half-space," *Journal of Applied Physics*, vol. 86, no. 6, pp. 3423–3427, 1999.

[140] D. Hurley and J. Spicer, "Line source representation for laser-generated ultrasound in an elastic transversely isotropic half-space," *The Journal of the Acoustical Society of America*, vol. 116, p. 2914, 2004.

[141] M. Dubois, F. Enguehard, L. Bertrand, M. Choquet, and J. Monchalin, "Modeling of laser thermoelastic generation of ultrasound in an orthotropic medium," *Applied Physics Letters*, vol. 64, no. 5, pp. 554–556, 1994.

[142] D. Hurley, "Laser-generated thermoelastic acoustic sources in anisotropic materials," *The Journal of the Acoustical Society of America*, vol. 115, p. 2054, 2003.

[143] B. Xu, Z. Shen, X. Ni, J. Lu, and Y. Wang, "Finite element model of laser-generated surface acoustic waves in coating-substrate system," *Journal of Applied Physics*, vol. 95, no. 4, pp. 2109–2115, 2004.

[144] B. Xu, Z. Shen, J. Wang, X. Ni, J. Guan, and J. Lu, "Thermoelastic finite element modeling of laser generation ultrasound," *Journal of Applied Physics*, vol. 99, no. 3, p. 033508, 2006.

[145] B. Xu, J. Feng, G. Xu, J. Wang, H. Sun, and G. Cao, "Laser-generated thermoelastic acoustic sources and lamb waves in anisotropic plates," *Applied Physics A*, vol. 91, no. 1, pp. 173–179, 2008.

[146] H. Al-Qahtani and S. K. Datta, "Thermoelastic waves in an anisotropic infinite plate," *Journal of Applied Physics*, vol. 96, no. 7, pp. 3645–3658, 2004.

[147] I. A. Veres, T. Berer, and P. Burgholzer, "Numerical modeling of thermoelastic generation of ultrasound by laser irradiation in the coupled thermoelasticity," *Ultrasonics*, vol. 53, no. 1, pp. 141–149, 2013.

[148] S. Hooker and C. Webb, *Laser Physics*. Oxford master series in physics: Atomic, optical, and laser physics, Oxford Univ Press, 2010.

[149] B. Childs, "The thermal expansion of anisotropic metals," *Reviews of Modern Physics*, vol. 25, no. 3, p. 665, 1953.

[150] H. Seiner and M. Landa, "Differential geometry of ray surfaces in anisotropic solids and its contribution to nde: Modelling and experiment," *Ultrasonics*, vol. 44, 2006.

[151] O. Cho and M. J. Lai, "A class of compactly supported orthonormal b spline wavelets," *Splines and Wavelets*, 2005.

[152] M. Andrle and L. Rebollo-Neira, "Cardinal b-spline dictionaries on a compact interval," *Applied and Computational Harmonic Analysis*, vol. 18, pp. 336–346, 2005.

[153] A. Bower, *Applied Mechanics of Solids*. CRC Press, 2009.

[154] G. W. C. Kaye and T. H. Laby, "Tables of physical and chemical constants," 1973.