All-Optical Scanning Acoustic Microscope

by Steve D. Sharples, MEng

Thesis submitted to The University of Nottingham

for the degree of Doctor of Philosophy, May 2003

The University of Nottingham

Contents

Abstract

In this thesis a new instrument, the all-optical scanning acoustic microscope (O-SAM) is presented, it is a non contact scanning acoustic microscope (SAM) which uses lasers to both generate and detect surface acoustics waves (SAWs)

The non contact nature of the O-SAM overcomes some difficulties associated with conventional SAMs because of the couplant and surface contact involved. This O-SAM also overcomes many of the problems associated with conventional laser ultrasound systems including those of sample damage and ablation, low signal to noise ratio and slow data acquisition. Furthermore, the instrument is *adaptive* enabling it to compensate to acoustic aberrations that can occur as a result of material microstructure. We believe this is a most significant feature that will greatly enhance its range of applications.

This thesis examines some of the key technological developments required to develop the O-SAM, including the use of tailored optical generation pattern realised through the use of a spatial light modulator (SLM) and the development of a novel high speed analogue data acquisition system.

This thesis presents the design and construction of the instrument and demonstrates its imaging capability on engineering materials using SAWs at 82 and 164MHz although the instrument is potentially capable of imaging at much higher frequencies. Images are presented on ceramic, steel and aluminium samples which demonstrate a range of contrast mechanisms and measurement techniques, including the interaction of the material with Rayleigh waves and Lamb modes for the purpose of defect detection and characterisation, and measurement of local variations in residual surface stress, changes in coating thickness, and plate thickness.

Acknowledgements

I'd like to thank Mike Somekh, my supervisor, for giving me the spark of enthusiasm for optics during my undergraduate course, which led to the opportunity to carry out the work presented in this thesis. I am very grateful for all the help I've received; the level of enthusiasm and support for my work he has given never ceases to surprise me, but I am very thankful for it.

I'm also very grateful to Matt Clark who, over the entire time I have been working with him, has constantly provided ideas, suggestions, solutions, support, and much inspiration. In addition, he has tirelessly and ceaselessly hassled me about writing up for longer than I can remember. Not that it's made any difference, of course.

Thanks also go to my colleagues past and present, amongst them Francis, Jason, Mark, See, Boon Hean, Yut, Hong, Azhar and José.

Amongst my friends, I'd especially like to thank Diana, who got me through a very rough patch during the course of my study, and for that I will always be grateful. Thanks also to Sara (another write-up hassler), Elli, and the Scooby Gang for providing interesting alternatives to writing up, and to Karen and Nick for just being ace when I needed it the most.

Most of all I'd like to thank Lisa, with all my heart, for pretty much everything.

Chapter 1

Introduction

The work presented in this thesis is concerned with the realisation of a totally non-contact method of acquiring images of the interaction between surface acoustic waves (SAWs) and a solid material. The system uses lasers to generate and detect the SAWs, and has been developed from an accurate non-contacting laser based system for surface wave velocity measurement [1]. The system shall be termed the *O-SAM*—an acronym of *optical scanning acoustic microscope*— for the duration of this thesis. Immediately prior to the work described in this thesis, the SAW excitation method used in the system had begun to evolve, to incorporate custom computer generated holographic zone plates [2] to increase the SAW amplitude, and thus to pave the way for SAW imaging [3] and methods of SAW frequency control [4]. Much of this work is covered in the PhD thesis of F. Linnane [5].

The work carried out, which will be described in this thesis, will be preceded with a brief overview of the role of ultrasonics in nondestructive testing (NDT), and the role of laser-based systems for excitation and detection of ultrasound within that context will be demonstrated. Furthermore, the more recent developments of imaging using laser ultrasound systems will be discussed in chapter 2.

1.1 Ultrasonics

Ultrasound refers to sound waves of frequencies above the normal range of human hearing, the upper limit of which is around 20kHz. Ultrasound can propagate in gases, liquids and solids. The waves are elastic in nature and their propagation is not affected by the optical properties of the material in question.

1.1.1 Use of ultrasound in NDT

Every few years the term ‘nondestructive testing’ and its abbreviation (NDT) undergo periods of unpopularity and are replaced by other terms. ‘Nondestructive inspection technology’ and ‘nondestructive inspection’ have both been fashionable in the past, and ‘nondestructive evaluation’ (NDE) seems to currently be in vogue. Nevertheless, the term NDT will be used here, since it sufficiently describes the process.

NDT involves testing an object for its integrity and fitness for the use for which it was designed, without damaging the object, and thus making it unfit for its purpose. There are many different general methods of testing objects nondestructively. These methods include optical inspection, radiographic testing, magnetic particle inspection, penetrant testing, holographic testing, acoustic emission and, of course, ultrasonic testing.

In reference [6], NDT is described as, “...an amalgam of three inseparable aspects: methods, instruments, and intelligence.” This can be illustrated by one of the oldest ‘nondestructive’ tests, and one from which the use of modern-day ultrasonics to test engineering materials is derived: that of hitting an object with a hammer, and listening to the noise that is emitted. For many simple shapes, solid homogeneous objects emit ‘clear, bell-like’ sounds, whilst defective objects emit ‘dull’ sounds, or sounds with many harmonics. The *method* of the test is to hit the object with a hammer, and listening to the emitted tone. The *instruments* are the hammer, and the ears of the person performing the test. The success of the test obviously depended on the *intelligence*—in the form of skill and experience—of the tester, and their ability to interpret the data acquired in the process of the test, namely the noise emitted. All three of the ‘three inseparable aspects’ are important for the realisation of a successful method of nondestructive testing. There is little use, in a practical sense, for an instrument that performs a certain kind of test and takes a certain type of measurement, that cannot usefully be interpreted.

The role of ultrasonics within the NDT community has increased significantly over the past fifty or so years. A completely non-scientific indicator of this is the amount of space given over to ultrasonics in reviews of NDT methods at various points in time. For instance, a book entitled *“The Testing and Inspection of Engineering Materials”* [7] published in 1964 devotes less than three pages out of more than 400 on the subject of ultrasonic testing; a book entitled *“Nondestructive Evaluation, A Tool in Design, Manufacturing and Service”* [8] (note the change of term) published in 1997 devotes more than 120 pages, out of 560, to it.

1.2 Methods of flaw detection and material characterisation

The role of ultrasonics in NDT can be split broadly into two areas; that of flaw detection, and material characterisation. There is also a slight crossover in the area of flaw characterisation. The interaction of ultrasonic waves with the material and features therein can be used to ascertain the presence of flaws (and perhaps their nature) or some physical or chemical characteristic of the material under test.

1.2.1 Flaw detection

A common requirement for an NDT process is to determine whether any flaws are present that will adversely affect the operation or lifespan of a component. Typical flaws include surface-breaking or interior cracks, weak joints, coating delamination or loss of adhesion between coating and substrate.

There are several methods of detecting flaws with ultrasound, and obviously different methods are appropriate for different types of flaws. In addition, different types of ultrasonic waves—in terms of longitudinal or shear waves, bulk or surface waves, wave mode, frequency, frequency spectrum, and amplitude—are generally more appropriate than others for different tasks.

In general, the higher the ultrasound frequency, the higher the resolution—both spatial and temporal—but with the disadvantage that the losses due to scattering and aberration are higher. The use of a broad or narrowband frequency spectrum of the waves is usually determined by the transducers used to excite—and also possibly detect—the ultrasound, as discussed in section 1.4.

There are almost as many techniques for examining sample properties acoustically as there are applications [9]. The aim here is not to go into too much detail, but to illustrate the basic principles of flaw detection.

In terms of inspection of flaws in the interior of materials, the techniques can be broadly split into two areas: the *reflection* or *pulse-echo* technique, and the *transmission* or *pitch-catch* technique [8]. These two broad areas can be further split into *normal beam* and *angle beam* techniques. Figure 1.1 illustrates schematically the different techniques. The insets show how the received signals relate to the defects present.

The use of a particular technique depends on many factors, including sample geometry, attenuation, and the characteristics of any possible flaws. In the pulse-echo techniques, the existence of a flaw is determined by the reflection of ultrasound, which is detected by the combined transmitter/receiver. For the normal-beam mode of operation, this will be in addition to the reflection from the back surface of the object under test; no such back reflection is present for the angle-beam mode. The location of the flaw from the transmitter can be calculated by the 'time of flight' of the received ultrasound if the velocity of the ultrasound in the material under investigation is known.

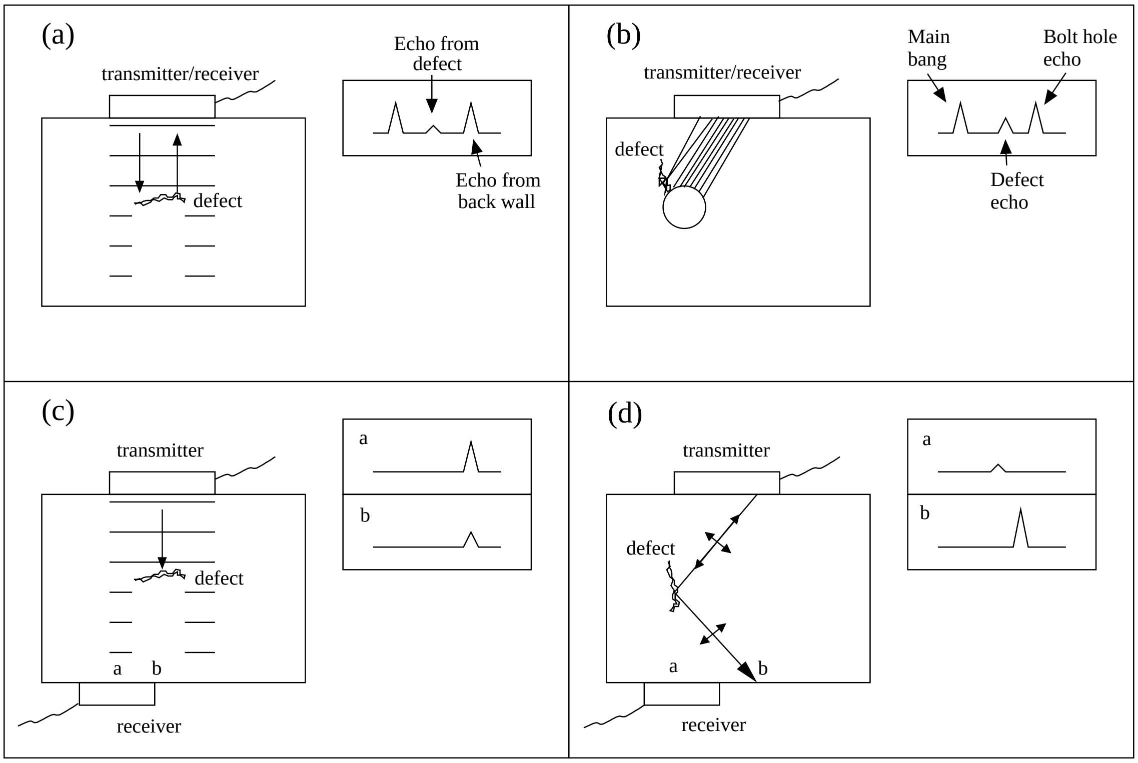

Figure 1.1: Illustrations of the four main techniques of detecting flaws in the interior or materials: (a) reflection, normal beam; (b) reflection, angle beam; (c) transmission, normal-beam; (d) transmission, angle beam. The type of signals received relating to the defect are shown as inserts.

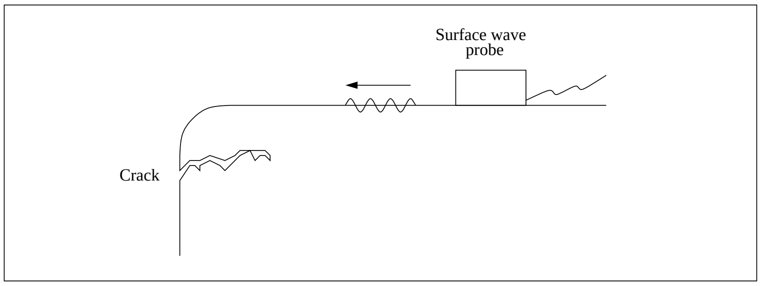

Figure 1.2: Detection of a surface-breaking crack by the reflection of a surface wave.

When the transmission technique is used, the existence of a flaw is indicated by an attenuation of the received signal.

The use of ultrasound to detect the presence of flaws is not limited to the interior of materials. Waves propagating on the surface of the material can be used to detect surface, or near-surface defects much more effectively than bulk waves. There is a lower limit to the depth over which bulk waves can provide useful information, related to the shortest pulse length that the transducer can achieve. Figure 1.2 illustrates the detection of a surface-breaking crack by the reflection of a surface wave.

A type of SAW of particular interest to NDT is the Rayleigh wave, the properties of which are described in [10] and [9]. They propagate at the boundary between a solid half-space and air. Most importantly for NDT applications, their penetration depth is approximately equal to their wavelength. This makes them particularly suitable for probing the near-surface of a material. They can propagate around curved surfaces with little or no attenuation, provided the radius of curvature is not 'too small'—radii of greater than twice the Rayleigh wavelength produce negligible attenuation [10]. They can therefore be used to probe areas that are physically difficult to get to.

Rayleigh waves may be described as guided waves, since they are limited to the surface of a material. Another form of guided wave is the Lamb wave, and these propagate in situations where two parallel surfaces are found, for instance sheets of metal, or pipes. An important property of Lamb waves is that they are *dispersive*, in that their frequency and the plate thickness determines their velocity. Changes in plate thickness—for instance due to corrosion—therefore affect the velocity of the detected waves.

1.2.2 Material characterisation

The characterisation of materials is another important area of NDT. Characterisation can include ascertaining various dimensions (for example the thickness of the steel hull of a boat, a subject

I have a particular personal interest in), coating thickness, porosity, residual stress, grain size in the case of multi-grained materials such as metals, or grain orientation in single-grained materials such as silicon. The characterisation may be required mid-process, in which case the ultrasonic inspection may provide a means of process control, or post-process, in which case it provides a means of quality management.

In terms of geometry characterisation, the pulse-echo method, described in the previous section is the most common form of ultrasonic inspection, whereby the time taken for an ultrasonic pulse to be reflected off the rear surface of the material is related to the thickness of the material. Knowledge of the wave velocity is necessary, but for many materials this has been well known for many years; the alternative of course is calibration. For thin plates, the thickness can be determined by the Lamb velocity.

The velocity of ultrasonic waves is the primary means of many forms of characterisation, since it is a function of the moduli of elasticity and density. Changes in velocity can therefore correlate with changes in the material properties such as porosity, residual stress or, in the case of surface waves, coating thickness. If the change in property is the result of a process that the material is being subjected to, then examination of the velocity may provide a means of monitoring the process concerned.

1.3 The range of information available from ultrasonic inspection

The previous section discussed the mechanisms of the interaction between the acoustic waves and features or properties of the material under inspection, and the various geometries employed to acquire information about this interaction. 'Time of flight' measurements have been briefly mentioned, and the insets of figure 1.1 give some kind of indication of one one way of 'viewing' the data.

The role of the ultrasonic transducer is to convert time-varying electrical energy to time-varying acoustic energy in the case of the transmitter, or back the other way in the case of the receiver. It is very easy to perceive viewing the signal that is output from the receiver, in which case the amplitude of the electric signal (representing the acoustic field) is shown on the y-axis, and the time subsequent to the input pulse is shown on a horizontal t-axis. This is generally known as an *a-scan*. Time on the t-axis is proportional to the propagation distance z into the material in the case of bulk waves, or either the x or y direction in the case of surface waves. Countless examples of a-scans exist in the literature—plucking an example out of the many that exist, is a paper in which

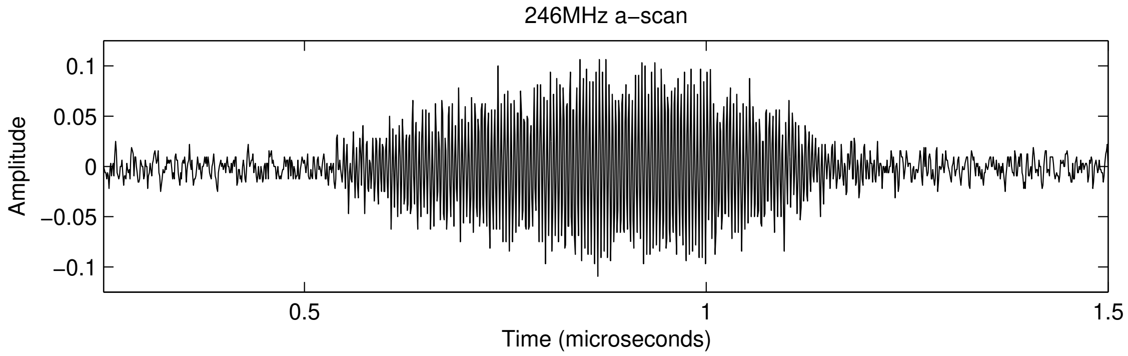

Figure 1.3: An a-scan of a 246MHz Rayleigh wave, acquired with the O-SAM instrument.

several different a-scans from different types of noncontact detector are illustrated and compared [11]. Figure 1.3 illustrates an a-scan acquired with the O-SAM instrument.

In the *b-scan,* time is plotted on one axis, the t-axis. The other axis, the y-axis, represents the distance y along the material under inspection, where y is a direction normal to the propagation direction z. The plot is usually represented by a grey-scale or colour image, whereby the intensity or colour of each point on the image is proportional to the amplitude of the received signal. Examples of b-scans can be found in a paper describing how they may be used for quantitative time-of-flight analysis of shear waves and longitudinal-shear mode-converted waves [12, 13] (reference [13] has the advantage of colour pictures but is very similar)—the paper is quite a good illustration of the role of b-scans, since it shows how a-scans relate to b-scans, and how trends may be spotted much more easily. Figure 1.4 shows two b-scans acquired with the O-SAM instrument.

The terminology and definitions get a little more complicated (and often contradictory) when we progress to *c-scans* and general imaging—that is, the representation of some acoustic parameter over an area of the material under inspection. As an illustration of the range of terms available, it is possible to perform 'c-scan imaging,' 'scanning acoustic microscopy,' 'scanned ultrasonic microscopy' and 'acoustic imaging.' It is generally the case that 'scanned ultrasonic microscopy' and 'scanned acoustic microscopy' can be interchanged, as can 'c-scan imaging' and 'acoustic imaging' [14]. In [14], it is argued that ultrasonic/acoustic microscopy creates 'magnified' images of objects, which involves acquiring data from many small areas, and that therefore "...either the interrogating acoustic beam must be focused, or the energy used to create the acoustic signal must be focused, in order to isonify each of these small areas independently." This is one interpretation of the terms and, if a distinction needs to be made between the different sorts of images, a reasonably good one. In any case, it conveniently lends some credibility to the title of this thesis, since the acoustic waves are usually focused to acquire images. Within the scope of this thesis, however, the

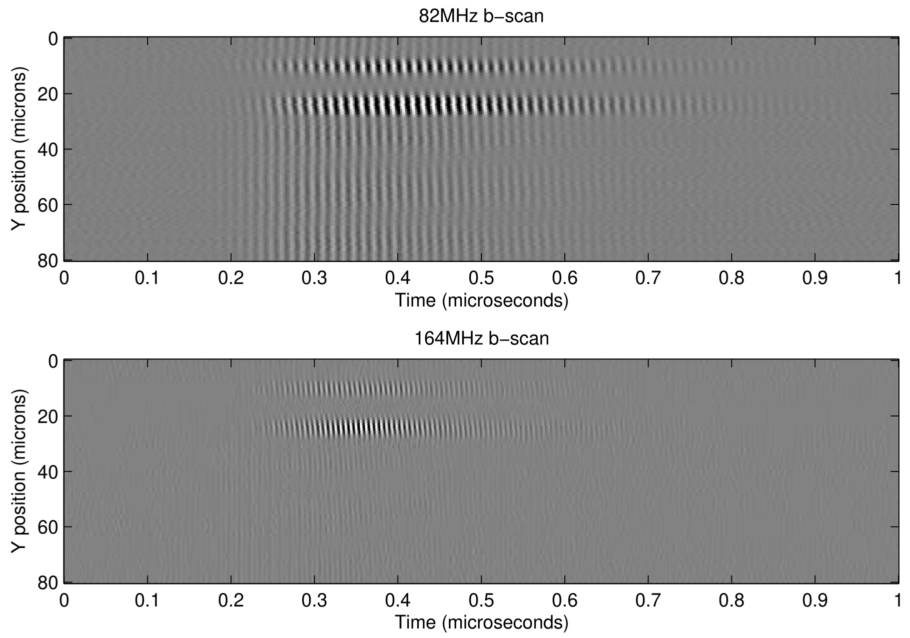

Figure 1.4: Two b-scans of 82MHz (top) and 164MHz (bottom) Rayleigh waves, acquired with the O-SAM instrument.

acoustic images will be referred to as c-scan images.

We then come to the problem of which parameter of the received ultrasound signal we are illustrating; examples include the gated peak amplitude, the amplitude of a certain frequency component, the phase of a certain frequency component, the arrival time of the first echo etc. In many cases, the complete waveform may be acquired and stored for each point in the image, in which case a number of different c-scan images, each showing a different acoustic parameter, may be used in analysis. This may be prohibitively slow (in the case of digital data acquisition systems) or require too much storage space or bandwidth to extract the data from the transducer to the storage medium to be feasible for medium to large scans, and so a decision must usually be made as to the type of data—to illustrate a certain contrast—that is to be acquired.

Arguably, the contrast mechanism may not necessarily arise from the direct conversion of acoustic energy to electrical energy—a paper from 1993 [15] images the SAW amplitude by dust patterns arising as a result of result of SAW-induced dust-particle removal from the surface. Many examples of acoustic imaging/microscopy exist in the literature, with several good ones in [14] and [16]. Figure 1.5 illustrates amplitude and phase c-scan images acquired with the O-SAM instrument. The images are of a boundary between a region of poorly-bonded aluminium and a region of well-bonded aluminium, and are discussed further in section 8.4.

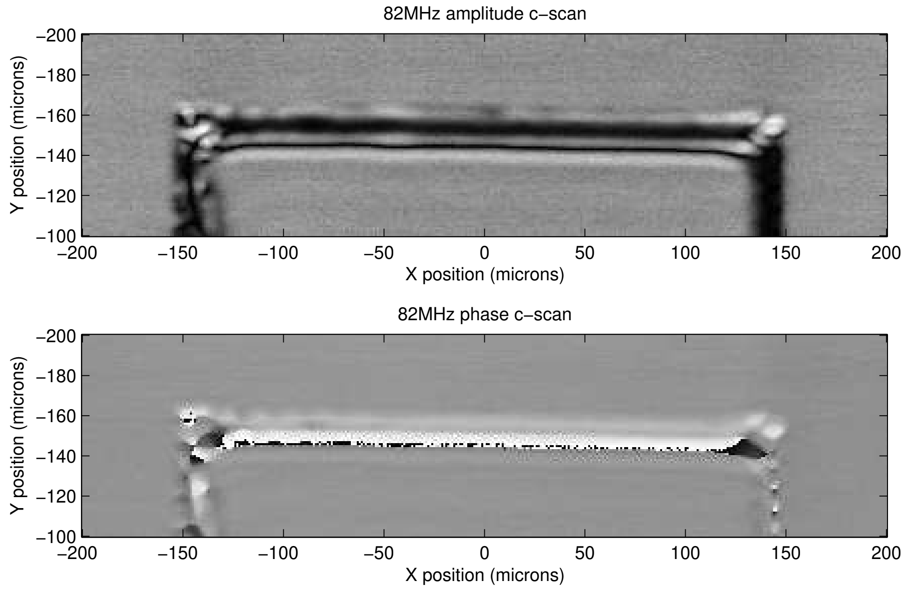

Figure 1.5: 82 MHz Amplitude (top) and phase (bottom) c-scans acquired with the O-SAM instrument.

The three general methods of illustrating acquired acoustic data—a-scans, b-scans and c-scans—each have their relative advantages and disadvantages. A-scans provide the 'most information,' since you effectively see everything that the receiving transducer receives. It is difficult to use this data as part of a larger group of data for similar points, and this is where the advantages of the b-scan area apparent, since it is easier to spot trends. Finally, c-scan imaging provides both a method of illustrating trends, and presents the data in a form which is easier for humans to interpret, since there is a direct mapping between spatial locations on the data and on the material under investigation.

Note that the method of illustrating the data is a different topic from the method of illustrating the *analysis* of the data. Using the data acquired—whether it be a single, one-, two- or three-dimensional array of a-scans, or a thresholded two-dimensional array of peak amplitude—many different forms of analysis can be performed. This can include spatial or temporal frequency analysis, trend plots, histograms, and dispersion curves. The data acquired from analysis may also be illustrated in a number of different ways, whether it be a 'green light = good sample, red light = bad sample' indicator, or a complex three dimensional image of residual strain over an object.

1.4 General techniques for excitation and measurement of ultrasound

Although the previous sections have described several methods of flaw detection and material characterisation, very little has been said of the means by which the ultrasound is generated and detected. That will be addressed in this section. The relative advantages and disadvantages of each of the techniques will be discussed at the end of the section.

1.4.1 Piezoelectric transducers

By far the most common method of exciting and detecting ultrasound within the field of non-destructive testing is the *piezoelectric transducer*. If a time-varying electrical potential is applied across two faces of a piezoelectric material, then the material will expand and contract with the electric field. Conversely, if the material is contracted and expanded by an external force, a time-varying electrical potential will be generated. Due to the large difference in acoustic impedance between the transducer and air, a liquid couplant must be used to efficiently transfer the ultrasound between the transducer and the material. The use of piezoelectric transducers is therefore very much a 'contact technique.' The couplant may be a thin layer of liquid (water or oil, for instance) or a tank of liquid in the case of immersion.

A vast range of piezoelectric transducers exist to service the requirements of the NDT industry, covering a range of geometries and frequency ranges, for focused and plane wave applications. Most transducers are built with longitudinal wave active elements—they produce waves normal to their front surface which radiate into the material. To generate shear waves for the angle beam configuration, if a shear wave transducer is not used, it is necessary to insert an angle block between the transducer and the material, with liquid couplant between each interface.

1.4.2 Scanning acoustic microscopy

Scanning acoustic microscopy [17] is a technique that uses piezoelectric transducers with a liquid couplant—as described in the previous section—but in one of its modes of operation uses different contrast mechanisms to acquire information about the material under investigation, and hence deserves a section of its own.

A schematic of the scanning acoustic microscope (SAM) in the $V(z)$ mode is shown in figure 1.6. The piezoelectric transducer is attached to an acoustic lens, which is in turn coupled to the material with a liquid couplant. The acoustic lens is positioned a certain distance z from the material surface. A tone burst of longitudinal waves couple at the sample surface into 'leaked'

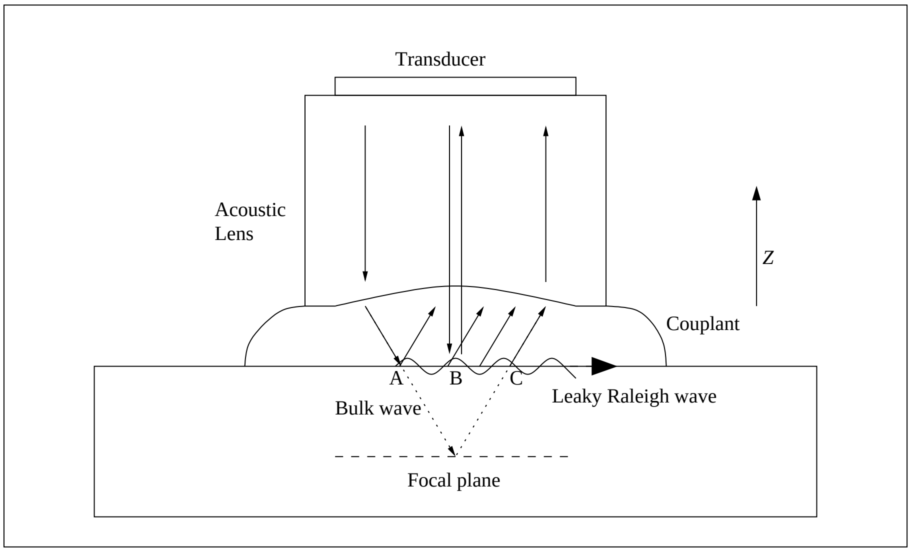

Figure 1.6: Schematic of the scanning acoustic microscope (SAM) in $V(z)$ mode.

Rayleigh waves at the edge of the acoustic lens—point A in the figure. These propagate to the far side of the lens (point C), and are recoupled back into longitudinal waves and are detected by the transducer via the acoustic lens. In addition, the directly reflected waves from the centre of the lens are detected at point B. If the tone burst is of a sufficiently long duration to allow interference between the directly reflected wave and the leaked wave, then the output is the vector sum of the two waves for the time period that they interfere. The contrast mechanism comes from the velocity of the leaked Rayleigh waves; as the material properties change, the velocity and hence the relative phase of the leaked waves and the direct reflection varies, producing an output whereby the change in amplitude is related to the change in velocity.

Furthermore, if the sample is defocused in the z direction an oscillating pattern will be produced due to the phase difference between the two rays. The periodicity of this oscillating pattern is simply related to the the surface wave velocity.

1.4.3 EMATs

EMATs (electromagnetic-acoustic transducers) use the principle that an electromagnetic wave incident on the surface of an electrical conductor induces eddy currents within the skin of the conductor [8]. The source of the acoustic waves are the Lorentz forces which are excited by a current-carrying coil close to the surface of the sample under investigation. Shear or longitudinal waves may be excited by rotating the direction of the constant magnetic field by 90°. The reverse process may be used to detect acoustic waves.

Although the current-carrying coil of an EMAT generates a Lorentz force in all electrically-conductive materials, the dominant force produced in ferromagnetic materials is due to the magnetostrictive effect. This force is parallel to the applied magnetic field (the Lorentz force is 90° to it) and so EMAT probes for ferrite materials are of different designs to those for nonferrite materials.

1.4.4 Lasers

Since the use of lasers to generate and detect ultrasound is one of the main topics of this thesis, a more in-depth study of their role as ultrasound transducers will be presented in chapter 2, particularly with respect to imaging systems. However, for completeness, they will be mentioned here, so that the relative strengths and weaknesses of the various techniques can be compared. A substantial review of the various laser ultrasound techniques can be found in [18].

The use of lasers to generate and detect ultrasound is different from the previously mentioned techniques in that different mechanisms are used in generation and detection.

For generation, the mechanism is the absorption of optical radiation to produce heating, and thus thermal expansion. If the incident optical radiation is time-varying, or moves spatially at a fast enough velocity, then acoustic waves will be generated and propagate.

Conversely, the detection of acoustic waves involves measuring some time-varying property of the surface of the material under investigation by the reflection of applied optical radiation. The property concerned could be angular displacement of the surface, localised surface velocity normal to the surface, or localised absolute displacement normal to the surface.

1.4.5 Summary of the various techniques

There are several very good reasons why contact transducers are the technology of choice for the majority of NDT applications where ultrasound is used. They are cheap (relative to the other techniques), easy to use, compact, durable, very efficient at coupling energy from transducer to material and back again—provided a liquid couplant is used—and safe. They can generate and detect a variety of different ultrasonic waves, longitudinal and shear bulk waves, and Rayleigh and Lamb waves can be generated and detected by coupling under the correct conditions. Special contrast mechanisms can be achieved using certain geometries (in the case of the SAM).

Their main limitation, however, is due to the fact that they are contact devices, and that they require a liquid couplant. There are many circumstances when this is undesirable or impossible to accommodate, for instance with fast-moving materials on a production line, materials with difficult geometries where transducers cannot reach, or materials in a radioactive or corrosive environment. Contact transducers do not cope well with elevated temperatures, making them unsuitable for process monitoring for many situations that would benefit from ultrasonic inspection. The fact that there is a couplant loading the surface of the material may well affect the measurement, or make quantitative (rather than qualitative) measurement difficult or impossible.

EMATs solve some of these problems, although they have their own unique problems. Although they are noncontact, they must generally be located very close to the sample surface. They only work on metals, and their coupling efficiency is relatively poor compared to contact probes, especially when physical size is taken into account.

Laser generation, and particularly laser detection of ultrasound also suffers from poor sensitivity. Although it is perfectly possible to generate acoustic waves of very large amplitudes using lasers, it is likely that damage will result due to the large localised input of energy at the surface. This may be tolerated in certain circumstances (inspection of rusty pipes) although it is technically not nondestructive and in most circumstances will be unacceptable. Detection sensitivity is generally much lower than contact transducers [18] and is affected by surface reflectivity. Certain laser detection techniques are very sensitive to surface finish, and an optically-smooth surface may be required in many circumstances. Finally, laser-based systems are relatively expensive, comparatively bulky, invariably more complicated to operate and align, and have the added disadvantage that they are potentially dangerous due to the high powers of optical radiation. One of the aims of the work presented in this thesis is to overcome some of these problems, and will be discussed in detail in chapter 2.

1.5 Objectives and layout of thesis

The background to the use of ultrasound in nondestructive testing has been illustrated in this chapter. The well-known relative merits of contacting and noncontacting transducers have been mentioned.

There is much active research in all areas of transducer development, which covers the 'instruments' area of the 'three inseparable aspects' of NDT described earlier in the chapter. There is of course a corresponding amount of effort applied in research into new 'methods' and the 'intelligence' to interpret the data acquired from the variety of techniques, new and old. Occasionally, the methods or intelligence from one sort of instrumentation can be applied to another form of instrumentation. Although the main thrust of the work presented in this thesis is concerned with the development of the 'instrumentation' aspect, the development of methods of performing measurements, and interpretation of the results are given appropriate attention.

The remaining chapters concern work associated with acquiring images of surface acoustic waves using a completely noncontact—and also nondestructive—approach. The next chapter describes the research into various laser ultrasound techniques, with particular relevance to imaging systems.

Starting at chapter 3 the O-SAM instrument developed to perform SAW imaging is described, and some of its capabilities in terms of control of the acoustic wavefront are demonstrated. Chapter 4 is concerned with the data acquisition system of the system; this is fundamental to the ability of the system to acquire useful high resolution images within a genuinely useful time period.

Chapters 5 and 6 present images and other results acquired by the O-SAM, with respect to Rayleigh and Lamb waves respectively, and these are discussed. Chapter 7 describes one of the many issues that arises from the work, which is the effects of material anisotropy on the ability of the instrument to acquire images, and methods of active wavefront correction to reduce or eliminate the problem. Finally, chapter 8 describes other work relating to the O-SAM instrument that has been done that does not relate directly to imaging, such as SAW velocity measurement techniques on anisotropic crystals. Future work is also discussed, and the final chapter finishes with a summary of the work presented in this thesis.

Chapter 2

Development of laser ultrasound imaging systems

2.1 Introduction

There are several potentially good reasons why the field of laser ultrasonics might lend itself particularly well to imaging, as a method of defect detection and material/defect characterisation. Primary amongst these is the non-contacting nature of laser generation and detection of ultrasound. Difficulties arising from the use of a couplant in contacting techniques are exacerbated when the transducer is scanned at a high speed across the surface, especially if the surface is rough. The requirement to maintain a good coupling tends to restrict scanning speeds, since the amount of couplant affects the attenuation and bandwidth of the transmitted and received ultrasound. Some of the problems can be reduced in severity by immersing the material to be tested in a liquid tank, although of course this may not be possible in all cases due to contamination or degradation worries, or simple geometry—try putting a whole aircraft into an immersion tank. Meanwhile EMATs, although not suffering from the problem of couplants, have their own problems because of changes in sensitivity whilst scanning, due to changes in geometry leading to probe/surface separation variations; and they only work on current-carrying materials. Laser beams, on the other hand, are much easier to steer; they do not perturb the surface, and so it is also possible to scan the sample underneath the generation and probe beams. This, combined with the high spatial resolution obtainable from a focused laser spot, offer the potential for very rapid detailed imaging of ultrasound over a surface by scanning.

However, even a cursory review of the literature will not reveal a vast array of laser-based ultrasound imaging systems—although there are some, and the prominent and interesting techniques will be discussed later in this chapter. The lack of developed systems is presumably due to a combination of the inherent difficulties associated with the generation and detection of ultrasound using lasers, and the availability of cheap, simple, safe, and sensitive contact transducers which, for the majority of NDT applications are perfectly adequate for the job. The absence of 'industrial pull'—often the case for 'solutions looking for problems'—did not hinder development of lasers during the 1960s, and presumably this is not the case with laser-based ultrasonic imaging systems.

Although 'laser ultrasound,' as a subset of the more general 'ultrasonics' research and development area has been reasonably well established for more than a decade, there is clearly the requirement for more innovation, new techniques and instrumentation development, before industry and even the NDT research community will embrace some of the more novel methods. Amongst their concluding remarks in *Laser Ultrasonics, Techniques and Applications* [18] published in 1990, Scruby and Drain state:

> Most of the essential fundamental work on laser generation and reception of ultrasound has been carried out. There are still some gaps in our understanding... but most of these are likely to be filled over the next few years. Laser ultrasonics is therefore ready, or at least close to being ready, for exploitation by industry...

This has more or less been the case, and the author feels that laser ultrasound *imaging* systems will at some point soon be approaching a similar degree of maturity, such that the techniques can be exploited.

There are several technical difficulties to be overcome in the realisation of a viable, genuinely-useful laser ultrasound imaging system. Two problems that are prominent amongst these are the difficulty in producing waves of sufficient amplitude to be easily detected without damaging the surface of the material, and the poor sensitivity of optical detection of ultrasound compared to other techniques. The two issues are intertwined, and the resulting problem therefore is essentially one of 'signal to noise ratio.' For single, multi-point or 'line scan' (essentially b-scan) measurements, this problem can be solved by coherent digital averaging. The inability to 'acquire the data' in a sensible time frame (for imaging) is arguably the predominant reason for the apparent lack of laser-based ultrasound imaging systems, compared to laser-based ultrasound systems in general.

The majority of the contents of this chapter is therefore a review of recent innovations and methods within the field of laser ultrasonics, particularly with regard to the development of ultrasonic imaging systems. A review of the current 'state of the art' is included in section 2.4.

2.2 Optical generation of ultrasound

In the 'generic laser ultrasound system' one laser—usually pulsed and of a relatively high intensity—acts as the 'source' of the acoustic waves, whilst another laser—usually continuous wave (cw) and generally of a lower intensity—acts as the 'detector' or 'probe,' although there are of course exceptions to both of these general cases. This section provides a background to the basics of laser-based generation of ultrasound, and some of the potentially useful and interesting techniques that could be suitable for imaging systems. The optical detection of ultrasound is dealt with in section 2.3

2.2.1 Mechanisms for laser-baser ultrasound generation

The mechanisms of converting incident electromagnetic radiation from a laser into ultrasonic waves has been studied for many years. Two good overviews [19, 18] cite several of the more influential papers; [20] is an example of a paper that compares theory to practise.

The dominant mechanism for the laser generation of ultrasound is the transfer of optical energy to thermal energy at the surface of the material. When an incident optical beam from a laser strikes a sample, some of the energy is reflected, and the rest is absorbed within the material close to the surface, causing heating. Depending on the power density of the incident energy and the absorption properties of the material, the thermal energy generates the acoustic waves by either a *thermoelastic* or *ablative* mechanism.

In the thermoelastic regime, the sudden increase in thermal energy from the incident laser pulse gives rise to local expansion, which is due to a sudden rise in the temperature of the surface layer. The expansion—a local change in volume—causes thermoelastic strains, which are proportional to the incident energy. The boundary conditions at the surface demand that there be zero tractions at the surface, and so the principal stresses are parallel to the surface; although there is also expansion in the vertical direction due to the non-zero thickness of the thermal source. Figure 2.1a schematically illustrates the thermoelastic generation of ultrasound by a pulsed laser source.

In the ablation regime, the power density of the incident energy from the laser is higher than in the thermoelastic regime. The surface temperature reaches the boiling point of the material, and ablation takes place. This results in the ejection of particles from the surface, and thus the incident energy is dissipated as kinetic energy. The ablation of material from the surface produces a net recoil force and strain normal to the surface. The process is illustrated in figure 2.1b.

By applying time-varying optical energy in the form of sharp laser pulses, the material expansion as the energy increases rapidly is accompanied by rarefaction as the incident energy rapidly decreases. This is the mechanism for the generation of acoustic waves in materials by an incident

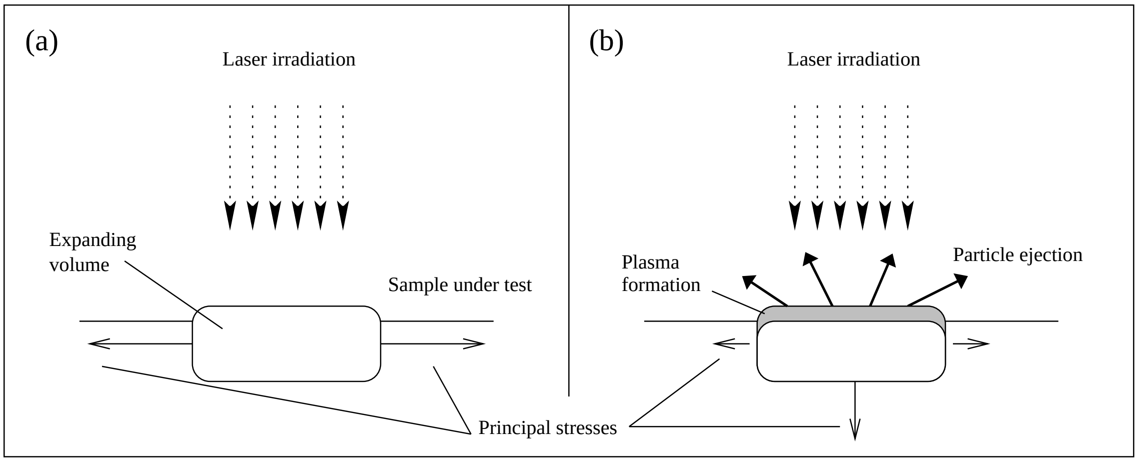

Figure 2.1: Schematic of thermoelastic (a) and ablative (b) mechanisms for the excitation of ultrasound by incident optical energy from a laser.

pulsed laser beam.

The two different regimes yield different stress and strain patterns within the material. In general, the generation of longitudinal bulk waves into the volume of the material is better suited to ablation. Both regimes generate surface waves, including efficient generation of Rayleigh waves. The ablation regime is, however, inherently not 'nondestructive,' in the sense that it damages the sample by ablating it. In some cases this may be deemed 'acceptable'—investigation of large objects such as steel plates for instance, where scorch marks from the laser may be viewed as 'cosmetic defects.' For the purpose of an all-optical acoustic wave imaging system, however, we shall consider *any* damage to the material as a result of the inspection process as unacceptable. We are therefore limited to the generation of acoustic waves in the thermoelastic regime. Surface acoustic waves are much more efficiently generated than bulk waves in the thermoelastic regime—due to the directions of the principal stresses—and thus any all-optical acoustic imaging system is likely to use SAWs as the principal mechanism for the interaction of the material properties with the acoustic waves.

The general mechanisms for the excitation of ultrasound by lasers have now been described. Several techniques will now be discussed that increase the potential maximum amplitude of the generated surface acoustic waves, without ablating the sample. This is particularly relevant to the realisation of a non-contact, completely nondestructive SAW imaging system, as discussed in the introduction to this chapter.

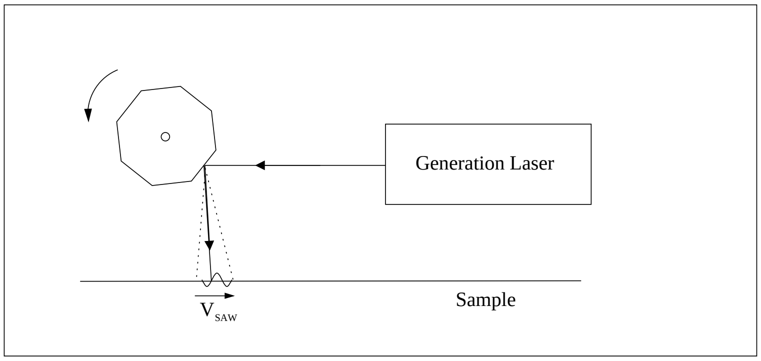

Figure 2.2: Illustration of the phase velocity scanning technique for SAW generation.

2.2.2 Phase velocity scanning

This technique—shown schematically in figure 2.2—involves scanning a laser beam over the surface of the material at a velocity equal to the phase velocity of the surface waves that are to be generated, and is described in detail in [21] and [22]. The laser beam acts as the excitation source of the SAWs and the amplitude of the SAWs increases linearly with scanning length [21], leading to an overall increase in the signal to noise ratio (SNR) achievable without damaging the sample. In addition, a relatively inexpensive continuous wave or long pulse length laser may be used as the source, the SAW frequency is continuously tunable (assuming the scanning velocity is also continuously tunable) with the upper frequency limit determined by the width of the focused laser beam, and the directionality of the SAWs is greatly increased—they essentially only propagate in the scanning direction.

2.2.3 Laser induced grating

This technique—illustrated in figure 2.3—involves imaging a set of optical fringes onto the sample, with the distance between the fringes equal to the wavelength of the SAWs that are to be generated. This can be achieved either by combining two time-coincident laser pulses at the sample surface [23, 24, 25], or with an optical diffraction grating [26]. Each of the optical fringes acts as an excitation source of SAWs, and counter-propagating surface waves are generated from each fringe if the laser beam is pulsed. If the spacing of the fringes is adjusted such that the inter-fringe spacing is equal to the SAW wavelength for the temporal frequency of generation—which is determined by the inter-pulse time period—then the SAWs will add up in phase.

The generation technique is inherently narrowband, and the frequency spectrum of the generated SAWs is related to the number of generation lines, as well as the number of pulses applied to

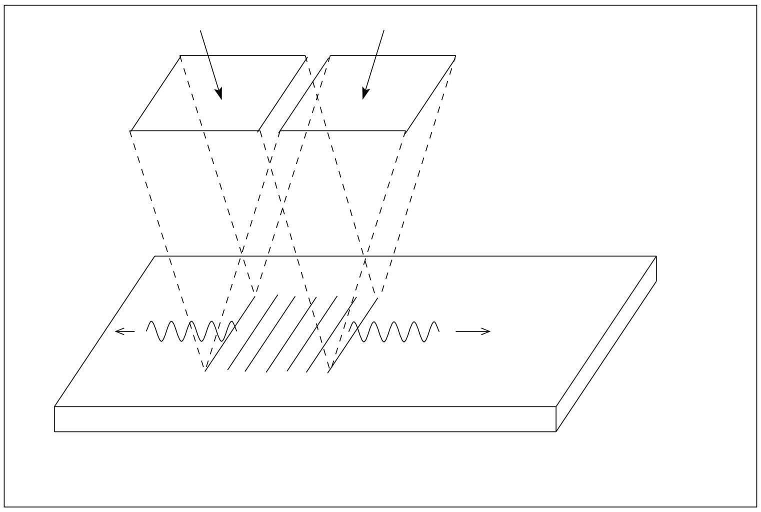

Figure 2.3: Illustration of the laser induced grating technique for SAW generation.

the material. The technique allows the overall amount of power that may be applied to the material to be increased by a factor equal to the number of fringes used, since the power is distributed over a much larger area, and the limit as to how much power may be applied to the sample is more concerned with the power density, rather than the total power applied.

2.2.4 Scanning interference fringes

The scanning interference fringes technique [27, 28] is a combination of the scanning phase velocity technique and the laser induced grating technique, and is shown schematically in figure 2.4. This retains the advantages of the scanning phase velocity approach (long laser pulse, excellent directivity, waves propagating in one direction only) with the features of the laser induced grating method (distribution of optical power, narrowband SAWs).

2.2.5 Focusing of surface acoustic waves

The principle of focusing light to a point to increase the intensity was used to good effect in the defence of Syracuse against the invading Roman fleet in the third century BC. At the direction of Archimedes, sunlight was focused onto the wooden hulls of the enemy ships, setting them alight. Although not quite as visually impressive, the amplitude of SAWs can be increased dramatically by focusing them to a point.

The technique has been developed in the area of contact ultrasonic transducers [29, 30, 31] to improve the signal to noise ratio of the detected acoustic waves and to give improved lateral

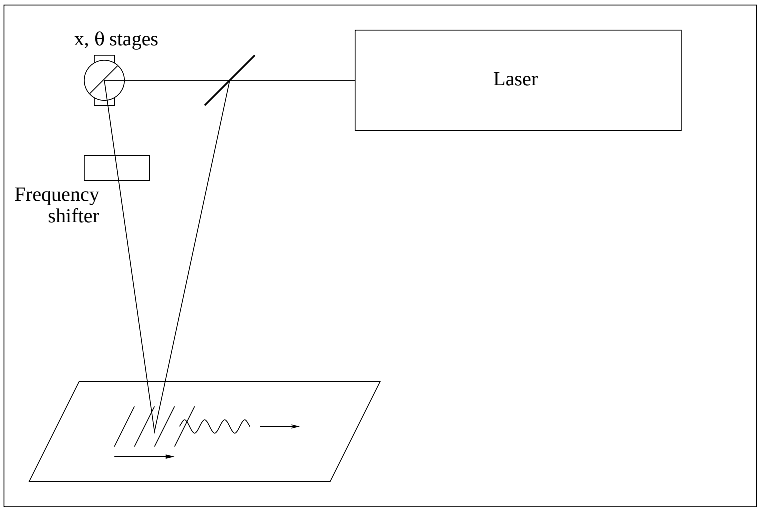

Figure 2.4: Illustration of the scanning interference fringes technique for SAW generation.

resolution. A noncontact version was developed around 1985 [32]. The method was based around an axicon lens, which focused a laser beam into an optical ring on the surface of the sample.

More recently, a customised zone plate was used to focus an arc onto the material [3]. This is discussed in detail in section 3.2, since this method is used (at least initially) in the O-SAM instrument.

Another method of generating an acoustic focus is to use a phased array of transducers. Again, these are quite common in the area of contact ultrasonic transducers [9] and a 16-element laser-generated phased array is reported in [33]. As with contact transducer phased arrays, focusing is achieved by introducing an appropriate time delay between each laser pulse. Amplitude gains at the focus of more than 10 were reported.

2.3 Optical detection of ultrasound

The means of detecting acoustic waves in materials by a laser beam is a completely different mechanism to the excitation of acoustic waves. Since the majority of materials to be inspected are opaque, some time-varying property of the surface must be measured.

There are many useful overviews of optical techniques for the detection of acoustic waves in solids—several chapters of *Laser Ultrasonics* [18] are devoted to the subject; there is a good review of the different techniques in the *Physical Acoustics* series by J. W. Wagner [34]; an oft-cited overview and comparison of different techniques was written by J.-P. Monchalin [35] in 1986; and a more recent paper [11] published in 2000 compares several different non-contact techniques for

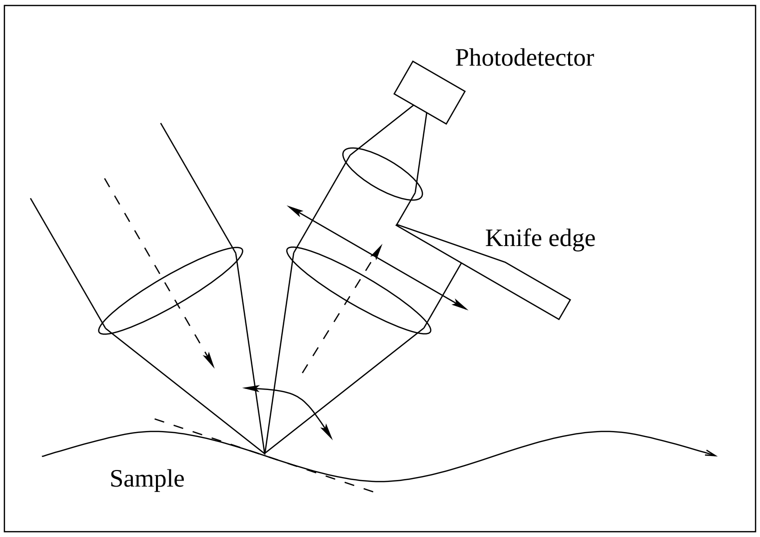

Figure 2.5: Schematic of the 'knife-edge' technique of SAW detection.

the detection of Rayleigh waves, including photo-emf detectors and EMATs. Each of the reviews cite several papers describing each of the methods, as well as the earlier reviews. Only the more popular general techniques, and those with particular relevance to imaging will be briefly described here to give context.

2.3.1 Knife-edge detection

Of the non-interferometric techniques—which includes the surface-grating technique and a technique based on changes of reflectivity [35]—the *knife-edge* or *optical beam deflection* technique is by far the most commonly used. The principle of its operation is illustrated schematically in figure 2.5.

An incident laser beam is focused to a spot of diameter less than half the SAW wavelength. As the SAWs propagate under the spot, the reflected beam is deflected by an amount equal to twice the angle of the surface underneath the spot. The amount of deflection is therefore proportional to the SAW amplitude, and also the SAW frequency. The angular displacement is translated to a spatial displacement by a lens, and this is then converted to a change in intensity by putting a 'knife-edge' in the path of the beam.

2.3.2 Displacement interferometry

The Michelson interferometer, shown in figure 2.6 is an example of an interferometric technique used to measure changes in the surface displacement; other techniques include the Mach-Zehnder

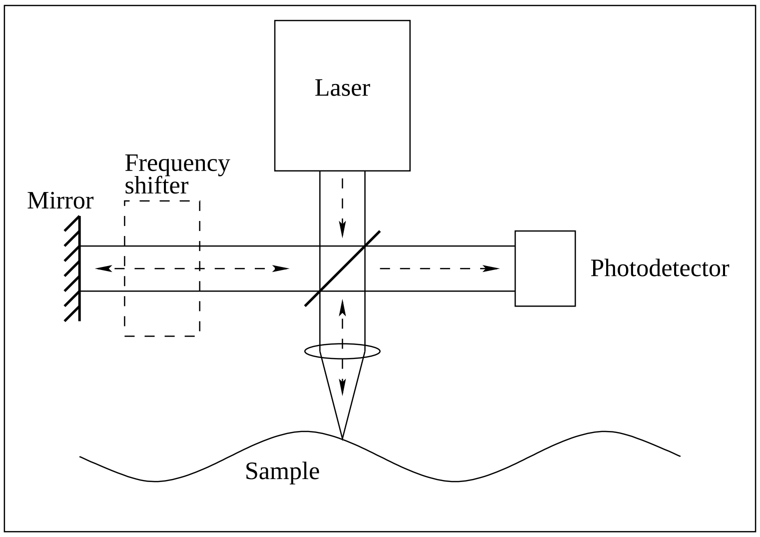

Figure 2.6: Schematic of a Michelson interferometer, an example of an optical heterodyning technique of SAW detection.

and Sagnac interferometers. They all work by 'beating' a wave reflected by the material surface with a reference wave. A change in displacement distance will cause the path length of one of the arms of the Michelson to increase or decrease. When the two beams are interfered with each other, this change in path length corresponds to a change in relative phase, and as such a change in amplitude. The amplitude of the output is therefore proportional to the displacement of the surface. If an optical frequency shifter is inserted into one of the arms of the interferometer [36], then the signal from the detector also includes a carrier at the shifting frequency and two sidebands. This 'heterodyne' (or sometimes called 'super-heterodyne') probe has advantages over the 'homodyne' case without the frequency shifter, since the displacement information contained within the phase information is now modulated, and thus the calibration of the amplitude can be performed by comparing the sideband amplitudes with that of the carrier.

2.3.3 Velocity interferometry

The *confocal Fabry-Pérot interferometer* (CFPI) is a very popular detection method, and for good reasons. The principle, as with all interferometers, is to interfere one wave with another, and in the case of the CFPI it is the interference between the optical beam reflected from the surface of the material, and a time-delayed wave from the same place. The system is based on the Doppler shift of the frequency of the reflected light from the material surface. As SAWs propagate on the material surface, local changes in the velocity normal to the surface cause a small shift in the reflect

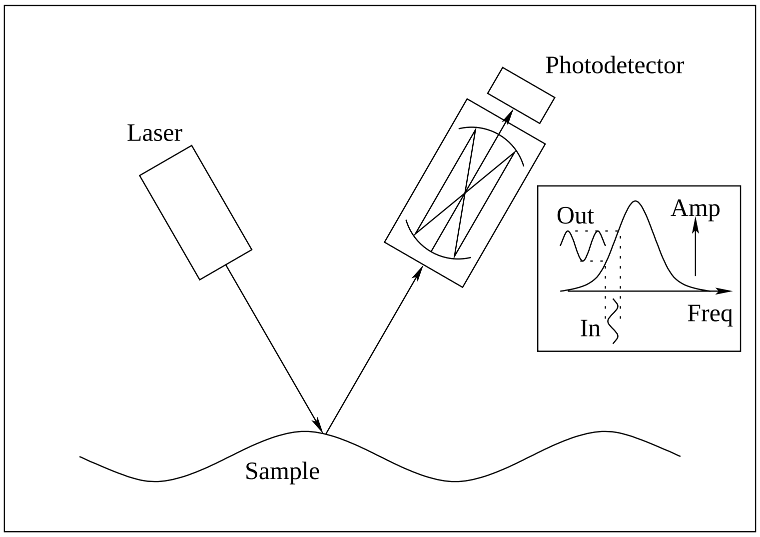

Figure 2.7: Schematic of a Fabry-Pérot interferometer, a technique based on the Doppler shift in frequency of the reflected wave.

light frequency, and hence wavelength. This change in wavelength is detected by a change in the transmittance or reflectance of light through a resonant cavity. Figure 2.7 illustrates the effect.

The partially-reflecting mirrors at either end of the cavity allow for interference between multiply-reflected wavefronts within the cavity. Maximum transmission occurs when the cavity length is an integer multiple of the wavelength of light within the cavity. By 'detuning' the cavity length to a point on the maximum 'slope' of transmittance, any change in the wavelength of light entering the cavity will result in a change in the transmittance (and also reflectance). The output of the Fabry-Pérot is therefore proportional to the velocity of the material surface.

2.3.4 Photo-emf detection

This is a relatively new technique, first described in 1986 by Petrov *et al.* [37]. The article is in Russian and has not been read by the author, although its use in laser-based ultrasound inspection is very well described in [11], [38] and [39]. The technique is realised by combining a reference beam from the detection laser with a probe beam from the material surface—which may consist of many speckles from a rough surface finish—on the surface of a non-steady-state photo-emf crystal. A periodic intensity pattern is formed, since the two beams are applied at a small angle to each other. Photo-generated charge carriers are produced that diffuse away from areas of high optical intensity, and become trapped in the low intensity regions. In the absence of any net change in the optical pattern no net current flows. However, a rapidly-moving pattern from the surface that ultrasound is propagating on causes an induced emf to be produced across the material that is dependent on the relative fringe displacement from the stationary pattern. This occurs if the movement is faster than the carrier recombination time within the crystal; the crystal therefore acts as a high pass filter, with decreasing sensitivity at low frequencies, typically of the order of 20–100kHz.

2.3.5 Strengths and weaknesses of the general detection techniques

Most of the general techniques described (knife-edge, Michelson, CFPI) have approximately the same sensitivity [35] of around 10^{-6} Å$(\mathrm{W/Hz})^{1/2}$ for a mirror-like surface; the sensitivity of a recent commercial photo-emf system is quoted as 3.9×10^{-5} Å$(\mathrm{W/Hz})^{1/2}$ [11].

The knife-edge technique is simple, robust, very inexpensive to implement, and is insensitive to vibrations. The use of an optical fibre bundle [40] to collect the light from the material surface, combined with fibre delivery of the incident beam, permits a very convenient optical setup. It is, however, very sensitive to the effects of rough surfaces, and requires a smooth, if not polished surface finish to be used effectively.

The Michelson interferometer, and others whose output is proportional to the displacement of the material surface, have a broader detection bandwidth than the knife-edge technique, and are all able to measure the amplitude of the SAWs directly and absolutely. Problems again arise with non-flat samples. There are additional complications as regards sensitivity to vibrations etc, with active feedback in the form of the reference mirror being mounted on a piezoelectric transducer to automatically adjust the path length in the case of homodyne probes, and heterodyne techniques require complex demodulation circuitry to detect the displacement information, and the overall bandwidth is limited by the frequency shifter.

The Fabry-Pérot interferometer, meanwhile, also requires active stabilisation of the cavity length, as well as a single-frequency laser source, but its strength lies in its ability to work effectively on samples with rough surfaces, since it is inherently able to interfere many speckles reflected from the material surface. The bandwidth can be relatively small compared to the other two general techniques, although this can be extended by 'optical sideband stripping,' a technique that involves using the CFPI in reflection mode [41].

The non-steady-state photo-emf detection technique is reasonably simple to implement, and works well with rough surfaces. Current systems are expensive, although this may well change in the future. The low frequency cut-off point of 20–100kHz could potentially allow the device to be used for scanning—some success has been achieved at detecting Lamb waves in paper samples moving at production speeds [42]. At present, however, the upper measurable frequency of photo-emf systems is limited to 'near 100MHz' [39] due to bulk-carrier recombination, which severely limits its use for high resolution imaging.

2.4 Review of laser-based ultrasonic imaging systems

This section describes the current 'state of the art' of laser ultrasound imaging systems. We shall begin with a very brief look at hybrid systems—those that consist of a contact transducer (usually the source) and a noncontact transducer—since the history of completely noncontact systems is intertwined with their hybrid counterparts.

2.4.1 Hybrid contact/noncontact systems

The early impetus for noncontact ultrasonic imaging systems was a desire to visualise the propagation of surface waves, and their interaction with materials and defects. The earliest paper to show this is by Adler *et al* [43], published in 1968. The system used a mirror-scanned continuous wave laser beam, scanning at television rates. A knife-edge detector was used to measure the angular displacement of the 8MHz surface waves. By mixing the output of the photodiode with a reference from the transducer drive—a contact device—an output proportional to the phase of the SAWs relative to the excitation is acquired for each scanning point. The television screen thus shows a stationary pattern of fringes representing the SAWs. Several impressive images of SAWs propagating are presented in this excellent early paper on the subject. A paper by Alers *et al* [44] in 1973 expanded on the technique to investigate the application to locating defects and studying the intrinsic properties of materials.

The scanning laser acoustic microscope (SLAM) [14] is another example of a hybrid system that uses a contact transducer as the source and a laser as the detector of ultrasound. Bulk waves are propagated through a material—solid or liquid—and are detected and imaged on the far surface by a raster-scanned laser. Sensitivity to bulk waves may then be achieved by using an inclined incidence, and the angular displacement measured with the knife-edge technique. The SLAM was successfully used in 1988 by Robbins *et al* [45] to image the interaction of SAWs generated by an interdigital transducer and thin surface films. The instrument was capable of vector contrast imaging, as well as being able to acquire an optical image, and so amplitude, phase and 'wavenumber reconstruction' (corresponding to velocity) images at 100MHz could be acquired and are shown in the paper, although the time taken to acquire the 256×256 pixel images is not stated. A SLAM was also used in 1997 by Anisimkin *et al* [46] to visualise the effectiveness of coupling bulk waves into SAWs using a groove grating fabricated on the surface of a non-piezoelectric material.

A phase-shifted Sagnac interferometer capable of imaging Lamb wave interaction with an artificial notch was demonstrated in a paper by Fomitchov *et al* [47] in 1997. The source of the Lamb waves was a piezoelectric transducer producing 5MHz tone bursts. The paper is mostly concerned with describing the Sagnac interferometer, and does not mention the resolution of the amplitude image presented, or how long it took to acquire.

A system is described by Hamilton and O'Donnell [48] in 1998 which measures the isonification of a wire sample in an immersion tank by reflecting an incident laser beam from a reflective pellicle and using a confocal Fabry-Pérot interferometer to measure the displacement. The paper is entitled "High frequency ultrasound imaging with optical arrays" although their interpretation of what an optical array consists of differs somewhat from the author's, since they simply acquire an acoustic image by scanning the detection point over the pellicle. The images presented are actually b-scans, although they did acquire data over a 2-D 'array.' Each point in the scan required coherent averaging to be performed; it is not stated how many averages, although other results presented in the paper are averaged between 200 and 1000 times. The data acquisition time for the 2-D 'array' was approximately 3 hours and 50 minutes.

Another paper by Hamilton *et al* [49] published in 2000 describes an interesting idea about how a plane-mirror etalon may be used as the basis of a optical array imaging system. No actual imaging results are presented, however, or details of the 'array.'

One of the few instances where the ultrasound source is a laser and the receiver is a contact device is described by Wickramasinghe *et al* [50] in 1978, in a transmission scanning acoustic microscope configuration. A Q-switched mode locked Nd:YAG laser was used to generate 840MHz waves, which were detected by a focused contact acoustic lens on the other side of the sample.

Finally, a very novel technique for optical imaging of the acoustic field has already been mentioned in section 1.3. Kolomenskii and Maznev [15] in 1993 imaged the acoustic fields on anisotropic crystals by examining the patterns made in dust by the propagation of surface waves. The source of the ultrasound was a Q-switched Nd:YAG laser. The samples were positioned vertically, and when the SAWs propagated along the surface with sufficient amplitude, the 'dust' (Al_2O_3 particles) was removed. Photographs were then taken after the experiment, and areas of high SAW amplitude are darker, due to there being less dust.

2.4.2 Early laser ultrasound imaging

Some of the earliest attempts at completely noncontact ultrasonic imaging were performed by Dewhurst *et al* in the summer of 1990, and were published a little later in 1991 [51] and 1992 [52]. A Q-switched Nd:YAG laser was used as the ultrasound source, and operated in the ablation regime. A 1W argon-ion laser was used as the detection probe, on the far side of an aluminium sample with a 4mm diameter hole drilled into the sidewall to imitate a defect. A Fabry-Pérot interferometer was used to detect the ultrasound. Transmission amplitude images of 19×15 pixels, consisting of 4 different grey levels, were produced. It is not stated how long the images took to be acquired, although it is stated that a 5 second delay was required between the measurement of the ultrasound signal and the dc light level for each pixel. Thresholded (binary) images of a circular laminar saw-cut were also presented in [52].

In some ways, the development of laser-based imaging systems has not significantly moved on in the intervening years. A paper published in December 1999 by di Scalea and Green, Jr. [53] presents a 21×21 pixel time of flight c-scan, with slightly more grey levels than the Dewhurst papers. Again, the source was a Q-switched Nd:YAG laser operating 'slightly in the ablation regime.' A Fabry-Pérot interferometer, powered by a 2W diode-pumped frequency-doubled Nd:YAG laser was able to detect a 6mm deep recess shaped as an "H" engineered into the back surface of a 12mm aluminium plate. The time required for each pixel was 0.5 seconds.

2.4.3 Commercial-scale systems

There do not appear to be very many commercial laser ultrasound imaging systems in existence, although the ones that are seem to have fairly reasonable specifications—at a price.

Perhaps the most well-known of these systems is the "laser ultrasonic inspection system" (LUIS) [54, 55] built by UltraOptec of Quebec. The company sells a range of systems, and they are typically used to inspect aircraft. The ultrasound source is a CO_2 laser producing 130ns long pulses of light at a wavelength of 10.6μm, with each pulse having up to 200mJ of energy. A long pulse Nd:YAG laser is used for detection, which produces 500W of power in the 50ms the system detects ultrasonic echoes. A Fabry-Pérot interferometer is used for detection. The system can scan up to 100 points a second, at scan steps of '0.2, 0.1 or 0.05 inch'—the maximum lateral resolution is therefore 1.3mm. A total of five PCs are used to run the lasers, control the scanning mirror, digitise the signal from the detector, and control the gantry robot which carries the generation laser and fibre-delivery detection system around. A Silicon Graphics workstation is used to set up the scans, process and archive the data. Scans can be up to 6 feet by 6 feet in size, and overall the system can scan structures as large as aircraft wings. The system is, quite frankly, a monster.

A very good review of a commercial-like system was presented in 1994 by McKie and Addison [56]. The system was able to acquire 100×100 pixel images—time of flight and pulse echo—in around five minutes using rapid scanning translation stages and a fast buffered digital acquisition system, which was able to store digital waveforms from an entire scan-line, and transfer the data to the host on fly-back. Two types of generation laser were investigated, and the relative merits of Nd:YAG (1.064μm wavelength) and CO_2 (10.6μm wavelength) lasers were examined. For rapid inspection of uncoated graphite/epoxy specimens the CO_2 laser source was preferred since the longer pulse length—compared to the Nd:YAG laser of the same energy per pulse— prevented thermal damage. A 470mW argon-ion laser was used in the Fabry-Pérot detection system. Again, lateral resolution was a very modest 1mm, but this was deemed more than adequate for the inspection process. Scanning acquisition rates of 33 points per second is a very good achievement, especially for the time the paper was published.

2.4.4 Synthetic aperture focusing technique

A different approach to laser-based c-scan imaging is the synthetic aperture focusing technique, or SAFT. This technique [57, 58] involves acquiring the complete digital waveform for each scanning point, and then post-processing 2-D arrays of adjacent points. The principal idea is that ultrasound reflected from a defect will radiate and be detected at a number of points, and arrive with a phase associated with the distance from the defect to the detector. Thus by coherent summation of a 2-D array of points for a given 'target,' the system can determine whether there is a reflection (corresponding to a defect) or not.

Although the entire waveform for each scan position is required—making the system inherently slow—post-processing of the data leads to an increase in SNR for defect detection by $\sqrt{N}$ where N is the number of points in the 2-D array [58], and so much lower laser powers may be used to acquire images of good quality.

2.4.5 Pulsed probe systems

Sugawara, Wright *et al* have very recently presented several papers [59, 60, 61, 62] in which acoustic waves, thermoelastically excited at a point by a $\approx$1ps pulse laser with a repetition rate of 80MHz (achieved by mode locking), are imaged as they propagate on the surface of isotropic and anisotropic crystals, or a substrate with a thin film coating. The detection technique uses a common path Sagnac interferometer [63]. Two very short ($\approx$ 200fs) pulses are applied to the sample at the same position, separated in time by around 500ps. The optical phase difference between the two pulses is measured by the interferometer. The output signal is therefore proportional to the out-of-plane velocity of the surface, and the interferometer is potentially sensitive to frequencies up to 1GHz. Many images of SAW wavefronts are presented in the papers, and the acquisition speed is impressive—a 100×100 pixel image can be acquired in 4 minutes, or a 400×400 pixel image in 22 minutes [62], with good spatial and temporal resolution. At these speeds it is possible to acquire many images and, by varying the time between the excitation pulse and the probe pulses, it is possible to view the propagation of surface waves in two dimensions. The laser and detection technology are relatively complicated and expensive compared to the usual methods of ultrasound generation and detection, but the results are undeniably impressive.

A pulsed probe laser has also been used to good effect in another very recent paper [64] by Glorieux *et al.* A grating interferometer is used to detect surface waves on aluminium coated glass. A grating spatially separates two incident probe beams, which after reflection from the surface are recombined by the grating, leading to a peak sensitivity to surface waves of a wavelength of twice the probe separation. The interferometer can be used in the full field configuration, by replacing the photodetector with a CCD (charge coupled device), and by synchronising the probe pulse laser with the excitation pulse laser and with the CCD refresh rate. The SNR is extremely poor as can be expected, with averaging required "during a few minutes per snapshot" (ie for each pump-probe separation time). Nevertheless, an image is presented where it is *just* about possible to discern two SAW wavefronts emanating from a scorch mark where they were excited by the pump laser.

2.4.6 The system developed at the University of Nottingham

In 1997, the laser ultrasound system from which the O-SAM instrument has been developed was undergoing changes in its configuration that would eventually allow rapid imaging to be performed. In March 1998 the images shown in figure 2.8 were acquired by the author—see also [5]. They are amplitude c-scan images of Rayleigh waves at 82MHz and 164MHz reflecting from cracks emanating from a Vickers hardness test indentation. The ultrasound source was a single arc of radius 10mm, imaged onto the sample using a computer generated hologram. The data was acquired using a digital storage oscilloscope, and 20 coherent averages were acquired at each point. The image consists of 151×71 pixels—each pixel represents an area of size $5 \times 5\mu$m, the area of the image is $750 \times 350\mu$m. With over ten thousand pixels, and each pixel taking approximately a second to acquire, the images took around 3 hours to obtain. The results in themselves were reasonably impressive for the time, but there was clearly much to be done to realise a genuinely useful laser ultrasound imaging system.

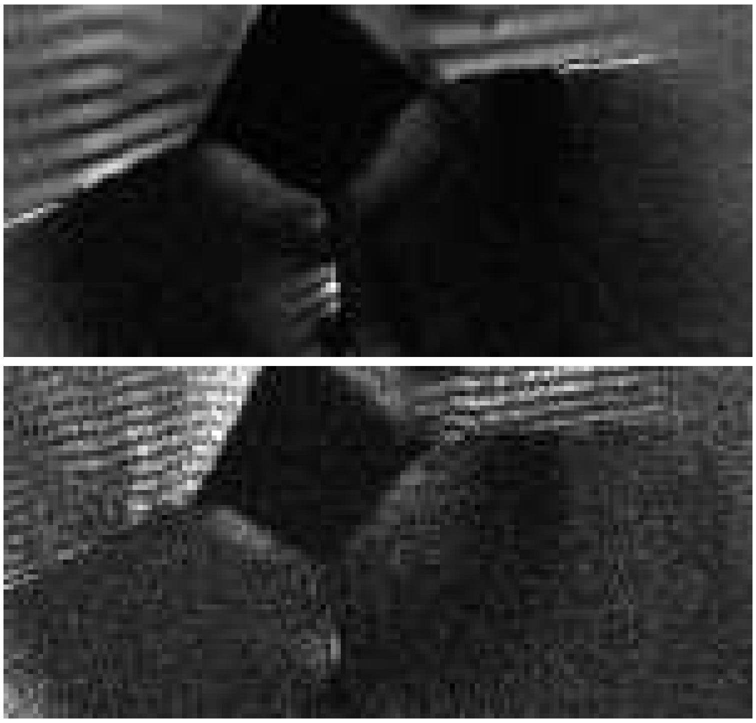

Figure 2.8: Rayleigh wave amplitude c-scans acquired in March 1998 at 82MHz (top) and 164MHz (bottom) of a Vickers hardness test indentation on silicon nitride. The image size is $750 \times 350\mu$m, the pixel size is $5 \times 5\mu$m. The images took around three hours to obtain.

Chapter 3

Instrumentation

3.1 Introduction

This chapter is concerned with the theoretical and practical requirements of the all-optical surface acoustic wave imaging system, including: general strategy for attaining surface waves of sufficient amplitude that may be detected optically without damaging the sample surface; control of the optical distribution of the generation source using either a computer generated hologram or a spatial light modulator; frequency control of the generated ultrasound; control of the spatial distribution of the generated ultrasound; optical detection of the surface acoustic waves; and instrument alignment issues. Although the information contained within this chapter is predominantly of a practical nature, it is important to describe it in detail, since the success of the instrument depends on the factors herein.

3.1.1 System overview

In line with other laser ultrasound systems described in the previous chapter, the all-optical scanning acoustic microscope (O-SAM) consists of two lasers—one for ultrasound generation, the other for detection of the ultrasound—some associated optics and mechanical stages (for adjustment and scanning) and some electronics and data acquisition equipment. A block diagram of the key parts of the system is shown in figure 3.1.

The energy from the pulsed generation laser is applied to the surface of the sample via optics that both focus the light onto the sample and control its spatial distribution. This control of the generation profile is necessary to obtain the large surface acoustic wave amplitudes required. The sample is mounted on automated mechanical stages to enable the formation of scanned images. The surface acoustic waves propagate on the sample surface, and are then detected by focusing

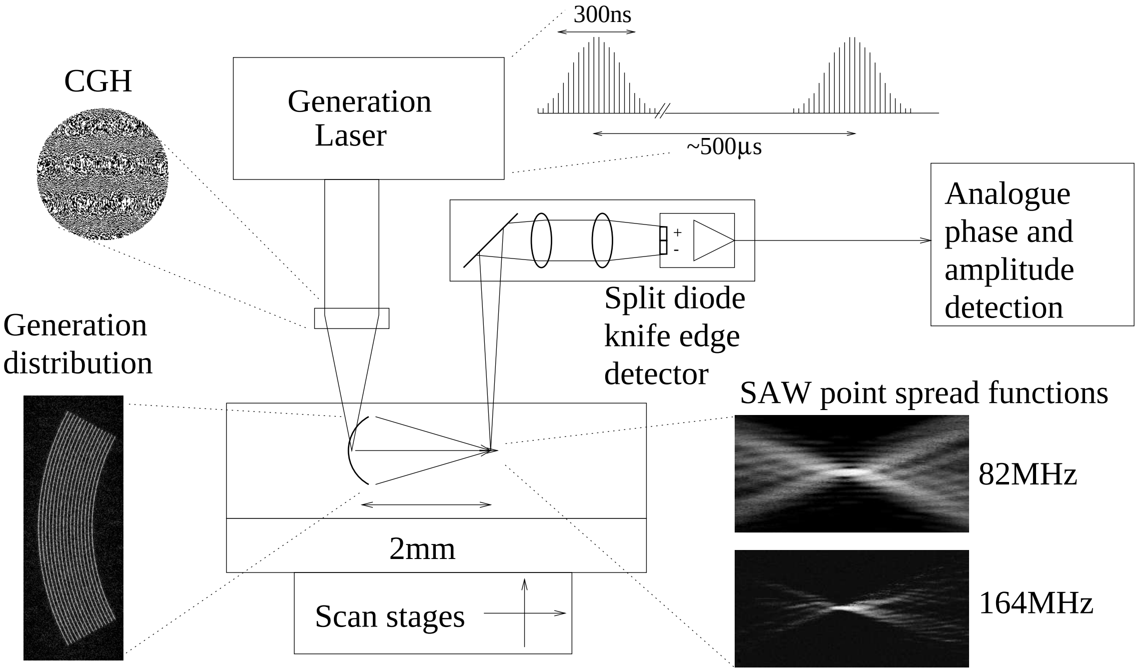

Figure 3.1: Block diagram of the all-optical scanning acoustic microscope (O-SAM).

a continuous wave laser onto the sample surface and measuring the angular deflection. Analogue amplitude and phase detection electronics then amplify and convert the signals from the detectors into a form suitable for acquisition by a host computer.

3.2 Control of ultrasound source

Section 2.2 described in general how surface acoustic waves may be generated optically; section 2.3 described how they may be detected optically. One of the inherent problems with an all-optical ultrasound system is that the optical detection sensitivity is generally much lower than for contact techniques.

Although it is possible to apply more and more optical power to generate ultrasound of greater amplitude, this tends to result in the sample being damaged. Whilst this may not be a problem on some materials (it may even be beneficial on, say, rusting pipes) in many circumstances this would be considered unacceptable.

In addition, the signal-to-noise ratio (SNR) of an all-optical system operating in the thermoelastic regime is proportional to $P_{gen}\sqrt{P_{det}}$, where P_{gen} is the optical power used to generate the surface wave, and P_{det} is the optical power used to detect it [18, 35]. This implies that there is much more to gain from increasing the generation power than from increasing the detection power.

The signal to noise ratio is particularly relevant if you wish to perform imaging with an all-optical system, where high speed data acquisition is required.

Much of the success of this instrument, as far as speed of operation and ease of use is concerned, is due to the technique of precisely controlling the optical distribution of the ultrasound source. This is done so that as much of the available optical power as possible is used efficiently and effectively to obtain surface acoustic waves of sufficiently large amplitudes for optical detection, but without causing any damage to the sample under investigation. There are two mechanisms by which this is achieved: by *focusing* the ultrasound to a finite region, producing amplitude gain at the focus; and by *distributing the power* of the incident laser energy. These two complementary techniques are combined to greatly enhance the maximum available surface acoustic wave amplitude, which is very important if a practical all-optical imaging system is to be realised.

This precise control of the optical distribution allows for precise control of the resulting ultrasound propagation in terms of modes, frequency content and spatial distribution. This may be used to attain other goals, other than the maximum possible SAW amplitude available. Practical examples of some of the possible optical distributions are presented later in this chapter.

Two methods have been used during the development of the O-SAM to control the optical distribution of the ultrasound generation profile: computer generated holograms, and a spatial light modulator. Before either of these two methods is described, it is useful to detail the characteristics of the generating laser used, and the general approach to the design of the generating optical distribution.

3.2.1 Generating laser source

The laser used to generate the ultrasound is a Q-switched, mode locked Nd:YAG laser with a wavelength of 1064nm. In this mode of operation, the mean output power of the laser is around 2W, and the peak energy from a pulse is of the order of 0.3MJ.

The Q-switch repetition rate may be adjusted from a few hundred hertz up to 10kHz, although with increasing frequency the peak power levels off and is eventually reduced. The Q-switch repetition rate is important, as this determines the upper measurement speed limit of the instrument. The Q-switch pulse duration is typically approximately 300ns, which implies that the pulse output has a bandwidth of approximately 3–4MHz.

The mode lock frequency is determined by the cavity length, and in the case of the laser used in this instrument, is 82MHz. This represents the fundamental frequency of the laser output, which is also rich in harmonics due to the very short duration of each individual output pulse, which is of the order of 200ps.

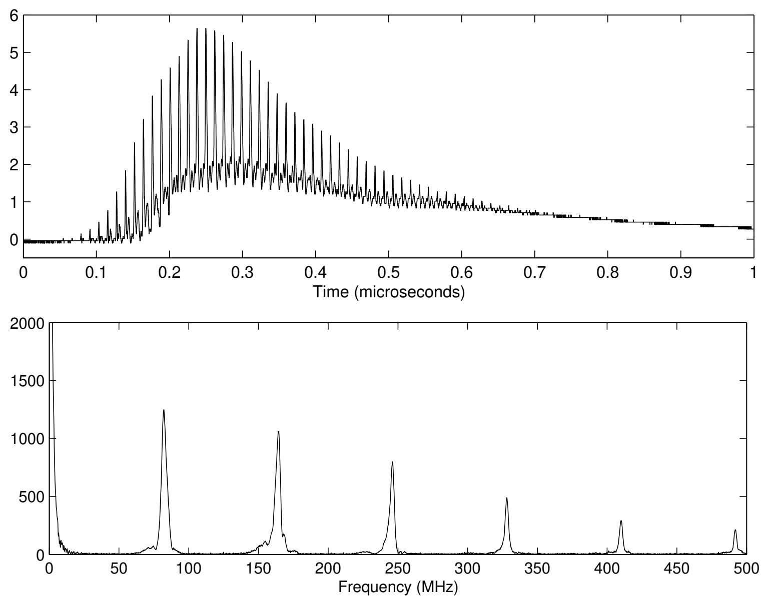

Figure 3.2: **(a)** The output of the generating laser source, as measured by a high speed photodiode. **(b)** Fourier transform of **(a)**.

Figure 3.2a shows the output of a high frequency photodiode on which a small proportion of the laser energy is allowed to fall. Figure 3.2b is the Fourier transform of the signal, confirming the bandwidth of the output due to the Q-switch duration. Each harmonic should in theory have approximately the same amplitude as the fundamental frequency, up to the GHz region; however, the photodiode and associated electronics have a finite bandwidth, and for this reason the amplitude trails off with frequency.

3.2.2 Focusing the surface acoustic waves

Here we will deal with the mechanism of focusing the surface acoustic waves, initially using a continuous arc.

Amplitude gain at focus

The amplitude of the surface acoustic waves may be significantly increased by focusing them to a finite region. For surface waves, the increase in acoustic amplitude by focusing the waves compared with the amplitude generated with a line source is:

$$A_{focus} = \frac{x}{\sqrt{2\lambda_{SAW}F}} \tag{3.1}$$

where x is the length of the line source or arc, λ is the acoustic wavelength, F is the working distance between the line or arc and the detection point (the focal length of the arc) and $x >> \lambda$ and $x > \sqrt{2\lambda_{SAW}F}$ so that the arc may be regarded as a true focusing element. This is derived using simple scalar diffraction theory [65] and conservation of energy. To give a practical example, a single arc with a focal length of 10mm and a half-angle of 30° would produce ultrasound of 12 times the amplitude of a single line source of the same length, for a SAW wavelength of 35μm (82MHz on aluminium).

Effect of varying the focal length and numerical aperture

By choosing different optical distributions, arcs of different focal lengths may be illuminated on the sample surface. The length of the arc may also be adjusted, which leads to a change in the numerical aperture (NA) for a given focal length—here the NA is defined as the sine of the half-angle of the arc. As with most design issues, it is usually possible to find a compromise between all the conflicting requirements for most situations. The choice of focal length of the arc—or arcs—is determined by a combination of many factors, including:

- Attenuation in the sample. If there is acute attenuation, then the shorter the focal length the better, unless it is the attenuation you wish to measure.

- Acoustic aberration. This issue is dealt with in detail in chapter 7. Briefly, if the sample is anisotropic, then the wavefront gets distorted as it propagates from source to detection point, resulting in a displaced or broken point spread function; hence the shorter the propagation distance (focal length) then the less the wavefront gets distorted.

- Instrument geometry. The chosen method of focusing the detection laser onto the sample surface and collecting the reflected light may determine whether a short or long focal length is the most convenient.

- Localised heating. Although the energy applied to the sample should be less than that required to cause permanent damage, it usually causes local heating and distortion of the sample. If the focal length is too short, this can perturb the detection region.

- Damage threshold. A shorter focal length implies a shorter arc length for any given numerical aperture. This implies a higher energy density, and hence less power may be applied to the sample, implying a lower surface wave amplitude. This issue is addressed in section 3.2.3.

- Sample features and geometry. For some samples, due to their size, a small propagation distance is desirable, or even necessary. In other cases, longer propagation distances may be desirable, for instance when changes in velocity need to be measured, or different SAW modes must be allowed to propagate. This is covered in detail in section 6.3.1.

- Power density at focus. The higher the numerical aperture, the higher the SAW amplitude gain at the focus, since the width of the point spread function (PSF) is diffraction limited.
- Ease of alignment. The higher the NA, the narrower the PSF, and the more difficult it is to find the focus with the detection laser. Alignment issues are covered in depth in section 3.6.

Aerial point spread functions

On an 'ideal' sample, the region around the focus of the surface acoustic waves from an arc source should represent a diffraction limited aerial point spread function, with a width of $\frac{\lambda}{NA}$.

The image on the left of figure 3.3 is an amplitude aerial point spread function (PSF) from a set of concentric arc sources, imaged onto a sample of aluminium-coated glass using a spatial light modulator (SLM) to control the optical distribution—see section 3.4. In this image, the amplitude of the 82MHz component of the surface waves is represented by the brightness of each pixel. The right hand image of figure 3.3 is a phase aerial PSF, where the brightness of each pixel is determined by the phase of the surface wave with respect to the excitation source.

The images were obtained by keeping the sample and source arcs stationary and scanning the detection point. Just one pulse 'event' was used to obtain the amplitude and phase information at each location. The image consists of 250000 points and, with high speed mechanical stages, could potentially be acquired in just over four minutes.

3.2.3 Distributing the optical power

Focusing the surface acoustic waves significantly increases their amplitude at the detection point, and further increases in amplitude may be obtained by carefully distributing the optical power over a large area on the sample surface, as described in section 2.2.3. This is because power density, rather than total power applied, dictates how much optical power may be applied to the sample before damage occurs.

Thus if the power may be distributed in such a way that it contributes to the amplitude of the surface acoustic waves, then an increase in optical power applied to the sample—and a resulting increase in SAW amplitude—is possible without causing damage. This may be achieved by using multiple lines—or arcs—as sources of the acoustic waves.

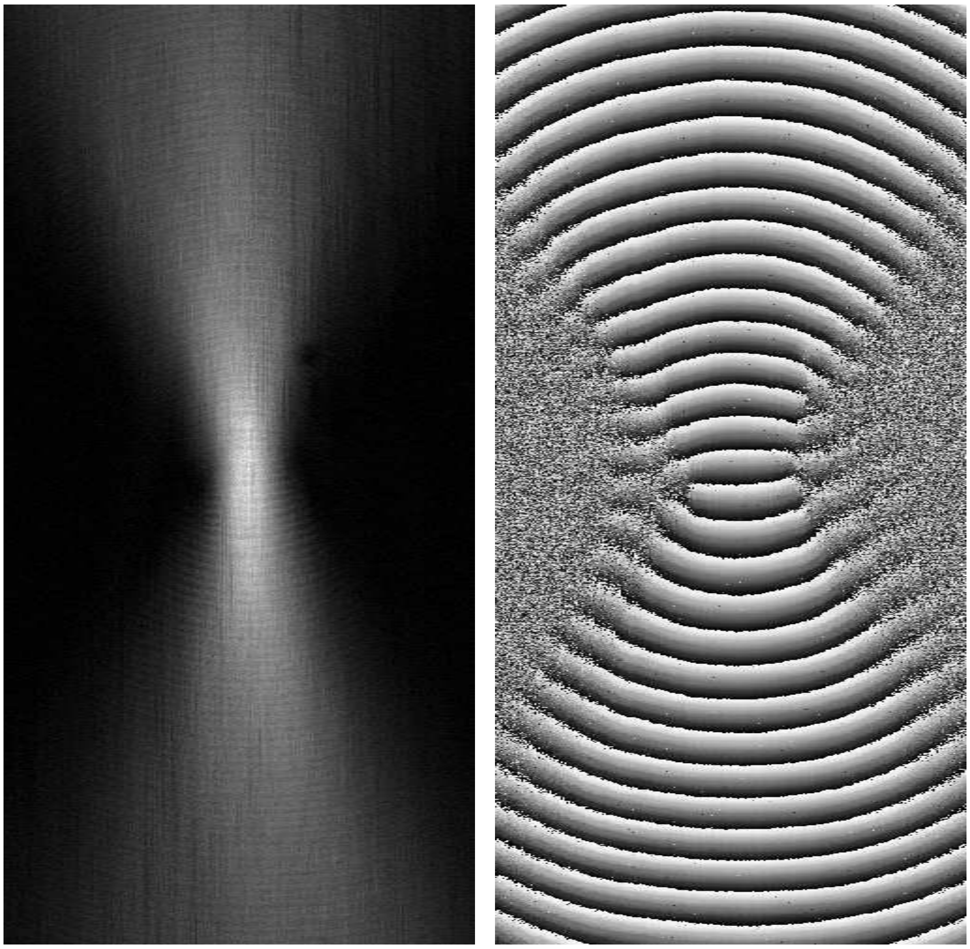

Figure 3.3: Amplitude (left) and phase (right) aerial point spread functions produced from a source of concentric arcs. The images were obtained using the O-SAM on an 'ideal' sample, by scanning the detection point relative to the generation arcs and the sample. The image size is 0.5×1mm, and each pixel represents an area of $2 \times 1\mu$m.

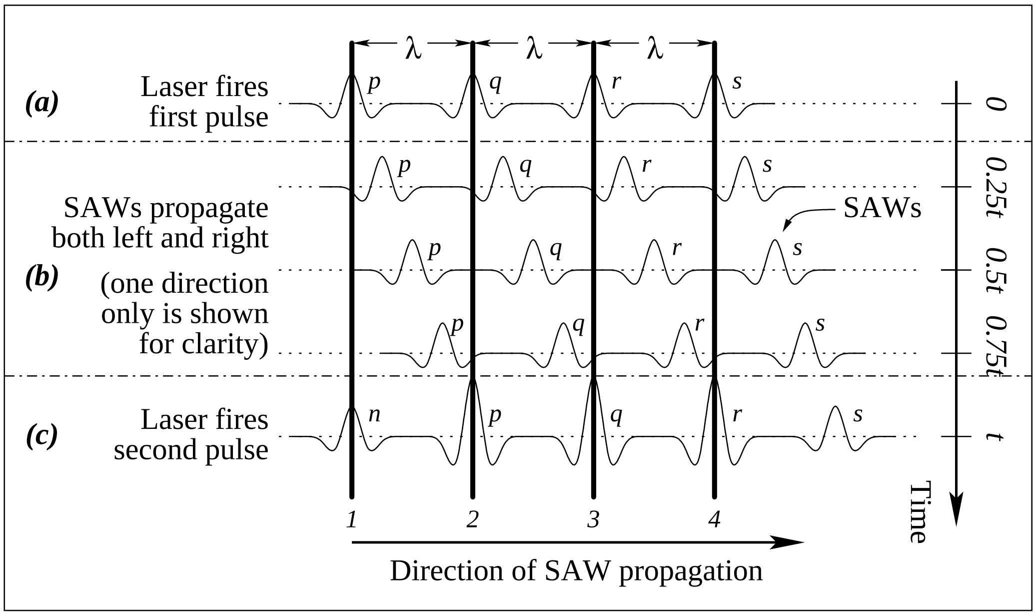

Figure 3.4: Schematic diagram showing how, by selecting the separation distance between excitation sources to be one wavelength, the excitation energy can be distributed such that all the acoustic energy from each of the sources adds up in phase.

Multiple lines

The key to the successful distribution of generation laser energy lies in ensuring that the acoustic waves produced from each part of the generation source add up in phase. This is achieved by careful choice of separation distance between the line sources.

Working from Krylov [66] and Karabutov [67], the amplitude of the SAW, a_r can be related to the source line spacing as follows:

$$a_r \propto \frac{1}{\omega_r} F_s(k_r) F_t(\omega_r) \tag{3.2}$$

where $F_s(k_r)$ is the component of the Fourier transform of the spatial distribution of the source intensity of the SAW wave number k_r, and $F_t(\omega_r)$ is the corresponding temporal frequency component of the source at the SAW frequency ω_r. This model is valid provided the thermal diffusion length (at the SAW frequency) and the optical absorption length are small compared to the SAW wavelength.

A simplified demonstration of the above theory will help to illustrate several points. Consider the optical distribution of four lines—labelled *1* to *4*—as depicted in figure 3.4, where the lines are one acoustic wavelength, λ, apart. At time 0 in part *(a)* of the figure the laser fires the first pulse, which displaces the surface of the sample along the four line sources and generates four surface waves, p, q, r and s. As time increases (going down the figure, in part *(b)*) the SAWs propagate both forwards and backwards along the sample surface—here just the forward-going waves are shown for clarity. At time t (part *(c)* of the figure), the laser fires a second pulse. Provided that $\lambda = vt$ (where v is the SAW velocity) some of the SAWs generated from the first pulse will be directly beneath the line sources. As shown in the figure, line *1* will generate a new SAW n, whilst the amplitude of SAWs p, q and r are increased by lines *2, 3* and *4*. Further laser pulses at multiples of time t will further increase the amplitude of any passing SAW—for the Q-switched, mode locked source used in the O-SAM instrument, there are approximately thirty such laser pulses in each pulse envelope as illustrated in figure 3.2a (which shows one such envelope). The fundamental frequency of these pulses is 82MHz.

This somewhat simplistic description illustrates several important practical points. The first is that the separation distance between the line sources is critical, and this is dependent on the frequency of the laser source and the SAW velocity—this implies that different line separations are required for different materials. This becomes more critical the more lines are used, and this point is illustrated in the optical zoom section. The second point is that the greater the number of lines that are used, the narrower the frequency content of the propagating SAWs, because more SAWs are generated by the separate sources. Our thirty pulse Q-switched, mode locked laser source may well produce thirty SAWs from a single line source, but will produce almost double that from thirty line sources, with a corresponding reduction in bandwidth. The third and final point is that the shape of the group of propagating SAWs is altered by the addition of extra line sources.

3.2.4 Frequency control

By varying the spacing of the multiple lines, it is possible to affect the frequencies of ultrasound that propagate on the sample surface. The envelope of pulses from the excitation source has a fundamental frequency of 82MHz but, as shown in figure 3.2b, this is rich in harmonics. In certain circumstances it may be desirable to suppress either the higher frequencies, or the fundamental frequency. This can be achieved by careful selection of the line spacing; interference between surface acoustic waves can be used to suppress different frequency components.

Band pass filter

Enhancing the harmonic signal and suppressing the fundamental can be important when trying to image at high frequencies where the signal at the fundamental is often much larger and tends to swamp the desired harmonic. A band pass filter is achieved by having a line spacing equal to the

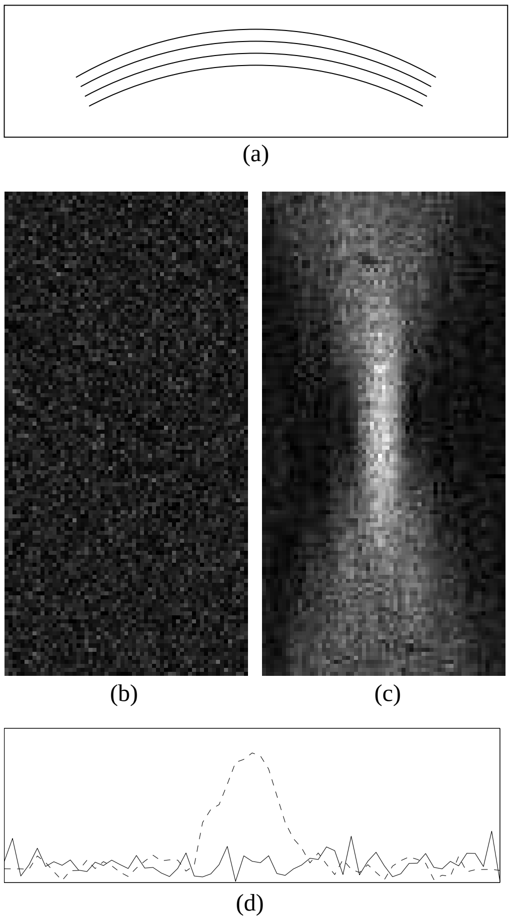

Figure 3.5: Realisation of a surface acoustic wave band pass filter (a), that removes the fundamental frequency of 82MHz inherent in the optical excitation energy. (b) and (c) show the aerial PSFs at 82MHz and 164MHz respectively. (d) is a line plot through the focus at both frequencies. Aerial PSF size is $600 \times 300\mu$m.

wavelength of the frequency that you wish to propagate. Frequencies at sub-multiples below this are then cut off.

Figure 3.5 shows the SAW PSF obtained at the focus of a four-arc distribution, which is shown schematically in figure 3.5a. The separation between adjacent lines is half a wavelength at 82MHz and thus a whole wavelength at 164MHz. The excitation from each 82MHz line cancels out, whereas the 164MHz signals add in phase. The distributions shown in figure 3.5 clearly demonstrate this, since no 82MHz signal can be detected above the noise in figures 3.5b and 3.5d, whereas a diffraction-limited focus for the 164MHz signal is observed in figures 3.5c and 3.5d.

Band stop filter

It is equally, if not more useful sometimes to remove an inherent harmonic frequency present within the optical excitation signal, so that only acoustic waves of the fundamental frequency are generated. This is a requirement of harmonic imaging, which is the process of exciting acoustic waves at a fundamental frequency and detecting the harmonic generated by non-linearities within the sample, and is discussed in detail in section 8.3. Suffice it to say for the moment that to successfully perform harmonic imaging it is essential that any harmonic content in the input signal be suppressed so that it is not confused with generated harmonic.

To ensure efficient suppression of the harmonic, the line spacing is chosen to be 1.5 acoustic wavelengths at the harmonic, thus effecting the cancellation of this signal. The separation at the fundamental frequency is therefore 0.75 acoustic wavelengths, so that the waves generated at each line add up in quadrature. The amplitude of the resulting surface acoustic wave at the fundamental frequency has an amplitude equivalent to $1/\sqrt{2}$ of the optimum case where all the contributing lines add up in phase, rather than in quadrature. Nevertheless, the harmonic signal is very efficiently suppressed, so that it is not visible above the noise.

The distributions shown in figure 3.6 clearly demonstrate this, since a diffraction-limited focus for the 82MHz signal is observed in figure 3.6b and 3.6d, whereas no 164MHz signal can be detected above the noise in figure 3.6c and 3.6d.

3.2.5 Spatial control using diffractive acoustic elements

The previous section has shown how interference between surface acoustic waves can be used to suppress different frequency components. This can be applied to both plane waves and focused waves. Controlling the frequency of the ultrasound is only one aspect of control, however; in many cases we may wish to control the spatial distribution of the propagating waves. At its simplest level, this involves tilting a line source to affect the propagation direction; tilting an arc source and altering its focal length allows the focus of the ultrasound to be moved to an arbitrary point. Further spatial control may be achieved by using diffraction.

In this section, the measured acoustic fields produced by a range of diffractive acoustic elements are compared with the predicted acoustic fields according to simulation. The predicted results were calculated by a model developed by Matt Clark, based on an angular spectrum propagation technique described by Goodman [68].

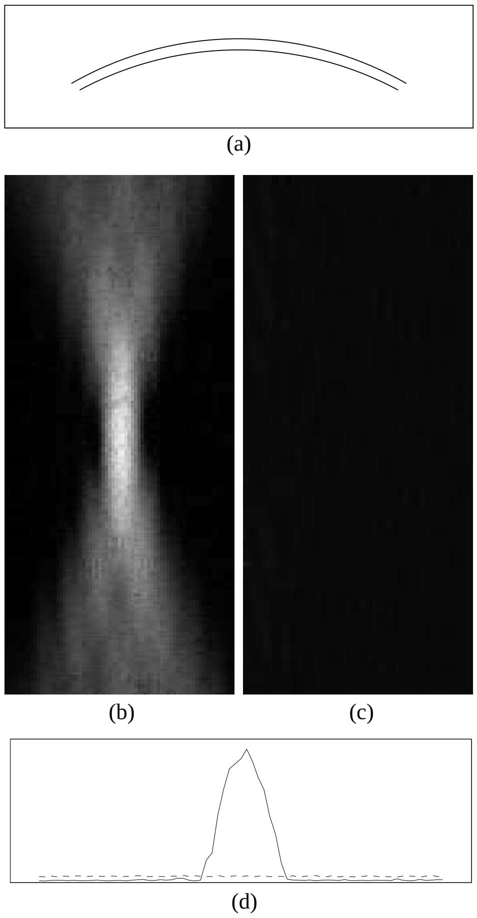

Figure 3.6: Realisation of a surface acoustic wave band stop filter (a), that retains the fundamental frequency of 82MHz inherent in the optical excitation energy and removes the 164MHz harmonic. (b) and (c) show the aerial PSFs at 82MHz and 164MHz respectively. (d) is a line plot through the focus at both frequencies. Aerial PSF size is $600 \times 300\mu$m.

Parallels between acoustic and optical elements

At this point it is useful to draw analogies between the acoustic elements described below and some common optical elements. A simple optical convex lens is described as a refractive element because it uses refraction to—for example—focus a collimated beam of light to a point. An 'arc' line source acoustic element generates a continuous acoustic wavefront which is focused to a point, so the device may be thought of as analogous to a refractive lens, even though there is no acoustic refraction occurring.

A zone plate or diffractive optical element may also be used to focus a collimated beam of light. Using the analogy between optical and acoustic elements, one can imagine that there is an equivalent of an optical zone plate in the acoustic domain. This would consist of discrete segments or zones, rather than a continuous line or curve. Drawing from these analogies between optical and acoustic elements, a device that produces a continuous wavefront would be termed a continuous wavefront acoustic element, whilst a device that produces a set of discrete wavefronts would be termed a diffractive acoustic element (DAE).

One of the simplest optical diffractive elements is the binary phase plate, which introduces 180 degree phase shifts between adjacent zones, so that it acts as a discrete lens. An array of short line excitation sources, offset from each other by half the Rayleigh wavelength forms a DAE which acts as a surface acoustic wave analogue of an optical zone plate. As with optical zone plates [69], the diffractive acoustic element produces a converging and a diverging beam corresponding to the +1 and -1 orders.

A set of computer generated holograms (CGHs) which acted as DAEs were designed by Matt Clark and Mike Somekh, and fabricated. The theory was later extended to produce images used by the spatial light modulator to exhibit the same degree of control over the acoustic wavefront.

Design considerations

DAEs, like their equivalent optical elements, are by their very nature highly dispersive. This dispersion can either be an inconvenience, or can be used to produce frequency-dependent surface acoustic wave focusing.

Unlike their optical counterparts, the DAEs used to excite surface acoustic waves are essentially one-dimensional—as opposed to two-dimensional—elements. To suppress the zero order in optical zone plates, it is only generally necessary to ensure that there is a 180 degree phase shift between adjacent zones; due to the large number of zones and the approximately equal areas in each zone—arising from the two-dimensional nature of optical elements—areas corresponding to each phase shift are naturally approximately equal. A little additional care is needed in DAE design to ensure

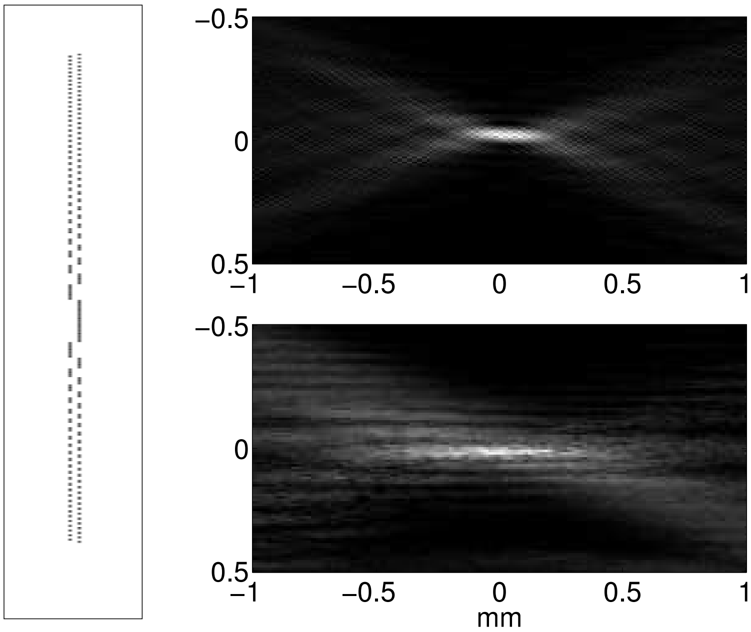

Figure 3.7: Diffractive acoustic element (expanded horizontal scale), predicted (top) and measured (bottom) psfs

that the total line lengths (and hence generated SAW power) corresponding to each phase shift are approximately equal, so that the zero order is effectively suppressed.

The temporal nature of the excitation source also needs to be taken into account when designing DAEs. Figure 3.7a shows a schematic of a DAE—the distance between front and back elements has been exaggerated—which was designed and fabricated as a CGH. The intended focal length of the acoustic element is 10mm for the +1 order. Figure 3.7b shows the predicted SAW amplitude distribution at the element's focus region, whilst figure 3.7c shows the measured amplitude distribution. The focus of the measured PSF is less well defined than the predicted PSF. This is not due to any failings of the acoustic element - rather it is due to the finite envelope length of the source and hence of the Rayleigh wave packet. The path length difference for acoustic waves generated at the centre and edge of the acoustic element is 1.4mm. At a Rayleigh wave velocity of $\approx$3100ms^{-1} this equates to a time delay of $\approx$450ns, or 37 cycles at 82MHz: the temporal length of the acoustic envelope at the predicted focus is more than doubled - see figure 3.8. Since the measured PSF in figure 3.7 is simply a measure of the *peak* of the pulse envelope at each point, the focus appears to be less well defined.

The temporal spreading of the resulting pulse envelope around the focus, in addition to reducing the peak amplitude, also reduces the amount of destructive interference taking place between

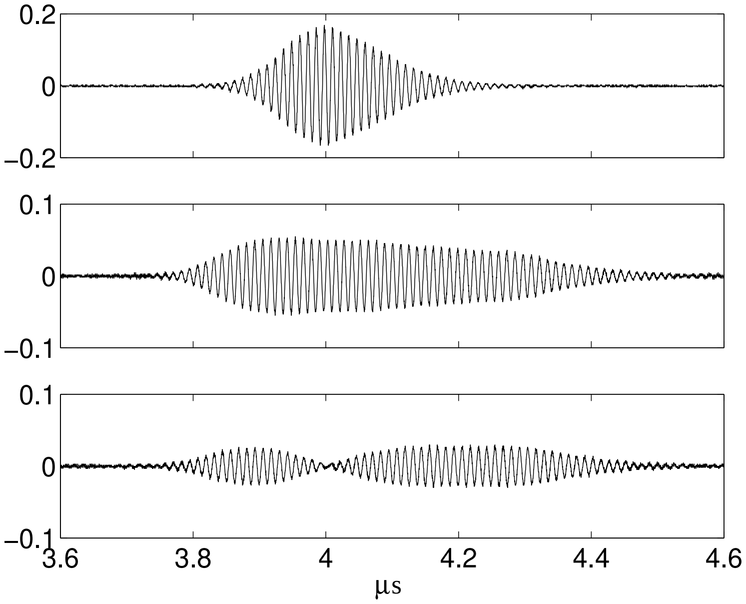

Figure 3.8: The first trace was obtained from the focus of an arc acoustic element and shows no envelope broadening. The second trace was obtained from the focus of the diffractive element shown in figure 3.7. The third trace was obtained from just above the focus of the diffractive element, within a 'dark lobe.' The horizontal scale denotes time in μs since Rayleigh wave excitation.

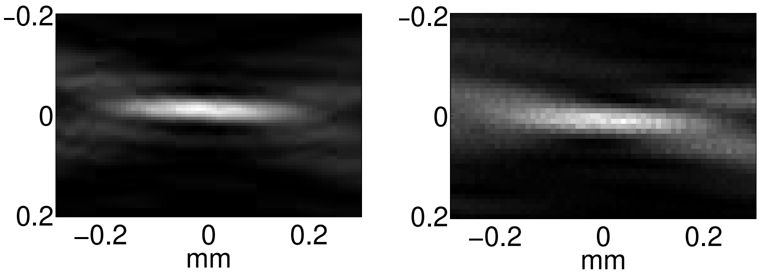

Figure 3.9: Predicted (left) and 82MHz filtered measured (right) PSFs at focus

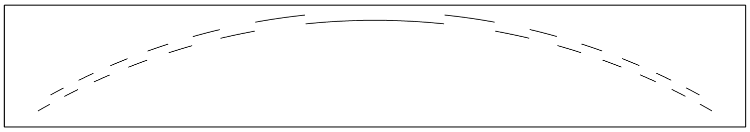

Figure 3.10: Hybrid continuous wavefront/diffractive acoustic element. The spacing between front and back segments has been exaggerated.

wavefronts originating from different zones. It is this destructive interference that leads to the characteristic dark lobes above and below the focus in the simulated PSF in figure 3.7. Since the Rayleigh wave packets are separated temporally (although not spatially) they do not interfere, as shown in the third trace of figure 3.8. This leads to a larger PSF.

To allow the pulse envelopes originating from different zones to interfere, the length of the generated pulse envelopes would have to be increased. Since this is not possible with the experimental setup, the equivalent is achieved by acquiring the full waveform at each scanning point, and taking the 82MHz component of the Fourier transform of the signal - effectively increasing the pulse envelope length by very sharp filtering. The results are shown in figure 3.9. Since the whole waveform must be captured at each point in order to obtain the 82MHz component - rather than just the peak amplitude - scanning takes much longer, hence the small scan area.

To address the temporal problems of the zone plate DAE, whilst still retaining aspects of diffraction to control the position of the Rayleigh wave focus, several hybrid continuous wavefront/diffractive acoustic elements have been developed. These elements have aspects of both a regular arc and the acoustic zone plate described above.

Hybrid continuous wavefront/diffractive acoustic elements

The first example of such a device, again a CGH, is shown in figure 3.10

The arc is broken into just 13 zones, the result of imposing a weak phase plate onto a continuous

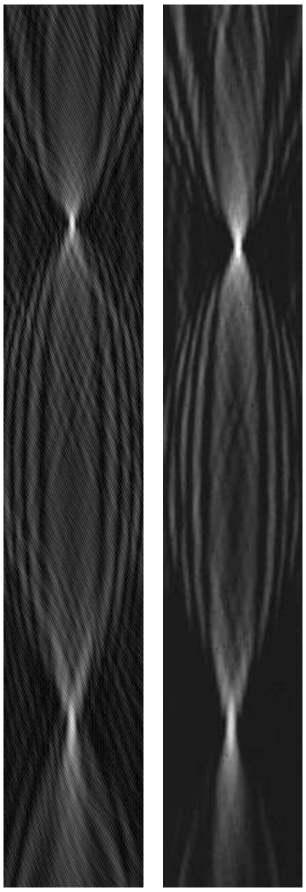

Figure 3.11: Predicted (a) and measured (b) SAW amplitude distributions at 82MHz produced by the hybrid continuous wavefront/DAE shown in figure 3.10. Image size is 0.8×6.4mm.

arc. The arc provides the dominant focusing effect, and the phase plate imposes a relatively weak perturbation. Due to the reduced number of zones, the maximum delay of wavefronts originating from the extremes of the element compared to wavefronts from the centre of the element is greatly reduced to six cycles, allowing effective interference—and therefore focusing—to take place from our 30-cycle excitation source. The relatively weak zone plate gives a +1 and -1 diffracted order such that two foci are formed—one inside the geometrical focus and one outside. If the focal length of the arc is 10mm and the focal length of the zone plate is 76mm, we expect the two foci to be located at 8.8 and 11.5mm. Figure 3.11a shows the predicted/simulated SAW amplitude distribution around the foci, figure 3.11b shows the measured distribution.

The half-wavelength phase delay at 82MHz imposed by the displaced elements gives a whole wavelength relative delay at 164MHz, so that the phase plate is "invisible" to the harmonic. The

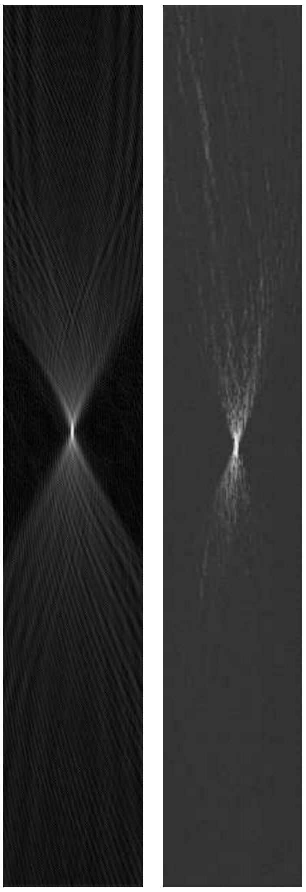

Figure 3.12: Predicted (a) and measured (b) SAW amplitude distributions at 164MHz produced by the hybrid continuous wavefront/DAE shown in figure 3.10. Image size is 0.8×6.4mm.

focus at 164MHz thus occurs at the geometrical focus of the arc. This is verified in figure 3.12, which shows the predicted (a) and measured (b) SAW amplitude distributions at 164MHz. The quantitative agreement between measured and predicted distributions may be determined from figure 3.13, which shows the measured and predicted amplitudes through the foci at 82MHz (3.13a) and 164MHz (3.13b).

In addition to controlling the focal position along the axial direction, a hybrid continuous wavefront/DAE can act as an acoustic analogue of a Nomarski objective in optics, which forms two adjacent foci. This is achieved by imposing a grating onto a continuous arc. The +1 and -1 orders are the equivalent of phase tilts on the beam, which displace the focus either to the left or to the right (if the beam is propagating down vertically). The element produced to achieve this is shown schematically in figure 3.14 The simulated and measured amplitude distributions at 82MHz

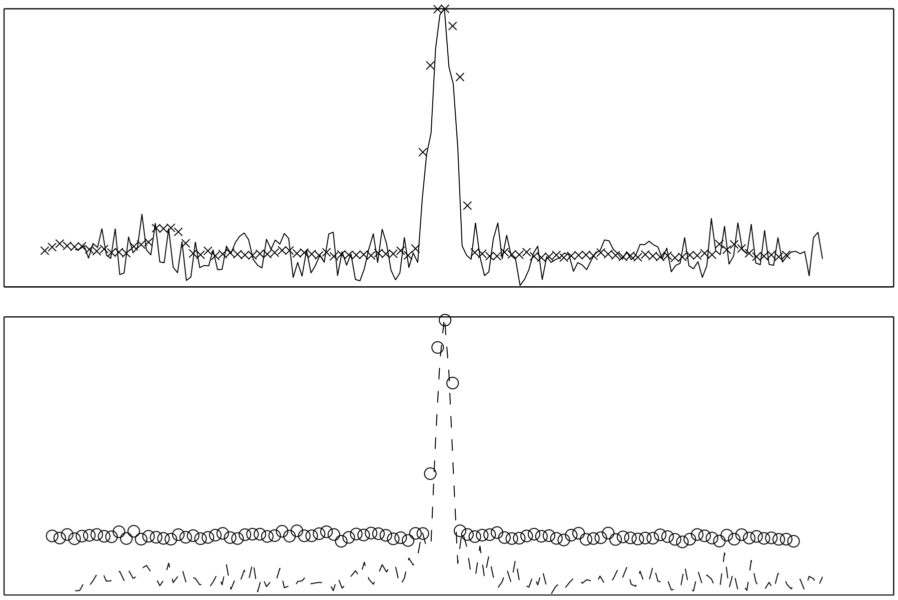

Figure 3.13: Predicted (-) and measured ($\times$) amplitude through the first 82MHz focus (a), and the predicted (- -) and measured ($\circ$) amplitude through the 164MHz focus (b)

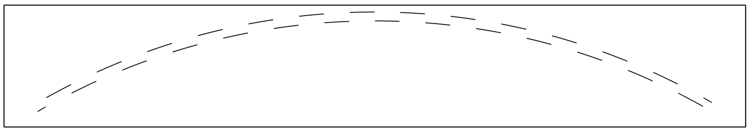

Figure 3.14: Hybrid continuous wavefront/diffractive acoustic element. The spacing between front and back segments has been exaggerated.

are shown in figure 3.15a and 3.15b respectively, a line plot through the foci is shown in figure 3.16. The agreement between predicted and measured response is very good.

A final demonstration of the flexibility of the hybrid element approach achieved using a CGH as the method of producing the surface acoustic wave excitation source (as opposed to using the SLM) is shown in figures 3.17, 3.18 and 3.19. In this element, the symmetry is no longer maintained, achieved by displacing the phase plate about the axis. The +1 and -1 orders now differ not only in their axial position, but also in their lateral displacement relative to the axis. Figures 3.18 and 3.19 show excellent agreement between the simulated and measured distributions, even showing extremely close agreement in the fine detail of the surface wave amplitude distribution between the two foci.

The DAEs and hybrid elements shown so far have all been imaged onto the surface of the

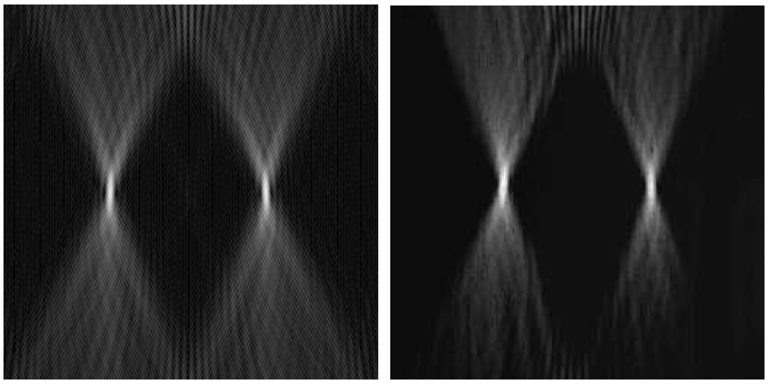

Figure 3.15: Predicted (a) and measured (b) SAW amplitude distributions at 82MHz produced by the hybrid continuous wavefront/DAE shown in figure 3.14. Image size is 2.4×2.4mm.

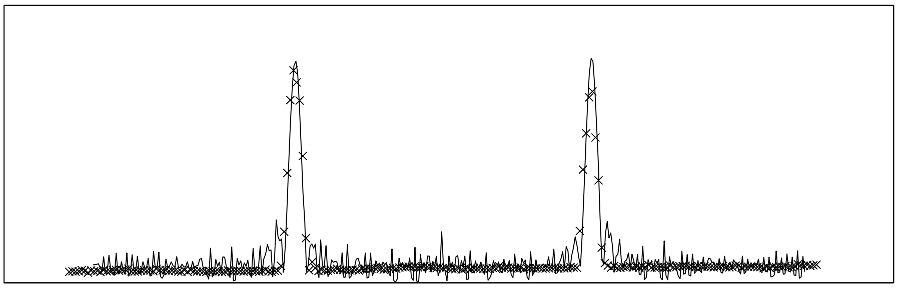

Figure 3.16: Predicted (-) and measured (×) amplitude through the two 82MHz foci.

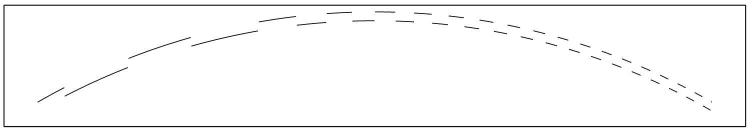

Figure 3.17: Hybrid continuous wavefront/diffractive acoustic element. The spacing between front and back segments has been exaggerated.

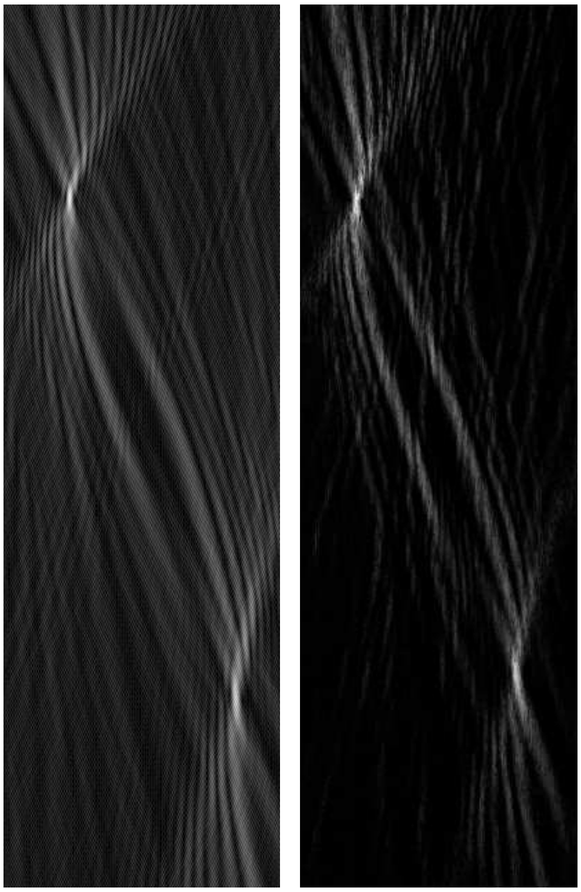

Figure 3.18: Predicted (a) and measured (b) SAW amplitude distributions at 82MHz produced by the hybrid continuous wavefront/DAE shown in figure 3.17. Image size is 2×6.4mm.

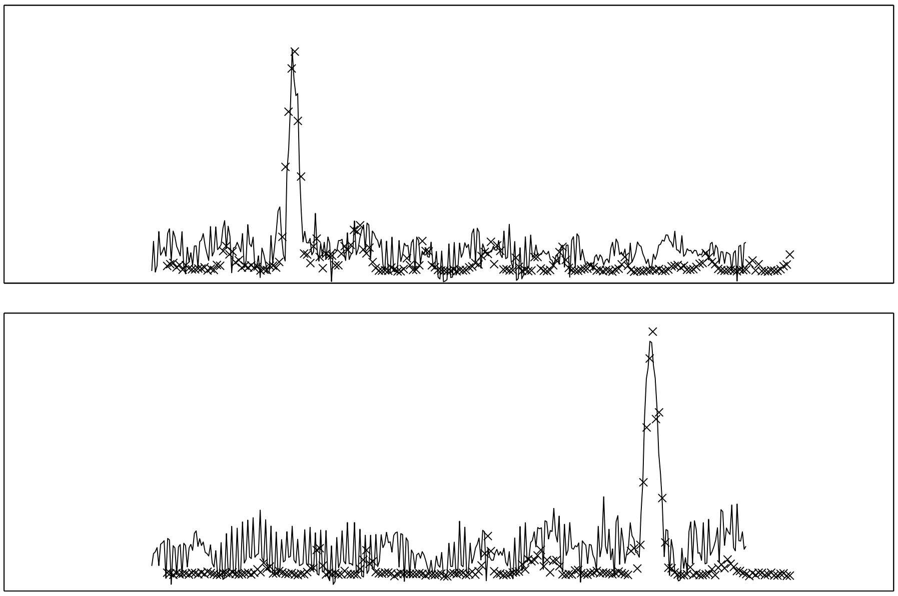

Figure 3.19: Predicted (-) and measured (×) amplitude through the two 82MHz foci.

material using computer generated holograms (one for each design, obviously). The equivalent designs, with some modifications, may also be imaged onto the surface of the material using a spatial light modulator. These have the advantage of being able to be changed and altered in real time.

Figure 3.20 demonstrates several aspects of control over the acoustic waveform propagation. On the left of each image of figure 3.20, the light distribution imaged onto the surface of the material by the SLM are shown—"SLM images"—and the corresponding SAW amplitude distributions are shown opposite each one on the right of figure 3.20. The fringe spacing has been increased for the illustration for clarity, but the patterns are otherwise unchanged. As can be seen from the amplitude distributions from top to bottom, the relative amplitudes of two foci, each displaced both axially and laterally, can be adjusted by changing the distribution of arcs on the SLM image. The images shown are experimental results, not simulations.

The method of calculating the SLM images is relatively simple. Consider a simple series of arcs focusing surface waves—they are sections of concentric circles, separated from each other by one acoustic wavelength. To produce one focus, the phase relative to that focal point is calculated for each point on the SLM image, and this phase is then quantised to two levels to produce the required excitation pattern. To generate two foci, the complex amplitude for each point on the

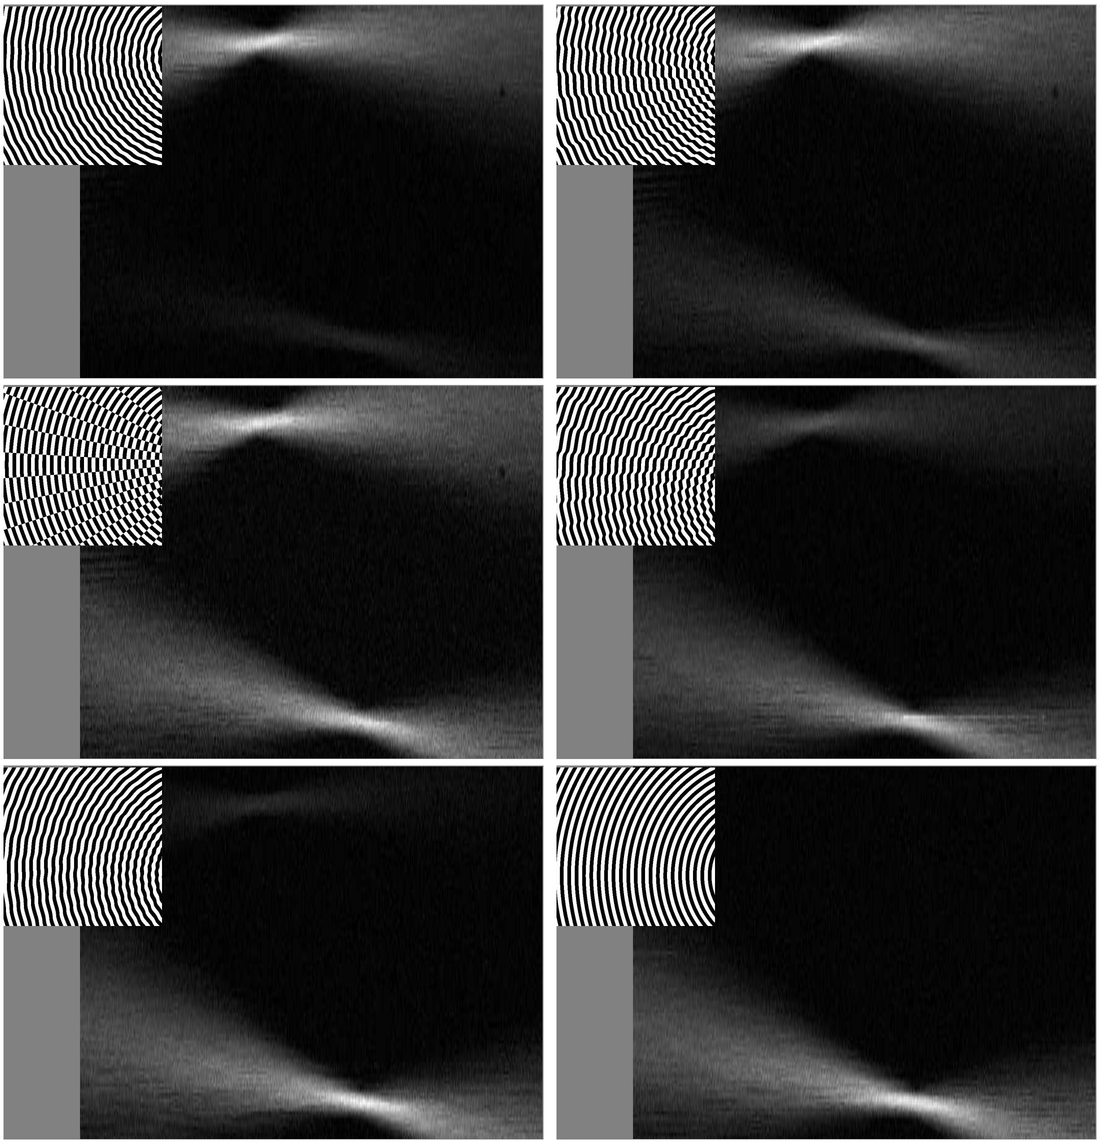

Figure 3.20: On the left of each image are distributions of light imaged onto the surface of the material by the SLM. Next to each one is the corresponding SAW amplitude distribution, as measured by the O-SAM instrument. The fringe spacing of the excitation images has been increased for clarity.

SLM image, relative to each focus, is calculated:

$$Amplitude = Ae^{\frac{i2\pi d}{\lambda}} \tag{3.3}$$

where A is the relative amplitude of each focus, d is the distance from the focus to each point on the SLM image, and λ is the acoustic wavelength. The two complex amplitude maps (one for each focus position) are then simply summed, and the phase quantised to two levels to produce the required excitation pattern.

3.3 Computer generated holograms to control the ultrasound source

Section 3.2 described the overall approach to generating ultrasound with the O-SAM instrument—focusing the surface waves, and spreading out the optical power. During the development of the instrument, the way of achieving this has been accomplished by two different methods. Computer generated holograms (CGHs), the first of these two methods, are described in this section.

The computer generated holograms used are binary phase diffractive optical elements. They are designed with the profile of the source laser beam in mind, and they shape and focus the beam on the surface of the material. They can even be designed to focus the beam over a three dimensional profile [2].

The CGHs were manufactured using the process of e-beam lithography at the University of Glasgow. A quartz optical flat is coated with chromium and photo-resist. The data from the CGH design algorithm is then used to control an electron beam-writer. The photo-resist coating is then developed and the chromium regions are wet-etched. The remaining photo-resist is removed and the exposed regions are dry-etched to ensure a π phase shift at the generation laser wavelength in the etched regions. After the removal of the chromium the process is complete.

3.3.1 Optical setup

The physical optical setup for using a CGH to image a SAW generation profile onto the surface of a material is relatively simple and straightforward. Figure 3.21 shows a schematic of the optical setup.

The generation laser's beam size and collimation is adjusted by the optical zoom system—described below—before it passes through the CGH itself. This focuses and shapes the beam, such that a sharp image—usually a set of concentric arcs— is imaged onto the sample surface. The distance between the CGH and the sample surface is manually adjustable using a mechanical stage for ease of alignment and focusing.

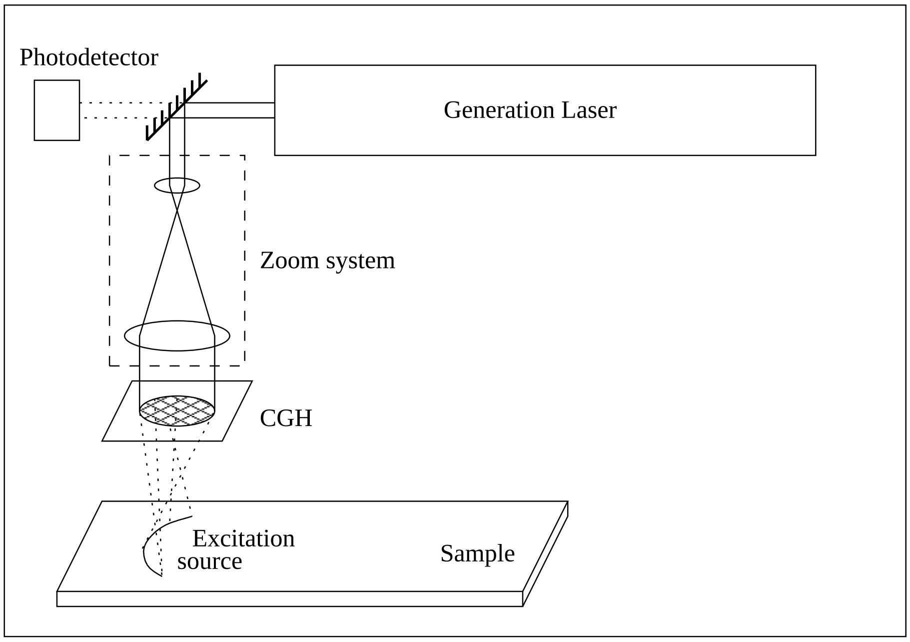

Figure 3.21: Schematic of the optical system used to image a SAW excitation profile onto the sample surface, using a computer generated hologram. The optical zoom system is also shown.

The problem of getting the line spacing correct for the material under investigation is referred to as ‘velocity matching,’ and it is a serious problem for arc or straight line sources consisting of 16 or 32 lines, since the first and last lines must add up in phase. Since it is impractical, costly and inconvenient to have a CGH manufactured for each specific arc or line spacing, an optical zoom system is used to either increase or decrease the size of the image formed on the sample surface by the CGH, thus altering the line spacing. This zoom system is placed between the Nd:YAG generation laser and the CGH as shown in figure 3.21.

The magnification factor is determined by the amount by which the beam is either converging or diverging as it passes through the CGH; this in turn is determined by the distance between the two lenses. As the magnification factor is adjusted, the distance between the CGH and a sharp image on the sample surface also changes. A further complication is that the beam as it exits the laser cavity is naturally divergent, and this divergence must be accommodated for in the determination of lens separation.

The constructed zoom system doubles the diameter of the laser beam at the point where it encounters the CGH, allowing the whole of the CGH to be illuminated, which ensures the most accurate reconstruction of the distribution of arcs or lines on the sample surface. The spacing of the lines produced by the CGH may be adjusted by around plus or minus 20% from the design spacing, allowing for a similar variation in material velocities.

3.4 Spatial light modulator to control the ultrasound source

The CGH method of imaging an excitation pattern onto the sample surface is efficient, elegant, flexible and simple. However, it features an obvious constraint in that one CGH design is required for each excitation pattern. The optical zoom system described in the previous section allows designs to be scaled, to a certain degree, to allow for different line spacings to match the velocity of the material, but this is a slow and cumbersome process.

A spatial light modulator (SLM) allows us to image any arbitrary pattern of light—within the constraints of the number of pixels on the device—onto a flat surface. This pattern can be changed in real time, at rates of around 1kHz. This flexibility allows us to adapt the generation profile to the material properties.

3.4.1 Details of the SLM

The SLM used in the O-SAM instrument is a 512 $\times$ 512 pixel multi-level/analog ferroelectric liquid crystal device made by Boulder Nonlinear Systems. The polarisation of incident light can be rotated by up to plus or minus 45 degrees on each pixel, and is then reflected. The device is described as 'analogue' since each pixel can have a voltage applied to it that corresponds to a phase rotation of between -45 and +45 degrees for reflected light, this voltage being derived from an 8-bit digital number. In theory, therefore, there should be 256 different phase angles possible between -45 and +45 degrees. Experience has shown us that the response is highly nonlinear, and in terms of discernible levels, there are significantly less than 256. In its role within the O-SAM instrument, the SLM is treated as a binary display device—sending each pixel either '0' or '255' rotates the polarisation of the reflected light by either plus or minus 45 degrees.

The pixels on the device have a pitch spacing of 15μm, with a fill-factor of 83%. Reflective losses occur in the glass window covering the device (despite it being anti-reflectivity coated for the wavelength of the generation laser) and coupling of reflected energy into orders other than the zero order. The manufacturers claim 65% reflected energy back into in the zero order and quick, rough tests carried out in the lab confirmed this value.

Up to 1024 images may be stored in an 'internal image buffer' of the card controlling the device, and these images may be uploaded to the device at a maximum rate of 1015Hz. New image uploads from the host computer to the SLM image buffer may be performed at a typical rate of around 100Hz.

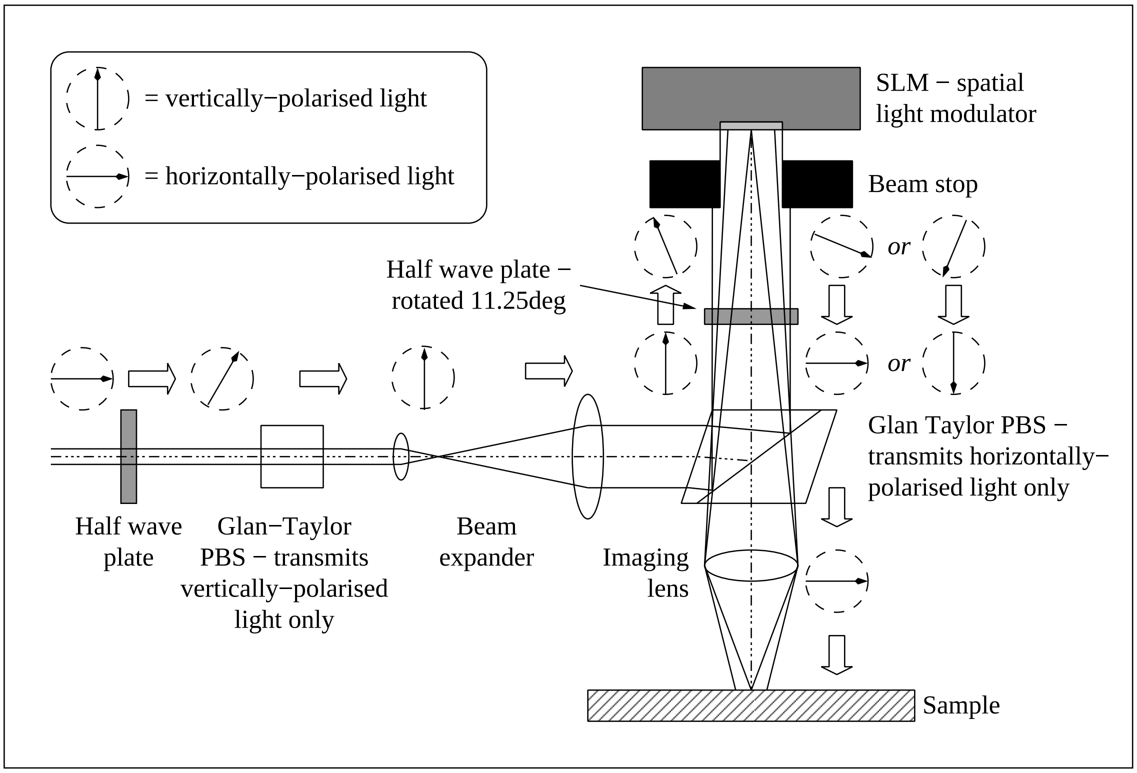

Figure 3.22: Schematic of the optical system used to image a SAW excitation profile onto the sample surface, using a spatial light modulator. The polarisation of the light is shown at various points.

The device needs to be 'd.c. balanced'—in practise, this means that the image on the SLM must be inverted within a certain time period, with a duty cycle between 'true' and 'false' images of exactly 50%. The minimum rate at which this can change, termed the 'repetition rate' of the SLM, is 0.1Hz. The maximum rate is 1015Hz however, since there is a constant period of time required to load the image onto the device and for the image to settle, then the higher the SLM repetition rate, the greater the chance that a pulse of incident optical energy from the generation laser will occur within one of these periods of uncertain phase output. The device interfaces to a host computer via a PCI card interface.

3.4.2 Optical setup

The optical setup for the SLM method of producing a generation profile within the O-SAM system is a little more complicated than the CGH method. Figure 3.22 is a schematic of the optical setup.

The laser beam from the generation laser is incident on a rotatable half-wave plate (left of diagram), and then a Glan-Taylor polarising beam splitter, which is oriented such that the sloped interface between the two halves of the beam splitter is from top to bottom of the optic, rather than left to right. In this orientation, the Glan-Taylor only lets vertically-polarised light pass; any horizontally-polarised component of the light is reflected downwards to a beam dump. By rotating the half-wave plate, the intensity of the [vertically-polarised] light emerging from the Glan-Taylor can be varied.

On emerging from the Glan-Taylor beam splitter, the beam passes through a beam-expander/ collimator. The size of the active area of the SLM device is 7.68mm square and it is important, from an ultrasound generation point of view, to illuminate this as uniformly as possible, without wasting too much laser energy on non-active areas of the device. The lenses are chosen to increase the diameter of the beam by a factor of approximately 3.2, and the distance between them is carefully adjusted to give a collimated beam.

This beam is then incident on another Glan-Taylor polarising beam splitter, and, because the incident beam is completely vertically-polarised (after emerging from the previous [rotated] Glan-Taylor) it is completely reflected to the left—towards the top of the figure—by the vertical interface between the two halves of the beam splitter. The beam splitter is designed such that the refraction incurred at the interface is cancelled by the rather strange angle of the end faces—the upshot of which is that the beam emerges at 90 degrees from its angle of incidence.

The beam then passes through another half-wave plate. This is orientated at 11.25 degrees from the 'vertical'—where 'vertical' in this case would be a position such that incident vertically-polarised light would not be rotated. The polarisation of the incident beam is rotated therefore by 22.5 degrees as it passes through on its way to the SLM, as shown in the figure.

Between the half-wave plate and the SLM is a 'stray light' guard, which is simply a square hole the size of the active area of the SLM cut into a piece of 10mm-thick aluminium that has been spray-painted black. This absorbs excess energy from the laser beam, and prevents it from hitting the area of the SLM around the active area. This is important, since the area around the active area is not as reflective as the active area; the incident energy would be absorbed within the device, causing it to heat up and malfunction, or even be destroyed.

The laser light—now polarised at 22.5 degrees from the vertical—hits the active area of the SLM and, depending on the state of each individual pixel, is rotated by either plus or minus 45 degrees. The light reflected from the SLM has components that are either polarised at +67.5 degrees, or -22.5 degrees from the vertical. This light is then rotated by a further 22.5 degrees as it passes through the half-wave plate so that, upon entering the polarising beam splitter, some components of the beam are horizontally-polarised, and some are vertically-polarised. The horizontally-polarised portion of the light passes straight through the beam splitter, whilst the vertically-polarised light

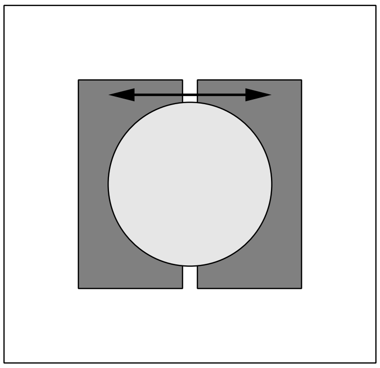

Figure 3.23: A split photodiode, as used in the modified knife edge detection system.

is reflected away from the sample by the beam splitter

Upon emerging from the beam splitter, the remaining [horizontally-polarised] light is imaged onto the surface of the material under investigation. The focal length of the lens is 40mm and is 167mm from the SLM device. By the Gaussian Lens Law, an image is formed on the surface when it is 52.6mm from the lens. This gives a magnification factor of 0.315; the image on the sample is therefore 2.42mm square.

Issues concerning the alignment of the optical setup for the O-SAM system employing an SLM are dealt with in section 3.6.

3.5 Detection of ultrasound using a 'modified' knife edge detector

Section 2.3 described the principles behind the knife edge detection technique for detecting surface acoustic waves, as well as other detection configurations. This section describes in some detail the 'modified' knife edge detection system employed in the O-SAM instrument.

3.5.1 Split photodiode technique

The knife edge detection technique described in section 2.3.1 is inherently wasteful of optical energy, since at the optimum sensitivity half the light is 'thrown away' at the knife edge. Since optical ultrasound detection techniques are relatively insensitive compared to their contact counterparts, as described in section 1.4.5, it important not to waste any available optical energy. With this in mind a split photodiode, as shown in figure 3.23 is used instead of a single photodiode coupled with a knife edge.

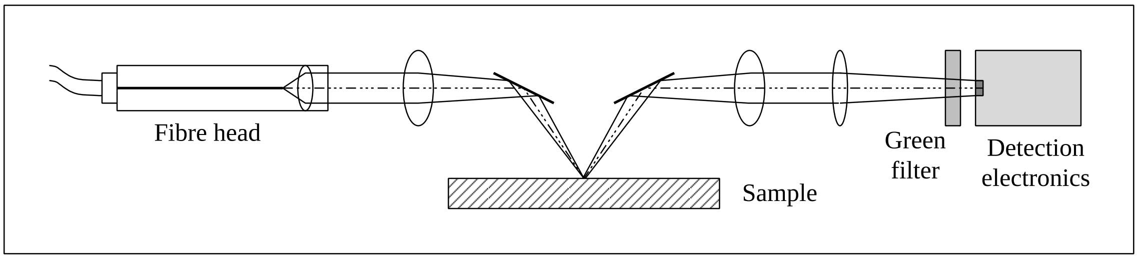

Figure 3.24: Optical configuration of the modified knife edge detection system.

As the surface waves cause the beam to be deflected left and right, so the intensity of light falling on one photodiode increases, as the other decreases. By measuring the difference in intensities, the angular deviation of the reflected beam may be measured.

The advantage of the system is more optical efficiency; the disadvantage is increased complexity in the electronics required to extract the signal. This will be described in section 3.5.4.

3.5.2 Optical laser source

The optical source for the detection system is a continuous wave doubled Nd:YAG laser, emitting around 100mW of optical power at a wavelength of 532nm (green). The light is coupled into a rugged single-mode optical fibre, allowing the laser source itself to be mechanically isolated from the O-SAM instrument—the laser has a cooling fan—and allowing the entire detection system to be mounted on movable stages. Around 40mW of optical power is emitted from the other end of the fibre.

3.5.3 Optical configuration

Figure 3.24 shows the optical configuration of the detection system in schematic form.

Light emerges from an optical fibre at the left of the figure, and is collimated by a fixed lens on the end of the fibre assembly. The diameter of the beam as it emerges from the collimating lens is 10mm.

The beam is focused onto the sample using a doublet with a focal length of 70mm. Approximately midway between the lens and the sample, the beam is reflected through an angle of approximately 140 degrees by a plane mirror, therefore hitting the sample at an angle of approximately 40 degrees from the normal. The spot size of the focused beam as it hits the sample is approximately 7μm in diameter. The beam is reflected from the sample at the same angle on the other side of the normal, is reflected by another mirror through the same angle of 140 degrees as the incident beam was. The beam is then collimated by another 70mm doublet.

The collimated beam, once again with a diameter of 10mm, then passes through a 'beam diameter reduction' lens with a focal length of 100mm. The split diode is located just forward of the focal plane of this lens. It is located here, rather than at the focus, for two very good reasons. Firstly, at the focus, there is no lateral deviation of the beam, since the beam effectively 'tilts' about the focal plane. Secondly, the larger the beam diameter is at the split photodiode, then the larger proportion of light falling on the diodes (as opposed to the gap between them) will be. Obviously, the position of the photodiodes relative to the lens is chosen such that the beam is of the optimum size—that is, just less than the combined width of the two halves of the detector. The split diode could alternatively be located just the other side of the focal plane, the only difference being the direction of lateral spot deviation with respect to angular deviation of the reflected beam from the sample.

As shown in the diagram, between the beam diameter reduction lens and the photodiodes is an optical filter. This inhibits any stray infrared light from the generation laser from reaching the split photodiode and causing an erroneous signal.

Adjustment requirements

This rather convoluted method of focusing a beam onto the sample surface is used to allow the generating laser source access to the sample along the normal to the sample surface. Recall that some of the generation profiles generated using computer generated holograms are of the order of 10mm wide. We additionally wish to move the detection point relative to the excitation region, both for adjustment and tuning, or to scan the aerial point spread function (see many of the previous figures in this chapter). We may also wish to scan the detection point over the excitation region, or even behind it. The use of the two mirrors to deflect the beam onto the sample and collect the reflected light allows for all these adjustments and scanning capabilities, whilst still keeping the detection system as a whole orthogonal to both the sample and the generation optics.

The entire detection system, as shown in figure 3.24, is constructed out of 'Microbench' (manufactured by Linos Photometrics) or compatible mechanical components. This acts as a robust mechanical basis for the optics, and allows the entire detection system to be mounted on motorised X-Y stages, with a manual Z-adjustment stage for focusing.

3.5.4 Electronic Configuration

It was mentioned in section 3.5.3 that the advantages of using a split diode technique (as opposed to a true knife-edge technique) were counter-balanced by increased complexity in the electronics. Those electronics are described in this section.

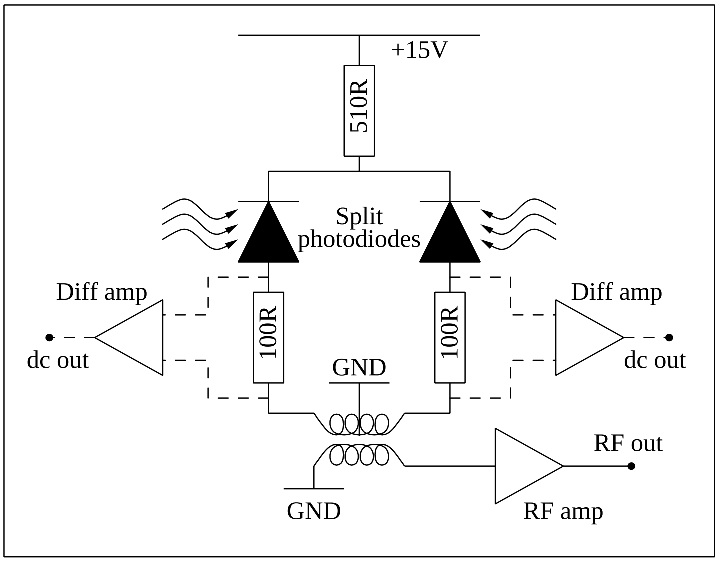

Figure 3.25: Electronic configuration of the split photodiode detection system.

Figure 3.25 shows the electronic detection system.

The split diode used is a Hammamatsu quad Si-PIN photodiode, model number S6058. Each of the four photodiodes measures 0.6mm $\times$ 1.2mm, and the gap between the photodiodes is 10μm wide. A quad photodetector was chosen in preference to a dual photodetector, since the photodiodes are smaller, and hence will have a lower junction capacitance and a correspondingly higher signal bandwidth. The data sheets on the diode specify their high frequency cut-off point at 150MHz with a reverse-bias voltage of 3V. Since we reverse-bias the photodiodes at 15V, this high frequency cut-off point is increased. Each photodiode shares a common cathode.

The photodiodes are reverse-biased through a 510Ω resistor—this, combined with the 100Ω resistor on each arm limits the maximum current through each photodiode to 25mA such that, in the event of a very bright laser source hitting the photodiode, it will not be damaged through over-current.

Each 'arm' of the system consists of a photodiode and a 100Ω resistor. The voltage across each resistor is naturally proportional to the amount of light falling on each photodiode. The voltage across each resistor is measured using a standard operational amplifier, with a reasonably low gain-bandwidth product—certainly orders of magnitude lower than the fundamental frequency of the excitation laser source—and provides a 'dc' level. The difference between the two voltages gives an indication of how 'centred' the detection spot is on the split photodiode; the sum gives an indication of how much light is reflected back from the sample onto the detector.

The rf signal, generated by the spot position oscillating between the two halves of the split detector, is measured as a change in the direction of the current flowing through the centre-tapped primary winding of an rf transformer. Although the overall current flow is from the positive supply rail to the centre-tap ground connection, there is a net alternating flow of current around the circuit consisting of the two 'arms' of the detector, as the beam oscillates from one photodiode to the other. This rf current is coupled into the secondary winding of the transformer. The rf transformer has a winding ratio 1:2 (primary:secondary); the impedance ratio is 1:4, since the impedance ratio is equal to the winding ratio squared. The secondary winding circuit therefore has an impedance of 50Ω, since the primary circuit impedance is 200Ω.

The signal is then amplified using a monolithic rf amplifier, before leaving the split detector module on its way to more amplification and either the analogue amplitude and phase detection (AAPD) system, or a digital oscilloscope for observation.

3.5.5 Calibration and absolute amplitude measurement

To calculate the relationship between amplitude of the surface acoustic waves and the amplitude of the signal emerging from the detector, several things must be known:

- The relationship between the amplitude of the acoustic waves and the angular deviation of the reflected beam.
- The relationship between the angular deviation of the reflected beam and the relative change in the amount of light falling onto each photodiode.
- The amount of light in total falling onto the photodiodes (affected by the different reflectivities of different materials).
- The relationship between the current generated in the photodiodes, and the voltage at the input of the first rf amplifier.
- The gain of the amplifier, and any subsequent gain or attenuation in the signal path.

Spatially, an acoustic wave of wavelength λ_{SAW} has an amplitude A at any given point z is given by:

$$A = A_{SAW} \sin \frac{2\pi z}{\lambda_{SAW}} \tag{3.4}$$

where A_{SAW} is the peak amplitude of the SAW. Differentiating gives the rate of change of the amplitude:

$$\frac{dA}{dz} = \frac{2\pi A_{SAW}}{\lambda_{SAW}} \cos \frac{2\pi z}{\lambda_{SAW}} \tag{3.5}$$

The maximum slope angle of the SAW, θ_{SAW} is therefore:

$$\theta_{SAW} = \tan^{-1}\frac{2\pi A_{SAW}}{\lambda_{SAW}} \tag{3.6}$$

Since the reflected angle is twice the slope angle of the SAW (by geometrical optics), the maximum reflected beam deviation, θ_{beam} is:

$$\theta_{beam} = 2\tan^{-1}\frac{2\pi A_{SAW}}{\lambda_{SAW}} \tag{3.7}$$

The reflected beam is collimated by a lens with a focal length F at this point the beam diameter is D. The angular deviation of the reflected beam is transposed to a lateral deviation d behind the lens:

$$d = F\tan\theta_{beam} \tag{3.8}$$

$$d = F\tan(2\tan^{-1}\frac{2\pi A_{SAW}}{\lambda_{SAW}}) \tag{3.9}$$

It is useful at this point to make some common sense approximations and simplifications. Since θ_{SAW} and θ_{beam} are extremely small angles, equation 3.9 can be re-written:

$$d = \frac{4\pi F A_{SAW}}{\lambda_{SAW}} \tag{3.10}$$

The lens used to reduce the beam diameter so that the beam will fit completely onto both sides of the split detector is inconsequential in the calculations, since d (the lateral deviation) and r (the beam radius) are always proportional, provided that the split detector is not within the focal plane of the beam reduction lens.

The area of the beam as a whole is obviously πr^2 and the area of the beam on one half of the detector (assuming the beam is central) is $0.5\pi r^2$. This generates a photocurrent I. If the beam moves a distance d at the maximum reflected beam angle, then the extra area of one half of the split detector that the beam now occupies is approximately equal to $2rd$ since, as d is several orders of magnitude less than r, the extra area can be considered a rectangle, with dimensions $d \times 2r$.

The extra photocurrent that the extra area of light generates (by more of the beam moving onto the detector in question), δI, is therefore:

$$\delta I = I\frac{2rd}{0.5\pi r^2} = \frac{4dI}{\pi r} \tag{3.11}$$

Substituting for d from equation 3.10 gives:

$$\delta I = \frac{16FIA_{SAW}}{r\lambda_{SAW}} \tag{3.12}$$

This is a useful, practical result, since we know F and r, we either know or can measure λ_{SAW}, and we know I, since this is measured (as the 'dc' light level), as described in section 3.5.4. This therefore gives the maximum change in photodiode current from one half of the detector at the maximum acoustic wave displacement, A_{SAW}.

The increase in current, δI, due to the beam moving more onto one detector, is matched equally by a decrease in current of the same magnitude in the other arm of the circuit. Considering the high frequency component only therefore, there is a net current of $+\delta I$ flowing to the transformer centre tap through one half of the primary winding, and $-\delta I$ flowing from the centre tap through the other half of the primary winding. Since the two currents are equal, the current flowing through the secondary winding of the transformer, δI_t is equal to:

$$\delta I_t = N\delta I \tag{3.13}$$

where N is the winding ratio of the transformer, which in our case is 0.5.

The circuit attached to the secondary winding has an impedance of 50Ω; the peak voltage generated at this point is therefore $50\delta I_t$, or $25\delta I$.

This voltage is amplified by a monolithic amplifier with a voltage gain G, assuming input and output impedance both equal 50Ω.

We now have an absolute relationship between the peak voltage, V_{pk}, and the peak surface acoustic wave amplitude, A_{SAW}, at any given acoustic wavelength, λ_{SAW}:

$$V_{pk} = G\frac{400FIA_{SAW}}{r\lambda_{SAW}} \tag{3.14}$$

Furthermore, the 'dc' light level on each photodiode, represented by a current through each 100Ω resistor, is measured using a transimpedance amplifier with a gain of G_{dc}, such that:

$$V_{dc} = 100G_{dc}I \tag{3.15}$$

Since F and r are known and fixed for our system ($F = 80$mm, $r = 5$mm),

$$A_{SAW} = \frac{V_{pk}\lambda_{SAW}}{V_{dc}}\frac{G_{dc}}{64G} \tag{3.16}$$

This assumes that the frequency response of the photodiode is uniform for both the dc and ac components—this will clearly not be the case for the higher SAW frequencies, where the photodiode response is reduced.

Mathematics proves the existence of a linear relationship between the SAW amplitude and the output of the detector; however in as far as we are interested in imaging, an absolute calibration has not been performed.

3.6 Alignment issues

There are two separate alignment issues for the instrument. The first is 'instrument alignment'—how the instrument is set up from scratch, or after a major change in method of operation, for example when switching the generation method from CGH to SLM—and 'sample alignment', which occurs each time a new sample is examined. These two issues will be dealt with separately in the next two sub-sections. It is important to document the alignment process, since it illustrates the practicalities and complexities of the instrument, the tolerances involved, and information necessary to duplicate the work.

3.6.1 Instrument alignment

The generation laser itself requires alignment of its cavity every so often, due to temperature drift, movement of mirrors or other elements on adjustable mounts within the cavity, or knocks from cleaners placing furniture on top of the laser. The use of the fast-response trigger photodiode coupled to a fast oscilloscope is necessary to tune the temporal response, to get a good Q-switched, mode-locked pulse envelope. The use of a CCD camera (with the infrared filter removed) to observe the spatial properties of the beam (projected onto a screen) is equally important. The use of polarising optics, such as the half-wave plate and Glan-Taylor analysing beam splitter, as described in section 3.4.2, in the path of the laser is particularly important for observing potential mode-hopping and polarisation integrity.

The generation beam diameter must be tuned to the ultrasound profiling method—either a CGH or the SLM. In the case of the SLM, this is simply a case of ensuring that the beam completely covers the entire active region of the SLM—or rather, the entire area of the 'stray light' guard. In either case, this is achieved using the optical zoom system. Ratios of lens focal lengths afford different magnifications of beam diameter, whilst fine-tuning of the inter-lens distance ensures collimation or, in the case of the CGH system, sufficient divergence or convergence to achieve the required line spacing on the sample surface.

For the CGH system, the distance between sample and CGH was adjusted using a measuring gauge of the correct length, and then fine-tuned by fine adjustment of the sample stage until the best signal occurred. A metal ruler, covered in black ink, was also a useful tool for ascertaining the approximate focal distance, since this would ablate at the point where the image was in focus.

The method of getting the image from the SLM in focus on the sample is a little more involved.

The SLM is mounted on a 3-axis adjustable stage, with additional tip-tilt adjustment. The polarising beam splitter, the half-wave plate, the lens and—for convenience—the stray light guard are all fixed, although the beam splitter orientation may be adjusted slightly to ensure orthogonality. Given the focal length of the imaging lens (40mm) and the approximate desired image magnification and/or clearance distance between the imaging lens and the sample (to allow space for the detection system), the approximate distances between the SLM and the lens, and the lens and the sample, can be acquired.

A CCD camera is mounted vertically above the beam splitter, pointing downwards towards a mirror angled such that the intended area of the SLM image can be seen. This chosen plane (a certain distance from the imaging lens) is brought into focus on the CCD camera by adjusting the CCD camera's own imaging lens position, whilst observing a pattern on a diffuse material such as paper. A distinct pattern is uploaded to the SLM and, by reducing the power of the generation laser source by rotating the first half-wave plate, this may be observed by the CCD camera. The distance between the SLM and the imaging lens is adjusted until the image is in sharp focus on the camera. This is easier to see if the material is moved slightly side to side, or rotated, to eliminate the effects of speckle caused by illumination by a coherent source.

At this point, the camera focus and SLM focus are in the same plane. We then position the detection system, such that the detection point is in focus on the diffuse material—this is achieved simply by moving the detector towards and away from the material, and observing the reflected speckle. The detector is in focus when the speckle is at its largest. Once the CCD camera, SLM and detector are all in focus in the same plane, no more adjustment of their relative positions should occur. The diffuse material may now be removed.

Other adjustments and alignments include rotating the second half-wave plate—between the second beam splitter and the stray light guard—until maximum contrast between dark and light portions of the SLM image occur; and ensuring that the reflected detection beam is of a sensible size, and in a sensible place, as it hits the split photodiode detector.

3.6.2 Sample alignment

Each time the sample under investigation is changed, the instrument must be adjusted to ensure that the sample is in the correct place, and to accommodate the material properties of the sample.

The sample must be in focus relative to the the detection point, and hence the in-focus image of the SLM, or in the Fourier plane of the CGH. This is achieved, as described in the previous section, by observing the reflected speckle from the detection beam. Experience has shown that this position is relatively easy to locate, even with highly polished, smooth samples, with very little speckle. Of course, the sample position should be adjusted, rather than the detector position, as this is used as a reference for the generation source focus.

With the majority of materials being tested, the approximate surface wave velocity will be known. It is important to get the excitation source line spacing as close to this velocity as possible, since the amplitude of any generated ultrasound is particularly sensitive to differences between the acoustic wavelength and the line spacing, especially so with excitation profiles that consist of a great number of lines, such as those generated by the SLM. However, for arc excitation profiles—as opposed to straight line profiles—the situation is a little more complex, since the detection may not be within the arc focus, for one reason or another. The ease with which the signal may be optimised and maximised depends on several factors, most notably sample reflectivity, degree of anisotropy—including crystal size and/or angle or rotation, surface finish, attenuation, ultrasound frequency, and the degree of accuracy with which the sample's SAW velocity is known. With skill and practice, the focus of the ultrasound can be located by moving the detection point relative to the generation region, and the velocity matched, within moments. If the velocity is not known very accurately, then by initially using straight line sources and using digital signal averaging on an oscilloscope, the line spacing can be tuned. Once it is known, arc excitation profiles may then be used, and the ultrasound focus located.

3.7 Scanning

The instrument is primarily used to generate—in ultrasonics terms—c-scans. In this configuration, the relative positions of the excitation source and the SAW detection point are kept constant, and the sample is scanned on x-y stages. The data is ideally collected using the analogue amplitude and phase detection (AAPD) technique, described in the next chapter—this is significantly faster than the alternative technique of capturing the data using a digital oscilloscope, since it is several orders of magnitude quicker, which is very significant for a scanning instrument.

The sample stages are controlled by a host computer, as are the stages on which the detection system is mounted. Data is collected during the scanning process at each point, also by the host computer. This data can come from either the digital oscilloscope via a GPIB card, at a maximum rate of approximately one point per second (slower with averaging), or from the AAPD system via an analogue to digital converter card, which will capture every 'event'—one event being one Q-switch envelope. The SLM image may also be controlled during scanning, to account for material aberrations, or for filtering purposes, as described in chapters 7 and 8. A flexible computer program, written in C, and running under the Linux operating system, allows efficient and flexible control of the stages, SLM, and capture of data from different sources. Aspects of this method of control, and future possibilities, will also be discussed in chapter 8.

Chapter 4

High speed amplitude and phase detection

4.1 Introduction

There are several key requirements of a practical imaging system: stability, speed, resolution, accuracy, and cost. Some of these factors (stability, resolution and, to some extent, cost) are determined largely by the optical configuration. The data acquisition technique profoundly affects the speed of the system, and significantly affects the cost. A genuinely useful ultrasound imaging system—one of the key aims of the O-SAM instrument—cannot be realised without a rapid data acquisition system.

This chapter describes the development of such a system. The work described herein is possibly the most important aspect of the development of the O-SAM system contributed by the author.

4.2 Digital data acquisition

In the early stages of the project, data acquisition had to be performed using a digital oscilloscope, for the simple reason that the signal to noise ratio of the detected signal was nearly always less than unity. The digital oscilloscope would sample the detected tone envelope, and digitally average the waveform over a number of pulses (typically 100–1000) from the generation laser. A photodiode detects the source laser tone burst, and the oscilloscope uses this as a coherent trigger source. Once the signal had been averaged, the relevant data (be it the amplitude, phase, or the entire averaged signal) could then be sent to, and recorded on, the host computer. Each measurement point would typically take between two seconds and a minute, depending upon the amount of averaging and the desired resolution. With the advances in the instrumentation of the system, as described in chapter 3, the signal to noise ratio was increased to a point whereby a single shot (i.e. not averaged) signal could be detected on the oscilloscope. However, since these instruments are not optimised for rapid data acquisition, the maximum throughput of data to the host computer was still of the order of two points per second—clearly too slow for routine high resolution imaging.

It is possible to get extremely fast dedicated digital data acquisition systems (as opposed to a multi-purpose digital oscilloscope)—however this was not considered a satisfactory solution for the following reasons.

Firstly is the consideration of cost. Whilst this does not initially appear to be a very relevant consideration given the cost of the system as a whole, it is important because during the development of the instrument in the future, multiple detectors—typically sixteen—will be required (this is explained in section 7.8). Therefore cost considerations are important if multiple detectors are taken into account.

Secondly is the consideration of bandwidth. Although the O-SAM currently routinely operates at 82 and 164MHz, it is expected that the frequency range will be pushed up towards the gigahertz region in the future. With this in mind, it is important not to build into the instrument now something that will limit its capabilities in the future. Although digital sampling at these frequencies is possible, the cost, complexity and speed of the overall system must be taken into consideration.

4.3 Analogue approach

The approach developed has been to use analogue electronics as a rapid, robust, and relatively inexpensive method for data acquisition. In order to successfully design the system, the following points need to be defined:

- The information available to the data acquisition system.
- The amount of available information we are able to extract.
- The information we are interested in.
- The input(s) to the data acquisition system from the previous sub-system: in this case the split diode detector.
  - Signal levels and dynamic range.
  - Frequency content.
  - Timing issues.

- The output(s) of the data acquisition system to the next sub-system: in this case an analogue to digital converter board inside the host computer.
  - How to represent the required information.
  - Signal levels.
  - Timing issues.

These issues are dealt with in the following sections.

4.3.1 Gathering information

It is important to differentiate between *information* and *signals*. Signals are quite specific, and can be defined in terms of how a voltage varies over time in the case of electronic signals. Information, however, is more abstract, and in general refers to attributes of—in this case—an electronic signal. It is perhaps best illustrated by an example.

Consider a signal, such that its voltage V is described by:

$$V = A\cos(2\pi ft + \phi) \tag{4.1}$$

where A is the amplitude, f is the frequency and ϕ is the phase. Given these three pieces of information, we can calculate V at any time t.

If we viewed the above signal on a digital oscilloscope, what we would obviously see is a sinusoidal waveform, acquired by sampling the voltage V at regular intervals to build up a 'snapshot' of the signal. This snapshot may consist of several thousand points of data. However, we are probably not interested in the voltage at certain discreet time intervals. If we wanted to reconstruct this (admittedly simple) signal, we just require the amplitude, frequency and phase information.

This brings about a second issue. The aim may not be to reconstruct the signal, we may just be interested in what its amplitude is. This can quite easily be measured without resorting to high speed sampling. Thus although it may take many thousands of sampled voltages to represent a signal, all that is required is one number that represents the amplitude. Thus the key to gathering information about the detected signal by analogue means is to decide what information is relevant and important, and to only collect that information.

4.4 System requirements

4.4.1 Information available

The characteristics of the detected signal are predominantly determined by the laser source used in generating the ultrasound, and secondly by the interaction of the ultrasound with the sample

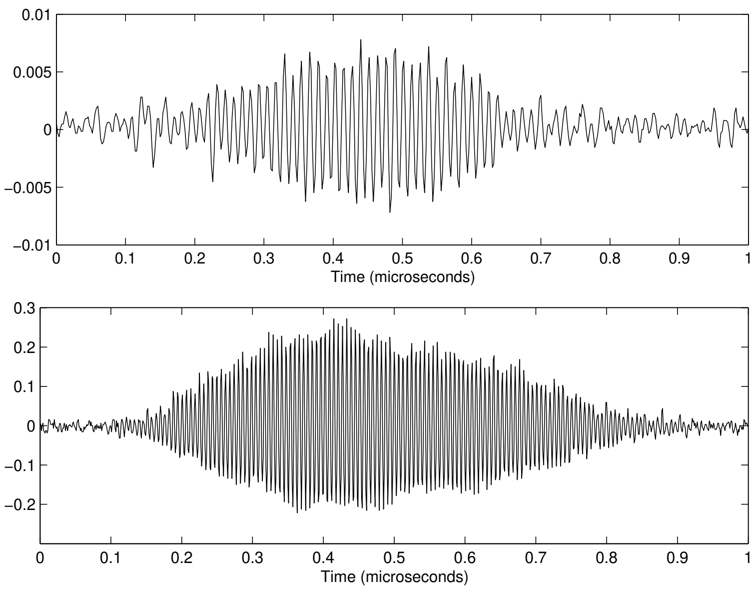

Figure 4.1: A typical 82MHz signal (top) and 164MHz signal (bottom) as received by the detector.

under investigation. As previously stated, the generating laser source has a fundamental frequency of 82MHz and is rich in harmonics. The bandwidth of each harmonic is determined by the Q-switcher characteristics (effectively how fast it switches state) and in our system is around 10MHz. We consider that the fundamental frequency and each harmonic may be treated separately—this is achieved in practice by electronic filtering. This does not make the signal single frequency, but it does make each component narrowband. We cannot therefore simply represent the signal in terms of amplitude, frequency and phase, since it is a little more complicated than that represented by equation 4.1.

4.4.2 Information we can extract

Consider two ‘typical’ detected signals as shown in figure 4.1. The upper signal in the figure has been filtered electronically around 82MHz; the lower signal has been filtered around 164MHz. The SAWs at each frequency have been excited by the 82MHz and 164MHz components of the generation laser excitation pulse, as shown in figure 3.2 on page 35.

We know that the two signals shown have fundamental frequencies of 82MHz and 164MHz. This is very useful, since it is possible to collect amplitude and phase information by analogue means. If we compare the detected signal to a reference signal of the same fundamental frequency, the relative amplitude and phase can be extracted.

4.4.3 Information of interest

The previous two sections have shown that, in general, there are two pieces of information that are of interest to us, namely phase and amplitude. These are two pieces of information that—in the majority of cases—can sufficiently describe the detected signal at a certain point on the sample. More accurately, when we say we wish to measure the phase of the signal, we are actually referring to the difference in phase between the 82MHz (or 164MHz) component of the signal and a coherent reference signal at the same fundamental frequency. The amplitude that we measure is actually the peak amplitude of the detected signal, since this will vary proportionally with the amplitude of any part of the signal—it is also a convenient point to measure.

4.4.4 System inputs

The information of interest in the signal has been defined. We must now consider the aspects of the signal that will affect the design of the analogue amplitude and phase detection (AAPD) system.

There are actually three input signals to the detection system electronics. These are the signal from the knife edge detector (KED) electronics, referred to as 'the detected signal' or simply 'the signal;' a coherent reference signal at the same fundamental frequency, referred to as the 'reference signal;' and a signal from the generation laser trigger photodiode—see figure 3.2, section 3.2.1—referred to as the 'trigger signal.' The following subsections refer to the detected signal.

Signal levels and dynamic range

The detected signal typically has a peak to peak amplitude of between 5 and 100mV as it is output from the KED electronics, at 50Ω impedance. The signal to noise ratio can vary from less than unity (in which case the signal will not be detectable using analogue electronics) to greater than 30.

Signal frequency content

The signal from the KED electronics is fed through an electronic bandpass filter, ideally centred on the frequency of interest—usually either 82 or 164MHz—and of sufficient bandwidth (10MHz) so as not to distort the tone burst profile. In the future, with sufficient optimisation of the optical detector and associated electronics, it is expected that higher harmonics of the 82MHz generation fundamental frequency, such as 328MHz and 410MHz will be detected with sufficient amplitude to be able to use analogue detection techniques.

Timing issues

The time between generation of the ultrasound at the arcs and detection depends on the Rayleigh velocity of the material, and distance between source and detector, which is usually the focal length of the generation arcs. Changes in the velocity of the material are usually very small, leading to very small changes in the position of the tone burst in time, much less than the length of the tone burst. However, if the detection point is scanned relative to the generation arcs, then the detected signal can be moved in time by as much as $\pm$200ns. This situation occurs if you wish to acquire an image of the aerial point spread function.

4.4.5 System outputs

The AAPD subsystem must output the information required (the amplitude and phase) in a form suitable for the next subsystem, which in this case is a multi-channel analogue to digital converter (ADC) board inside the host computer.

Signal levels

The ADC can acquire voltages in the range -10 to $+10$ volts, and has 12-bit accuracy. To maintain accuracy and resolution, the output from the AAPD should make as much use of this dynamic voltage range as possible, without exceeding it.

Timing issues

The particular ADC board currently used is an Eagle Technology PC30G. It has 16 channels of single-ended analogue inputs, and a selection of any or all of these may be sampled (sequentially) at a maximum rate of 100kHz. The sampling may be polled by an external gate—in this case the gate must have a duration of 20μs for each channel to be sampled to be certain of acquiring the voltage on that channel.

These specifications affect the timing specifications of the outputs from the AAPD. The time taken to sample one channel of data (20μs) is clearly much longer than the duration of the tone burst envelope ($\approx$300ns total). This, in addition to the fact that the channels are sampled sequentially rather than simultaneously implies that the voltages representing the amplitude and phase information must be held for a suitable length of time to be certain of it being accurately acquired.

4.5 Analogue phase detection using double balanced mixers

4.5.1 Theory

The standard method of 'extracting' the phase of an analogue signal is to multiply the signal by another signal of the same frequency. Consider two signals, V_{ref} and V_{sig}:

$$\begin{aligned} V_{ref} &= A_{ref}\cos(\omega_{ref}t+\phi_{ref}) \\ V_{sig} &= A_{sig}\cos(\omega_{sig}t+\phi_{sig}) \end{aligned} \tag{4.2}$$

If these two signals are multiplied together, then by trigonometry:

$$\begin{aligned} V_{out} = V_{ref}\times V_{sig} &= \frac{1}{2}A_{ref}A_{sig}\cos[(\omega_{ref}-\omega_{sig})t+(\phi_{ref}-\phi_{sig})] \\ &+\frac{1}{2}A_{ref}A_{sig}\cos[(\omega_{ref}+\omega_{sig})t+(\phi_{ref}+\phi_{sig})] \end{aligned} \tag{4.3}$$

If $\omega_{ref}=\omega_{sig}(=\omega)$, then

$$V_{out} = \frac{1}{2}A_{ref}A_{sig}\cos(\phi_{ref}-\phi_{sig})+\frac{1}{2}A_{ref}A_{sig}\cos[2\omega t+(\phi_{ref}+\phi_{sig})] \tag{4.4}$$

The second half of the result has a $2\omega t$ dependence, and may be filtered out using a low pass filter. If the amplitude of the reference signal is constant, this leads to:

$$V_{out} = kA_{sig}\cos(\phi_{ref}-\phi_{sig}) \tag{4.5}$$

where k is a constant.

The multiplication of two high frequency signals may be easily achieved using a double balanced mixer.

4.5.2 Amplitude and phase response

Equation 4.5 is an important one, and describes the filtered output of a double balanced mixer when used as a phase detector, assuming the reference signal has a fixed amplitude. Figure 4.2 represents the transfer function of a mixer when used in this way.

Two things are noteworthy if this device is to be used as a phase detector:

1. The output voltage is related to the amplitude of the signal, as well as its phase.
2. There is a sinusoidal relationship between output voltage and phase of the signal rather than a linear one.

The first of these points makes the mixer sensitive to changes in the amplitude of the input signal, which initially appears problematic in the case of detecting ultrasound, where the input

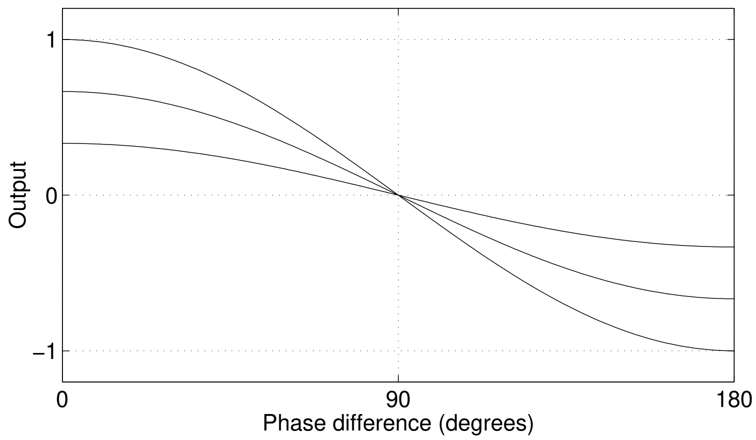

Figure 4.2: The transfer function of a double balanced mixer when used as a phase detector, assuming a sinusoidal reference signal at the same frequency as the input signal. The graph shows the output voltage with respect to phase difference for three different input signal amplitudes.

signal has a huge dynamic range. The method by which this amplitude relationship is usually removed is by saturating or limiting the input signal, so that its amplitude is constant.

The second of these points shows that there may be problems in extracting the phase of the signal due to the nonlinear relationship between voltage and phase, the insensitivity to changes in phase around the 0° and 180° points, and the fact that the output voltage is not unique to any given phase between 0° and 360°, only between 0° and 180°, and between 180° and 360°.

All of the problems described above can be overcome by using two mixers in quadrature.

4.5.3 Using two mixers in quadrature

The sinusoidal variation of output voltage with respect to input phase is actually very useful if two mixers are used in quadrature. This is achieved by applying the same reference and input signal to two mixers, and delaying either the reference or the input signal by 90° in one of them; see figure 4.3.

The outputs of the two mixers are therefore

$$\begin{aligned} V_{out1} &= kA_{sig}\cos(\phi_{ref} - \frac{\pi}{2} - \phi_{sig}) \\ V_{out2} &= kA_{sig}\cos(\phi_{ref} - \phi_{sig}) \end{aligned} \tag{4.6}$$

The two outputs (V_{out1} and V_{out2}) now represent the real and imaginary parts of the complex amplitude of the 82MHz component of the input signal (V_{real} and V_{imag}). Since ϕ_{ref} is the same

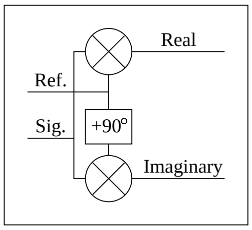

Figure 4.3: Block diagram showing how two mixers may be used in quadrature.

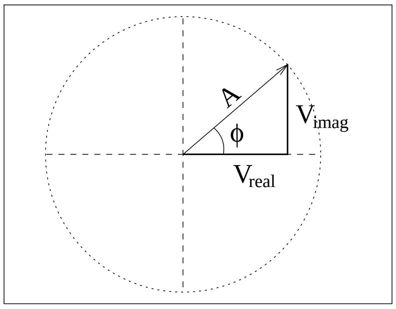

Figure 4.4: The *complex amplitude* of a signal consists of a real part and imaginary part. From equations 4.8 and 4.9 one is able to extract the absolute amplitude and phase.

in both cases, equation 4.6 may be re-written:

$$\begin{aligned} V_{real} &= kA_{sig}\sin\phi_{sig} \\ V_{imag} &= kA_{sig}\cos\phi_{sig} \end{aligned} \tag{4.7}$$

These real and imaginary components may be used to extract the amplitude (A) and phase (ϕ) of the input signal:

$$A = \sqrt{{V_{real}}^2 + {V_{imag}}^2} \tag{4.8}$$

$$\phi = \arctan\frac{V_{imag}}{V_{real}} \tag{4.9}$$

This is illustrated in figure 4.4.

4.6 System design: overview

The method of using two double balanced mixers in quadrature to acquire the amplitude and phase of the detected signal has been described. Due to the system output requirements as described in section 4.4.5 the outputs of the mixers cannot be directly used. This section describes the necessary subsystems of the AAPD required to attain the amplitude and phase information.

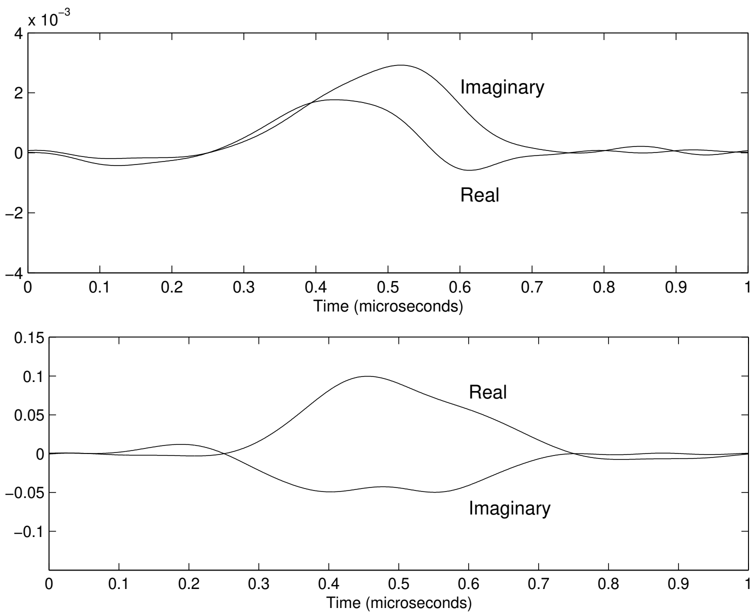

Figure 4.5: Theoretical outputs of two mixers operating in quadrature for two typical input signals

4.6.1 Mixer outputs

The filtered output characteristics of two mixers operating in quadrature have been defined in equation 4.7, and it is instrumental to investigate the theoretical outputs of the two mixers for typical input signals. This will precipitate the design of the electronic systems following on from the mixers, to produce the required outputs as described in section 4.4.5.

Figure 4.5 shows the real and imaginary mixer outputs for the signal in figure 4.1. The reference signal in this case is a pure 82MHz signal of constant amplitude.

We note that the mixer outputs follow the general envelope shape of the tone burst, as expected by the proportionality between output voltage and input signal amplitude. The polarity and precise shape of the output envelope is determined by the phase of the signal relative to the reference frequency - if this phase is constant for the duration of the tone burst, then the output envelope will follow the tone burst envelope precisely. The implication of this is that if we measure the output voltage at the peak of the envelope for each of the two 'arms' of the system (the real and imaginary parts) then we will be measuring the amplitude and phase of the signal at the same point in time.

The peak amplitude of each mixer output is proportional to both the amplitude of the detected signal and the reference signal. The value of the constant of proportionality—k in equation 4.7—for a typical mixer operating linearly is of the order of 0.8. This implies that the maximum output

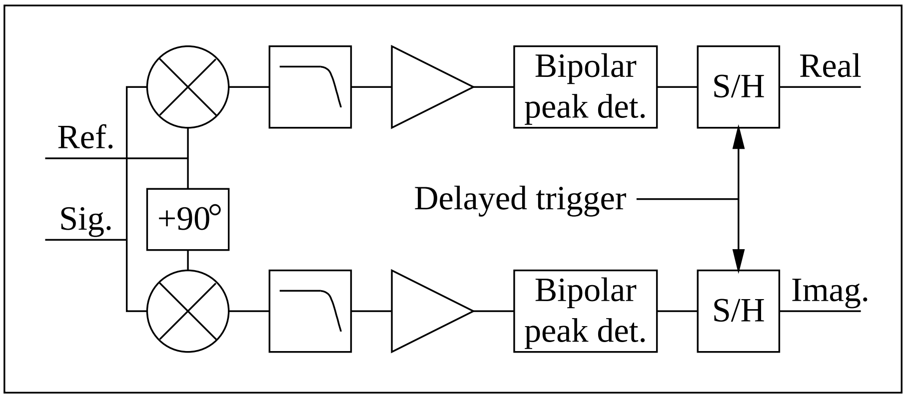

Figure 4.6: Analogue amplitude and phase detector—block diagram of various system components.

voltage from each mixer will be of the order of ±80mV. The signal has a bandwidth from d.c. to around 10MHz.

4.6.2 The need to detect and hold the peak

If we sample the output from the two mixers simultaneously at the peak of the output envelope then we will have the necessary information to acquire the phase and amplitude of the detected signal. Unfortunately we do not know when in time (relative to the trigger pulse) the peak of the signal will occur, for the reasons described in section 4.4.4. In addition, the particular ADC board used to acquire the voltages for the host computer would neither be able to sample the voltage with anywhere near enough timing accuracy, nor can it acquire the voltages simultaneously. For these reasons, the *peak* of the output envelope from each mixer must be detected and held for a duration sufficiently long enough for the ADC to acquire both voltages. Since each mixer output is bipolar, the peak detector must also be bipolar, i.e. it must detect either the maximum or minimum voltage, depending on whether the envelope is positive or negative.

4.6.3 Analogue amplitude and phase detector—block diagram

With the considerations of the previous two sections taken into account, the figure 4.6 is a block diagram of the proposed system for acquiring the amplitude and phase of the detected signal.

The reference signal is split, with one arm delayed by 90°; the detected signal is split without any delay. Thereafter, each arm of the system is identical. The reference and detected signals are mixed, and the output filtered with a 10MHz low pass filter. The signal is then amplified, and its peak is detected by a bipolar peak detector. Finally a sample and hold device ensures that the detected peak voltage is held for a sufficiently long time to enable accurate capture by the ADC.

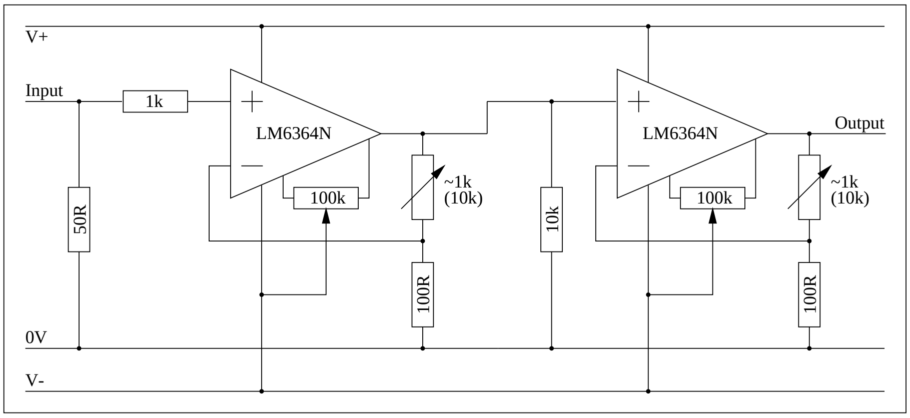

Figure 4.7: Operational amplifier circuit. Two op-amps are used in series, due to the bandwidth of the signal and high gain required.

4.7 System design: electronic design of AAPD subsystems

The design of the electronics to perform the functions of all the subsystems identified in the previous section is far from straightforward. The signal as it leaves the low pass filter has a bandwidth that goes all the way from d.c. to 10MHz. Whilst 10MHz is not a particularly high frequency as far as r.f. electronics is concerned, a bandwidth that large that must include the d.c. component can be troublesome, due to the accumulation of noise, possibility of offset voltage drift, the need for d.c. coupling along the length of the signal path, and the necessity for all components to operate linearly within that bandwidth. If these considerations are not taken into account then it is very easy to design a system that is either unstable, or does not meet the requirements for accuracy and certainty.

4.7.1 Amplification

The *maximum* output of each mixer was described in section 4.6.1 as being ± 80mV—this corresponds to a detected signal with a large amplitude and a phase of either 0° or 180°. For different phases, and detected signals of a smaller amplitude, the output envelope will be smaller. If the peak of the envelope is not accurately held, this will correspond to an error in either the measured amplitude, phase, or both. It is therefore useful to amplify the signal after the low pass filter, to make use of the full voltage range available from the power supply rails.

This is achieved in practice by two high speed operational amplifiers used in series, each having a gain of around 10–12; see figure 4.7. The first op-amp has an input impedance of 50Ω to match the output impedance of the mixer/filter. The second op-amp has an input impedance of $10\text{k}\Omega$,

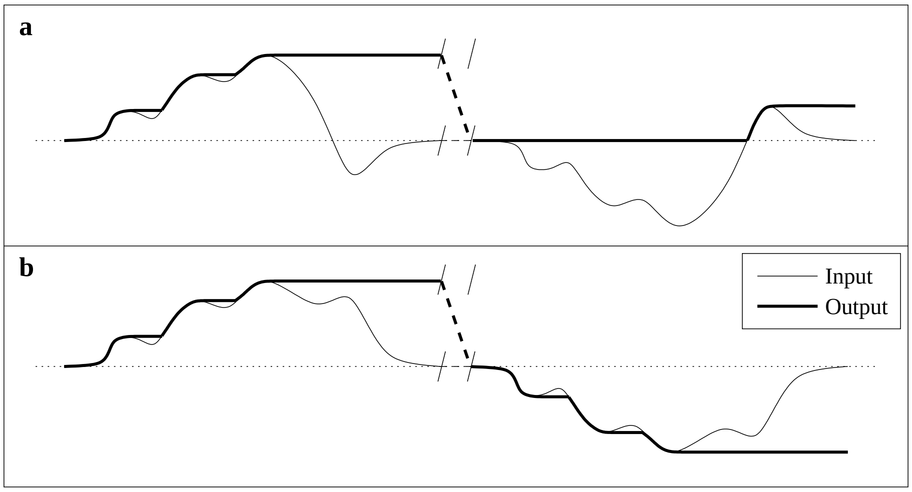

Figure 4.8: Expected outputs from **(a)** a unipolar peak detector, and **(b)** a bipolar peak detector.

so that the first op-amp does not have to work as hard driving a low impedance.

The op-amps used are National Semiconductor LM6364N's. They have a gain-bandwidth product (GBWP) of 175MHz, and are stable down to a gain of 5. Offset null adjustment is available and is used, since any d.c. offset present here will propagate through the rest of the circuit and result in significant measurement errors.

4.7.2 Bipolar peak detection

When we say we wish to 'detect' the peak of a signal, what we actually mean is that we would like an output that follows the input precisely if the value of the input is greater than, or equal to, any previous value of the input. If the input is less than the greatest value attained so far, we would like the output to remain fixed at this large value, until such a time as the system is reset: see figure 4.8a.

This is how a unipolar peak detector works. A *bipolar* peak detector, however, would detect either the greatest value, or the smallest, depending on the polarity of the input signal: see figure 4.8b. Obviously a bipolar peak detector will not be able to cope with inputs that go both above and below the initial starting value, since its output would be undefined.

General approach

The general approach to making a bipolar peak detector in practice is to use two unipolar peak detectors—one of which detects the maximum voltage, the other detects the minimum voltage—and summing their outputs. This works well provided that the input voltage only goes either above

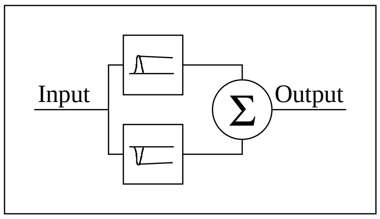

Figure 4.9: Realisation of a bipolar peak detector by using two unipolar peak detectors of opposite polarities and summing their outputs.

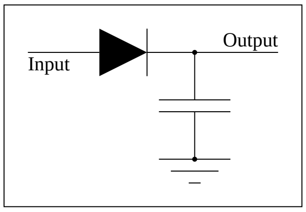

Figure 4.10: Simple unipolar peak detector.

or below the initial starting voltage, which is usually zero volts. This is the approach used in the AAPD bipolar peak detector subsystem: see figure 4.9.

The effect of diode forward voltage drop on measurements

The general principle of peak detection as described in section 4.7.2 is realised in analogue electronics by the input signal charging a capacitor through a diode, which prevents the capacitor from discharging when the signal drops to below its maximum: see figure 4.10.

This has several disadvantages: variable input impedance, susceptibility to low impedances following on from the circuit and, most notably, the diode forward biased voltage drop makes the circuit insensitive to input voltages of less than around 0.6V, and is inaccurate for voltages greater than this.

This has a profound effect on the accuracy of the system, particularly with regard to phase errors in the AAPD system around the quadrature points, even if low voltage drop germanium diodes are used. One arm of the system will have a large amplitude (and hence a reasonably low error), but the other arm will be at zero volts until the voltage rises above the diode voltage drop—the angle at which this occurs is:

$$\phi = \arctan\left(\frac{d}{A-d}\right) \tag{4.10}$$

where ϕ is the phase error, d is the diode drop, and A is the amplitude of the input voltage. The

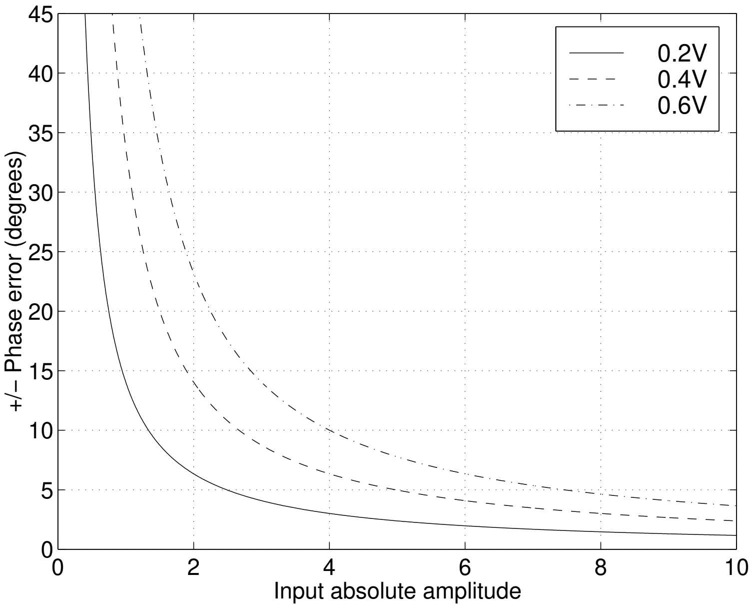

Figure 4.11: The maximum phase errors for a given input signal amplitude for three different diode voltage drops.

maximum phase errors for a range of diode voltage drops and input amplitudes are plotted in figure 4.11.

Given the possible dynamic input range of the input signal given in section 4.4.4 it is obvious that even small voltage drops lead to unacceptably high phase errors for small signals.

Elimination of diode forward voltage drop

This problem is overcome by placing the diode within the feedback loop of an op-amp voltage follower: see figure 4.12. The voltage on the charged capacitor is isolated from any following circuitry by a second op-amp voltage follower.

Whilst this circuit overcomes the problems of input and output impedance and the diode voltage drop, it is now affected by the limitations of the first op-amp. The three main limitations are maximum slew rate, maximum output current, and input bias current.

Consider that the output of the op-amp is in negative saturation for all the time when the input is below the peak output. When the input next rises above the peak, the output has to swing all the way from the negative power rail to follow the input voltage again—hence the maximum slew rate actually needs to be much higher than expected with a peak detector used in this configuration. The output current drive capabilities affect the rate of change of voltage across the capacitor, which again affects the rate at which the output can follow the input. The input bias current of

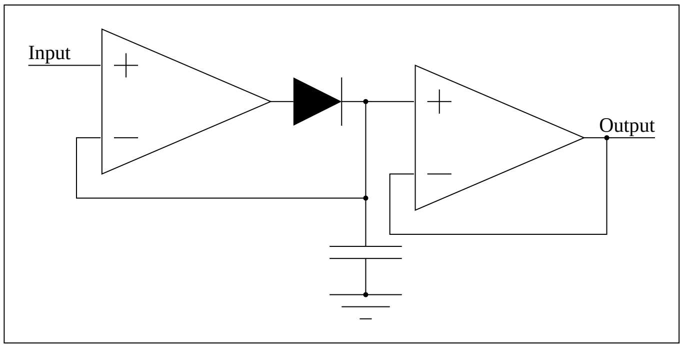

Figure 4.12: Modified version of the simple peak detector shown in figure 4.10. The diode is within the feedback loop on an op-amp, thus overcoming the forward voltage drop.

the op-amp affects the amount by which the output ‘droops’ as the capacitor is discharged, as does the input impedance of the following op-amp. Using a capacitor with a relatively high value reduces the droop, but also reduces the maximum slew rate.

The problem of the op-amp output being in negative saturation for much of the time is a serious one when a high bandwidth (10MHz) is required. A partial solution to this particular problem is to clamp the negative output of the op-amp to prevent it going to the negative power rail. This is achieved with the circuit in figure 4.13, which is an active half-wave rectifier circuit from Horowitz and Hill [70] modified to be a peak detector. It is the circuit actually used in the AAPD system. The circuit shown detects negative-going peaks. Its complement detects positive-going peaks—the two diodes are reversed and V_{ref} is set to an equal but opposite value.

The output of the circuit is actually the inverse of the peak of the input. The circuit works as follows. When the input voltage is below V_{ref}, the circuit acts as an inverting amplifier, with the gain set by R_f/R_{in}—in this case 1. Capacitor C is charged by the op-amp through diode D_1. When the input voltage rises, the op-amp cannot sink the current to equalise the voltages on its two inputs because D_1 is back-biased, and so the charge remains stored on the capacitor. The op-amp output is clamped to one diode voltage drop below V_{ref} by D_2. Hence, when the input next goes more negative than the previous minimum, the op-amp output only has to swing from one diode drop below V_{ref} to the peak voltage.

The disadvantage of this circuit compared to the simplified version in figure 4.12 is that the capacitor C is now able to discharge through resistor R_f to the ‘virtual’ voltage V_{ref} at the op-amp's inverting input. Unfortunately this feedback resistor value cannot be made arbitrarily high—to quote the National Semiconductor LM6171 data sheet:

> Large values of feedback resistors can couple with parasitic capacitance and cause

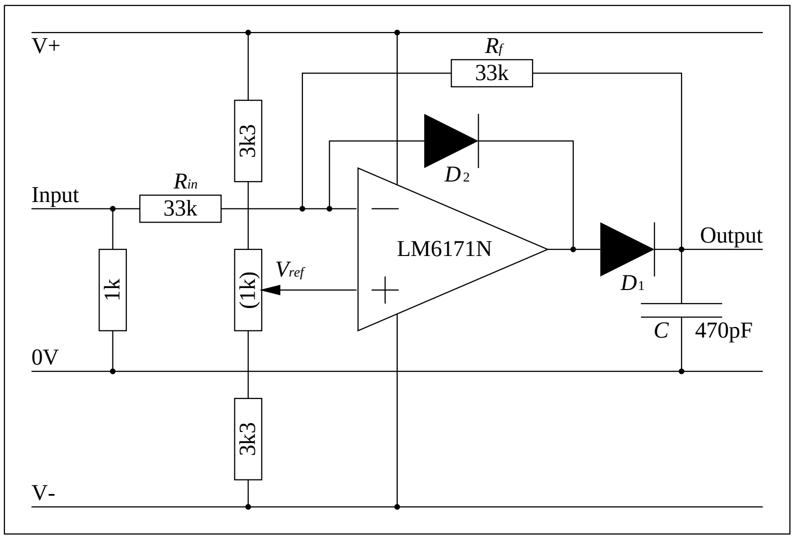

Figure 4.13: Modified version of the diode-compensated peak detector shown in figure 4.12. The op-amp output is clamped to one diode voltage drop below zero volts when the input is less than the peak output. This reduces the maximum slew rate required from the op-amp.

undesirable effects such as ringing or oscillation in high speed amplifiers. For LM6171, a feedback resistor of 510Ω gives optimal performance.

As ever the choice of components used comes down to compromise. Ideally we require an op-amp with a very high slew rate, high output current to charge the capacitor, very low input bias current (to reduce droop) that is stable for high value feedback resistors and happy to drive a capacitive load.

The National Semiconductor LM6171 was found to be the best compromise for this task. Its main attractive feature is a high output current drive capability, which is 100mA maximum. To follow a 10MHz bandwidth envelope from the previous subsection of 10V amplitude requires a slew rate in excess of $100\text{V}\,\mu\text{s}^{-1}$. The rate of change of voltage across a capacitor is given by:

$$\frac{dV}{dt} = \frac{I}{C} \tag{4.11}$$

which implies that the capacitor must have a value of less than 1nF. If a value of 470pF is used, and with an input bias current of 4μA, the rate of discharge of the capacitor is $8.5\text{mV}\,\mu\text{s}^{-1}$. The op-amp was found to be stable with $R_f = 33\text{k}\Omega$, which gives an RC time constant of 15.5μs. The op-amp's quoted maximum slew rate of $3600\text{V}\,\mu\text{s}^{-1}$ is more than sufficient.

A voltage divider network is used to define V_{ref} rather than setting it to zero volts for two reasons. Firstly, the op-amp does not have any input voltage bias adjustment capabilities. Secondly,

V_{ref} is actually set to around +100mV, so that diode D_1 is always at least partly forward-biased (unless the input is a positive pulse, in which case the complementary peak detector will be active) - this reduces some of the non-linearities of the variable impedance of the forward-biased diode, which the op-amp cannot completely cope with. Due to setting V_{ref} in this way means that the output will be 100mV greater than the inverse of the input; however, the complementary peak detector's output will always be 100mV less than the inverse of the input, so it all evens out.

Summing amplifier

The outputs from the two peak detectors described above are applied to the inputs of a standard op-amp summing amplifier, whose output is actually the inverse of the sum of the two inputs. As described in the previous section, the droop rate of capacitor C partly depends on the resistance to ground, and so the summing amplifier has very large input impedances.

A feedback resistor R_f is adjustable to trim the gain of the circuit, to ensure that the peak output is always less than 10 volts for the maximum input signal to avoid clipping when the voltage is acquired by the ADC. V_{ref} is also trimmable to ensure that the output is zero volts for no input signal. It is important to 'balance' the two arms of the system (real and imaginary outputs) and to remove any d.c. offset as it has a profound effect on errors.

4.7.3 Sample and hold

The RC time constant of the peak detector is of the order of 13μs, in addition the output droops at a constant rate of around 8mV μs^{-1}. The droop/discharge rate cannot realistically be made any lower given the current limitations of op-amp capabilities. However, it is still too short for the ADC inside the host computer to sample directly, as there is a timing uncertainty of 20μs, and two channels to acquire sequentially.

Two possible solutions are to either follow the outputs from each unipolar peak detector with a further peak detector with a lower bandwidth and a much reduced decay rate, or to sample and hold the output of the output of the bipolar peak detector subsystem. The first option has the advantage of not requiring any timing signals, making it completely insensitive to the position of the detected signal with respect to the trigger signal. The second option has the advantages of convenience, reduced susceptibility to noise, and the benefit that since both arms of the system will be sampled at the same instant, there is no error due to sampling the two channels sequentially.

Since the sample and hold method is preferred for reasons of simplicity and stability, it is necessary to examine the amplitude errors due to the movement in time of the detected signal with regard to the trigger signal. With the O-SAM operating in normal c-scan mode, this is a fraction of the tone burst duration, since velocity changes are generally small. However, as noted in section 4.4.4, when aerial point spread functions are acquired, the detected signal can move by as much as $\approx \pm 200$ns (assuming a Rayleigh wave velocity of 3000ms^{-1} and a scanning distance of 1.2mm). The fractional error in voltage from one side of the scan to the other is therefore:

$$e^{\frac{-t_1}{RC}} - e^{\frac{-t_2}{RC}} \tag{4.12}$$

where t_1 and t_2 are 400ns apart. The error varies slightly depending where on the decay curve we acquire the voltage; ideally it is as close to the start as we can be certain of acquiring the peak, allowing for changes in velocity and timing jitter, and is in practice around 500ns after the detected pulse. This gives a voltage amplitude error of 2.9% (from one side of the aerial PSF to the other).

In the c-scanning mode, even if the velocity changes by 10% (quite a large amount), this corresponds to a timing error of $\approx$70ns over 2mm (a typical surface wave propagation distance from generation arcs to detection point) on aluminium, which corresponds to a maximum amplitude error of 0.5%.

Given that these errors are both small and predictable, the benefits of using a sample and hold device to get round the problem of short peak hold time outweigh the disadvantages, and is therefore the chosen method.

The sample and hold device used is the popular LF398 device. Acquisition time, droop rate and hold step are the main trade-offs in the selection of hold capacitor value. With the capacitor used, acquisition time is $\approx 4\mu$s to within 0.1% of the final value, and droop rate is of the order of a few millivolts per millisecond.

The 'trigger' signal for the sample and hold device is derived from the photodiode trigger signal, and is provided by an external commercial pulse generator at present—although there is no reason why the equivalent cannot be constructed in-house and incorporated into the final design. The sample and hold trigger signal—the so-called 'delayed trigger'—is delayed by a time set manually depending on the velocity of the sample under investigation and the surface wave propagation distance; it is typically 0.3–1μs. This delay time corresponds to the time it takes the surface wave to propagate from the generation arcs to the detection point, plus a few hundred nanoseconds to allow for delays in the electronics and to be sure of capturing the peak of the detected envelope. The hold period can be as long as required, although it must obviously be shorter than the time between successive tone envelopes (determined by the Q-switch repetition rate).

4.7.4 Miscellaneous circuit construction issues

Due to the reasonably high bandwidth of the signals propagating through the analogue amplitude and phase detector, care must be taken with construction and layout, as well as choice of

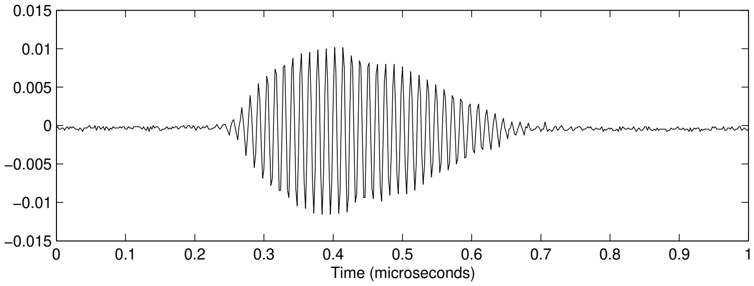

Figure 4.14: A typical simulated detector signal used to test the AAPD.

components, to minimise the risk of oscillations and non-linearities.

All the op-amps used in the AAPD are made by National Semiconductor, with different models being used for different parts of the circuit according to specifications. They are all made using a so-called VIP ('vertically-integrated PNP') process—the upshot of which is that they are fast, stable, and are genuine voltage-feedback amplifiers, and are not restricted to small power supply ranges.

All active components are decoupled using good quality ceramic capacitors, usually 100nF in value, at each power supply pin, with additional 10μF tantalum capacitors nearby.

4.8 System performance

In the preceding sections concerned with the electronic design of the analogue amplitude and phase detector, several theoretical sources of error and uncertainty have been discussed, which has affected the eventual design of the system. It is obviously important to assess the final design in a practical sense. To this end, the performance of the AAPD was compared to that of a digital oscilloscope for the entire range of signals (in terms of amplitude over the entire dynamic range and all phases) that the system would be subjected to.

An 82MHz continuous reference signal was provided by a signal generator. This is equivalent to the 41MHz coherent reference signal provided by the mode locker, after it has been doubled and filtered. A simulated "detected signal" was generated by mixing a pure 82MHz sine wave with a pulse resembling the envelope of a typical detected signal. The envelope was provided by a commercial pulse generator, and had a variable repetition rate (resembling the repetition rate of the Q-switcher). A typical simulated signal is shown in figure 4.14.

The test setup is shown in figure 4.15. The reference and simulated signals were applied to both a digital oscilloscope and the AAPD system in parallel. The two signal generators (each

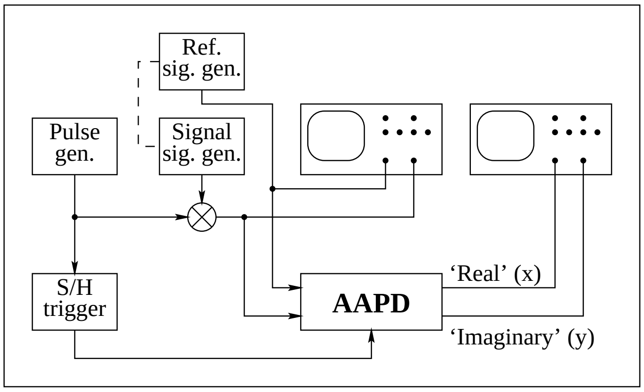

Figure 4.15: Test equipment used to test the amplitude and phase response of the AAPD.

generating a pure 82MHz tone) were incoherent with each other, and due to slight differences in frequency (of the order of 10Hz or so) the relative phase between the two was constantly changing. The simulated signal pulse repetition frequency was set to a rate of around one pulse per second. This ensured that the two oscilloscopes—one of which recorded the AAPD inputs, the other of which recorded the AAPD outputs—both recorded the same event. The amplitude of the reference signal was kept constant, and the amplitude of the 'simulated signal' signal generator—and hence the simulated signal—was varied to test the accuracy of the system over the entire dynamic range.

The measurement equipment and signal generators were controlled via GPIB by a host computer, substantially automating the test process.

4.8.1 Amplitude response

The amplitude of the input signal is proportional to the absolute amplitude, derived from the voltages of the real and imaginary arms of the AAPD system, as described in equation 4.8. The constant of proportionality is determined by many factors:

- The gain in any amplification prior to the mixers.
- The specifications of the mixer (known as the figure-of-merit, M, or efficiency).
- The gain in the amplifiers following the mixers.
- Any gain in the bipolar peak detector circuit (for instance at the summing amplifier).
- The droop rate of the bipolar peak detector, with respect to the point at which the sample and hold device is triggered.

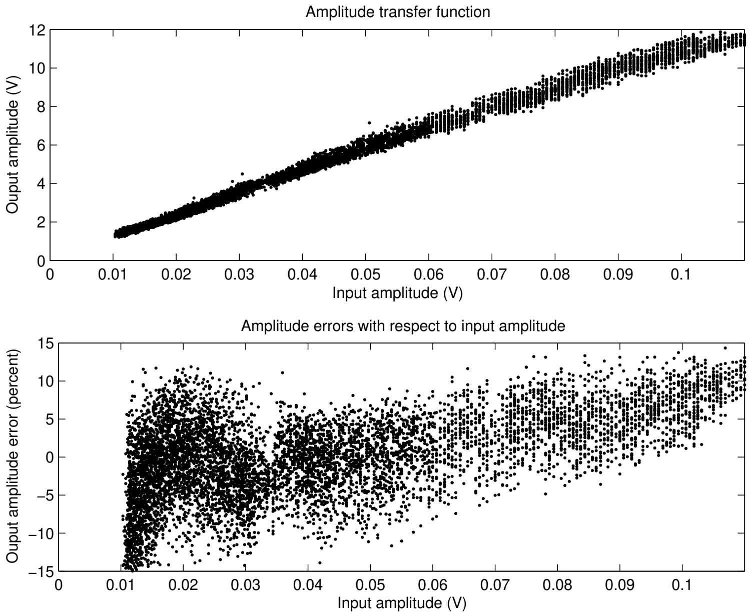

Figure 4.16: The amplitude transfer function (top) and output amplitude errors with respect to input amplitude (bottom).

If these specifications are known, then the theoretical system gain can be calculated. In practice, it is much easier to do analytically, since the system is optimised to use the entire dynamic range of the system following the AAPD, given the dynamic range of the input signal.

The upper plot of figure 4.16 shows the amplitude response of the system. The phase for all the points was random. Ideally the transfer function would be a straight line, implying a purely linear response. The percentage error from the expected output, given a perfect response, can be calculated, and this is shown in the lower plot of figure 4.16.

The error distribution is significantly greater at the lower input amplitudes, as might be expected in an analogue system subject to noise. Overall the errors are within ±15% for an input amplitude dynamic range of 25dB. Although the uncertainty is lower for the higher amplitudes—around ±8%—there is clearly a slightly nonlinear amplitude response, denoted by the upward trend in the lower plot of figure 4.16. This is almost certainly caused by the output current drive capabilities of the op-amps in the bipolar peak detection circuit. By limiting the dynamic range of the input signal, the nonlinearity can be minimised at the expense of increased uncertainty.

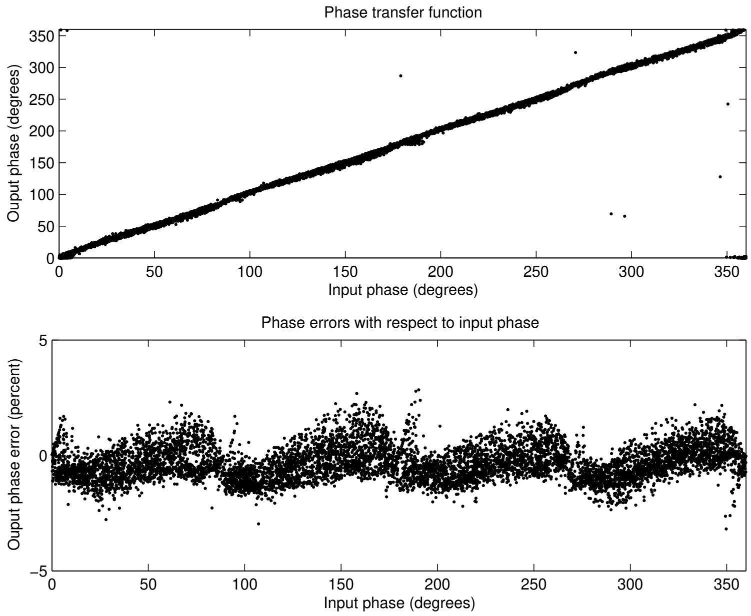

Figure 4.17: The phase transfer function (top) and output phase errors with respect to input phase (bottom), represented as a percentage.

4.8.2 Phase response

The phase transfer function is shown in the top part of figure 4.17, and the phase errors with respect to input phase are shown in the bottom part of figure 4.17. The data points represent random input amplitudes, in the same 25dB dynamic range as the amplitude response above.

The phase response is interesting, since the error is clearly dependent on the input phase. The errors have 90° symmetry, and this is due to the problems of diode forward voltage drop, as discussed in section 4.7.2. Nevertheless, the results of the phase errors are impressive, with the error consistently being less than ±2% for the entire input dynamic range.

4.9 Summary

This chapter has dealt with the concept of the difference between signals and information, and the most useful information available from the O-SAM instrument is the amplitude and phase (relative to the excitation source) of the detected ultrasonic signal, at the fundamental frequency of 82MHz and its harmonics. A high speed analogue amplitude and phase detection system has been described that can obtain this information for every single pulse envelope, provided the SNR of the input signal is greater than one. The system is not limited by SAW frequency, unlike an all-digital system: provided the mixers can produce the real and imaginary pulse envelopes, the amplitude and phase can be obtained. The system is also very inexpensive to construct, which is very important if multiple channels at multiple frequencies are to be constructed. Although the uncertainty and errors in the amplitude response is disappointing, it is important to note that the phase response is within $\pm 2\%$ for the entire input dynamic range. This implies that amplitude/phase crosstalk is very low. Furthermore, a $\pm 2\%$ error in phase corresponds to an error in velocity measurement of just 0.02% when measuring velocities over a distance of 2mm at 164MHz on typical engineering materials such as aluminium.

Chapter 5

Imaging using Rayleigh waves

5.1 Introduction

The O-SAM instrument, and corresponding data capture technique has been described in the previous two chapters. This chapter, and the next, describe some of the imaging capabilities of the instrument in terms of using different contrast mechanisms to either measure or infer various material properties, using either Rayleigh waves (this chapter) or Lamb waves (chapter 6). Several examples of the different sorts of contrast mechanisms will be presented in both chapters. The very act of imaging, which involves acquiring lots of information either over a large area, or in very fine detail, often brings to light extra information or insight into different acoustic mechanisms that might otherwise be missed with simple point-by-point or line-scans. Examples of this are presented in the following two chapters and chapters 7 and 8, which describe some of the issues arising from use of the instrument, and ideas for further contrast mechanisms that have not yet been exploited.

The various contrast mechanisms that the instrument is capable of will be described and illustrated by way of example images obtained by the O-SAM instrument.

5.1.1 A note about surface finish

As described in section 2.3.5, the knife-edge detection technique used by the O-SAM instrument is sensitive to the surface finish of the material under investigation, specifically the roughness of the surface. The practical criterion for the instrument to successfully produce acoustic images is that the material should be sufficiently smooth such that 'you can see your face in it.' Although it is possible to image rougher materials, the image quality is severely degraded due to the amount of optical speckle.

All the acoustic images in the following chapters are taken from samples that were either naturally smooth enough to permit good quality imaging, or have been polished.

5.2 Some properties and uses of Rayleigh waves

Some of the basic properties and uses of Rayleigh waves have already been touched upon in chapters 1 and 2. A relatively early and comprehensive review of the physical theory of Rayleigh waves, and some applications, can be found in Viktorov's book [10]. They propagate at the boundary between—usually in terms of NDT—a solid half-space and air. Their amplitude characteristically decays with depth, and so they may be used to probe the surface or near-surface properties of materials. They are nondispersive, and their velocity is dependent on the density and elastic constants of the solid.

Their interaction with edges [71], slots [72], and cracks and discontinuities [73, 74, 75] have been extensively studied. The scanning acoustic microscope (SAM) has already been described in section 1.4.2; Briggs [76] wrote a good summary in 1991 of the SAM's use of changes in Rayleigh wave velocity to probe the elastic properties of the surface, for instance in media with a surface layer, or anisotropic materials. Work has been done relatively recently [77] to achieve *quantitative* imaging of Rayleigh wave velocity with a SAM. For many years there has been much interest within the NDT community in measuring the Rayleigh wave velocity to measure [78, 79] or image [80] the surface or near-surface residual stress in engineering materials. Changes in velocity may also be used for thin-film characterisation [45]. The orientation of single crystals may be determined by SAW velocity measurements [16, 81] and work is described in section 8.2 using the O-SAM instrument to perform rapid measurement of velocity on single crystals [82].

Finally, no introduction to the properties and uses of Rayleigh waves would be complete without mentioning SAW devices [83, 84]. Traditionally, these devices consist of a series of interdigital metallic lines fabricated on the surface of a polished piezoelectric plate. There may be two or more of these interdigital transducers (IDTs) on a SAW device, and these serve the dual role of excitation/detection and time/frequency signal processing. It is possible to make—amongst other things—bandpass filters, matched filters, delay lines, autocorrelators, and chirp filters. The range of applications in electronic devices for SAW filters is truly staggering [84].

5.3 Detection of cracks and discontinuities on isotropic media

A 'classic' example of a practical use of ultrasonics is flaw detection, and this is certainly possible with the O-SAM instrument. Considering first of all cracks and discontinuities, the O-SAM is

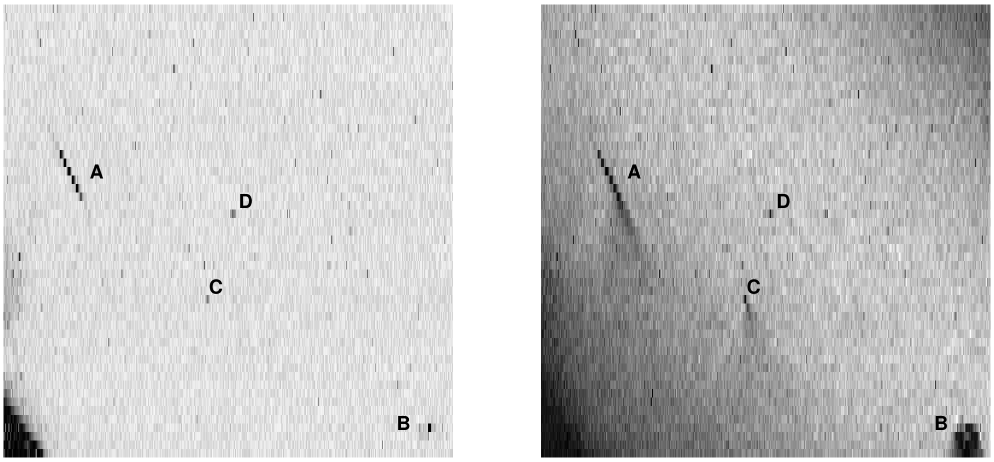

Figure 5.1: Optical (left) and 82MHz SAW amplitude (right) images of a block of silicon nitride, which contains a scratch, and small and large Vickers hardness test indentations. Image size is 10.5×10.4mm, resolution is $20\times200\mu$m

ideally suited to detecting features within a wavelength or so of the material surface, since this is the penetration depth of Rayleigh waves, and is thus the region giving active contrast.

5.3.1 Large-area scans for flaw detection

Surface—or near-surface—cracks act as a barrier to the progress of surface acoustic waves. The waves are reflected off the cracks, either completely, or partially. The amplitude of the waves behind the cracks is either zero if the crack is deep and wide, or attenuated if the crack is thin and/or shallow. If some of the SAW energy gets past the crack, then there will usually be an associated phase shift.

These large changes in detected wave properties are useful to detect flaws, since they introduce a form of contrast that does not exist optically. Figure 5.1 demonstrates this.

The two images in the figure are of the same area on a sample of silicon nitride—in fact the images were acquired simultaneously. The left-hand image is an 'optical' image, which was acquired by measuring the sum of the amount of 'dc' light falling onto each half of the split photodetector. Although of obviously lower quality than images that may be acquired with a standard microscope, these sorts of optical images—acquired by the O-SAM in parallel with the acoustic images—are useful for analysis.

The image on the right is an '82MHz amplitude image;' in other words the intensity of the pixels is proportional to the peak amplitude of the 82MHz component of the propagating surface acoustic waves at the detection point. In this image, as is the case with all c-scans within this thesis unless explicitly stated otherwise, the ultrasound source is to the top of the page, and the direction of SAW propagation is downwards.

The image is of a reasonably large area (considering that the 'M' in O-SAM stands for 'microscope'), being 10.5×10.4mm. The resolution is a very coarse $20\times200\mu$m, however this implies that the c-scan may be performed extremely quickly—the type of scan that could be performed to check for the *presence* of faults, rather than for the *characterisation* of faults.

Four 'features' of the images are labelled A, B, C and D in figure 5.1. Feature A is a coarse, deep scratch in the surface of the material, produced using a diamond tip. By examination under a microscope, the crack is approximately 1mm long by 0.05mm wide. This feature is apparent in both the optical image and the SAW amplitude image—on the optical image as a line of dark pixels, and on the amplitude image as a dark region beneath the scratch.

Feature B is a large *Vickers hardness test* indentation. A Vickers hardness test is a way of ascertaining the hardness, or density of a material by applying a certain known force to an inverted diamond pyramid onto the material, and measuring—usually with a microscope—the distance between opposite corners of the resulting inverted pyramid indentation [7]. The softer the material, the further the diamond pyramid will penetrate, and hence the greater the distance between opposite corners. Vickers hardness tests performed on hard ceramics, such as the case here with silicon nitride, often yield fine stress cracks from the corners of the indentation. Referring to the two images of figure 5.1, the Vickers test feature is apparent on both the optical image and the SAW amplitude image—on the optical image as a cluster of [easily-missed] dark pixels on the optical image, and as a shadow region beneath the feature on the acoustic amplitude image. Feature C is a somewhat smaller indentation caused by a Vickers test; again, it is visible on both the images, although is much easier to spot on the amplitude image due to its large shadow.

Feature D is apparent only on the optical image, and not on the acoustic amplitude image, or at least there is no 'acoustic shadow' This implies that the feature does not impede SAW waves at the frequency used. In fact feature D is almost certainly a speck of dust.

The important observation is that on the optical image, it is difficult to differentiate between features C and D, whereas on the amplitude image, it is immediately apparent that feature C is a flaw, whereas feature D is not. Feature A is an easily-observable flaw on both optical and SAW amplitude images; however this feature, along with features B and C, are easier to observe acoustically.

These observations are by no means new revelations, nor are they in any way unique to the O-SAM instrument—they may be observed using any number of the SAW generation and detection techniques described in Chapter 1. But the demonstration that the O-SAM may be used in this way is an important step. The ability to obtain optical—as well as acoustic—images is a useful

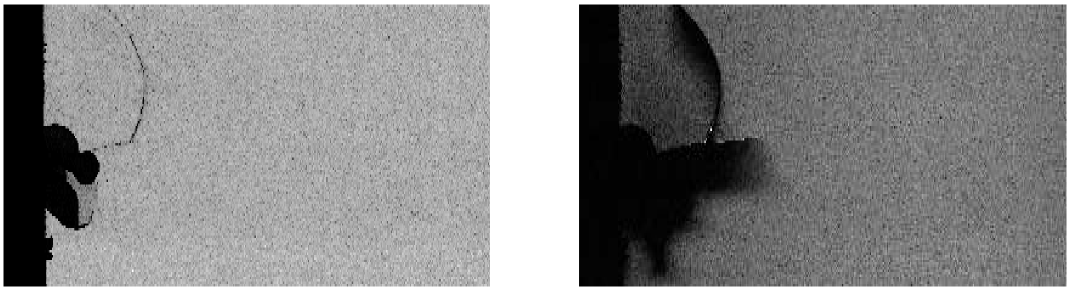

Figure 5.2: Optical (left) and 82MHz SAW amplitude (right) images of a large block of silicon nitride, which contains an optically-visible crack, as well as a finer crack that is invisible at this resolution. Image size is 20×11.5mm, resolution is 50×100μm

'built-in' feature of laser ultrasound techniques in general, since the more contrast mechanisms there are available, the higher the probability of correctly ascertaining the material properties in question.

Figure 5.2 demonstrates the concept further. Once again, the optical image is on the left, whilst the 82MHz SAW amplitude image is on the right. There is an obvious flaw in the material on the left hand side—a very obvious crack, probably caused by stresses whilst the material was being cut, is visible on both images. However, the SAW amplitude image also highlights a much finer crack, emanating from a corner of the larger crack; this is not visible on the optical image. A higher resolution image of this finer crack is shown in the next section.

5.3.2 Detailed scans for characterisation

In addition to the detection of cracks and discontinuities the O-SAM instrument may be used to characterise them, in terms of location and approximate depth. This can be achieved by examining both the phase and amplitude images, at two or more SAW wavelengths.

Figure 5.3 is a high resolution scan of the area around the edge of the fine crack in figure 5.2. Each pixel in the figure represents a distance of 2μm; however the spot size of the knife edge detector is approximately 8μm, which ultimately determines both the optical resolution, and the resolution of the ultrasound images.

The higher resolution of this image, as opposed to the large area images previously shown, obviously affords more detail, and interesting features can be observed. Reflection fringes are observed in the SAW amplitude images at both wavelengths. These fringes are the result of interference between incident and reflected surface waves, and only occur because each 'event' consists of a packet of waves. Since in the normal c-scan regime the propagation distance between

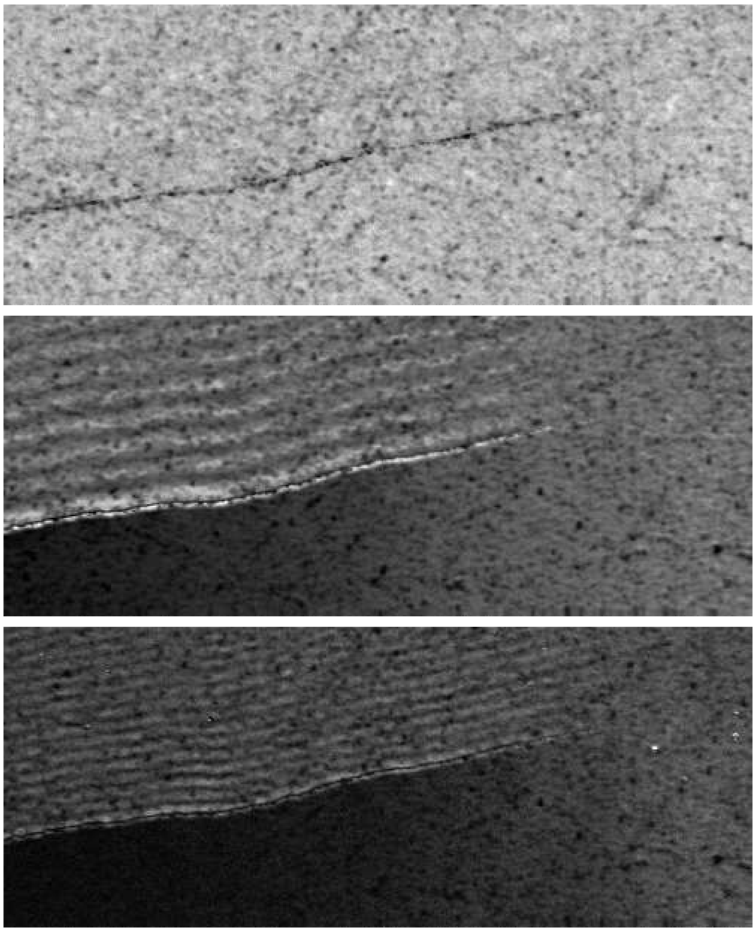

Figure 5.3: Optical (top), 82MHz SAW amplitude (centre) and 164MHz SAW amplitude (bottom) images of a fine crack in a silicon nitride sample. Image size is 1×0.6mm, resolution is 2×2μm

source and detector is constant whilst the sample is moved, a change in sample position of a distance of half a wavelength implies a change in propagation distance of the reflected wave of one wavelength. The fringe spacing is therefore equal to half the Rayleigh wave wavelength. This wavelength dependence explains why the spacing of the fringes in the 164MHz amplitude image is half that of the fringes on the 82MHz image. Very similar effects have been observed with a water-coupled scanning acoustic microscope [73].

We note that the depth of the fringes is proportional to the amplitude of the reflected waves with respect to the amplitude of the incident waves. At a 'perfect' boundary—one that incurs no losses and has 100% reflection—the forward- and backward-propagating waves would interfere perfectly, such that the maximum amplitude was twice the amplitude of the incident wave, and the minimum amplitude was zero. This is unlikely to occur at distances more than a couple of wavelengths from the crack, because the shape of the pulse envelope means that the amplitude of each wave is different. The number of pulses in the envelope determines the number of visible fringes.

The fringe modulation depth can in part be used to infer some information about the location and depth of the crack. The penetration depth of Rayleigh waves is of the order of a wavelength—cracks of a depth less than this will not reflect 100% of the wave. Similarly, the edge of a crack can be determined by a combination of the edge of the 'acoustic shadow' behind the crack and the edge of the fringes in front of it, with greater accuracy than an optical image of the same resolution—if indeed the crack is visible at all.

By scanning at two or more frequencies, a clearer picture about the crack profile can be produced; SAWs of higher wavelengths are more 'sensitive' to cracks and discontinuities due to their smaller penetration depth into the material. This can help with the detection and characterisation of even very shallow cracks.

Figure 5.4 contains three high resolution images, all of the same area. It is a 0.8mm square area of silicon nitride, with a large Vickers hardness test indentation in the centre, similar to the one described in section 5.3.1. Figure 5.4a is the optical image; figures 5.4b and 5.4c are the SAW amplitude images at 82MHz and 164MHz respectively.

The subject in these images is obviously more complicated than a single crack. Initially, however, we will concern ourselves with the two cracks emanating from the sides of the indentation, and examine some of the other features of the images later on.

Cracks can be seen emanating from the corners of the inverted pyramid indentation—the two horizontal cracks are particularly visible in the SAW amplitude images, since the ultrasound propagation direction is—as usual—from the top of the image to the bottom of the image. The cracks from the top and bottom corners are also visible.

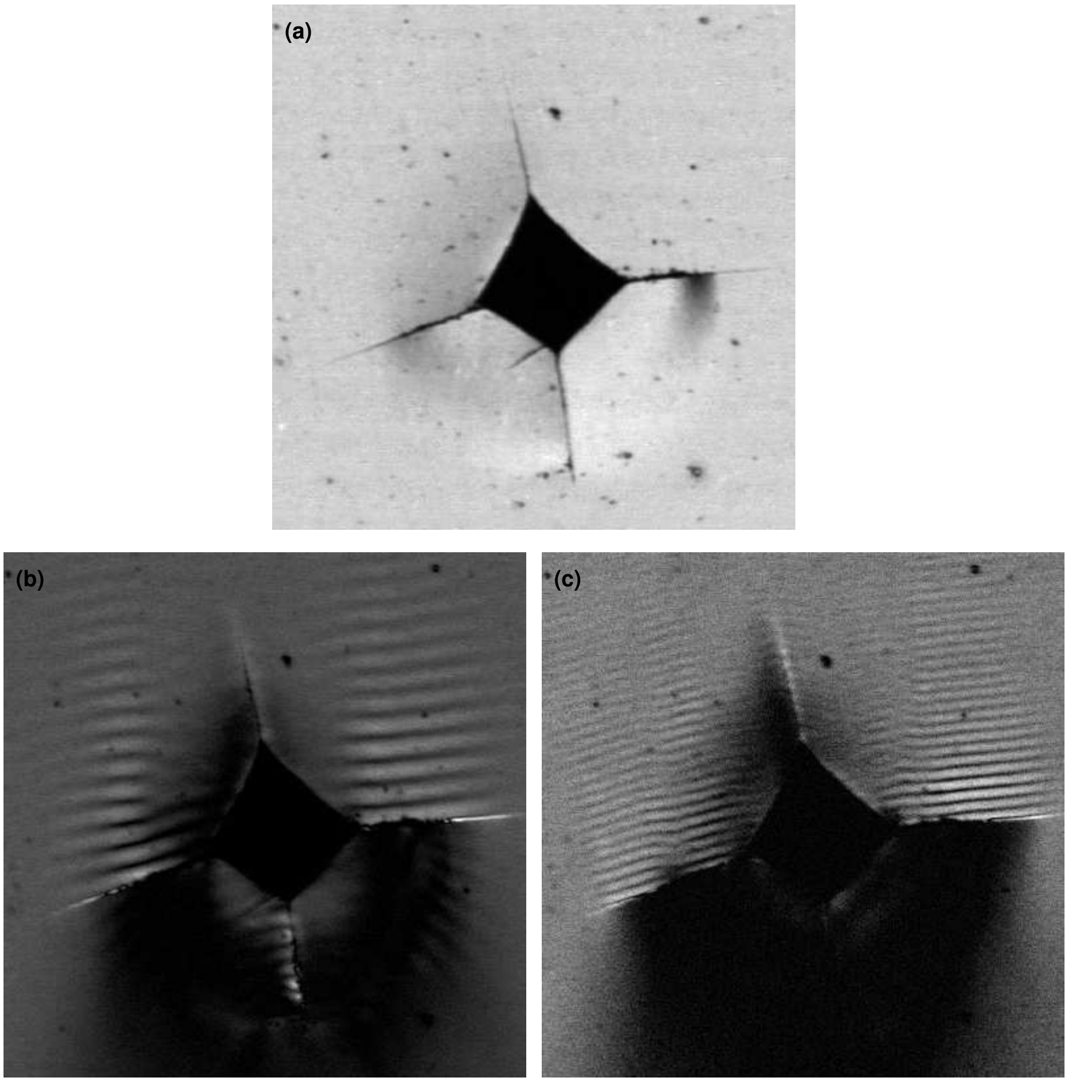

Figure 5.4: 0.8×0.8mm images of a Vickers hardness test indentation on a silicon nitride sample. The three images are as follows: (a) optical image; (b) 82MHz SAW amplitude image; (c) 164MHz SAW amplitude image. The resolution is 1×1μm

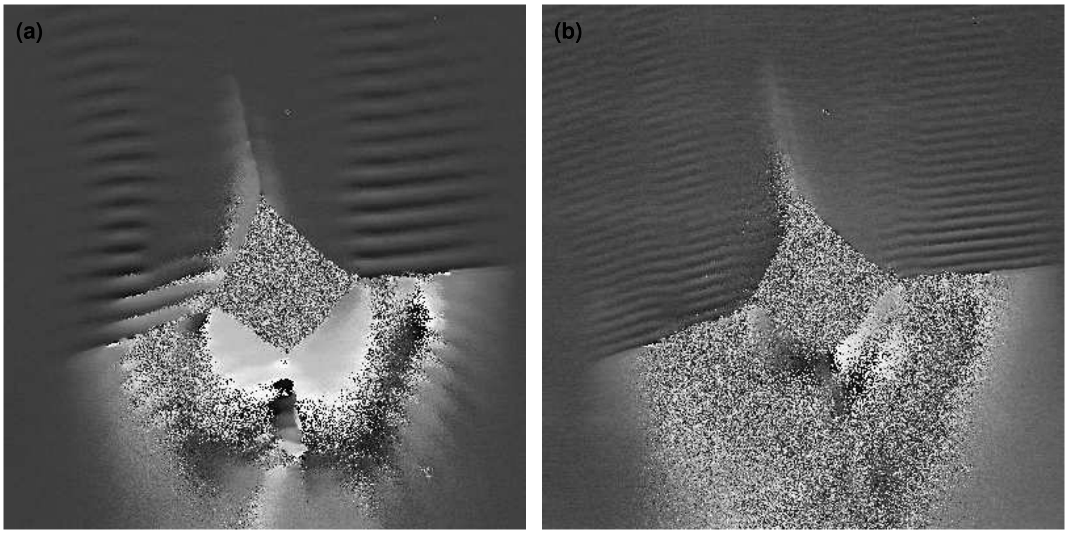

Figure 5.5: 0.8×0.8mm SAW phase images of a Vickers hardness test indentation on a silicon nitride sample, covering the same area as the images in figure 5.4. (a) 82MHz SAW phase image; (b) 164MHz SAW phase image. The resolution is 1×1μm

The modulation depth of the reflected fringes can give some kind of indication as to the profile of the crack—whether it comes to an abrupt end, or whether it gets gradually less deep. Unfortunately, differences in amplitude and signal to noise ratios between the two frequencies (82MHz and 164MHz) can lead to difficulties in interpreting the data correctly.

However, the SAW amplitude image is not the only tool available to us for determination of crack features. SAW phase images, especially at two or more frequencies, can often be an invaluable source of information in this respect. Figure 5.5 contains two SAW phase images covering the same area as figure 5.4. Figure 5.5a is at 82MHz and figure 5.5b is at 164MHz.

"Reflection fringes" are observed in the phase images—this is to be expected, since this occurs whenever forward- and backward-propagating waves interfere. There is a distinct phase change observed at the cracks, and this may be used to infer some kind of information about the crack profile.

Figure 5.6 shows a collection of line scans of the detected phase before and after the right hand crack, at various places along the crack, and at 82MHz and 164MHz. We note that close to the Vickers indentation—where the crack is presumably relatively deep—the observed phase change is similar at both frequencies. As we move further to the right, we note that the phase jump at the lower frequency begins to decrease rapidly, whilst the phase jump at the higher frequency decreases more slowly. This is due to the relative penetration depths of the two frequencies of acoustic waves.

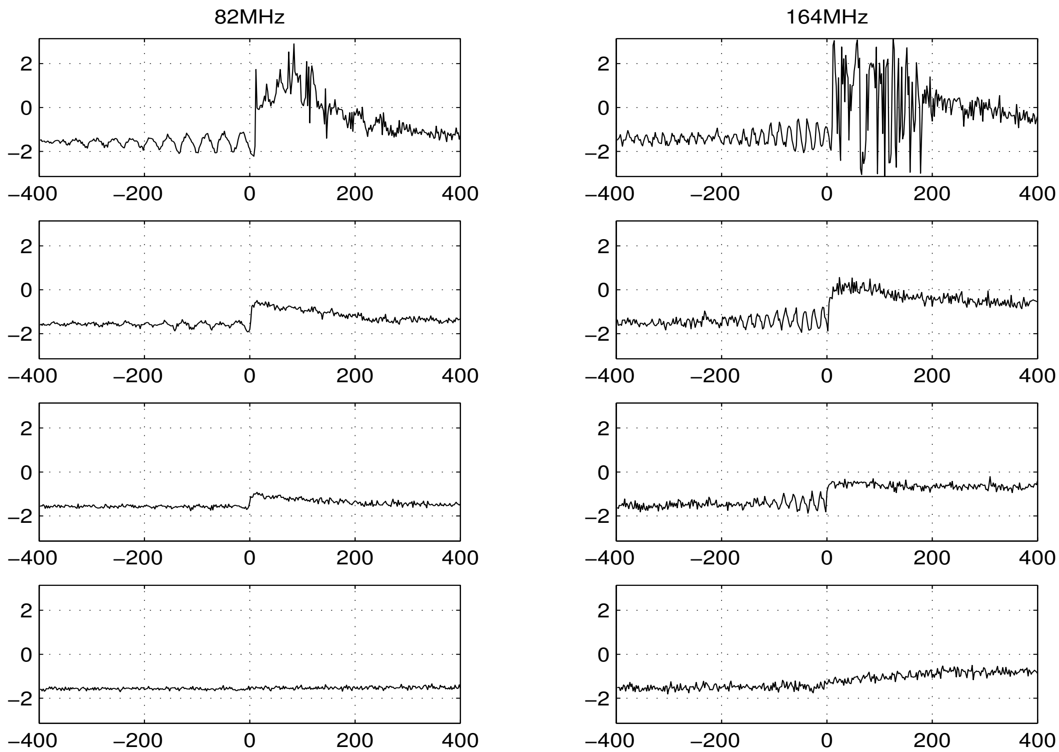

Figure 5.6: On the left are line scans of the 82MHz SAW phase at various points along the right-hand crack emanating from the Vickers indentation; on the right are 164MHz SAW phase line scans at the same positions. The horizontal and vertical scales are the same on all the plots: distance (x-axis) is in microns; phase (y-axis) is in radians.

The 164MHz SAWs, having a wavelength half that of the 82MHz SAWs, has a correspondingly smaller penetration depth, and is therefore more sensitive to very shallow cracks.

There are several other interesting features of the images presented in figures 5.4 and 5.5. There is a lot of information contained within the two images, as they are of a reasonably complex artifact. We shall return to the analysis of these images, and similar features of a different artifact, in section 8.4.

5.4 Determination of surface stress and/or porosity

The previous section has shown that the Rayleigh wave amplitude and phase information, particularly when acquired at more than one frequency, can contain useful information and infer certain defect properties, such as the location and depth of discontinuities close to the surface. Additionally, SAW phase images may be processed to provide information about the change in velocity of the Rayleigh waves. This change in velocity is associated with a change in some characteristic of the material properties, typically surface stress, or porosity.

5.4.1 Determination of mean SAW velocity

There are essentially two ways of expressing the velocity information derived from the phase information. It is relatively easy to derive the *mean velocity* between the source and detector from the phase information. It is a little trickier to extract the *local velocity* at any given location, although this may be more appropriate where abrupt changes in velocity are required to be measured. This second method is described in section 5.5.1. For measurements of changes in surface stress over relatively large areas, the mean velocity is perhaps more appropriate if spatial resolution less than the propagation distance is not required. This is because, as will be explained in section 5.5.1, the calculation of the local velocity requires not only knowledge of the local change of phase (ie at the detector) but also the exact separation distance between the excitation source and the detection probe, and knowledge of the local velocity at the point of the excitation source, which is itself the result of a previous calculation. The presence of high frequency spatial noise in the system, due for instance to a poor signal to noise (SNR) ratio in the received ultrasonic signal, can adversely affect the calculations.

The calculation of the mean velocity, v_{mean}, between the SAW source and detection point is relatively straightforward. We know that

$$v_{mean} = f\lambda_{mean} \tag{5.1}$$

where f is the generation frequency and λ_{mean} is the mean acoustic wavelength. We also know that

$$d = \lambda_{mean}\left(N + \frac{\theta}{2\pi}\right) \tag{5.2}$$

where d is the distance between source and detection, N is the number of complete wavelengths, and θ is the phase at the detection point (in radians). It therefore follows that

$$v_{mean} = \frac{fd}{N + \frac{\theta}{2\pi}} \tag{5.3}$$

In terms of measurements, there are three unknowns in the equation: v_{mean}, N and d, although d is usually known to some degree of accuracy by the design geometry—the mean focal length of the generation arcs, for instance. For cases where the velocity variation is small enough that the maximum change in phase is less than 2π radians, the following method may be used to ascertain a 'guide' mean velocity, v_{guide} and N. Once these unknowns are acquired for a given sample or material, then real numbers may plugged into equation 5.3 to obtain the mean velocity from a

SAW phase image. Note, however, that the mean velocity can only be known with the accuracy with which d is known. Although this appears to be a serious limitation, the *change* in mean velocity can be ascertained to a much higher degree of accuracy. In the determination of changes in surface stress, this is the most important parameter.

To experimentally obtain values of v_{guide} and N, the following procedure may be followed. A flat acoustic wavefront is generated—either by a line, or a series of lines—and the phase of the acoustic waves detected. Since the waves are not focused, their amplitude is significantly reduced, and so digital averaging may have to be used. The detection point is moved in relatively small, known fixed steps in a direction perpendicular to the generation lines, such that the phase of the detected signal slowly increases as the detection point moves further from the generation region. The detected phase will 'wrap around' at 2π intervals, however this can be 'unwrapped' easily. In this case, equation 5.2 becomes simply

$$d = \lambda \left(\frac{\theta}{2\pi} \right) \tag{5.4}$$

$$\hookrightarrow v_{guide} = \frac{2\pi f}{g} \tag{5.5}$$

where g is the mean gradient of θ plotted with respect to d.

We now know roughly what the mean velocity is, and this allows us to calculate N. Once N is known, we can substitute all the known values back into equation 5.3 to calculate v_{mean} for any source to detector region.

5.4.2 Relating SAW velocity to material properties

A change in the SAW velocity can be used to infer certain material properties—for instance, experiments where a material has been placed in a stress jig have generally yielded an approximately linear relationship between the force applied and the SAW velocity, at least for small changes in velocity. This is due to the well-known relationship between the fractional change in wave velocity and the stress. For a Rayleigh wave travelling in the direction x, the relationship between velocity and stress may be described thus:[80]

$$\frac{\Delta v_x}{v_x} = B_{\parallel}\sigma_x + B_{\perp}\sigma_y \tag{5.6}$$

where σ_x and σ_y are the surface stresses in the x and y directions, and $B_{\parallel}$, $B_{\perp}$ are the acoustoelastic constants propagating parallel to the stress field and perpendicular to the stress field, respectively.

Unfortunately it is sometimes difficult to pin down the particular material property that is causing the change in SAW velocity, unless some prior information about a particular material sample is known; for instance the local porosity of the material also affects the SAW velocity.

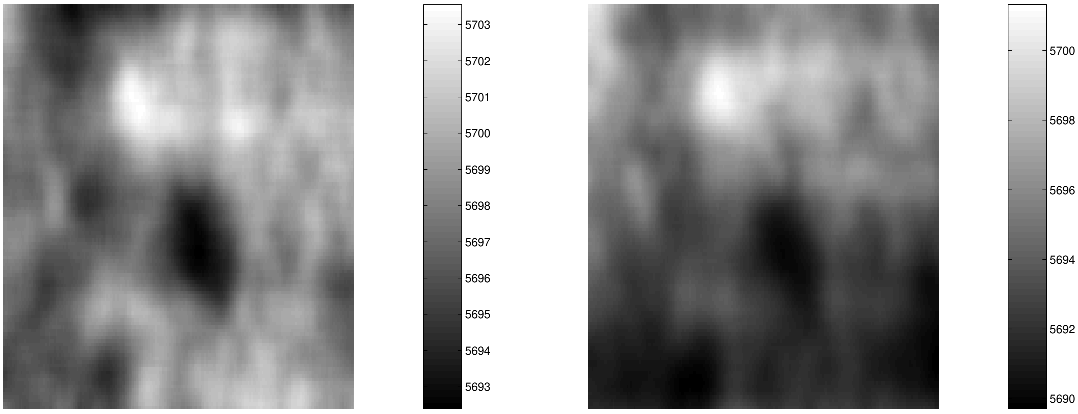

Figure 5.7: 11.5×16mm mean velocity images of a silicon nitride sample. The left hand image is derived from 82MHz SAW phase, the right hand image is derived from 164MHz SAW phase. The source → detector separation is 2mm, and the data has been spatially filtered to approximately this resolution.

5.4.3 Measurement of mean SAW velocity on silicon nitride

The large block of silicon nitride shown in figure 5.1 is used to demonstrate the O-SAM instrument's capabilities to detect small changes in SAW velocity. SAW phase scans at 82MHz and 164MHz were acquired of a relatively large (11.5mm × 16mm) area where there were no 'obvious features' such as cracks.

Figure 5.7 shows the mean velocity over the area at the two SAW frequencies.

The source to detector distance d is 2mm. The source is 16 concentric arcs imaged onto the material surface using a CGH. The data has been spatially filtered with a low pass filter to the same approximate spatial resolution as the propagation distance.

Changes in mean velocity are noted in both SAW frequencies of around ±6ms^{-1}. This could correspond to, for example, changes in surface stress, or changes in porosity. Since no knowledge of the material's history is known, it is difficult to ascertain precisely which material property is being measured. It is interesting to note that, whilst the velocity profiles of the two SAW frequencies are in many ways similar, there are differences in the degree of velocity change at different locations—this is presumably related to the penetration depth of the surface waves—but this is clearly not the main effect.

It is clear that the use of surface wave velocity measurement alone is not adequate to determine which particular material property is changing over the area of the sample if no prior information about the sample is available; the technique may be combined with other techniques such as x-ray diffraction [8] to determine this.

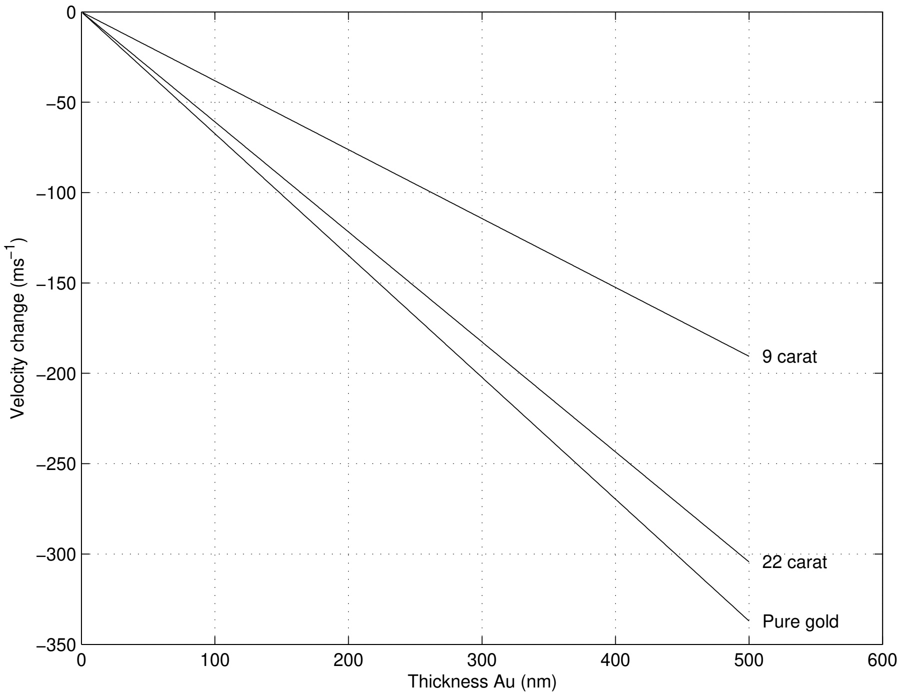

Figure 5.8: Change in Rayleigh velocity for different thicknesses of different types of gold, layered on silicon nitride

5.5 Determination of coating thickness

The velocity of surface acoustic waves propagating along a material surface can be affected by a thin coating of a different material due to mechanical perturbation. The relationship between the material properties of the bulk and layer materials and the change in SAW velocity has been examined by Tiersten [85]. For very thin coatings, the linear relationship described by Bromwich [86] may be used. To give some indication of the type of variation that can occur, figure 5.8 shows the change in velocity for three different thicknesses of different types of gold on silicon nitride.

This relationship—that between SAW velocity and coating thickness—can be exploited to measure coating thicknesses of a known material using the O-SAM instrument. Furthermore, changes in thickness can be measured and quantified.

5.5.1 Relating local velocity to SAW phase

The method of calculating the mean SAW velocity has been explained in section 5.4.1. This technique is adequate and useful for spatially 'slow'—with respect to propagation distance—changes in velocity. To measure abrupt changes in velocity—for instance when the thickness of a coating

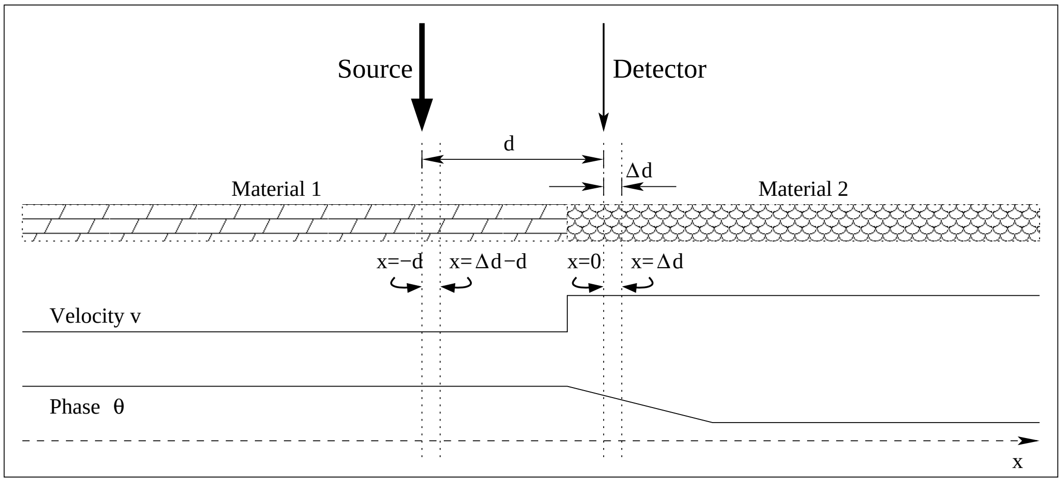

Figure 5.9: Illustration of how the phase measured by a probe a certain distance from a source is affected by an abrupt change in the SAW velocity, as the probe and source move across the region of abrupt velocity change.

changes suddenly over a very short distance—a more involved technique is required.

Consider figure 5.9, which represents a sample over a certain distance x. Initially the detector, which measures the phase θ of the received acoustic wavefront, is at point $x = 0$. The source of the acoustic wavefront is at point $x = -d$, where d is the source/detector separation distance.

Suppose then that the source and detector move a distance Δd in the positive x direction, such that the detector is now at point $x = \Delta d$ and the source is now at point $x = (\Delta d - d)$ —i.e. the source and detector are still a distance d apart. This is equivalent to moving the sample a distance Δd to the next scan point in a real c-scan experiment.

We wish to ascertain the velocity of the acoustic wave between points $x = 0$ and $x = \Delta d$, which we will term $v_{(\Delta d)}$. Since we know the distance between these two points, which is the step distance in the real experiment and Δd in figure 5.9, we could calculate the velocity between these two points if we knew the time taken for the wave to travel over this length of material. We wish, therefore, to know $\left[t_{(\Delta d)} - t_{(0)}\right]$ where the subscripts denote the position x.

We know that, in general, the time taken for an acoustic wave to travel from source to detector, t, is:

$$t = \frac{N + \frac{\theta}{2\pi}}{f} \tag{5.7}$$

where N is the number of whole wavelengths traversed, θ is the measured phase at the detection point and f is the frequency of the wave.

We therefore know that:

$$\left[t_{(0)} - t_{(-d)}\right] = \frac{N + \frac{\theta_{(0)}}{2\pi}}{f} \tag{5.8}$$

$$\left[t_{(\Delta d)} - t_{(\Delta d - d)}\right] = \frac{N + \frac{\theta_{(\Delta d)}}{2\pi}}{f} \tag{5.9}$$

Subtracting equation 5.8 from equation 5.9 gives:

$$\left[t_{(\Delta d)} - t_{(0)}\right] = \frac{1}{2\pi f}\left(\theta_{(\Delta d)} - \theta_{(0)}\right) + \left[t_{(\Delta d - d)} - t_{(-d)}\right] \tag{5.10}$$

If we have knowledge of the velocity of the wave as it travels between $x = -d$ and $x = \Delta d - d$, which we will term $v_{(\Delta d - d)}$ then we can solve equation 5.10:

$$\frac{1}{2\pi f \Delta d}\left(\theta_{(\Delta d)} - \theta_{(0)}\right) = \frac{1}{v_{(\Delta d)}} - \frac{1}{v_{(\Delta d - d)}} \tag{5.11}$$

This is an important equation, since it allows us to find $v_{(\Delta d)}$ from things that we know (d, Δd, f), a quantity we can measure experimentally (θ) and a quantity that we have calculated previously ($v_{(\Delta d - d)}$). It is a simple matter to solve equation 5.11 for $v_{(\Delta d)}$:

$$v_{(\Delta d)} = \frac{2\pi f \Delta d v_{(\Delta d - d)}}{2\pi f \Delta d + v_{(\Delta d - d)}\left(\theta_{(\Delta d)} - \theta_{(0)}\right)} \tag{5.12}$$

We note that we require knowledge of the specific velocity 'at a point $x = \Delta d - d$,' which is in reality the velocity between the finite region between $x = -d$ and $x = \Delta d - d$. Alternatively, if $v_{(\Delta d)}$ is known, then we can calculate $v_{(\Delta d - d)}$ —either way, knowledge of the velocity a distance d from the intended velocity measurement region is required, where d must be the distance between the ultrasound source and the detection point. This can obviously lead to difficulties in many situations: for instance, on samples where the velocity is changing rapidly over the entire c-scan region, it will be very difficult, if not impossible, to find initial values for the velocity for all the regions between the source and detector, to allow calculation of the velocities in other areas. If, however, there is a region where the velocity does not change between the source and detector, then initial values of the velocity at all points in between is simply the mean velocity, which can always be calculated using equation 5.3. Areas such as these can usually be found by looking for areas where the phase does not change over the source-to-detector distance; however it is possible that this region could plausibly contain changes in velocity that have a spatial period equal to the source-to-detector separation, as illustrated in figure 5.10.

If a region of uniform velocity is not available, then a different method is required to obtain the local velocity at enough points such that equation 5.12 may be used to extract the local velocity from the phase information of a regular c-scan. The method is very similar to the method of ascertaining a 'guide' mean velocity, described in section 5.4.1. A flat wavefront is made to

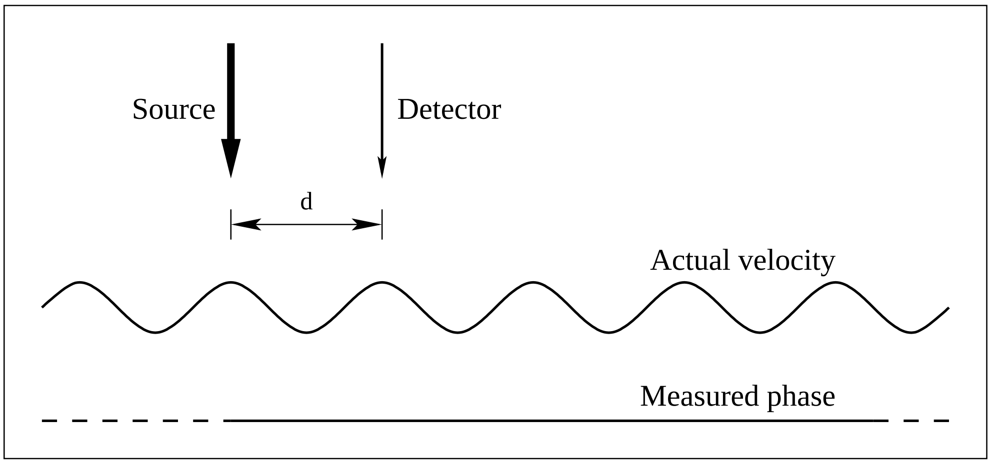

Figure 5.10: Illustration of the danger of assuming that a region of constant phase measurement corresponds to a region of constant velocity.

propagate along the material surface, and the detector is moved in small, known steps away from the source region. If we define the velocity between a point $x = -\Delta d$ and $x = 0$ as $v_{(0)}$, then:

$$v_{(0)} = \frac{2\pi f \Delta d}{\theta_{(0)} - \theta_{(-\Delta d)}} \tag{5.13}$$

Since this measurement of a local velocity at a certain 'point' does not involve any knowledge of the velocity at any other 'point,' it solves the problem of needing to know an initial map of velocities over a distance at least as large as the propagation distance.

The issue is further complicated by two other factors. The first is the scenario of having multiple sources of the surface waves—for instance by projecting many straight lines onto the sample as a grating, perhaps to spread out the optical power so as to avoid damaging the sample. As the sample is moved with respect to the generation region and detection point—to perform a regular c-scan—it is difficult to determine exactly which point corresponds to $x = -d$. Moreover, as the generation region 'passes over' a region where the velocity changes *within* that region, then the wavefronts from different parts of the grating no longer interfere completely constructively, leading to a change in phase as well as a loss of overall generation efficiency.

The second factor is concerned with generation geometry, particularly in the case of the O-SAM instrument where the surface wave source may be an arc, or a series of concentric arcs, rather than a straight line. This causes two problems in itself; firstly, the detected wavefront is a summation of all parts of the wavefront from the extremes of the source arc(s), so the wavefronts travel through the material at various different angles. This leads to further problems in determining where exactly the point '$x = -d$' lies. The second problem with having a focused wavefront is that,

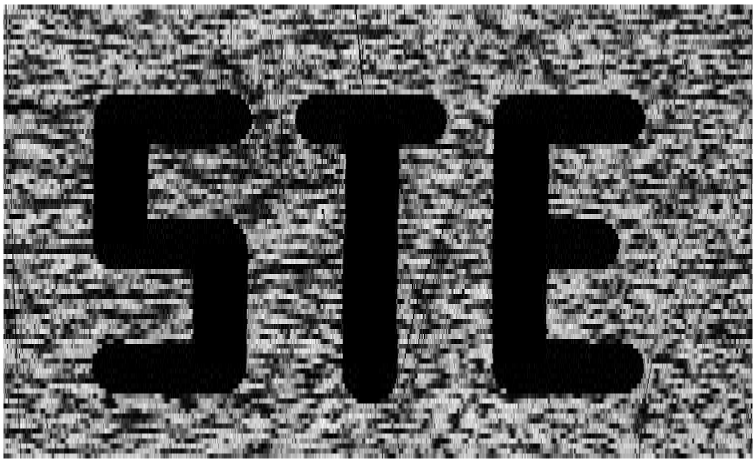

Figure 5.11: Optical image—taken using the O-SAM instrument—of the metal mask, through which 31.6nm of gold was evaporated onto a silicon nitride sample. The size of the image is 15 × 9mm

with all focusing systems, an extra phase change of $\pi/2$ is observed over the depth of focus. By observing the SAWs 'at the focus,' an extra arbitrary phase with a value between 0 and $\pi/2$ is measured. This leads to a [hopefully] consistent error in the velocity calculation.

5.5.2 Determination of thickness of gold on silicon nitride

To demonstrate the technique, 31.6nm of gold was evaporated onto a large slab of silicon nitride through a metal mask with several features on it. The thickness of the gold was measured by the instrument performing the coating, and has a reported accuracy of ±0.2nm. Figure 5.11 shows the features of the metal mask.

The mask was clamped on to the surface of the silicon nitride so that the edges were reasonably well defined.

The metal mask was then removed, and 153.7nm of gold was evaporated onto the sample, completely obscuring the thin pattern of gold. Figure 5.12 is an optical image of the coated sample covering the area over which the mask was placed to coat the extra layer of gold. The image was acquired using the O-SAM instrument, and the pattern beneath was not in any way visible to the naked eye.

A SAW vector contrast image was then acquired of the area in figure 5.12 at 82MHz. The signal to noise ratio was very poor, and this is due to the fact that gold is highly reflective at 1064nm optical wavelength—the SAW generation frequency—and so very little optical energy is absorbed in the sample to be converted to acoustic energy. This reason was confirmed by scratching away a

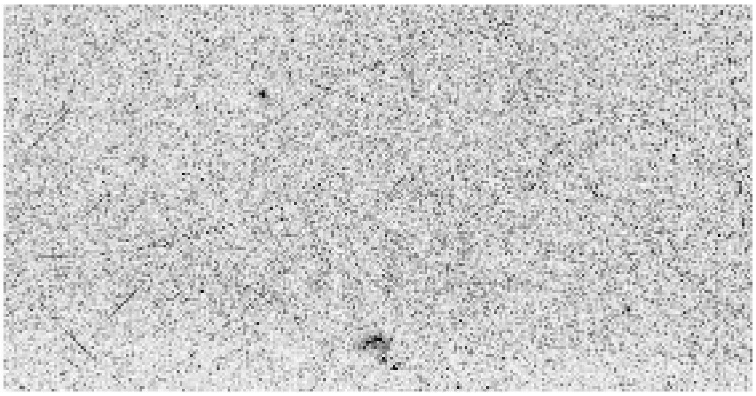

Figure 5.12: Optical image—taken using the O-SAM instrument—of the silicon nitride sample, after it had been coated with gold. The thicker areas of gold are completely invisible to the naked eye. The size of the image is 14 × 7.2mm

small region of gold from the surface, generating the SAWs in this region, and detecting them on the gold-plated surface a little distance away. This led to a very good signal to noise ratio. The top image of figure 5.13 shows the SAW amplitude across the region, and the bottom image of figure 5.13 shows the SAW phase across the region. As always, the propagation direction is from the top of the page to the bottom.

It is possible, by careful examination of the amplitude image in figure 5.13, to conclude that you can just about see a pattern. This is probably due to the difference in impedance for the travelling Rayleigh waves, leading to an almost indiscernible decrease in amplitude in the thicker regions. This is a little unfortunate, since the desired aim of the experiment is to demonstrate that the phase can give rise to information that is not contained within the optical image or the amplitude image. Nevertheless, it is obvious that the change in phase is much more apparent than the change in SAW amplitude.

The phase image in figure 5.13 does not directly give rise to the areas where the gold is thicker, it must first be converted to a velocity. If we used the mean velocity method, as described in section 5.4.1, then the velocity image would look very similar to the phase image (except the light and dark regions would be inverted). This requires the more complicated method described in section 5.5.1.

Firstly we note that, due to the poor signal to noise ratio, the phase image contains much high frequency spatial noise. This is important, because the algorithm to extract the local velocity information is extremely susceptible to this noise. This is because the calculated velocity at any point may be used to calculate the velocity at another point a distance d away. Errors soon add

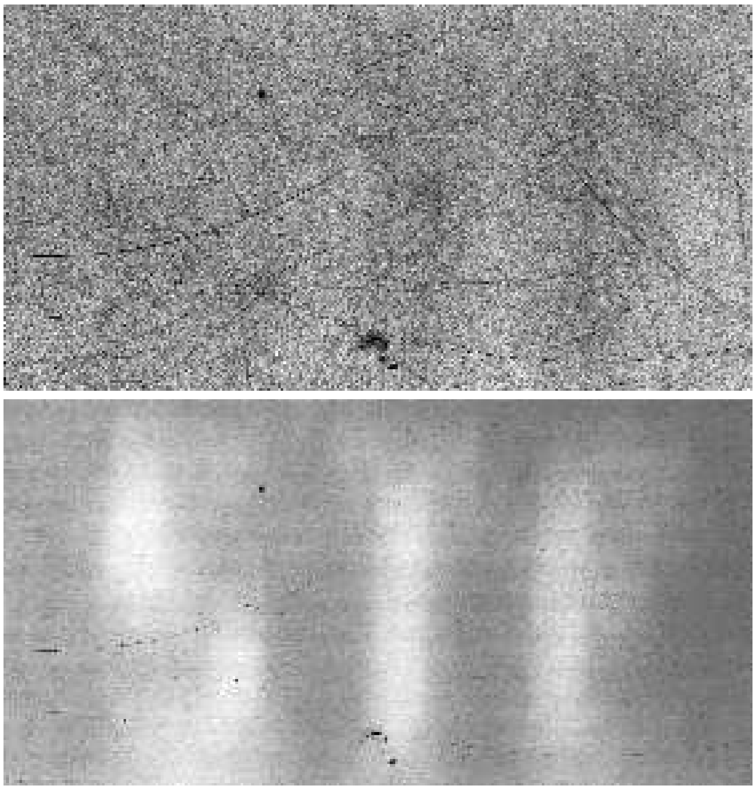

Figure 5.13: 82MHz SAW amplitude (top) and SAW phase (bottom) c-scans of the region of gold-coated silicon nitride that an extra layer of gold had been applied to beneath the top layer of gold. The size of the image is 14 $\times$ 7.2mm

Figure 5.14: 82MHz SAW phase image, covering the same area as figure 5.13b, spatially filtered in the direction of propagation (vertical). The size of the image is 14 × 7.2mm

up, and the small 'real' velocity changes are soon buried beneath the noise. Therefore, some simple low pass spatial filtering is applied to the phase image along the propagation direction, before the local velocity is calculated. Figure 5.14 is an image of the spatially filtered phase information.

The local velocity information is then acquired from the filtered phase data using equation 5.12. This is shown in figure 5.15.

This velocity information may then be directly converted to information representing differences in gold thickness. Figure 5.16 presents this data, with the outline of the mask overlaid on top.

Obviously, the image representing the extra gold thickness does not tally very well with the outline of the mask. Most of the reasons for this are outlined towards the end of section 5.4.1; namely uncertainties and errors caused by the propagation geometry (generation by several arcs, propagating to a point), inaccurate determination of the source-to-detector propagation distance (which is itself difficult to define since several lines are used in generation), and low signal to noise ratio. A single straight line source would eliminate the first few problems, but it is extremely doubtful that any SAWs could be detected at all, without ablating the sample by providing too much localised optical energy during SAW generation.

5.6 Summary

This chapter has dealt with the issue of using Rayleigh waves to determine different forms of information about samples, using a variety of different contrast mechanisms. This includes the detection of cracks by attenuation of the Rayleigh waves to detect surface and near-surface cracks; phase changes at discontinuities at different SAW wavelengths to infer information about the

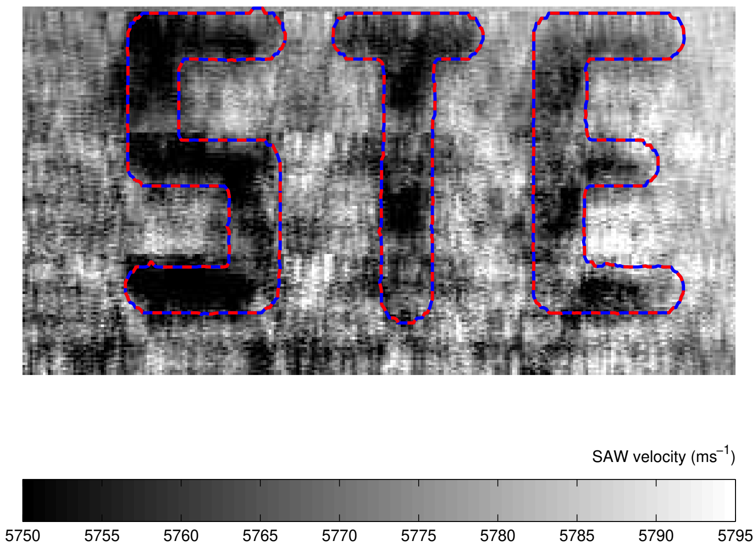

Figure 5.15: SAW velocity image, acquired from the filtered phase information shown in figure 5.14. The size of the image is 14 × 7.2mm

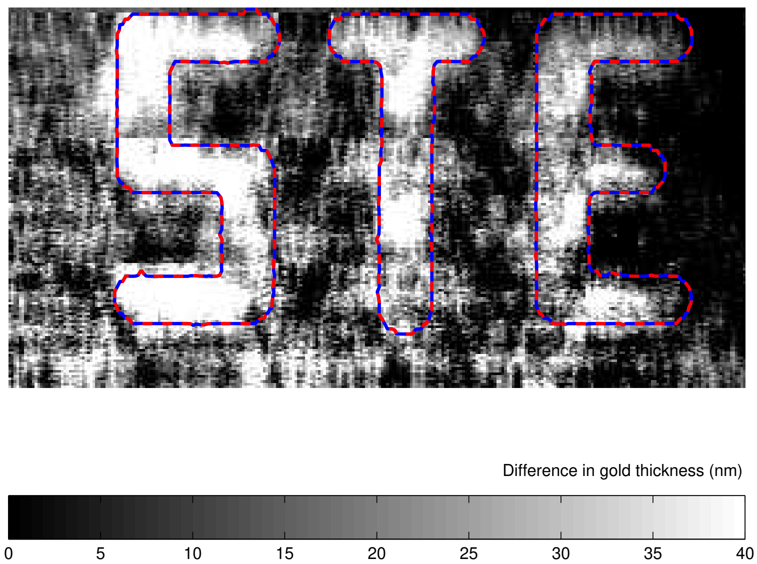

Figure 5.16: Image representing the difference in gold thickness at different points on the area shown. An outline of the mask edges has been artificially added for clarity. The size of the image is 14 × 7.2mm

likely depths of shallow surface cracks; measurement of changes in stress/porosity from mean velocity variations; and determination of coating thickness differences, by more spatially-detailed variations in Rayleigh wave velocity. Clearly, the O-SAM instrument gains some advantages over either traditional laser ultrasound systems (by obtaining much higher signal to noise ratios without any surface damage) and the contact scanning acoustic microscope (by its non-contact nature and system geometries). It is also clear that it is not a win-win situation... extra complexity in the generation profile and propagation direction can lead to confusing or erroneous results if it is not fully taken into account; and simple issues—such as gold is a very good reflector of light at 1064nm, and thus it is very difficult to generate any Rayleigh waves at all—can sometimes unfortunately put the instrument at a disadvantage compared to contact techniques.

Chapter 6

Imaging using Lamb waves

6.1 Introduction

The previous chapter demonstrated some of the imaging capabilities of the O-SAM instrument, using the interaction between the generated Rayleigh waves and the sample to either determine or infer some physical property or properties of the near-surface structure of the sample under investigation. Rayleigh waves are relatively easy to generate using a laser-based generation system, especially one that allows for the generation profile to be adjusted to suit the material properties and the desired modes of the generated waves. They (Rayleigh waves) are also easily detectable using the modified knife edge detection technique.

After a brief introduction on some of the more pertinent properties of Lamb waves, this chapter will describe how they may—in certain circumstances—be excited by Rayleigh waves generated by the O-SAM system, and directly detected using the modified knife edge detector already described. Images will be presented of Lamb waves propagating on thin surfaces, and these images will be used to acquire information about the physical structure of the material under test, including contrast mechanisms that extend beyond the capabilities of contact-based scanning acoustic microscopes.

6.2 Brief review of some Lamb wave properties

There are many descriptions of the properties of Lamb waves in the literature [10, 8] but a brief description will be given here for completeness and context.

Whereas Rayleigh waves propagate at the boundary between—usually in terms of nondestructive evaluation—a solid elastic half-space and air, Lamb waves propagate in situations where two parallel surfaces are found; typically a solid plate, or a layer, with free boundaries. As the Lamb

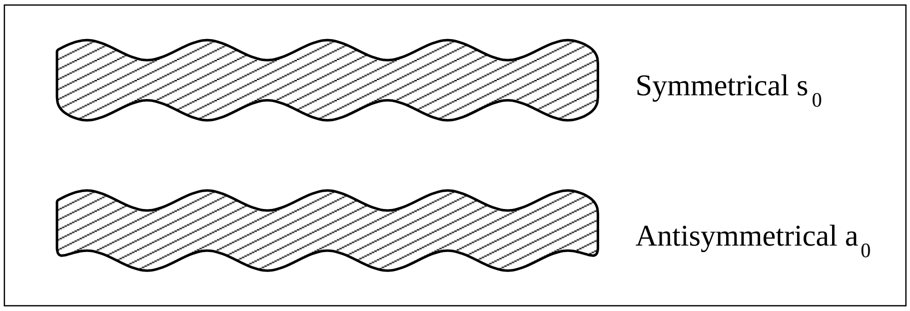

Figure 6.1: Representation of the first symmetrical (s_0) and antisymmetrical (a_0) Lamb wave modes. The amplitude of the waves is significantly exaggerated in the figure for illustration purposes.

waves propagate, displacements occur in the direction of wave propagation and perpendicular to the plane of the plate, on both surfaces of the plate.

Typically, several Lamb wave modes exist as the waves propagate, and these are split into two groups: *antisymmetric* and *symmetric* modes. Figure 6.1 shows representations of the first mode of each type, the so-called a_0 and s_0 modes.

An important property of Lamb waves is that they are *dispersive,* in that the excitation frequency and plate thickness determines their propagation velocities. For each frequency and thickness combination, the group and phase velocities of the symmetric and antisymmetric modes may well be different; the velocities of higher modes (a_1, s_2 etc.) will also be different. These phase and group velocities may be plotted as a function of frequency and plate thickness for a given set of material properties, and these plots are known as dispersion curves. Figure 6.2 shows the phase velocity dispersion curves for the a_0 and s_0 Lamb modes at a fixed frequency (82MHz) on aluminium.

The calculation of the curves is relatively common-place with the advent of computers, and will not be described in any more detail here, since it is covered in the literature [87]. Since the velocities are dependent on both the frequency and the plate thickness, the x-axis is often the product of these two variables in many dispersion curves.

It is noted from figure 6.2 that as the thickness increases, the velocities of the fundamental modes approach that of Rayleigh waves in that material.

6.2.1 Conversion of Rayleigh waves to Lamb waves

It is possible to generate Lamb waves directly using laser ultrasound, in the same way that we excite Rayleigh waves on 'thick' materials—that is, by matching the fringe spacing of the incident

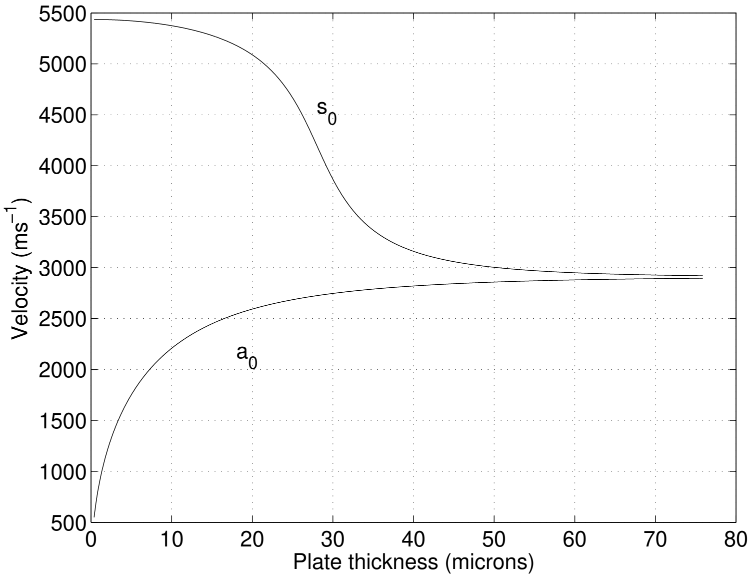

Figure 6.2: Phase velocity dispersion curves for the a_0 and s_0 Lamb modes at a fixed frequency (82MHz) on aluminium.

laser light to the Lamb wavelength of the required mode, given a certain pulse frequency.

It is also possible to excite Lamb waves using Rayleigh waves at boundaries between 'thick' and 'thin' media. In this case, by 'thick' we mean having a depth of greater than about two times the Rayleigh wavelength; by 'thin' we mean less than this. The knife edge detection system employed in the O-SAM system is sensitive to the angular displacement of the top surface of the material as the waves propagate, and is thus ideal to measure the complex amplitude of the Lamb modes as they propagate in thin materials. This, in turn, can yield important information about either the structure or properties of the material.

6.3 Depth profiling using Lamb mode conversion

The O-SAM instrument is a useful tool therefore to study the propagation of Lamb modes on thin layers; not least because the SAW source is narrowband in nature (albeit at several discrete frequencies) and so the results are more easily interpreted than if the source of the ultrasound is wideband. For a given source frequency—82MHz, 164MHz etc—the velocity of the Lamb modes depends only on the plate thickness for a given material.

A series of experiments was performed to investigate the feasibility of using the O-SAM instrument to generate and detect Lamb waves, and to use the acquired complex amplitude to

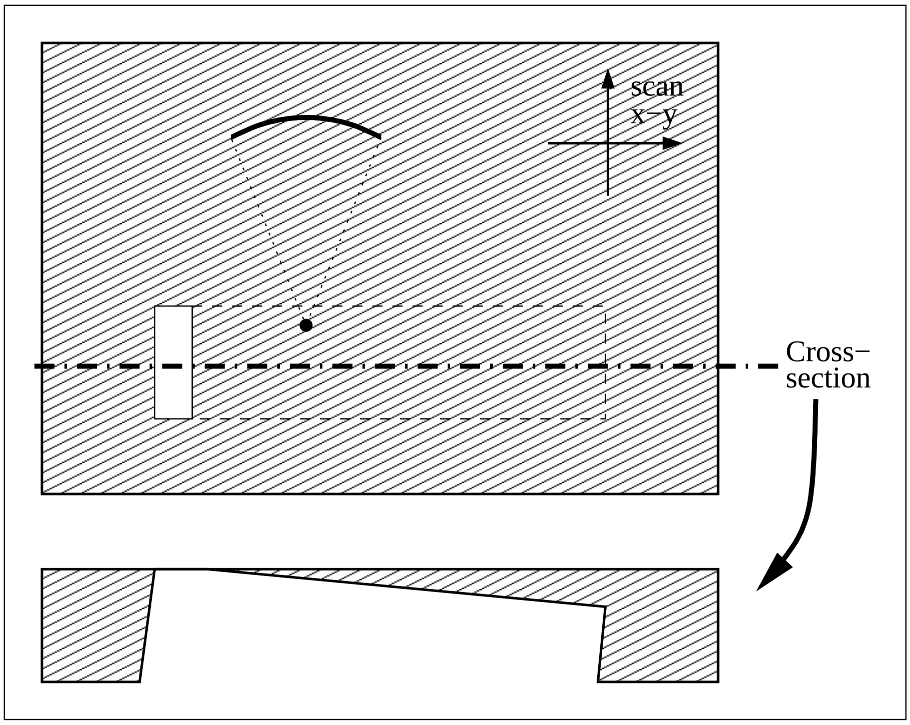

Figure 6.3: Schematic of the '1° slit experiment.' A slit is milled into base of a piece of polished aluminium at an approximate 1° angle. The sample is scanned in the x-y directions to achieve a normal c-scan. The vertical dimension and slit angle in the cross-section diagram have been greatly exaggerated.

calculate the velocities of the two lowest Lamb modes (a_0 and s_0). A variety of different sample geometries were used, and both focused and plane SAWs were used in the experiments. All the results presented were performed on aluminium samples that had had their top surface polished to accommodate easier detection using the modified knife edge detector.

6.3.1 '1° slit experiment'—using focused waves

A piece of aluminium approximately $60 \times 60 \times 5$mm in size was cut from stock. With what was intended to be the 'top side' placed faced down, the sample was tilted at an approximate angle of 1° from the horizontal. A slit was then milled into the 'back' of the sample—i.e. the upper side as it sat on the milling machine—with a 1mm diameter milling bit. The overall length of the slit is approximately 30mm, and the milling bit broke through the 'top side' at one end. The sample was then removed from the milling machine, and the 'top side' was polished. Figure 6.3 is a schematic of the sample, with the thickness and slit angle greatly exaggerated in the cross-section.

The figure also gives an indication of the experimental layout. The sample was scanned in the usual c-scan mode of operation; that is, the positions of the surface wave source and detector were kept fixed, whilst the sample was raster-scanned using x-y mechanical stages, to build up a c-scan image point by point of the complex amplitude of the SAWs at 82MHz.

The general aim of the 'depth profiling experiments' is to propagate Rayleigh waves from a 'thick' area of the sample towards the 'thin' area of the sample; we shall refer to this area from now on as simply 'the slit.' At the slit edge the Rayleigh waves excite Lamb modes, which are then measured with the O-SAM instrument.

The ultrasound source for the $1°$ slit experiment was an image produced by the spatial light modulator. This section describes the results for focused waves, and so the SLM pattern imaged onto the sample was a series of concentric arcs, with a mean focal length of 2.5mm, measured from the centre of the SLM image. The half-angle of the SAW focus was therefore approximately $26°$ from the centre of the SLM image, and ranged from $18°$ to $43°$, due to the fact that the projected image was square. Before the experiment itself was carried out, the mean Rayleigh velocity was ascertained by the method described in section 5.4.1, 'Determination of mean SAW velocity,' and was found to be 2945ms^{-1}, giving the Rayleigh waves a wavelength of 35.9μm.

Knowledge of the velocity is important, because of the way the O-SAM measures the SAW complex amplitude in c-scanning configuration. Since the source to detector distance is kept constant, then over an area of constant velocity—for example the Rayleigh velocity—the measured phase stays steady. The measured phase of waves travelling at different velocities—for example Lamb modes—varies along the path of the waves, as a function of wave velocity. We can measure the *difference* in velocity between each of the modes and the Rayleigh velocity using Fourier analysis, but without knowledge of the Rayleigh velocity, we cannot determine the *absolute* velocities of each of the modes, and cannot directly match the mode velocities to those on the dispersion curves.

Because the slit is relatively thin, it is likely that the predominant modes will be the a_0 and s_0 modes. These modes then propagate on the slit, and can be detected by the O-SAM instrument. The Lamb modes interfere with each other if they have different velocities, and create fringes in the amplitude response, and can produce complex variations in the phase response. With knowledge of the dispersion curves, this can be easily modelled for various different thicknesses at a given excitation frequency; this has been done in figure 6.4.

In the figure, the x-axis represents the plate thickness, and the y-axis represents the amplitude or phase as measured by the O-SAM over a certain distance, which in this case is 1mm in total. The model was created by setting up two waves of the appropriate velocities of the a_0 and s_0 modes for a range of thicknesses (between approximately 25 and 60μm). The model assumes that the velocity of the wave prior to the plate is the Rayleigh velocity, and the appropriate phase associated with

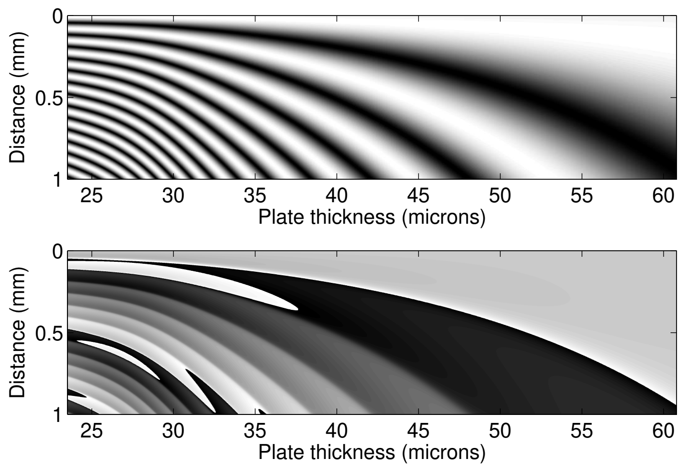

Figure 6.4: Simulated amplitude (top) and phase (bottom) responses from the O-SAM system for a_0 and s_0 Lamb modes propagating from top to bottom on plates of various thicknesses.

this is subtracted from the sum of the Lamb modes, to give the amplitude and phase responses shown. Note that the response is calculated for each given thickness separately, with no regard for any 'adjacent' thicknesses. In other words, the model does not attempt to truly predict the response of the O-SAM instrument for waves propagating over a plate of varying thickness (from left to right), but rather illustrates the response of plane waves travelling over uniform plates of various thicknesses, plotted such that one vertical line corresponds to the response on a plate of the corresponding thickness. Nevertheless, this very simple model can provide a very useful indication of the expected response of the O-SAM system.

The c-scan images obtained for the experiment as described are illustrated in figure 6.5.

The upper part of the figure shows the optical image obtained whilst the 82MHz c-scan was being performed. Added to the image are two horizontal dashed lines, indicating the approximate boundaries of the subsurface slit. The aluminium is thinner on the left hand side, and at the very edge of the image one can see that the surface is sufficiently disturbed such that the optical response drops away. This is due to the [relatively] crude manufacturing process involved in making the slit.

The centre and bottom images of figure 6.5 show the 82MHz amplitude and phase responses of the c-scan respectively. There are several interesting features worthy of note in these two images.

Contrast between the two Lamb modes is clearly illustrated by a series of characteristic dark fringes in the amplitude image. The fringe spacing is greater on the right hand side of the image, indicating that the two modes have similar velocities. Towards the left hand side of the amplitude image, the fringe spacing decreases, implying velocities that are further apart. By referring to the dispersion curves in figure 6.2 we see that these situations—large and small fringe spacing—arise

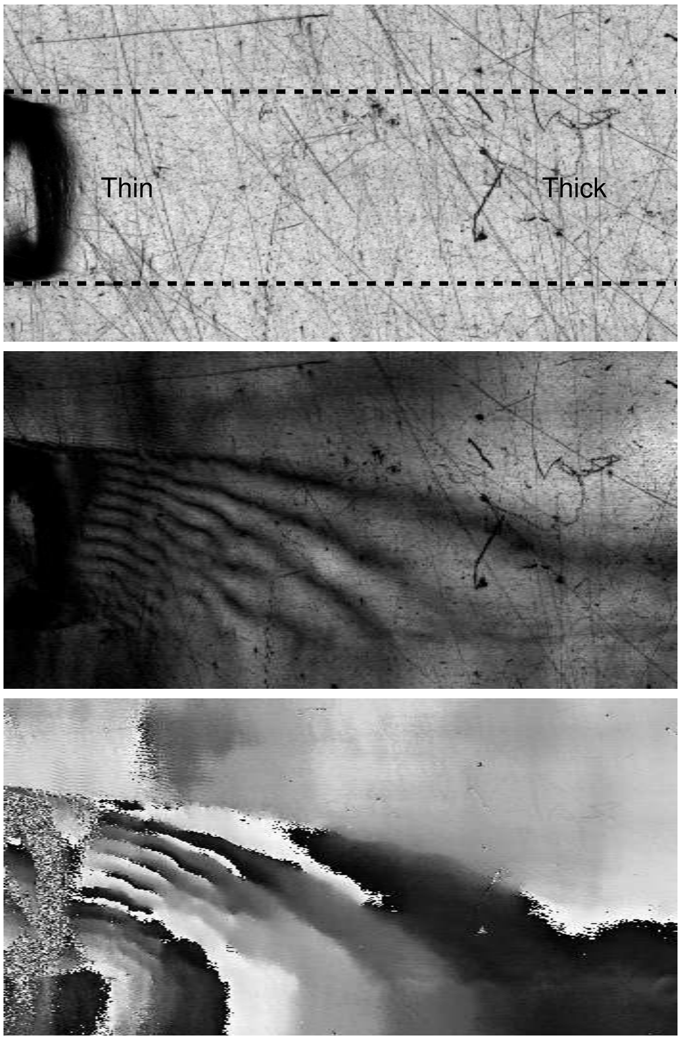

Figure 6.5: Optical (top), 82MHz SAW amplitude (centre) and 82MHz SAW phase (bottom) c-scan images of a 3.5 × 1.75mm area of the top surface of a piece of polished aluminium. The location of the slit is shown by dashed lines in the optical image. The slit breaks the surface on the left of the image, where the thickness of the aluminium is at its thinest.

for thick and thin plates respectively.

The phase image also indicates contrast between the two Lamb modes, although the results are a little more difficult to interpret from the phase image alone. There is clearly more than one kind of mode present, since the pattern is complicated and there is not simply a gradual change in phase (as would be the case if only one mode was present).

The amplitude and phase images bear a reasonable similarity to the images of the simple model presented in figure 6.4, although there are differences also. This is due in part to the simplicity of the model, as already stated. Differences of this nature include the unevenness of the dark fringes on the amplitude image, which in several places appear to split or join, rather than being continuous as with the model. Unevenness in the phase image also illustrates this. Possible reasons for this disparity will be discussed and investigated in the following two sub-sections.

Other observations of note include the relative effects of the surface finish on the amplitude and phase responses, in that surface scratches and other cosmetic defects are more apparent on the amplitude image than the phase image. This is to be expected, since the amplitude of the surface wave measurement is proportional to the received optical amplitude at the detector, whereas the phase is not. It is comforting to note that a reasonably large change in amplitude caused by the scratches does not produce an appreciable change in measured phase, indicating good amplitude-phase crosstalk immunity of the system.

One can locate on the amplitude and phase images the boundary of the thick and thin regions of the aluminium, i.e. between the upper edge of the slit and the thick region, by the Rayleigh wave reflection fringes at this boundary, which are particularly evident where the slit is very thin on the left hand side of the images. These fringes are similar to those observed at cracks, as described in section 5.3.2, 'Detailed scans for characterisation.' In this case, they indicate that not all the energy from the Rayleigh waves is being mode converted into Lamb waves.

The final observation of note concerns the area above the slit, i.e. the top quarter of the images. Here the SAWs are propagating on a 'thick' piece of material, and so are Rayleigh waves. There are appreciable variations in the amplitude and phase of the measured waves in this region, which cannot be explained simply by variations in optical reflection (which is clearly not evident), coating thickness (there is no coating), or residual surface stress/porosity (since both the amplitude and phase are affected). In fact it is due to wavefront aberrations caused by material anisotropy in the form of random grain structures, and we have termed this effect *acoustic speckle*. This topic is worthy of an entire thesis in itself, but will be discussed in some depth in chapter 7.

Although some analysis of the amplitude and phase images shown in figure 6.5 can be useful, the true power of the measurement system occurs because the full complex amplitude of the SAWs is captured. The presence of individual modes can be shown by Fourier transforming the complex

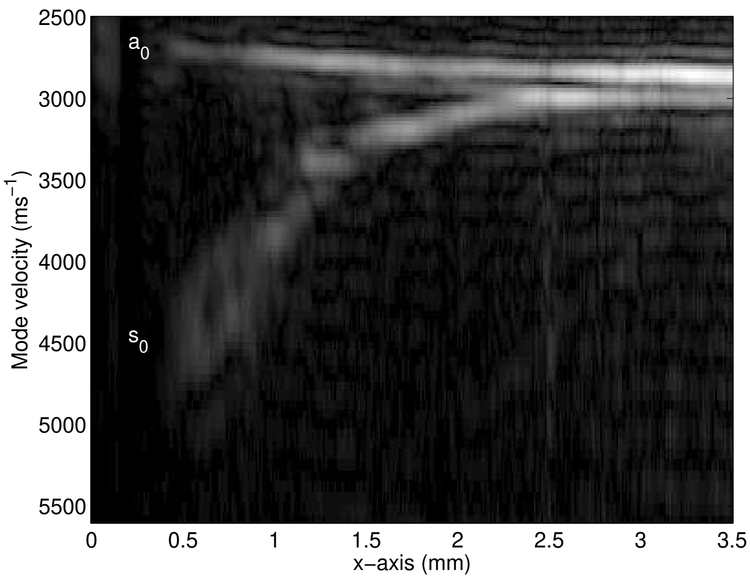

Figure 6.6: Fast Fourier transform (FFT) of the complex amplitude of the SAWs propagating over the slit shown in figure 6.5. The result has been transformed to show velocity rather than spatial frequency. The a_0 and s_0 modes are labelled.

amplitude distribution along the direction of propagation. This separates the velocity components in spatial frequency space and allows the velocity of each component to be measured. This is illustrated in figure 6.6.

The data used was the SAW complex amplitude on the area of thin aluminium, in other words just the slit itself. The FFT was taken along the direction of propagation. The y-axis was transformed to show velocity rather than spatial frequency, and the x-axis shows position along the length of the slit in millimetres. The a_0 and s_0 modes are marked on the image. On the left hand side, where the plate is thin, the velocities of the two modes are substantially different. This difference in velocities decreases toward the right hand side of the figure, where the plate gets thicker, and tends towards the Rayleigh velocity of 2945ms^{-1}. The general trend of the two modes agrees with the dispersion curves in figure 6.2.

The velocity data shown in figure 6.6 can now be used to ascertain the thickness of the plate along the x-axis. The velocity corresponding to the maximum amplitude for each mode—at each point along the x-axis—is used to calculate the thickness using the dispersion curves. The result is shown in figure 6.7.

The graph shows both the data points and two best-fit curves of the calculated plate thickness for positions along the x-axis. The two curves correspond to each of the two observable Lamb

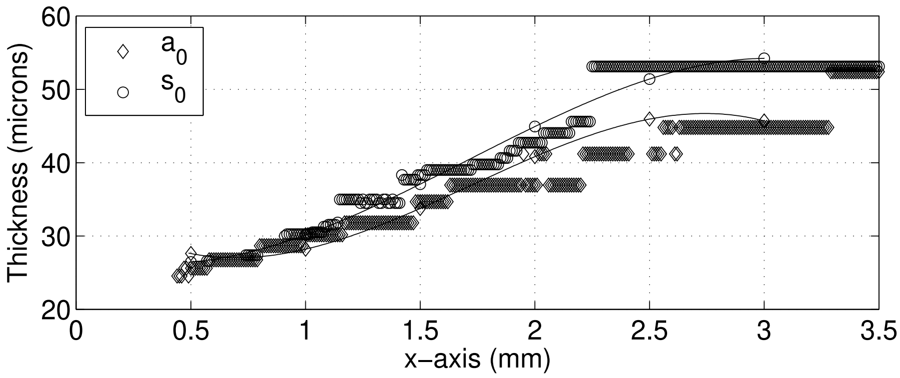

Figure 6.7: Calculated thickness of the slit at positions along the x-axis. The thickness is calculated by comparing the FFT data in figure 6.6 with the dispersion curves in figure 6.2. Because there are two observable modes, it is possible to calculate the thickness using either of them, and both are shown.

modes that were used to calculate the thickness. Not all the points along the x-direction are shown on the graph; a mask was applied to the data, such that if the amplitude of the chosen point on the FFT plot in figure 6.6 is below a certain threshold, then the thickness will not be calculated or plotted.

It is noted that the thickness as calculated using one mode does not match the thickness calculated using the other mode. This can be explained by differences in the physical constants used to calculate the dispersion curves, and the physical constants of the piece of aluminium on which the experiment was performed. An alternative explanation is that an error in the measured Rayleigh velocity would lead to errors in the 'measured' Lamb mode velocities, since this is used to translate the spatial frequencies into velocities. An 'adjustment' of just 20ms^{-1}—less than 1%—to the Rayleigh velocity would ensure that the thicknesses as calculated by the symmetric and antisymmetric modes would tally.

It is observed that the thickness appears to be discretized, seriously so for larger thicknesses. This is due to the maximum resolution attainable from the FFT process in converting the complex amplitude into spatial frequencies, which in turn is due to the number of points available in the data. Propagating over larger distances, or taking more data points in the y-direction would improve the results in this sense. The effect is more noticeable at larger thicknesses because the gradient of the dispersion curves, where the velocity is plotted with respect to thickness, is very shallow; a small change in measured velocity indicates a large change in calculated thickness.

However, even though the calculated thickness is discretized in this way, it is apparent that

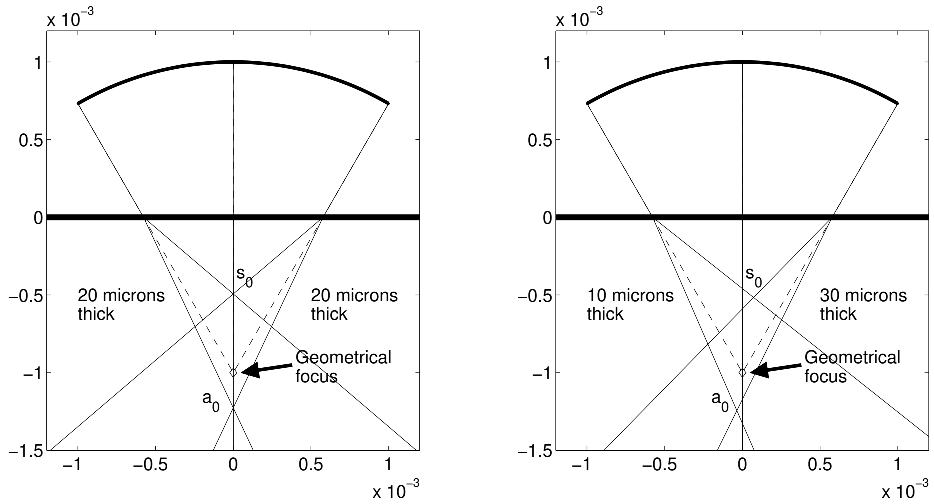

Figure 6.8: The left hand image illustrates the effect of attempting to focus SAWs though a boundary between uniform thick and thin layers. The faster s_0 waves come to a focus before the geometrical focus, whereas the a_0 waves come to a focus after it. The right hand image illustrates the additional effects of propagating through a medium of varying thickness along the x-direction; the foci are now displaced in the x-direction also.

the overall trend does not appear to be very smooth; overall, the best-fit lines are generally not straight. A possible reason is discussed in the following sections.

6.3.2 Effects of using focused waves

It has already been stated that the model used to generate the predicted response of the O-SAM system of imaging a plate of varying thickness is a relatively simple one. It takes no account of either the geometry of the slit itself, or of the geometries of the propagating surface waves.

Considering the first of these two points, when a plane wave of finite width propagates on a plate with varying thickness over the width of the wavefront, different parts of the wave travel at different velocities, and the wavefront becomes ‘tilted.’ The direction of propagation of the SAWs also affects the response. If the direction of wave propagation is not perpendicular to the boundary between the thick and thin regions of material, then refraction occurs due to differences in velocity between the Rayleigh waves and the Lamb modes. Figure 6.8 illustrates the concept.

The left hand figure follows the two outer ray paths of a focused acoustic wave as it refracts at a thick-thin boundary, where both regions have a uniform thickness. In this case, the thin layer has uniform thickness. The faster s_0 waves refract inwards, coming to a focus above (i.e. before) the geometrical focus. The slower a_0 waves refract outwards, coming to a focus after the geometrical

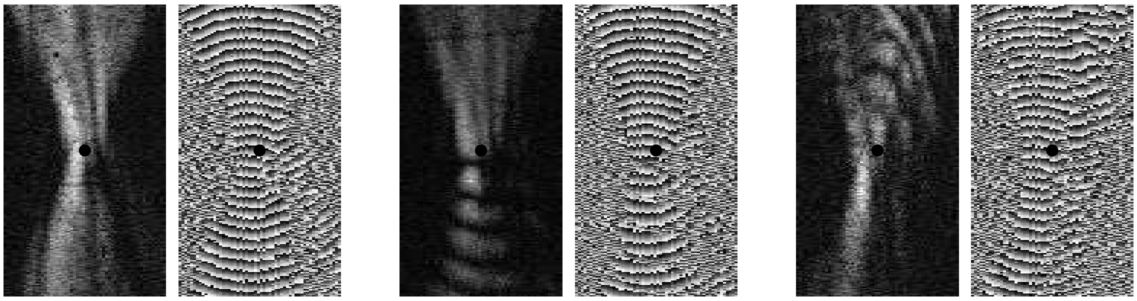

Figure 6.9: Amplitude and phase aerial PSFs at three scan positions for the c-scan shown in figure 6.5. The left hand PSFs are for a position above (i.e. before) the slit. The centre pair of images are for a point just on the boundary of the upper edge of the slit and the thick aluminium. The right hand PSFs are for a point on the slit itself. The small circle in the centre of each image indicates the location of the detection point.

focus.

The right hand figure illustrates the effect of the thin layer having a thickness gradient from left to right. The ray entering the thin layer on the left is subject to more refraction than the ray on the right, due to the greater change in velocity between the Rayleigh wave and each of the Lamb modes at that point. The two foci, according to the two outer rays of each, now occur displaced either to the left or to the right of the geometrical focus.

Since the detection point in the O-SAM system is located at the geometrical focus, it is unclear from figure 6.8 exactly what will be detected.

Figure 6.9 shows three aerial point spread functions (PSFs) at locations before the slit (i.e. on thick material), at the slit boundary, and on the slit itself. Amplitude and phase images are shown for all three PSFs. A small black circle in the centre of each PSF gives an indication of the location of the detection point, although the size of the circle has been exaggerated for clarity.

The left hand PSF, taken completely on the thick region of aluminium, is a relatively good approximation of an ideal PSF as shown in chapter 3. The focus is not quite symmetrical, and the reason for this is the effect of aberrations mentioned earlier in this section, that will be discussed in detail in the next chapter. The PSF in the centre of the figure is interesting, since its focus is at the boundary between the thick and thin regions. This is analogous to a point source—albeit with a limited field of divergence—on the boundary of a thick and thin plate. The dark interference fringes between the a_0 and s_0 modes can clearly be seen in the lower half of the amplitude image, and the wavefront is beginning to break up into two directions in the phase image. The right hand PSF is taken in a region that is entirely on the thin plate, and the foci of the two modes are now clearly spatially separate; the focus of the a_0 mode is to the lower left of the aerial PSF, and it is unclear where the focus of the s_0 mode is. The dark circle superimposed on the image is now no longer at the location of either of the two foci, and is in fact predominantly picking up the amplitude and phase from the side lobe of the focus of the a_0 mode.

It is suggested that this refraction of the two modes is disturbing the measurement of the complex amplitude, and may be partially responsible for both some of the inconsistencies between the measured response and the simple model (which does not take propagation direction, and hence refraction, into account), and possibly for some of the step-like changes in the calculated thickness.

6.3.3 '1° slit experiment'—using plane waves

An experiment was performed, very similar to the one described above in every way, with the exception that the SLM projected an image consisting of straight fringes onto the sample to act as the surface wave source, rather than an image of curved fringes. The source to detector remained the same, set at 2.5mm from the centre of the SLM image; in fact the relative positions of the source and detector were not changed in any way. The same area on the slit was scanned, at the same resolution. The results for this experiment—the '1° slit/plane wave' experiment—are presented along side the results for the '1° slit/focused wave' experiment for ease of comparison.

The amplitude c-scan results are shown in figure 6.10, and the phase results are shown in figure 6.11. As well as the results for the plane wave propagation, the simulated response and the focus results are also shown in both figures.

The third, rather dark image in figure 6.10 is identical to the second image in the same figure, with the exception of the way in which the data is scaled for presentation. In the second image, the maximum brightness corresponds to the maximum SAW amplitude. In the third image, the maximum brightness corresponds to the maximum SAW amplitude for the *focused* response—the bottom image in the figure—i.e. the two images are on the same scale. This graphically illustrates the advantage—in terms of signal amplitude—of focusing the acoustic waves and detecting them at a point, rather than propagating plane waves; the signal is approximately eight to nine times larger when focusing in this instance.

The plane wave amplitude and phase images obviously suffer from a lower signal to noise ratio, and this is reflected in the higher spatial frequencies present in the images. The acoustic speckle, briefly mentioned in section 6.3.1, is also present, although it appears to be more 'streaky' than 'blobby.'

It is noted that in the comparison of the amplitude images, the dark interference fringes appear much smoother, and do not appear to break up or join together to the same extent that the

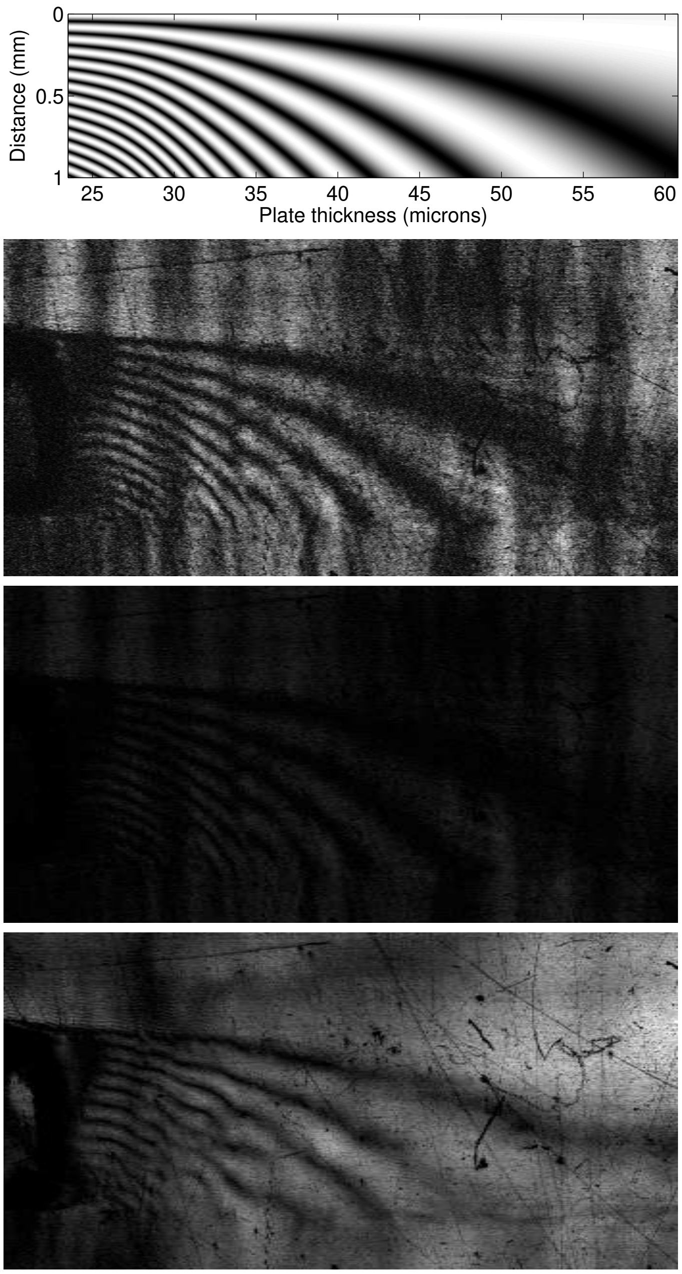

Figure 6.10: Amplitude response for Lamb modes propagating over a plate of varying thickness from left to right. The top image is the simulated response, also shown in figure 6.4. The second image is the response when the SAW excitation source is a set of straight fringes producing a flat wavefront. The third image is identical to the second image, except that the amplitude is on the same scale as the bottom image, which is the response when the SAWs are excited by a concentric series of arcs, producing a focusing wavefront.

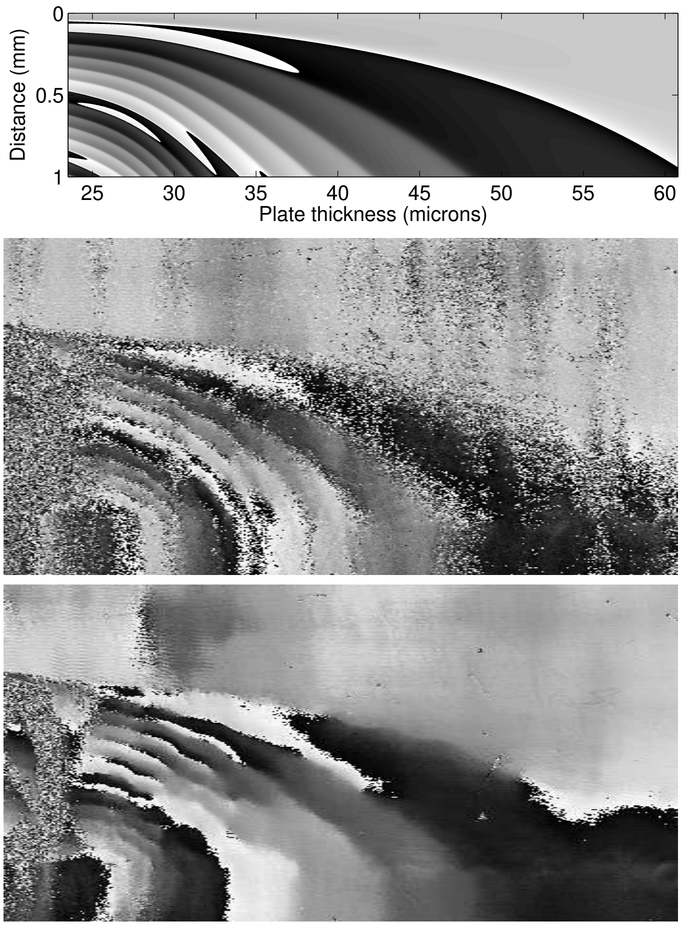

Figure 6.11: Phase response for Lamb modes propagating over a plate of varying thickness from left to right. The top image is the simulated response, also shown in figure 6.4. The centre image is the response when the SAW excitation source is a set of straight fringes producing a flat wavefront. The bottom image is the response when the SAWs are excited by a concentric series of arcs, producing a focusing wavefront.

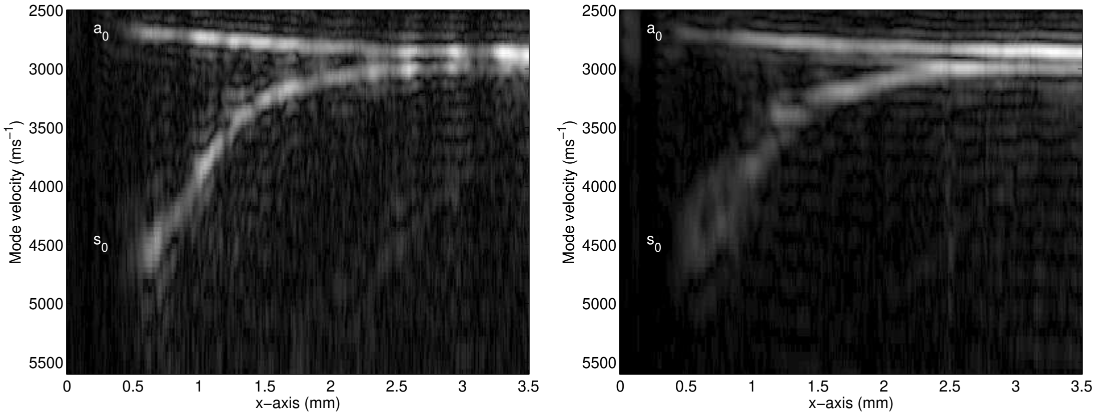

Figure 6.12: Fast Fourier transform (FFT) of the complex amplitude of the SAWs propagating over the 1° slit. The left hand image is from the data acquired by propagating a plane wave, the right hand image is from the data acquired by propagating a focused wave, and is identical to figure 6.6. The a_0 and s_0 modes are labelled.

amplitude response of the focused waves does. In fact only one fringe dislocation occurs on the plane wave image, and this could plausibly be due to the effects of the grain structure causing the acoustic speckle; it occurs on the edge of a 'dark speckle'. The plane wave phase image also appears to have a more regular geometrical structure, viewed purely qualitatively. The area below the slit—where the Lamb modes are converted back into Rayleigh waves—also show a more constant phase in the y-direction, indicated by the more vertical nature of the areas where the phase differs sharply along the x-direction.

As with the focused waves experiment, the complex amplitude data acquired on the slit was Fourier transformed along the y-direction to allow the velocity of each mode to be measured. This is illustrated in figure 6.12.

The data from the plane wave experiment is presented on the left, and the data from the focused wave experiment is presented on the right for comparison.

It is immediately apparent that the spatial frequencies along the direction of the FFT are much more clearly defined for the plane wave experiment. This is despite a *much* lower SNR for the complex amplitude data. In particular the location of the s_0 mode is much more precisely defined to the left of the image, where the slit is at its thinnest. In addition, the s_0 mode appears to be more continuous and smooth for the plane wave data.

As with the focused wave experiment, the data in the left hand image of figure 6.12 was used to calculate the thickness of the slit at various points along the x-axis by fitting the measured mode velocities to the dispersion curves. The result is shown in figure 6.13.

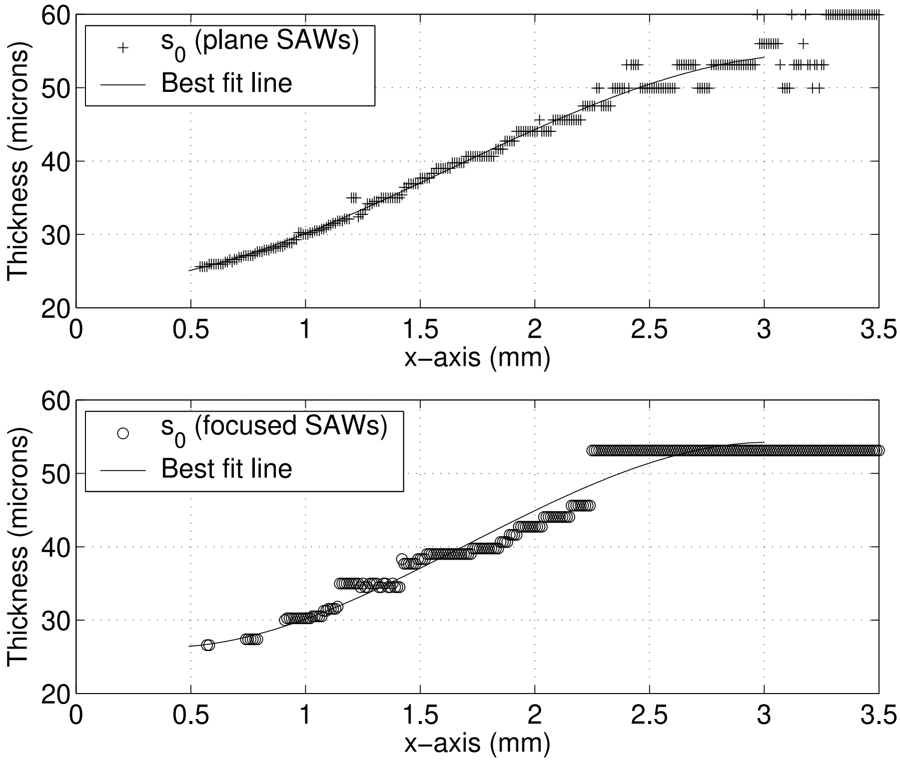

Figure 6.13: The upper plot is the calculated thickness of the slit using data from the plane wave experiment, the lower plot is the thickness calculated using data from the focused wave experiment. In both plots, the s_0 mode was used to calculate the thickness. Best fit lines are also shown on each plot.

The upper plot is clearly less discretized than the lower plot, and the calculated thickness profile is much smoother, and much more likely to be a true representation of the actual thickness. This is particularly true at the extremes of the plot, in regions where the plate was either comparatively thin or thick. It is noted that the best fit lines through both sets of data have very similar properties, however the maximum deviation is much greater in the lower plot.

6.3.4 'Circular profile slit experiment'—using focused waves

A similar experiment to the one carried out in section 6.3.1 ('1° slit experiment'—using focused waves), was performed on a sample with a slit of a slightly different geometry. Instead of a slit being milled into the rear surface of a piece of aluminium, a slit of circular geometry was cut into the rear of the sample, using a circular cutting disc. The diameter of the cutting disc used to make the slit is unknown, but it is of the order of 10cm. The width of the slit was approximately 2mm. Figure 6.14 indicates the experiment geometry. Once again, the vertical dimension in the cross-section diagram has been greatly exaggerated.

This experiment was carried out at a considerably earlier time than the 1° slit experiments, and the source of the surface waves was a set of 16 concentric arcs with a mean focal length of 2mm, projected onto the sample using a computer generated hologram (CGH). The half-angle of the arcs was 30°. Nevertheless, the rest of the details of the experiment were the same. The c-scan images obtained for the experiment are illustrated in figure 6.15.

As before, the optical image, acquired at the same time as the complex amplitude of the SAWs at 82MHz, is shown in the top image. The slit is completely optically invisible from the top surface, as the slit did not break the surface.

Two features of the amplitude image that should be mentioned are the dark patch in the upper left hand corner, and the horizontal bright line to the right of this. Both these 'features' are the result of the occasional random oscillations that plagued the high frequency electronics at the time the images were acquired. They do not reflect any actual change in amplitude (or phase) of the SAWs.

Moving on to more general observations, the amplitude and phase images are beautiful—if strange—images. They are also a little tricky to immediately interpret. The amplitude c-scan clearly shows dark interference fringes of varying pitch, indicating propagation of Lamb modes over a thin area of varying thickness. The region 'behind' the slit, i.e. in the lower half of the image, has much reduced amplitude. This can be explained by the fact that the source to detector distance is on average 2mm, which is also the width of the slit. Hence, for regions up to 2mm behind the slit, the source of the SAWs—sixteen concentric arcs imaged onto the sample using a

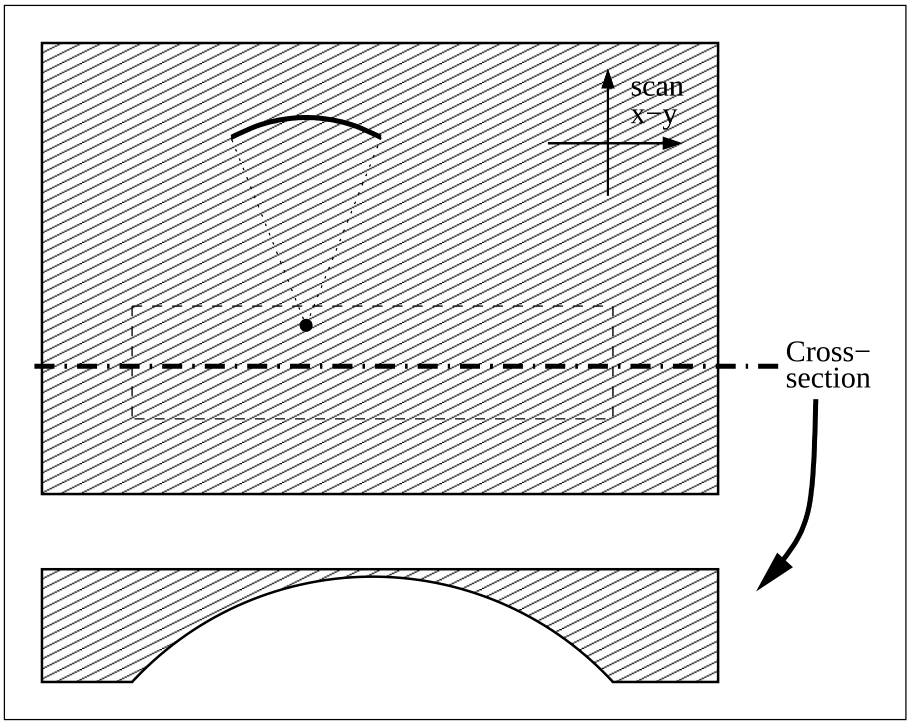

Figure 6.14: Schematic of the 'circular profile slit experiment.' A slit is cut into base of a piece of polished aluminium such that the thickness profile follows the radius of a circle. The sample is scanned in the x-y directions to achieve a normal c-scan. The vertical dimension and slit angle in the cross-section diagram have been greatly exaggerated.

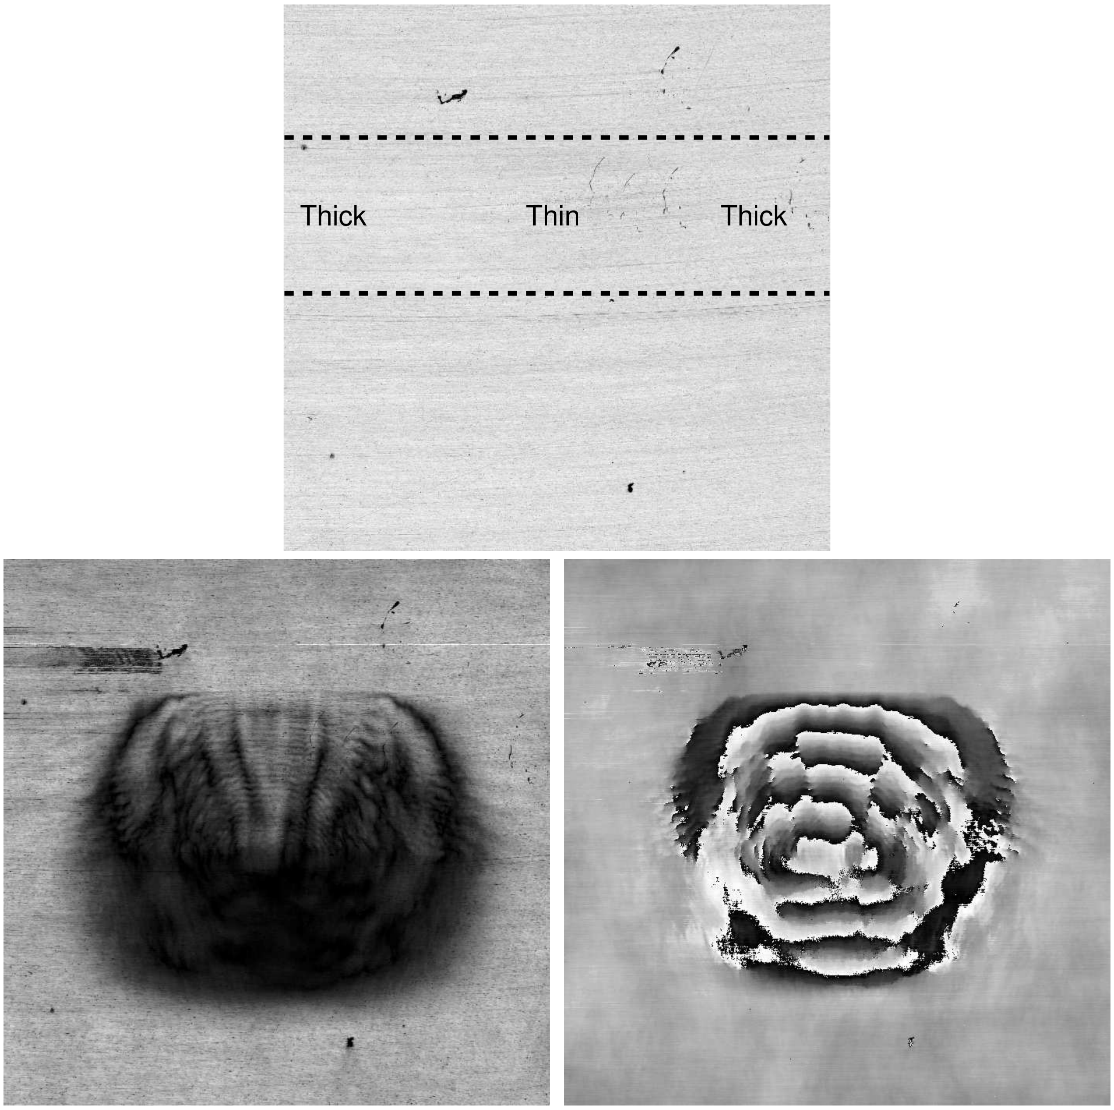

Figure 6.15: Optical (top), 82MHz SAW amplitude (lower left) and 82MHz SAW phase (lower right) images of a 7 × 7mm area of the top surface of a piece of polished aluminium. The location of the slit is shown by dashed lines in the optical image. The slit is thinest at the centre (horizontally) of the image.

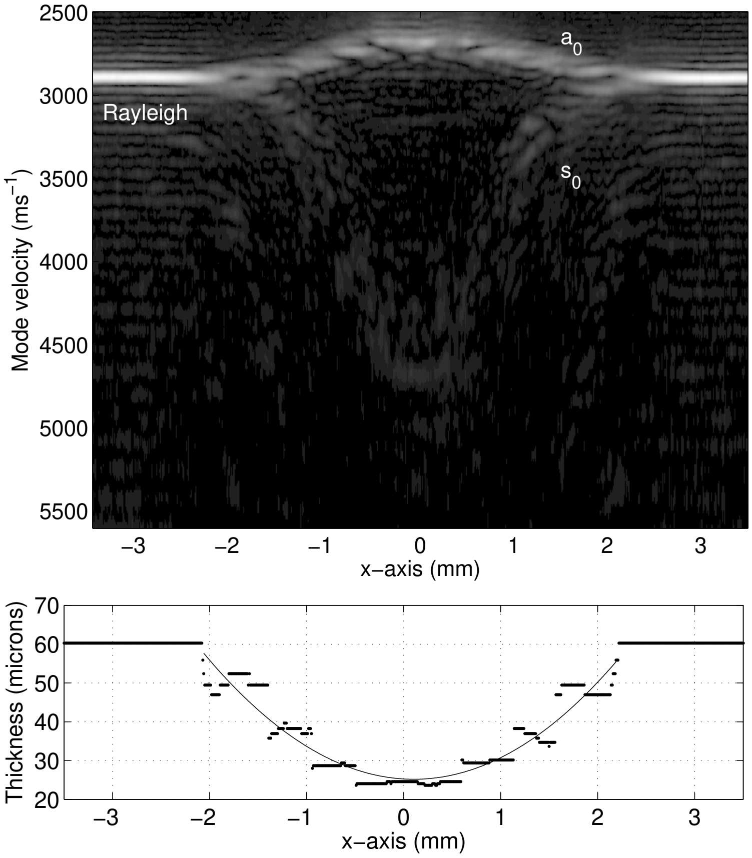

Figure 6.16: The upper plot is a representation of the Lamb mode velocities at different points along the slit shown in figure 6.15, calculated by FFT analysis of the complex amplitude over the slit area. The lower plot shows the thickness as calculated from the a_0 mode.

CGH—are on the slit itself. Since the spacing of the arcs is set for generation of 82MHz Rayleigh waves, which have a different velocity from either of the fundamental Lamb modes, then very inefficient SAW generation will take place, leading to the dark 'shadow' behind the slit.

However, SAWs are clearly being generated, since complex phase changes are observed in this region. This is because, even though the detector is not on the slit itself, the phase is affected by all the material between the source (the arcs) and the detector which, up to 2mm behind the slit, includes a portion of the slit itself.

Once again, the complex amplitude data acquired on the slit itself was Fourier transformed along the y-direction to allow the velocity of each mode to be measured. From the FFT image, the thickness of the slit along the x-direction was calculated by fitting the measured mode velocities to the dispersion curves. Figure 6.16 shows both the FFT and the thickness plot.

Considering first of all the FFT plot, one can see three modes in total—the Lamb a_0 and s_0 modes, and also the Rayleigh mode at the left and right hand sides. The s_0 mode is very weak in this sample.

The thickness plot, which was calculated using the stronger a_0 mode, is very interesting, in that the calculated thickness consists of three or four very steep steps on each side of the thinnest portion of the slit. These step levels are not due to the discretisation of the thickness due to the FFT resolution, they simply follow the maximum amplitude of the calculated a_0 mode on the FFT plot. It is assumed that the calculated thickness profile does not follow the true thickness profile, due to knowledge of the sample preparation technique. One can conclude that the fact that the SAWs were focused to a point has certainly contributed to this, as it did in the focused wave 1° slit experiment.

The next logical step in the investigation would be to propagate plane waves across the slit, as was done for the previous experiment. This was not possible at the time the original data was obtained because it was not possible to propagate plane waves of sufficient amplitude that they could be detected with the analogue vector contrast imaging system, at least not without damaging the sample due to ablation. It has not been possible to redo the experiment using the SLM to image the surface wave generation source, because the sample has since been damaged.

6.3.5 Propagating along the slit

An experiment was performed on the same sample as the previous section, i.e. the aluminium plate with a slit with a circular profile cut into the back of it. Focused SAWs were propagated over the slit, however this time the waves were propagated in a direction perpendicular to the direction of propagation in the previous experiment—the waves were propagated along the slit, rather than across it. Since the thickness of the slit now varies along the direction of propagation, the results are much more difficult to analyse quantitatively.

The amplitude and phase images acquired, along side amplitude and phase images acquired by propagating the SAWs in a perpendicular direction, are shown in figure 6.17. Note that the sample orientation has been fixed, whilst the direction of SAW propagation has now changed. The images cover roughly the same area, and the images may be lined up using the optical defects—scratches—apparent in both amplitude images.

Analysis by FFT along the direction of propagation will not yield much in the way of useful information. There is likely to be a large range of velocities along the duration of the slit, as each of the modes speeds up or slows down according to the plate thickness, which is constantly changing along the propagation direction in this experiment.

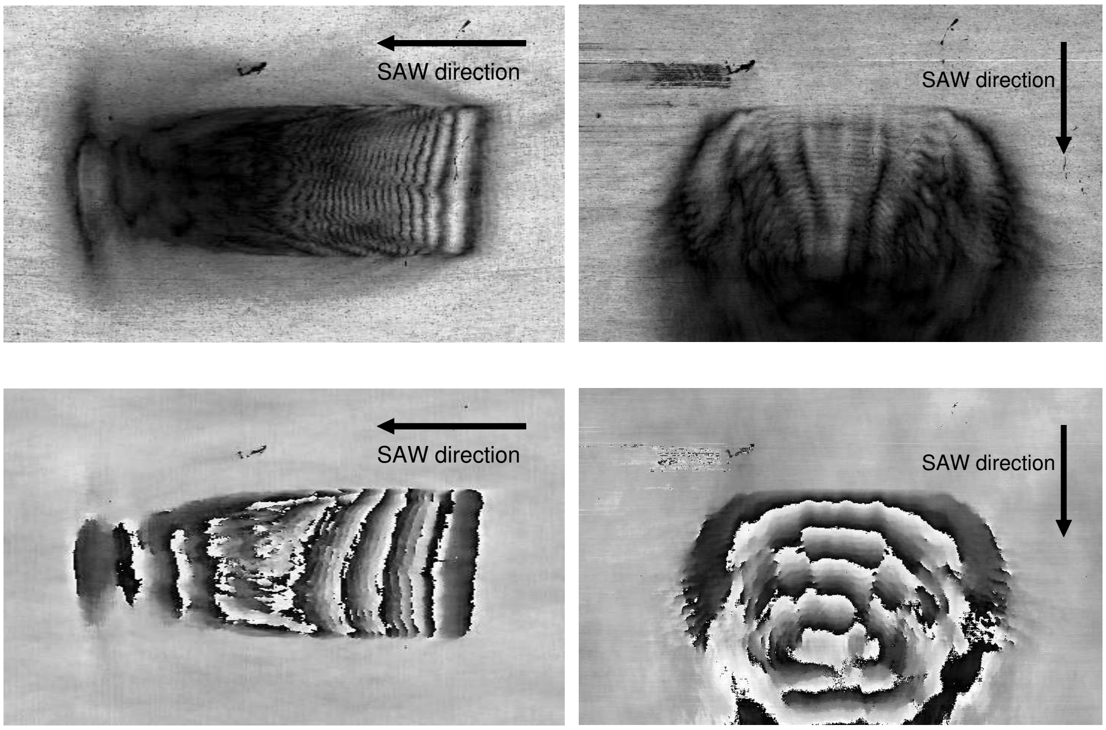

Figure 6.17: SAW 82MHz amplitude (top) and phase (bottom) images of the 'circular profile slit' shown in figure 6.15. In the left hand images, the SAWs are travelling from right to left, as indicated by the arrow. In the right hand images, the SAWs are travelling from top to bottom. The left hand images are 7.5 × 4.5mm, the right hand images are 7 × 4.5mm.

It would be interesting to investigate the guiding properties of the slit on each of the Lamb modes. Some evidence of this appears in the distorted wavefront of the propagating waves, and in the reduction in amplitude of the Rayleigh waves to the top and bottom of the slit, although this could certainly be the result of propagating to a focus, rather than propagating a plane wavefront.

The waves are travelling from right to left in the images, and there is a distinct overall variation in the detected amplitude of the Lamb modes from right to left. This is due to inefficiencies in the generation mechanism, since the SAWs are generated using 16 concentric arcs of mean radius 2mm, imaged onto the sample using a CGH. These arcs have a spacing designed for the wavelength of Rayleigh waves on the material, and as the excitation source moves onto the slit, the generation efficiency will decrease. This could easily be achieved using the spatial light modulator to image a 'custom' fringe pattern for regions where the excitation source is on the thin plate. One or more modes could then be generated by adjusting the SLM image.

6.3.6 'Hole experiment'

In a similar process to the fabrication of an aluminium sample with a slit milled into the rear surface at a 1° angle, a series of 1mm diameter holes were drilled into the back of a sample. The holes were drilled to a fixed depth with respect to the horizontal, but because the sample was seated at a 1° angle, the holes were at different depths. The holes were drilled by a milling bit with a flat end, and so the holes were reasonably flat-bottomed, albeit at a 1° angle. After polishing the front surface, some holes broke through, whilst others remained optically invisible from the front surface. One such hole was examined using the O-SAM instrument. Figure 6.18 is a representation of the schematic of the experiment.

At the time the experiment was performed, the analogue vector contrast detection system had not been built, although an analogue amplitude detection system had. This worked by rectifying the amplified signal from the knife edge detector, passing it through a low pass filter, and then detecting the peak. The upshot of this, is that only the SAW amplitude image is available.

The source of the surface waves was a set of 16 concentric arcs with a mean focal length of 2mm, projected onto the sample using a computer generated hologram (CGH).

The image on the left hand side of figure 6.19 shows the 82MHz amplitude c-scan of the hole. A dashed circle has been superimposed on top of the image to indicate the approximate edge of the hole. The image on the right of the figure shows the predicted amplitude response for a 1mm diameter hole of uniform thickness and 36μm deep.

The fringes are formed by interference between the a_0 and s_0 modes; however their pattern is made more complicated by the refraction at the interface between the thick and thin regions. The

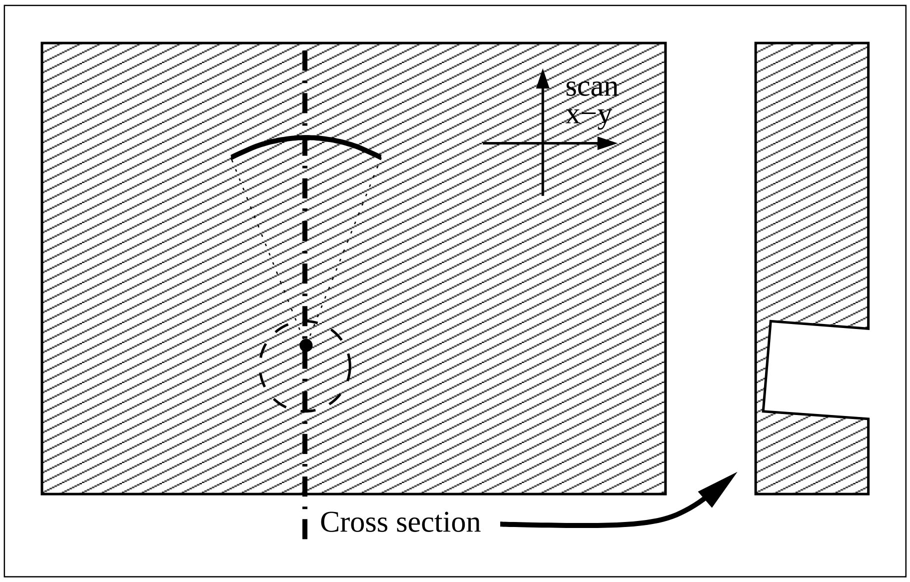

Figure 6.18: Schematic of the 'hole experiment.' A 1mm diameter hole is drilled into base of a piece of polished aluminium at an approximate 1° angle. The sample is scanned in the x-y directions to achieve a normal c-scan. The vertical dimension and slit angle in the cross-section diagram have been greatly exaggerated.

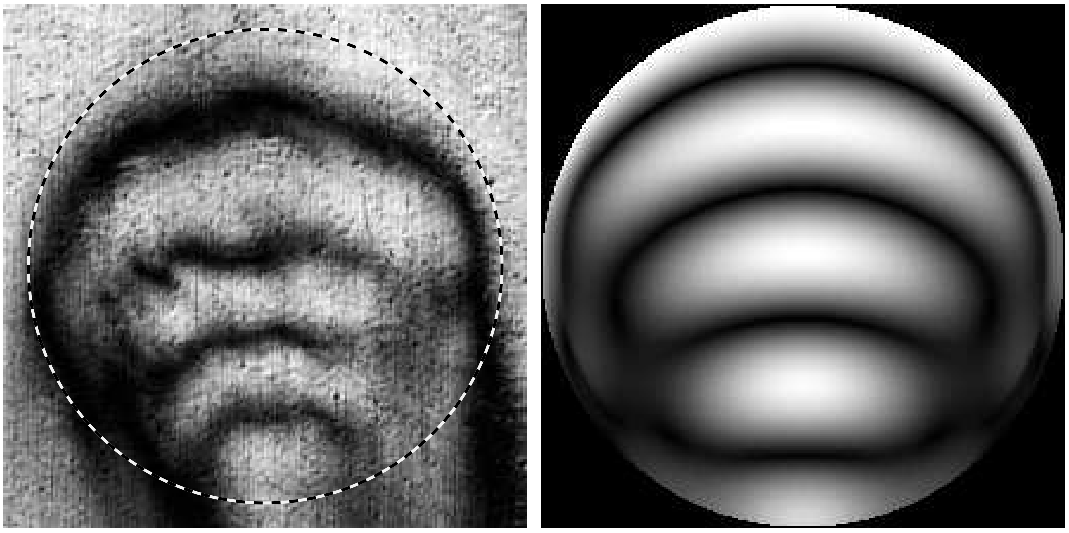

Figure 6.19: The left hand images in an 82MHz SAW amplitude c-scan of a 1.1×1.1mm area of polished aluminium, under which lies a 1mm diameter hole. The approximate location of the hole is marked by the dashed circle. The right hand image is the predicted amplitude response.

a_0 mode is refracted towards the normal, due to its velocity being lower than the Raleigh velocity, and the s_0 mode is refracted away from the normal for the opposite reason. The interference pattern is a function of the semi-angle of the excitation arc, the thickness of the material, and the hole diameter.

The amplitude response was modelled using a ray-tracing approach, which associated a phase to each ray. These were then combined for each point within the area of the hole.

The amplitude response was modelled using a ray-tracing approach, as follows. Each point within the circle defect corresponds to a point on a c-scan image, and for each of those points the excitation arc will be at a fixed point relative to the detection point. The points at which a number of 'rays' from around the excitation arc intersect with the edge of the circle are calculated and, since the distance travelled from arc to intersection point is known, their phase at the intersection point can be calculated. The rays are then refracted in two directions, corresponding to the two Lamb modes. The phase along each of the two rays, for all the rays emanating from the excitation arc, may be calculated for the given point on the circle, since it is the phase at the point along the ray that is perpendicular to the given point with respect to the propagation direction. The contributions of all the rays are then summed, to give the relative amplitude of the detected SAW at each point within the circle as shown.

The experimental data and the model bear a good resemblance to each other, although there are differences. Much of the difference in appearance can be attributed to the fact that the bottom of the hole is at a 1° angle from the top of the image to the bottom. The fringe spacing is greater at the top of the image (implying a thicker plate) than the bottom. Since the model was simulated with a plate thickness of 36μm, it would be fair to estimate that the thickness of the material varied from around 40μm to around 35μm from top to bottom in the image. Nevertheless, the model gives a good prediction of the response, including regions where there is a loss of fringe contrast, such as to the lower left and lower right, close to the edge of the circle.

6.4 Determination of subsurface features by Lamb mode dampening

The previous section described various experiments whereby Lamb waves were used to calculate the thickness of a thin plate. It is a powerful technique, made even more so by the range of contrast mechanisms. For instance, the fringe spacing in the amplitude response can itself be used to calculate the thickness, or the velocity of either of the two Lamb modes obtained by spatial Fourier analysis.

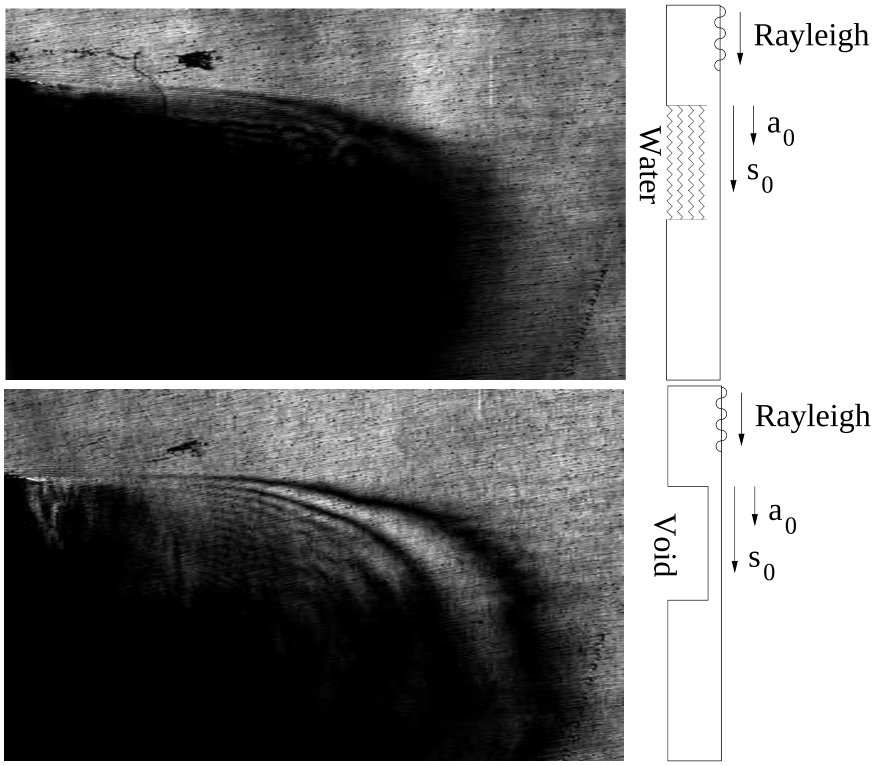

Figure 6.20: The upper 82MHz SAW amplitude image is of a 1° sub-surface slit, with the void behind the slit filled with water. At least one of the Lamb modes is severely dampened in this case by the water. The lower image is of the same region on the sample, but with the water drained out.

An experiment was performed to determine whether the O-SAM instrument was capable of determining not only the thickness of a plate, but also if the nature of any material on the 'rear side' of the plate could be inferred. An additional aim of the experiment was to give some kind of indication of the advantages of the O-SAM instrument compared to traditional fluid-coupled SAMs.

A sample, very similar to the one used in the '1° slit experiment' was imaged using the O-SAM. At the time, only the SAW amplitude could be detected, as with the 'hole experiment' described in the last section. The difference between this experiment and the previous slit experiments is that the void milled out to form the slit was filled with water. The area around the thinnest portion of the slit was then imaged. The water was released from the void on the rear of the sample, and then the same area was imaged again.

The results of the experiment, as well as schematic diagrams illustrating the cross section of the slit at the time the two images were take, are illustrated in figure 6.20. The upper image shows the

82MHz SAW amplitude c-scan taken with the water-filled void, the lower image shows the same region without the water.

The Lamb modes in the water-backed plate are heavily attenuated by the presence of the water, and the fringes are seen to die out quickly across the void. At least one of the modes is heavily damped by the presence of the water on the rear surface. Contrast is severely reduced, and it would be difficult to use the information for any quantitative assessment of the plate profile, even if the full complex amplitude was available, especially for the thin regions. After the water was drained and the area around the slit re-imaged, the usual contrast between the two fundamental Lamb modes can clearly by seen.

The experiment illustrates that not only is the O-SAM instrument capable of measuring the thickness of a thin plate by Lamb wave contrast, it is also capable of inferring the nature of the substance on the rear plate. It also illustrates one of the advantages of the system over a fluid-coupled SAM; since a liquid on the rear surface of the plate severely dampens at least one of the Lamb modes, it is sensible to assume that the coupling fluid on the front surface of the sample would have a similar effect.

6.5 Summary

This chapter has dealt with the use of the O-SAM instrument to perform quantitative and qualitative assessments of thin plates using Lamb waves. Dispersion curves, common in the literature and relatively easy to calculate, may be used to measure the thickness of the plates with good accuracy. The accuracy depends on the characteristics of the dispersion curves at the given plate thickness—the greater the slope of the curves, the greater the accuracy. Different modes may be used to measure different plate thicknesses, the higher order modes being more suitable to thicker plates. The acquisition of the complex amplitude of the surface waves allows information about all the present modes to be extracted by spatial Fourier analysis. Simple models have been used to predict or verify results, or provide a means of quantifying difficult to interpret patterns (in the case of the circular defect). An investigation into the effects of using focused acoustic waves, as opposed to plane waves, has been carried out, and has led to an understanding of how this affects the measurements taken with the instrument. The nature of substances on the rear surface of samples can be inferred from whether any loading of that surface occurs, and to what extent this affects the attenuation of one or more Lamb modes. Contrast has been demonstrated that would be difficult, if not impossible to achieve with a fluid-coupled scanning acoustic microscope, due to the lack of couplant, and also the flexible geometry of the system, and the ability to propagate over relatively long distances if necessary (several millimetres) to allow for greater contrast, either with focused acoustic waves or, if the signal to noise ratio is sufficiently high enough, with plane waves. In addition to the advantages over contact acoustic methods, there are of course the usual advantages over contact probe profilometers in that because the O-SAM is completely noncontact, it is totally damage-free and can potentially be used in harsh environments, such as with very hot materials.

In addition, the issue of acoustic aberrations and acoustic speckle has been touched upon very lightly, and this will be the subject of the following chapter.

Finally, by the illustration of the many images obtained in this chapter, the true capabilities of the instrument as a tool for presenting data in a manner in which analysis is easier, and detail is enhanced, have been shown. Many of the images presented contain 100,000 points or more, and were acquired in minutes, rather than hours or days. Not only are some of the images produced truly beautiful in their own right, it is a powerful enabling tool to explore both the capabilities and limitations of the technique, and to examine interesting features of samples, materials, and the propagation of surface waves on them.

Chapter 7

Adapting to the nature of the material

7.1 Introduction

This chapter deals with the subject of the effects of material anisotropy, in particular that caused by random grain structure in materials such as aluminium and steel. The effect this can have on measurements and images obtained by the O-SAM instrument has already been briefly mentioned in section 6.3.1 in the examination of an aluminium sample. An explanation of the cause of the aberrations and the effects on surface waves is followed by a discussion of techniques to reduce or eliminate the problem.

7.2 Acoustic speckle - effects of anisotropy

The effects of material anisotropy on the quality of the data acquired by the O-SAM instrument first became apparent soon after the first high speed analogue data acquisition system was developed. Figure 7.1 shows two examples of early attempts to perform SAW amplitude c-scans on aluminium.

The left hand image in the figure is of a 'hole defect' that breaks the surface in the lower third of the image; the right hand image is of a '1° slit' defect, very similar to the one described in section 6.3.1. In both the images, SAW reflection fringes may be seen at the edge of the defect, and fringes caused by interference of two Lamb modes on the thin material are observable in the right hand image. It is immediately apparent that there is significant degradation in the quality of the images, with dark patches or speckles present at random positions across the images. These speckles are caused by the anisotropy and grain structure of the sample.

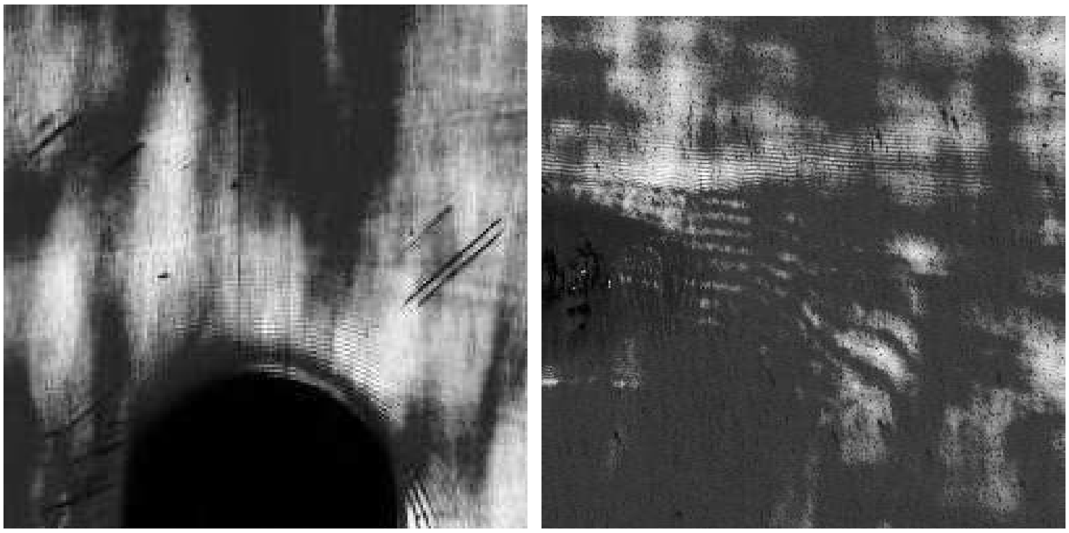

Figure 7.1: Two 82MHz SAW amplitude images illustrating the effects that material anisotropy has on the quality of the c-scan images obtained with the O-SAM instrument.

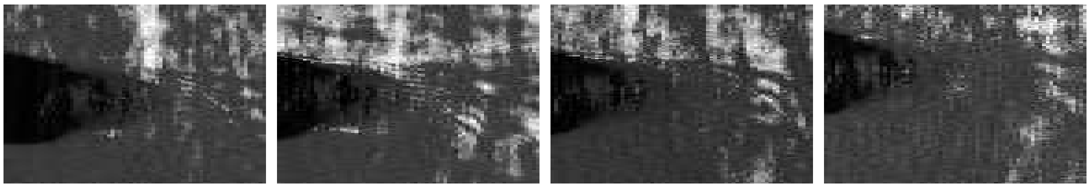

Figure 7.2: Four amplitude images of the same region on an aluminium sample. The only difference between the images is the location of the detection point with respect to the SAW generation region during the c-scans.

Figure 7.2 shows four SAW amplitude images of the same area as the right hand image of figure 7.1; apart from the images being of slightly different areas, the only difference between the four images is the position of the detection point with respect to the SAW generation region. In each case, the location of the dark speckles, which have a distinctly ‘cloudy’ feel to them, is different.

7.3 Wavefront aberration

The problem is similar to the problem of optical imaging through a turbulent atmosphere [88], or misplaced sonar readings due to fluctuations in the density of water [89], although in the case of SAW imaging on solids the aberration is fixed in time and changes when the sample moves. The signal drop-out occurs not because the SAWs are absorbed in the material in the dark regions, but because the acoustic wavefront is being aberrated as it propagates through the randomly-orientated grains of the material. It is aberrated because in anisotropic media, the velocity of the acoustic wave is dependent on the grain orientation.

The effect this has on the acoustic wavefront depends on the size of the grains (relative to the acoustic wavelength), the propagation distance, and the degree of anisotropy. For single crystals, "geometric" aberration is produced, with no scattering losses. With knowledge of the slowness surface of the material, the aberration is predictable and in many cases correctable, and will remain constant no matter which part of the crystal is examined, provided the orientation remains the same. Scattering and attenuation occurs when ultrasound propagates through materials with grains of similar or smaller dimensions to the acoustic wavelength. In general, the wavefront propagates to the expected destination, albeit attenuated. When the grains are longer than the acoustic wavelength moderate attenuation occurs, but the principal source of signal loss and experimental error in this regime is aberration of the wavefront. Different parts of the wavefront travel through different grains with different velocities, and the wavefront has a tendency to either end up in a different location to that expected, or to break up altogether—hence the degraded images in figures 7.1 and 7.2.

It is difficult to visualise the distortion of SAW wavefronts as they propagate using contact techniques, although it had been observed as early as 1973 using a hybrid contact/noncontact technique [44]. Here the wavefronts of SAWs propagating on titanium were seen to suffer from "severe distortions and interference effects." It noted that when the sample was subjected to ultrasonic pulse echo inspection techniques, it produced "a multitude of small reflected signals which gave the appearance of a high noise floor." It is this effect that is normally noted in the literature. The whole phenomenon of acoustic aberrations due to material inhomogeneities has led to an interest in the time reversal of ultrasonic fields [90, 91] to attempt to correct for the aberrations; with time reversal techniques, however, not only are you limited by the aperture of the time reversal mirror, but more importantly you only see the effects of the inhomogeneities on the wavefront after it has been aberrated, not as it is propagating.

The O-SAM is an excellent instrument with which to observe this effect, since the amplitude and phase of the surface waves can be imaged as they propagate from the excitation region to the point where the detector would normally be, as we have already seen in the previous chapter. Figure 7.3 illustrates the concept of how wavefront aberration affects focused surface acoustic waves.

The left hand images illustrate the propagation of focused SAWs on an isotropic medium—the lower left hand image is an aerial PSF taken on aluminium-coated glass. The location of the intended SAW detection point is also marked in the figures, although its size has been exaggerated for clarity. The right hand images illustrate the propagation of focused SAWs on a multi-grained

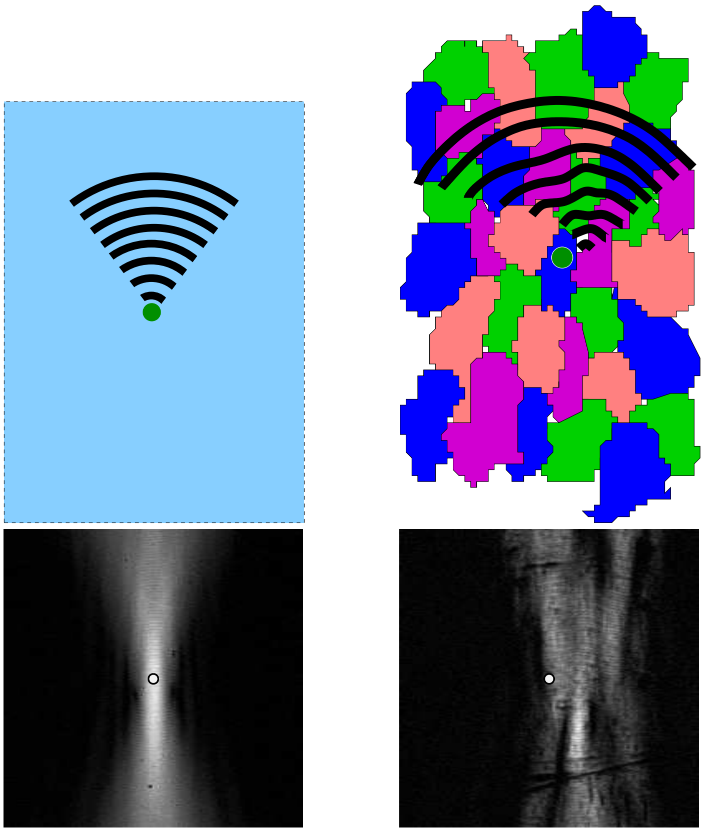

Figure 7.3: The left hand images illustrate the propagation of focused SAWs on an isotropic medium—the lower left hand image is an aerial PSF taken on aluminium-coated glass. The right hand images illustrate the propagation of focused SAWs on a multi-grained anisotropic medium—the lower right hand image is an aerial PSF taken on aluminium

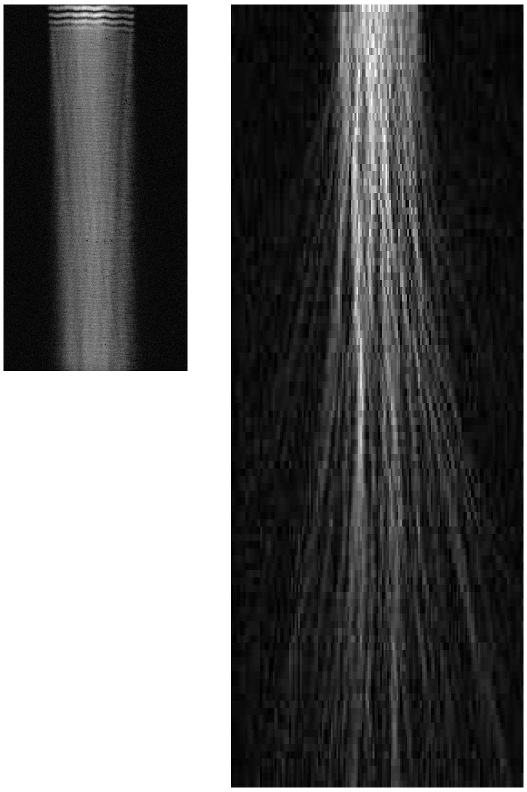

Figure 7.4: The left hand image shows the amplitude of 82MHz SAW plane waves propagating along the surface of aluminium-coated glass. The right hand image shows plane waves propagating along the surface of hardened steel.

anisotropic medium—the lower right hand image is an aerial PSF taken on aluminium. Not only does the majority of the acoustic energy miss the geometric focus of the generation arcs, but the point spread function has clearly broadened and dispersed.

This type of effect is not limited to focused waves, as plane waves are also aberrated. Figure 7.4 illustrates the effect.

The left hand image shows the amplitude of 82MHz SAW plane waves propagating along the surface of aluminium-coated glass, over a distance of 10mm. No discernible aberration occurs, and the wavefront continues to be flat for the entire length of the propagation distance observed. The right hand image of the same figure shows plane waves propagating along the surface of hardened steel, over a distance of 20mm. The wavefront can clearly be seen to break up as it propagates; caustics appear as the wavefronts refract at grain boundaries, and the energy spreads out spatially.

This aberration of plane wavefronts has obvious implications for the measurement of SAW velocities in aberrating materials. Since the wavefront clearly breaks up over distance, then any

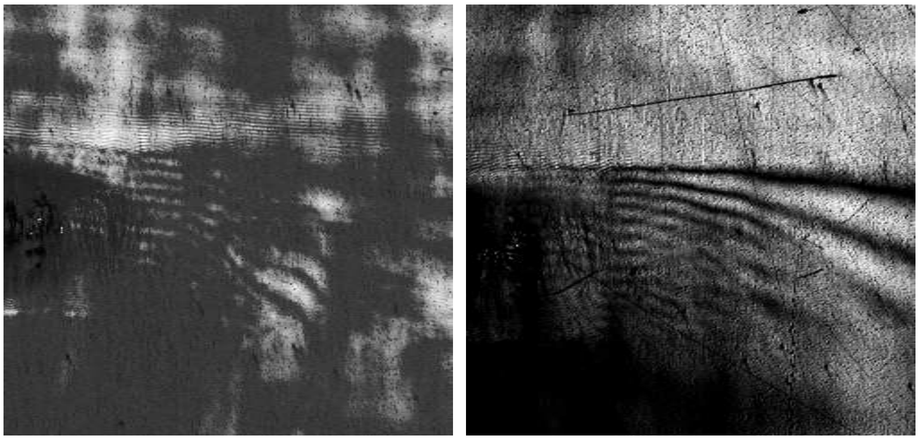

Figure 7.5: The two images are 82MHz SAW amplitude c-scans of the same region on the same sample, the only difference between them being the focal length of the generation arcs. The focal length is 10mm in the left hand image, and 2mm in the right hand image.

measurement of the phase of the SAWs over that distance is likely to be affected by the aberration. This effect was noted in the 'early days' of the O-SAM instrument, before high speed imaging was possible and when only point-by-point measurements were taken. On attempting to measure the velocity of SAWs on a sample similar to the one imaged in figure 7.4, the measured velocity was found to vary significantly, depending on where exactly the optical detection point was located. Without the ability to acquire rapid, high resolution images, the cause is difficult to ascertain.

7.4 Correcting for the aberration: shorter focal lengths

In terms of c-scan imaging using focused surface waves, a simple method of reducing the effects of aberration is to reduce the distance travelled by the SAWs between source and detection regions; this is done by reducing the focal length of the generating arc profile. This is illustrated in figure 7.5.

The two images 82MHz SAW amplitude c-scans of the same region on the same aluminium sample. The source of the SAWs in the left hand image is a set of 4 concentric arcs with a mean focal length of 10mm, imaged onto the sample using a computer generated hologram. The surface wave source in the right hand image of figure 7.5 is a set of 16 concentric arcs with a mean focal length of 2mm, again imaged onto the sample using a CGH. The amount of acoustic speckle is considerably reduced in the image with the shorter focal length. However, this approach may not be suitable for all samples and all experiments, since one of the key advantages of the O-SAM over

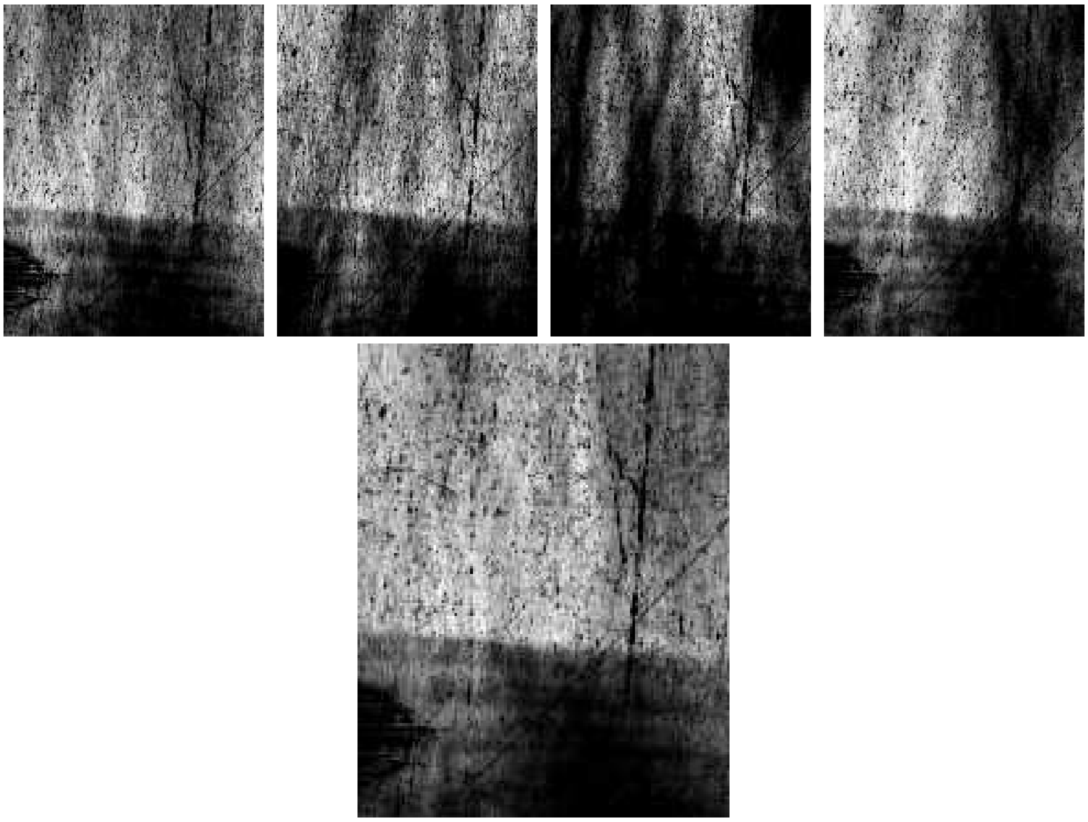

Figure 7.6: The four upper images are 82MHz SAW amplitude c-scans of the same region near a defect on a piece of aluminium, the only difference between them being the location of the detection point with respect to the generation region. The lower image is the summed average of the top four images.

the contact SAM is the increased flexibility in propagation geometry. As discussed in chapter 6, the accuracy in calculating plate thickness using Lamb mode velocity is increased by propagating over larger distances. At the other extreme, the propagation distance cannot be made arbitrarily short, since the heating caused by the generating source may perturb the measurements at the focus.

7.5 Averaging/Multiple detection points

In certain circumstances, it may be appropriate to spatially 'average out' the acoustic speckle noise. Figure 7.6 illustrates the technique.

The upper four images in the figure are all 82MHz SAW amplitude images of the same area, on a piece of aluminium with a defect in the lower third of the image. As with figure 7.2, the difference between each of the four images was the location of the detection point with respect to the arc generation source; that is, for each c-scan the relative positions were kept fixed and the sample was scanned, but between each scan the detection point was moved slightly, to detect different acoustic speckles. The images were then spatially lined up using features common to all the optical images that were acquired at the same time, and the amplitude images were summed to produce the lower image of figure 7.6.

Although the image contains a lower degree of acoustic speckle than any of the other four images in the figure, a certain amount of resolution is lost in the spatial averaging process. In a way, the process is similar to reducing the frequency of the SAWs, in that there is a reduction in resolution and an increase in the tolerance to material anisotropy; however the frequency—and hence the penetration depth of the SAWs—has not been changed, and so they are still sensitive to the material closest to the surface of the sample, which may be the most important issue in some circumstances.

7.6 Source tilt correction

The previous two subsections detailed passive measures to help reduce the effects of material anisotropy on the measurements taken. In each case the source of the ultrasound is of a fixed profile for each experiment undertaken, and the gain in data consistency is balanced by the loss of either spatial resolution in the case of averaging, or the detectable change in velocity in the case of reducing the focal length. By actively correcting for the aberration, resolution and desired propagation geometry may be maintained.

The simplest, first order control over the propagation of the acoustic wavefront can be achieved by 'tilting' the generation source. Figure 7.3 illustrated that the focus of acoustic waves propagating on the surface of an aberrating material can occur in a location adjacent to the detection point, resulting in 'dark clouds' on the c-scan image. If the source could be tilted slightly, then the acoustic energy would be redirected back to the detection point.

To correct for the wavefront distortion introduced by the random grain structure of a material, it is first necessary to acquire information about the acoustic wavefront at the intended detection point. We can obtain the information required for tilt correction by scanning the detection probe across the PSF for a selection of detection points. Figure 7.7 illustrates the method.

The left hand image of parts *a* and *b* of the figure shows an 82MHz SAW amplitude c-scan image of an area of aluminium, with a schematic of the generation arcs superimposed. The white horizontal line at the focus of the arcs indicated the location across which a PSF was acquired.

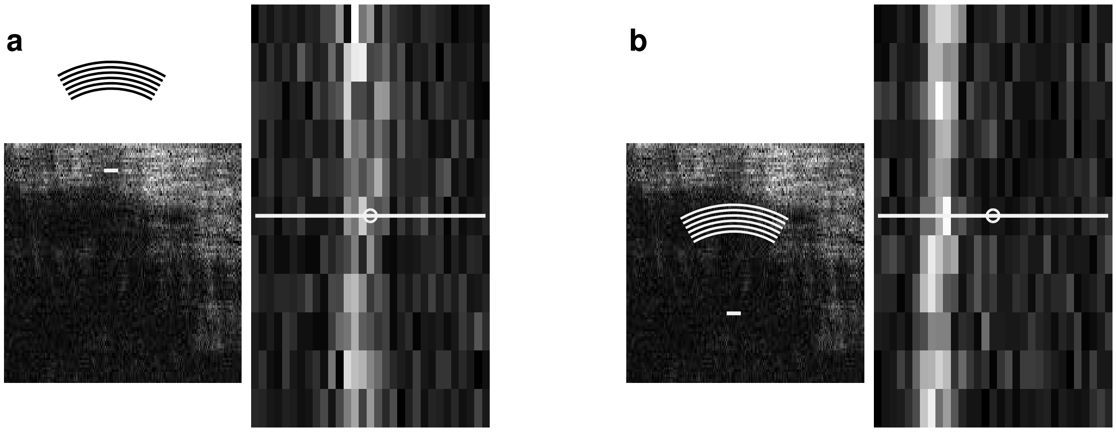

Figure 7.7: The left hand image of parts *a* and *b* shows an 82MHz SAW amplitude c-scan image of an area of aluminium, with a schematic of the generation arcs superimposed. The white horizontal line at the focus of the arcs indicated the location across which a PSF was acquired. This PSF is shown in the right hand image of parts *a* and *b*, along with the PSFs at scanning positions above and below the current marked position. These images are referred to as *w-scans.* Part *b* shows that although SAWs are propagating on the sample, the SAWs 'miss' the detection point used to make the c-scan image, leading to dark regions.

This PSF is shown in the right hand image of parts *a* and *b*, along with the PSFs at scanning positions above and below the current marked position. These images are referred to as *w-scans,* since they indicate how the PSF wanders from its intended central location. Part *b* shows that although SAWs are propagating on the sample, the SAWs 'miss' the detection point used to make the c-scan image, leading to dark regions.

Figure 7.8 shows the w-scans for ten vertical slices within the 5mm square region of polished aluminium shown on the left hand side of figure 7.7a and 7.7b. It is clear that for many of the scan positions, the SAW PSF is still intact, but displaced laterally. This implies that a tilt correction can be obtained by setting the intended focus of the excitation arcs to be an equal—but opposite—distance from the measured displacement. In this way the SAW waves will be subject to substantially the same aberrations that caused the PSF to wander, and should now propagate back to the desired detection point.

Two c-scans were performed over the area where the w-scans in figure 7.8 were obtained; the first without any tilt correction, the second using the data obtained from the ten w-scans. Figure 7.9 illustrates the improvement in the data obtained.

The left hand and centre images are the amplitude and phase c-scan images before tilt correction

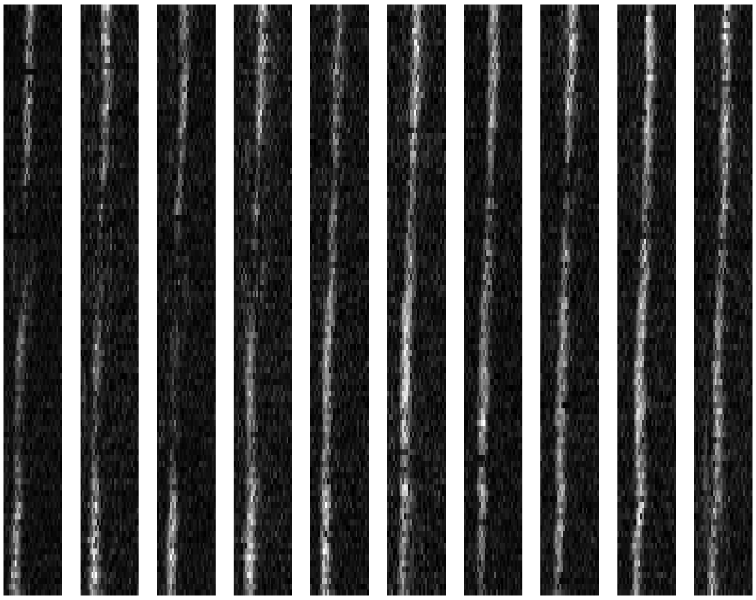

Figure 7.8: The figure shows ten w-scans acquired over the region of the c-scan shown in figure 7.7. The data from these w-scans is used to correct for the acoustic aberration by 'tilting' the generation arcs to steer the acoustic waves back towards their intended focus.

was applied (top) and after tilt correction was applied (bottom). The right hand images show the regions where the amplitude was greater than 25% of the maximum amplitude of the uncorrected data.

The area of valid data has now substantially increased, from around 24% for the uncorrected image to approximately 80% for the corrected image. It should be noted also that the correction data was only obtained for ten discrete slices (due to the time taken to acquire the data by detection probe scanning), and so the further from the measured slice, the less the data is a valid correction. The areas where the signal could not be recovered correspond to areas where the PSF breaks up further (referring to figure 7.8) due to higher order aberrations.

7.7 Higher order correction

The so-called 'tilt,' or 'first order' correction technique described in the previous section works well for materials and situations where aberration is relatively weak, and the PSF has a tendency to wander off from the detection point rather than break up completely. It has the advantages of the correction involving negligible computation, and correction may be applied by physically tilting the optics used to produce the generation profile, if a sophisticated adaptable optical image source

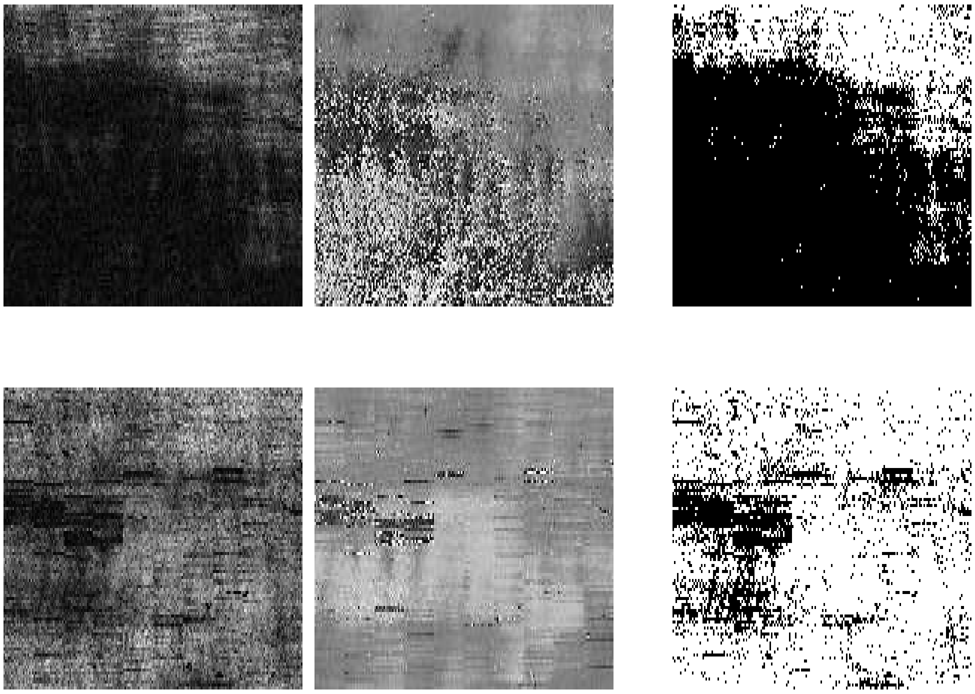

Figure 7.9: The left hand and centre images are the amplitude and phase c-scan images before tilt correction was applied (top) and after tilt correction was applied (bottom). The right hand images show the regions where the amplitude was greater than 25% of the maximum amplitude of the uncorrected data.

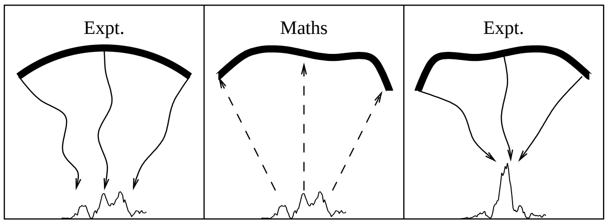

Figure 7.10: The figure describes the process of higher order wavefront correction. In the left hand figure, the aberrated wavefront is acquired. This is back-propagated to the excitation region (centre), and the difference between this and the geometric ideal is then used as the excitation pattern (right hand side).

such as a spatial light modulator is not available. Where the aberration is more severe and the wavefront breaks up altogether, a higher order correction technique is required, and the principle is shown in schematic form in figure 7.10.

The principle of this technique is to back-propagate the measured wavefront to the generation region, where the phase error is calculated compared to the geometric generation profile. This error is then used to generate a new generation profile to correct for the material aberrations.

This back-propagation is achieved by using an angular spectrum propagation technique, as described by Goodman [68]. The complex amplitude of the measured PSF is represented as an angular spectrum of plane waves using Fourier analysis. These plane waves are propagated to the generation region by adding an appropriate propagation phase to each of them. The complex amplitude in this region can then be acquired by taking the inverse Fourier transform.

Figures 7.11 and 7.12 illustrate the technique in practice. Figure 7.11 shows the aberrated aerial PSF acquired with the O-SAM instrument, when surface waves were excited using an image of concentric arcs; the amplitude across the centre of the PSF is also shown. Figure 7.12 shows the corrected aerial PSF, again acquired with the O-SAM instrument, in exactly the same area as the previous figure. The only difference is the SAW excitation pattern, which for the second image is based on the errors between the back-propagated wavefront and the ideal, geometric excitation pattern.

To further demonstrate the effectiveness of the two techniques, 2mm of a particularly aberrating piece of aluminium with average grain size $600\times200\mu$m was scanned—the long axis of the grains being in the direction of SAW propagation—and at each point along the scan-line the complex amplitude of the PSF was acquired. The data is combined in the left hand image of figure 7.13 to generate a w-scan, which shows the amplitude of the wavefront along the line of the nominal focus as the sample is scanned. The w-scan shows that the PSF is of a low quality for much of the

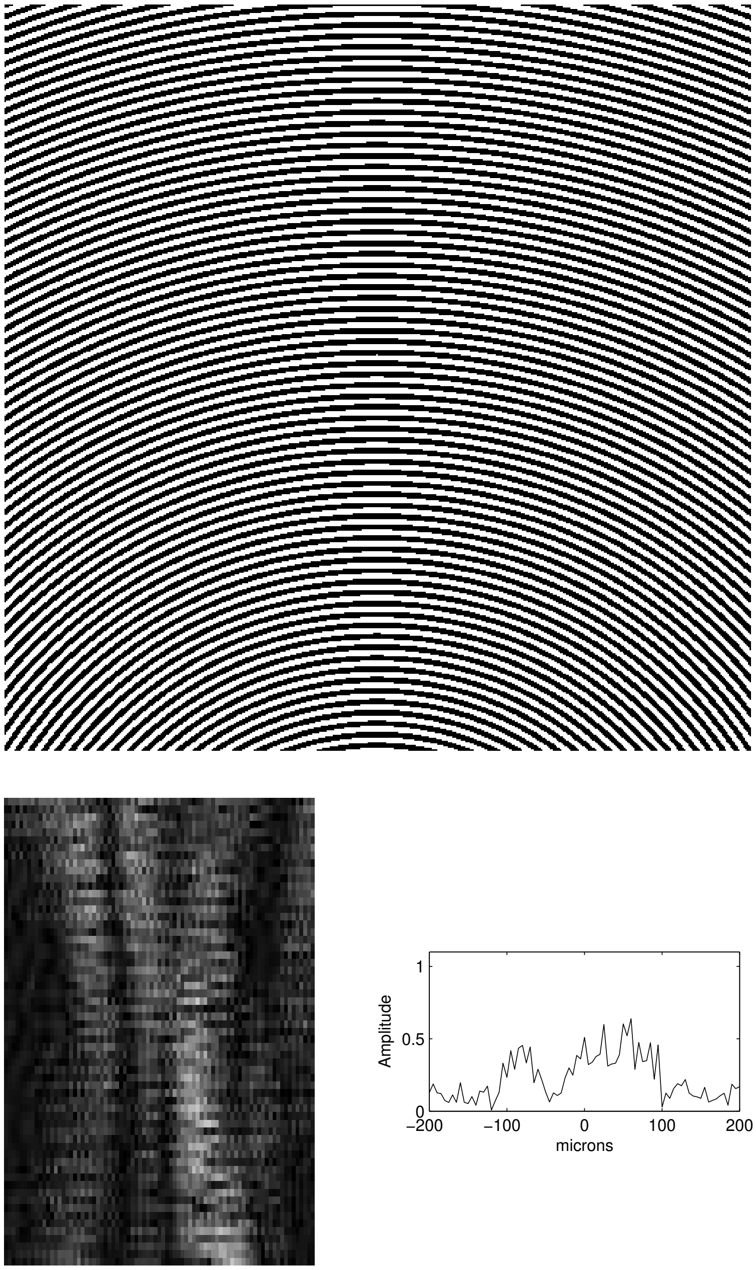

Figure 7.11: The figure shows an aerial PSF which shows degradation due to aberration. The SLM image used to generate the SAWs is shown in the upper part of the figure, and is a concentric set of geometric arcs with a focal length of 2mm. The amplitude PSF is shown in the bottom right.

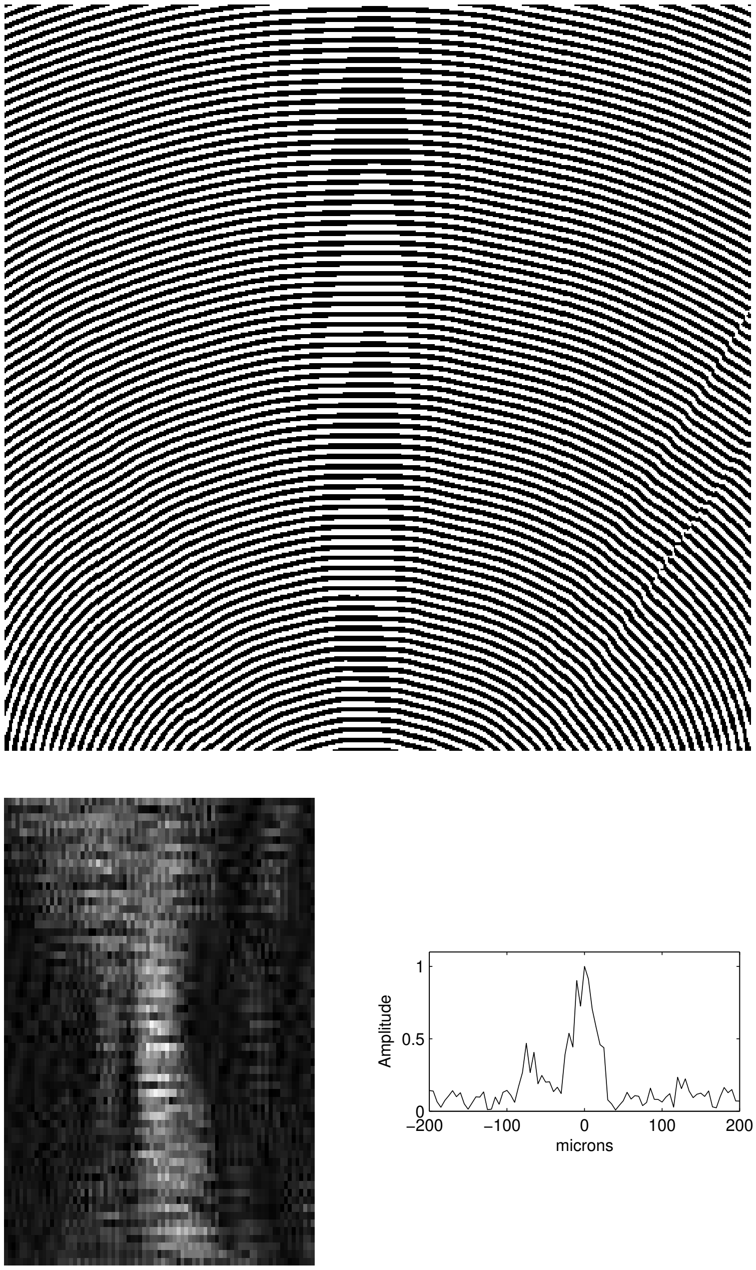

Figure 7.12: The figure shows an aerial PSF which has been improved by applying higher order correction to the generation profile, shown in the upper part of the figure. The PSF was taken at exactly the same location as the PSF in the previous figure, the only difference being the image applied to the sample via the SLM.

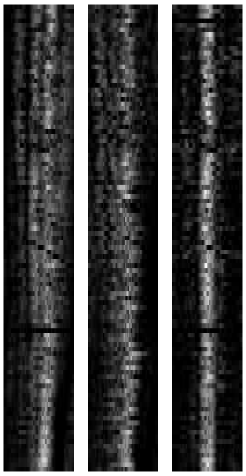

Figure 7.13: Images of the amplitude PSF as an aluminium sample is scanned along a line (w-scans). The left hand image shows the uncorrected w-scan, the middle image shows the result of tilt correction, and the right hand image shows the result of higher order correction. Image size is 0.3mm$\times$2mm.

scanned line, since it is broken and of a reduced amplitude. The effect of applying tilt correction is shown in the middle image of figure 7.13—here the 'peak' of the PSF is defined as the first moment of the amplitude. Although the PSF is, in general, steered back towards the ideal position (a vertical line along the centre of the w-scan), the PSF is still broken and attenuated. The right hand image in figure 7.13 shows the result of higher order correction. Not only has the PSF been moved back to the correct place, its quality is greatly improved and the amplitude increased. The majority of the high frequency noise on all three images is due to optical speckle resulting from a poor surface finish.

A more rigorous approach to defining the "quality" of the PSF is to take the first and second order moments of the intensity of the surface acoustic waves for each position along the scan-line, and this is done in figure 7.14.

The upper plot shows the normalised first order moments of the PSF at each point along the scan-line (going left to right in figure 7.13) for the uncorrected, tilt-corrected and higher-order-corrected cases. The tilt-corrections produce better results than the uncorrected scan. The higher-order corrections move the PSF even closer to the ideal position. The lower plot of figure 7.14 shows the second order moments, which indicate how 'wide' the PSF is. Here, the data acquired after higher order correction has been performed indicates a much better quality PSF, although we note that there is some significant variation in the first and second order moments. In general, for all three cases, the errors are more severe at the 'start' of the scan line—corresponding

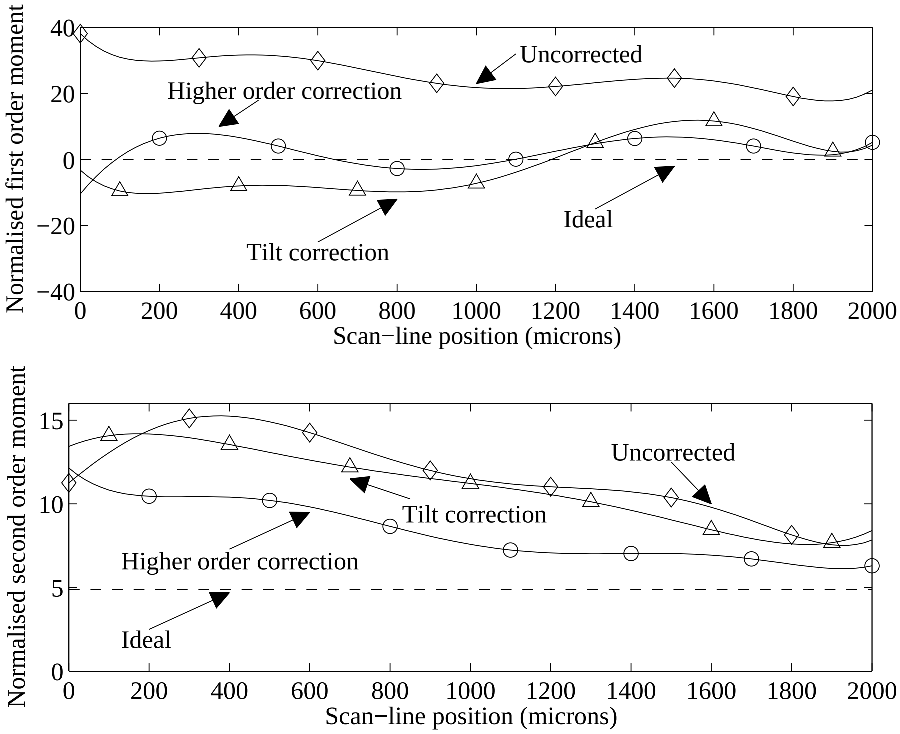

Figure 7.14: The upper plot shows normalised first order moments of the acoustic intensity for uncorrected (⋄), tilt-corrected (△) and higher-order-corrected (∘) line-scans. The ideal (- -) is also shown. The lower plot shows the second order moments of the same data. Both plots are best-fit lines through experimental data.

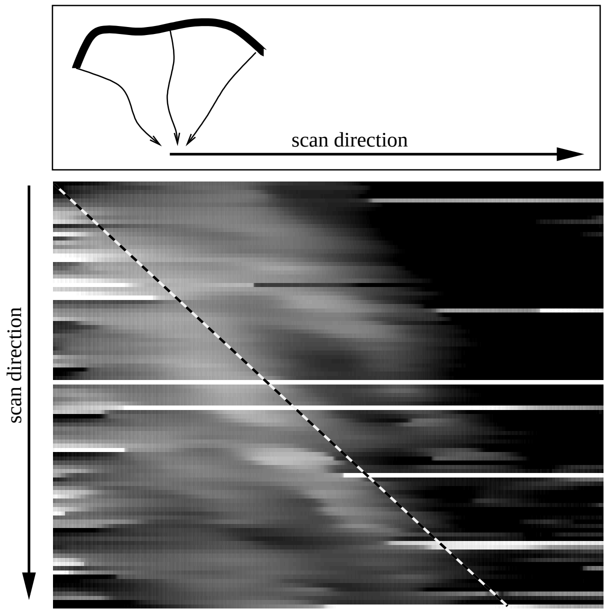

Figure 7.15: The upper schematic diagram indicates the wavefront propagation and scanning directions. The intensity of the lower image represents the phase errors (horizontal axis) for each point along the scan-line (vertical axis). The added diagonal dashed line indicates the position of a reference point. Size of image is 2.6mm$\times$2mm.

to the left of figure 7.14 and the top of figure 7.13. With reference to figure 7.13 we see that the aberration is generally more severe at the start of the scan line.

It is interesting to investigate the phase errors between the back-propagated aberrated wavefront and the ideal geometric case, as this allows us to build up a map of the aberrations experienced by the surface acoustic waves. As the sample is scanned, there should be correlation between the errors (and therefore the correction applied) between adjacent points, and indeed there is. This allows us to re-use much of the collected data for correction at nearby scan-points. The principle is described in figure 7.15. In this case, the scan-line is in a direction perpendicular to the direction of acoustic propagation. As you would expect, the various features of the phase errors move by an amount equal to the distance travelled along the scan-line, as confirmed by the diagonal dashed reference line on the image.

7.8 Real-time adaptive acoustic system

It is clear that to correct for the effects of aberrations it is necessary to detect the acoustic wavefront at the intended detection region. For the results presented in the previous two sub-sections,

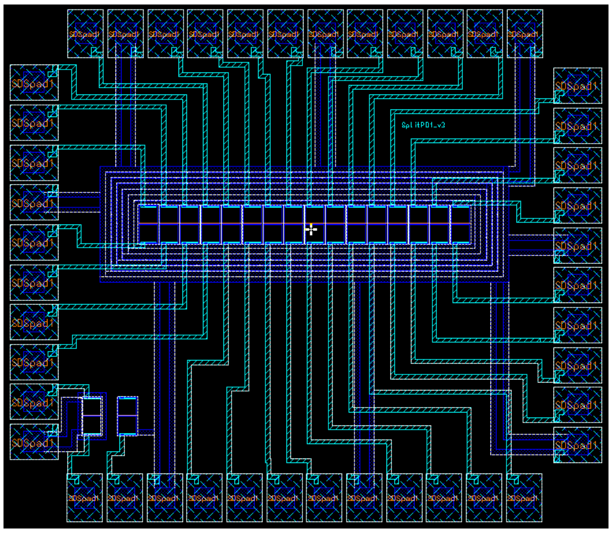

Figure 7.16: Design for a 16-channel optical acoustic wave front sensor (AWFS).

concerning first order and higher order correction techniques, the acoustic wavefront was detected by mechanically scanning the single-point optical detection probe across the point spread function focus. This adds a significant overhead in terms of scanning times, as well as increasing the complexity of the scanning operation.

To make a real-time adaptive system viable, an acoustic wavefront sensor is required. This would measure the complex amplitude of the surface waves along a line perpendicular to the SAW propagation direction, rather than at a single point as is the case now. To this end, construction of a multiple-point acoustic wave front sensor (AWFS) system has recently begun.

At the heart of the AWFS system lies a 16 channel linear modified knife edge detector—effectively a 16 channel version of the detection system already described in chapter 3. This is an array of 16×2 photodiodes, designed and manufactured using a standard CMOS process. Figure 7.16 shows the layout of the chip.

The design is obviously very simple, and contains—at this stage—no integrated electronics. The device as it stands is simply a custom photodiode array. However, because the design was made using a standard CMOS process, it is easy to further integrate the associated electronics with the optical detector, with substantial savings in overall system complexity and cost.

Overall, the adaptive O-SAM, based on a spatial light modulator to correct the SAW generation profile, and an acoustic wavefront sensor to detect the aberrated wavefront, will resemble the schematic shown in figure 7.17.

A cylindrical lens will focus a beam from the detection laser (not shown) onto the sample in the

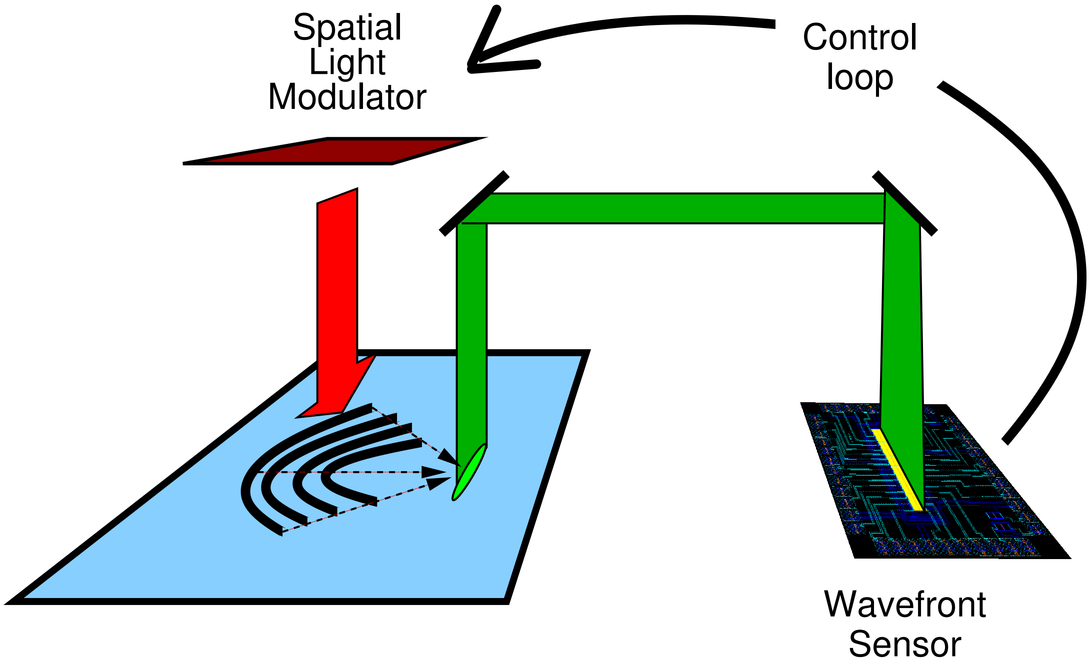

Figure 7.17: System design for rapid wavefront sensing and higher order correction.

shape of a thin line, rather than a diffraction limited spot as in the present system. SAWs passing beneath the line will perturb portions of it. The line is imaged onto the AWFS device along the centre of the array of split photodiodes, thus each 'pair' of photodiodes will measure one sixteenth of the wavefront.

It should be pointed out that in the higher order correction experiment shown in the previous subsection, only 16 equally-spaced points along the measured point spread function were used in the back-propagation algorithm, thus giving a good representation of what should be achievable with the acoustic wavefront sensor, although of course the amount of optical power available to each photodiode will be reduced accordingly. This is partly compensated for by using a detection laser with a higher output power.

It is envisaged that the system will eventually be able to operate in a 'real time' closed loop higher order correction configuration. At the beginning of the scan, higher order correction is applied to the first c-scan point, using data acquired by the AWFS system, performing several iterations if necessary. As demonstrated in the previous section, if the amount by which the sample moves between each scan point is relatively small compared to the propagation distance and the width of the excitation region, then the surface waves are subject to substantially the same aberrations, and the corrections applied to one point will relate to the corrections applied to another point, usually in some kind of semi-predictable way. This information can then be used in a control loop, to adapt 'on the fly' as a c-scan is performed.

A question that repeatedly gets asked is, "If you are correcting for the aberrations, correcting for the 'errors,' then what are you actually measuring? Is there not a danger that you'll 'correct out' the features you are trying to measure?" The reason the answer is 'no' is relatively straightforward, when you consider that in performing the correction for each point in the c-scan, you have *knowledge of the correction applied.* Rather than measuring the change in characteristics of a detected surface wave due to some feature of the material, you are measuring the necessary changes in the generation profile required to keep the SAW characteristics constant (or as constant as possible). Instead of, for instance, measuring the complex amplitude of the SAWs in an open loop system with no correction and noting a change in phase and a reduction in amplitude—which could imply any manner of things—you would note the amount of phase correction required to keep the measured phase constant, and also note that the general direction of propagation needed to be altered by a certain degree, that the profile geometry needed to be altered in a certain way, and the amplitude compared to other points in the c-scan, all of which may tell you something much more interesting.

Another important point about correcting for the material structure is concerned with the detection of small defects. The possibility of 'missing' small features is significantly increased in the case of aberrating materials, as demonstrated in section 7.5. Adaption ensures that the acoustic energy is focused to a known, small region, and so any defects there will interact much more strongly with the acoustic waves.

Chapter 8

Issues arising and further work

8.1 Introduction

The c-scan images and other results from the previous three chapters illustrate some of the powerful contrast mechanisms available from the O-SAM instrument. However, there is still much research to be done in the areas such as optical generation and detection of surface acoustic waves, exploring different contrast mechanisms, interpretation of data, and general instrumentation and control. Significant inroads into many of these issues have been made over the last two or three years, in terms of extending the capabilities of the O-SAM instrument itself, by using the instrument to explore the relationship between surface waves and material structure, and verifying models to that effect.

8.2 Velocity measurement—the O-SAM as a material characterisation tool

The majority of the work presented in this thesis has concerned the use of the O-SAM instrument as an imaging tool—whether it be the amplitude, phase or velocity of Rayleigh waves, or interference between, and velocities of, various Lamb wave modes. But the flexibility of the design and implementation of the instrument allows it to be used for many other purposes.

The measurement of surface wave velocity has already been referred to in chapter 1 as a useful method of material characterisation, and is one that is well suited to laser ultrasound techniques in general, and arguably the O-SAM instrument in particular. Its strengths lie in several factors. Firstly, generation and propagation of discrete, narrowband frequencies of waves over a large overall bandwidth—from 82MHz up to and potentially beyond 1GHz—means less difficulties in interpretation of the data compared to broadband systems, whilst maintaining the ability to tailor, to a certain extent, the wavelength to the application. Secondly, the ability to easily control the generation profile, to either adjust the wavelength for optimum excitation, or to correct for the aberrating material characteristics, means that new methods can be used to successfully perform velocity measurements for a wide range of materials.

8.2.1 Rapid velocity measurement by changing the grating spacing

It was stated in section 3.4, concerning the use of a spatial light modulator to control the ultrasound source, that "...this flexibility allows us to adapt the generation profile to the material properties." As well as altering the profile of the excitation source to correct for aberrations caused by the material grain structure as described in the previous chapter, this allows us to match the fringe spacing to the wavelength of the SAWs at a certain frequency (82MHz, 164MHz etc), and thus to the SAW velocity of the material. This matching of the fringe spacing to the material velocity gives us a very simple, powerful and accurate method of velocity measurement. An investigation into using this technique was initiated by the author relatively recently, and has been continued by Hong Yi, a research student in his first year with the group.

A set of straight fringes is imaged onto the material whose Rayleigh phase velocity is to be determined, and the detection point is located a short distance away from the edge of the excitation region—typically 2mm from the centre of the region. The spacing of the fringes is changed progressively, and the amplitude of the surface waves are measured for each different fringe spacing. This can be done at several frequencies (over a range of fringe spacings) and the amplitude can be measured by a variety of means, including the analogue vector contrast measurement system for rapid acquisition in cases of good signal strength, or by calculating the amplitude of a given frequency component (say, 82MHz) from the FFT of a waveform acquired using a digital storage oscilloscope, using coherent digital averaging if necessary in cases of extremely poor signal strength. The amplitude will reach a peak when the fringe spacing of the excitation source matches the Rayleigh wavelength at the frequency being measured, from which the phase velocity is obtained. Figure 8.1 illustrates the kind of response obtained. This plot of SAW amplitude with respect to fringe spacing was taken along an angle of 22° on (111) oriented silicon at 82MHz.

The method is particularly useful for the measurement of angular dispersion of single crystals, from which slowness measurements can be derived. In this configuration, the peak of the amplitude response denotes the surface wave phase velocity for a given orientation of the crystal; the crystal is then rotated by a small amount, and the peak of the amplitude response is again found. The velocity for each orientation of the crystal is obtained in this way.

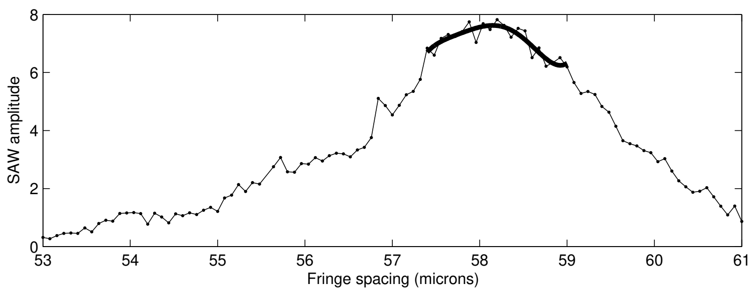

Figure 8.1: Amplitude response with respect to fringe spacing for a certain orientation of a single crystal, from which the phase velocity may be calculated. A polynomial curve is fit around the peak.

In fact the crystal need not be rotated at all, since the SLM can produce images of straight lines in any orientation. The detection point would in this case have to be moved for each angular direction of propagation, but its exact location is not critical, and any errors and inconsistencies in the mechanics of the moving components do not affect the measurement accuracy. Figure 8.2 illustrates the operation of the velocity measurement technique in this configuration.

A problem arises with this technique, in that the knife edge detection system employed as the SAW detector is sensitive to only one orthogonal component of the SAW direction, and hence the detected amplitude decreases significantly as the angle of propagation is adjusted. This problem could be eliminated by the design and construction of a detector sensitive to angular displacement in two dimensions, which would employ a four-element photodiode to do this.

Overall, this technique, whether achieved by rotating the sample or by rotating the excitation fringe image, offers a simple, rapid and convenient method of velocity determination without recourse to mechanical movement. Combined with the phase gradient method of determining the SAW velocity, as described in section 5.4.1, could potentially offer extremely good measurement precision.

8.2.2 Measurement of phase gradient on aberrating materials

Figure 7.4 in the previous section illustrated the propagation of plane waves over the surface of steel, which is composed of many randomly-orientated grains, and consequently aberrates the acoustic wave front. Figure 8.3 illustrates the effect that the variation in velocity due to aberrations has on the measurement of phase.

The images are of the residual phase along the direction of propagation for a given wave velocity; that is, the expected phase (given a certain velocity) minus the measured phase for each

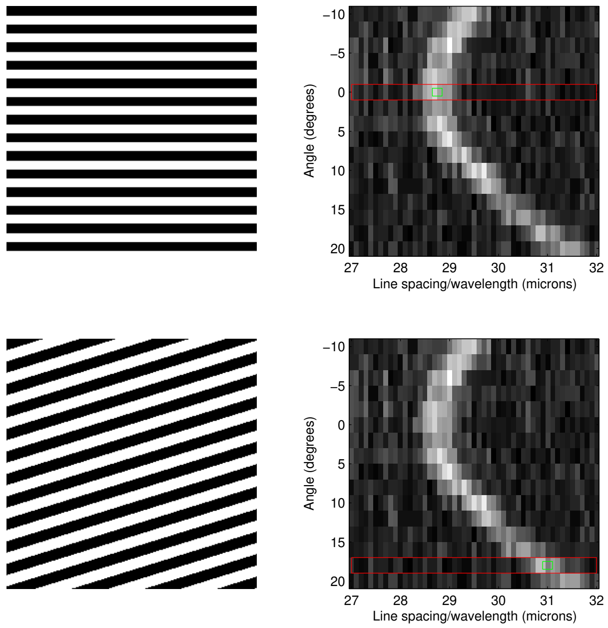

Figure 8.2: Rotating fringe spacing phase velocity measurement technique. Rather than rotating the sample, the image of the fringes is rotated instead.

position along the direction of propagation. The left hand image shows propagation on an isotropic medium, and the residual phase is reasonably constant—variations are assumed to be the result of mechanical errors in the moving stages. The right hand image shows propagation on hardened steel, and not only is there considerable variation in the residual phase, but there are several locations where the phase along a vertical line jumps by 2π, or a multiple thereof. Both of these factors have implications for the successful measurement of phase velocity using the phase gradient technique.

Work has recently been undertaken by Matt Clark to investigate whether it is possible to produce a region of high spatial coherence of the propagating acoustic wave suitable for performing good velocity measurements, by applying a 'custom' excitation pattern based on measured Green's functions. The Green's functions are measured by 'scanning' a point-line source using the spatial light modulator, and measuring the complex amplitude across a line perpendicular to the intended SAW propagation direction; see figure 8.4.

The Green's functions, $G(x, y, S)$, are then used to optimise the system for velocity measurement, with the aim to produce a complex amplitude distribution $U(x, y)$ in the intended measurement region that has a high amplitude, low noise, and a flat wavefront. This is achieved by adding

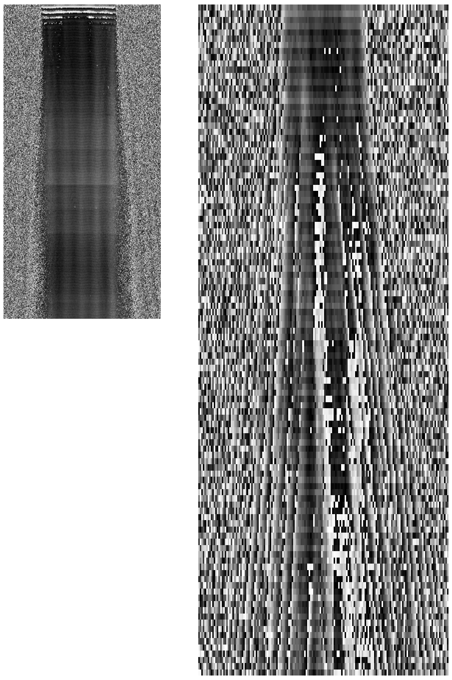

Figure 8.3: The left hand image shows the residual phase of 82MHz SAW plane waves propagating along the surface of aluminium-coated glass. The right hand image shows plane waves propagating along the surface of hardened steel.

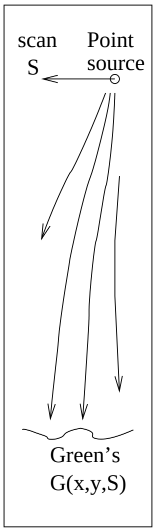

Figure 8.4: Measurement of the Green's functions by scanning a point-like source and measuring the material response.

a carefully chosen phase, $\phi_d(S)$, since

$$U(x,y) = \int_S G(x,y,S)e^{-i\phi_d(S)}dS \tag{8.1}$$

.

The method yields good results, with projected regions of SAWs with high spatial coherence and a large 0-order component for velocity measurement. Even though the phase gradient may vary within this region, any variation will be due to the effects of the individual grains within that region, rather than the accumulated phase errors due to the aberration associated with the wavefront before it gets to the measurement region, which causes the apparent measured velocity to vary so much as the measurement position is changed. Improvements in the measurements can be achieved by increasing the length of the spatially coherent measurement region, thus averaging over a greater number of grains and so reducing the intrinsic velocity variation.

8.3 Nonlinear imaging

The subject of nonlinear and harmonic detection of ultrasound was discussed in chapter 1, and is a very active area of research in the world wide ultrasonics community in general, since it can potentially yield much greater information about material properties, structures, and faults. The use of non-contact techniques in general offer an advantage over fluid coupled techniques from the start, since any nonlinear activity in the material under investigation is likely to be swamped by the intrinsic nonlinearity of the liquid couplant.

The O-SAM instrument has many inherent strengths in the quest to measure large nonlinear activity characteristic of coating flaws and disbonds, and it is hoped that intrinsic material nonlinearity will be measurable in the long term. Focusing the acoustic waves to form a small region of high intensity is an ideal way to excite waves of sufficient amplitude to ensure a measurable harmonic amplitude, and to ensure that the generated harmonic arises from a well defined localised region. The ability to suppress any harmonic frequencies from being generated is very important for this type of experiment, and this is possible with the O-SAM instrument. Section 3.2.4 described how the frequency of the propagated waveforms can be controlled by adjusting the line spacing of the excitation source; section 3.2.5 described how the different SAW frequencies may be spatially separated by using diffractive acoustic elements. An adaptive source—the SLM—allows us to combine the two techniques, such that any harmonic activity detected can be attributed to nonlinear behaviour in the material, rather than harmonics generated by the source itself.

There are two further issues that will make the demonstration of harmonic imaging using the O-SAM feasible. The first is the availability of samples that exhibit strong nonlinear behaviour—

the most obvious choice for these would be samples with kissing bond cracks. The second is the ability to ensure the generation of high amplitude SAWs without harmonic contamination, and for this reason it may be necessary to increase the power of the excitation source—a 'bigger laser,' or an optical amplifier, may be required.

8.4 Interpretation of the acoustic information

Recall that in chapter 5, acoustic images of a Vickers hardness test indentation were presented. The 82MHz acoustic images, and the optical image, are illustrated here in figure 8.5 as a reminder.

The feature is a complicated one, and although analysis of parts of the images can give good qualitative and quantitative information (see section 5.3.2), there is still much to be done in the full analysis of the data. It may be that the 'complete story' of the interaction of the feature with the surface waves cannot be completely understood, without recourse to other testing methods.

There are two features in the acoustic images that are both interesting and difficult to explain. The first is the area 'behind' the Vickers indentation—that is, in the lower half of the images, beneath the black diamond. There is clearly some acoustic activity in this region, and this is presumably because the surface waves propagate on the indented surfaces of the test point. SAWs cannot be observed in this region because, as can be seen from the optical image, the reflected light from the detection probe is deflected away from the photodetectors. The waves have gained extra phase, due to the increased propagation distance. Of most interest is the set of bright fringes in the amplitude image on the left hand side of the lowermost crack. In this region there is also a change in phase. As we have seen in chapter 6, this could indicate two or more modes travelling at different velocities, such as Lamb a_0 and s_0 modes. It is possible that this region is delaminated. It is, however, puzzling that the fringes appear to be limited to the region along side the crack.

The other interesting feature occurs in the amplitude image, and is the bright fringe at the end of the right hand crack (and also, to a lesser extent on the left hand crack). The amplitude of this fringe is significantly greater than the incoming Rayleigh wave, and greater even than the interference fringes between forward- and backward-propagating waves. The bright fringe is spatially limited to a region much less than one acoustic wavelength and, due to the limited optical resolution, it is very difficult to state on which side of the crack the bright fringe lies.

Bright fringes such as these have also been observed on a quite different sample. Figure 8.6 shows the optical (a) 82MHz SAW amplitude (b) and SAW phase (c) of a $400 \times 450\mu$m region of a piece of quartz covered in a thin layer of aluminium. The square, evident in all the images, represents the boundary between a region of poorly-bonded aluminium (inside the square) and aluminium bonded well (outside the square).

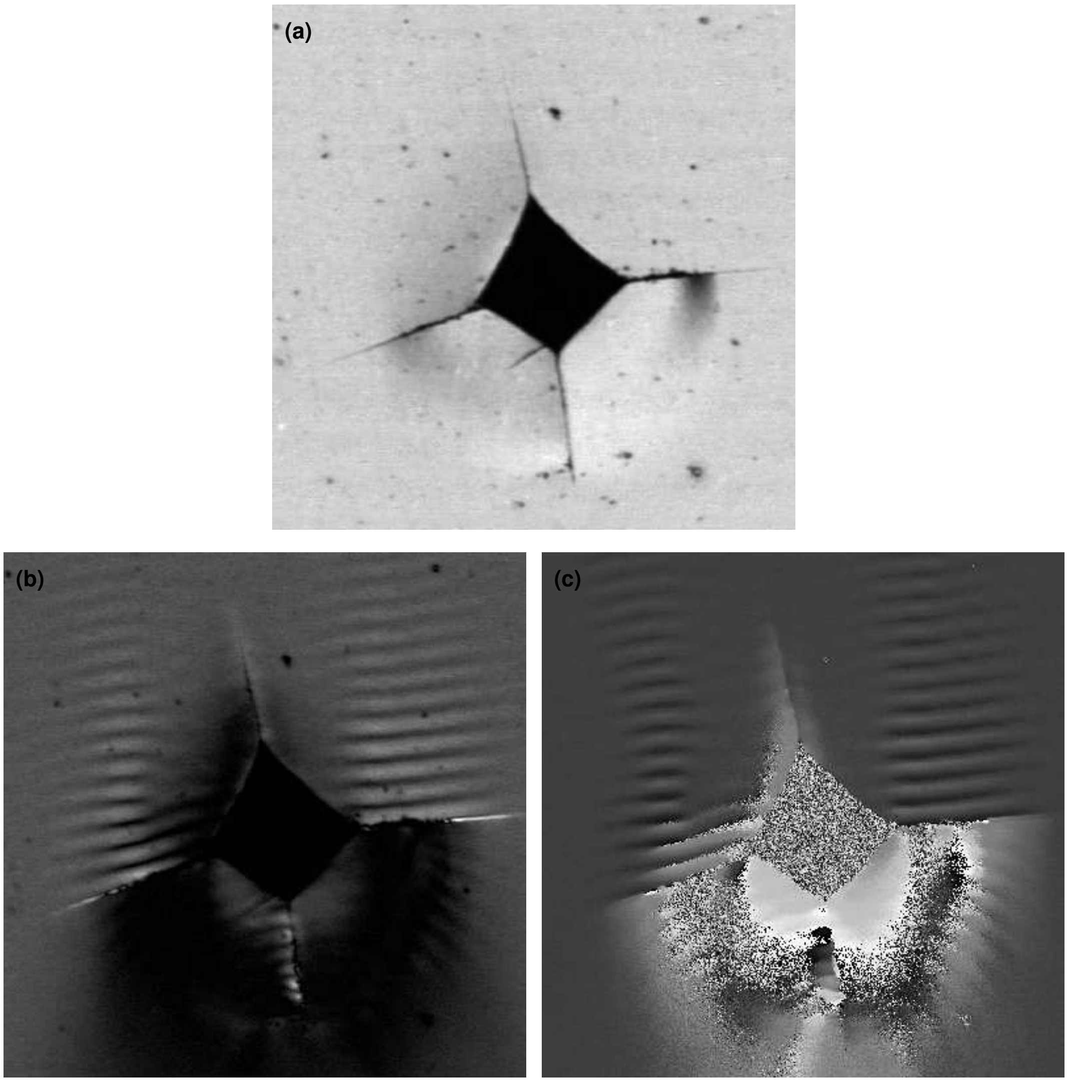

Figure 8.5: Optical (a), 82MHz SAW amplitude (b) and phase (c) images of the Vickers hardness indentation on silicon nitride, previously presented in figure 5.4 in chapter 5. Some of the features of the acoustic images are difficult to interpret.

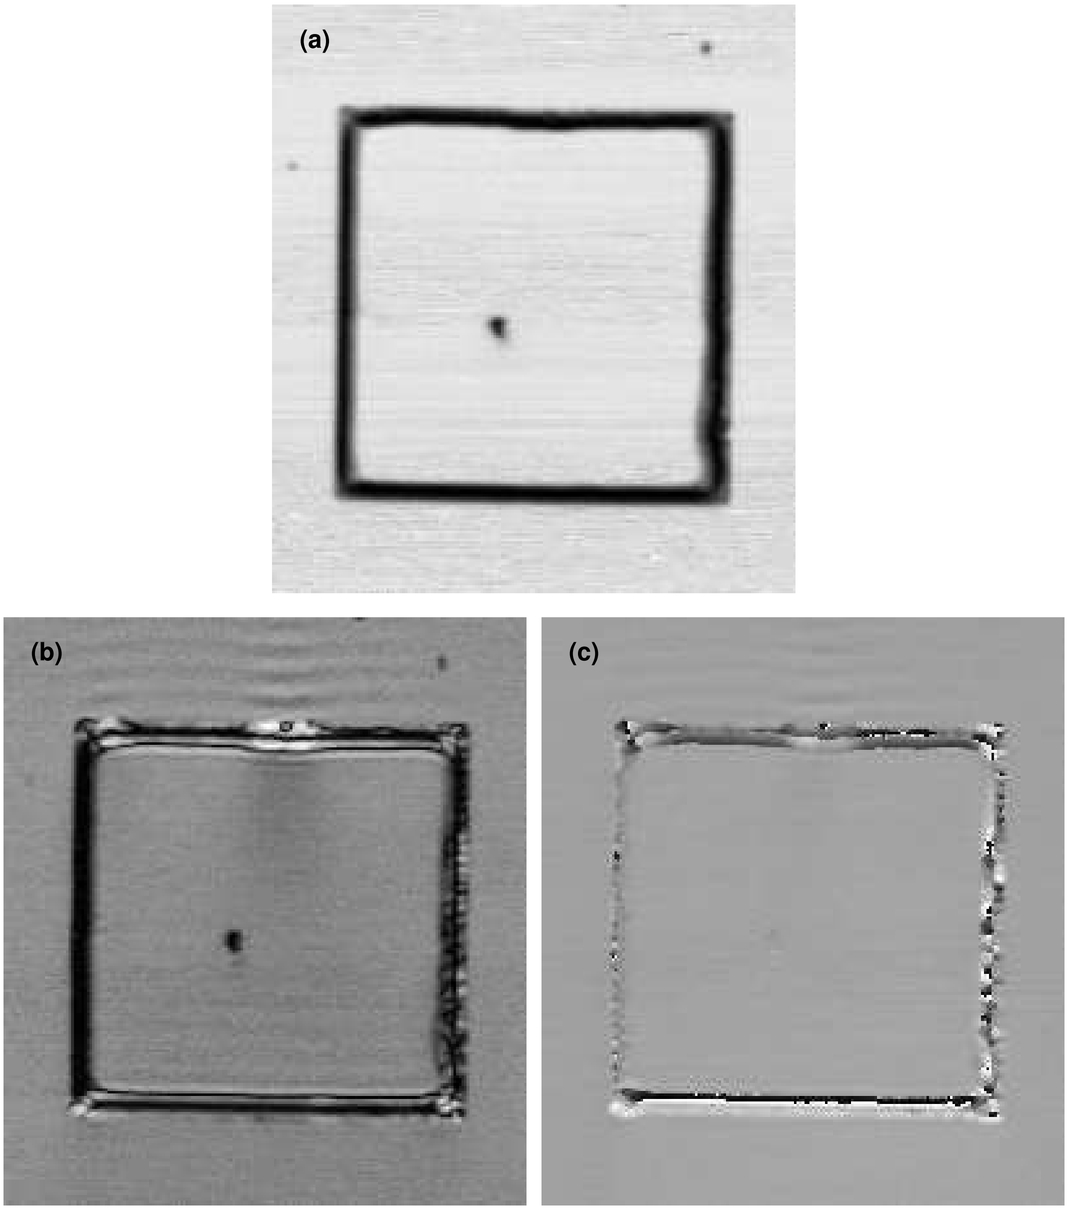

Figure 8.6: Optical (a), 82MHz SAW amplitude (b) and phase (c) images of an aluminium-coated sample. The area inside the square is designed to be poorly bonded to the substrate.

Small reflection fringes can be seen in the amplitude and phase images from the top and bottom edges of the square feature, although there is significant transmission between through the 'edges' of the square. Again, very bright fringes are noted at the top and bottom edges that are difficult to explain.

It is clear that there is some form of interesting interaction occurring in these regions, and this requires further investigation, but is beyond the scope of the work presented in this thesis. Possible causes of the bright fringes could be some sort of resonance at the interface, and it is possible that the bright fringes in the Vickers images are not related to the bright fringes in the coated quartz, since the geometry of the samples is very different. It is clearly localised to the interface—either a crack edge or a coating boundary—and needs to be examined further.

Chapter 9

Conclusions

The work presented in this thesis has concerned the realisation of an all-optical scanning acoustic microscope, the O-SAM. Its place within the ultrasound NDT community as a whole, and non-contact techniques involving lasers specifically, has been stated and demonstrated. Ultrasound imaging is a relatively common technique in NDT; the use of lasers to generate and detect ultrasound has been common for many years. Although laser ultrasound *imaging* systems exist, in the main they are either prohibitively slow, complicated, of an inadequate resolution or simply too expensive. The 'missing link' appears to a high resolution, genuinely useful, *rapid* means of acquiring images of the interaction of acoustic waves with materials.

The instrument has been described in great detail in chapter 3, and its ability to control both the frequency content and spatial distribution of the surface wave excitation source has been demonstrated to good effect. At the heart of the instrument's ability to do this is either a computer generated hologram (CGH) or spatial light modulator (SLM). This allows the surface waves to be focused—localising the acoustic interaction and increasing the SNR—and the energy of the excitation laser to be distributed over a wide area such that damage is avoided. A wide range of control of the acoustic wavefront has been demonstrated, and this enhances the instrument's capabilities and potential contrast mechanisms. The rapid analogue vector contrast data acquisition system has been described, and this important aspect of the instrument enables high quality, high resolution images to be obtained at rates potentially equal to the generation laser pulse repetition rate.

The capabilities of the O-SAM to acquire a range of information about the properties of a material have been demonstrated, in particular the interaction of the material with Rayleigh waves (chapter 5) and Lamb modes on thin plates (chapter 6). Surface and near-surface defect detection, some degree of crack characterisation, the ability to detect changes in surface stress or porosity, and differences in coating thickness by a change in velocity have all been demonstrated in chapter 5. Plate thickness measurement by Lamb mode interaction has been demonstrated in chapter 6, and some of the effects of focusing acoustic waves on thin plates has been discussed.

The effects of material anisotropy due to random grain structures and distributions has been observed, and techniques have been described in chapter 7 to correct for these effects. Amongst the techniques described and suggested, the higher order correction of the spatial distribution of the SAW excitation profile yields the greatest improvement in the consistency of the acquired data. This technique, which involves detecting the aberrations to the acoustic wavefront caused by material anisotropy with an acoustic wavefront sensor, and using the information collected to generate a new excitation profile tailored to a specific point on a specific material has been demonstrated in principle by measuring the acoustic wavefront by a scanned single point detector. An active area of research is to develop an acoustic wavefront detector, to allow higher order correction to be performed 'on the fly' during scanning, to increase the quality of the acoustic images acquired on aberrating materials.

Some of the other uses and potential uses of the O-SAM instrument—apart from imaging the interaction between surface waves and materials—have been presented or discussed in this chapter. These include a rapid method of measuring the velocity of anisotropic crystals at different orientations, and a technique to improve the reliability of velocity measurement on aberrating materials.

Of the future work that could be done with the instrument, nonlinear imaging is perhaps the most exciting. The adaptive nature of the SAW excitation source, the ability to control the frequency content of the excited surface waves, and the ability to spatially control their distribution at different frequencies all contribute to a strong possibility of not only detecting nonlinear interaction between surface waves and materials, but also to image this interaction.

Finally, it is worth stressing the usefulness and importance of imaging, and specifically the ability to acquire high resolution, detailed images in a genuinely-useful time period. This capability has led to research into the propagation of acoustic waves in anisotropic media, and the O-SAM's ability to probe the surface waves as they propagate is an extremely useful tool for model development and confirmation. Although some of the acoustic interactions detected with the O-SAM remain as yet not completely explained, the ability to image the nature of these interactions—combined with more detailed a-scans and b-scans—significantly increases the chances of finding out.

Bibliography

[1] M. G. Somekh, M. Liu, H. P. Ho, and C. W. See, "An accurate non-contacting laser based system for surface wave velocity measurement," *Measurement Science and Technology*, vol. 6, pp. 1329–1337, 1995.

[2] M. Clark and R. Smith, "A direct-search method for the computer design of holograms," *Optics Communications*, vol. 124, no. 1-2, pp. 150–164, 1996.

[3] M. Liu, H. P. Ho, M. G. Somekh, and J. M. R. Weaver, "Noncontacting optical-generation of focused surface acoustic waves using a customized zoneplate," *Electronics Letters*, vol. 31, no. 4, pp. 264–265, 1995.

[4] M. Clark, F. Linnane, S. D. Sharples, and M. G. Somekh, "Frequency control in laser ultrasound with computer generated holography," *Applied Physics Letters*, vol. 72, no. 16, pp. 1963–1965, 1998.

[5] F. Linnane, *Investigation of an all-optical laser ultrasound SAW system with respect to velocity measurement and imaging.* PhD thesis, The University of Nottingham, 1998.

[6] E. P. Papadakis, *Ultrasonic Instruments and Devices I*, vol. 23 of *Physical Acoustics*, ch. 3, pp. 193–274. San Diego, CA, USA: Academic Press, 1999. (Editors: R N Thurston and Allan D Pierce).

[7] H. E. Davis, G. E. Troxell, and C. T. Wiskocil, *The Testing and Inspection of Engineering Materials.* McGraw-Hill, 1964.

[8] D. E. Bray and R. K. Stanley, *Nondestructive Evaluation, A Tool in Design, Manufacturing and Service.* CRC Press, 1997.

[9] J. Krautkramer and H. Krautkramer, *Ultrasonic Testing of Materials.* Springer-Verlag, 1990.

[10] I. A. Viktorov, *Rayleigh and Lamb waves: physical theory and applications.* New York: Plenum Press, 1967.

[11] A. S. Murfin, R. A. J. Soden, D. Hatrick, and R. J. Dewhurst, "Laser-ultrasound detection systems: a comparative study with Rayleigh waves," *Measurement Science and Technology*, vol. 11, pp. 1208–1219, 2000.

[12] P. R. Murray and R. J. Dewhurst, "Application of a laser/EMAT system for using shear and LS mode converted waves," *Ultrasonics*, vol. 40, pp. 771–776, 2002.

[13] P. R. Murray and R. J. Dewhurst, "A laser/EMAT system for thickness monitoring applications using shear and LS mode-converted waves," *Measurement Science and Technology*, vol. 12, pp. 1651–1659, 2001.

[14] R. S. Gilmore, *Ultrasonic Instruments and Devices II*, vol. 24 of *Physical Acoustics*, ch. 5, pp. 275–346. San Diego, CA, USA: Academic Press, 1999. (Editors: R N Thurston and Allan D Pierce).

[15] A. A. Kolomenskii and A. A. Maznev, "Phonon-focusing effect with laser-generated ultrasonic surface waves," *Physical Review B*, vol. 48, no. 19, pp. 14502–14508, 1993.

[16] J. P. Wolfe, *Imaging Phonons.* Cambridge, UK: Cambridge University Press, 1998.

[17] C. F. Quate, "Acoustic microscopy: recollections," *IEEE Transactions on Sonics and Ultrasonics*, vol. SU-32, p. 132, 1985.

[18] C. B. Scruby and L. E. Drain, *Laser Ultrasonics, Techniques and Applications.* Bristol, UK: Adam Hilger, 1990.

[19] C. B. Scruby, R. J. Dewhurst, D. A. Hutchins, and S. B. Palmer, *Research Techniques in Nondestructive Testing*, vol. 4, ch. 5, pp. 281–327. San Diego, CA, USA: Academic Press, 1982. (Editor: R. S. Sharpe).

[20] C. M. Scala and P. A. Doyle, "Time- and frequency-domain characteristics of laser-generated ultrasonic surface waves," *Journal of the Acoustical Society of America*, vol. 85, no. 4, pp. 1569–1576, 1989.

[21] E. V. Dan'shchikov, V. A. Dymshakov, A. M. Dykhne, F. V. Lebedev, and B. P. Rysev, "Characteristics of a surface acoustic wave generated by a moving laser beam," *Soviet Physical Acoustics*, vol. 33, no. 6, pp. 602–605, 1987.

[22] K. Yamanaka, Y. Nagata, and T. Koda, "Selective excitation of single-mode acoustic waves by phase velocity scanning of a laser beam," *Applied Physics Letters*, vol. 58, no. 15, pp. 1591–1593, 1991.

[23] K. A. Nelson, R. J. D. Miller, D. R. Lutz, and M. D. Fayer, "Optical generation of tunable ultrasonic waves," *Journal of Applied Physics*, vol. 53, no. 2, pp. 1144–1149, 1982.

[24] 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.

[25] Y. Shen and P. Hess, "Real-time detection of laser-induced transient gratings and surface acoustic wave pulses with a michelson interferometer," *Journal of Applied Physics*, vol. 82, no. 10, pp. 4758–4762, 1997.

[26] J. Huang, S. Krishnaswamy, and J. D. Achenbach, "Laser generation of narrow-band surface waves," *Journal of the Acoustical Society of America*, vol. 92, no. 5, pp. 2527–2531, 1992.

[27] K. Yamanaka, O. V. Kolosov, Y. Nagata, T. Koda, H. Nishino, and Y. Tsukahara, "Analysis of excitation and coherent amplitude enhancement of surface acoustic waves by the phase velocity scanning method," *Journal of Applied Physics*, vol. 74, no. 11, pp. 6511–6522, 1993.

[28] H. Sato, H. Cho, H. Nishino, H. Ogiso, and K. Yamanaka, "Evaluation of surface defects using surface acoustic waves generated by phase velocity scanning of laser interference fringes," *Japanese Journal of Applied Physics*, vol. 35 part 1, no. 5B, pp. 3066–3069, 1996.

[29] H. Köymen and A. Atalar, "Focusing surface waves using an axicon," *Applied Physics Letters*, vol. 47, no. 12, pp. 1266–1268, 1985.

[30] B. G. Kim, J. O. Lee, and S. Lee, "New design and analysis for point-focusing of surface waves in contact testing," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 40, no. 2, pp. 162–166, 1993.

[31] B. G. Kim, S. Lee, M. Enoki, and T. Kishi, "Development and aplication of ultrasonic transducer for intensive focusing of Rayleigh waves," *Ultrasonics*, vol. 36, pp. 825–836, 1998.

[32] C. K. Jen, P. Cielo, J. Bussiere, F. Nadeau, and G. W. Farnell, "Phase variation of focused surface acoustic wave," *Applied Physics Letters*, vol. 46, no. 3, pp. 241–243, 1985.

[33] M.-H. Noroy, D. Royer, and M. Fink, "The laser-generated ultrasonic phased array: analysis and experiments," *Journal of the Acoustical Society of America*, vol. 94, no. 4, pp. 1934–1943, 1993.

[34] J. W. Wagner, *Ultrasonic Measurement Methods*, vol. 19 of *Physical Acoustics*, ch. 5, pp. 201–266. San Diego, CA, USA: Academic Press, 1990. (Editors: R N Thurston and Allan D Pierce).

[35] J.-P. Monchalin, "Optical detection of ultrasound," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 33, no. 5, pp. 485–499, 1986.

[36] J.-P. Monchalin, "Heterodyne interferometric laser probe to measure continuous ultrasonic displacements," *Review of Scientific Instruments*, vol. 56, no. 4, pp. 543–546, 1985.

[37] M. P. Petrov, S. I. Stepanov, and G. S. Trofimov, "Non-stationary emf in the heterogeneously illuminated photoconductor," *Pisma V Zhurnal Tekhnicheskoi Fiziki*, vol. 12, no. 15, pp. 916–921, 1986.

[38] A. S. Murfin and R. J. Dewhurst, "Estimation of wall thinning in mild steel using laser ultrasound Lamb waves and a non-steady-state photo-emf detector," *Ultrasonics*, vol. 40, pp. 777–781, 2002.

[39] M. Klein, B. Pouet, and P. Mitchell, "Photo-emf detector enables laser ultrasonic receiver," *Laser Focus World*, vol. 36, no. 8, pp. 25–27, 2000.

[40] B. A. Williams and R. J. Dewhurst, "Differential fibre-optic sensing of laser-generated ultrasound," *Electronics Letters*, vol. 31, no. 5, pp. 391–392, 1995.

[41] J.-P. Monchalin, "Broadband optical detection of ultrasound by optical sideband stripping with a confocal fabry-pérot," *Applied Physics Letters*, vol. 55, no. 16, pp. 1612–1614, 1989.

[42] J. H. Jong, P. H. Brodeur, E. F. Lafond, J. P. Gerhardstein, and B. M. Pufahl, "Laser ultrasonics for noncontact measurement of Lamb waves in static and moving paper," *Journal of Pulp and Paper Science*, vol. 26, no. 10, pp. 358–366, 2000.

[43] R. Adler, A. Korpel, and P. Desmares, "An instrument for making surface waves visible," *IEEE Transactions on Sonics and Ultrasonics*, vol. SU-15, no. 3, pp. 157–161, 1968.

[44] G. Alers, M. A. Tennison, R. B. Thompson, and B. R. Tittmann, "Visualisation of surface elastic waves on structured materials," *Ultrasonics*, vol. 11, no. 4, pp. 174–177, 1973.

[45] W. P. Robbins, R. K. Mueller, and E. Rudd, "Thin-film characterization using a scanning laser acoustic microscope with surface acoustic waves," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 35, no. 4, pp. 477–483, 1988.

[46] V. I. Anisimkin, V. A. M. Luprano, M. Penza, G. Montagna, and A. I. Morozov, "Characterization and imaging of SAW grooved transducer acoustic field by SLAM," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 44, no. 6, pp. 1232–1238, 1997.

[47] P. A. Fomitchov, S. Krishnaswamy, and J. D. Achenbach, "Compact phase-shifted sagnac interferometer for ultrasound detection," *Optics and Laser Technology*, vol. 29, no. 6, pp. 333–338, 1997.

[48] J. D. Hamilton and M. O'Donnell, "High frequency ultrasound imaging with optical arrays," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 45, no. 1, pp. 216–235, 1998.

[49] J. D. Hamilton, T. Buma, M. Spisar, and M. O'Donnell, "High frequency optoacoustic arrays using etalon detection," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 47, no. 1, pp. 160–169, 2000.

[50] H. K. Wickramasinghe, R. C. Bray, V. Jipson, C. F. Quate, and J. R. Salcedo, "Photoacoustics on a microscopic scale," *Applied Physics Letters*, vol. 33, no. 11, pp. 923–925, 1978.

[51] J. B. Hoyes, Q. Shan, and R. J. Dewhurst, "A non-contact scanning system for laser ultrasonic defect imaging," *Measurement Science and Technology*, vol. 2, pp. 628–634, 1991.

[52] R. J. Dewhurst and Q. Shan, "Through-transmission ultrasonic imaging of sub-surface defects using non-contact laser techniques," *Optics and Lasers in Engineering*, vol. 16, pp. 163–178, 1992.

[53] F. L. di Scalea and R. E. Green, Jr., "High-sensitivity laser-based ultrasonic c-scan system for materials inspection," *Experimental Mechanics*, vol. 39, no. 4, pp. 329–334, 1999.

[54] C. J. Fiedler and T. Ducharme, "The laser ultrasonic inspection system (LUIS) at the Sacramento Air Logistics Center," in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 16, (New York), pp. 515–522, Plenum Press, 1997.

[55] 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.

[56] A. D. W. McKie and R. C. Addison, Jr., "Practical considerations for the rapid inspection of composite materials using laser-based ultrasound," *Ultrasonics*, vol. 32, no. 5, pp. 333–345, 1994.

[57] P. W. Larraine, R. A. Hewes, and D. Drolet, "High resolution laser ultrasound detection of metal defects," in *Review of Progress in Quantitative Nondestructive Evaluation*, vol. 16, (New York), pp. 555–562, Plenum Press, 1997.

[58] D. Lévesque, A. Blouin, C. Néron, and J.-P. Monchalin, "Performance of laser-ultrasonic F-SAFT imaging," *Ultrasonics*, vol. 40, pp. 1057–1063, 2002.

[59] Y. Sugawara, O. B. Wright, O. Matsuda, M. Takigahira, Y. Tanaka, S. Tamura, and V. E. Gusev, "Watching ripples on crystals," *Physical Review Letters*, vol. 88, no. 18, pp. 185504–1–185504–4, 2002.

[60] Y. Sugawara, O. B. Wright, O. Matsuda, and V. E. Gusev, "Spatiotemporal mapping of surface acoustic waves in isotropic and anisotropic materials," *Ultrasonics*, vol. 40, pp. 55–59, 2002.

[61] O. B. Wright, Y. Sugawara, O. Matsuda, M. Takigahira, Y. Tanaka, S. Tamura, and V. E. Gusev, "Real-time imaging and dispersion of surface phonons in isotropic and anisotropic materials," *Physica B*, vol. 316–317, pp. 29–34, 2002.

[62] Y. Sugawara, O. B. Wright, and O. Matsuda, "Real-time imaging of surface acoustic waves in thin films and microstructures on opaque substrates," *Review of Scientific Instruments*, vol. 74, no. 1, pp. 519–522, 2003.

[63] D. H. Hurley and O. B. Wright, "Detection of ultrafast phenomena by use of a modified Sagnac interferometer," *Optics Letters*, vol. 24, no. 18, pp. 1305–1307, 1999.

[64] C. Glorieux, K. Van de Rostyne, J. D. Beers, W. Gao, S. Petillion, N. Van Riet, K. A. Nelson, J. F. Allard, V. E. Gusev, W. Lauriks, and J. Thoen, "Acoustic waves at interfaces studied by laser ultrasonics," *Review of Scientific Instruments*, vol. 74, no. 1, pp. 465–469, 2003.

[65] M. Born and E. Wolf, *Principles of Optics.* Pergamon Press, 1964.

[66] V. V. Krylov and V. I. Pavlov, "Thermo-optical generation of surface acoustic waves in a solid," *Sov. Phys. Acoust.*, vol. 28, no. 6, pp. 493–494, 1982.

[67] A. A. Karabutov, "Laser excitation of surface acoustic waves: A new direction in opto-acoustic spectroscopy of a solid," *Sov. Phys Usp.*, vol. 11, no. 28, pp. 1042–1051, 1985.

[68] J. W. Goodman, *Introduction to Fourier Optics.* McGraw-Hill, 1968.

[69] E. Hecht, *Optics.* Reading, Massachusetts: Addison-Wesley, 1998.

[70] P. Horowitz and W. Hill, *The Art of Electronics.* Cambridge, UK: Cambridge University Press, 1989.

[71] F. C. Cuozzo, E. L. Cambiaggio, J.-P. Damiano, and E. Rivier, "Influence of elastic properties on Rayleigh wave scattering by normal discontinuities," *IEEE Transactions on Sonics and Ultrasonics*, vol. SU-24, no. 4, pp. 280–289, 1977.

[72] F. Nadeau and D. A. Hutchins, "A study of the interaction of surface acoustic waves with slots using non-contact laser generation and detection of ultrasound," in *1984 IEEE Ultrasonics Symposium*, (Piscataway, NJ, USA), IEEE and UFFC, IEEE, 1984.

[73] K. Yamanaka and Y. Enomoto, "Fringe pattern around surface crack observed with scanning acoustic microscope," *Electronics Letters*, vol. 17, no. 18, pp. 638–640, 1981.

[74] M. G. Somekh, H. L. Bertoni, G. A. D. Briggs, and N. J. Burton, "A two-dimensional imaging theory of surface discontinuities with the scanning acoustic microscope," *Proceedings of the Royal Society of London A*, vol. 401, pp. 29–51, 1985.

[75] J. A. Cooper, R. A. Crosbie, R. J. Dewhurst, A. D. W. McKie, and S. B. Palmer, "Surface acoustic wave interactions with cracks and slots: a noncontacting study using lasers," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 33, no. 5, pp. 462–470, 1986.

[76] A. Briggs, "Acoustic microscopy—a summary," *Reports on Progress in Physics*, vol. 55, pp. 851–909, 1992.

[77] S. Sathish and R. W. Martin, "Quantitative imaging of Rayleigh wave velocity with a scanning acoustic microscope," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 49, no. 5, pp. 550–557, 2002.

[78] D. Husson, S. D. Bennett, and G. S. Kino, "Measurement of stress with surface waves," *Materials Evaluation*, vol. 43, pp. 92–100, 1985.

[79] C. M. Sayers, D. R. Allen, G. E. Haines, and G. G. Proudfoot, "Texture and stress determination in metals by using ultrasonic Rayleigh waves and neutron diffraction," *Philosophical Transactions of the Royal Society of London A*, vol. 320, pp. 187–200, 1986.

[80] S. W. Meeks, D. Peter, D. Horne, K. Young, and V. Novotny, "Microscopic imaging of residual stress using a scanning phase-measuring acoustic microscope," *Applied Physics Letters*, vol. 55, no. 18, pp. 1835–1837, 1989.

[81] H. Coufal, K. Meyer, R. K. Grygier, P. Hess, and A. Neubrand, "Precision measurement of the surface acoustic wave velocity on silicon single crystals using optical excitation and detection," *Journal of the Acoustical Society of America*, vol. 95, no. 2, pp. 1158–1160, 1994.

[82] Y. Hong, S. D. Sharples, M. Clark, and M. G. Somekh, "Rapid measurement of surface acoustic wave velocity on single crystals using an all-optical adaptive scanning acoustic microscope," *Applied Physics Letters*, vol. 83, no. 16, pp. 3260–3262, 2003.

[83] H. Matthews, *Surface Wave Filters.* New York: John Wiley and Sons, Inc, 1977.

[84] F. S. Hickernell, *Ultrasonic Instruments and Devices II*, vol. 24 of *Physical Acoustics*, ch. 3, pp. 135–207. San Diego, CA, USA: Academic Press, 1999. (Editors: R N Thurston and Allan D Pierce).

[85] H. F. Tiersten, "Elastic surface waves guided by thin films," *Journal of Applied Physics*, vol. 40, no. 2, pp. 770–789, 1969.

[86] T. J. I. Bromwich, "On the influence of gravity on elastic waves, and, in particular, on the vibrations of an elastic globe," *Proceedings of the London Mathematical Society*, vol. 30, pp. 98–120, 1898.

[87] M. J. S. Lowe, "Matrix techniques for modeling ultrasonic waves in multilayered media," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 42, no. 4, pp. 525–542, 1995.

[88] R. K. Tyson, *Principles of adaptive optics.* Academic Press, 1991.

[89] R. W. B. Stephens, *Underwater Acoustics.* Wiley, 1970.

[90] M. Fink, "Time reversal of ultrasonic fields- parts I,II,III," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 39, no. 5, pp. 555–591, 1992.

[91] A. Derode, P. Roux, and M. Fink, "Robust acoustic time reversal with high-order multiple scattering," *Physical Review Letters*, vol. 75, no. 23, pp. 4206–4209, 1995.