Inspection of Additively Manufactured Components by Spatially Resolved Acoustic Spectroscopy

Paul Dryburgh, MEng

Thesis submitted to The University of Nottingham for the degree of Doctor of Engineering

March 2022



Abstract

This thesis presents the development of spatially resolved acoustic spectroscopy (SRAS) for the materials characterisation of new alloys and additive manufacturing specimens. SRAS is an acoustic microscopy technique that uses laser-generated and detected ultrasound to probe the material properties of crystalline specimens. Laser generation allows the surface acoustic wave velocity to be measured across the specimen and is a valuable tool for highlighting the grain structure. In anisotropic materials, inversion of the SAW velocity in multiple directions is used to determine the crystalline orientation.

SRAS measurements are presented in two notable additive manufacturing techniques, wire-arc additive and laser powder bed fusion. This section is focused on the study of prepared surfaces and utilises the powerful forward model for calculating SAW velocities to reveal pertinent information about material microstructure and, in turn, the build process. Examples include detecting the elimination of prior- β microstructures in wire-arc specimens and the measurement of crystalline texture in high-silicon steel. These results demonstrate the power of SRAS as a materials characterisation tool in additive manufacturing, providing the ability to map the grain structure and extract critical crystalline information rapidly.

This analysis relies on accurate knowledge of the material elastic constants, and further results demonstrate the determination of crystallographic orientation is also highly sensitive to these values. Therefore, in this work, this inversion process is extended to allow the simultaneous determination of both orientation and elasticity from SAW velocity as measured by SRAS. In contrast to existing methods, this technique can be applied to polycrystalline specimens rather than specifically prepared single crystals, which are usually studied. Experimental results are presented for single-crystal nickel, along with more interesting polycrystalline CMSX-4 and titanium; the results show good agreement with electron backscatter diffraction measurements of crystallographic orientation and existing literature values for elastic constants. We believe this feature will significantly enhance the range of applications of SRAS and have a tangible and significant impact across materials and manufacturing sciences.

Having developed the technical case for the use of SRAS within additive manufacturing as a characterisation tool, the industrial case is considered. This focuses on the monetary and temporal consequences of a hypothetical SRAS inspection. This methodology is the basis for outlining the future potential of SRAS within the additive manufacturing sector. In the case of complex components where builds tend to fail (potentially due to an incorrect microstructure), there is a strong case for the integration of SRAS system in the build machine from a monetary viewpoint, along with the technical case developed above. Nevertheless, much work remains to be done in this field to realise SRAS within a functioning additive build system in the case of powder bed fusion.

Acknowledgements

I want to thank Keith, Peter and Nina at the RCNDE for facilitating my EngD and the countless opportunities the CDT has provided. I pay thanks to the AMTC team at Renishaw, particularly to Nick Jones, for their help and sponsorship of this work.

I am indebted to my supervisors, Adam, Richard and Matt, for their patient guidance and support throughout my EngD - along with the gentle encouragement to write when needed. Their breadth of knowledge in the field is boundless. Every day is an opportunity to learn from them, and I am fortunate to have had the opportunity to do so. I hope their sense of curiosity and commitment to worthy scientific advancement has rubbed off on me a little. I am also grateful to have had the opportunity to have been supervised by Steve Sharples, who offered invaluable advice and suggestions for my project and research more broadly, which I will always hold in my mind.

I have greatly enjoyed the opportunity to be a member of the Optics and Photonics group, and I'm unsure I would have made it through without the help of the 202-clique specifically. I am very grateful to Wenqi and Rikesh for their support and help as supervisors, fellow researchers and friends. I would particularly like to thank my deskmate, Don, for helping me not soften my accent one iota over the last five years. Moreover he has always being a kindred mechanical mind to talk a problem through, and it has been a pleasure to slowly break and have to fix most of the lab with him through our EngD's. I am also grateful to Matthias, who supervised my first year, made the transition to Nottingham exponentially easier, and always ensured ample coffee was to hand.

I want to thank Kally, who has remained patient, supportive and loving throughout my unusual lab hours and days spent preoccupied with my work. Despite the stress of work and COVID, I will always treasure the moments of fun and joy spent with her.

I am eternally grateful to my family, including George, Alfie and Reggie, for their continued unconditional support and giving me the platform to follow my ambition. In particular, I'd like to thank my Mum, her love, selflessness, and determination are an endless source of inspiration to me, and I wouldn't be the person I am without her.

Contents

| 1 | Introduction | | |
|---|--------------|---|----|
| | 1.1 | A brief background to the problem | 1 |
| | 1.2 | Outline of work | 2 |
| | 1.3 | Additive Manufacturing | 5 |
| | | 1.3.1 The need for inspection \ldots | 7 |
| | 1.4 | Physical properties of materials | 8 |
| | | 1.4.1 Hooke's law | 9 |
| | | 1.4.2 Stress and strain | 9 |
| | | 1.4.3 Crystallography | 11 |
| | | 1.4.4 Elastic constants | 12 |
| | | 1.4.5 Microstructure | 14 |
| | 1.5 | Non-destructive evaluation and materials characterisation $\ldots \ldots \ldots \ldots \ldots$ | 15 |
| | | 1.5.1 The use of non-destructive evaluation | 15 |
| | | 1.5.2 Industrial NDE techniques | 17 |
| | | 1.5.3 Materials characterisation | 20 |
| | 1.6 | Ultrasonics | 22 |
| | | 1.6.1 Resonance | 28 |
| | 1.7 | Laser ultrasonics | 28 |
| | | 1.7.1 Types of generation source | 29 |
| | | 1.7.2 Generation mechanisms | 30 |
| | | 1.7.3 Optical detection | 31 |
| | 1.8 | Surface wave measurements | 35 |
| | 1.9 | Summary | 37 |
| 2 | Sur | vey of existing literature and methods | 39 |
| | 2.1 | Introduction | 39 |
| | 2.2 | Additive manufacturing techniques and properties | 40 |
| | | · · · · · | |

| | | 2.2.1 Laser powder bed fusion $\ldots \ldots 4$ | 40 |
|---|-----|---|-----|
| | | 2.2.2 Wire-arc additive manufacturing | 46 |
| | | 2.2.3 Materials | 48 |
| | 2.3 | NDE and characterization tools for additive manufacturing | 52 |
| | | 2.3.1 X-ray computed tomography 5 | 58 |
| | | 2.3.2 Other traditional NDE approaches | 61 |
| | | 2.3.3 EBSD and diffraction techniques | 61 |
| | | 2.3.4 Ultrasound | 63 |
| | | 2.3.5 Laser ultrasound $\ldots \ldots \ldots$ | 65 |
| | 2.4 | Elastic constant measurements | 69 |
| | | 2.4.1 Measurement methods | 73 |
| | 2.5 | Spatially resolved acoustic spectroscopy | 79 |
| | | 2.5.1 Spatially resolved acoustic spectroscopy for additive manufacturing 8 | 83 |
| | | 2.5.2 Calculation of surface acoustic wave velocity | 84 |
| | 2.6 | Summary | 86 |
| 3 | Spa | ially Resolved Acoustic Spectroscopy 8 | 89 |
| 0 | 3.1 | Early SBAS - k -SBAS | 89 |
| | 3.2 | The EMDA system and f-SRAS | 91 |
| | | 3.2.1 Generation optical path | 92 |
| | | 3.2.2 Detection optical path | 96 |
| | | 3.2.3 Signal detection and acquisition | 97 |
| | | 3.2.4 Realising a scanning instrument | 98 |
| | | 3.2.5 Data processing | 100 |
| | | 3.2.6 Spatial resolution | 101 |
| | | 3.2.7 Velocity resolution | 103 |
| | | 3.2.8 Typical data | 104 |
| | 3.3 | Optically rough surfaces | 108 |
| | 3.4 | Multi-angle scanning approaches | 111 |
| | | 3.4.1 Mask Rotation | 111 |
| | | 3.4.2 Sample Rotation | 111 |
| | 3.5 | Orientation results | 113 |
| A | NT | anical emperature of CDAS measurements | 115 |
| 4 | | Introduction | 11F |
| | 4.1 | Colouisting the SAW velocity forward problem | 110 |
| | 4.2 | Calculating the SAW velocity - forward problem | 110 |

| | | 4.2.1 | The wave equation | 116 |
|---|-----------|----------|--|------|
| | | 4.2.2 | Bulk wave velocities | 117 |
| | | 4.2.3 | Surface acoustic wave velocity | 122 |
| | | 4.2.4 | Calculation time | 133 |
| | | 4.2.5 | Limitations | 136 |
| | 4.3 | Inversi | on | 138 |
| | | 4.3.1 | Objective function $\ldots \ldots \ldots$ | 139 |
| | | 4.3.2 | Finding the orientation | 141 |
| | | 4.3.3 | Finding the elasticity | 143 |
| | | 4.3.4 | Determining unknown elastic constants and orientation $\ldots \ldots \ldots \ldots$ | 145 |
| | | 4.3.5 | Inversion summary | 145 |
| | 4.4 | Simula | tion of SRAS measurements | 146 |
| | | 4.4.1 | Time domain signals | 146 |
| | | 4.4.2 | Velocity surfaces | 147 |
| | | 4.4.3 | Simulated polycrystalline specimens | 148 |
| | 4.5 | Summ | ary | 154 |
| 5 | Me | asurem | ents in additive manufacturing specimens | 155 |
| 0 | 5.1 | Wire-a | arc additive manufacture | 156 |
| | | 5.1.1 | Materials and Methods | 157 |
| | | 5.1.2 | Side-wall measurements $(x - z \text{ plane})$ | 159 |
| | | 5.1.3 | Cross-section measurements $(y - z \text{ plane})$ | 161 |
| | | 5.1.4 | Examination of Interface | 164 |
| | | 5.1.5 | Crystallography | 164 |
| | | 5.1.6 | Minimum stiffness direction | 172 |
| | | 5.1.7 | Extracting build process metrics from acoustic data | 173 |
| | | 5.1.8 | Preliminary as-deposited results | 178 |
| | 5.2 | Laser | powder bed fusion | 183 |
| | | 5.2.1 | Ti-6Al-4V | 183 |
| | | 5.2.2 | Nickel superallov | 185 |
| | | 5.2.3 | Fe-Si | 187 |
| | 5.3 | Summ | ary | 195 |
| c | D! | dim = 41 | a single equated electricity and the coinstation | 100 |
| U | r ind | ung th | le single-crystal elasticity and the orientation | 100 |
| | 0.1 | Introd | | 199 |
| | 6.2 | Accura | acy of orientation determination in titanium | -200 |

| | | 6.2.1 | Representing orientation data | 202 |
|---|-----|----------|--|-----|
| | | 6.2.2 | Orientation dependent accuracy | 204 |
| | | 6.2.3 | Velocity resolution | 206 |
| | | 6.2.4 | Number of measured propagation directions and interval | 207 |
| | | 6.2.5 | Orientation accuracy - experimental | 211 |
| | 6.3 | Elastic | c constant determination | 217 |
| | | 6.3.1 | Alogrithm | 217 |
| | | 6.3.2 | Influence of elasticity on ultrasonic velocity | 218 |
| | | 6.3.3 | Experimental results | 223 |
| | | 6.3.4 | Analysis of error and uncertainty | 229 |
| | | 6.3.5 | Implementation | 231 |
| | 6.4 | Summ | ary | 233 |
| 7 | The | e case f | for SRAS in additive manufacturing | 238 |
| | 7.1 | Introd | luction | 238 |
| | 7.2 | Existi | ng cost-time models | 239 |
| | | 7.2.1 | Early work | 239 |
| | | 7.2.2 | Modelling failure and defects | 240 |
| | 7.3 | System | n costing | 241 |
| | | 7.3.1 | Additive build system | 241 |
| | | 7.3.2 | SRAS | 242 |
| | | 7.3.3 | X-ray computed tomography | 244 |
| | 7.4 | Defini | tion of defects and the impact on processing routes | 245 |
| | 7.5 | Cost a | and time model \ldots | 248 |
| | | 7.5.1 | Sample build | 248 |
| | | 7.5.2 | Time | 249 |
| | | 7.5.3 | Cost based on annual production | 255 |
| | 7.6 | Build | failure | 259 |
| | | 7.6.1 | Scrap rate | 259 |
| | | 7.6.2 | Opportunity to repair | 265 |
| | | 7.6.3 | Cost of repair | 266 |
| | | 7.6.4 | Preliminary demonstration of repair by reworking surface defects | 266 |
| | 7.7 | Implei | mentation and improvements | 273 |
| | | 7.7.1 | Scan rate and time | 273 |
| | | 7.7.2 | Example system | 275 |

| | 7.8 | Summary | 275 | |
|--------------|---|--|-----|--|
| 8 | Fut | ure work 27 | | |
| | 8.1 | Introduction | 279 | |
| | 8.2 | Instrumentation development | 279 | |
| | | 8.2.1 Random access scanning | 279 | |
| | | 8.2.2 Multi-angle scanning | 280 | |
| | | 8.2.3 Spatial resolution | 282 | |
| | | 8.2.4 Velocity resolution | 284 | |
| | | 8.2.5 Temperature control | 285 | |
| | 8.3 | High temperature measurements | 285 | |
| | 8.4 | Verification of elastic constant measurements | 287 | |
| | 8.5 | Additional information for inversion | 288 | |
| | | 8.5.1 Multiple wave modes | 288 | |
| | | 8.5.2 Amplitude | 289 | |
| | 8.6 | Spatially resolved properties | 291 | |
| | 8.7 | Hybrid solvers | 292 | |
| | | 8.7.1 Artificial neural network classifier | 293 | |
| | | 8.7.2 Optimisation algorithms | 293 | |
| | 8.8 | Textured materials | 294 | |
| | 8.9 | Closed loop-feed back for microstructure control in additive manufacturing $\ . \ . \ .$ | 297 | |
| 9 | Cor | nclusions | 299 | |
| A | Ela | stic constants in cubic materials | 304 | |
| в | B Elastic constants in hexagonal materials | | | |
| \mathbf{C} | Ger | neralised anisotropic elastic modulus relations | 308 | |
| D | D Generalised anisotropic bulk wave velocity and elasticity relations | | | |
| \mathbf{E} | E R-value calculation | | | |
| P | Pibliography 91 | | | |
| D | Bibliography 314 | | | |

Chapter 1

Introduction

1.1 A brief background to the problem

In recent years additive manufacturing has captivated the interest of high-value manufacturing industries. Driven by increased design freedom, reduced lead times, functional optimisation and small batch production, sales of commercial additive systems have increased dramatically, driving a boom in academic research. In some respects the technology has already matured, with initial adopters in medical and dental industries utilising the 'free' customisation to offer implants tailored to each patient [1]. However, the ability to confidently certify the performance of additively manufactured components remains a stumbling block to up-take. The new paradigm of characterisation tools required to meet the demands of additive manufacture presents a significant scientific challenge.

A vast body of literature is dedicated to describing a variety of failure mechanisms which plague the technology. Clearly, without improving the quality control the notion of seeing additive components in truly safety critical applications (such as nuclear, pressure vessels or aero-engines) seems a significant way off.

For a tangible example that demonstrates both the advantages and fragility of additive manufacturing, the reader need only look to the cycling event at the recent Tokyo Olympics. The Australian team made use of handlebars which were made from the common aerospace alloy Ti-6Al-4V fabricated by laser additive manufacturing, using a Renishaw AM250 system. By using AM the handlebars can be designed to fit measurements of each individual, allowing the rider to attain a more aerodynamic position, reducing drag and increasing efficiency. Producing low volumes of performance components with a high degree of customisation between each component is a challenge that additive manufacturing is ideally suited to. However, catastrophic failure of these handlebars occurred during the race, snapping into two separate pieces and injuring the rider. Whilst there are a lack of examples of failures in more critical applications, this merely reflects the conservative approach manufactures are adopting, preferring to slowly introduce additive processes in non-critical applications, such as cockpit and fuselage interiors. Whilst sensible, the 'walk before you can run' approach is not a sustainable outlook for the industry, only by manufacturing these high-value components can the large capital investment in additive manufacturing development begin to be recouped [2]. Ensuring confidence in the performance of additive manufacturing components is an essential step towards this.

Non-destructive evaluation has an important role to play in meeting this challenge, allowing the material properties of additive components to be probed, as a step towards both performance certification and feedstock back into process control and improvement. Within the field of nondestructive evaluation, ultrasound is an attractive tool for this problem; firstly, detecting the presence of ultrasound through a material can reveal structural anomalies such as cracks. Secondly, the velocity of acoustic waves are a function of the underlying material properties, offering the potential to determine salient material characteristics, in a bid to certify the performance of each component. Laser ultrasound techniques are particularity promising thanks to their non-contact nature, opening up the potential to inspect the additive build process and 'certify as you build' [3].

1.2 Outline of work

This thesis presents development of the laser ultrasound technique spatially resolved acoustic spectroscopy (SRAS) towards meeting the demands presented by additive manufacturing. The work presented in this thesis concerns the development of a method to elucidate the physical properties of specimens produced by additive manufacturing, by exploiting the interaction between surfaces acoustic waves (SAWs) and solid materials.

The primary aims of this work are summarised as:

- Investigate the information contained within SRAS measurements on additive manufacturing specimens, with focus on microstructure information.
- Advancing the current inversion methodology to facilitate the determination of the elastic constants in polycrystalline materials. To contextualise these results, develop a framework to understand the accuracy of orientation and elastic constant determination.
- Consider the appropriateness of SRAS within the industrial additive manufacturing context, considering the both the technical and commercial challenges.

The overarching theme of this work is to investigate the sensitivity of SRAS to the physical state of additive manufacturing specimens - covering both the microstructure and defect formations in these materials. To this end, this goal is firstly realised through experimental results in two contrasting additive manufacturing methods, giving focus to the microstructure information that can be discerned from SRAS measurements. There is a wealth of commentary that could be provided on the specimen from the ultrasonics data that has been acquired, but where possible attention is paid to metrics which reflect the fitness for purpose (or lack thereof) of the material.

Through the course of the project, the scope of the work broadened, as a logical extension to the previous aim, to include the development of the underlying methods to allow the extraction of the elastic properties in polycrystalline materials - an issue long pre-dating additive manufacturing but ever more pressing thanks to the ability to easily vary the composition of additive components insitu. This represents a significant step forward for SRAS and ultrasonics more broadly, naturally a significant portion of the work contained in this thesis is given over to the development of methods for extracting information from the acoustic data, with a view to these new functions being eminently applicable to additive manufacturing. Underpinning these results are of course a broad body of numerical techniques along with an established experimental system, and whilst not the primary focus of the work, these are given appropriate attention in stand-alone chapters, including developments which have made the realisation of non-contact method of acquiring images of the SAW velocity and the subsequence development of a fast SRAS system, capable of determining crystallographic orientation. These advances are well explained in the theses of S. D. Sharples [4] and W. Li [5].

The initial objective of this work was to develop an industrially relevant laser ultrasound technique to non-destructively inspect additive manufacturing specimens in-situ. However, given the significant challenges of integrating a SRAS system with the build environment, it became evident that a strong scientific case was first required to justify this pursuit. This is covered by the first two objectives of this thesis, but the economic question remained. As such, the final aim of this thesis pivoted away from creating an in-situ system, to considering if such as system would be viable from an industrial perspective - focussing on the economic cost and benefits a SRAS system could provide through measurements of the microstructure.

As a relevant aside, funding has recently been secured to integrate a SRAS system within an additive build system - providing the opportunity for the techniques developed under the first two aims to be applied in anger.

The following is a brief synopsis of the content contained within each chapter of this thesis:

- Chapter 1 contains a brief introduction to additive manufacturing, specifically focussing on the need for materials characterisation and inspection. This is followed by an introduction to the fundamentals of materials characterisation and non destructive testing, a summary of viable common techniques is provided with particular interest given over to laser ultrasonics.
- Chapter 2 reviews the state of the art in materials inspection and characterisation, within additive manufacturing, in particular, focussing on the development of techniques for probing the microstructure and elastic properties of additive components. To properly summarise the field of elastic constant determination, results and techniques beyond the field of additive manufacturing are discussed, including ultrasonic techniques. A review of surface wave measurement and modelling for the determination of crystalline properties concludes this chapter.
- Chapter 3 is concerned with the experimental laser ultrasound system used throughout this thesis to measure the surface acoustic wave velocity. This includes a description of the system, summarised results to provide an overview of its performance and a brief foray into the accompanying data processing.
- Chapter 4 outlines a suite of numerical tools that have been developed in lockstep with SRAS. The first half of the chapter describes the process of calculating the SAW velocity. Whilst the latter half is mainly concerned with the inversion of the measured surface acoustic wave velocity to determine the elastic constants and/or crystallographic orientation in the probed material.
- Chapter 5 presents experimental results from specimens fabricated using two of the most common additive manufacturing methods, wire-arc (WAAM) and laser powder bed fusion (L-PBF). The primary focus is on using SRAS measurements to extract useful material information. In the case of WAAM, this focusses on the detection of the undesirable large prior-β grain structure in Ti-6Al-4V specimens. Three L-PBF materials are considered, titanium, nickel and high-silicon steel. These materials have been selected as they each present distinct microstructures and crystallographic texturing. The analysis of titanium commence with a simple discussion on surface and near-surface defects measured by SRAS. In high silicon steel SAW velocity measures are used to infer the crystallographic texture of the specimen. Both sections are broadly structured such that the complexity of results and analysis for each technique increases as the chapter progresses.
- Chapter 6 details orientation and elasticity measurements using SRAS. The chapter begins with a discussion on the accuracy of crystallographic orientation determination in hexagonal

crystal structure materials, focussing on titanium. Simulated results are presented giving an indication of the accuracy of determination across a wide range of experimental parameters, with experimental results provided for validation. This section concludes by hypothesising that uncertainty in the elastic constants of the forward model (introduced in Chapter 4) is the greatest source of error. This leads into the latter part of this chapter which discussed the ability to simultaneously determine the single crystal elastic constants (C_{ij}) and crystallographic orientation in polycrystalline materials, from measure surface acoustic wave velocities. Again, simulated and experimental results are presented, with focussed paid to the sensitivity of elastic constants to both orientation and experimental parameters.

- Chapter 7 is particularly focussed on the industrial and commercial aspects of this work. Based on a commercial additive manufacturing system a cost-time model for inspection is developed. This model is then applied to assess the case for SRAS within the production cycle. Ancillary investigations also consider the technical merits of in-process inspection not covered in Chapters 6 and 7, including the ability to detect and then 'correct' defects, this is then related back to the afore mentioned cost-time model to conclude the chapter.
- Chapters 8 and 9 conclude the thesis. This commences by outlining a body of potential future work, specifically designed to address issues arising from the results presented in chapter 5 and 6. Before presenting a final overview of the work performed in the thesis, summarising the noteworthy outcomes and novelty of the work.

1.3 Additive Manufacturing

Modern day additive manufacturing can trace its lineage back to 1984 and Chuck Hull's patent for 'Apparatus for production of three-dimensional objects by 'stereolithography', in essence the first 3D printer. Hall's work began the field of rapid prototyping, the idea that the technique could be used to quickly iterate design. This has never been truer, as designers, engineers and hobbyists around the world now have access to 3D printers at their desks. Later, the idea of using 3D printing to create functional parts rather than models gained traction. Additive manufacturing has become a popular term for this field, particularly when related to methods using metal feedstock and will be used throughout this thesis.

There is now a multitude of metal-based techniques (and an even greater list of synonymous names for these, with little standardisation of the nomenclature). To make some sense of these, and to place the two additive techniques studied in thesis within the context of wider field, it is possible to classify AM techniques based on the energy source and feedstock form, as shown in fig. 1.1. Experimental results presented in this work focus on laser powder bed fusion and wire-arc fabricated specimens, these two techniques were selected for study, in part because they share neither feedstock or energy source and as a result have quite distinct challenges for inspection. A comprehensive review on these techniques will be presented in Chapter 2.



Figure 1.1: Overview of AM process by feedstock and deposition type. Experimental results in specimens manufactured by L-PBF and WAAM are presented in this thesis.

Additive manufacturing promises to herald a new age in fabrication, by improving utilisation of raw materials, removing many of the design constraints found with traditional techniques, and offering 'free' customisation and complexity by removing tooling costs. The cost of suitable feedstock and commercial metal additive systems is significant, therefore AM is generally considered well suited to 'high-value' industries such as aerospace [6], tooling [7] and medical applications [8]. In medical applications the principal benefit is the capability to tailor components to each patient. For heavy industries though, it is the a removal of the machining constraints and improved utilisation of raw materials which are of greatest interest. The need to ensure tooling access imposes significant limitations on designers and work to prevent the realisation of the optimal component geometry [9]. Thus, the ability to create more complex geometry with fewer design constraints will allow the realisation of optimal topology of components, improving overall design efficiency in performance components. Indeed, components optimised for additive manufacturing look remarkably different to their subtractive counterparts [10, 11].

Further benefits of additive manufacture include, reduction in material wastage in comparison to subtractive techniques, due to the near-net shape process. This offers large potential monetary savings for high performance materials, such as the nickel superalloy CM247-LC used in turbine hot stages, due to the high billet cost associated with these exotic alloys [12].

1.3.1 The need for inspection

With additive manufacturing the number of process variables has increased greatly, compared to traditional fabrication methods. This is made more complex because the properties of the final component are a function of the build process, component design and the feedstock material, to say a part has been manufactured by powder bed fusion in titanium is an inadequate description of the underlying physical properties. Building a structurally sound and reliable component requires specification of an optimum set of process variables that affect the transient temperature fields, geometry and cooling rates, however selecting ideal parameters for a new build is not straightforward - instead an iterative 'trial and error' process is required. Additionally, the proliferation of powder-feedstock additive manufacturing techniques has opened the door to easy in-situ alloying, allowing new alloys with bespoke proprieties to be created rapidly [13].

A multitude of defect mechanisms threaten the integrity of a build, these can include obvious geometrical errors to the more subtle internal porosity [14], micro-cracking [15] and heterogeneous microstructure formation [16]. To expand the final point, additive manufacture specimens exhibits variation in chemical composition across the build [17], and small changes in the build strategy are capable of causing dramatic changes in elastic anisotropy of the manufactured component [18]. The spatially varying microstructure makes it essential to probe larger areas of the microstructure as homogeneity cannot be assumed. Tools to measure the elastic and crystallographic characteristics of additive components are of particular importance, as these define functional properties including Young's modulus and are typically unknown for an individual part due to this spatially varying chemistry.

Traditional statistically-based qualification processes for metallic materials is well defined [19], requiring extensive testing which can take several years to complete with costs running into the millions [20]. This is typically based around mechanical testing to establish the components yield and ultimate tensile strength [21], performed over a large number of specimens to obtain high confidence in performance across each build set. For instance, when producing a billet of titanium for the aerospace industry with a fixed process, this style of testing allows both the average performance and the part-to-part variation in performance to be captured. This then feeds into the design process, ensuring the weakest of components is still suitably safe - using 'overdesign' to guarantee reliability. Of course, from an industrial standpoint, the value of testing is to demonstrate that there is a minimal variation in performance. By way of compromise between the intrinsic material variations and 'overdesign' to maintain safety, the aerospace industry follows a 95% rule to control material variation [22]. Referring back to our titanium example to illustrate this, when following this legislation 95% of the billets produced by this process would be expected to possess

a yield strength ≥ 240 MPa, for example.

Much of this data is of course made obsolete when there is a change to the process or input, promoting conservatism within the manufacturing industry. These approaches have served us well, evidenced by the rarity of catastrophic failures across aerospace, nuclear and civil engineering industries. However, additive manufacturing is at its most efficient when used for low volume production, even producing one-off components [23], making established qualification regimes an unsuitable approach [24]. Turning the problem of low-run qualification on its head also presents a significant opportunity; if the material that sees service can be characterised there is an opportunity realise a new paradigm in material efficiency, removing the need for material design tolerances to ensure safety.

One aspect of additive manufacturing which is perhaps less well acknowledged is the length of time builds run for, with many components taking over thirty hours to fully build. The implication being that defects and conformity cause the whole time to be wasted. On this point, the layer-by-layer nature of additive manufacturing does however provide an opportunity. If the formation of defects and microstructure inhomogeneities can be detected during the build process then there is the potential to, at a minimum, cancel the build to save the time and powder. More desirably, the information obtained from inspection could be used to implement correction, be it 'infilling' voids [25] or varying the build strategy to break up grain growth [26].

As such, it is agreed the lack of ability to assure these physical properties conform to the design parameters is the greatest roadblock in the uptake of additive manufacturing, a new suite of tools is essential if the potential benefits additive manufacturing offers are to be realised. As others have concluded there is a demand for inspection tools which can verify the build quality (without impacting the build itself).

Having set out the need and opportunity for a tool cable of probing the microstructure of additive manufacturing specimens, the remainder of this chapter introduces the fundamentals of the elastic properties of materials and basis concepts of inspection and characterisation of materials - the building blocks upon which the technical work of this thesis is based.

1.4 Physical properties of materials

This thesis uses the experimental measurement of surface acoustic waves (SAW) velocity to determine the elastic properties of anisotropic materials. Chapter 4 presents a numerical method for calculating the SAW velocity and relating this to crystalline materials properties. To express these ideas, we first need to introduce the theory of linear elasticity and the requisite nomenclature to describe anisotropic crystals. The principles summarised in the subsequent sections are elementary concepts continuum mechanics, and great depth of detail on their derivation can be found in many textbooks, including the work of Cowin [27]. It is important to note this section (and this chapter on the whole) does not claim this to be novel.

1.4.1 Hooke's law

In 1660, Robert Hooke observed that a linear relationship existed between the force applied to a spring and the subsequent extension it experiences. This elementary concept is often introduced as

$$F = kx \tag{1.1}$$

where F is the applied force, x is the displacement and k is commonly called the spring constant, essentially a proportionality factor. Of course, many materials and objects beyond springs obey this law, materials behaviour which can be approximated by this are known as linear-elastic (or Hookean) materials.

Hooke's law is valid in a materials elastic regime, where non-linear effects are small and implies deformations are reversible. In order to generalise Hooke's law to three dimensions the fundamental concepts of stress and strain need to be introduced.

1.4.2 Stress and strain

The force per unit area generated when a force is applied to an object of a finite cross-section is known as stress. These stresses can be generated perpendicular, normal, or parallel to the surface. In general, the stress σ_{ij} describes the force per unit area corresponding to a force applied in the *i* direction onto a plane with its normal in the *j* direction. A normal stress is the result of the force direction and the plane normal lying parallel, conversely a shear stress is generated when the force direction is perpendicular to the plane normal. The stress tensor σ_{ij} is symmetric such that $\sigma_{ij} = \sigma_{ji}$, implying only six of the nine components are required to describe the stress tensor.

With respect to a reference position, a body under stress experiences a deformation in its shape. This deformation may be normal, fig. 1.3a, or shearing, fig. 1.3b. The strain tensor, ϵ_{kl} , describing the deformation is then given as

$$\epsilon_{kl} = \frac{1}{2} \left(\frac{\partial u_k}{\partial l} + \frac{\partial u_l}{\partial k} \right) \tag{1.2}$$

This is valid for infinitesimally small strains, within the elastic limit of a material.



Figure 1.2: Definition of stress, including both normal, σ , and shear, τ stresses.



(a) Normal Strain, the displacement u_x increases with increasing x (independent of y).

(b) Shear strain, the displacement u_x increases with increasing y.

Figure 1.3: Definition of the two forms of strain.

Returning then to the definition of Hooke's law, equation 1.1, we can use this to formalise the relationship between stress and strain for three-dimensional linear-elastic bodies, giving equation 1.3.

$$\sigma = E\epsilon \tag{1.3}$$

where σ and ϵ are stress and strain, respectively. E is the familiar term Young's modulus, replacing the spring constant. For a body under bulk or shear stress the Young's modulus is replaced by the bulk modulus, K, or shear modulus, G, respectively. In essence, these stiffness terms provide a linear map of proportionality between stress and strain. The form of equation 1.3 is for an isotropic body, where the stress-strain response is independent of direction. However, many materials exhibit direction dependent properties due to the symmetry of their crystal lattice. Thus, the final step is to develop a stress-strain relation for bodies that are not isotropic, by way of introducing the study of crystallography.

1.4.3 Crystallography

Crystallography is the study of the structure and arrangement of atoms in crystalline materials. The crystal structure strongly influences the properties of a material. There are many examples of allotropes, materials of the same chemical composition but different crystal structure, including iron and carbon but it is most simply demonstrated by considering tin. Cooling pure tin to below room temperature (13.2°C) induces a transformation from the ductile β – Sn cubic phase (colloquially known as white tin) to the brittle α – Sn tetragonal phase (grey tin). When left unchecked this process, known as tin blight, will lead to the disintegration of tin to a powder. This change in symmetry of the underlying atomic structure dramatically changes the mechanical performance of the material.

Isotropic materials are those that have physical properties which are independent of direction. Therefore, describing the properties of such materials requires few parameters. For example, the Young's modulus and Poission's ratio would be sufficient to completely describe the linear elastic behaviour of an isotropic material. Young's modulus is readily measured in linear-elastic materials through tensile testing, where a component is loaded under tensile stress and pulled, the gradient of the measured stress-strain curve is then the Young's modulus. This is sufficient to describe the elastic response of an isotropic material, however, a significant proportion of engineering materials are anisotropic, that is to say their response is a function of the direction of examination. This necessitates a tool for finding the stress in any direction as a function of the strain in any direction and a given proportionality. In a general form we can express this as

$$\sigma_{ij} = \mathcal{C}_{ijkl} \,\epsilon_{kl} \tag{1.4}$$

where C_{ijkl} is a tensor representing the elastic constants, σ_{ij} is the stress in the directions i,j related to the strain, ϵ_{kl} , in the directions k,l. Since both stress and strain are second-order tensors, it follows that C_{ijkl} is a fourth order tensor, which consists of $3^4=81$ material elements if symmetry is not assumed. In practice, materials form in a regular lattice of atoms, of which there are seven possible systems, with a maximum of 21 unique constants in the triclinic system. From this relation we can see the anisotropy arises from the rotation of C_{ijkl} to different crystallographic directions, altering the values of the tensor elements.

To define the particular crystallographic plane of observation, the Miller index convention is adopted. From the lattice vectors $ha_1 + ka_2 + la_3$ the plane (hkl) and normal direction [hkl] are defined, normally written in terms of the lowest integers. This convention is shown for four planes in fig. 1.4. {hkl} is the family of planes equivalent to (hkl), similarly <hkl> is the family of directions equivalent to [hkl]. When Miller indices are used a final parameter determining the rotation on the plane (hkl) is required, this is usually denoted by ϕ_1 .



Figure 1.4: Definition of Miller indices, showing primary planes and demonstrating plane families.

1.4.4 Elastic constants

These elements of C_{ijkl} are commonly known as the elastic constants. For clarity, it is important to differentiate between the different forms of elastic constants and moduli which are discussed throughout the literature. The nomenclature proposed by Ledbetter et al. makes a sensible attempt to standardise the terminology and will be used in this work [28], as defined in fig. 1.5, although the reader should be aware this is by no-means consistent across the art. When, for example, the single crystal elastic constants and engineering moduli are discussed in close proximity they will be given their full names, in other cases the shorthand of elastic constants is used to refer to the single crystal elements.



Figure 1.5: Definition of elastic constant nomenclature, all of these terms are a subset of the general term elastic constants.

It would be difficult to overstate the importance of the elastic constants to understanding the behaviour of a material. Fundamentally, the elastic constants relate stress to strain [29], and provide a bridge between the micro- and macroscopic worlds. From an engineering perspective, knowledge of the elastic constants and orientation in anisotropic crystalline materials is vital for understanding the in-service mechanical performance and facilitates the calculation of engineering elastic parameters including: Voigt's modulus, Reuss's modulus, Hill's modulus, shear modulus, Young's modulus, bulk modulus, Poisson's ratio, and unveils other physical properties including strength, hardness, wear, plasticity and melting temperature [30]. For the materials scientist, measured elastic constants can communicate information on many important phenomena in solids including magnetic and electron-lattice interactions, phase-transitions and mode softening [31]. In short, elastic constants have an essential role to play in the calculation of many key physical quantities and in turn our wider understanding of materials on the whole.

The elastic constants, if available, are typically quoted for a material composition, temperature and pressure. In reality, the presence of trace impurities, atomic dislocations, residual stress, a magnetic field and chemical inhomogeneity all impact the true values of the specimen in question [28]. These effects, which can be difficult to quantify, lead to discrepancies in the published values for many materials [32]. The state-of-the-art in the measurement of elastic constants is reviewed in §2.4, including a tabulation of the values for materials pertinent to this thesis. Knowledge of the single crystal elastic constants also allows the engineering moduli to be determined. As an example of this, Young's modulus, E is found in a given direction $E = \frac{\sigma_1}{\epsilon_1}$. This can be solved by taking the inverse of the stiffness tensor, known as the compliance tensor, equation 1.5. Given these moduli are more 'accessible' it was once common to see materials described by their compliance tensor, although this has became less common with the advent of techniques to directly measure the single crystal elastic constants.

$$S_{ijkl} = C_{ijkl}^{-1}.$$
(1.5)

From equation 1.5, Young's modulus can be found as $E = 1/S_{11}$, such calculations are quickly computed and allow elastic constants to be related back to the physical properties used to describe isotropic materials above, when then combined with the rotation of C_{ijkl} these properties can then be solved in an arbitrary direction.

In cubic materials the degree of anisotropy can be described by the Zener ratio, η , such that

$$\eta = \frac{2C_{44}}{C_{11} - C_{12}} \tag{1.6}$$

Isotropic materials have a Zener ratio of $\eta = 1$, and it follows anisotropic materials have $\eta < 1$ or $\eta > 1$. Most engineering cubic materials, including all presented in this work, have $\eta > 1$. Potassium chloride (KCl) is an example of a cubic material with $\eta < 1$.

1.4.5 Microstructure

We can observe the presence of crystallography in many of the everyday objects that surround us, a simple example being the large differently coloured areas that appear on lamp-posts. This formation of discrete areas of granularity is known as the microstructure. Materials which exhibit these distinct regions are known as polycrystalline, in whom the presentation of this microstructure (size and orientation for example) is a cardinal factor governing performance.

The crystallographic orientation refers to the rotation of the crystal lattice, relative to a set of reference directions. The orientation of a microstructure can be visualised in a number of ways, of these inverse pole figures are the most intuitive. Inverse pole figures are based around the notion that the full range of crystallographic orientations, relative to the reference direction of the observer, can be uniquely represented on a colourscale. Therefore, a point in the unit triangle represents a plane (hkl) with normal aligned to the observers reference frame.

An example of an inverse pole figure for cubic nickel is given in fig. 1.6a, the large grain structure is clearly observed and the colour of the pixels describes the plane observed when looking down on the specimen, in reference to the key fig. 1.6b. this retains spatial information, however the overall distribution of orientations is not immediately obvious. The crystallographic texture refers to this notion of material having preferential orientations and is more readily viewed in a fig. 1.6c, where the intensity indicates the strength of a given orientation within the specimen.

Knowledge of a materials microstructure morphology, crystallographic orientation and the elastic constants are vital information for engineers and material scientists in the design, manufacture and on-going testing of materials. Indeed, a cursory glance at almost any in the field of materials science is likely to make reference to all three parameters. Unsurprisingly then, the development of techniques to probe these properties is distinct research field, falling under the broad heading of materials characterisation. Quite obviously, there is particular interest in techniques which can measure these properties with either limited or no damage to the specimen material.

1.5 Non-destructive evaluation and materials characterisation

1.5.1 The use of non-destructive evaluation

By definition, non-destructive evaluation (NDE) is the inspection of a specimen, that does not risk inducing damage, with the aim of detecting discrete defects and material properties. This is in contrast to 'destructive testing', whereby techniques such as tensile testing can be used to experimentally determine material parameters, through the deliberately induced failure of the component. The notion of assessing the integrity of an object, particularly where life is at risk, has long been an integral part of human life. There is of course a simple visual check but the idea of tap testing, where the a sharp clear sound is emitted from a solid component when tapped, compared to the dull response of a defective component was widely used in early manufacturing. Unfortunately, an awareness of the need for inspection has often been preceded by catastrophic engineering failures [33].

NDE is usually applied as a part of the manufacturing phase and offers the advantage of allowing direct measurements of the parts going into service rather than relying upon statistical models. The common example is the inspection of welds to ensure they are free from defects which have the potential to negatively impact performance. The term condition monitoring, rather than NDE, has traditionally been used to describe the ongoing assessment of integrity once a component is in-service. Today the terms are used somewhat interchangeably, the inspection of the pipework found in nuclear power plants is certainly an example of NDE of an asset in-service. Working 'in the field' to inspect in-service can introduce a number of challenges, such as limited access, high temperature and time restrictions [34].





(a) Spatial map of inverse pole figure, showing the crystallographic orientation of each grain within the specimen.

(c) Inverse pole figure of the specimen.

Figure 1.6: inverse pole figures are an intuitive method to visualise crystallographic orientation. In this figure the grains in a nickel specimen are coloured relative to the orientation key in (b), indicating the crystallographic plane which is observed looking down on to the specimen. Six specific grains (A-F) marked on both to allow the mapping to be followed. Rather then presenting the information in the orientation map (a), the density of orientations may be directly mapped to the inverse pole figure (c). This form of visualisation looses spatial information, but allows the texture to be inferred.

Most texts cite Wilhelm Conrad Röntgen's and his discovery of the x-rays in 1895, as the birth of the modern non-destructive evaluation industry, indeed his first paper cited flaw detection as a possible use. The onset of the second world war gave real impetus to the field, with the need to assure the integrity of submarines and airplanes, with many new techniques developed and refined. Today's NDE technician now has access to over one hundred variations of inspection, dependent upon the nature of the test. The majority of inspections are achieved by the so-called 'big-five' of ultrasound, magnetic particle, eddy current, dye-penetrant and radiography [35], contributing to a global market value exceeding \$5B [36]. The maturing of the industry has brought with it a high degree of standardisation and regulation governing its application [37].

Broadly, we may divide NDE techniques into two categories, passive and active. In a passive inspection the inherent emission energy or spectra of a component is observed to reveal its integrity. More commonly, in an active inspection some signal is injected into the inspected component, this signal then interacts with the material in such a way that by measuring the returned signal a characteristic of the component is revealed.

The purpose of subjecting a component to such a test falls within two fields, defect detection and materials characterisation. Defect detection is concerned with the detection and sizing of flaws within the components which will adversely affect the performance or lifespan of the component. This includes the cracking and porosity often found in additive manufacturing components, as introduced in §1.3.1. The field of materials characterisation covers a lot of ground, in the simplest case a pipe may be tested to ascertain its wall thickness (which may decrease due to corrosion). More complex evaluations are then concerned with evaluating the physical properties introduced in §1.4, for example measuring the components residual stress or the crystallographic orientation [38]. These more complex measurements are more akin to the results discussed in this work, however, we should note the field of materials characterisation is dominated by techniques that, while not strictly NDE, are important to introduce at this juncture. Hence, the following section will briefly introduce a number of pertinent inspection techniques across both fields, their use in additive manufacturing will be explored in the following chapter.

1.5.2 Industrial NDE techniques

Visual inspection

An often overlooked tool in NDE, visual inspection is simply the process of checking a component for clear and obvious defects. Flaws indicated by the inspection can either be used to inform further inspection via another technique or directly lead to the quarantining of the component. Inspections are often augmented with tools to improve access, such as the use of a borescope for internal pipe inspection, and sensitivity, for example using a magnifying glass. It is prudent to always make use of a visual inspection where possible, as this is generally the quickest method and has minimal cost, whilst being aware that only the most conspicuous of defects will be observed.

Acoustic Emission

Acoustic emission testing does not supply energy to the component under test, but instead measures the acoustic spectrum released by the component either during service or manufacture [39]. For example, stressed steel joints emit acoustic waves due to the spontaneous release of energy within the material. Much of the equipment used for acoustic emission monitoring is used in ultrasonic testing.

Optical Emission

Just as the acoustic spectrum emitted by a process can be studied to provide insight to an ongoing process, so can the optical spectra. Optical emission spectroscopy captures the wavelength of electromagnetic radiation emitted from a process which induces plasma formation, for example welding. The measured spectra can then be used to infer the elemental composition of the material, this can be used to assess the properties of the welding arc [40].

Dye penetrant

One of the most basic inspection techniques used in industry is the dye penetrant method. In this method following rigorous cleaning a liquid penetrant, normally in solvent suspension, is applied to the test surface. Following a suitable dwell time to allow the penetrant to flow into any cracks, the excess penetrant is removed and a developing agent applied. This developer then works to draw the remaining penetrant out of the surface cracks, giving clear indications of surface discontinuities in the test specimen.

Magnetic Particle

Similar to dye penetrant, magnetic particle inspection (MPI) is a well established technique for the inspection of surface integrity. When a magnetic field is induced into a component, either directly by a permanent magnet or by electric current, surface breaking defects shall cause magnetic flux leakage. In MPI ferromagnetic particles, either dry or in suspension, are then applied to the test surface, when magnetisation is induced these particles are drawn towards any areas of magnetic flux leakage, resulting in discontinuity indications. Whilst generally considered superior in comparison to dye penetrant inspection for surface breaking defects, MPI is limited to ferromagnetic materials.

One of the key points of consideration with both dye penetrant and magnetic particle inspection compared to visual inspection is operator fatigue. Whilst skilled inspectors will often be able to detect, without an augmentation, the same surface breaks that would be highlighted by dye penetrant or magnetic particle inspection in a one off scenario, MPI and dye penetrant are shown to maintain high probability of detection over repeated inspections [41]. In contrast, using a purely visual technique operator performance is seen to fluctuate over increasing time periods.

Radiography and computed tomography

Electromagnetic waves with wavelengths of 0.01 to 10 nm are known as X-rays, any electromagnetic radiation of a shorter wavelength is given the name gamma ray. By bombarding a test piece with directional high frequency electromagnetic radiation, from a source emitting either gamma or X-rays, it is possible to measure the density through a component by use of a radiation-sensitive detector. Radiography is commonly used in medical and industrial testing as it gives the ability to visualise internal features of solid opaque objects. Clearly one of the primary considerations in radiography is user safety, consider the harmful effects of radioactive exposure. Reference [42] includes a comprehensive discussion on safety in radiographic testing.

In X-ray computed tomography a process known as filtered back projection is used to take scans at varying orientation angles and reconstruct the 3D volume [43], the resolution of the reconstruction is measured by the voxel size, the three dimensional equivalent of the pixels captured in each two dimensional slice. This can be a slow process requiring several seconds of exposure for each image, and more than 720 images to form a 3D model [44], however the outstanding 3D data acquired means computed tomography is widely recognised as the current 'gold-standard' for volumetric imaging.

Eddy Current

Eddy current testing is based upon the principle of electromagnetism. Passing an alternating current through a coil results in a magnetic field generated around the coil, perpendicular to the coil. When placed on a test sample, an eddy current is induced in the sample in opposition to the coil current, which is then monitored by the inspection equipment. When the coil is placed in the presence of a defect, the phase and amplitude of the eddy current shall vary, indicating the presence of a defect. Eddy currents are advantageous in coated systems due to their ability to detect defects through surface coatings, and the speed of operation. However, the use of the eddy current technique is limited by the depth of penetration. The depth of penetration, can be assessed by equation 1.7, effectively the deeper into the test piece that is probed, the less sensitivity the inspection becomes.

$$\delta = \frac{1}{\sqrt{\pi f \mu_r \mu_0 \sigma}} \tag{1.7}$$

Fig. 1.7 visually displays the results from Equation 1.7, it is clear in order to increase the depth of penetration, a lower frequency, permeability or conductivity is required. Generally, conductivity and permeability are fixed properties of the test sample, meaning the only option is to reduce the inspection frequency.



Figure 1.7: Schematics explaining the relationship between penetration depth and inspection frequency in eddy current testing.

1.5.3 Materials characterisation

Optical approaches

Under bright field illumination incident light on a flat polished specimen will be reflected uniformly, without influence from the crystallographic orientation. However, there are two routes to observe the microstucture optically, either through novel illumination or preparation of the specimen surface. Considering the first route, by using techniques such as polarised light or differential interference contrast, the optical contrast can be improved. Whilst non-destructive this must usually be combined with surface preparation to ensure a smooth finish.

In chemical etching atoms are dissolved by attacking the surface with an oxidizing acid. As the rate of atomic dissolution is dependent upon the crystallographic orientation, the final etched specimen shows improved optical contrast between grains, allowing the microstructure to be imaged through normal optical methods to capture the full-field. This orientation-dependent dissolution is regularly exploited in silicon wafer manufacture. Chemical etching provides a relatively simple and quick method, although its use is accompanied by safety considerations particularly when working with hazardous enchants such as Hydrofluoric acid. Etching as described here cannot be used to recover the quantitative crystallographic orientations of the specimen.

Recent advances in the parallel technique electrolyte jet processing have a demonstrated the final etch surface displays a pitting formation characteristic of the crystalline orientation. Through processing the surface topology, with a complementary optical technique such as coherent scanning interferometry, the crystallographic orientation can then be calculated. Orientation determination is based around measurements of surface gradient and periodicity, so far this method has been shown to only determine orientations rounded to the closest principal direction in cubic material, with more advanced interpretation of the surface topography required to elucidate full orientation information [45].

Electron backscatter diffraction

Electron backscatter diffraction (EBSD) is a powerful technique for mapping the crystalline orientations of specimens, making it a now commonplace tool in the development of additive manufacturing and the most popular method of determining crystallographic orientation in material science more generally. EBSD is commonly used to produce datasets as introduced in §1.4.5.

EBSD is based around the interaction of a high power electron beam with a specimen. When the incident electrons collide with the specimen, a number of electrons will be inelastically scattered by the specimens atoms. Some of these electrons are at an angle to satisfy the Bragg equation

$$n\lambda = 2dsin\theta \tag{1.8}$$

where d is the interplanar distance of the lattice, θ is the Bragg scattering angle, λ is the wavelength of incident electrons and n is an integer representing the diffraction order.

γ

Electrons satisfying this equation are then ejected back from the specimen and form Kossel cones. Tilting the specimen to between 60° and 70° relative to the beam, allows more electrons to be diffracted and escape towards the detector. The image then observed at phosphor screen is then made up of Kikuchi bands. The position of these bands is then digitally extracted, usually by Hough transform. The interplanar spacing and angles that derive from the width of the measured bands are compared to simulated patterns to identify most probable crystallographic orientation.

Simply by scanning the electron beam across the specimen it is then possible to image the specimens surface orientation and microstructure. Along side the imaging capability, the popularity of EBSD also stems from its unparalleled accuracy and spatial resolution. In optimal conditions, current systems can resolve features down to 10 nm and with orientation determination error of $\pm 0.5^{\circ}$.

EBSD is conducted using a scanning electron microscope instrument, where the specimen is held under vacuum to prevent attenuation of the electrons in air. Limits on the size of vacuum chambers, combined with the acquisition rate of EBSD generally limit the size of area which can be interrogated. Typical diffraction pattern acquisition rates extend up to around 1 kHz, however faster detectors are available, as these see wider uptake the remarked figure will undoubtedly increase.

We should also note the two forms of image distortion which blight EBSD. Firstly, the specimen tilt angle causes trapezoidal distortion, where relative distance between sample points varies across the specimen surface to give an effective trapezoidal sampled area, rather than the intended rectangle. Secondly, as typical scans run for many hours drift in position of the specimen relative to the detector also becomes an issue, tending to stretch the data in one axis [46]. Correcting these distortion effects is possible but not trivial and most pertinently to this work, make different orientation imaging modalities challenging to compare.

X-ray diffraction

As with other electromagnetic waves, X-rays can be scattered, reflected, interfered. X-rays are well suited to use in diffraction studies, given inter-atomic spacings are typically a few Ångstroms (on the order of ~ 100 pm). By accelerating electrons into a metal target, typically tungsten, a collimated beam of X-rays can be generated. When incident on the surface of a crystalline specimen, this beam is scattered by the crystal lattice. By placing a detector the scattering pattern can then be observed. Intensity peaks present at points of constructive interference when Bragg's Law, equation 1.8, is satisfied. In order to ensure Bragg's equation is satisfied there are two approaches. Firstly, the specimen may be slowly rotated relative to the monochromatic X-ray beam, hence the θ term in the Bragg equation is varied, this is monochromatic method. Laue's method on the other hand uses a 'white' source, in essence allowing the λ parameter in the Bragg equation to vary until diffraction is observed. The imaged diffraction pattern is then processed using a Greninger chart.

X-ray diffraction is traditionally used for orientation determination of single crystals, and rarely applied to image the microstucture due to its poor resolution and relatively low rate of acquisition and processing. The beam diameter of a typical system is around 2 mm, this prevents grains smaller than this from being resolved. Currently, reduced beam diameters are inaccessible in standard laboratory machines as this results in insufficient flux through the focussing pinhole. Whilst the orientation can be determined with an accuracy of $\pm 1.5^{\circ}$. Typical surface preparation does not extend beyond ensuring the specimens surface is exposed, removing coating layers and containments. Rough surface specimens are readily measured without polishing.

1.6 Ultrasonics

An acoustic wave is generated by the vibration of a particle through gases, liquids or solids. Ultrasonics, the term for the field of acoustics concerned with devices generating waves of frequencies higher than the audible human range. The building blocks of modern acoustics can trace it's



Figure 1.8: Schematics describing the operation and working principle of EBSD analysis, as used to obtain the orientation information of the material under test.

lineage back to the previous millennia, when the Greek scholars first studied the harmonics of musical instruments. The discovery of the piezoelectric effect in the late 1800s was a paradigm shift, paving the way for the first transducers and giving ultrasound a tangible industrial use. Rapidly, applications such as sonar and flaw detection gained traction. Modern-day ultrasonic devices can refer to anything in the range of 20 kHz to several GHz [47].

Ultrasonic testing for NDE has now been used in industry for over 70 years. Traditionally, short ultrasonic pulse waves are sent into a material and the back wall echo signal is used to obtain information about the integrity of the part under test. The established method of ultrasonic test would typically use a piezoelectric transducer in contact with the component to create an elastic wave propagating through the depth of the component. A couplant gel is usually required to maintain consistent contact between the transducer and the component. This elastic wave can either be detected on the far side by a second transducer in a pitch-catch configuration, or the wave can be detected by the original driving transducer after reflection from the back wall in a pulse-echo configuration. Both approaches and the corresponding signals from a simple flaw are shown in fig. 1.9. The sign of a flaw depends on the inspection configuration used, for example, in a pulse-echo system it is the unexpected middled echo that indicates the presence of the flaw, caused by the reflection the acoustic wave at the interface with the flaw. Conversely, in the pitch-catch



configuration it is the absence of signal that points towards a defective object.

Figure 1.9: Basic outline of an ultrasonic inspection for flaw detection in NDE. The schematic shown mimics both through-transmission and pulse-echo set-ups by plotting the time-domain signals that would be received at both elements. In the pulse-echo configuration, the acoustic reflection between the wall echoes indicates the presence of a defect. In the through transmission configuration in the initial position (i) transducer (B) observed a large amplitude acoustic signal. As the transducer is moved to position (ii) the amplitude of this signal falls, with the energy being reflected back towards (A) by the defect.

The signals presented in fig. 1.9 plot the amplitude of the time-varying signal as measured at the detector, this intuitive visualisation is known as an a-scan. Fig 1.10 demonstrates b-scan and c-scan representations of an ultrasonic dataset. In the b-scans, time is plotted against position with the intensity of the colourscale indicating signal amplitude. There is no strict definition of c-scan, merely that it maps in x and y some acoustic parameter. In fig 1.10, the maximum amplitude at each position is indicated; however, from the same dataset parameters such as arrival time or phase of a given frequency component could equally be viewed as c-scans. Throughout this work c-scans, particularly those showing the acoustic velocity measured across a specimen, and to a lesser extent a-scans will be presented.

In materials where the ultrasonic velocity, v, is known a range of further information can be deduced. For example, the wall thickness of a pipe, d, regularly measured where corrosion is a concern, can be easily be determined by $d = \frac{1}{2}vt$, where t is the arrival time of the back wall echo, and v is the acoustic velocity of the material - an approach know as time-of-flight. The discussion thus far has been concerned with the interaction of the acoustic wave with a discontinuity, specimen boundary or flaw, however as the acoustic velocity is primarily dependent on a materials density and elastic contacts, the time-of-flight also provides a insight to the state of all the material



Figure 1.10: Examples of b- and c-scans in ultrasound, the b-scans show amplitude as a function of time and position, whilst the c-scan on the top surface shows maximum amplitude as a function of position. There are many alternative c-scans that could be plotted, for example, the amplitude at a specific time index.

through which it propagates. Intuitively, where the dimensions of a specimens are well measured, the time-of-flight can then reveal the materials properties such as porosity and elastic moduli.

Bulk waves

In the simplest case of an isotropic solid two wave modes can propagate through the medium the longitudinal and transverse modes. It is sensible to first introduce longitudinal waves as these include the sound travelling through air, the acoustic wave most people will be familiar with. These longitudinal waves induce particle vibration in the direction of wave propagation, producing sinusoidally varying areas of high and low particle density. In an isotropic medium the velocity of the longitudinal wave can be found by

$$v_l = \sqrt{\frac{E(1-\nu)}{\rho(1+\nu)(1-2\nu)}},$$
(1.9)

where E is Young's modulus, ρ is density and ν is Poission's ratio. This relation may be expressed in terms of Lamé's parameters. Longitudinal velocities in engineering metals range between 5600 to $6500 \,\mathrm{ms}^{-1}$.

As with longitudinal waves, the velocity of the transverse wave can be expressed in terms of material parameters:

$$v_t = \sqrt{\frac{E}{2\rho(1+\nu)}} \tag{1.10}$$

Transverse velocities in engineering metals range between 3000 to $3300 \,\mathrm{ms}^{-1}$.

Lamb waves

An elastic wave propagating in a solid layer which is free to displace both in the direction of propagation and perpendicular to the surface, is called a Lamb wave. The existence of such waves was reported by Horace Lamb in 1917 [48], Viktorov's later detailed reporting on Lamb waves remains an essential text [49]. An important property of Lamb waves is dispersion, the frequency and thickness of the plate determine the acoustic velocity. This raises the question of the definition of a plate; fig. 1.11 shows the dispersion curves of aluminium, we see as the product of thickness, d, and inspection frequency, f, increases the velocity tends towards the Rayleigh wave velocity. This is showing as the frequency of inspection increases the waves are less influenced by the far side of the plate and become waves propagating along a half-space. There are two principal groupings of Lamb waves modes, symmetric or antisymmetric, depending on the symmetry of the mode shape about the centre-line of the plate - the lower half of fig. 1.11 shows the mode shape of the two fundamental Lamb modes.

As Lamb waves travel through the entire thickness of the material, interrogation of both surfaces and the bulk of the material is possible. Imagine a conventional inspection of a plate, as introduced in fig. 1.9, the probe would need to be scanned over the entire surface. A Lamb wave inspection allows the length of the plate to be probed from a single generation location [50].

A further family of surface waves exist, the Love wave, named after the A. E. H. Love who predicted their existence in 1911 [51]. Love waves propagate along the surface of a solid with transverse particle motion in the plane of the surface. These waves are not easy to generate by laser ultrasound, and will not be discussed in any further detail but more information can be found in [52].

Rayleigh waves

Rayleigh waves are a special case of guided waves which propagate along the free surface of semiinfinite elastic half-space. In practice, as Rayleigh waves are surface bound, the medium can be considered semi-infinite when $d >> \lambda$. The amplitude of a Rayleigh wave decays exponentially away from the free surface, as shown in fig 1.12b.

Having first been hypothesised by the eponymous Lord Rayleigh in the 1880's, we now understand the Rayleigh wave to be one of the destructive components in earthquakes. Rayleigh waves



Figure 1.11: Dispersion curves in aluminium. Each curve in this figure represents a separate wave mode, these can be grouped as either symmetrical or antisymmetrical, based on the displacement symmetry about the x-axis. To illustrate this, the mode shapes of fundamental symmetrical (s_0) and antisymmetrical (a_0) modes are shown in the lower half of the figure.

first gained traction as an NDT tool in the 1950's with researchers exploiting Rayleigh waves for material inspection.

Viktorov introduced a useful empirical approximation for the Rayleigh wave velocity in isotropic media, equation 1.11. It is important to note this equation suggests the Rayleigh wave velocity is a fraction of the transverse wave velocity, where this fraction is a function of only Possion's ratio, ν . Given a typical Poisson's ratio of 0.33, the Rayleigh wave velocity would be expected to fall in the range of 2700 to $3100 \,\mathrm{ms}^{-1}$. It can be readily seen that as neither equation 1.11 or c_t contain a frequency term, the Rayleigh wave is non-dispersive.

$$c_r = c_t \frac{0.862 + 1.14\nu}{1 + \nu} \tag{1.11}$$

Beyond Rayleigh waves it is also necessary to consider the presence of so-called pseudo surface acoustic waves (PSAWs). Unlike true surface waves these PSAWs attenuate along the propagation length as they lose energy into the bulk medium. However, this attenuation effect is usually small allowing PSAWs to be measured along with RSAWs over short propagation distances.



Figure 1.12: Schematic of Rayleigh surface acoustic wave propagation and the associated displacement components.

1.6.1 Resonance

Solid objects have natural frequencies at which they wish to vibrate when mechanically excited. These natural frequencies are a property of the specimens shape, size, density and - most pertinently - elasticity. Resonant ultrasound spectroscopy takes advantage of this relationship to determine the specimens elastic constants, the most commonplace approach to measuring the elastic constants of a material currently [53]. Whilst not strictly concerned with wave propagation, the measurement of resonance primarily occurs at ultrasonic frequencies, with a typical range of 5 kHz to 5 MHz, allowing similar equipment that used in ultrasonic inspections to be deployed - hence its place in this section.

1.7 Laser ultrasonics

Physicist Theodore Maiman is widely credited with the invention of the laser in 1960. The applications and potential impact of lasers has not always been clear, indeed Maiman once famously commented the laser was "a solution looking for a problem". However, the unique properties of the laser, namely the ability to produce high intensity well collimated beams and of coherent light, at well-defined narrowband wavelengths, have made them eminently useful for a range of applications from welding to range-finding to long-distance communication infrastructure. Modern lasers offer high stability, pulses down to the attosecond range (1×10^{-18} seconds) and powers from nanowatts to petawatts [54, 55]. These advances have launched several distinct research fields, spanning manufacture, clinical and industrial diagnostics through to fundamental physics, in addition to their ubiquity in everyday consumer devices.
The field of laser ultrasonics introduced in this section refers to the process of using a laser to generate and/or detect the propagation of acoustic waves within a specimen [56]. The advent of this field can be traced back to preliminary quantitative measurements in the early 1980s [57]. The realisation of practical laser ultrasound systems has allowed the benefits of both ultrasonic - probing the material properties through the volume with point-like generation - and optical techniques - non-contact and non-destructive - to be realised in harmony.

1.7.1 Types of generation source

The generation of acoustic waves caused by the absorption of electromagnetic radiation in a material is known as the photoacoustic effect. In order to generate ultrasound in a material the incident irradiance must be temporally varying. This can be achieved through modulation of the source or using short pulses. There are numerous methods available to generate and control such pulses. Consider first the simplest method to achieve this, modulating the beam of a continuouswave source, this can be done by using an optical chopper [58] or applying pseudo-random noise modulation to the laser bias current [59]. These approaches are relatively easy to implement and can allow damage reduction in fragile specimens as the peak optical power is much smaller than of pulsed lasers, however the generation efficiency is small and requires sensitive lock-in detection schemes to be used.

More commonly a Q-switched pulsed laser will be used to achieve the generation of ultrasonic waves. The term Q-switch stems from the process of varying the inter-cavity losses (and hence the Q factor) of the laser resonator and can be done passively or actively. In passive Q-switching a saturatable absorber is located in the laser resonator, the absorption coefficient of this absorber decreases with increasing intensity of incident light. Once the saturation level has been reached the absorber is effectively transparent, allowing the pulse to be formed and emitted. Again, this approach is relatively simple and cost-effective to implement as it requires no opto-mechanical or electronic control components, however the build up of gain is limited and a large variance in the time between pulses (known as laser jitter) can be problematic.

Active Q-switching, the most commonly used approach to generated pulses, typically uses an acousto-optic modulator, known as a Q-switch, to control the loss modulation within the resonator. When inactive the Q-switch prevents reflection of light back to the grain medium by diffracting the first order beam out of the cavity. This prevents feedback and, as pumping continues, the energy density approaches a maximum, set by spontaneous emission and other loses. The Q-switch is then activated to allow feedback to the gain material and a high intensity laser pulse is produced. This approach allows for low jitter, short, high energy pulses. The duty cycle on such sources

is often relatively low ($\ll 1\%$) and therefore average powers are often deceptively small, however peak powers can exceed a megawatt.

1.7.2 Generation mechanisms

When applying laser irradiance to a material surface there are two possible regimes of laser generated ultrasound. The optical power along with pulse width and specimen absorptivity determine the mechanism exploited.



(a) ablative regime.

(b) thermoelastic regime.

Figure 1.13: Schematic of the two generation mechanisms (a) ablative and (b) thermoelastic for the excitation of ultrasound by the optical energy indent on a surface from a laser.

Ablation

In the ablation regime (fig. 1.13a), high electric field generated due to laser light is enough to remove electrons from the bulk sample. The generated free electron collides with the atoms of the bulk sample, in which transfer of energy occurs. This leads to the heating of the surface, which is followed by vaporization of the surface material. The incident energy is dissipated as kinetic energy through particle ejection, creating a responsive stress primarily normal to the surface. Therefore, applications that wish to utilise ultrasonic bulk waves are suited to ablative generation.

Depending on strictness of the definition ablation can be considered either destructive or nondestructive, as the material removed is small and localised to the surface. However, practically this can create an optically rough surface and reduce detection efficiency, as discussed later.

Thermoelastic

The second generation regime of thermoelastic excitation occurs at a lower laser irradiance and is strictly non-destructive (fig. 1.13b). The incident laser pulse causes a sudden rise in thermal energy at the specimen, leading to a rise of temperature at the surface in turn causing localised thermal expansion. Principal stresses are then directed parallel to the free surface, creating surface waves, with a small component directed into the specimen. Therefore, the thermoelastic regime is quite suitable when working with surface acoustic waves. Alternatively, the boundary conditions can be varied to more efficiently generate, for example immersion of the specimen, or the optical power can be increased as the amplitude scales linearly with incident energy, up to the ablation limit.

1.7.3 Optical detection

Having generated an acoustic wave in a specimen, the challenge of measuring this wave now presents itself. This section provides a brief overview of approaches to optical detection, given the numerous exotic detection methods available, only techniques applicable to the detection of SAWs will be discussed here. [60] provides an excellent overview of this field. Whilst not a recent publication [61] also covers the topic in detail, with detection schemes considered from an industrial perspective, where relevant.

Knife-edge detection

Knife-edge (also known as optical beam deflection) style detection is regularly implemented as it is relatively simple and requires little specialist equipment whilst providing high sensitivity. A continuous laser beam incident on the specimen surface is focussed to a spot in the propagation path of the SAW. The propagation of the SAW causes a perturbation of the surface inducing a deflection in the reflected beam, proportional to the local surface gradient. By placing a 'knifeedge' in the path of the beam this displacement can be sensed as a change in intensity by a suitable photodetector. This principal is shown schematically in fig. 1.14. The name knife-edge is derived from the traditional implementation where a knife-like edge was used to block half of the deflected beam. Instead, by balancing the reflected beam between two photodiodes and measuring the difference between the two, the same measurement can be made whilst utilising all the reflected light and suppressing common noise.

When working with optically smooth surfaces knife-edge detection has very few drawbacks,



Figure 1.14: Schematic of knife-edge style detector, with split photo-diodes and requisite optical components. As the SAW propagates under the probe beam the majority of the returned beam fluctuates between photodiode A and B, therefore by taking the difference between the channels the acoustic signal can be recovered.

offering excellent sensitivity, low inherent noise and adapt quickly to surface changes, allowing single shot measurements to be realised in many cases. However, the performance of knife-edge detectors falls off rapidly as the surface roughness increases as the reflection becomes diffuse.

Optical surface roughness

The surface of most unprepared industrial specimens, including those manufactured additively, are optically rough, scattering incident light in many directions. The surface roughness creates a distorted wavefront reflection, with the interference of many components of difference phase and amplitude - the result is a speckle pattern.

It is important to think of polished and rough surfaces as two separate applications and consider that a different detection scheme may be appropriate for each. In the case of applications on polished surfaces the simplicity, wide detection bandwidth and insensitivity to background vibrations make the knife-edge style detector a sensible choice. However, in a rough specimen the speckle pattern means the image formed at the detector is no longer Gaussian, containing many bright and dark spots, the sum of which roughly cancels out. Therefore, a different detection scheme is essential when considering measurement on a rough surface - this section concludes with the introduction of suitable techniques for the optical detection of ultrasound.

Interferometry

All of the following techniques rely on the interference of a distorted wavefront reflected from the specimen with a reference wave. The Michaelson interferometer is the simplest implementation of this, fig. 1.15. The incident beam first passed through a beam splitter, sending the beam in two orthogonal directions, one directed to the specimen under measurement and one to a reference mirror. Again, the presence of SAWs perturb the specimen surface and therefore periodically vary the path length. Both paths are then recombined and projected to a photo-detector, where the interference allows the amplitude and phase of surface motion to be deduced. The approach is sensitive to surface motion down to the order sub-nanometre and is less sensitive to surface finish than the knife-edge approach.



Figure 1.15: Schematic of a Michaelson interferometer.

Two wave mixing

In two wave mixing the probe beam is again split into two paths by a beam splitter, with one leg directed to the specimen surface and one reference leg [62]. The reflected beam from the specimen is directed into a photo-refractive crystal, where it is combined with the reference leg. This gives rise to an interference pattern of bright and dark fringes, allowing the photo-refractive nature of the crystal to be exploited. Through a diffusion of electrons and the resulting electric field causes a spatially varying refractive index grating to be stored in the crystal. This then imposes the same pattern on the reference beam, giving two spatially identical beams, this allows demodulation of a transient phase shift encrypted in a speckled wave, such as that reflected from an optically rough

surface. This principle is shown schematically in fig. 1.16a.

There are two primary drawbacks of the two wave mixing approach. Firstly, the cost, commercial systems typically cost upwards of £50,000. Secondly, the speed of adaptation, the photorefractive effect takes time to set-up the refractive grating pattern and diffract the reference beam, which is sensitive to continued variations in the specimen beam, rendering this approach unsuitable in applications with scanning (as opposed to point measurements).



Figure 1.16: Schematics of common interferometry techniques for signal detection in laser ultrasonics, particularly when working with rough surfaces.

Confocal Fabry-Pérot Interferometry

In velocity interferometry a distorted beam is interfered with a time-delayed version of itself, exploiting the Doppler shift of the frequency of the reflected light from the interrogated specimen surface. The propagation of SAWs gives rise to a velocity component in the perturbation of the surface, this induces subtle variations in the frequency, and therefore wavelength, of the reflected light. One of the most commonly used devices to measure this effect is a Confocal Fabry-Pérot interferometry (CFPI). In a CPFI, two curved highly reflective mirrors ware placed facing each other, the incident beam is reflected by these mirrors many times, causing constructive interference of the reflected beams inside the cavity. The intensity of the transmission, and hence reflectance, of the cavity varies significantly with this small change in wavelength of the light due to SAW perturbation. This principle, is illustrated in fig. 1.16b.

As the as the reference beam for interference is a time delayed version of the incident beam itself, CPFI is inherently insensitive to speckle and distortion from vibrations, making CFPI highly suitable when working with optically rough surfaces or in industrial environmental - although, despite stabilised versions being available, performance degrades in environment where the surface changes significantly. That is to say, CFPI is suitable for rough surfaces, these surfaces must be consistent, else the system needs constant recalibration. Clearly, extremely stable cavity lengths are requisite to making meaningful measurements, a non-trivial challenge concomitant with a the significant cost of suitable commercial systems (> \pounds 80,000). Additionally, the detection bandwidth of CPFI is relatively limited compared to other techniques, which may be viewed as a significant drawback when considering the broadband generation possible in laser ultrasonics.

1.8 Surface wave measurements

The ability to measure the amplitude, phase and velocity of surface acoustic waves is an important tool in NDE. Notable uses include detecting surface breaking cracking [63], characterising damage induced by surface finishing [64], fatigue [65] and imaging material microstructure [66]. For the purposes of this work it is the interaction between the Rayleigh, specifically its velocity, and the microstructure of the probed sample which is of interest. The measurement of surface acoustic waves is a promising approach for the measurement of the single crystal elastic constants and orientation in polycrystalline specimens thanks to the ability to localise the acoustic signal to individual grains.

Traditionally, surface wave measurements would be made by an acoustic microscope, similar to the schematic shown in fig 1.17. Using a acoustic lens the plane wave of piezoelectric transducer can be focussed at a specimen. The tone burst from the transducer creates both bulk waves, along the centre axis, and 'leaky' SAWs, when the ray satisfies the Rayleigh wave critical angle, θ_{CR} . The recorded signal is the result of interference between these two wave modes. As the Rayleigh wave changes with the microstructure, the phase and amplitude varies, allowing qualitative microstructure imaging.

If the lens is kept over a single point on the specimen and moved towards the specimen from the focal position, the signal can undergo a series of oscillations, due to the phase difference between the rays. This behaviour is visualised as a so-called V(z) curve, where the period of oscillation can be extracted. The Rayleigh velocity, v_R , can then be found by equation 1.12, allowing a quantitative measurement [67].



Figure 1.17: Scanning acoustic microscope in V(z) configuration, the acoustic source map be focussed to a spot or a line for line-focus acoustic microscopy.

$$v_R = v_0 \left[1 - \left(1 - \frac{v_0}{2f\Delta z} \right)^2 \right]^{-\frac{1}{2}}$$
(1.12)

Where v_0 is the velocity in the coupling medium, f is the inspection frequency and Δz , is the periodicity of the V(z) curve. This approach it is rarely used for imaging, as it is too time consuming to repeat the defocus at each x - y position.

The propagation of acoustic waves in polycrystalline anisotropic materials is complex. To shed light on the problem, consider a plane surface wave propagating away from a generation source across multiple grains, which are small enough such that different parts of the wave experiences different grains. As each grain has a distinct elastic response, the wave experiences a change in amplitude and phase, given the random nature of neighbouring grain orientation, the propagated wavefront contains random bright and dark 'speckles' - these effects are termed 'acoustic aberrations' [68, 69].

This is further complicated by the impedance mismatch between grains, which causes the refraction of the forward-propagating wave and the reflection of a portion of the acoustic energy, creating a backward propagating field. Of course, this backward propagating wave can encounter further grain boundaries and be reflected again, and so forth. As the grain size becomes small or the propagation distance large, this can happen many times. The effect is an acoustic field with a spread of speckles. The salient point being these acoustic aberrations and scattering effects make reliable measurement of the SAW velocity in polycrystalline materials difficult by direct

time-of-flight or phase measurements [70].

Spatially resolved acoustic spectroscopy (SRAS) is an important surface acoustic wave measurement technique, developed at The University of Nottingham, which uses laser ultrasound to both generate and detect the surface acoustic wave. By structuring the generation beam such that a grating pattern is imaged on to the surface of the interrogated sample, narrowband generation of a non-dispersive Rayleigh wave is achieved. The presence of the acoustic wave may then be measured by any of the techniques discussed in §1.7.3. Fig. 1.18 shows a schematic of a simplified SRAS instrument and the generation of surface acoustic waves in polycrystalline media.

By sampling either the k-vector or frequency domain, the SAW velocity as a function the properties under the generation patch can be determined. This makes SRAS very tolerant to acoustic aberrations and the result being single-shot measurements are possible, making the technique exceptionally fast [71].



Figure 1.18: Schematic of (a) SRAS experimental system and (b) surface wave generation and measurement in a polycrystalline material. It is important to note the generation and detection of the acoustic wave across a grain boundary, is the measured velocity in this situation is only a function of the properties under the generation patch and not affected by the grain boundary crossing, making SRAS an attractive tool for the imaging of polycrystalline microstructures.

1.9 Summary

A great deal of attention has been paid to conveying the motivation of this work in this chapter, highlighting the demand for a characterisation tool for additive manufacturing. Surface acoustic wave imaging is a promising approach for this, scanning acoustic microscopy and SRAS have been briefly introduced, effectively representing the past and the present of the field, respectively.

The dominant tools of NDE and materials characterisation have been introduced, with their

well-covered relative merits noted. Ultrasonics in particular plays a vital role in NDE and materials characterisation now, with the availability of scientific lasers only amplifying this. This leads in to the following chapter which describes the state of the art in materials characterisation and evaluation within additive manufacturing and the measurement of elastic constants, paying particular attention to acoustic and ultrasonic approaches.

Chapter 2

Survey of existing literature and methods

2.1 Introduction

The main objective of this thesis is to understand the materials characterisation capability of spatially resolved acoustic spectroscopy (SRAS) measurements in additive manufacturing specimens, including work to measure the single-crystal elastic constants in polycrystalline materials. As such, to give context to the work at hand, this chapter will provide a brief review of the existing experimental methods to characterise additive manufacturing components or measure the singlecrystal elastic properties. As these are currently disparate fields, both require a suitable review and discussion. As with the previous chapter, this chapter aims to set the scene for the results presented later in thesis by drawing together examples of the current state-of-the-art, and does not yet present the novel contribution of this thesis.

The microstructure formation observed in additive manufacturing is quite unlike traditional processing routes; as such, the first part of this chapter is given over to the materials and additive processes used in this thesis. This includes a presentation of typical defects and microstructure formations - essentially the features which need to be assessed - this leads into a brief review of current AM inspection technology.

A review of the state-of-the-art in the measurement of elastic constants follow, with a focus primarily given to capable acoustic techniques. For the sake of brevity, tabulated elastic constants, as used in this thesis, are presented in appendices A and B.

figures presented in this chapter that are not the author's own work are reproduced with publisher permission.

The chapter concludes by discussing the state of the art in SRAS §2.5, in relation to additive manufacturing and materials characterisation more generally, bringing together the two main themes of this chapter. Given the reliance upon the forward model, this section also reviews procedures for calculating the surface acoustic wave (SAW) velocity.

2.2 Additive manufacturing techniques and properties

As described in the previous chapter, the results presented in this thesis focus on two distinct AM techniques, and the specifics of both are detailed in the following section. Where necessary, detail is also provided on interesting features, such as the defects and microstructures typically generated from these two techniques.

2.2.1 Laser powder bed fusion

Process

Laser-powder bed fusion (L-PBF) is the use of a high power laser to irradiate a bed of metallic powder feedstock. Compared to other AM techniques, L-PBF is useful for creating complex geometry components with intricate and internal features, although the build volume is relatively limited, with low deposition rates [72]. A comprehensive guide to L-PBF can be found in [73].

Fig. 2.1 shows a schematic of a typical powder bed fusion build system. The build area operates within a vacuum chamber backfilled with argon to prevent oxidization. This atmosphere is transparent to the near-infrared wavelength of the build laser, typically a ytterbium laser of 1070 nm wavelength. A thin layer of powder (10 - 100 μ m) is deposited and spread evenly across the bed. A high-power laser source is then focussed on to the powder bed, irradiating the virgin feedstock, inducing melting of the powder, which then solidifies, allowing a part to be formed. This process repeats (often thousands of times in a single build), building up the final specimen layer-by-layer.

The laser-feedstock interaction is defined by a number of processing parameters, chiefly the laser power, scan velocity, hatch spacing (the separation of neighbouring scan paths), and scanning strategy - essentially the order in which the powder bed is irradiated. The build laser is usually focussed to an extremely small spot ($<100 \,\mu\text{m}$ in diameter), allowing complex features, such as internal cooling passages and lattice structures, to be created with unparalleled accuracy. This small spot size also has the benefit of producing components with a superior surface finish, compared to other additive techniques. To be specific, the surface roughness of L-PBF components are lower than any other powder based technique [74]. Wire-based techniques, to be introduced in §2.2.2 may have a smoother feel but have a high roughness due to low frequency periodicity



Figure 2.1: Schematic of laser powder bed fusion system.

from the wire layers; therefore, L-PBF components require less surface finishing and most closely resemble the designer's vision.

Typical features and industrial considerations

As L-PBF is more expensive than conventional manufacturing routes, it is viewed primarily as a candidate for high-value industries, particularly where either a high degree of customisation is required or to create parts of improved performance through features that are hard or impossible to create with existing fabrication techniques - internal structures are an obvious example of this. Customisation makes additive manufacturing attractive in medical industries, whilst design flexibility in high performance components are attractive for aerospace and nuclear sectors. Examples of reported uses include, skeletal implants, fuel nozzles and lightweight brackets [75, 76].

The feedstock material found in L-PBF is often a complex alloy. These materials tend to have relatively wide temperature ranges over which liquid and solid phases can coexist and often solidify as columnar grains under the conditions imposed by AM. The resulting material can therefore contain both a columnar microstructure and numerous cracks. Due to the high cooling rates, a finer microstructure can be observed in L-PFB than in conventional manufacturing methods. The rate of solidification seen in L-PBF is an order of magnitude greater than that of traditional casting approaches promoting epitaxial growth, which in turn results in columnar grains that grow in the <001> directions that are normal to the deposited powder layers. Fig. 2.2 provides examples of the typical features found in L-PBF produced specimens. This microstructural alignment results



in the strongly anisotropic performance of as-built components.

Figure 2.2: Key features of laser powder bed fusion specimens.

Of course, to be of use to these high-value industries, L-PBF must deliver components with exceptional mechanical performance - unfortunately, this is currently only partially true. Whilst near defect-free parts can exhibit mechanical properties that are superior to traditional wrought equivalents, there exists an intolerable variation in performance. In a detailed review, Sanchez et al. reported a variation of 85% and 100% in the respective tensile and yield strengths of AM nickel components from contemporary literature. The authors conclude that whilst some of this unpredictability in performance can be attributed to the propensity of L-PBF to form defects, the primary driver is the variability in the as-built microstructure [77]. Pertinently, the authors conclude by lamenting the lack of testing methods suitable for L-PBF; traditional testing method-ologies assume components are monolithic and are therefore ill-suited to capturing the heterogeneous properties of L-PBF components. Thus, the authors state that new tools for mapping the wide-area response of L-PBF components are required.

It is essential to convey that in many cases that this variable does not stem from 'incorrect' builds but due to the number of processing parameters available to the user. Van Elsen states over fifty parameters are available to the user [78]. A simple example of this effect is build orientation; keeping all other parameters fixed, horizontally orientated parts exhibit superior tensile performance compared to vertically orientated equivalents.

The challenge in implementing this approach is that the linkages between geometry, alloy chemistry, process, and microstructure are convoluted, and can lead to spatially heterogeneous microstructure. For example, coarse grain regions can form at melt pool centres, with fine grain regions at melt pool boundaries, as demonstrated in fig. 2.3a [79]. Similarly, variations in geometry cause unexpected microstructure and defect formation because the current generation of build systems to do not vary scan paths to account for differences in heat distribution as a result of changing part geometry [80]. Undesirable microstructures can also arise from errors in the printing-slicing process [81], reflecting the 'black box' nature of the commercial systems.

Whilst it is early days, there is clearly a recognition that the problem of process-microstructure influence can be turned on its head, with a view to creating 'designer' microstructures with L-PBF. An elementary example, as demonstrated by Niendorf et al., shows high power exposures (1000 W) will create typical columnar large grains structures, whereas reduced power exposures (400 W) create a finer equiaxed structure, allowing functional grading of the component [82]. At the vanguard of this field, Lei et al. have demonstrated the ability to generate a wide range of distinct grain structures, including 'zigzagged' interlocking grains [83].

Beyond controlling the grain size and shape, there has been interest in controlling the crystallographic orientation of L-PBF specimens. The earliest example of this is the work of Dehoff et al., who used a variety of scan strategies to create the letters DOE at a different orientation to the bulk of the sample [84]. Whilst less visually interesting, Sun et al. have also demonstrated the control of the microstructure to generate single-crystal structures [85], fig. 2.3c. Most recently, this area has been advanced by Sofinowski et al., who used a library of scan strategies to create a QR code within the microstructure of the component [86], fig. 2.3d, using the three primary orientations (<001>, <101> and <111>) to create a trinary colour coding system. A particularly illustrative example of this concept of controlling the orientations to create images conveying intelligible information in the inverse pole figure, is the work of Plotkowski et al., who recreated the famous Mona Lisa painting in the microstructure of a nickel additive component created by powder bed fusion [87], fig. 2.3e. Whilst this is of limited practical use, the purpose is to demonstrate that if artwork can be recreated within the microstructure of powder bed components, achieving functional grading is clearly within our reach. It should be noted these orientation control strategies have only been demonstrated in small, simple geometry specimens of single-materials.

Returning to the example of the QR code embedded in the microstructure, whilst in this case, the QR code merely directs the reader to the website of the research group, embedded data for security and material traceability has been an area of significant research. Whilst traditional manufacture is not immune to counterfeiting or tampering [88, 89], the digitisation of manufacture has created significant security concerns. Unlike, traditional subtractive manufacturing, where milling and turning machines can define only the object's shape, adjusting AM process parameters can influence the material's micro-structure, thus affecting the part's physical properties, remembering the expectation these components will be used almost exclusively in sectors such as aerospace and nuclear [90, 91, 92]. Previously researchers have proposed using the stochastic nature of microstructures as a fingerprint [93], with L-PBF information can be embedded into the microstructure [94]. Critically, embedding this information into the microstructure has the advantage that it can only be separated or altered by damaging the part. Sola et al. noted that for these traceability mechanisms to be realised, then capable wide-area mapping techniques are essential and not presently readily available [95] - this being the salient point for the work at hand.

Powder bed manufacture has also had a profound effect on material design and creation through the ability to combine multiple feedstocks allows flexible in-situ alloying [96]. This combining process may take the form of a simple mixture [97] or more complex satelliting [13]. The components formed from these feedstocks may be monolithic, or the composition may be spatially varying (in two or three dimensions), known as functional grading [98]. Alloying of the feedstock may be used to modify the elastic properties of the component [99]. In-situ alloying provides a route to quickly generate new materials, noting that the single-crystal elastic properties of these materials will be unknown.

Current-class L-PBF build systems can produce high-density near defect-free components; however, a relatively narrow processing window provides this, and there is no established systematic route for determining these parameters [100]. This leads to components with volatile mechanical properties, making them unsuitable for use in high-integrity applications, necessitating the provision of many trial-and-error builds to refine the processing parameters, particularly as geometries and material feedstocks become more complex [101]. Each of these iterations requires a full treatment of characterisation to inform the subsequent trial. This directly contradicts the promised benefit of improved material usage and agile manufacturing, particularly when developing components with small-production runs, common in the aerospace industry [102].

The forms of defect observed in L-PBF are summarised in table 2.1. An interesting point to note from table 2.1 is the wide range of possible defect dimensions. This shall necessitate a sensing system that can operate over a wide resolution range whilst maintaining a high probability of detection. One of the most common defects to occur in powder bed processes is porosity, which act as stress concentration sites to the detriment of fatigue [103] and elongation [104] performance in L-PBF components [105]. Furthermore, porosity negates the mechanical benefit offered by the finer lamellar microstructure characteristic of L-PBF by reducing the yield strength [104]. It has been shown that these voids can present in a range of sizes, 5–500 μ m [106], and can either spread throughout the bulk or be located primarily between the internal hatching area and the external border [107]. Further investigations have developed a classification system for porosity based on morphology which relates defects to processing parameters; spherical pores caused by



Figure 2.3: L-PBF opens the door to spatially heterogeneous microstructures, the structure of which is highly sensitive to the build process variables. This can result in both undesirable [79](a) and unexpected [80](b) grain structures; however, there is now growing interesting in harnessing this power to deliver 'designer' microstructures. This includes the manufacture of single crystal components [85](c), components with data structures, such as QR codes, embedded within the microstructure [86](d) or even the reproduction of famous artwork [87](e).

gas entrapment or acicular pores due to lack of fusion between layers. As the gas entrapment is a by-product of the powder production process, and Du Plessis et al. showed that the ensuing porosity this caused had less an adverse impact on the mechanical properties compared to other defect forms - therefore, this is generally a lesser concern [108]. Other sources of defect formation include a lack of material fusion due to incorrect laser power, impurities in the feedstock and improper gas flow [109, 110, 111].

To summarise, L-PBF is the vanguard of additive manufacturing. As researchers move beyond recreating traditional wrought components and begin to utilise the power of L-PBF, have been truly awe-inspiring results, particularly in the control of microstructure, creation of lattice structures and alloy in-situ, not with traditional fabrication methods. However, it presents an unprecedented challenge for the field of materials characterisation. The complex process-material interactions

| Name | Description | Size | Reference |
|----------------|--|----------------------------|-----------------|
| Balling | Molten material is not a flat layer, but | 10 - $350~\mu\mathrm{m}$ | [112, 113, 114] |
| | instead creates large spherically shaped | | |
| | particles on the surface | | |
| Acicular | Pores in between layers of the AM pro- | 50 - $500~\mu{\rm m}$ | [115, 107, 116] |
| Pores | cess | | |
| Spherical | Entrapped gas pores within the bulk of | 5 - $20~\mu\mathrm{m}$ | [115, 107] |
| Pores | the material. Material dependent | | |
| Cracking | Cracks occurring due to residual | 50 - $500~\mu{\rm m}$ | [117, 118] |
| | stresses caused by heat sink effect | | |
| Inclusions | Originating from trace presence of im- | 10 - $200~\mu{\rm m}$ | [119] |
| | purities in build powder | | |
| Lack of Fusion | The melt pool varies in size and unfused | 100 - $150~\mu{\rm m}$ | [120] |
| | powder is present | | |
| Geometric | Curling of the bed and dimensional in- | 1 - $100~\mu\mathrm{m}$ | [121] |
| inaccuracies | accuracies | | |

 Table 2.1: Defect table, adapted from Everton et al. [106]

results in components with properties heavily dependent on process variables. Compared to traditional manufacturing techniques, the defect density in components remains obstructively high. Understanding the potential impact of defects in L-PBF components and the stringent performance requirements of high performance or safety-critical industries highlights the importance of component certification in AM. The need to characterise these parts must be balanced against the high-specific cost of creating L-PBF components, making batch testing undesirable. Therefore, for L-PBF to realise its potential beyond a laboratory novelty, a new suite of appropriate characterisation techniques are essential.

2.2.2 Wire-arc additive manufacturing

Process

Wire-arc additive manufacturing (WAAM) uses a wire feedstock and welding arc to build up multiple layer structures [122], see fig. 2.4 [123]. In comparison to powder-bed fusion, WAAM is able to produce larger parts at significantly faster deposition rates, up to 10 kg/h compared to 50-200 g/h [124]. However, these techniques generally cannot produce the intricate featuring seen

in powder-bed fusion, making WAAM well suited to less complex, large build-volume components [123]. The main drivers for WAAM's adoption are a reduction in manufacturing cost thanks to minimal feedstock waste, and a reduction in lead times [125]. Further detailed information on the WAAM process can be found in Martina et al. [123].



Figure 2.4: Diagram of automated welding equipment used in WAAM.

Typical features and industrial considerations

The drive to use WAAM comes from its ability to quickly create large components thanks to a high deposition rate, with relatively low associated costs thanks to a high material utilization and simple equipment requirements. Williams et al. provide several industrial examples of the use of WAAM, which convey the size of components discussed. For example, a titanium wing spar of 1.2 m was produced for BAE Systems [125].

Components manufactured by the wire-arc process are characterised by a strongly anisotropic microstructure, primarily columnar grain growth in the build direction [72]. In WAAM processed Ti-6Al-4V, this leads to < 001 > alignment in columnar prior- β grains, which gives rise to the strong texturing of α grains [126]. This is known to contribute to mechanical anisotropy [127], inferior fatigue life [128], geometric distortion and crack initiation at prior- β grain boundaries [129]. Martin et al. have suggested that one way of producing a microstructure more suitable to engineering applications and preventing the induced residual stress is to promote equiaxed grain growth [18]. The use of inter-pass rolling has also been trialled in WAAM to induce plastic strain to relieve tensile stresses and enact grain refinement [130, 131]. This approach was found to

significantly reduce the prior- β grain size and weaken the texture of both the prior- β and final α structures [26, 126]. The resulting texture leads to components with more isotropic properties.

Recent literature focused on the industrialisation of WAAM has shown that the effectiveness of the rolling process is dependent upon the roller's geometry, as well as the rolling strategy, rolling load and/or the spacing between adjacent rolling passes [132]. Ensuring this grain control is the paramount challenge in the qualification of WAAM components [125]. Therefore, a technique to monitor grain refinement would allow for closed-loop control of the rolling force, minimise parameter optimisation time and, crucially, enable quality control facilitating the industrialisation of WAAM. Currently, inspection of the microstructure requires the specimen to be sectioned and prepared for diffraction techniques such as electron backscatter diffraction or through etching and microscopy. Neither option can be applied in-line with the WAAM process non-destructively.

As WAAM is essentially a welding process, it can suffer from the same defects which blight traditional welding, namely porosity, cracking, delaminations, and oxidisation [133]. Of these, porosity is again the most common, however rather than being process-driven, this is primarily due to poor quality wire feedstock or mishandling [134]. This has primarily been an issue for aluminium components and is not a main concern in the titanium components [15]. Additionally, there is evidence that inter-pass rolling can close these voids to improve density.

Aside from traditional defects, an important issue for WAAM is controlling the geometry of the component; operating outside of the processing window can cause unexpected and significant variations of the geometry of the deposited bead. As such, the development of in-line monitoring techniques has primarily focused on optical and thermal methods, with a view to controlling the build geometry, with little work on microstructure measurement being reported [106].

2.2.3 Materials

This thesis presents results in three materials, titanium, nickel (and its superalloys) and highsilicon steel. The following section briefly summarises each material, highlighting their use in an industrial context. Microstructural features which are encountered later in this work are remarked upon.

Nickel

Rarely used in its pure form in engineering, nickel is instead almost exclusively used as an alloying element. An exceptional combination of corrosion resistance, high-temperature strength, and toughness makes nickel superalloys ideal for the hot section of turbine engines.

Since Frank Whittle's first jet engine utilised the Nimonic nickel-superalloy and henceforth, the

development of the two have been intertwined. Rapid development in the intervening period has led to materials that can tolerate temperatures of up to 1450 K - 90% of the melting point (it is this ability to operate at a high fraction of its melting point which is fundamental to the superalloys). Nearly two tonnes of nickel superalloys are found in modern-day jet engines, with other common modern-day uses, including pressure vessels and chemical processing components [135].

Nickel and its superalloys form a face-centre cubic structure, with various secondary phases such as γ' and γ'' possible, providing strengthening through precipitation hardening. Unfortunately, the properties which make many nickel superalloys desirable also make them difficult to process, with these materials demonstrating a high propensity of cracking. Processing routes that allow components to be manufactured in these superalloys without defects would allow their exceptional properties to be fully exploited [136].

Nickel is unique in that it is regularly used in industry in three distinct crystal forms: equiaxed polycrystalline, directionally solidified (columnar grains) polycrystalline, and finally, single crystals. As the grain boundaries are sites for damage accumulation at high temperatures, the blades in the early stages of the turbine are typically single crystals. In contrast, the blades in the later stages of the turbine are fabricated from equiaxed nickel alloys [135].

Unfortunately, typical processing parameters in additive manufacture induce columnar grain growth, generating a strong [001] fibre texture in the build direction, reducing the range of potential application of additive nickel materials. Unsurprisingly, this microstructural anisotropy results in strongly anisotropic mechanical performance [77].

With these points in mind, Attallah et al. summarised the outstanding issues in the pursuit of additive superalloy components as: creating high density and crack free components, particularly using 'unweldable' feedstock; creating equiaxed microstructure materials, and control of the microstructure more generally; creating components with isotropic mechanical performance at least equivalent to their traditional wrought counterparts; and of less importance to the work at hand, improving the communities understanding of residual stress formation in additive nickel[137]. Developing tools for rapid qualification of the microstructure is an important step towards realising all of these targets.

Titanium

Although not widely used until the later part of the twentieth century, titanium is now ubiquitous in aerospace, pressure vessels, surgical implants and maritime vessel hulls. Possessing high strength, corrosion resistance and ductility whilst maintaining a relatively low density, approximately 60% of nickel, titanium fulfils the demands of innumerable engineering demands. Given that titanium is also relatively easy to process and widely available, it is clear to see why it is a popular choice

across engineering [138]. At room temperature, pure titanium exhibits a close-packed hexagonal structure, known as α -Ti, at 888°C titanium experiences an allotropic transformation to a cubic structure, β -Ti. The alloying of titanium allows access to the β phase at room temperature. Ti-6Al-4V is an $\alpha + \beta$ titanium alloy, favoured thanks to its high strength-to-weight ratio and good corrosion resistance; this gives it uses in aerospace, pressure vessels, gas turbines and surgical implants [138].

The high processing temperatures found in additive manufacturing promotes the formation of prior- β grains at high temperatures, before rapid cooling to form α' , and α in a Widmanstätten pattern [134, 139]. Burgers orientation relationship defines the transformation between the high temperature cubic β -phase and the low temperature hexagonal α -phase, given in Equation 2.1. This allows twelve α orientations to form within a prior- β grain [140].

$$\{110\}_{\beta} / (0001)_{\alpha} \qquad \langle 111 \rangle_{\beta} / \langle 2\bar{1}\bar{1}0 \rangle_{\alpha} \tag{2.1}$$

In practice, for a known orientation of the parent β -phase, g^{β} , the twelve possible α -variants can be predicted by Burger's relationship. This process has been explained in detail by Davis [141]. This limited number of orientation possibilities leads to large microtextured regions, much larger than the size of the individual α -phase grains [142]. Sinha et al. showed that the crystallographic orientation and size of the large microtextured regions are associated with facet-initiation sites of fatigue cracks [143].

High silicon steel

A soft magnetic alloy, silicon steel is used for low to medium frequency applications in electric machines and transformers [144]. The addition of silicon to iron offers increased electrical resistance and improved magnetic softness. High silicon steel alloys in particular offer high electrical resistivity, near-zero magnetostriction, and low magnetocrystalline anisotropy, making them promising for magnetic device and electrical machine applications [145]. However, in the case of alloys > 2%silicon, embrittlement of the material occurs. This has meant that despite the excellent properties of 6.5% high silicon steel, its use beyond the laboratory has been minimal. Ouyang et al. suggested that if a suitable manufacturing route could be developed, high silicon steel would rapidly see industrial acceptance [144].

Goss, whom the texture shown schematically in fig. 2.5a is named after, demonstrated that the properties of silicon steel could be improved by controlling the grain orientation [146]. By aligning the easy axes of magnetisation, fig. 2.5b, the permeability is increased, and iron losses minimised - compared to an isotropic equivalent [147]. This has made the material an excellent candidate for high-voltage power transformers. Several routes are possible to achieve this texturing, including rapid solidification and annealing after cold rolling. The target crystallographic textures for alignment of these easy axes of magnetisation are known as Goss and cube textures; a cube texture aligns the $\{001\}$ planes with the <001> directions in all axes. However, in the literature, it is not uncommon for authors to report a cube texture relative to a perpendicular bi-sector in the material [148], and therefore the difference between these two textures becomes arbitrary, and we may consider them one and the same for the purpose of this thesis.



Figure 2.5: In electrical steels, strong crystallographic alignment allows access to the easy axes of magnetisation, meaning magnetic saturation can be reach for a lower applied field (right). For the body centre cubic lattice of silicon-steel (centre) the easy axis of magnetisation is the <001> direction. Therefore, it is desirable to create a Goss texture (as shown of the left) or cube texture, in the fabrication of silicon-steel components.

Garibaldi et al. demonstrated that near-fully dense 6.9% silicon steel specimens could be produced by powder bed fusion [148]. Crucially, the melt pool depth was shown to control the final texture formed, with a strong cube texture formed when a deep melt pool was used. The melt pool geometry is not directly controlled; instead, the laser power and scan speed are varied in conjunction to achieve the desired effect.

High silicon steel provides an interesting example material as controlling the microstructure is essential and highly process dependent, with a clearly defined industrial application. Particularly as the control of the orientation of the crystal lattice occurs in three axes, rather than two of a typical fibre texture. Therefore, silicon steel is a good candidate for SRAS investigations, and results in this material are presented in chapter 6.

2.3 NDE and characterization tools for additive manufacturing

To date, a lack of well-suited validation techniques has been cited as one of the most significant inhibitors to the industrial pick-up of AM. Commonplace NDE techniques were introduced in §1.5.1, this section extends this to reported applications of these techniques within additive manufacturing.

There is much work concerned with process-based monitoring with additive manufacturing, which do not directly detect the physical properties of the specimen. However, as these are not analogous to the aims of the work at hand, this should not be considered a comprehensive collection. The past decade has seen a wealth of literature in this domain, concerned with both characterising materials and the process. Several authors have conducted thorough literature reviews into the state of the art, including Everton et al. and Grasso and Colosimo, which cover process-centric measurements in greater depth than discussed here. The reader is referred to the review conducted by Du Plessis et al. [149] for an up-to-date and exhaustive report on process monitoring technologies.

Virtually all in-situ NDE techniques to gain traction within the industry to date are non-contact methods. The logic behind this is that the challenging environment within the build chamber, loose powder spray, acoustic noise and space limitations all make measurement more difficult. Due to the need to use non-contact stand-off devices, the majority of techniques trialled thus far have been for surface inspection rather than volumetric techniques.

As stated by Du Plessis et al. in their recent review of non-destructive testing for AM, the assessment of texture and microstructural inhomogeneity is 'not something usually required for NDT, usually done destructively using representative samples' [149]. It is with this commentary in mind that the following review attempts to cover the place of both defect detection and microstructure assessment within additive manufacturing.

Optical

Taking observations of the component or process is the easiest way to determine the success of the manufacturing process. The earliest and most straightforward method of monitoring the build was to install a camera to observe the build chamber. Foster et al. provide a good summary of this approach, providing several examples of observed defects, including cracking, uneven powder distribution and recoater blade damage [150]. Generally, this approach is limited to observing gross defects and complications. As such, this has fallen from common use in favour of the more advanced methods discussed below, which provide enhanced sensitivity.

Within additive manufacturing, there are now two primary applications of optical imaging. Scanners based on structured light, photogrammetry and laser scanning are now common for assessing the geometry correctness are now possible in-situ. Where sufficient resolution is available, these techniques can also be used to measure the surface roughness of components. Several authors have also demonstrated that failing to achieve an even powder distribution leads to dimensional inaccuracies of the finished part [151]. As such, these same optical systems have been deployed to measure the evenness of the powder bed prior to irradiation [152].

High-speed imaging systems, with frame rates \gg 1,000 Hz, have also been extensively used to study additive manufacturing. Again this allows the correctness of the build to be monitored, as demonstrated by Kleszczynski et al., who implemented a camera onto the external body of an EOS build chamber. The authors found the technique was able to measure feature sizes across the powder bed and could therefore be used as a technique to measure geometric correctness insitu. Furthermore, the method was also shown to be capable of detecting voids in the powder bed following bed re-coating, which could result in porosity defects occurring in that layer. However, one common issue with high speed imaging is the sheer volume of data generated, the majority of which goes unused. For example, Berumen et al. trialled a high-speed camera to measure the dimensions of the melt pool with a frame rate of 16,666 frame/s, giving a data output of 75.1 GB/s for full bed imaging [153]. Thus, whilst high-speed imaging remains a valuable technique for understanding the processing dynamics within the build chamber during fabrication, due to the volume of data generated, the difficulties associated with defect recognition and being limited to inspection of the component surface mean high-speed imaging is more suitable to monitor the melt-pool interactions or spatter dynamics. This approach has been used to good effect by Bidare et al. for studying the particle and fluid dynamics in L-PBF within an open architecture build system to better understand the fundamentals of the process [154, 155].

A novel optical microscopy based approach has been proposed by Wittwer and Seita, who recognised that EBSD is not suitable for characterising samples with long-range microstructure heterogeneity, and a more appropriate tool is required for additive manufacturing microstructures [156]. When a material is chemically etched, the preferential dissolution of crystallographic planes can generate topographical surface features which are linked to crystallographic orientation [45]. These features reflect visible light preferentially at certain angles, producing an anisotropic reflectance effect. Adjusting the position of the source allows characteristic directional reflectance spectra to be captured, which can then be fitted against simulated spectra to determine the crystallographic orientation [157]. Comparison with EBSD in L-PBF nickel specimens showed excellent

quantitative agreement (no values quoted) and was acquired two orders of magnitude faster. Indeed, this method was later used to capture the inverse pole figure of the nickel superalloy with the embedded QR code example, as shown fig. 2.3d. The principal drawback to this method is the need to etch the specimen surface to generate the characteristic reflection profiles, this will prevent the measurement from being deployed in-situ, and the etched surface may not be desirable, necessitating additional surface treatment after inspection. Furthermore, there are safety concerns with chemical etching, particularly in titanium-based alloys.

Acoustic emission

The study of the acoustic emission during the laser powder fusion process was reported by Wang et al. [158], who placed an acoustic sensor on each end of the substrate and monitored the emission during the fabrication process. This sensing technique shows some promise for crack detection, as clear signals were generated by the presence of cracks in the sample, but the characterisation of these cracks proved challenging. Whilst it is stated that the signals are collected and post-processed, the steps required in signal processing is not clearly outlined. It is likely that the process of laser melting is acoustically noisy and the SNR of defect signals to background noise is low. Furthermore, this technique would be limited to isotropic homogeneous materials, where the acoustic velocity is well known and constant. Strantza et al. [159] have assessed the potential to use acoustic emission as an inspection technique and found in L-PBF manufactured Ti-6Al-4V samples locate and monitor the propagation of cracking.

Pouet et al. investigated the use of a random quadrature demodulation system for recording the acoustic emission from laser welding, a process with a high degree of commonality with L-PBF [160]. By following the melt pool, the authors were able to detect a lack of fusion and partial penetration defects from the time trace. Other defects appeared to generate weaker or noisy acoustic emissions, but there is scope for additional post-processing to improve the viability of this data. From a practical viewpoint, such detector systems cost ~ \pounds 100k, which may be prohibitively high - this will be considered further in Chapter 7.

Koester et al. have further developed this approach for DED, again mounting piezoelectric transducers to the underside of the build plate [161]. The electromagnetic emission spectrum was found to range between 150 kHz to 2 MHz, spectral feature-based clustering method was used to create 'signature' signals for the build process, with the authors suggesting deviations from these the could be used to indicate defects in the build process.

The great challenge of acoustic emission data is interpreting the collected signals, which are necessarily complex. Artificial intelligible algorithms have found popularity for decoding these datasets and providing a more reliable defect detection method. This was first demonstrated by Shevchik et al. [162]. The authors used a fiber-Bragg grating for the detection of airborne acoustic signals, citing their high sensitivity across a wide spectral range, compared to using piezoelectric receivers. A 4 mW laser was used to continuously irradiate the fibre, with the reflected signals from the grating then detected by a photodiode. Building one L-PBF specimen with varying the build parameters throughout, the authors created areas that are well built and regions with high defect densities. The specimen was repeatedly sectioned, and optical microscopy was used to locate defects, the position of which was then correlated to the acoustic signal. The authors classified each sectioned layer as low, medium and high quality based on the number of defects. The final trained network correctly classified signals to these categories with an accuracy of > 80%. Similar approaches have since between trialled with structure-borne acoustic signals, detected using piezoelectric sensors, to classify overall part density [163]. The obstacle to using neural network techniques widely has been the need to train the network with signals specific to that build, meaning it is only applicable to the material and build parameters of the training data; furthermore, it can only provide broad classifications (e.g. low quality). To see an uptake in industry, generalised algorithms, at least capable of changing build geometry, and event detection will be required.

Optical emission

By far and away, the most common monitoring approach to additive manufacture is the use of photodiodes to record the optical emissions from the build process. Fig. 2.6 shows the electromagnetic emission spectra from the L-PBF process when using a 1064 nm build laser. The information in the first, green, region can be used to measure the existence of visible plasma formation. The second, red, region allows the build laser power to be monitored. Finally, the information in the third, yellow, area allows the temperature of the melt pool to be monitored [164]. The primary advantage of this technique is the data is 'free' and is spatially indexed. As the melt-pool dynamics are directly related to microstructure of the final component and defect formation, monitoring the optical emissions allow indirect sensitivity to these characteristics.

An early implementation of such a system by Concept Laser used a photodiode to monitor the powder bed for over-melted areas and vary the laser powder based on this feedback [165, 153]. This was previously a patented system, but the patent has recently been lost, and many machine manufacturers have now implemented similar systems.

Renishaw's InfiniAM Spectral system tracks the laser input energy and the response for two photodiodes with a maximum responsivity at 1070 nm (laser emissions) and > 1100 nm (melt pool emissions), respectively. The capability of this system was evaluated by Egan and Dowling [166]. The authors demonstrated a convincing linear relationship between the input energy, the photodiode signal level and the strength of the part under compressive testing. The authors



Figure 2.6: Irradiating the feedstock of an additive process with a high-power energy source emits an electromagnetic spectrum. These emissions can be used to as a signature of the build, with deviations indicating errors in the build process.

system.

suggested this could be used to create signatures for commonly built components.

Eriksson et al. [164].

Away from specific IR imaging, Berumen et al. have shown it is possible to measure the electromagnetic radiation across the build by a photodiode alone, giving information on the thermal gradient across the build. The wide temperature range seen around the melting area in L-PBF components (25 to 3000 K) means the resolution of the system has to be relatively low to capture such a wide window. This technique is patented by Concept Laser to offer closed-loop feedback. Clijsters et al. attempted to further develop this technique to allow direct defect detection and classification. Following image processing, artificial discontinuities of 1 mm have been detected, but typical L-PBF defect sizes could not be detected using this technique [167]. A generalisation of this technique has been suggested as an approach for all in-line measurements by Grasso and Colosimo [168].

Thermography

An extension of optical emission monitoring is to focus on the long wavelength range which carries temperature information. Monitoring thermal gradients can highlight discontinuities which affect the thermal conductivity of the specimen. This can be applied in two fashions, passive tomography to study the build process or induction thermography for the detection of defects, *ex post* manufacture [169].

Hooper demonstrated a two-wavelength high-speed camera for monitoring the melt pool [170], results are presented in fig. 2.7a. Stability and dynamics in powder bed fusion procedures are affected by the evolved temperature fields in the melt pool. Additionally, in the case where the temperature field is adverse, defects are prone to appear. Similar approaches to this have infrared camera [171] or a scanning pyrometry [172]. As the melt pool dynamics are the cardinal factors in determining the final properties of the competent, these approaches provide valuable information on the manufacturing process.

Traditionally, thermography has no place in materials characterisation - injection thermal energy to a component cannot reveal the microstructure. However, Raplee et al. exploited the relation between melt-pool temperature to correlated scan strategies to the final microstructure [173]. The authors demonstrated whilst line scanning generated typical columnar grain growth of an (001) orientation, using short point exposures could create areas of equiaxed microstructure with a significantly more random crystallographic texture, with point exposure regions reaching temperatures of 1200°K, 50°K hotter than the line exposure regions, fig. 2.7c. This is part of the same broad body of work that developed many of the microstructure control techniques introduced in §2.2.1. The relationship becomes more complex when the geometry becomes complex, whilst the authors introduced a heuristic for indicting if a region is more likely to be columnar or equiaxed they conclude by stating further trials in different materials and build processes are required.

To date, limited results have shown the ability of thermal monitoring to detect errors within the L-PBF build process; currently, thermography is focused on passive techniques monitoring the melt pool. As an example, fig. 2.7d shows the results of an IR monitoring study by Krauss et al. [174]. The authors added artificial defects to the build to assess the ability of melt-pool monitoring to detect in-defect components caused by insufficient heat dissipation during the build process. The temperature distribution across the build chamber bed could be observed using an IR camera at a long wavelength banding and 50 Hz sampling rate. Due to pace restrictions and the need to protect the monitoring system from the debris within the build chamber, the camera was mounted externally, limiting the field of view to 160 x 120 mm, 30% of the total build area. It was shown that process irregularities could be detected so long as they occurred at a rate less than the sampling frequency, and defects of $\geq 100 \ \mu m$ could be detected. However, in comparison to the typical defects seen in L-PBF, table 2.1, a significant number of defects typically occur at less than 100 μm in size. Furthermore, in reference to fig. ,2.7 the defect position is not immediately apparent, and classification or sizing of the defect is near impossible.

Cerniglia et al. compared laser ultrasound (see §2.3.5) to laser thermography for the detection defects in additive manufacture specimens and concluded the robustness and the whilst ultrasound was superior for the detection and evaluation of flaws, the robustness and the easy set-up of thermography would make it attractive [175]. However, in this work, the authors needed to paint the surface of the specimen, giving a matt finish to ensure greater absorption optical energy – clearly an undesirable extra step.

2.3.1 X-ray computed tomography

X-ray computed tomography (XCT) is generally seen as the gold standard for NDE of additive components. There are two primary uses for XCT in additive manufacturing, porosity measurements and dimensional metrology; for the purposes of this thesis, only the first topic is of relevance. Thompson presented a detailed review of X-ray computed tomography for additive manufacturing, and the reader is referred to this text for a comprehensive review [176].

When considering the overall density of the specimen, the Archimedes method generally provides superior accuracy and is, of course, much simpler to conduct [177]. The value of XCT is, therefore, the ability to localise and characterise specific defects. This has proved particularly advantageous when combining XCT with mechanical measurements, allowing the failure initiation at porosity sites to be observed. To this end, Carlton et al. demonstrated that the distribution of porosity through the component has a greater influence on the tensile strength than the bulk density that would be measured by the Archimedes method, reinforcing the importance of defect localisation [178]. Through XCT, it is now understood that in samples with large inhomogeneous void distributions, failure will be dominated by individual flaws as opposed to bulk density, with cracks consistently initiating at pre-existing large voids. Authors have demonstrated these detected defects can be correlated to build events such as misaligned overlap regions between contour and hatched sections [179], or high-density inclusions due to powder feedstock contamination [180].

Tammas-Williams et al. demonstrated the ability of XCT to both detect and extensively characterise porosity [107], based on size, morphology and location, within L-PBF components. An example dataset from this work is presented in fig. 2.8, showing the map of porosity through a full specimen (a) and two high resolution zoomed images, highlighting the range in pore size and morphology within the component. The quality, and ease of interpretation of this data should be



Figure 2.7: (a) Thermographs of meltpool, showing the temperature distribution and morphology of the meltpool at 10µs intervals over a single laser pulse [170] (b) Typical equipment set-up to make thermography measurements in L-PBF. The chamber door may need altered to minimise the viewing angle [174]. Alternately, the equipment may be mount in-line, much like the optical emission digram in fig. 2.6. (c) Meltpool temperatures may be correlated to the resulting microstructure formation; point melt regions reach a higher temperature and induce an equiaxed microstructure, whereas the typical lien melt produces the expected columnar grains structure [173]. (d) Thermograph and accompanying micrograph of additive specimen surface, showing artificial flaws in the build. Areas around the defect sites are less able to dissipate heat, producing clear indications in the thermograph [174].

noted, in comparison to techniques introduced later in this section, as this has been a principal factor in the prominence of XCT within additive manufacturing, thus far. It was found that whilst larger defects could be reliably detected using a macro XCT system the lower detection limit in terms of pore size was $\sim 25 \ \mu\text{m}$. Whilst defects below this resolution limit do occur in additive specimens, authors have argued away this point by suggesting only defects > 25 μ m are fatigue critical, and thus smaller defects can be ignored [181].



Figure 2.8: (a) Example of XCT dataset, captured with a voxel size of 9.9μ m in powder bed Ti-6Al-4V. The red markers indicate voids throughout the volume, consistent with a high level of entrapped porosity. (b) Zoomed volumes of the same specimen, machined from the locations indicated by blue boxes in (a). Using smaller volumes allows the voxel size to be decreased to 2.1μ m, hence allowing the smaller porosity to be visualised [107].

The fundamental drawback of XCT is the volume of specimen which can be interrogated. For example, titanium specimens are generally limited to a thickness of 40 mm. Therefore, smaller coupons from the same build are usually studied as a proxy for measuring the actual component. However, this has raised questions on the representativeness, given the importance of geometry and build parameters on the final structure and defect formation, as discussed in §2.2.1.

From X-ray μ CT over repeated scans of multiple X-ray spectra, it is possible to reconstruct the microstructure distribution of materials [182]. Initially, this was limited to the measurement of material phase distributions. However, contemporary research has demonstrated the ability to determine each grain's full crystallographic orientation and morphologies in the measurement volume. McDonald et al. used this method to track the grain formation during the sintering of copper powder [183]. Whilst this method is subject to the same restrictions as 'classical' XCT, the potential applications of non-destructive 3D orientation measurements in additive manufacturing are abundant, and it is sure that forthcoming publications will explore this to provide seminal insights into the microstructure formation in additive components.

2.3.2 Other traditional NDE approaches

Despite their ubiquity in the NDT industry as part of the 'big-five', the use of eddy current, dyepenetrant and magnetic particle testing for additive manufacturing has thus far proved unsuitable, albeit in a limited number of studies. In all techniques discussed in this section, there is no reporting of the ability to measure or characterise the microstructure of additive components.

The as-built rough surface of additive parts makes dye-penetrant inspection challenging as the dye is captured in many features on the surface. The rough surface must be machined to allow testing, following manufacture, to be suitable for the detection of surface-breaking cracks [149]. To date, despite many authors dismissing magnetic particle inspection, there appears to be a lack of actual studies. In all likelihoods, magnetic particle inspections would be plagued by the same issues as dye penetrant and with the additional limitation of only being applicable to ferrous components, for which there is less of a market within additive manufacturing - compared to nickel and titanium.

Eddy current is of interest as it is well suited to the detection of surface-breaking and nearsurface defects, particularly cracking. Use use of eddy current testing for surface crack inspection in L-PBF components has previously been reported [184]. It was found as-deposited surface roughness contributed to a lower signal to noise ratio due to the lift-off effect. An exciting development has been the work of Todorov et al., who have demonstrated the first in-situ measurements by eddycurrent. By mounting the probe to the recoater blade to maintain contact with the surface of the component, measurements can be made layer-by-layer. Detect deliberate lack-of-fusion type defects ~ 10 mm in diameter. This approach is limited to defects > 1 mm and relies on calibration (to eliminate the effect of the grain structure) using previous builds of the same material with known properties [185].

Returning to the problem of counterfeiting and traceability, introduced in §2.2.1, Eisenbarth et al. proposed embedding voids within components to create unique signatures of average porosity, which could be detected by eddy-current testing as an anti-counterfeiting measurement [186]. The authors demonstrate a prototype system that can validate parts within a few seconds. However, the density of porosity required to generate unique and measurable signals will significantly degrade the functional performance of the component.

2.3.3 EBSD and diffraction techniques

The study and understanding of microstructure in additive manufacturing relies heavily on the use of EBSD. It is customary for new materials and processes to be subjected to several rounds of EBSD to understand the microstructure formation. The traditional downsides of EBSD have been the speed of acquisition and limited area of study, as the specimen must fit with the confines of the vacuum chamber. These issues are present for any materials researcher but particularly within additive manufacture for two reasons. Firstly, the propensity for heterogeneous microstructure results in wider areas having to be mapped to characterise the structure fully. Secondly, given batch runs and a high degree of customisation, it is counter-intuitive to have to destroy the specimen to characterise the microstructure.

Most commercial detectors offer acquisition rates in the range of hundreds of points per second. Developments in detector technology have increased this to 3000 and even 5000 points per second, but these remain uncommon beyond academic pursuits. Tong et al. have shown that forescatter electron images (acquired an order of magnitude faster) can highlight the grain structure. Therefore, Kikuchi patterns and the subsequent EBSD analysis only need be performed once per grain; in a test specimen, this was shown to make the process $100 \times$ faster [187]. As detector speeds increase, it is becoming more common to interrogate large areas by EBSD - however, this is still on the scale of a few mm²[188], and usually requires stitching together several discrete maps [189].

A particularly challenging example is the study of Ti-6Al-4V, where the user wishes to study the prior- β structure. To recover a map of this cubic β structure, the orientation of each α -phase lathe must be measured and fed into the Burgers orientation relationship [190]. As outlined by Seifi et al., the scan must have a high enough resolution to record α orientations from the lath $(1 - 2\mu m$ width) whilst probing an area that is large enough to record data from multiple prior β grains for reliable statistics, recognising that each β grain can be several millimetres in size [24]. Additionally, it may be a challenge to simply fit a sufficiently large sample within the vacuum chamber, such that many multiple prior- β grains are present, given the size limitations of EBSD (upper limit of ~ 20mm in specimen length). In the case these conditions are satisfied, the resulting scan can have up to 100 million measurement points, necessitating big-data/cloud computing approaches for appraising the dataset. This is an overly time, data and resource-intensive process, given that it is generally the texture or average orientations of the α -lathes over a meaningful area that is of interest, rather than the orientations of each grain; however, the nature of EBSD means each grain must be resolved.

The research groups of Collins and Pollock are making particularly exciting developments in the use of EBSD for AM, delivering the first 3D microstructure maps of additive specimens. The recent work of DeMott et al. presents, for the first time, the EBSD analysis of additive Ti-6Al-4V in three dimensions. Note that whilst 3D-EBSD has been available for over twenty years, it remains uncommon (~ 1% of EBSD publications per year make use of this) [191]. The authors used the established technique of sectioning the specimen by focussed ion beam, repeating the EBSD acquisition at each sectioned layer. Pollock et al. have used a similar system with the addition of a femtosecond-laser to form the so-called tri-beam system [192]. The femtosecond laser offers faster material removal rates than the focussed ion beam. This method was used to study the solidified melt pool of a powder bed nickel superalloy specimen [193]. The strong columnar grain growth, aligned to the < 001 > direction, was seen away from the melt pool, within the melt pool volume, grains form radially around the centre of the melt pool near the top surface, with significant misorientation to the columnar growth. Previous modelling has failed to capture this microstructure formation. Thus, this experimental technique has paved the way for a new microstructure modelling - this is a vital step for enabling the forms of microstructure control discussed in §2.2.1. Characterisation in 3D facilitates a comprehensive analysis of the solidification process beyond that which conventional 2D metallography allows and is, therefore, a valuable new technique for understanding the physics of additive processes. Of course, in both methods, the studied volume is wholly destroyed by the sectioning process.

In-situ synchrotron imaging has been used extensively by Lee et al. to study the volumetric metal pool dynamics [194]. This body of work has demonstrated that porosity is due to the formation of a keyhole in unstable melt pools. This creates a road map for reducing porosity by optimising processing parameters. Similarly, Zhao et al. have developed a real-time monitoring system using high-speed x-ray imaging for powder-bed fusion [195] for the characterisation of the spatter origination mechanism. However, the laser scan path is typically limited to a single continuous track, again limiting to the study of process dynamics and not inspection of additive components.

2.3.4 Ultrasound

Ultrasound, induced via a piezoelectric transducer located under the baseplate, has been trialled as an in-situ volumetric inspection technique. Rieder et al. monitored the increasing build height by observing the shift in back wall echo location as a function of time. The authors demonstrated this arrangement was capable of detecting the presence of a 2 mm diameter sphere embedded within a cylinder of diameter 20 mm and 10 mm tall [196]. However, the results show limited defect detection capabilities. The authors suggest this may be a result of transducer misalignment but suggest future studies with phased array may be more effective. This was progressed further by manufacturing a half-cylinder with a radius of 50 mm and width of 30 mm, containing nine internal cylinders, 2 mm, 3 mm and 4 mm in diameter, increasing with depth.

In traditional NDE, it is common to see researchers test the capability of their new technique by examining a specimen with side-drilled holes. The development of NDE for AM is no different, apart from the ability to build these defect-like features as part of the AM process, rather than placing them after manufacture. The literature usually names these build-in or nested defects, and the latter will be used henceforth. There are two competing arguments on the validity of this approach. On the one hand, real AM features, such as the rough surface, will change the scattering matrix of the defect compared to machined holes [197] - meaning they are not equivalent acoustic reflectors. Conversely, by nesting the defects as part of the build process, the interpretation of results must also consider the accuracy of the build process, for example, under-sizing of hole diameter. This means the measurement of the defect size is a function of both the fidelity of the ultrasonic measurement and the AM build process; thus, isolating the sensitivity of the ultrasonic measurement necessitates a complementary measurement of the defect size, usually by x-ray or similar. A number of studies have demonstrated the capability to detect such nested defects using laser ultrasound techniques. Javaldi et al. have attempted to tackle this problem by presenting a method for reliably creating defects of a known size by placing tungsten carbide balls within the build, claiming these were more representative of defects than side-drilled holes and were of a known size allowing calibration for sizing [198].

The use of imaging algorithms are an obvious route to improving ultrasonic measurement utility and fidelity. When working with ultrasonic phased arrays, it is now common to employ the total focusing method. In order to assess the suitability of the total focusing method for AM, Javadi et al. have reported a series of studies focused on the detection of defects in wire-arc specimens [199]. The authors compared the performance of single-element wheel probe, conventional phased array methods and the total focusing method for the detection of side-drilled holes with diameters between 0.5 mm and 3 mm. The single element probe failed to detect even the 3 mm holes. The conventional phased-array image highlighted the majority of 3 mm defects but missed those close to the surface. The total focusing method proved capable of detecting down to the 0.5 mm holes; however, the frequency of inspection needed to be increased to 10 MHz to obtain indications of the near-surface defects. No sizing of the defect was attempted. The primary limitation of this approach is the need to use a gel/liquid between the transducer and the inspected specimen to maintain acoustic coupling, particularly impractical for inspection during fabrication. To address this, the authors reinvestigated the air-coupled transducer wheel, using a 64-element array rather than a single element as previous reported, with the expectation that it would provide the dual benefits of the improved imaging capability of a phased array whilst negating the need for couplant [200]. This approach was shown to be capable of detecting defects of 1 mm diameter, without the need to machine the inspected surface. However, as contact needs to be maintained between the transducer roller and the deposited specimen, the temperature must either be allowed to cool after deposition or the roller be capable of working at temperatures in the hundreds of degrees Celsius, the researchers citing this as a key area for future development.
Huang et al. used ultrasound to study the microstructure-property relationships in additively manufactured stainless steel [201]. An increase in attenuation of the acoustic wave was correlated to increasing porosity in the specimens, whilst this does not allow for the detection of specific voids, it does facilitate statistical based monitoring of porosity levels. The acoustic wave velocity is a function of both grain size and porosity, but as Stanke and Kino demonstrated [202], the relationship between grain size and wave speed is complex, due to the scattering and propagation effects (as introduced in §1.8 and are expanded upon in §2.5), even without considering the effects of porosity. An iterative solution to this unified theory for bulk wave propagation in polycrystalline materials was previously prosed by Kube, and utilised in this work to predict velocity and attenuation coefficients, noting this requires knowledge of the single-crystal elastic constants of the material [203]. Given the complexities of this method, it is without surprise then that attempts to correlate wave velocity with porosity and grain size were inconclusive. Nevertheless, this work represents the contemporary zenith in the application of ultrasound to the determination of microstructure-mechanical properties in additive manufacturing - beyond SRAS.

Whilst it has not been applied to additive manufacturing as yet, the work of Lan et al. is worth noting. The authors use measurements of acoustic wave velocity in multiple directions through a specimen to recover the crystallographic texture [204]. At the time of writing, this is the only way to non-destructively measure the bulk microstructure of metallic specimens. Application of this technique to additive specimens is likely a worthwhile pursuit. Tant et al. use full-waveform inversion to characterise bulk microstructures, but this is primarily to suppress the scattering effects of the microstructure for improved imaging for flaw detection in ultrasonic inspections [205].

2.3.5 Laser ultrasound

The non-contact nature has made laser ultrasound a popular candidate for inspection in additive manufacturing. It has been common to use to use time-of-flight methods to detect potential defects in additive manufacturing specimens. Klein and Sears reported the first study of additive manufacturing components by laser ultrasound in 2004, launching this field of research [206]. Using a standard pulse-echo set-up with a commercial two-wave mixing detector, the authors aimed to detect the presence of near-surface holes. In similar pulse-echo set-ups, the simplest way to confirm the presence of a defect would be to look at the arrival time of the reflected longitudinal wave; it be expected to arrive before the back wall echo. However, in this work, the authors stressed the need to detect defects within the last few deposited layers. In which case, the propagation distance of the longitudinal wave to a defect is relatively short $(10 - 100 \,\mu\text{m})$, and therefore the echo may arrive within the initial front-wall echo. As such, the authors exclusively studied the use

of shear and surface waves from mode conversion and the broadening of the arrival content (due to the presence of multiple modes). This, in the words of the authors, 'bulging' of the Rayleigh wave arrival was used to detect the presence of near-surface defects. Despite the authors presenting further wavelet analysis, consistent detection and sizing capability remains undemonstrated.

Yu et al. used a similar pulse-echo arrangement, with the addition of a rough surface detector allowing the as-deposited surface to be interrogated [207]. Four holes, with a diameter between 0.4 and 2 mm were placed 2.5 mm from the inspected surface. An example b-scan of this specimen is given in fig. 2.9a; amplitude drop due to the presence of the holes, whilst the presence of the 0.4 mm cannot be discerned. Building upon this work, Davis et al. used a through-transmission laser ultrasound arrangement to detect nested 1 mm diameter holes, centred 5 mm from both detection and generation surfaces [208]. An example b-scan from this work is reproduced in fig. 2.9b. In this case, a drop in amplitude of the received acoustic signal corresponded to the presence of the nested holes.

Everton et al. conducted a more realistic study by varying the build parameters to generate regions of under-melting, over-melting and increased hatch spacing between melt lines - leaving regions of unsolidified powder. The authors then used bulk wave laser ultrasound to detect the presence of the ensuring defects. Clear indications in the time-domain signal due to the presence of defects were reported; however, the ability of the technique to size, accurately locate or characterise defects was not assessed and appears limited based on current results [209].

Lévesque et al. used laser ultrasound in the destructive ablation regime to study WAAM-like specimens. The synthetic aperture focusing technique was used to correct focussing distortion and improve the resolution of the collected b-scans. In the processed b-scans, the interface between the substrate the deposition could clearly be seen, whilst indications of large internal porosity was confirmed with x-ray tomography. Cerniglia et al. demonstrated the first in-process results for WAAM-like specimens; by attaching the ultrasound system to an industrial robot, alongside the deposition tool, the deposit could be examined seconds after fabrication. Results demonstrated the ability to detect voids within nickel superalloy specimens, of 0.1 mm near the surface and 0.8 mm deeper in the material, defects smaller than 0.1 mm were consistently missed.

Millon et al. presented b-scans of additive components with machined notches on the examined surface, the morphology of which resemble naturally occurring lack of fusion defects. Measurements are made by generating acoustic waves with a line source on the top surface of the specimen and detection by a two-wave mixing interferometer, with the authors focussing on the generation and detection of Rayleigh waves. There is a wealth of literature on the use of SAW for surface cracking, but it is notable that detection was possible despite the rough surface of the specimen [210, 211]. Zhang et al. again used a pulse-echo arraignment and scanned across the surface of the specimen, allowing c-scan imaging in the pursuit of crack-like defects. The authors demonstrated sizing capability with 1 mm notch sized with 1% accuracy. Finally, it was demonstrated that increasingly rough surfaces (within the typical AM operating window) decreased the signal-to-noise ratio of measurement, noting that the inhomogeneous distribution of the SNR is particularly detrimental to defect identification [212].

In most applications using a piezoelectric transducer as a phased array is preferred, thanks to the capability to focus the ultrasonic beam on areas of interest. To arrange 32 or 64 piezoelectric elements within a package with a useful form factor is trivial and does not significantly increase the cost. Whilst it is possible to use multiple laser sources to create such an array [213], to purchase 32 or 64 lasers is unrealistic, not to mention the low portability of such as system. It is possible to use fibre delay lines to pattern the beam meaning only once source is required [214], but this necessitates a complex optical arrangement and the delivery of optical power to the specimen is limited by the power the fibre can withstand. Therefore, to replicate this for laser ultrasound imaging, the most effective approach is to use capture the waveform at every generation-detection pair and use imaging algorithms to create the array in post-processing synthetically. Pieris et al. build upon previous work by Stratoudaki et al. [215], where simulated phased arrays where creating using laser ultrasound, allowing the total focusing imaging algorithm to be used [216]. An aluminium specimen, manufactured by L-PBF, was made with deliberately placed voids of 0.5 and 1 mm diameters, at a depth of 26 mm.

The technique demonstrated excellent capability to detect nested flaws within the component, as shown in fig. 2.9c, comparing well to conventional piezoelectric phased-array imaging. The generation and detection of ultrasound was made from the substrate face of the specimen as this has a 'smoother' surface, but there is no reason why a suitable ultrasonic detector would not make measurements on an as-deposited surface possible. The scan time remains an issue, with the measured component taking 3 hours to measure, although Lukas et al. have demonstrated this can be significantly reduced by considering only the generation-detection positions, which yield high sensitivity to the same location [217]. This is termed selective matrix capture, as opposed to the full matrix captured originally used by Pieris et al. this approach was shown to reduce the scan time by a factor of 10. With a view to integrating within a build system, the primary drawbacks are cost and the fact that defects are observed far from the free surface, the argument being that it is then too late to enact any remedial action.

To provide a brief summary at this juncture, all of the laser ultrasound results discussed thus far are concerned with the detection and sizing of nested or relatively gross defects, reflecting the priority of this field. Whilst additive manufacturing presents some unique challenges, it is inherently a metallic structure, and it should be of no surprise that well-established laser ultrasound



Figure 2.9: Example results of laser ultrasonic studies in additive manufacturing. B-scan data in specimens with side-drilled holes is shown from (a) pulse-echo [207] and (b) through transmission arrangements [208]. (c) The normalised total focussing method image, using shear-shear wave arrival, in the LIPA study of L-PBF aluminium specimen with nested defects. Clear indications of all but one defect are present [216].

techniques for detecting deep porosity in wrought materials remain viable. Much of the interest in additive manufacturing the opportunity, and it is therefore notable that few works have proposed a technique capable of detecting near-surface defects, providing the opportunity to correct the build. In contrast, little weight has been given in the research to the use of laser ultrasound to characterise the microstructure of additive components.

In the domain of material characterisation, Dong et al. staged through transmission experiments in additive aluminium specimens, with the aim of correlating the measurements to the grain size [218]. The authors demonstrated through experimental results and simulation that the attenuation rate and the grain size was well correlated. Although, due to the directivity pattern of ultrasound generation, the results become less reliable as the detection beam is moved away from the epicentre of generation. It should be noted that the estimation of grain size by measuring the attenuation of the acoustic waves using laser ultrasound was first proposed in the 1980's (see Scruby et al. [219]) and is now well established, with recent work tending to focus on applying these measurements in-situ. For example, the recent work of Malmström et al. in the measurement of grain sizing of steels, as part of the hot strip mill process, provides a good example of this state-of-the-art in this field away from additive manufacturing [220].

Zhan et al. investigated the use of laser ultrasound for the measurement of residual stress in titanium DED components; the authors had to first find the acoustoelastic coefficient (the parameter which relates the change in acoustic velocity to applied stress) by tensile loading. The authors showed laser ultrasound results showed good agreement with the accepted hole-drill method for residual stress measurements. However, the need to determine the acoustoelastic coefficient makes this an impractical technique to apply. Furthermore, the numerical framework used in this case assumes isotropy, an assumption that does not hold for most additive components [221].

Ma et al. have recently presented a novel hybrid technique that combines laser ultrasound with optical emission to study the microstructure of WAAM components offline [222]. The ablative generation of ultrasound creates plasma with the characteristic optical spectra, reflecting the elemental composition within the generation area. The authors were able to map variations in the titanium content across the specimen surface, correlating this to a region of anomalous as measured by EBSD. Again, the measured attenuation coefficient from the acoustic data was used to infer the grain size in the material - these results showed some correlation to baseline measurements with EBSD. The authors state the technique is well suited to in-process inspection within WAAM and present a possible implementation, although no in-process results have yet been reported [223]. The authors attempt to argue around the need to calibrate the acoustoelastic curves as the elemental composition may be tracked, and they assume a database of pre-calibrate curves for each material will be available).

As with 3D-EBSD, volumetric imaging is an exciting area of much interest. Therefore, with an eye on the future, it is worth recounting the work of Thréard et al. where by measuring the time-domain Brillouin scattering of the transverse wave, the grain structure is revealed in a volume of $30 \times 30 \times 8 \mu m$ [224]. Whilst this reveals only the acoustic contrast and not crystallographic orientation, work is underway to tackle this inverse problem. Development of this and derivative techniques will be a significant step in the study of additive manufacture and materials characterisation more broadly.

2.4 Elastic constant measurements

As stated in §1.4.4, it would be difficult to overstate the importance of elastic constants within engineering and physics. To reiterate for the reader who remains unconvinced by this statement, a few pertinent examples of the uses will lay bare their significance. Fundamentally, the elastic constants relate stress to strain [29], and provide a bridge between the micro- and macroscopic worlds. From an engineering perspective, knowledge of the elastic constants and orientation in anisotropic crystalline materials is vital for understanding the in-service mechanical performance and facilitates the calculation of engineering elastic parameters including: Voigt-Reuss-Hill moduli, shear modulus, Young's modulus, bulk modulus, Poisson's ratio, and unveils other physical properties including strength, hardness, wear and melting temperature [30]. For the materials scientist, measured elastic constants can communicate information on many important phenomena in solids including magnetic and electron-lattice interactions, phase-transitions and mode softening [31]. In the search for new alloys, ensuring crystal stability is of paramount importance, requiring knowledge of the elasticity. Moreover, they are requisite in any acoustic materials inversion scheme (including determination of the bulk texture or grain size). In short, elastic constants have an essential role to play in the calculation of many key physical quantities material behaviour under loading, and in turn, inform our wider understanding of materials.

Primarily the need to measure the elastic constants is driven by the evolution of material processing techniques and novel alloy development. However, the progress in measuring techniques has fallen out of step with the rate of development in materials. This has resulted in a situation where the majority of techniques for determining the elastic constants are only viable with single crystals specimens. However, it is infeasible for the vast majority of compounds to be grown into single crystals, making such techniques unsuitable. Du and Zhao estimate this applies to around 99% of the possible 160,000 distinct solid compounds, in addition to many solid-solution compounds used in common structural alloys [225]. Additionally, the proliferation of powderfeedstock additive manufacturing techniques has opened the door to easy in-situ alloying, allowing new alloys to be created rapidly[13]. In such cases, the elastic constants for the pure element are used, often resulting in significant errors in the calculated mechanical properties.

The role of elastic constants within NDE is not always immediately obvious. For example, the velocity of a longitudinal wave in an anisotropic medium can be expressed as $c_l = \sqrt{\frac{C_{11}}{\rho}}$ (it is not uncommon to see this given as $c_l = \sqrt{\frac{E}{\rho}}$ but this refers specifically to a wave in a rod), all ultrasonic measurements are in some sense probing the elasticity of the material. Indeed, calibrating the material velocity is an essential step in any ultrasonic flaw detection or thickness measurement.

Prior to giving an overview of the experimental techniques, as will be given in §2.4.1, some general comments and a potted history of elastic constant measurement is necessary. One of the first practical experiments undertaken by a mechanical engineer will be the measurement of tensile stress in a bar in response to strain, which allows Young's modulus to be found. Through this elementary method, Young (of Young's modulus) and Riccati took the first steps on the road towards characterising the elasticity of a material in the late 18th century. However, it was a further century before the single-crystal elasticity matrix was measured. The requisite breakthrough came in the form of the eponymously named Bridgman method, which provides a technique for growing single-crystal metals of controllable orientation. This allowed for single crystals to be studied and the relationships above to be exploited in the pursuit of elastic constant measurement.

The period between the early 1950's and late 1970's saw a wealth of investigation into the determination of elastic constants, particularly as ultrasonic techniques became increasingly viable. Many of the measurements made in this time period are still regularly cited in today's literature. Hearmon tabulated elastic constants for a variety of sources in the Landolt-Bőrnstein series [226], later updated by Every, and his remains an essential reference for elastic constant values. Similarly, the compilation of iron and nickel elastic constants by Ledbetter and Reed, despite it being nearly fifty years since its original publication is of significant value [28]. Is it values from this Landolt-Bőrnstein series for which Hearmon is often credited in modern literature, however, in the majority of cases these are the averaged results of other authors results. Averaging to find the values assumes the variations stem from experimental inaccuracy, however a brief glance through such complied tables reveals notable variation in the reported values for many materials. To demonstrate this, fig. 2.10 visualises the reported elastic constant values of pure nickel, as reported to date, on a series of violin plots. There is seen to be significant variation in the reported values of all three elastic constants, with notable outliers.

Given many of these measurements are over 50 years old, one may assume the variation is a function of measurement error that may reasonably be expected to improve with time. However, whilst the case of errors is not specifically dealt with by many authors, the tables of Appendices A and B summarises these values where reported. In most cases, authors estimate a measurement error of no greater than 1%. If this is taken at face value, then it does not explain the range in reported values for many pure materials. Unfortunately, the considerable development in the experimental techniques, inversion processes and a more pressing demand to characterise new materials over this period has taken precedence over repeated studies of the elastic constants in ubiquitous materials, such as nickel. This leaves the true source of these variations somewhat unaddressed.

Ledbetter and Reed suggest the variations in elastic constants reported for the 'same' material are less a symptom of poor measurement accuracy and more a failure to rigorously characterise the materials studied [28]. Variations in the measured elastic constants can arise from chemical impurities, residual stress, processing route and magnetic saturation. Trace impurities influence elastic constants through an impurity-dislocation interaction; dislocations contribute an additional strain upon stress application. Similarly, small residual stresses affect measured elastic constants. Therefore, when measuring the elastic constants, it is preferable to study a low-impurity, and well-annealed specimen as both residual stresses and dislocation density are reduced by annealing.



Figure 2.10: Violin plots of reported elastic constants values for pure nickel, values taken from Ledbetter and Reed [28]. For each constant the white dot indicates the median, horizontal lines are the mean, the height represents the range, the width is the amount of data (i.e. horizontal positioning of data points is done only to prevent overlapping and aid readability), and the vertical lines in the middle are the interquartile range. The y-axis is common to the three plots.

The actual experimental impact of chemical impurities is not well reported; Speich et al., who determined the variation in the engineering elastic moduli for iron due to alloying elements, provide one of the few works in this area [227].

Values of elastic constants for various cubic and hexagonal materials are given in table A.1 and B.1, respectively. Values are reproduced as given in the original literature. Where possible, uncertainties and errors are listed, however to date many studies have omitted these values. Where error bounds are given, these are often the error associated with the measurement of the direct property, offered with the caveat that the errors in the final calculation of the elastic constants will be no less than the initial measurement error. Furthermore, the tools for testing the validity of elastic constants extend to only checking the stability criterion for the tensor. Beginning with the eponymously named Born criterion for cubic crystals [228]:

$$C_{11} - C_{12} > 0,$$
 $C_{11} + 2C_{12} > 0,$ $C_{44} > 0$

and later extended for other crystal systems [229], including hexagonal:

$$\begin{split} C_{11} > |C_{12}|, \qquad C_{33}(C_{11}+C_{12}) - 2C_{13}^2 > 0 \\ C_{11}C_{13} - C_{12}^2 > 0, \qquad C_{44} > 0. \end{split}$$

Alers and Neighbours used this approach to detect errors in reported elastic constants for brass and gold [229]. Of late, Kube et al. have encountered similar issues with incorrectly tabulated property constants and have taken up this challenge with the development of their propSym tool [230]. This tool is primarily for higher-order physical tensors as it calculates the number of independent members, this is well known for the fourth-order elasticity tensor (but less well known for other properties such as the piezoelectricity tensor); unfortunately, it does not provide insight into the validity of elastic constants, beyond checking they satisfy the symmetry of the tensor. All of the constants listed in Appendices A and B satisfy the stability conditions; beyond that, it is impossible to draw further conclusions on their validity. Alongside the advancement of technology to measure elastic constants, further tools for validation are also a necessary development.

2.4.1 Measurement methods

The following section introduces and discusses the merits and drawbacks of the four common techniques for determining the single crystal elastic constants of a specimen.

Static measurements

Static measurements are a general family of time-independent approaches which rely on measuring the engineering elastic moduli and converting these into the single crystal elastic constants. As discussed in §1.4.3. the inverse of the stiffness tensor, C_{ijkl} , is the compliance tensor, S_{ijkl} . The compliance tensor relates directly to engineering moduli which can be directly measured by compression, tension, bending or torsional testing. General forms of E, ν and G, as a function of orientation are provided in appendix .C, with selected examples returned to stiffness terms in equation 2.2, away from these primary directions writing the moduli in terms of elastic constants becomes arduous. The bulk modulus is independent of crystallographic direction thus only one expression is necessary.

$$E_{[001]} = C_{11} - 2 \frac{C_{12}^2}{C_{11} + C_{12}}$$

$$E_{[101]} = 4 \frac{C_{44}(C_{11}^2 + C_{11}C_{12} - C_{11}2C_{12}^2)}{2C_{11}C_{44} + C_{11}^2 + C_{11}C_{12} - 2C_{12}^2}$$

$$G_{[001]} = C_{44}$$

$$B = \frac{1}{3} (C_{11} + 2C_{12})$$
(2.2)

In order to determine the single-crystal elastic constants, the specimen must be a single crystal of known orientation. In such samples, three independent measurements (for example, E measured in the directions [001], [101] and [111]) allow the elastic constants to be determined by the inversion of the compliance matrix. These techniques were widely used to determine the elastic constants, including those of brass [231], nickel [232] and cobolt [233], for example. The primary experimental sources of error are the precise calibrated measurement of stress and strain. During popular use, up to the 1950's, these calibration errors were usually quoted to around 1% [28], and it is realistic to imagine this could be further improved with modern experimental equipment.

Unfortunately, these errors become significant due to the propagation of errors when inverting the compliance tensor. This was succinctly demonstrated by Hearmon, who showed in cubic materials an error of just 2% in the determination of S_{11} and S_{12} would lead to errors of 7.2% and 13.4% in the determination of C_{11} and C_{12} , respectively. It should be noted when working with materials that exhibit positive C_{12} , representing the vast majority of engineering materials, inversion of the stiffness tensor does not lead to greater errors in the compliance tensor. Therefore, by measuring the elastic constants, the engineering constants can be reliably calculated, whilst the reverse is not true. These approaches were the first commonly used elastic constant techniques, primarily because the experimental methods were already well established; however, ultrasonic measurements have generally superseded the use of static methods over the past fifty years, thanks to the improved accuracy.

More recently, nanoindentation modulus measurements coupled with orientation measurement has been used to obtain orientation-dependent Young's modulus and gives a viable approach to applying these techniques to polycrystalline materials [234]. The method has not yet demonstrated the capability of evaluating the full elastic constant values, but there is no reason why this is not feasible. Similarly, atomic force microscopy (AFM) based methods are well-established for the elasticity measurement of biological specimens, producing extremely high spatial resolution qualitative modulus maps but extraction of accurate elastic constants is still far from a reality, especially for hard materials such as metals [235].

Bulk wave velocity

By preparation and careful alignment of the principal crystal axes with those of the experiment, it is possible to derive relatively simple equations describing the acoustic velocity as a function of elastic constants. Table 2.2 shows the relations of the bulk wave velocities to elastic constants in cubic materials and hexagonal materials.

Similar relationships can be approximated for other planes; however, the expressions become much more complex, the relations for arbitrary directions are given in appendix D. For cubic

| Cubic | | | |
|----------------|---------------------|------------------------|---|
| Plane (hkl) | Direction of motion | wave mode | Relation to ρv^2 |
| (001) | [001] | longitudinal | C_{11} |
| | [001] | shear | C_{44} |
| (101) | [001] | longitudinal | $\frac{1}{2} \left(C_{11} + C_{12} \right) + C_{44}$ |
| | [001] | fast shear | $\frac{1}{2} \left(C_{11} - C_{12} \right)$ |
| | [001] | slow shear | C_{44} |
| Hexagonal | | | |
| Plane (Φ) | wave mode | Relation to ρv^2 | |
| 0° | longitudinal | C ₃₃ | |
| | shear | C_{44} | |
| 90° | | longitudinal | C_{11} |
| | fast shear | | C_{44} |
| | slow shear | | C_{66} |

Table 2.2: Example relationships between elastic constants and bulk wave velocities, for planes in cubic and hexagonal materials.

materials, such as aluminium and nickel, if the (001) plane is exposed, the measurement of the three bulk wave velocities are sufficient to determine the three unknown elastic constants, with relatively simple expressions. In hexagonal materials however, at least two distinct crystal cuts are required to determine the five unknown elastic constants. Measurements along the prism plane are sufficient to determine C_{11} , C_{12} and C_{44} , however these measurements have no dependence on the value of C_{13} or C_{12} , therefore an second cut at an arbitrary angle ($0^{\circ} < \Phi < 90^{\circ}$) is required.

The simplicity of this method has made it the technique of choice for many authors, meaning the elastic constants of most pure metals have been measured by this technique [236, 237, 28]. Briefly disregarding the demands of specimen preparation - covered in the next paragraph; the actual acoustic measurement of velocity makes no special demands compared to other studies concerned with precise velocity measurements, for example wall thickness determination. The point being that the measurement of acoustic velocity can be made with errors of less than 0.001% and that there is no reason why these cannot be translated to the measurement of elastic constants [31].

The primary drawback of this method is the need to carefully prepare a single crystal of the specimen, cut at a known orientation. Usually, faceting is required to expose multiple crystal planes for measurement, and therefore this is certainly a destructive process. Ideally, the crystallographic planes of measurement are high-symmetry directions, such that simple relations as in table 2.2

may be derived. However, in the case of low-symmetry measurements (such as those required in hexagonal materials, as described above), then algebraic inversion is generally infeasible and numerical optimisation schemes must be employed. These optimisations have the general form of a least-squares minimisation

$$F = \frac{1}{N} \sum_{i=1}^{N} \left(v_i^{\text{meas}} - v_i^{\text{calc}} \right)^2$$
(2.3)

where v^{meas} is the experimentally obtained velocity, N is the number of sampling points in the measurement and F is the value of the objective function. The task is then to vary the calculated velocity, v^{calc} , as a function of C_{ij} to minimise F.

Surface wave measurements

As with the bulk wave velocities, the velocity of SAW depends on the materials elastic constants. However, it is not feasible to derive closed form expressions relating the SAW velocity to the elastic constants in arbitrary directions, therefore a forward and inverse model is required making this a generally more complex approach. The general advantage of surface wave techniques is acoustic localisation, allowing measurements to be made within specific grains. As surface waves provide greatest sensitivity to the shear constants C_{44} and $\frac{C_{11}-C_{12}}{2}$, however, they are less sensitive to longitudinal elasticity, leading to an acknowledged indeterminacy between C_{11} and C_{12} [238].

Developed in the late 1990's acoustic spectro-microscopy (ASM) was developed to measure the acoustic wave velocity with a view to extracting the elastic constants [239]. This approach was applied to coarse textured cubic pure nickel, and interestingly used to simultaneously find the crystallographic orientation, to date this is the only reported method to solve for both elasticity and crystalline orientation - although the approach is limited to cubic structure materials. However, beyond a limited publication history, it appears the technique was a abandoned shortly there after due to the impracticality of measurement. An acoustic couplant was required and the acoustic measurement system required with the separate optical microscopy unit. The measurement is also highly sensitive to eccentricity when rotating, requiring further challenging alignment. Despite this, the technique did facilitate high measurement accuracy, constants within 5% and orientation within 1°. Kim et al. have used line-focus acoustic microscopy to measure the elastic constants of thin films and exotic alloys [240]. Much like ASM, this requires the specimen to be immersed for acoustic coupling, limiting the application of the technique [240].

In a collection of works by Seiner and Landa, transient grating spectroscopy has been used to measure the elastic constants in shape memory alloys, where stress induces phase changes in the materials crystal structure [238]. The results . To improve the sensitivity to C_{11} and C_{12}



Figure 2.11: Schematic of ASM system, as used to localise large grains and measure acoustic slowness surfaces for inversion to determine the elastic constants and/or crystallographic orientation. The grain center is found on the microscopy arm of the system before translating the specimen to the acoustic arm. The velocity surface is captured by rotating the specimen through the ϕ stage, whilst the incident angle is controlled by the θ stage. Measurements are made for each (ϕ, θ) pair, in order to find the Rayleigh critical angle. Acoustic coupling is provided by immersion of the sample [239].

the authors are investigating the possibility of using both surface wave modes and the quasilongitudinal wave in the inversion to provide a more deterministic solution [241].

Du et al. have developed a method where-by a grating film of polydimethylsiloxane is deposited on the sample surface at the area of interrogation [225]. Illumination of the film by laser pulse then excites a surface acoustic wavelength equivalent to the grating spacing, extracting the frequency of the measured acoustic wave then allows the SAW velocity to be found accurately. This was repeated in several directions to capture slowness surfaces, and used to determine the elastic constants of a number of metallics including nickel and aluminium. However, there are two primary drawbacks. Firstly, the film needs to be deposited and replaced for each measurement position and propagation direction. Secondly, the proposed inverse methods again requires the orientation of each grain to be measured carefully beforehand by EBSD.

Very recently a laser pump-probe measurement was developed to extract elastic constants from a coarse-grained polycrystalline stainless steel sample, without knowledge of the grain orientations, by measuring the SAW [242]. The method is only applicable to cubic crystal structures with the assumption of a random distribution of the crystal orientations and requires measurements over 1000 individual grains; thus it can only be applicable to samples satisfying these stringent conditions. Other methods are developed to extract elastic constants from highly textured samples only [243].

A method of determining the elastic constants without the need for couplant, film, or painstaking alignment of a grain centre to the acoustic source is extremely desirable. Particularly if an accompanying inversion scheme could be used without prior knowledge of - or strict assumptions on - the crystalline orientations in the specimen.

Resonant ultrasound spectroscopy

The general working principle of resonant ultrasound spectroscopy (RUS) for the measurement of elastic constants is outlined in fig 2.12. Typically, specimens are precisely cut as a parallelepiped with care taken to ensure opposite faces are perpendicular and normal to the principal crystallographic axes. In practice, this is achieved by time-consuming iterations of machining and Laué or electron backscatter diffraction to determine the orientation.

The forward problem to predict the natural frequencies for a specimen of known dimensions and elasticity is relatively simple to implement. However, the desire is to calculate the elasticity from measured resonances - the inverse problem. No analytical solution exists to the inverse problem; thus, an iterative numerical approach must be taken. By comparing the measured frequencies to those computed from the forward model for a range of elastic constants and minimising the residual, the experimental elastic constants can be found. The relatively simple apparatus used lends itself to making measurements at elevated temperature measurements. Leisure and Willis provides a comprehensive outline of the 'standard working' concepts of RUS in [244]. Further to this, notable, with respect to the work at hand, improvements and use cases of RUS are discussed below.

RUS is well established and has been widely used in the study of elastic constants in industrial metallics and earth minerals in the general configuration outlined above. The technique is to accurately determine the elastic constants, assuming the measured inputs (dimensions, density and orientation) to the forward problem have first been precisely measured; otherwise, the resultant errors are significant. Simply machining a single crystal specimen along the crystal growth direction is insufficient, as only a few degrees of misalignment can lead to unacceptably large uncertainty in modulus estimates.

Notable recent developments in the field include the work of Pollock et al., who have proposed Bayesian inference for the inverse problem [245]. Most inverse schemes provide only point estimates, that is to say, a single value. What is lacking is the associated error or confidence in the result. As discussed in §2.4, this is a notorious problem for elastic constants, where multiple values are for the same material without indication of the associated accuracy. In most cases, users of the values are left to propagate sources of measurement error to arrive at a 'best-case' accuracy. Computing a Bayesian posterior yields probability distributions, meaning the confidence of the result is inherent to the inversion method. In this work, the authors report the elastic constants for the Ni superalloy CMSX-4 to a maximum accuracy of ± 4 GPa. Whilst this is larger than the 1 GPa/1% quoted by many authors (see appendix A), it is a more realistic estimate of the accuracy as it considers the inverse process rather than just measurement errors. Notably, this group have also then applied this technique in the study of additive specimens.

The resonant frequencies of a polycrystalline material depend on the crystallographic texture of the grains through the specimen - this means there is no longer a method to extract the singlecrystal elastic constants. However, it is possible to determine 21 independent elastic constants of the textured elasticity matrix. This also allows the texture coefficients to be extracted, similar to the work of Lan et al. discussed in §2.3.4. This approach was taken by Rossin et al. to determine the texture coefficients and elastic constants in additively manufactured cobalt-nickel alloys [246]. This allows the complete bulk elastic response of the specimen to be characterised but relies on accurate knowledge of the single-crystal elasticity. In previous work, the authors measure the elastic constants of single crystal cobalt-nickel specimens, enabling this later work on the additive material [247].

2.5 Spatially resolved acoustic spectroscopy

Spatially resolved acoustic spectroscopy (SRAS) is an important SAW measurement technique, allowing the local velocity to measured across the specimen surface. The technique has been developed at The University of Nottingham and can trace its roots back to earlier work in acoustic microscopy and adaptive optics. In 2003, Sharples presented the optical scanning acoustic microscope [4], leveraging the non-contact and non-destructive nature of laser ultrasound to create an alternative system that removed the need for couplant and surface contact by a transducer, as used in scanning acoustic microscopy systems, which were common at the time. This opened up a new research domain for laser ultrasonic imaging. Rayleigh- and Lamb-type waves were used to image defects and materials characterisation, such as coating thickness. Having established the O-SAM instrument for isotropic media, Sharples et al. conducted a series of experiments to develop a technique capable of accurately measuring the SAW in polycrystalline media, in spite of acoustic aberrations. The authors first set out the problem's scale, describing the possible steps to achieve improved measurement fidelity in anisotropic polycrystalline media [68].



Figure 2.12: Working principle of resonant ultrasound spectroscopy for elastic constant extraction. The input transducer is driven by a sinusoidal signal, this in turn drives vibration of the specimen under test. The frequency components of this vibration are function of the specimen the dimensions and elastic properties. In traditional operation the elastic constants are extracted from a single crystal material of known orientation. A polycrystalline specimen may also be studied to extract the textured elasticity matrix, in which case the single crystal elastic constants must be known for the material under examination.

When using a focussed acoustic beam, it is possible to compensate for aberrations by using an adaptive source. By first measuring the plane's point spread function perpendicular to the propagation direction at the desired point of detection, it was possible to obtain a new excitation profile that would redirect the acoustic waves towards the desired detection point. This adaptive correction could take two forms, lower-order 'tilt' correction, where the generation image is rotated slightly to refocus the acoustic wave at the detection position, or higher-order correction shaping the wave-front, as shown in fig. 2.13.

It was also shown that by measuring Green's function response of the material and using this as a basis for optimising the generation of the SAWs that it was possible to create regions of high spatial coherence, even when propagating the surface wave through anisotropic polycrystalline media [69]. This allows SAW velocity to be robustly determined across the sample surface by tracking the phase gradient of the zero-order of diffraction of the measurement region, provided a good measurement of the Green's function response could be captured for the material.

Concurrent with this body of work, Hong et al. used the O-SAM system to accurately measure



Figure 2.13: (a) The propagation of SAWs on an isotropic media, the acoustic wave is focused at the detection point. (b) The propagation of SAWs on an anisotropic polycrystalline media, the acoustic wave is focused away from the intended detection point. Lower order correction can be applied by simply tilting the arc to return focus to the detection point. (c) The aberrated wavefront as measured when using the uncorrected generation image, this can be used to apply higher order aberration correction, as in (d). Experimental amplitude c-scans at 82 MHz on $2 \times 2 \text{ mm}$ polycrystalline aluminium acquired before (e) and after (f) aberration correction [248].

the velocity surface on the (001) plane of nickel and (111) plane of silicon single crystals [249]. A spatial light modulator was used to generate the fringe pattern. Having measured the SAW at several acoustic wavelengths, a Gaussian curve was fitted to the discrete data in the k-vector space, the maximal amplitude was then taken to correct λ , allowing the SAW velocity to be found by the elementary $v = f\lambda$, where f is the mode-lock frequency of the generation laser. This method provided an extremely accurate measure of SAW velocity without damaging the sample, as had been the case with existing laser ultrasound techniques. Whilst being initially developed for single crystal materials, it was discovered this technique would allow measurement of the SAW velocity in polycrystalline materials. This is thanks to the non-dispersive nature of Rayleigh waves, meaning the wavelength of the signal was then insensitive to the problematic grain boundary crossings. This led to the filing of a patent in 2005 [250] and the first publication demonstrating the spatially

indexed measurement of SAW velocity in 2006 [71]. In this work, Sharples et al. measured two titanium alloy specimens with spatial resolution of $\sim 800 \mu m$. The contrast in measured SAW velocity was used to map the specimens microstructure. It was a remarkable step forward to image the grain structure this way, particularity when compared to existing techniques as the need for a vacuum chamber was removed, in theory allowing any size of specimen to be probed.

Visualising microstructure the contrast mechanism is useful, but for in-depth materials characterisation, the orientation of each grain is required. The first steps towards extracting crystallographic orientation were taken when Li et al. presented SAW velocity maps with waves propagating in perpendicular directions, highlighting the sensitivity of the SAW velocity to the anisotropy of the microstructure [251]. The measurement of velocity surfaces was then combined with the ability to measure the SAW velocity across the surface of polycrystalline materials in order to determine the crystallographic orientation of polycrystalline cubic materials [252]. To achieve this, Li et al. proposed a forward model solver to calculate the theoretical SAW velocity, a review of approaches to this problem is presented in §2.5.2, and a complimentary inverse solver to determine the crystallographic orientation from the measured SAW velocity. Agreement with diffraction measurement of orientation was quoted as better than 4.1°.

Results presented through to 2014 were made using instrumentation broadly similar to the original O-SAM system. Smith et al. then developed the first modern SRAS instrument [253], this system is used throughout this thesis and will be introduced in detail in Chapter 3. Briefly, the instrument provided a step-change in scanning rate, with approximately 1000 points per second now possible, making it practical to capture multi-megapixel images of the microstructure whilst shrinking the footprint of the instrument from $\sim 6m^2$ to $< 0.5m^2$.

Subsequent developments have focussed on validating SRAS orientation measurements against EBSD [254] and investigating new materials beyond the industrial metals, including silicon [255]. The power of SRAS comes from measuring the SAW velocity in multiple directions; this requires multiple scans of the sample, significantly increasing the scan time compared to a single velocity map. Therefore, in a bid to realise faster orientation imaging, Patel et al. recently investigated coded grating patterns which would allow the velocity surface to be captured from a single measurement. This was shown to be a promising approach but relied on exacting positioning of the detection beam relative to the coded grating - challenging to realise in the experimental system consistently [256].

2.5.1 Spatially resolved acoustic spectroscopy for additive manufacturing

The following section summarises the use of SRAs in AM, prior to the commencement of this EngD programme and does not report any of the novel contribution of the author. SRAS, and its forerunner systems, have been extensively used to study additive manufacturing specimens, beginning with the use of the O-SAM system to measure variations in acoustic amplitude across the material, which could be correlated to porosity and lack of fusion defects in Ti-6Al-4V specimens. In the same work, the authors presented SRAS measurements in the nickel superalloys MERL76 and Waspaloy, with the contrast in SAW highlighting grain boundary locations [257]. These measurements were all on specimens manufactured by the laser engineered net shaping process, but all subsequent studies have focussed on the L-PBF process.

Smith et al. conducted a thorough investigation into the use of SRAS with Ti-6Al-4V specimens. One particularly interesting outcome was the detection of subsurface porosity thanks to the mode conversion of the surface wave to Lamb wave in a small region, where the thin material acts as a plate-like structure. The measured SAW velocity in this region falls far below the expected range for titanium, $> 300 \,\mathrm{ms}^{-1}$, indicating the measurement of the fundamental antisymmetric mode. The authors also demonstrated a relationship between the mean SAW velocity and the sample porosity.

Hirsch et al. furthered the investigation in the understanding of the microstructure in additive manufacturing, studying four 10 mm cubes made manufactured from the nickel superalloys CM-247LC [81]. CM-247LC is a member of the 'unweldable' set, which suffer from cracking during rapid cooling and solidification; therefore, developing L-PBF build processes which can create crack-free specimens is an active area of research (see §2.2.3) [136]. The authors demonstrated the acoustic dataset could be used to detect and size surface-breaking cracks in the material, with good correlation with measurement by optical microscopy. More value was extracted from the acoustic data by correlating the SAW velocity maps to the approximate crystallographic orientation, i.e. closer to the (001) plane.

Throughout these works, the authors discussed at length the feasibility and problems associated with integrating a SRAS instrument within the build process, with the surface roughness of additive components acknowledged as the primary challenge. Previous SRAS measurements had all been made on polished AM specimens; this was then addressed by Patel et al. in 2018, who presented SAW velocity measurements in 'as-deposited' Ti-6Al-4V specimens [258]. To cope with the surface roughness ($S_a \sim 6 \mu m$), the knife-edge detector was replaced by a speckle knife edge detector. The acoustic wavelength was decreased to 250 µm, compared to 24 µm used in the smooth surface

studies. The SAW velocity maps reveal consistent texture across the specimen surface, consistent with the initial results of Smith et al. in the same material. In this work, the authors also presented a demonstration SRAS system installed within a decommissioned L-PBF system. The generation and detection beams were scanned using galvanometer mirrors rather than translating the specimen by linear stages, as in other SRAS experiments.

Measurements of SRAS signals on rough surfaces in this system illustrated that a SRAS instrument could be integrated within an L-PBF build system and be able to make measurements. The authors concluded by noting the next steps in this domain would require the integration of a SRAS system within a working L-PBF build system.

2.5.2 Calculation of surface acoustic wave velocity

Much of the power of SRAS is found in its ability to determine the crystallographic orientation of materials. This relies upon accurate knowledge of the theoretical SAW velocity, and as such, it is appropriate to briefly summarise approaches to this calculation. The initial steps in this field were made by Lord Rayleigh, who first derived the explicit secular equation for the eponymous Rayleigh waves in isotropic media in 1885 [259], where the velocity is the a root of equation 2.4.

$$R(V) = \left(2 - M_2^2\right)^2 - 4\sqrt{(1 - M_1^2)(1 - M_2^2)} = 0$$
(2.4)

Where $M_1 = \frac{V}{c_L}$ and $M_2 = \frac{V}{c_T}$. In the case of $0.265 < \nu < 0.5$ one real root exists corresponding to the SAW velocity.

However, the calculation of SAW velocity in non-isotropic materials is more complex. When deriving a solution to the problem, most authors, this work included (see chapter 4 for greater detail), follow the Stroh formalisation, which presents an eigenvalue value problem to determine the wave propagation constant that satisfies the boundary conditions. Although alternative approaches have demonstrated this is not the only viable formalisation, it remains by far the most common. This approach was first demonstrated by Stonely, who derived the explicit secular equation for orthotropic materials [260] and a closed form solution, equation 2.5, of the SAW phase velocity propagating in the [100] direction on the (001) plane in cubic materials [261]. Stonely used the Newton–Raphson method to determine the roots of the characteristic equation. Similar solutions may be derived for the (011) plane [262], however as Stonely remarked, the algebraical and arithmetical work involved in deriving such expressions becomes extremely cumbersome. In any case, the resulting eigenvalue problem is not analytically tractable for an arbitrarily rotated crystal of arbitrary crystal symmetry.

$$\left(1 - \frac{C_{11}}{C_{44}}R\right) \left[1 - \frac{C_{12}}{C_{11}}^2 - R\right]^2 = R^2(1 - R)$$
(2.5)

Where $R = \frac{\rho c^2}{C_{11}}$

Destrade reported an explicit form of the characteristic equation for monoclinic symmetries [263], which was later generalised fro any crystal symmetry by Ting [264]. This approach is attractive as the SAW velocity can be determined without the need to iteratively minimise the determinant. However, the derivation of these explicit equations is again extremely challenging, even with modern symbolic math packages. Additionally, these forms do not allow for the calculation of the pseudo-surface wave modes, these are of practical importance as they are experimentally measured in certain crystallographic directions. An interesting approach to the problem has been proposed by Stoklasová et al. using a Rayleigh-Ritz formalisation to calculate the velocity of SAWs. This allows the non-linear equation to be transformed into a linear eigenvalue problem, as is widely used in resonant ultrasound spectroscopy. However, this method does induce quantifiable errors in the calculation of the pseudo modes and does not allow the calculation of the relative amplitude of the wave modes [265].

We, therefore, return to the eigenvalue problem of the Strohl formalisation. The current implementation will be described in detail in chapter 4; it is, however, worth making a few remarks presently. In general, most approaches solve the eigenvalue problem numerically - this can trace its roots back to the method of conducting an iterative search across the velocity range to find values that satisfy the boundary conditions and governing differential equations, as first demonstrated by Farnell [266]. This approach is suitably generalised such that the SAW may be found in any direction on any arbitrary plane, may be applied to any crystal and allows the calculation of all wave modes which propagate on the surface - not just the Rayleigh wave. The reader is encouraged to consult the work of Farnell for greater detail. Despite being over fifty years since its first publication, this seminal work remains the authoritative text on the calculation of surface acoustic wave velocities.

Along crystalline planes where one surface mode can propagate, this approach is particularly elegant as we need only find the single minimum of the determinant, a problem easily tackled by modern numerical minimisation algorithms. However, in directions where multiple modes may propagate, all local minima must also be found. Li et al. reported modifications to the technique, which sort based on relative displacement of the mode in sensitive directions (vertical and in-plane) [252]. This allows the calculation of a compound surface wave from the dominant mode at each propagation direction. The calculation of the displacement follows the method outlined by Farnell, but this is thought to poorly reflect the reality of thermoelastic generation, inducing errors in the transition between dominant modes. This approach has been the basis of SRAS measurements of crystallographic orientation.

Further notable implementations of a numerical search include Cherry's method, where a limiting velocity was introduced, reducing the numerical search range to that where two or more of the eigenvalues become real [267]. This reduces the search space speeding up the calculation and preventing discontinuities in the determinant curve, in theory making minima search algorithms more stable. However, this approach is not generalised to find all possible surface modes.

This iterative numerical search is non-trivial, with a significant temporal cost associated with the calculation. Therefore, most approaches have pre-computed for SAW velocities for a given set of elastic constants and orientation once to use as a forward solution in the inverse problem. This has prevented the of elastic constants and orientation, with the study of orientation favoured as the elastic constants do not need to be varied for simple materials.

A modern implementation of Farnell's method, with modifications to improve computation time and calculate the displacements to which the experiment is sensitive, is described in Chapter 4. The solver is now suitably fast that it may be used as a forward input for the simultaneous inverse determination of elastic constants and crystallographic orientation, as demonstrated in Chapter 6.

2.6 Summary

This chapter has built upon the fundamental concepts introduced in chapter 1 to conduct a review of contemporary literature across fields pertinent to this thesis. This has been done with the aim of carving out the place of this thesis within the wider fields of additive manufacturing, materials characterisation and elastic constant measurement.

Over the past 20 years, additive manufacturing has evolved from a prototyping tool to one capable of delivering production components. The promise of unprecedented design freedom has created a great deal of interest in AM across high-value industries. Two distinct AM techniques, laser powder bed fusion and wire-arc additive manufacturing, have been introduced, including their relative merits and the primary challenges facing industrial uptake of both methods. Recent research on the processing-microstructure relationships has demonstrated that high density, defectfree components are possible in AM. However, there remains wide variability in the performance of AM components - primarily due to the sensitivity of the microstructure to process variables, including geometry and build orientation - with a spatially heterogeneous grain structure commonly observed. Authors have naturally proposed to take advantage of AM to deliver 'bespoke' microstructures, even encoding information within the grain structure.

Of course, to be a viable manufacturing method for these aforementioned safety-critical sectors,

the quality of the AM component needs to be ensured. The use of numerous NDE and materials characterisation methods within AM has have been described, including ultrasound. Many of the techniques introduced are focused on monitoring the process, whilst research into those methods capable of making direct measurements of the specimen have concentrated on detecting void-like defects. This has meant EBSD remains the only widely used tool for characterising the microstructure of AM components. The limitations of EBSD mean measurements cannot be made in larger components that hope to see service, and established destructive approaches are a poor match given the inherent variability and low volume production pf AM. Therefore, there is a clear place for a tool to probe wide areas rapidly to confirm the presence of the desired microstructure — particularly one that has the potential to be applied either in-situ or with little surface preparation.

The measurement of elastic constants is relevant to this work for two reasons. Firstly, knowledge of the elasticity of a material is integral to many of the acoustic inversion schemes used to derive useful properties of the specimen under test - indeed, such approaches will be used throughout both chapters 5 and 6 of this thesis. Secondly, the advent of AM has made it trivial to create functionally graded multi-material parts or novel alloys in-situ. In such instances, the elastic constants may vary across the material or simply not be known for the new alloy. This raises the question of how uncertainty in the elastic constants affects the acoustic inversion schemes or if techniques are available to measure the elastic constants in such specimens directly.

The dominant approaches for the measurement of elastic constants have been described, starting with the earliest static methods, followed by bulk wave measurements which became popular with the arrival of ultrasonic testing and, more recently, resonant ultrasound spectroscopy. Surface wave measurements have also been introduced, although these have not seen widespread use like the other three techniques. Disagreements in the elastic constants of common materials have been noted, with potential error sources discussed, but this remains an unsolved topic. The measurement of elastic constants in polycrystalline materials remains a significant challenge, particularly when the crystallographic orientation of each grain is unknown. A technique that could make accurate measurements of the elastic constants (and the crystallographic orientation) in real polycrystalline specimens would have a broad-ranging impact across materials science and make acoustic inversion schemes applicable to the multi-material dynamic alloys possible with AM.

The conclusion of the chapter has focused upon spatially resolved acoustic spectroscopy measurements, firstly reviewing the technique more generally then proceeding to specifically concentrate on results on AM specimens, before finally introducing methods of calculating the SAW velocity in anisotropic materials - an essential step in finding the crystallographic orientation. The following two chapters will continue this narrative in order, with chapter 3 firstly presenting the SRAS instrument used to obtain experimental results in this thesis, and chapter 4 then method of SAW velocity calculation used in this work, as part of a wider discussion on the numerical schemes used in conjunction with SRAS measurements.

Chapter 3

Spatially Resolved Acoustic Spectroscopy

By the turn of the millennium, the use of lasers to generate and detect ultrasound had been common for many years. However, laser ultrasound imaging systems were, in the main, either prohibitively slow, complicated, of inadequate resolution or simply too expensive. The 'missing link' was a high resolution, rapid means of acquiring images of the interaction of acoustic waves with materials. The pinnacle of this technique, named spatially resolved acoustic spectroscopy (SRAS), builds upon extensive research into laser ultrasound at The University of Nottingham [268]. SRAS is an acoustic microscopy technique that allows the accurate measurement of the local surface acoustic (SAW) velocity across the surface of a specimen, making it a valuable technique in materials characterisation. The first publication of the technique dates back to 2005, building upon previous work in adaptive optics and leading to an active patent [250].

The general history of SRAS has been introduced in §2.5. This chapter builds upon this by providing an in-depth description of the SRAS instrumentation used throughout this thesis to obtain experimental results (Chapter 5 and 6). It is im[portant to note that whilst the system has been rebuilt by the autho over the course of the EngD, At the conclusion of this chapter, the reader should be familiar with the EMDA instrument, signal processing and typical SAW velocity maps resulting from SRAS measurements.

3.1 Early SRAS - k-SRAS

The first embodiment of the SRAS technique to be explored is known as k-SRAS, so named because k-vector spectrum is sampled. A schematic of the optical arrangement and instrument

of this system, as used to conduct early SRAS experiments, is shown in fig. 3.1. A mode-locked Q-switched Nd:YAG laser with a fundamental frequency of 82 MHz is used for the generation of surface waves. By passing the beam of the generation laser through a spatial light modulator, a grating pattern of known spacing is created and subsequently imaged on the sample surface, fig. 3.1(inset (a)).

Through thermoelastic absorption of the laser pulse, a local stress-strain field is created, in turn generating a surface acoustic wave packet. Fig. 3.1(inset (b)) describes the process of matching the temporal properties of the laser with the spatial properties of the sample in the generation of the surface waves. If the grating spacing matches the SAWs wavelength, such that $\lambda_g = vt$ (where v is SAW velocity and t is the time between laser pulses), the generated wave reaches the subsequent grating position at the time of the next laser pulse. This is repeated for the number of pulses within a cycle of the generation laser. Hence, through constructive interference, large-amplitude surface waves can be generated. In contrast, the frequency response has a broader bandwidth and a smaller amplitude when the wavelength does not satisfy $\lambda_g = vt$. Crucially, Rayleigh surface waves are non-dispersive; thus, this frequency of propagation does not change once generated and is not affected by grain boundary crossings or variations in the propagation distance. Therefore, the SAW velocity can be measured for each generation point across the surface of the specimen. This method is unlike traditional time-of-flight measurements and is immune to acoustic aberrations [71]. This SAW then propagates away from the generation patch, causing a perturbation in the surface. This perturbation is then probed with a second laser in continuous-wave operation - as the SAW passes below the location of the probe beam, the reflection of the probe beam deflects. This deflection is then measured using a knife-edge style detector.

Therefore, the problem at hand is that the velocity and frequency of wave packet are fixed (by the material and the generation laser, respectively), and the task is to determine the grating spacing λ_g which maximises the amplitude of the frequency response, and in turn, evaluate the SAW velocity from the elementary equation $v_s = f_s \lambda_g$, where. This is realised by performing multiple scans across the range of probable wavelengths to determine the wavelength of the SAWs. The wavelength is a property of the measurement pixel and will vary with the inherent SAW of the material at that location. By way of illustration, if a nickel specimen is examined (where the SAW velocities range from 2300 ms⁻¹ to $3650 ms^{-1}$) with an operating frequency of 82 MHz, then the range of determined wavelengths should be bounded by $28.04 \,\mu\text{m}$ and $44.51 \,\mu\text{m}$. Each measurement provides a discrete point in k-vector, and Gaussian fitting is typically used to determine the maximum position. The obvious drawback to this approach is the need to conduct multiple scans in order to build up the k-vector spectrum. It is also worth considering the mode-locked generation laser and spatial light modulator are relatively expensive components, making this approach



difficult to realise beyond the laboratory.

Figure 3.1: Schematic of k-SRAS system, showing the generation of SAWs in a polycrystalline specimen from the successive pulse of the mode-locked laser. Inset (a) shows the grating pattern as imaged on to the specimen surface, during laser emission. Inset (b) describes the relationship between the acoustic wavelength, λ_g , and the fundamental frequency, f_s , of the generation laser. As the laser fires a first pule at t₀, SAWs are created at each fringe position. As time passes, moving from top to bottom in this figure, these SAWs propagate forwards and backwards - only one direction is shown here for simplicity. Finally, the laser fires a second pulse at t₁. When the correct wavelength is chosen SAWs generated from the first pulse will be directly beneath the line sources at the next laser pulse.

3.2 The EMDA system and *f*-SRAS

The current implementation of SRAS, known as f-SRAS, addresses the pitfalls of k-SRAS by allowing the entire spectrum to be captured in a single measurement, greatly increasing the acquisition speed. The mode-locked generation laser and spatial light modulator are also replaced by a fixed mask and a broadband Q-switched laser, reducing the system footprint and cost whilst simultaneously improving the general usability of the system.

The principle of matching the temporal properties of the laser with the spatial properties of the sample remains, however, the generation of SAWs is now achieved by imaging an amplitude mask of fixed fringe spacing, λ_g , on to the specimen surface, fig. 3.2a - replacing the spatial light modulator. Unlike the mode-locked laser where the generation of the SAWs occurred over the multiple laser pulses in each pulse envelope, with the Q-switched laser, the SAWs are generated from a single pulse. In fig 3.2b a range of possible frequencies are shown relative to the position of the generation fringe locations - only the frequency which corresponds to the characteristic velocity of the sample is in-phase at all grating positions (shown in red). The number of generation fringes controls the bandwidth of the signal in the frequency domain - whilst frequency components far from f_s fall out-of-phase and are destroyed within a couple of generation fringes, components similar to f_s take several fringes to be destroyed. In effect, the grating pattern acts as a filter on the generation, creating narrowband excitation from a broadband source. As previous, the generated surface acoustic wave propagates at a frequency, f_s , which is determined by taking the position of the maximum in the frequency domain, having applied a Fourier transformation to the time domain signal. Hence, in f-SRAS the acoustic wavelength λ_g is known (determined by the fringe spacing on the amplitude mask and the system magnification) and f_s has been measured, allowing the SAW velocity to be found by $v_s = f_s \lambda_q$.

To restate what has already been said, the fundamental advantage of SRAS is that the velocity can be determined by only knowing λ_g is the critical point. The measurement of SAW velocity is independent of arrival time or distance between the acoustic source and detection position. Furthermore, the frequency content of the generated wave packet depends on the properties of the material under the fringe only, allowing the measurements to be spatially resolved. It is these characteristics that make SRAS a uniquely capable tool for imaging the SAW velocity.

The first modern SRAS experiment and deployment of the f-SRAS method, named EMDA, was built in 2011 - the construction of the system was funded by the now defunct East Midlands Development Agency, hence the system acronym. EMDA marked a significant leap forward in the capability of SRAS; alongside providing a system capable of rapid velocity measurements, with the ability to work on large specimens, the design also prioritised minimising the footprint of the system. The system has been rebuilt since its initial construction, with changes to the enclosure and beam paths to improve laser safety, but this aside, the majority of the system has remained unchanged. The system in its current form is shown in fig. 3.3.

Given that all of the subsequent SRAS systems have been developed in EMDA's image, as it were, and EMDA has been used extensively throughout this work, it is important to characterise the system.

3.2.1 Generation optical path

EMDA utilises an AOT-YAG-10Q as the generation laser, with a custom miniature head allowing it to be directly mounted to the optical breadboard. The source itself uses Nd:YAG crystal as its gain media, producing linearly polarised light at 1064 nm, with active Q-switching.





(a) The generation patch (red) and detection spot (green) are imaged on the surface of the sample. The imaged distance between fringes determines the acoustic wavelength, λ_g . Positioning the detection spot in either shown location does not change the frequency of the measured acoustic wave.

(b) The distance between fringes, λ_g , determines the frequency, f_s , of the SAW wave packet. Only the frequency component corresponding to the acoustic velocity of the sample is in-phase at all grating locations.

Figure 3.2: The SAW is generated within the centre grain of (a), it can then be detected within the grain or any subsequent position along the propagation path (as shown in the figure) without changing the measured frequency.

By varying the laser pump current, the pulse length and energy can be varied between $\sim 1 \text{ ns}$ to 3 ns and 80 µJ to 200 µJ, respectively. In practice, the system is now almost exclusively used at a pulse width of 1.5 ns at 2 kHz repetition rate, giving a pulse energy of 160 µJ, with a peak power of approximately 0.1 MW.

To efficiently generate SAWs, the pulse width must be suitably short to contain sufficient frequency content within the requisite measurement spectrum. Fig. 3.4 shows the temporal and frequency characteristics of a number of nanosecond pulses. The short pulse length of the EMDA generation source provides a wide operating window that can span from tens to hundreds of megahertz.

To enact power control for laser safety and to prevent ablative damage of vulnerable specimens,



(a) Photo of EMDA system.

(b) Detailed optical diagram of EMDA system.

Figure 3.3: EMDA system, the primary SRAS system for on the measurement of smooth-surface specimens with a view to determining crystallographic orientation.

a rotatable half-wave plate and Glan-Taylor polarizer are mounted directly after the laser head. Whilst the same effect, at a lower cost, could be achieved with a polarising beam splitter, the high peak power of the source exceeds the damage threshold of commercially available polarising beam splitters. A Glan-Taylor avoids this by having an air gap between the prism faces. Adjusting the diode pump current as a means of power control is undesirable as this also alters the pulse length.

After power control optics and safety aperture, the beam is folded by two mirrors, allowing the beam to be centred in the propagation plane. All mirrors used in this system are made from fused silica, and those found in the generation path are coated to provide reflectance > 95% across the near-infrared spectrum.

Following the mirrors, the next components are the mask and aperture - these elements structure the light and create the pattern to be imaged onto the specimen surface. The mask is a fused quartz substrate, onto which regularly spaced chrome lines are deposited. The aperture is an extremely thin aluminium foil with a centred hole. A range of mask spacings and aperture diameters are available, depending on the specific application the user prioritises - e.g. spatial resolution (smaller aperture) or bandwidth of generation (larger aperture or higher frequency). Typically, an aperture diameter of 1 mm is used with a mask spacing of 100 μ m. The aperture is placed directly atop the



(a) Temporal response of nanosecond width laser pulses.

(b) Frequency content of nanosecond width laser pulses.

Figure 3.4: A set of laser pulses with pulse widths between 1 and 10 ns. The frequency content of the pulse can be controlled by changing the pulse width - with shorter pulse widths required for higher frequency inspections.

aperture, such that the assembly can be considered a single optical element in terms of diffraction. This assembly is mounted in a small rotation stage, allowing the direction of acoustic propagation to be controlled. Care is taken to ensure the aperture is centred on the mask, otherwise as the assembly rotates, the imaged generation patch is displaced across the specimen surface.

The tube lens of the system is an f100 achromat located 100 mm from the mask, as a result the mask is at the back focal plane of the lens. When combined with the objective lens - a long working distance, infinity corrected Mitutoyo M plan APO - which has an effective focal length of 20 mm a 5:1 demagnification is achieved. This premium lens produces a well-focused image of the grating pattern, with a working distance of 34 mm. Usefully, three other objective lenses of the same product range are also available for use, $\times 5$, $\times 20$ and $\times 50$, these lenses share a parfocal length of 95 mm. Therefore, by changing the objective, the corresponding demagnification can be achieved without changing or replacing any other optical elements, giving the system to access acoustic wavelengths between $8 - 72 \,\mu\text{m}$.

This structured light is absorbed thermoelastically, creating the SAW. Light of 1064 nm wavelength is well absorbed by most engineering metals, materials that appear silver in colour are generally suitable. The specimen surface reflects a small proportion (< 10%) of the light. This reflected light can be exploited to monitor the generation patch on the specimen - valuable for the alignment of the system.

The source is of linear polarization, and a polarising beam splitter (PBS) is mounted such that on the first pass, the beam is unreflected. The PBS is followed by a quarter-wave plate mounted with the fast axis at 45° , which changes the polarization state from linear to circular. When the specimen reflects the beam, the handedness of the beam is reversed. Upon the second pass of the quarter-wave plate, the reflected beam is converted back to linearly polarised light, orthogonal to the original input. Therefore, the beam is then mirrored by PBS_{IR} to a second f100 achromatic lens, producing an image of the generation patch on the specimen at a camera at a 1:1 scale with the mask. Whilst the CMOS sensor of the monitoring camera has low quantum efficiency in the NIR range, the relatively high optical power of the reflected generation beam means the grating image can still be well resolved, despite this inefficiency.

3.2.2 Detection optical path

In order to enact one of the suitable detection schemes for the measurement of SAW velocity, a second laser is required to probe the sample surface. In EMDA, this source is a continuous wave 532 nm laser, achieved by the frequency doubling of a Nd:YAG crystal, with a maximum optical power of 420 mW. The head of this laser is mounted away from the main breadboard, partially to save space and partly because the head has a propensity to overheat - necessitating additional cooling. As such, the beam is fibre coupled into the system; the optical power is reduced to $\sim 200 \,\mathrm{mW}$ after the collimator.

The first optical component the beam interacts with is a PBS, again mounted such that transmission is maximised on the first pass. As in the generation path, this is followed by a quarter-wave plate mounted with the fast axis at 45°, which changes the polarisation state from linear to circular. A piezo-actuated steering mirror is used to displace the beam in the x-z and z-y axes - this allows the position of the detection spot on the specimen to be adjusted relative to the orientation of the grating image. A long-pass dichroic mirror then turns the beam towards the objective lens, where a well-focused spot is formed on the specimen surface.

The specular reflection from the mirror-like surface of the specimen is captured by the objective lens, re-collimating the beam. The dichroic mirror offers a reflection-transmission ratio of approximately 99:1, and this small transmitted portion is then turned by the PBS onto the camera, allowing both the grating image and detection spot to be observed. The majority of the beam instead follows back the initial path to the quarter-wave plate and PBS, where the polarisation is changed back to linear, orthogonal to input, and the beam reflected onto the detection arm of the

system by the PBS.

A second piezo-actuated steering mirror reflects the beam towards an f60 achromat that focuses the beam towards the detector. Finally, a hot-mirror (short-pass) is placed before the detector to prevent undesirable reflections from the generation beam from reaching the photodiode. Without this component, the measured signal is initially dominated by the generation pulse. The surface of the segmented photodiode is displaced slightly from the focal plane of the final achromat to ensure the maximum possible area across both segments is filled. The piezo-actuated steering mirror allows fine control of the beam in order to achieve this positioning, ensuring that both segments of the photodiode are balanced. The acoustic wave can then be measured by subtracting the voltage on one photodiode segment from the other. The KED is a gradient detector, that is to say, it is sensitive to the rate-of-change of the acoustic wave. Thus, it is inherently more sensitive to higher frequencies.

During an experiment, three alignment steps of the detection path are required each time the propagation direction of SAWs is to be changed. Firstly, the detection spot must be repositioned on the specimen surface to be in the propagation path. The KED detector is most sensitive when the defection of the beam is perpendicular to the split between the two segments. To maintain this, the KED itself must be rotated by the same angular increment as the mask. Finally, the position of the beam incident on the detector must be realigned to balance the photodiode segments and maximise the measured acoustic signal. These steps remain unautomated.

3.2.3 Signal detection and acquisition

A schematic of the signal chain is shown in fig 3.5. Depending on the scan parameters and specimen, the passive components may be changed but typically have the following values: a high pass filter with frequency cut-off of 10 MHz, an RF attenuator of 6 dB and a bandpass filter of 90 MHz to 250 MHz. Two powered amplifiers are also placed on either side of the attenuator. This filtered signal is then measured on a digital oscilloscope.

The oscilloscope is triggered by the synchronisation signal from the generation laser, indicating the time the Q-switch device activates, forming the pulse. The oscilloscope has a maximum sampling rate of 4 GHz with a bit depth of 8.

At the commencement of an experiment, a single trace is captured to sample the coherent noise in the system. The primary noise source comes from the Q-switch of the generation laser - the measured frequency of this falls within the measurement bandwidth, corrupting the measured signal. To combat this, the measured noise trace is subtracted from each experimental measurement to suppress this noise source as far as possible. In specimens where large amplitude surfaces waves



Figure 3.5: Schematic of signal acquisition chain between the knife-edge detector and digital oscilloscope, used to filter unwanted noise and amplify the SAW signal.

can be generated, such as the shown in fig. 3.6a, subtraction of the noise is of less importance as the amplitude of the acoustic wave is ~ 2 orders of magnitude greater than that of the noise. In contrast, when working with challenging industrial materials such as fine-grain titanium, subtracting the initial background noise makes a significant improvement to the signal-to-noise (SNR) ratio of the measurement. In the example waveform of fig. 3.6b, captured in highly textured and fine-grained titanium, subtracting the initial noise increases the SNR of the acoustic measurement from 15 dB to 19 dB.

3.2.4 Realising a scanning instrument

The specimen is mounted on two linear stages, allowing the optics to remain fixed and the specimen to be scanned, with sub-micrometre repeatability.

As SRAS measurements can be made single-shot, it is possible to scan 'on-the-fly' to increase the speed of acquisition significantly. There is a subtle but important difference here between how EMDA scans and normal 'point-to-point' scanning operates; when wishing to capture data at an array of points $(a_0, a_1, ..., a_n)$, a logical approach is to ask the scanner to travel $a_0 \rightarrow a_1$, followed by $a_1 \rightarrow a_2$ and so forth. EMDA does not function like this. Instead, the linear stage of EMDA is instructed to travel $a_0 \rightarrow a_n$ in one motion; by assuming this journey is completed at a constant velocity, we can infer the position and execute scanning from a temporal trigger. In practice, this 'on-the-fly' scanning is achieved by gating of the trigger pulse from the laser with a signal



Figure 3.6: Capturing an initial noise measurement and subtracting this from subsequent waveforms allows the systematic noise, primarily from the Q-switching of the generation laser, to be excluded. This is particularly important in materials with lower acoustic amplitude. Here acoustic waveforms are plotted with noise traces subtracted.

from the y-axis stage, which reports when the stage has accelerated to the appropriate velocity – a small extra distance before the specimen is specified to allow the y-stage to begin accelerating, such that the stage reaches its travel velocity at exactly the first measurement pixel. When the stage has reached constant velocity, each synchronisation pulse from the laser is then passed to the oscilloscope, triggering the acquisition of a new waveform. When a scan line is complete the stage returns to its initial position, whilst the x-axis is incremented.

By continually moving the specimen, the time to accelerate and decelerate when moving from one measurement point to another is eliminated; in theory, the upper acquisition rate is limited only by the repetition rate of the generation laser. In practice however, specimens are scanned unidirectional, as this allows time for the acquired data to be passed from the oscilloscope to the pc and prevents potential data misalignment due to the linear stage having direction dependent acceleration. Scanning like this approximately halves the acquisition rate compared to bidirectional scanning. The y-axis generally scans at speeds in the range of 20 mms^{-1} to 500 mms^{-1} . The velocity is selected to match the step-size between measurement pixels to the repetition rate of laser - i.e. for a step size of 50 µm at a repetition rate of 2 kHz the stage must travel at 100 mms⁻¹. This is $\ll 0.001\%$ of the SAW velocity, and over a pulse width of 1.5 ns would move the stage by < 1 nm, which is negligibly small compared to the acoustic wavelength. Example experimental data is presented in §3.2.8 and includes illustrative calculation of scan time based on these values.

The oscilloscope is used in segmented mode, allowing a scan line to be buffered into the oscilloscope memory and then downloaded to the experimental PC as the stage position resets to begin a new line. Where a small step size or large specimen is scanned, the time taken to download the data can exceed the reset time of the stage - the following scan line cannot start until the data is completely transferred.

The specimen is mounted on a tip-tilt stage and a z-stage, both of these motorized. This is necessary to ensure the surface remains normal to the incident laser beam across the scan area and the grating image remains in focus. Fig 3.8 shows the impact on the imaged grating pattern due to small perturbations in the z-position.



(a) 50 µm below focal plane. (b) in-focus. (c) 50 µm above focal plane.

Figure 3.7: The vertical position of the specimen is important for ensuring the imaged generation pattern is in focus. As the specimen moves out of focus the generation fringes no longer have sharp edges and the imaged wavelength changes.

3.2.5 Data processing

Fig. 3.8a shows a typical time-domain signal captured by the system; the number of cycles in the waveform correspond to the number of grating fringes imaged onto the sample's surface. The frequency-domain response, fig. 3.8b, is revealed by calculating the fast Fourier transform (FFT) of
the time domain signal. The SAW velocity is then simply the frequency of the highest amplitude. Previously, a Gaussian-fit was applied to the frequency domain data to calculate the peak; however, with the high acquisition rates of the EMDA system, this became a bottleneck.



(a) Time domain waveform showing a wave packet of a fast and slow grain.

(b) Frequency content of SRAS measurements in a fast and slow grain.

Figure 3.8: Typical SRAS waveform and frequency response for a fast and slow wave. It can be seen from (b) the two waves have dominant frequency of 104 MHz and 146 MHz, from an acoustic wavelength of $24 \,\mu\text{m}$ these correspond to velocities of $2500 \,\text{ms}^{-1}$ and $3500 \,\text{ms}^{-1}$, respectively. A velocity range of $1000 \,\text{ms}^{-1}$ is typically found in highly anisotropic metals, such as nickel.

The process of zero-padding to increase the length of the time domain waveform prior to taking the FFT is typically used for two purposes. Firstly, FFT algorithms are generally more efficient when the signal length is a power of two. Secondly, this allows the frequency spectrum to be interpolated, allowing the maximum frequency - and hence velocity - to be found with greater precision. It is important to state plainly that zero-padding has no impact on the upper-frequency limit of the measurement, merely proving greater granularity between frequency bins.

3.2.6 Spatial resolution

The spatial resolution of SRAS is primarily a function of the size of the generation patch - understanding this spatial resolution relative to the size of features in the specimen is an essential part of interpreting the SAW velocity maps. Consider the spatial resolution in the direction of propagation, fig. 3.9 shows the measured velocity when scanning between two dissimilar grains. In the case of a 200 μ m patch, fig. 3.9a, the velocity response shows a transition length of appropriately the size of the patch and the measured velocity transitions from 3400 ms^{-1} to 2800 ms^{-1} beyond the midpoint, 500μ m. In other words, the measured velocity. This suggests the spatial resolution is approximation half the patch size, in this case, 200μ m. As the process of extracting the peak in the frequency domain is non-linear, the transition between grains can have hard edges, suggestive of a higher spatial resolution than reality.

Further complexity is introduced when considering a $100 \,\mu\text{m}$ patch between the same two grains, fig. 3.9b. In the previous example, the $200 \,\mu\text{m}$ patch contains 8 generation fringes, giving a narrowband frequency response - this allows two distinct frequency envelopes to be resolved. Reducing the patch size and making the signal more broadband means the frequency envelopes now overlap and cannot be resolved. Therefore, the measured velocity in the transition region is a weighted mean of the velocity of the two grains.



Figure 3.9: Simulated velocity response at a grain boundary crossing for (a) a 200 μ m generation patch and (b) a 100 μ m patch. The acoustic wavelength is the same in both simulations, thus, reducing the patch size reduces the number of generation fringes, in turn making the signal more broadband.

These two examples lead to the counter-intuitive conclusion that sharper edges are more likely

observed when a larger generation patch is used, despite the system having a lower spatial resolution. These sharp edges are also likely to be observed between grains with greater differences in velocity. This simple simulation assumes the frequency amplitude of the two grains is identical; in reality, the generation efficiency varies from grain to grain. Furthermore, the component generated on the first grain will be steered away from the detection spot at the grain boundary.

Therefore, to improve the spatial resolution, a smaller grating image must be used. This can easily be achieved by using a smaller aperture in the generation optical path, but this has the undesirable effect of reducing the number of cycles in the waveform, in turn broadening the frequency and increasing the likelihood of observing this mean measured velocity around features. Instead, reducing the acoustic wavelength is a more practical avenue to improve the system's spatial resolution.

Currently, the detector has an upper bandwidth of $\sim 300 \text{ MHz}$, which is the primary limiting factor in increasing the working spatial resolution of the system. Taking 3000 ms^{-1} as the average SAW velocity in most relevant metallic specimens, a working frequency of $\sim 300 \text{ MHz}$ corresponds to an acoustic wavelength of 10 µm and spatial resolution of approximately 40 µm.

This has dealt with the

3.2.7 Velocity resolution

A 4 mm thick piece of glass coated with a thin layer of aluminium is an ideal calibration specimen for the EMDA system. The sample is completely isotropic, has no grain boundaries and offers an extremely high reflection of the detection beam. Fig. 3.10 shows the measured velocity across a section of this specimen, the measured velocity ranged from $2860 \,\mathrm{ms^{-1}}$ to $3010 \,\mathrm{ms^{-1}}$. The true velocity of the sample is $2933 \,\mathrm{ms^{-1}}$, having been measured with an earlier - already calibrated -SRAS system.

The histogram of measured SAW velocity is shown in fig. 3.11a, giving a standard deviation of 13.6 ms^{-1} (0.5% variation). Fig. 3.11b shows the impact of averaging on the velocity standard deviation; as the number of averages increases from 1 to 1024, the velocity standard deviation decreases to 0.5 ms^{-1} (0.02% variation).

As with the spatial resolution, the velocity resolution is also influenced by the number of fringes in the grating image. Fig. 3.12a compares the standard deviation of measured SAW velocity as a function of the number of fringes and noise levels. The trivial case of zero noise (an SNR of Inf dB) is not shown on the log-log scale as the standard deviation is 0 ms^{-1} . Each line exhibits a gradient of approximately -1.7, demonstrating that for a fixed wavelength, increasing the number of fringes and in turn, the size of the patch significantly improves the velocity resolution. As the



Figure 3.10: Velocity map from isotropic and homogeneous glass-aluminium sample, variations in the measured velocity arise from system noise.

size of the generation patch is increased, this improvement in the velocity resolution comes at the cost of reducing the spatial resolution of the measurement.

As well as the velocity standard deviation, the bandwidth of signal in the frequency/velocity domain is also a key parameter in determining the velocity resolution of the system. The frequency domain response (shown in subsequent figures as velocity, following $v = f\lambda$) can be considered analogous to a probability distribution. In other words, the position of the maximum is the velocity of greatest probability; a narrowband signal is more immune to noise in the measurement, allowing the true velocity to still be determined. Controlling the bandwidth of the velocity domain response is also particularly important for creating a convex solution space in the inversion procedures that will be introduced in §4.3. Fig. 3.12b plots the full width at half maximum of the signal in the velocity domain as a function of the number of grating fringes. In summary, increasing the number of fringes provides a more narrowband measurement with a lower velocity standard deviation.

3.2.8 Typical data

Drawing together the information of §3.2.6 and §3.2.7, two extreme cases can be envisaged that require significantly different scan parameters. Firstly, the inspection of large grain isotropic materials, such as annealed aluminium, would prioritise high velocity accuracy; conversely, fine-grained nickel, a relatively anisotropic material, would demand spatial resolution precedence. In reality, most materials fall somewhere between these two extremes, and the scan parameters given in



eraging.

locity in 'ideal' sample as a function of temporal averages.

Figure 3.11: Standard deviation of velocity from experimental measurements in the isotropic and homogeneous glass-aluminium sample that is used for calibration. All measurements are for an 8 fringe generation patch.

table 3.1 have proven to be suitable default for most materials, striking an acceptable balance between spatial and velocity resolution.

Table 3.1: User-defined parameters and resultant in-direct parameters, for typical experimental SRAS scanning using the EMDA system.

| Wavelength | $24\mu m$ | Fringes | 8 |
|------------------|-----------------------------|--------------------|----------------------------|
| Signal bandwidth | $\sim\!500\mathrm{ms}^{-1}$ | Spatial resolution | $\sim\!100\mu\mathrm{m}$ |
| Noise std. | 0.2 | SNR | $15\mathrm{dB}$ |
| Averages | 1 | Velocity std. | $\sim\!35\mathrm{ms}^{-1}$ |

Fig. 3.13 shows experimental velocity maps of two large polycrystalline specimens. The contrast in velocity shown in these images indicates the crystal microstructure - regions of homogeneous velocity are the grains. Fig. 3.13a shows a large-grained pure titanium specimen; the grains are well resolved and the specimen has a SAW velocity range of $2500 \,\mathrm{ms}^{-1}$ to $3100 \,\mathrm{ms}^{-1}$. In contrast, fig. 3.13b shows a Ti-6Al-4V specimen with many small grains and features. Whilst some of the



locity a function of number of fringes.

(b) FWHM of velocity/frequency domain response as a function of number of fringes.

Figure 3.12: Simulated data demonstrating the influence of the number of fringes in the grating image on the velocity resolution of the measurement. Increasing the number of fringes reduces both the velocity standard deviation and the FWHM of the frequency/velocity response.

grains remain well resolved, some grains are now below the spatial resolution of the measurement; in this case, the measured velocity is a convolution of the elastic properties of all the grains under the generation patch.

Chapter 7 revisits the speed of image acquisition within the context of an industrially viable system. Therefore, it is worthwhile to calculate the time taken to acquire these images to illustrate the speed of the current EMDA system as a reference point. From §3.2.4 the upper rate of acquisition is equivalent to the repetition rate of the generation laser; therefore, the acquisition time may be estimated by equation 3.1. The estimated and actual scan times for the specimens of fig. 3.13 are given in table 3.2, along with the scan parameters. It is seen that the estimated time, whilst close to reality, underestimates the total scan time in both cases. The true distance travelled by the y-axis in each scan line is slightly greater than the measured length (usually by 10 mm), as the stage needs space to accelerate and decelerate at the start and end of the line, respectively. Additionally, as noted in §3.2.4 when very fine step-sizes are used, the time taken to return to the initial position is not sufficient to finish downloading the data from the oscilloscope, resulting in a brief delay before acquisition of the following scan line commences.



(a) Velocity map from α -Ti.



(b) Velocity map from Ti-6Al-4V specimen.

Figure 3.13: SRAS velocity maps in pure and alloyed titanium.

$$t_{\rm scan} = \underbrace{\left(\frac{x}{x_{\rm step}} \frac{y}{y_{\rm step}}\right)}_{\rm scan \, points} \times \underbrace{\frac{1}{1000}}_{\rm seconds \, per \, point}$$
(3.1)

The measurement of these surface acoustic wave velocities maps is the main aim of the SRAS

| | $\alpha\text{-Ti}$ (fig. 3.13a) | Ti-6Al-4V (fig. 3.13b) |
|--------------------------|---------------------------------|------------------------|
| x length | $13.2\mathrm{mm}$ | $160\mathrm{mm}$ |
| y length | $23.5\mathrm{mm}$ | $96\mathrm{mm}$ |
| $x_{ m step}$ | $20\mu{ m m}$ | $100\mu{ m m}$ |
| $y_{ m step}$ | $20\mu{ m m}$ | $25\mu{ m m}$ |
| Total measurement pixels | 778,800 | 6,144,000 |
| Estimated time | 15 minutes | 1 hour 45 minutes |
| Actual time | 22 minutes | 2 hours 5 minutes |

Table 3.2: Scan time estimation for velocity maps presented in fig. 3.13.

measurement, however two complimentary 'free' datasets are also produced. Firstly, the optical power recorded at the detector can also be plotted as a map, as shown in fig. 3.14a. The resultant image is analogous to an optical micrograph of the sample surface. Such maps are useful for creating data masks to remove the background in velocity maps, as used in fig. 3.13b. In specimens where surface defects are also a concern, as in additive manufacturing, the optical image can give a clear indication of cracking or porosity. In fig. 3.14a the grain structure can also be faintly seen at points, along with surface defects.

The maximum amplitude of the signal in the frequency domain can similarly be plotted, fig. 3.14b. This can either be plotted raw or divided by the measured optical amplitude at each pixel to remove the influence of varying light return. The amplitude of surface acoustic wave measured is not constant across the sample; instead, it is clear the acoustic amplitude is a function of the grain orientation (predicting this amplitude is complex and is discussed in greater detail in §4.2.3).

3.3 Optically rough surfaces

The majority of data presented in this thesis has been captured from specimens that have been polished to a mirror finish, allowing them to be measured with the EMDA system. The EMDA system is limited to smooth surfaces by the knife-edge detector; as the surface becomes less mirrorlike and the reflection diffuse, the KED is unable to observe a change in intensity across the split photodiodes from the incident speckles. As-deposited additive surfaces in additive manufacturing present significant roughness, and there is a challenge in making laser ultrasonics measurements from these. Therefore, rough surface capable systems have required a new detector. Two of the



(a) Optical amplitude.



(b) Acoustic amplitude.

Figure 3.14: Optical and acoustic amplitude datasets are also generated by SRAS and can be visualised as spatial maps. The amplitude of both figures is not calibrated and is instead plotted on a relative scale. The maps shown in this figure correspond to the SAW velocity map presented in fig. 3.13a.

rough surface detector solutions described in 1.7.3 have been trialled with success in derivative f-SRAS systems. These are the speckle knife-edge detector [269] and random quadrature demodu-

lation [270]; velocity maps from wire-arc specimens captured by random quadrature demodulation are presented in Chapter 5. The process of unravelling the speckle pattern by both common detection methods has been detailed in §1.7.3, but from a practical standpoint, using either detector is requires modification from the EMDA system. The following draws attention to several noteworthy differences between these systems and the described EMDA system.

The first practical challenge in rough surface detection is to maximise the amount of collected light and thus increase the sensitivity - therefore, a large optical étendue is required. To achieve this, systems use high numerical aperture lenses, with the objective lens working close to the specimen surface to maximise the cross-section of the reflected cone which can be captured. For example, a 2", f = 100 mm collection lens is used to capture part of the diffuse reflection in the SKED system. Furthermore, this problem is also tackled by increasing the optical power of the detection laser; in both systems, the 420 mW detection laser of EMDA is replaced by a 1.5 W CW laser.

Despite this, neither system is capable of making single-shot measurements, with temporal averaging required to obtain reasonable SNR signals. The leap-forward in acquisition rate of EMDA is due to its ability to scan 'on-the-fly', as described in §3.2.4; hence, it is no surprise that these gains are lost in the rough surface systems.

The spatial resolution of these systems are both lower than EMDA, primarily because of the upper bandwidth limit of these detectors, ranging from 25 MHz to 60 MHz. This has several knockon effects. Firstly, a lower working frequency requires an increase in the acoustic wavelength, $80 \,\mu\text{m} - 250 \,\mu\text{m}$ are common in these rough surface systems. Therefore, to maintain 8 fringes in the grating pattern, the corresponding spatial resolution falls to $320 \,\mu\text{m} - 1 \,\text{mm}$, comparative to a typical spatial resolution of $100 \,\mu\text{m}$ in the EMDA system. Using the generation laser detailed in §3.2.1 to imaging the grating pattern at a lower magnification results in a lower peak power density, the principal factor in the generation of acoustic waves over short time-scales [56]. Thus, to maintain the generation of large amplitude surface waves, the power of the generation laser is increased to one of a peak power of 0.27 MW, over double that of the EMDA generation laser.

Whilst in practice, the frequency limit is due to the bandwidth of the detectors, operating at lower frequencies is also preferable to minimise potential attenuation of the surface wave. The effect of SAW attenuation due to the as-deposited surface in L-PBF is not yet well understood. Attenuation and frequency shift of SAWs has been explored at length by Maradudin and various collaborators [271], primarily through perturbation theory. This body of work concludes with two primary findings, SAWs propagating across a rough surface are subject to 1) attenuation and 2) dispersion, the degree to which these effects occur is a complex dependency on the acoustic wavelength, surface height and spatial frequency. When the wavelength of inspection is large compared to the surface roughness, that attenuation varies as the fifth power of the frequency, but increases to a sixth power relation when the wavelength approaches roughness length. This suggests that for high-frequency Rayleigh wave propagation, the surface roughness may play a critical role [271]. Ruiz and Nagy have previously measured these effects in aluminium specimens by laser ultrasound [272].

3.4 Multi-angle scanning approaches

As should become clear over the course of this work, the ability to propagate SAWs in multiple directions is intrinsic to the function of SRAS. The section briefly describes the current methods of varying the propagation direction of SAWs to capture velocity surfaces.

3.4.1 Mask Rotation

In the EMDA system, the rotation of the propagation direction is achieved in a rather ad-hoc manner. Whilst the generation patch may be automatically rotated, both the placement of the detection spot and the rotation of the KED are done manually. The result is that the light balance on the KED must subsequently be adjusted prior to each scan - clearly, doing this 180 times per specimen to capture velocity surfaces is unrealistic. This has limited previous reporting using the EMDA instrument to 18 directions (0° to 170° at 10° degree increments).

Furthermore, when the amplitude mask is rotated, the imaged wavelength at the specimen varies slightly (on the order of 1%). Experimental results demonstrating this effect are shown in fig. 3.15. This effect is thought to be due to the mask not being mounted level. When measuring subtle variations in the SAW velocity - necessary to calculate the elastic constants and orientation, this effect can begin to dominate and limit the accuracy.

3.4.2 Sample Rotation

An obvious alternative is to keep the optical train static and rotate the specimen. Fig. 3.16 shows velocity surfaces captured in isotropic aluminium and single-crystal nickel by this method.

In this work, several velocity surfaces will be presented with data captured at 1° increments, achieved by placing the specimen on a rotation stage, allowing the angular response to be measured, whilst keeping the optics fixed - ensuring the wavelength stays constant across all propagation directions. This is a more automated method as alignment is only required once per velocity surface, rather than at every propagation direction, allowing velocity surfaces to be captured at 1° increments. This approach is acceptable for homogeneous materials - including the ideal sample



Figure 3.15: Effective wavelength as measured at different rotations of the generation mask.



(a) Velocity surface from isotropic aluminium.

(b) Velocity surface from (001) plane of nickel.

Figure 3.16: Velocity surfaces in captured by rotation of the specimen, the radial velocity scale in aluminium is compressed to highlight variations in the velocity.

and single crystals but proved to be impractical for repeated measurements in polycrystalline materials beyond a few grains, where the centre of rotation had to be precisely aligned with the probed grain.

3.5 Orientation results

Capturing the velocity surface in multiple directions is an essential step in determining the crystallographic orientation or elastic constants of a specimen. The following chapter will provide a detailed review of the processing of determining the crystallographic orientation. Still, it is worth briefly demonstrating such maps and discussing any instrumentation related features before continuing.

Fig. 3.17 compares a SRAS orientation map with the EBSD data acquired in the same nickel superalloy specimen, these figures are orientation plots on an inverse pole figure key, as introduced in §1.4.5. The pixel size in the SRAS data is $25 \times 7 \,\mu\text{m}$ whilst the EBSD data is $5 \times 5 \,\mu\text{m}$, and this introduces two important points. Most datasets have square pixels because the temporal penalty of reducing the step size is equal in both directions. However, in EMDA, because the temporal penalty of adding acquisition points in the *y*-axis in the *y*-axis is less than the *x*-axis, it is often advantageous to scan with a greater number of pixels or a smaller step size in the *y*-axis - this results in somewhat unusual rectangular pixels. Furthermore, the spatial resolution of SRAS is approximately two orders of magnitude below that of EBSD. These factors mean the comparison of orientation data from different modalities is not trivial.

Overall, the SRAS dataset compares well to the EBSD results, both in terms of the grain morphology and orientation determination. The determination of crystallographic orientation in cubic materials, with known elastic constants, is well established and the performance characterised [252]. The orientation determination in hexagonal materials or any materials where the elastic constants are unknown is more complex and will be addressed in Chapter 6.

This chapter has introduced and characterised the EMDA instrument as used throughout this thesis, providing examples of typical data generated by the system. The primary concern of the following chapter is demonstrating how these measurements can be used to solve inverse problems to determine elastic constants and/or crystallographic orientation.



Figure 3.17: Comparison of crystallographic orientation between SRAS (left) and EBSD (right). At each pixel the velocity surface was captured in eighteen directions.

Chapter 4

Numerical apparatus of SRAS measurements

4.1 Introduction

This chapter introduces the numerical procedures which accompany the spatially resolved acoustic spectroscopy (SRAS) experimental technique. In isolation, the surface acoustic wave (SAW) velocity measurements allow the microstructure to be visualised easily. However, to quantitatively understand the physical properties of the specimen, the measured velocity must be inverted. More specifically, this chapter begins by describing in detail the calculation of the theoretical SAW velocity, then proceeds to demonstrate how this information, in combination with experiential measurements, can be used to calculate the crystallographic orientation and elastic constants of measured specimens. The chapter concludes by describing methods of generating synthetic measurement data. Results presented in this chapter are primarily shown to demonstrate how a given procedure works, results themselves, both through experimentation and simulation, will be discussed in greater detail in Chapters 5 and 6. It it worth clearly stating the method for calculating the SAW velocity in SRAS and solving the inverse problem for orientation was the work of Wenqi Li's thesis [5]. The method is reported here for firstly for completeness and secondly to provide a background to the speed of computation improvements achieved in this EngD, one of the key achievements of this project.

4.2 Calculating the SAW velocity - forward problem

Fig 4.1 conveys the link between the three factors (i) acoustic velocity (ii) elastic constants and (iii) crystallographic orientation, whereby normally two are needed to calculate the third. These phenomena are exploited throughout our engineered world, from medical imaging [273] to seismology [274]. As indicated in the figure, the measurement of the SAW velocity is the only phenomenon directly measurable by acoustic methods, necessitating solving inverse problems to reveal the useful material properties of crystallographic orientation and elasticity. From this, three problems can be defined, this chapter will describe appropriate solutions to each, in turn.



Figure 4.1: The measured SAW velocity is a function of the elastic constants and crystallographic orientation of the specimen. The elastic constants and crystallographic orientation are informative material properties however they cannot be measured directly by acoustic techniques, thus to determine these properties inverse problems must be solved.

In this work, the experimental measurement of surface acoustic wave velocity is used to determine the physical properties of anisotropic materials, including additive manufacturing specimens. Therefore, it is imperative to provide a comprehensive introduction to the theory of acoustic waves in anisotropic media.

4.2.1 The wave equation

Stress, strain and Hooke's law were introduced in §1.4; in summary, Hooke's law, equation 1.4 (reprinted below), provides a relation between strain in any direction and the resultant direction-

dependent stress experienced. Thermoelastic excitation, as used to create acoustic waves, generates a thermal stress field. Therefore, from Hooke's law, we can see this stress field will induce a displacement (strain) in the material, resulting in a new stress field in the surrounding material and so forth. This stress-strain relation describes the propagation of an acoustic wave through a medium.

$$\sigma_{ij} = \mathcal{C}_{ijkl} \epsilon_{kl}$$

The wave equation will now be derived by also considering Newton's second law, F = ma. Consider an infinitesimally small volume of dimensions (dx, dy, dz), with the sum of all external forces acting on the object $\Sigma F_x = 0$ in the x direction

$$ma_{x} = \Sigma F_{x} = 0$$

$$V = dxdydz$$

$$m\frac{\partial^{2}u_{j}}{\partial t^{2}} = V\frac{\partial\sigma_{xx}}{\partial x} + \frac{\partial\sigma_{xy}}{\partial y} + \frac{\partial\sigma_{xz}}{\partial z}$$

$$\rho = \frac{m}{V}$$

$$\rho\frac{\partial^{2}u_{j}}{\partial t^{2}} = \frac{\partial\sigma_{ij}}{\partial x_{j}}$$

$$(4.1)$$

Replacing the second term, defined by stress, with the definition of strain and elasticity tensor from Hooke's law gives the final anisotropic wave equation

$$\rho \frac{\partial^2 u_j}{\partial t^2} = C_{ijkl} \frac{\partial^2 u_k}{\partial x_i \partial x_l} \text{ for } i, j, k, l = 1, 2, 3.$$
(4.2)

Where ρ is the material density, u_i the displacement in the x_i axis and C_{ijkl} is the materials fourth-rank elasticity stiffness tensor. This formalises the relationship between crystallographic orientation, elasticity and SAW propagation.

4.2.2 Bulk wave velocities

The wave equation will now be solved to calculate acoustic velocities as a function of the elasticity tensor. It is first illustrative to detail the more straightforward process of calculating bulk wave velocities in anisotropic media before extending this to the more complex surface wave case.

Waves along the *x*-axis

Assuming the medium is infinite in all directions, one possible solution to the wave equation, equation 4.2, is

$$u = \alpha e^{jk(\boldsymbol{l}\boldsymbol{x} - vt)} \tag{4.3}$$

with direction cosines $\mathbf{l} = (l_x, l_y, l_z)^T$ of the propagation vector, v is the wave velocity, k is the wavenumber (such that $k = 2\pi/\lambda$) and α is the eigenvector of displacement.

Substituting equation 4.3 into equation 4.2 gives equation 4.4. From equation 4.4, particle displacement and phase velocity can be calculated for arbitrary directions, when C_{ijkl} is known.

$$(C_{ijkl}\boldsymbol{l} - \rho v^2 \delta_{ik})\alpha_k = 0 \text{ for } i, j, k, l = 1, 2, 3$$

$$\text{where } \delta_{ik} \text{ is the Kronecker delta } \delta_{ik} = \begin{cases} 0, & i \neq k \\ 1, & i = k \end{cases}$$

$$(4.4)$$

Equation 4.4 is the Christoffel equation and may be written as

$$\Gamma \begin{pmatrix} u_x \\ u_y \\ u_z \end{pmatrix} = 0 \tag{4.5}$$

where Γ , the Christoffel matrix, is given by

$$\mathbf{\Gamma} = \begin{pmatrix} \Gamma_{11} - \rho v^2 & \Gamma_{12} & \Gamma_{13} \\ \Gamma_{12} & \Gamma_{22} - \rho v^2 & \Gamma_{23} \\ \Gamma_{13} & \Gamma_{23} & \Gamma_{33} - \rho v^2 \end{pmatrix}$$
(4.6)

The elements of Γ are written in their expanded form as

$$\Gamma_{11} = l_x^2 C_{11} + l_y^2 C_{66} + l_z^2 C_{55} + 2l_x l_y C_{16} + 2l_x l_z C_{15} + 2l_y l_z C_{56}
\Gamma_{12} = l_x^2 C_{16} + l_y^2 C_{26} + l_z^2 C_{45} + l_x l_y (C_{12} + C_{66}) + l_x l_z (C_{14} + c_{56}) + l_y l_z (C_{25} + C_{46})
\Gamma_{13} = l_x^2 C_{15} + l_y^2 C_{46} + l_z^2 C_{35} + l_x l_y (C_{14} + C_{56}) + l_x l_z (C_{13} + C_{55}) + l_y l_z (C_{36} + C_{45})
\Gamma_{22} = l_x^2 C_{66} + l_y^2 C_{22} + l_z^2 C_{44} + 2l_x l_y C_{26} + 2l_x l_z C_{46} + 2l_y l_z C_{24}
\Gamma_{23} = l_x^2 C_{56} + l_y^2 C_{24} + l_z^2 C_{34} + l_x l_y (C_{25} + C_{46}) + l_x l_z (C_{36} + C_{45}) + l_y l_z (C_{44} + C_{23})
\Gamma_{33} = l_x^2 C_{55} + l_y^2 C_{44} + l_z^2 C_{33} + 2l_x l_y C_{45} + 2l_x l_z C_{35} + 2l_y l_z C_{34}$$
(4.7)

For non-trivial solutions of \boldsymbol{u} the determinant of equation 4.6 must equal zero.

In equation 4.7 Γ is dependent on the value of elements in the elasticity matrix. Whilst the tensor contains 36 terms, the material's symmetry reduces the number of independent terms. For example, the elasticity matrix of cubic materials is given by equation 4.8, reducing the number of independent terms to 3.

$$C_{c} = \begin{bmatrix} C_{11} & C_{12} & C_{12} & 0 & 0 & 0\\ C_{12} & C_{11} & C_{12} & 0 & 0 & 0\\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0\\ 0 & 0 & 0 & C_{44} & 0 & 0\\ 0 & 0 & 0 & 0 & C_{44} & 0\\ 0 & 0 & 0 & 0 & 0 & C_{44} \end{bmatrix}$$
(4.8)

In the case of a wave propagating along the x-axis, $\boldsymbol{l} = (1,0,0)^T$, the Christoffel equation reduces to

$$\begin{pmatrix} C_{11} - \rho v^2 & 0 & 0\\ 0 & C_{44} - \rho v^2 & 0\\ 0 & 0 & C_{44} - \rho v^2 \end{pmatrix}$$
(4.9)

for cubic materials, note this is independent of the value of C_{12} in this direction. The determinant equality, $\Delta\Gamma$, may then be written as

$$\Delta\Gamma = \underbrace{(C_{11} - \rho v^2)}_{A} \underbrace{(C_{44} - \rho v^2)}_{B} \underbrace{(C_{44} - \rho v^2)}_{C}$$
(4.10)

The three bulk wave velocities are then given by

$$A = 0 : v_L = \sqrt{\frac{C_{11}}{\rho}} B = 0 : v_{T,1} = \sqrt{\frac{C_{44}}{\rho}} C = 0 : v_{T,2} = \sqrt{\frac{C_{44}}{\rho}}$$
(4.11)

in this case $v_{T,1} = v_{T,2}$ meaning the fast and slow transverse waves have the same velocity along the x-axis. The symmetry of a cubic structure means $\boldsymbol{l} = (1,0,0)^T = (0,1,0)^T = (0,0,1)^T$.

Waves across the x - y plane

Without a great deal of adjustment, this formalisation can be generalised to allow the calculation of the bulk velocities in any direction on the plane x - y. The direction vector l can be written as a 2D direction cosine $l = (\cos \theta, \sin \theta, 0)$, hence the Chrisstoffel equation becomes

$$\begin{pmatrix} \cos^2 \theta C_{11} + \sin^2 \theta C_{44} - \rho v^2 & \cos \theta \sin \theta (C_{12} + C_{44}) & 0 \\ \cos \theta \sin \theta (C_{12} + C_{44}) & \cos^2 \theta C_{44} + \sin^2 \theta C_{11} - \rho v^2 & 0 \\ 0 & 0 & (\cos^2 \theta + \sin^2 \theta) C_{44} - \rho v^2 \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix} = 0$$

$$(4.12)$$

again the solutions of this equation are non-trivial only when the determinant is zero. Usefully, equation 4.12 is in the form of the eigenvalue equation, allowing the more compact form $(\mathbf{A} = \lambda I)\mathbf{u} = 0$, where $\lambda = \rho v^2$. Of course, the calculation of the eigenvalues of \mathbf{A} is somewhat tedious by-hand but found with ease by any modern software computation package. The characteristic polynomial of \mathbf{A} is then given by

$$-\lambda^{3} + \lambda^{2} \operatorname{Tr}(\boldsymbol{A}) + \frac{1}{2}\lambda(\operatorname{Tr}(\boldsymbol{A})^{2} - \operatorname{Tr}(\boldsymbol{A}^{2})) + \det \boldsymbol{A}$$
(4.13)

Where Tr is the trace of the matrix **A**. Therefore, the wave velocities are found by $v = \sqrt{\frac{\lambda_{1,2,3}}{\rho}}$, giving the final equations

$$v_L = \sqrt{\frac{C_{11} + C_{44} + \sqrt{(C_{11} - C_{44})^2 \cos^2 2\theta + (C_{12} + C_{44})^2 \sin^2 2\theta}{2\rho}}$$
(4.14)

$$v_{T,1} = \sqrt{\frac{C_{11} + C_{44} - \sqrt{(C_{11} - C_{44})^2 \cos^2 2\theta + (C_{12} + C_{44})^2 \sin^2 2\theta}{2\rho}}$$
(4.15)

$$v_{T,2} = \sqrt{\frac{\mathcal{C}_{44}}{\rho}} \tag{4.16}$$

This set of equations provides a method to calculate the velocities of the bulk acoustic modes on the x-y plane. These can be extended for the calculation in an arbitrary crystal direction by expressing l as the three-dimensional direction cosine, such that $l = al_x + bl_y + cl_z$. However, another approach is first to rotate the elasticity matrix, C, such that the crystallographic plane of interest is aligned with the plane x - y. The rotation of the fourth-rank tensor is given by equation 4.17 - the rotation must be performed on the original elasticity tensor and not the second-rank Voigt representation.

$$C_{\text{rotated}} = CR_{rot}^T \tag{4.17}$$

Where R is the rotation matrix, given by

$$R_{rot} = \begin{pmatrix} a^2(1-\cos\theta) + \cos\theta & ab(1-\cos\theta) - c\sin\theta & ac(1-\cos\theta) + b\sin\theta \\ ab(1-\cos\theta) + c\sin\theta & b^2(1-\cos\theta) + \cos\theta & bc(1-\cos\theta) - a\sin\theta \\ ac(1-\cos\theta) - b\sin\theta & bc(1-\cos\theta) + a\cos\theta & c^2(1-\cos\theta) + \cos\theta \end{pmatrix}$$
(4.18)

where $R_a = [abc]$ is the rotation axis between the (001) plane and the desired plane and θ is the angle between the plane normals. For example, to rotate from the (001) to (011) plane the normals are a = [0 0 1] and b = [0 1 1], hence

$$\theta = \arccos \frac{a \cdot b}{|a| \, |b|} = 45^{\circ}$$

$$R_a = [001] \times [011] = [-100]$$

$$R_{rot} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$
(4.19)

From this the three bulk wave velocities may be calculated on any plane and in any direction. Considering nickel as an example, the elastic constants have been reported as $C_{11} = 247 \text{ GPa}$, $C_{12} = 153 \text{ GPa}$, and $C_{44} = 122 \text{ GPa}$, with density of $\rho = 8912 \text{ kgm}^{-3}$ [275]. The three bulk wave velocities on the planes (001) and (111) are shown in fig. 4.2.



Figure 4.2: bulk wave velocities on two planes of nickel, both the fast shear wave $v_{T,1}$ and the longitudinal wave v_L are isotropic on one of these planes.

The (001) plane exhibits 90° symmetry, derived from the terms $\cos^2 \theta$ and $\sin^2 \theta$ terms in equations 4.14 and 4.15. The fast transverse wave, $v_{T,2}$, is isotropic on this plane as equations 4.16 has no angular dependence. On the (111) plane the v_L velocity is isotropic, whilst the transverse waves exhibit six-fold symmetry.

4.2.3 Surface acoustic wave velocity

Figure 4.3: Definition of SAW axis.

The surface acoustic wave can be thought of as a vector sum of a longitudinal wave and a shear wave, and hence it propagates with elliptical motion. From fig. 4.3, considering a surface wave propagating in the x-direction on the x - y plane, the wave has motion in both the x and z directions, independent of the y-direction. Unlike the bulk waves described above, as SAWs are surface-bound, the amplitude decays into the material, and thus the component in the z-direction must also be considered. A solution to this is given by equation 4.20

$$u_i = \alpha_i \underbrace{e^{jkl_z z}}_{\text{amplitude wave propogation}} \underbrace{e^{jk(l_x x + l_y y - vt)}}_{\text{wave propogation}}$$
(4.20)

Physically, this describes a wave propagating in the x - y plane with direction $l_x \hat{i} + l_y \hat{j}$ and a z dependence determined by the value of l_z travelling with phase velocity v. Conceptually, the z dependence may be considered as a part of the amplitude term. As such, the propagation vector is assumed to be parallel to the free surface regardless of a real component in l_z . With known direction cosines, the Christoffel equation in the surface wave case is then

$$\boldsymbol{\Gamma}(l_x, l_y, l_z, v)\boldsymbol{u} = 0 \tag{4.21}$$

where the Christoffel matrix is again given by equation 4.6, with the members given by

$$\Gamma_{11} = \cos^{2}\theta C_{11} + \sin^{2}\theta C_{66} + l_{z}^{2}C_{55} + 2\cos\theta\sin\theta C_{16} + 2\cos\theta l_{z}C_{15} + 2\sin\theta l_{z}C_{56}$$

$$\Gamma_{12} = \cos^{2}\theta c_{16} + \sin^{2}\theta C_{26} + l_{z}^{2}C_{45} + \cos\theta\sin\theta (C_{12} + C_{66}) + \cos\theta l_{z}(C_{14} + C_{56}) + \sin\theta l_{z}(C_{25} + C_{46})$$

$$\Gamma_{13} = \cos^{2}\theta C_{15} + \sin^{2}\theta C_{46} + l_{z}^{2}C_{35} + \cos\theta\sin\theta (C_{14} + C_{56}) + \cos\theta l_{z}(C_{13} + C_{55}) + \sin\theta l_{z}(C_{36} + C_{45})$$

$$\Gamma_{22} = \cos^{2}\theta C_{66} + \sin^{2}\theta C_{22} + l_{z}^{2}C_{44} + 2\cos\theta\sin\theta C_{26} + 2\cos\theta l_{z}C_{46} + 2\sin\theta l_{z}C_{24}$$

$$\Gamma_{23} = \cos^{2}\theta C_{56} + \sin^{2}\theta C_{24} + l_{z}^{2}C_{43} + \cos\theta\sin\theta (C_{25} + C_{46}) + \cos\theta l_{z}(C_{36} + C_{45}) + \sin\theta l_{z}(C_{44} + C_{23})$$

$$\Gamma_{33} = \cos^{2}\theta C_{55} + \sin^{2}\theta C_{44} + l_{z}^{2}C_{33} + 2\cos\theta\sin\theta C_{45} + 2\cos\theta l_{z}C_{35} + 2\sin\theta l_{z}C_{34}$$

$$(4.22)$$

Boundary conditions

Again, for non-trivial solutions $\Delta \Gamma = 0$. Writing this as a sixth-degree polynomial in l_z allows the roots to be determined for any value of v to satisfy the anisotropic wave equation. As the direction cosines in the x - y plane are real, there exist three pairs of complex conjugate roots for each velocity. However, the roots must also satisfy the boundary conditions of surface waves, table 4.1.

Table 4.1: Boundary conditions which must be satisfied for a SAW propagating on an semi-infinite half-space, with the free surface at z = 0.

| $z = -\infty$ | z = 0 |
|---------------|-----------------|
| $u_x = 0$ | $\tau_{xz} = 0$ |
| $u_y = 0$ | $\tau_{yz} = 0$ |
| $u_z = 0$ | $\sigma_z = 0$ |

The roots found in the positive half of the complex plane describe waves of amplitude tending toward infinity as $z \to -\infty$, this fails to satisfy the boundary condition of decaying amplitude with depth. Accordingly, only the three lower half roots satisfying this condition need to be considered. Therefore, the solution of the wave equation, equation 4.20, may be re-framed as the sum of contributions from the three roots, equation 4.23

$$u_i = \sum_{n=1}^{3} C_n \alpha_i^{\langle n \rangle} e^{j(l_x x + l_y y + l_z^{\langle n \rangle} z - vt)}$$
(4.23)

where $\alpha_i^{\langle n \rangle}$ is the eigenvector corresponding to the root $l_z^{\langle n \rangle}$. A secondary boundary constitution of zero traction at the free surface must also be satisfied, table 4.1. The three weighting factors C_n must now be found, such that the boundary conditions are satisfied across the free surface.

A further set of equations can be derived by substituting the solution, equation 4.23, into the boundary conditions. We may call this the Rayleigh boundary value matrix, R. Again the determinant of this set of equations is zero for a non-trivial solution, $\Delta R = 0$. The members of Rmay be found by equation 4.24.

$$r_{nm} = \mathcal{C}_{m3kl} \alpha_k^{\langle n \rangle} l_l^{\langle n \rangle} \tag{4.24}$$

Therefore, a velocity that can satisfy these conditions is a valid surface wave solution. In the case of isotropic symmetry or certain high symmetry directions in anisotropic crystals, it is possible

to write an explicit equation for the phase velocity. However, in general, the complexity of these non-linear equations make it impractical to directly calculate the phase velocity. Rather it is easier to determine the value of ΔR for a range of velocities.

When searching numerically, the condition $\Delta R = 0$ shall never be satisfied by the RSAW. In addition, pseudo-surface waves (PSAW) can exist with a non-zero determinant, which can only fulfil the boundary conditions by shedding energy in the form of a bulk wave that leaks into the solid. Thus, PSAWs can only propagate with attenuation. Nevertheless, along certain directions on specific planes, these waves are observed in preference to true surface waves. Therefore, it is necessary to look for velocities that minimise $|\Delta|$, where each minimum corresponds to a separate wave mode.

The process of conducting an iterative search across the velocity range was first demonstrated by Farnell [266], and the approach described thus far follows the procedure outlined in his seminal text. A comprehensive review of SAW velocity calculation is provided in §2.5.2.

Calculating the displacement

The framework detailed thus far gives a relatively simple approach to calculating the SAW velocities of all possible wave modes in an arbitrary direction. However, simply calculating the velocity of each wave mode lacks any indication of how 'measurable' these modes are experimentally. The dominant mode at each direction may be found by calculating the relative displacement of each and selecting velocity corresponding to the mode of maximum displacement. This then provides a single 'measurable' velocity surface.

Having determined a velocity that satisfies the surface wave boundary conditions, the partial wave solutions to equation 4.23 can be found.

The knife-edge style detector is sensitive to the gradient of the surface and is thus primarily sensitive to the out-of-plane displacement, u_z . However, considering the elliptical motion of SAWs, it can be seen the wave motion is also driven by the displacement in the longitudinal direction, u_x . Hence, the measured wave displacement may be considered to be the vector sum of displacements in the x and z directions, given by

$$u_{v,\text{total}} = \sqrt{u_x^2 + u_z^2} \tag{4.25}$$

where $u_{v,\text{total}}$ is the total displacement component to which the detector is sensitive. For a SAW propagating along the x-axis in the plane x - y, as the propagation direction is rotated towards the y-axis, the longitudinal component becomes a vector sum of contributions in x and y-axes. The transverse component of displacement is omitted.

Summary of numerical method

This generalised numerical search procured for calculation of the SAW phase velocity, as a function of crystallographic orientation and elastic constants, is summarised in the flowchart of fig. 4.4. To perform the computation of SAW velocities the following information must be defined, step 'A'.

- the material density, ρ .
- the crystal symmetry of the material, symmetry of C_{ijkl}.
- the elastic constants, members of C_{ijkl}.
- the velocity search step size, v_{step} .
- the range of the orientation space over which the saw velocity is to be calculated.

The density of a material can be found in literature, and knowledge of the material's crystal symmetry is implicit with knowledge of the elastic constants of the material. Typically, it is not necessary to define the range of velocities over which to search, instead upper and lower bounds may be defined relative to the second slow shear wave velocity, $V_{T,2}$, equation 4.16. To find the minima corresponding to both Rayleigh and pseudo-surface waves the search range typically spans the range $0.5v_{t,2} \leq v \leq 1.1v_{t,2}$. In most circumstances, the full orientation space is to be searched, and we may also neglect to define this term. It is therefore true to say in most cases, only the elastic constants, density and velocity search resolution need to be defined, as reflected in step 'A'.

In step 'B' the defined elasticity tensor C_{ijkl} is rotated from (001) to the pane of interest. This repeats for every modelled plane, forming the C_{ijkl} - loop.

Next, step 'C', the propagation direction of the SAW on the plane x - y is rotated between 0° and 180° , at one degree increments, by varying the direction cosines l_x and l_y . This forms the θ -loop.

The previous steps allow the propagation direction-specific Christoffel equation, equation 4.21, to be defined. In step 'E', the roots of the determinant in l_z may now be found from $\Delta\Gamma = 0$. Having determined the three lower-half space roots, the corresponding eigenvectors, α , are then found from 4.23 in step 'F'.

Having determined values of l_z and α for the trial velocity, the determinant value of the boundary condition ΔR may now be found, step 'G'. The *v*-loop then returns to step 'D' and increments the search velocity.

Subsequent to constructing the complete determinant curve, the minima and their corresponding velocities are found in order, step 'H'. For each minimum, the measurable displacement is



Figure 4.4: Flowchart outlining the process of numerically calculating the SAW phase velocity as a function of elastic constants and crystallographic orientation.

then calculated in step 'I', having determined the relative displacement of all modes, the 'dominant' mound is found, step 'J'. Finally, steps 'B' through 'J' repeat for each plane and associated propagation direction required, until the full SAW velocity forward model is calculated, step 'K'. Nickel

bic materials.



Figure 4.5: The search range of the forward model for cubic materials is defined in Miller indices, covering the space between the planes (001) to (101)/(011) to (111). The corresponding Euler angles are given by the definition in (b).

As in the bulk wave case, numerical examples for nickel shall now be provided. Fig. 4.5b shows the unit sphere where the area shaded in red contains the normals of the planes calculated by the forward model. Due to the crystal symmetry, all possible velocity surfaces are contained within the shaded region. The corresponding Euler angle rotations, in reference to Miller indices, are found by the definition in fig. 4.5a for a cubic crystal.

Fig. 4.6 shows the value of $|\Delta R|$ for the (001) plane in nickel. Four distinct modes can exist simultaneously at certain propagation directions: RSAW, PSAW and the fast and slow transverse waves. The 90°/four-fold symmetry of the (001) plane can be seen. The SAW velocity between 0° to 30° and then 60° to 90° are of typical RSAW type. However, between 30° to 60°, the wave velocity exceeds the slow transverse wave velocity, $v_{T,1}$, the wave is therefore classified as a supersonic PSAW. The PSAW will radiate energy into the bulk of the specimen and attenuate.

Fig. 4.8 is a hybrid plot, first showing the velocity of the two primary wave modes on the (001)



Figure 4.6: Value of the Rayleigh determinant, $|\Delta R|$, on the (001) plane on nickel. The position 0° is equivalent to [100] and 90° is equivalent to [010], the data over this range then repeats in the subsequent quarters due to the four-fold symmetry. The shear wave velocities, marked in the second quarter, are the same as those shown in fig. 4.2a. Minima correspond to velocities which satisfy the boundary conditions. Note the colourscale is inverted from normal use, such that 'bright' regions (yellow) correspond to minima.

plane and then plots the relative displacement of both modes across the plane. Between 24° and 66° the PSAW mode dominates, giving a discontinuity in the calculated velocity, with a difference of $\sim 220 \,\mathrm{ms}^{-1}$ between the RSAW and PSAW modes. Around 22° the amplitudes of both have the same amplitude, and in such cases, it is unclear which wave mode will be measured - this will be discussed further in §4.2.5.

Fig 4.9 shows the calculated SAW velocity for the three principal planes of nickel. Two wave modes are observed on the (001) plane (fig. 4.9a), with the aforementioned mode transition visible. In contrast, the calculated SAW velocity is of RSAW on the (101) plane (fig. 4.9b), however between 0° and 67° has a 4th root which satisfies the boundary conditions. In the (111) plane (fig. 4.9b). The SAW velocity ranges from approximately 2300 ms^{-1} to 3100 ms^{-1} across the whole orientation space for nickel.



(a) Vertical displacement. (b) Longitudinal displacement. (c) Transverse displacement.

Figure 4.7: The individual relative displacement components, of the two surfaces modes, on the (001) plane in nickel. The transverse component does not contribute to the 'measurable' displacement.

Other cubic materials

The velocity surfaces of two further cubic materials: tungsten and potassium chloride (KCl) provide useful examples to comment upon the general properties of cubic surface wave velocities.

Fig. 4.10 plots a useful method of visualising the elastic constants of cubic materials by normalising C_{11} and C_{12} to the value of C_{44} . Diagonal lines indicate the anisotropy ratio of the material. It can be seen that nickel ($\eta = 2.60$) and tungsten ($\eta = 1.23$) are examples of engineering materials of high and low anisotropy, respectively. KCl may also be described as an anisotropic material ($\eta = 0.38$), but is of 'inverse' isotropy to nickel.

The SAW velocity range in tungsten is just 4 ms^{-1} around 2647 ms^{-1} , Fig 4.11(a), the forward model is solved with $v_{\text{step}} = 5 \text{ mms}^{-1}$ in order to adequately resolve the anisotropy of the velocity surfaces. For comparison, the phase velocity calculated by Viktorv's approximation is 2645 ms^{-1} . This is in good agreement with the forward model, confirming materials of $\eta \approx 1$ may be treated as isotropic. In all directions, the dominant mode is an RSAW.

KCl exhibits a velocity range of around 350 ms^{-1} ; however, SAW velocity on the (001) plane in KCl is shown to be near independent of the propagation angle, which is counter-intuitive given the large anisotropy of the crystal. However, the slow transverse wave on this plane does not satisfy the boundary conditions for a surface wave - it is not possible for the SAW to degenerate into a PSAW. In fact, only in materials with $\eta >> 1$ does the surface wave degenerate into the pseudo surface wave on the (001) plane. On the (101) plane, the dominant surface wave transitions from



Figure 4.8: Hybrid plot, the top section shows the relative displacement of the RSAW (red) and PSAW(blue) on the (001) plane of nickel. The lower section plots the velocity of both modes, with the dominant mode across the plane indicated by *****. Ellipses are illustrative of the particle motion at the given propagation angle.

an RSAW to a PSAW around 70° , unlike in nickel, where the wave remains of Rayleigh-type across the plane.

Clearly, the fidelity of SRAS measurements depends on the variation in velocity between crystal ordinations; small variations in tungsten would require high velocity accuracy to allow comment on the microstructure. Therefore, the anisotropy ratio is a helpful tool for informing the required experimental accuracy.

Titanium

The elasticity matrix of hexagonal materials is given by equation 4.26, giving five independent elastic constant terms. For pure α -Ti the elastic constants are reported as $C_{11} = 160$ GPa, $C_{12} = 90$ GPa, $C_{13} = 66$ GPa, $C_{33} = 181$ GPa, and $C_{44} = 46.5$ GPa, with density of $\rho = 4506$ kgm⁻³ [226]. The nomenclature and Euler-Bunge angle rotation of hexagonal crystals is defined in fig. 4.12a.



Figure 4.9: Calculated SAW velocities on the three principal planes of nickel between 0° and 90° , directions are also noted in reference to Miller indices. The three planes are chosen for easiest visualisation on the unit cells (upper) from a single view point, plane families, such as $\{001\}$, also exhibit these velocity surfaces.

$$C_{h} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0\\ C_{12} & C_{11} & C_{13} & 0 & 0 & 0\\ C_{13} & C_{13} & C_{33} & 0 & 0 & 0\\ 0 & 0 & 0 & C_{44} & 0 & 0\\ 0 & 0 & 0 & 0 & C_{44} & 0\\ 0 & 0 & 0 & 0 & 0 & C_{66} \end{bmatrix}$$
(4.26)

where $C_{66} = (C_{11} - C_{12})/2$.

Titanium exhibits transverse isotropy, meaning the material has the same proprieties through a plane normal to the c-axis but different properties at any titled plane. Fig. 4.12b illustrates this issue, in reference to the unit sphere, the area searched by the forward model does not depend on the angle of rotation around the z-axis. The area searched corresponds to the space $0 \leq \Phi \leq 90^{\circ}$, where any value of ϕ_2 generates the same response.

Consequently, this leads to a few cases where elastic properties, and subsequently SAW velocity, do not vary with the crystal orientation. For example, the velocity surface seen on the basal plane



Figure 4.10: Colourmap of constant relative velocity (v_R/v_s) for surfaces waves propagating along the [100] axis on the (001) plane of cubic crystals. Lines of anisotropy are also indicated, where $\eta = 1$ is an isotropic material, for $C_{12} > 1$ the constants below contours are aligned with the lines of anisotropy. Constant velocities are calculated from the explit form of the surface wave velocity in this direction, as derived by Stoneley [261]. The space $C_{12} > C_{11}$ is not stable as the eigenvalues of the elasticity matrix must be non-negative [228].

 $\{0001\}$ is isotropic - rotation of the velocity surface does not change the SAW velocity, such as on the basal plane. Fig. 4.13 plots the determinant value across the basal plane of titanium, where all wave velocities, including bulk modes, are seen to be isotropic. RSAWs are not found on the basal plane in titanium; instead, the dominant surface wave mode on this specific crystallographic plane is a PSAW, often called the supersonic wave, as the velocity exceeds that of the slow transverse wave [276]. The practical impact of this is that linear acoustic techniques cannot determine ϕ_2 .

Fig. 4.14 plots the dominant velocity on the planes $\Phi = 0^{\circ}45^{\circ}, 90^{\circ}$, in titanium. As the cutplane tilts away from the basal plane ($\Phi \rightarrow 90^{\circ}$), the dominant supersonic mode gives way to the subsonic RSAW, which first appears at $\Phi = 28^{\circ}$. This Rayleigh wave contributes to an increasing proportion of the velocity surface until the supersonic mode is no longer measured, as only the Rayleigh wave is seen on the prism plane, fig. 4.14c. The acoustic velocity surfaces of the principal prism planes $\{10\overline{1}0\}$ and $\{11\overline{2}0\}$ are identical, and therefore only one is calculated and shown.

The velocity of titanium ranges from approximately $2600 \,\mathrm{ms}^{-1}$ to $3200 \,\mathrm{ms}^{-1}$. In general, a higher velocity indicates an orientation closer to the basal plane as the dominant SAW velocity



Figure 4.11: The anistropic ratio of cubic materials has a consequential impact on the velocity surfaces. Tungsten is near isotropic and thus exhibits a small SAW velocity range across the orientation space.

slows as $\Phi \rightarrow 90^{\circ}$. This idea has previously been exploited to produce velocity vector maps, indicting the approximate c-axis position from two orthogonal SAW velocity measurements [277].

The area calculated by the forward model in hexagonal materials is between $0^{\circ} \ge \Phi \ge 90^{\circ}$; therefore, 91 planes are calculated at a 1° increment, in comparison ~ 2700 planes would need to be calculated in cubic materials to achieve an equivalent angular resolution. In reality, the cubic models are calculated at larger steps as a compromise between computation time and resolution. With this in mind, now is a sensible juncture to comment on the computational efficiency of the forward model.

4.2.4 Calculation time

The method detailed allows the calculation of the SAW phase velocity - on an arbitrary plane and direction - when the elastic constants are known. When working with established materials where the elastic constants are well defined and there is a limited range of materials to be evaluated, computation speed is inconsequential as the forward model need only be computed once per



(a) Definition of Euler-Bunge angles in hexagonal materials. The third angle, ϕ_2 , does not vary the SAW velocity = planes (11 $\overline{2}0$) and (1010) are equivalent.



(b) Unit sphere, shaded region contains the plane normals calculated by the forward model. Note the azimuth angle is infinitesimal but is enlarged here for clarity.

Figure 4.12: Definition of Euler angles and search space of the forward model, for hexagonal crystal system. Only the planes $0^{\circ} \ge \Phi \ge 90^{\circ}$ need to be calculated as hexagonal materials are transversely isotropic.

material.

However, with additive manufacturing, the range of materials and elastic constants grows rapidly, and by extension, the computation speed of computing of the forward model becomes a critical factor. This is particularly true if both unknown elastic constants and orientations are to be determined by inversion - one of the primary aims of this work.

The values of table 4.2 provide some essential context to the scale of the full inverse problem and expected computation time. These parameters would be typical of such optimisation problems [239]. This requires the calculation of 121,000 velocity surfaces, and the original method would take around 450 seconds to calculate the velocity in one plane of nickel, thus the complete calculation would take approximately 560 days. Therefore, the speed of the forward model was identified as the first obstruction to solving the full inverse problem.

Constructing the determinant curve $|\Delta \mathbf{R}|$ is the most time consuming task of the forward



Figure 4.13: Value of determinant on the isotropic basal plane of titanium. Minima correspond to velocities which satisfy the boundary conditions. In the case of this plane, the faster PSAW dominates - for titanium this continues to be the case until $\Phi \approx 25^{\circ}$, when the RSAW becomes measurable near $\phi_1 = 90^{\circ}$. Note the colourscale is inverted from normal use, such that 'bright' regions (yellow) correspond to minima.

| Parameter | resolution |
|------------|---------------------------------------|
| v_{step} | $2\mathrm{ms}^{-1}$ |
| (hkl) | 0.1 to (011) to (111) |
| C_{11} | $2\mathrm{GPa}$ over $20\mathrm{GPa}$ |
| C_{12} | $2\mathrm{GPa}$ over $20\mathrm{GPa}$ |
| C_{44} | $2\mathrm{GPa}$ over $20\mathrm{GPa}$ |

Table 4.2: Forward model search resolution

model, the total computation time is strongly dependent on the velocity range (v_{max} and v_{min}) and the velocity search step (v_{step}), as demonstrated in fig. 4.15a. Following the process outlined in



Figure 4.14: Calculated SAW velocities on the three principal planes of titanium. The SAW velocity surface of the basal plane is supersonic and isotropic. As the cut plane tilts away from the basal plane the PSAW gives way to the RSAW. Note for the $\Phi = 45^{\circ}$ plane, the plane in Miller-Bravis notation depends on the a/c ratio of the crystal.

fig. 4.4, the determinant curve would need to be computed for every propagation direction. Whilst observing the mode transitions in a final dominant velocity surface gives the impression the wave velocity is discontinuous, the velocity of each individual mode is slowly varying. Therefore, having determined the velocity of the minima for one propagation direction, in the succeeding direction the determinant need only be searched in a small velocity round around the minima of each mode - such that, $v_{\text{track}_\text{max},i} - v_{\text{track}_\text{min},i} << v_{\text{max}} - v_{\text{min}}$, where *i* indicates the wave mode and v_{track} is the velocity of the *i*th mode at the previous propagation direction.

Furthermore, the computation time may be further reduced by compiling the calculation of the determinant curve in C code. The overall time saving is demonstrated in fig. 4.15b.

These refinements have facilitated a speed increase of a factor of 10 for cubic materials and a factor of 100 for hexagonal materials - this makes a brute-force approach a viable solution to the full inverse problem.

4.2.5 Limitations

One known drawback of the forward solver is the calculation of the dominant modes; currently, the solver discriminates from the calculated relative displacement of each mode. However, this is


Figure 4.15: The run time of the forward model is an important factor, the computation time can be minimised by selecting an appropriate velocity resolution (a) and using the new compiled tracking method (b).

an imperfect model as it does not account for the generation mechanism. Some of the possible wave modes are inefficiently generated by thermoelastic absorption. Fig. 4.16 shows three velocity surface spectra on the (001) plane of nickel. The simulated spectrum in fig. 4.16b uses the calculated displacement to weigh the amplitude at each direction. This suggests the PSAW mode has a lesser amplitude than the RSAW mode, with the amplitude diminishing to zero near $45^{\circ}/45^{\circ}$. However, the measured data in fig. 4.16c shows that the wave is indeed measured in this direction. Although the PSAW wave amplitude is smaller than that of the RSAW, the relative amplitude does not match that predicted by the forward model.

Fig. 4.8 has shown the relative amplitudes of the two modes are very similar around 22° on the (001) plane, incorrectly calculating these relative amplitudes leads to uncertainty in the position of the mode transition.

Furthermore, the forward solver only allows for the wavenumber in the z-direction to be complex; however, it is possible to also consider complex numbers in the x-direction, where the complex term can be interpreted as a dampening of the wave amplitude as it propagates along the surface.



Figure 4.16: Velocity surface spectrum on the (001) plane in nickel, the measured relative amplitudes do not agree with the simulated amplitudes. The colourscale in all figures moves from blue to yellow with increasing amplitude.

As described above the damping is true of PSAWs. The thesis of Jan Sermeus, considers this problem and concludes the additional computation time, due to having to minimise ΔR in the complex velocity plane (2D) rather than just the real velocity line (1D), introduced from the complex method is not worthwhile. This conclusion is reached as on certain planes, for example the (111) plane, the solution is purely real and in cases where a complex solution exists the discrepancy is only a few ms⁻¹ [278]. It is worth bearing in mind this may need revisited in the pursuit of the high accuracy elastic constant measurements, as will be discussed towards the end of chapter 6.

4.3 Inversion

Whilst the ability to image the structure and contrast across a material's microstructure is useful, the true power of SRAS lies in its ability to map the crystallographic orientation in crystalline materials. The process of determining the orientation and/or elastic constants from the SAW velocity is not straightforward. If any two of: crystal orientation; SAW velocity or elastic constants are known, then in principle, the third can be computed. However, determining either physical parameter from the velocity is an ill-conditioned problem that does not lend itself to a tractable analytical solution. The lack of unique mapping of $v \to (h, k, l\theta)$ means a different approach is required.

4.3.1 Objective function

Method A - Least squares

The first step is to construct an objective function, F, to minimise, a least squares approach is commonly used for such problems [225], as given in equation 4.27.

$$F = \frac{1}{N} \sum_{i=1}^{N} (v_{\exp}(\theta_i) - v_{calc}(X, \phi_{1,i}))^2$$
(4.27)

where v_{\exp} is the experimentally measured velocity at angle θ_i , with N samples in the velocity surface. $v_{\operatorname{calc}}(h, k, l, \phi_{1,i})$ is the velocity calculated by the forward model. The solution is then the argument of calculated velocity surface which maximises F, such that $O(X, \phi_1) =$ $\operatorname{argmax}_{X,\phi_1} F(X,\phi_1)$ where ϕ_1 is the rotation of the velocity surface and X represents some crystallographic parameters, such as elasticity and/or orientation.

This approach compares the velocity taken from the maximum amplitude at each angle in the velocity surface spectrum and compares it to the velocity calculated by the forward model, by looping through every plane and rotation the orientation may be determined.

The performance of this approach in terms of R-value (a standard method for comparing the misorientation between cubic crystals, see appendix E), as a function of scanned directions and velocity standard deviation, is given in fig. 4.17. The method is seen to perform well, R-value $< 2^{\circ}$, when an angle increment of less than 5° is used. A similar performance was seen when determining the orientation in experimental data (fig. 4.16c) to an of R-value, of $< 1^{\circ}$.

The forward model is calculated to at a discrete plane and rotation, and it is therefore probable that the plane and rotation of a measured grain will fall between two calculated planes. Thus, in the above simulation, the rotation is misaligned by 0.5° (the forward model is calculated at 1° increments), and the 'experimental' plane is (0.6232 0.1336 1) (planes are calculated at increments of 0.05). Despite this, the algorithm performs well, even as the number of scanned directions is reduced.

The obvious drawback of method A is that it takes no account of the signal-to-noise ratio of SAW velocity measured - put simply, all data is considered good data by this approach. This could be addressed by introducing a weighting term, w_i , which contains the measured amplitude of the signal i.

However, method A is only suitable for comparing a single measured velocity in each propagation direction; in reality, two modes are regularly measured around the 'mode-transition'. This would not be an issue if the location of this measured transition matched the forward model perfectly, but - as discussed in §4.2.5 - the forward model does not capture the reality of the generation mechanism, so the transition is not expected to match reality perfectly. In §4.2.3, we see the rela-



Figure 4.17: The R-value misorientation performance of the method A inverse solver in nickel, as a function of the number of propagation directions and measurement velocity standard deviation. Each value is the mean of 1000 repeats. See appendix E for detail on the definition of R-value used.

tive amplitudes of the surface wave modes are slowly varying and are of similar amplitude around the 'mode-transition'. In cases where the amplitude (in the frequency domain) of both wave-modes are equal, selecting a dominant velocity is somewhat arbitrary. In cases where the mode selected is different from the forward model, a significant error in the sum of the squares will result. For single velocity surfaces, the measured directions close to mode transitions can be omitted from the inversion by hand. However, in megapixel (or larger) orientation images, this is clearly not practical. Rather, an inversion algorithm that can cope with the measurement of multiple modes is preferable.

Therefore, a least squares approach is valuable, thanks to the relative ease and speed of calculation, when working with model data when insights on the shape of the solution space are sought. However, the inability to cope with two (or more) measured wave modes make it unsuitable for experimental data.

Method B - Overlap function

Rather than perform the inversion on just the velocity, taken from the location of the maxima in the velocity spectrum, it is desirable to the inversion against the full velocity surface spectra. This method, the so-called overlap function, has been previously reported by Li et al. for determining the crystallographic orientation in cubic materials [279].

For each pixel in the specimen, the acoustic measurement provides a plot of signal amplitude against velocity as a function of propagation direction, θ (as shown in fig. 4.16). Thus the measured signal can be defined as $A(v(X, \phi_1), \theta)$, where $0^\circ \leq \phi_1 < 180^\circ$ and X represents some crystallographic parameters, such as elasticity and/or orientation.

The velocity database calculated from the forward model, $v_c(X, \phi_1)$ can then be transformed into a binary matrix I_X by equation 4.28. I_{hkl} has the same velocity dimension, N_v , as measured signal A (varies depending on the zero-padding in Fourier transform), and is twice the length of A in the rotation dimension - this allows the rotation of A relative to the forward model to be determined.

$$I_{\rm hkl}(v,\phi_1) = \begin{cases} 0, & v \neq v_c(X,\phi_1) \\ 1, & v = v_c(X,\phi_1) \end{cases}$$
(4.28)

The overlap between I_{hkl} and A is now determined by calculating the sum of the element-wise product as the lag of I_{hkl} (with respect to A) is varied, as defined by equation 4.29.

$$S_X(p,q) = \sum_{v=1}^{N_v} \sum_{\theta=1}^{180} A(v,\theta) I_X(v-p,\theta-q)$$

$$-(N_v-1) \le p \le N_v - 1$$

$$360 < q < 360$$
(4.29)

The vertical index in the result, S represents velocity offsets, that is to say, making the measured wave linearly faster or slower. However, the user defines the acoustic wavelength when setting up the experiment; therefore, only p = 0 is of interest. The q value reflects the rotation of the measured data relative to the $\phi_1 = 0^\circ$ definition in the forward model, and this can again be limited to the range $-180^\circ < q < 180^\circ$, outside of this range the velocity surfaces do not fully overlap and maxima should not be found.

By comparing to the full velocity spectrum, measuring more than one wave mode won't corrupt the inversion. Of course, constructing the binary matrix and calculating the pixel-wise overlap is more computationally demanding than method A, but this is essential when working with experiential SRAS data.

4.3.2 Finding the orientation

The objective function for determining the orientation is then found by equation 4.30, where S_{hkl} is the output of equation 4.29 for a given plane (hkl).



(a) Binary matrix of velocities from the forward model for a single plane.

(b) Value of the overlap function, S, when the binary matrix is compared to the experimental measurement. The position of the maxima is indicated.

Figure 4.18: A binary matrix of the velocity, as calculated by the forward model, is overlapped with the experimental velocity surface spectrum, the result is S. The colourscale in all figures moves from blue to yellow with increasing amplitude.

$$F_O(h, k, l, \phi_1) = \max(S_{hkl})$$
(4.30)

The steps defined in equations 4.28 to 4.30 are then repeated for each orientation to evaluate the objective function for each possible orientation. Finally, the calculated orientation (h,k,l,ϕ_1) is given by the arguments which are found to maximise F_O , equation 4.31.

$$\underset{(\mathbf{h},\mathbf{k},\mathbf{l},\phi_1)}{\arg\max} F_O(h,k,l,\phi_1) \tag{4.31}$$

An example of the value of F_O for the plane (0.10 0.95 1) is shown in fig. 4.19. The value of F_O varies as a function of the plane (hkl) and rotation (ϕ_1). A maximum is found at the position (0.10 0.95 1) with rotation 30°.



Figure 4.19: Inverse fitting result for the plane (0.10 0.95 1). (a) Stacked figures of merit, F_O across the orientation space, z-slices represent rotation on the plane. (b) maximum figure of merit, for a rotation angle of 30°, the best fit is indicated. (c) the measured velocity surface spectrum with the fitted velocity overlaid *. The colourscale in all figures moves from blue to yellow with improving fit.

4.3.3 Finding the elasticity

Similarly, this approach can be re-framed to find the elastic constants. In this case, the objective function value is then found by equation 4.32, where $S_{C_{ij}}$ is the output of equation 4.29 for a given elastic constant set at a single orientation.

$$F_E(\mathcal{C}_{ij}) = \max(S_{\mathcal{C}_{ij},\phi_1}) \tag{4.32}$$

Again, the steps defined in equations 4.28 to 4.30 are then repeated for each permutation of the elastic constants and the objective function evaluated. Finally, the calculated elastic constants C_{ij} are given by the arguments which are found to maximise F_E , equation 4.33.

$$\underset{\mathcal{C}_{ij}}{\arg\max} F_E(\mathcal{C}_{i,j}) \tag{4.33}$$

Fig. 4.20 shows the value of F_E for an 8 and 32 fringe generation patches in the C₁₁-C₁₂ plane, respectively. Increasing the number of generation fringes makes the solution space more convex, highlighting the maxima position. However we see the value of F_E is similar along the diagonal $C_{11}-C_{12} = 70$ GPa. Chapter 6 investigates the required accuracy to accurately determine the position of the maximum, in turn determining the elastic constants of the specimen.



(a) Magnitude of F_E , 8 fringes in generation patch.

(b) Magnitude of F_E , 32 fringes in generation patch.

Figure 4.20: Value of F_E in the plane $C_{11} - C_{12}$, increasing the number of fringes in the generation patch creates a more narrowband signal in the velocity domain, in turn making the solution space more convex. Both results are noise free. The colourscale in both figures moves from blue to yellow with improving fit.

4.3.4 Determining unknown elastic constants and orientation

A primary objective of this work is to demonstrate the elastic constants of a specimen may be determined by SRAS measurement without prior knowledge of the crystallographic orientation.

For a given pixel, P, the determined orientation has no influence on a neighbouring pixel, Q. However, we shall assume the elastic constants are a global property of the material such that $C_{ij}(P) = C_{ij}(Q)$.

The objective functions of §4.3.2 and §4.3.3 can be combined to provide a new objective function, F_{EO} , which depends on both the elastic constants and orientation, equation 4.34.

$$\overline{F_E}(C_{ijkl}) = \prod_{1}^{N_g} F_{EO}(C_{ij}, W, X, Y, Z)$$

$$(W(C_{ij}), X(C_{ij}), Y(C_{ij}), Z(C_{ij})) =$$

$$\underset{(h,k,l,\phi_1)}{\operatorname{arg\,max}} F_{EO}(C_{ij}, h, k, l, \phi_1).$$
(4.34)

In this work, equation 4.29 has been repeated for every modelled elastic constant set and orientation; thus, the figure of merit for the full inverse problem is $F_{EO}(C_{ij}, h, k, l, \phi_1)$, for a single pixel. The ability to simultaneously determine unknown elastic constants and crystallographic orientation from just SAW velocity measurements is explored through simulation and experimental measurements in Chapter 6.

4.3.5 Inversion summary

The inverse problem considers the process of crystallographic orientation and/or elastic constants from experimentally measured velocity surfaces. Two approaches for the inversion have been introduced based on a least-squares method and overlap integral. The least-squares approach has been shown to be quite robust to noise, both with simulated and experimental data, but performance significantly degrades when the number of scanned directions is reduced, noise added, and the evaluated velocity surface falls between the calculated planes in the databased. For example, the forward model is usually completed at 1° in the plane; therefore, a measured plane of $\Phi = 45.5^{\circ}$ is inherently less well represented in the database than a measured plane of $\Phi = 45.5^{\circ}$. This case reflects the reality of much of the data captured experientially by SRAS. Thus, an alternative optimisation algorithm, the overlap function, more appropriate for this non-ideal data, was introduced.

Rather than using just the calculated SAW velocity (i.e. the position of maximum amplitude in the frequency domain), the overlap between the forward model and the full experimentally measured spectrum is used for the inversion; this is unlike most schemes that use just the single velocity measurement [252]. This approach was found to be more robust to experimental noise and a reduced number of scanned directions. The correct plane is that with the highest summed correlation value. Examples of the solutions space are provided for both the orientation and elasticity problems, with a detailed example of the orientation calculation given.

The drawback of the overlap method is the slow speed of calculation. Thus, if the priority is insights into the shape of the solution space or if well sampled and low noise experimental data is available, the least-squares approach may be of use.

4.4 Simulation of SRAS measurements

In order to investigate the hypothetical performance of a SRAS instrument, it is often useful to generate synthetic SRAS measurements. This allows the performance over a range of experiential parameters to be estimated, including noise and the number of generation fringes. The techniques introduced here will be used primarily in Chapter 5, where the accuracy of ordination and elastic constants determination is of interest.

4.4.1 Time domain signals

A time-domain SRAS signal may be simulated by the convolution of a confined Gaussian window and a sinusoid, at a fixed frequency, as defined by equation 4.35.

$$S_0(t,f) = \sin(2\pi f_0 t) \exp(-\frac{(t-k)^2}{k^2})$$
(4.35)

where $k = \frac{c_n}{f_0}$, c_n is the number of cycles, f_0 is the centre frequency and S_0 is the SRAS waveform. Fig. 4.21 shows a simulated signal and the Fourier transform, showing the velocity spectrum, for a waveform of SAW velocity $3000 \,\mathrm{ms}^{-1}$ at a wavelength of $24 \,\mu\mathrm{m}$, giving $f_0 = 125 \,\mathrm{MHz}$

To fully simulate experimental measurements typical incoherent noise must also be considered, this may be simulated by

$$S_n(t, f, \sigma_S) = S_0 + \delta S(\sigma_S) \tag{4.36}$$

where δS is drawn from a normal distribution, as defined by

$$\delta S(\sigma_S) \sim \mathcal{N}(0, \sigma_S^2). \tag{4.37}$$



Figure 4.21: Simulated SRAS signal in time and frequency domains.

This adds noise to the amplitude of the SRAS signal. Simulated noise and that which is measured experimentally is compared in fig 4.22a; the simulated noise shows good agreement with the experimental data when a standard deviation of 0.2 is used. The simulated noise has been bandpass filtered with a passband of 50 MHz - 200 MHz to match the response of the experimental system.

Fig 4.22b shows a simulated SRAS signal of 8 generation fringes and a noise component generated by a standard deviation of 0.2, giving a signal-to-noise ratio of 8 dB.

4.4.2 Velocity surfaces

When working with velocity surfaces (rather than velocity surface spectra), a more straightforward route to simulate measured data is to apply a perturbation to the calculated velocity.

Starting with a pre-calculated velocity surface, from the forward problem, the 180 (or fewer) velocity measurements around the plane, v, are perturbed by a small velocity, δv . To produce a final velocity surface that mimics that measured experimentally. This process can be shown by



(a) Comparison of experimentally measured and simulated noise signals.

(b) Simulated SRAS signal of 8 generation fringes, with a noise component of 0.2, giving an SNR of 8 dB.

Figure 4.22: Noise, generated from a normal distribution, is added to the simulated SRAS signal to recreated an experimental measurement.

equation 4.38.

$$v + \delta v = v_{\text{measured}} \text{ where } \delta v \sim \mathcal{N}(0, \sigma_v^2).$$
 (4.38)

A simulated velocity surface of $\delta v = 20 \, ms^{-1}$ is compared to an experimental measurement in fig 4.24. The velocity surfaces shown are comparable, confirming this simulation method produces representative experimental data.

4.4.3 Simulated polycrystalline specimens

The output of the forward model can be visualised as a velocity distribution, fig 4.25(a) (blue distribution). This can be compared to a measured velocity distribution and, in effect, compares the concentration of crystallographic orientations within a measurement to an untextured specimen - that is to say, a measured velocity distribution in good agreement with that of the forward model suggests the absence of preferential orientations. This concept is more formally known as the study of texture.



(a) Simulated b-scan of SRAS signals with monotonically increasing SAW velocity, across the velocity range found in titanium specimens.



(b) Set of experimentally measured waveforms measured in titanium specimen, sorted by velocity. The velocity ranges from approximately 2400 ms^{-1} to 3200 ms^{-1} .

Figure 4.23: Comparison of simulated and experimentally measured b-scans across the velocity range typical of titanium specimens.

Many engineering specimens develop microstructures that have strong crystallographic texturing. Usually, these are well characterised for a given process and material, and therefore it is useful to assess the impact this texturing has on the observed SAW velocity. For a given specimen, an orientation distribution function can be used to express the relative weight of each orientation, g.

$$ODF(g) = \frac{1}{V} \frac{dV(g)}{dg}$$
(4.39)

Where V is the length dimension of the whole specimen, depending on the dataset, this may be area or volume, and dV is the length occupied by a given orientation.

There exist a number of typical crystallographic textures, for example, cube and Goss texture components, as defined by table 4.3. These textures are commonly found in rolled sheets and are of particular importance in magnetic materials (see §2.2.3). Goss and cube textures are demonstrated here, but this approach may be used with any appropriate texture or superposition of texture components.



(a) Experimentally measured velocity surface.

(b) Simulated velocity surface, with a velocity standard deviation of $10 \,\mathrm{ms}^{-1}$

Figure 4.24: Comparison of experimentally measured and simulated velocity surfaces on the (001) plane of nickel. In this case, the simulated measurement takes the velocity surface from the forward model and applies a velocity perturbation from a normal distribution. Fig. 3.11a shows an experimentally measured velocity distribution.

From the texture definitions it is possible to generate a hypothetical orientation distribution function, which is then sampled a number of times to build a velocity distribution representative of that texture. Fig. 4.25 shows compares the full velocity distribution of nickel with a hypothetical specimens showing a strong Goss and cube textures, respectively, sampled in the 0° propagation direction.

Finally, if required, this statistical velocity information can be assembled into an image analogous to a SAW velocity map. Voronoi tessellations are a valuable tool for replicating the morphology of polycrystalline microstructures [280]. When a crystalline material begins to solidify, it does so from an initial nucleation site, radiating outward. This growth continues until another grain is encountered, forming a grain boundary.

The process of calculating Voronoi tessellations is well explained by Okabe et al. [281] and is now implemented on most numerical computation packages. The mathematical nomenclature required to describe the calculation of the tessellations is somewhat involved, and therefore it is

| Texture | plane and direction | | | |
|---------|---------------------|--|--|--|
| Cube | $\{001\} < 100 >$ | | | |
| Goss | $\{011\} < 100 >$ | | | |



 Table 4.3: Definitions of Goss and cube crystallographic textures.

(a) Goss texture velocity distribution in nickel.

(b) Cube texture velocity distribution in nickel.

Figure 4.25: Velocity distributions for full nickel orientation space compared with the common cube and Goss textures. The presence of texture leads to unequal weighting in the measured velocities.

satisfactory for the purposes of this work to simply describe the final properties of the diagram.

For a number of random seeds points, N, distributed across a grid in two dimensions, an equal number of convex polygons, C, can be generated, such that each polygon contains only one seed point. The region bounded by the polygon C_i is then closest to the seed point N_j , by Euclidean distance than all other seed points.

In the Poisson-Voronoi method, seed pixels are distributed at random with no condition on the proximity; this can lead to a high degree of geometric anisotropy. Fritzen et al. suggested the 'hardcore' Voronoi method was more appropriate when modelling crystalline microstructures[282]; the geometric anisotropy is reduced by only allowing one seed pixel within a given radius. Fig. 4.26 shows an example Voronoi diagram, calculated by the 'hardcore' method, which can be used to represent a specimen microstructure. This approach generates equiaxed grains, but modifications are possible to incorporate grain diameter distributions. Additionally, the distance is equal in all directions such that grain will grow outwards circularly; the physical implication is the cooling rate in the specimen is isotropic. Instead, an effective cooling field may be specified to generate columnar microstructures with growth directions rotated towards the heat source [283].



Figure 4.26: Simulated grain structure by Voroni tessellation, markers indicate the centres of each region - effectively the initial nucleations seeds of each grain.

Fig. 4.27 shows simulated SRAS velocity maps (the grain morphology in each map is that of fig. 4.26). As in reality, several neighbouring grain regions have similar velocities and thus appear as one grain; the grains can only be separated by observing contrast across both velocity maps.

As additive manufacturing specimens exhibit strong texturing, these approaches will be used in the following chapter to relate the experientially measured SAW velocity and the underlying crystallographic orientation and texture.



Figure 4.27: Simulated velocity maps of the 0° and 45° propagation directions for strong Goss and cube textures, from the grain structure in fig 4.26. A velocity standard deviation of 30^{-1} is added to better replicate experimental data. All data is plotted to the same velocity scale.

4.5 Summary

This chapter has introduced the numerical toolbox associated with SRAS; leveraging these tools allows properties, such as crystallographic orientation, to be determined from SAW velocity measurements. The primary concern within that has been the calculation of SAW velocities, known as the forward problem, and it is upon this which all subsequent methods depend. The process for calculating the SAW phase velocity as a function of the crystallographic orientation and elastic constants has been described, results in cubic nickel and hexagonal titanium have been presented in detail, with brief results in tungsten and potassium chloride given to highlight the impact of lattice anisotropy on the velocity surface in cubic materials.

The ability to solve the rapidly forward problem lends itself to running multiple brute-force style solutions to inverse problems. As such, the computational performance of the forward model is reported as a function of the velocity search step. The latest iteration of the forward model has reduced computation time by a factor of ~ 100 in titanium and ~ 10 in nickel. This has made searches over the elasticity space viable for the first time.

The overlap method is introduced for the inversion of experimental velocity surface spectra to determine the crystallographic orientation and/or elastic constants. A least-squares approach is also shown to be viable with datasets where only one velocity is measured in each propagation direction. Compared to the overlap method, the reduced computation time makes this a practical approach for future simulations.

Methods of generating synthetic measurement data are outlined. Time-domain signals, with and without noise, are generated and compared to experimental measurements. Velocity surfaces, as generated by the forward model, are shown to compare well with experimental data when perturbations following a normal distribution are applied.

Finally, the velocities calculated by the forward model represents the entire span of the orientation space. This space can be sampled non-uniformly to represent crystallographic textures. This has been shown for the common cubic textures of Goss and cube - which were introduced in Chapter 2 - but may be extended for any arbitrary texture. These velocity distributions can then be assembled into synthetic SAW velocity maps, mimicking experimental measurements. In specimens with well-characterised texture, such maps can be compared to experimental data to relate the measured velocity to crystallographic orientation - even when the individual grains are not resolved.

This suite of techniques, along with the instrument described in the previous chapter, will now be exploited in the following two results chapters, firstly to explore the information in additive manufacturing specimens which can be extracted from SRAS measurements, Chapter 5.

Chapter 5

Measurements in additive manufacturing specimens

The current spatially resolved acoustic spectroscopy (SRAS) instrument and corresponding techniques to calculate the surface acoustic wave (SAW) velocity and perform inversions has been described in the previous two chapters. This chapter reports on one of the primary thrusts of this thesis, understanding SRAS measurements of additively manufactured specimens. Results are split into two primary sections, firstly results on wire-arc additive specimens (§5.1), followed by results with laser-powder bed fusion(§5.2). Only one material in wire-arc specimens is studied, Ti-6Al-4V, as such results and background to these specimens are presented in-depth. For laser powder bed fusion, three materials are presented, however the study of high silicon steel (Fe-Si) is most interesting, and therefore a greater breadth of information is provided for these specimens (§5.2.3).

The small grain size typically found in additive manufacturing is below the spatial resolution of the EMDA instrument. Therefore the SAW velocity measured is a function of all of the grains at the generation position. The impact of this is that the inversion method described in the previous chapter cannot be used to determine the crystallographic orientation of each crystallite. With this in mind, the overarching theme of this chapter is to use the SAW velocity measurements of the textured specimen, in combination with the material forward model, to still deliver some degree of useful information on the microstructure of the material.

5.1 Wire-arc additive manufacture

In the titanium alloy wire-arc specimens presented, the primary concern is the control and elimination of microtexture regions, which form as the weldment cools. Components manufactured by an additive process are characterised by a strongly anisotropic microstructure, primarily columnar grain growth in the build direction [72]. In WAAM processed Ti-6Al-4V, this leads to [001] alignment in columnar prior- β grains, which gives rise to the child's strong texturing α grains [126]. This is known to contribute to mechanical anisotropy [127], inferior fatigue life [128], geometric distortion and crack initiation at prior- β grain boundaries [129]. Martin et al. have suggested that one way of producing a microstructure more suitable to engineering applications and preventing the induced residual stress is to promote equiaxed grain growth [18]. The use of inter-pass rolling has also been trialled in WAAM to induce plastic strain to relieve tensile stresses and enact grain refinement [130, 131]. This approach significantly reduced the prior- β grain size and weakened the texture of both the prior- β and final α structures [26, 126]. The resulting texture leads to components with more isotropic and desirable properties. Therefore, ensuring this rolling process has been enacted efficaciously is vital to the industrial uptake of WAAM components.

This study utilises the SRAS technique to capture SAW velocity maps from undeformed and rolled WAAM specimens, in both polished and as-deposited states. To date, the effects of dualphase materials on SRAS imaging have not been reported. In effect, as the α -phase is probed, it is unclear if features from the prior- β grains can be captured. The three primary aims of this first section are:

- Understand the extent to which the texture and other features originating from the fabrication process can be probed through the use of SAWs.
- Investigate if SAW velocity maps can be used to measure the efficacy of rolling.
- Establish the feasibility of capturing these measurements in-line as a step towards inspection during fabrication.

This work demonstrates off-line inspection of prepared WAAM fabricated specimens using SRAS. The detection of prior- β grains and the microstructural differentiation between undeformed and rolled specimens has been successful, and an explanation to relate the underlying microstructure to the measured acoustic response has been provided. Further investigation of the angular response allows the anisotropic features to be studied, which has been utilised to extract additional information such as the texture of the substrate and minimum stiffness direction. Finally, the challenges to overcome for on-line inspection are explored in detail.

5.1.1 Materials and Methods

WAAM Specimens

All specimens reported in this study haven been produced using the WAAM process using Ti-6Al-4V welding wire feedstock, on to a titanium substrate. A pulsed gas tungsten arc welding torch, with argon shielding, was used to deposit structures made of 20 single-bead layers, giving an approximate final width of $\sim 6 \text{ mm}$ and height of $\sim 24 \text{ mm}$. Following the deposition and cooling of each layer, a 100 mm diameter roller was passed over the top surface to create the deformed specimens. Rolling applied vertical forces were of 50 or 75 kN, monitored by a load-cell. Key build parameters are given in Table 5.1.

Specimens were sectioned using electrical discharge machining and mounted. For each roller force, both the x - z and y - z planes were exposed for analysis. All specimens were prepared by standard metallographic preparation for titanium, and etched with Kroll's reagent for optical micrographs. Additionally, sections of the undeformed and 75 kN rolled specimen have been left unprepared so as to allow imaging in the as-deposited state to simulate measurements performed in an industrial environment during the manufacture of a component.

| Build parameter | Value | Unit | Build parameter | Value | Unit |
|---------------------|-------|------|-------------------------------|-------|-------------------------------------|
| Average arc voltage | 12 | V | Torch stand-off | 3.5 | mm |
| Average current | 110 | А | Torch shield gas flow rate | 10 | $l \cdot min^{-1}$ |
| Peak current | 150 | А | Trailing gas flow rate | 20 | $l \cdot min^{-1}$ |
| Background current | 70 | А | Travel speed | 270 | $\mathrm{mm}\cdot\mathrm{min}^{-1}$ |
| Wire diameter | 1.2 | mm | Wire feed speed | 1.6 | $\rm m\cdot min^{-1}$ |

Table 5.1: WAAM deposition parameters, used to fabricate specimens for this study.

Microstructure analysis and preparation

To establish a baseline of specimen microstructure, EBSD orientation images were captured using a JEOL 7100F FEG-SEM scanning electron microscope, with an Oxford Instruments EBSD system and Aztec acquisition package. The MATLAB package MTex was used for processing and preparation of pole figures, with a 5×5 median window used to smooth the collected EBSD data. Representative optical macrographs were captured using a Leica M205 FA stereo microscope and



Figure 5.1: (a) Normal SRAS experiment in large grain materials, five grains of different stiffness shown in this diagram, where the generation patch, D_g is much smaller than the grain, D_{grain} . In specimens like this, the orientation of each grain region can be found. (b) Diagram of SRAS experiment in specimen with fine Widmanstätten pattern; at each fringe the acoustic wave velocity is a function of the elastic properties of all lamella, C_{ijkl} , in that fringe width, λ_g . It is possible to have acoustic impedance mismatch due to differing crystal orientations.

stitched together in post-processing. Finally, high resolution images and surface profiles of the asdeposited side-walls have been captured with a focus variation microscope, Alicona InfiniteFocus G5, with $10 \times$ objective lens, giving a lateral resolution of $3.9 \,\mu\text{m}$.

SRAS measurements

The velocity of these SAWs, v_{saw} is a function of the elastic properties of the specimen, furthermore, the calculated velocity is a property of the specimen area under the generation patch only. In large grain materials [284] where the acoustic wavelength $\lambda_g \ll D_{\text{grain}}$ (where D_{grain} is the major axis length of the grain), each measurement point can be treated as a single crystal, fig. 5.1a. The angular response of such grains is well understood. However, in fine microstructures such as WAAM processed Ti-6Al-4V, where $\lambda_g \gg D_{\text{grain}}$, the acoustic velocity is a function of the elastic properties of all lamella in that fringe width, as shown in fig. 5.1b. When generating across multiple orientations or phases, the measured wave packet contained multiple frequency components. The measured velocity by the SRAS technique relates to the most common SAW frequency within the whole generation patch, selected by taking the maximum amplitude in the frequency domain. Increasing the number of phases or orientations under the generation patch will increase the bandwidth of the signal, decreasing the velocity resolution and signal-to-noise ratio.



Figure 5.2: Results from prepared side-wall (x-z). Scale bars indicates 2 mm. (a) SRAS velocity map from undeformed specimen, annotated lines have been added to indicate vertical prior- β grain growth. (b) SRAS velocity map from 75 kN rolled specimen, large prior- β features can no longer be discerned.

Various schemes exist for estimating the bulk elastic properties of anisotropic polycrystalline materials, the most common of which is the Voigt-Reuss-Hill approximation. Using the texture distribution obtained by EBSD, the stiffness tensor for the hexagonal α -phase was transformed into the orthorhombic macroscopic tensor. This allowed a prediction of the SAW velocity, as measured by the SRAS technique, when generating across a region with many dissimilar grain orientations. A full introduction to the Voigt-Reuss-Hill method and its application to SAWs can be found in Kube et al. [285].

5.1.2 Side-wall measurements (x - z plane)

Acoustic velocity maps captured with SRAS for varying rolling force are shown in fig. 5.2, for the (a) undeformed and (b) 75 kN rolled specimens respectively, after polishing.

Considering first the undeformed specimen, fig. 5.2a, columnar structuring, running the length of the specimen in the z-direction and around 2-3 mm in width, can be discerned in the acoustic image. These large features are consistent with the formation of prior- β grains. Within the prior- β grains themselves, there is little acoustic variation; this is representative of large microtextured regions. For a single grain the velocity varies with crystallographic orientation by $\pm 250 \text{ ms}^{-1}$ for Ti-6Al-4V; the observed variation is lower as the average velocity is measured for the grain population

underneath the patch. However, velocity variations of ~ 100 ms^{-1} can be seen between prior- β grains. A band of ~ $500 \,\mu\text{m}$ in height can be seen at the very top of the specimen; this represents the final deposited layer, which has not had the opportunity to recrystallise from subsequent deposition. Grain growth is found to be at a mean angle of $74.2^{\circ} \pm 2.2^{\circ}$ to the x-direction, as indicted by the dashed lines in fig. 5.2a. This 16° rotation from ideal vertical growth is caused by the direction of deposition (left to right), creating a thermal gradient in the x-direction. This is consistent with the findings of Donoghue et al. [126].

Fig. 5.2b, shows the acoustic velocity map for the 75 kN rolled specimen. No columnar structuring can be discerned, suggesting the acoustic response has detected the effects of inter-pass rolling. There is no larger scale structuring, on the millimetre scale, but significant velocity variations can be seen across the specimen. This suggests a more equiaxed microstructure, with a more varied microstructural texture. The velocity variations occur in a spatial range of hundreds of micrometers, compared to millimetres in the undeformed specimen, consistent with a refined β -grain size.

For comparison an enlarged macrograph, after etching, taken from the 75 kN rolled specimen is shown in fig. 5.2c. There is good agreement between structure shown in the SRAS velocity map and the etched image. The macro-regions shown are consistent with the formation of prior- β grains, which in the rolled specimens do not have a fibre texture. Dashed lines indicate representative macro-structure features which can be seen in both fig. 5.5b and fig. 5.5c.

Angle dependent response

Acoustic maps have also been captured at varying source rotations in order to probe possible anisotropic effects. SRAS velocity maps for acoustic propagation angles of 0, 45 and 90° for the three prepared side-wall specimens are shown in fig. 5.3a-i. Again, a clear distinction can be drawn between the undeformed and rolled specimens; whilst the rolled specimens show little angle dependency, the undeformed specimen exhibits variations of up to 100 ms^{-1} . For example, in the 0° map there is little contrast between prior- β phases, however in the 90° map a variation of around 100 ms^{-1} can be seen between prior- β grains. This is suggestive of close crystallographic alignment between prior- β grains in the z-direction, and misalignment in the x-direction. A variation of $\sim 100 \text{ ms}^{-1}$ due to varying propagation angle is suggestive of anisotropic elasticity, typical of large microtextured regions.

Additional features are revealed in the angle scans, such as the high velocity response that runs as a line from the top left of the specimen to the bottom in the 45° scan, fig. 5.3b. Interestingly, some prior- β grains are seen to split off as the deposition continues, triggered by the formation of thinner grains within. A repetitive banding can be seen in the z-direction; this likely corresponds to the layers of deposition. This feature is most clear in the 90° scan, because the axial resolution is normally significantly better than the lateral resolution due to the nature of the generation patch in SRAS [253]. The spatial resolution of the EMDA system is discussed in §3.2.6.

In both the 0 and 45° degree scan in the undeformed specimen high velocity regions are seen at the interface between some of the prior- β grains. From inspection of the measured acoustic waves, this effect is believed to be an anomalous velocity caused by significant acoustic impedance mismatch between the prior- β grains, causing reflection of the acoustic wave across the generation patch, convoluting the final measurement in a small area. This effect will only be noticed in areas of high β -phase concentrations.

Clearly from these maps the prior- β features can be discerned, an interesting outcome as the acoustic measurement is probing the α -phase. The acoustic response throughout these prior- β features appears consistent. Of further interest is the inconsistency of the anisotropy of the prior- β grains. For example, in the undeformed specimen some prior- β show no measurable acoustic variation with propagation angle.

5.1.3 Cross-section measurements (y - z plane)

Additional SRAS maps have been captured from the prepared cross-section, y - z plane in the undeformed and 75 kN rolled specimens. Results are shown in fig. 5.4. An etched macrograph of the undeformed cross-section is shown fig. 5.4a and the corresponding SRAS acoustic map in fig. 5.4d. Acoustic maps for this specimen were captured at propagation angles of 0 to 170°, in 10° increments. Please note the map shown is from 110°.

The most prominent feature seen in the SRAS velocity map fig. 5.4d, is the region in the lower left of maximum 6 mm in height, which extends across the specimen at $\sim 150 \,\mathrm{ms}^{-1}$ slower than the bulk of the specimen. Inspection of the macrograph, fig. 5.4a, shows clear parallels with the acoustic map fig. 5.4d, foremost the low velocity region can be clearly seen in both images. Again, this is attributed to the formation of a prior- β grain.

To further understand the interface between the regions, a detailed micrograph at the boundary is shown in fig. 5.4c. The etched surface reveals a typical $\alpha + \beta$ microstructure, with a clear interface running diagonally. α -phase lathes can be seen on either side of the interface but at a different orientation. The interface observed optically (i) agrees well with the boundary seen in the SRAS velocity map, and is consistent with the formation of misorientated prior- β grains. For comparison a similar micrograph has been captured toward the top of the specimen fig. 5.4b, where no acoustic boundary can be seen, this region is the final layer of deposition which has



Figure 5.3: Varying the propagation direction of the SAW, by rotation of the optical mask and detection spot - as illustrated at the top of the figure, relative to the specimen allows the effects of material anisotropy to be probed. SRAS acoustic velocity maps captured for three WAAM specimens (a)-(c) undeformed, (d)-(f) 50 kN, (g)-(i) 75 kN rolling force, for propagation angles of 0, 45 and 90 °, relative to the z-axis, as indicated by fringe schematics at the top of figure. Scale bar shown in (a) applies to all figures and indicates 2 mm.



Figure 5.4: Results from y - z plane cross-section specimens, (a) - (d) undeformed specimen (e) - (f) 75 kN rolled specimen. (a)Etched macrograph,(i) showing prior- β formation, (ii) interface at layers of deposition, (iii) fusion boundary and (iv) final deposited layer which has not been reheated. (b) and (c) Micrographs taken after etching, scale bar indicates 250 µm, (c) shows the interface between two prior- β grains. (d) SRAS velocity map corresponding to (a). (e) 45° SRAS velocity map from 75 kN rolled specimen and (f) 90° SRAS velocity map from 75 kN rolled specimen. All scale bars indicate 2 mm, except in (b) and (c) where the scale bar indicates 250 µm.

not been reheated (iv). In line with this, a consistent Widmanstätten pattern is seen with no substantial change in orientation.(ii) Boundries between depsoited layers are again observed. (iii) Close inspection of the etched image, also shows repeated curved bands which correspond to the fusion boundary segregation band, these features are also weakly seen in the acoustic map. Ho et al. have shown these bands are caused by the development of a transient solute boundary layer at the solidification front when it first accelerates, causing chemical segregation [139]. Such a change in chemical composition explains why this region can be discerned in the acoustic map.

The y - z cross-section of the 75 kN specimen has also been prepared and scanned, fig. 5.4e shows the acoustic velocity map from a propagation angle of 90°. Notably, the distortion of the build in the direction can clearly be seen. Whilst little structuring can be seen towards the centre of the specimen, in-line with significant grain refinement, larger features (up to 2 mm in length) are observed towards the edge of the specimen. This is indicative of the strain field induced by the

163

roller not being evenly distributed though the specimen; indeed, grain refinement is more efficacious towards the centre of the cross-section. The boundaries due to layer deposition are again seen.

5.1.4 Examination of Interface

In WAAM, the layer-wise deposits are built up from a substrate; for titanium alloys an annealed $\alpha + \beta$ titanium base is used. Fig. 5.5 shows the acoustic velocity maps for the interface between the substrate and the deposited structure at propagation angles of 0, 45 and 90° fig. 5.5 (a)-(c), captured at the area indicated in the extended 0° velocity map, fig. 5.5 (d). Firstly, it is clear that texture is much more homogeneous in this area directly above the substrate, suggesting a finer grain size in this region. As this effect occurs only in this isolated area, it is likely to due to the heat sink effect of the substrate causing rapid cooling of the deposition, preventing epitaxial grain growth. Compared to the substrate, the deposited area near the interface shows little acoustic variation, either within a single scan or at varying propagation angles, fig. 5.5 (b)-(d). The little contrast seen suggests the texture is consistent across this region of deposition. Furthermore, the measured velocity is around 50 ms⁻¹ away from the calculated mean velocity, suggesting there is a strong non-random texture in this region.

One additional result that may be of interest is the acoustic response of the substrate. From fig. 5.5b-d, variations of up to 250 ms^{-1} in the acoustic response suggest crystallographic anisotropy. fig. 5.5 shows a velocity vector map for the substrate area, which amalgamates the 0 and 90° scans to show the alignment of the c-axis in large grain hexagonal materials [253]. In this instance, the vector map fig. 5.5e suggests a strong aliment with the prism plane (basal plane in x-axis).

Finally, some texture inheritance can be seen to propagate from the substrate into the initial few layers of the deposition, in the order of tens of micrometres. Given previous investigations have shown that the orientation of the substrate influences the orientation developed in the additive part [286], the substrate texture may be important for control of the deposition in future. Donoghue et al. described the 'ratcheting process' that occurs in the undeformed WAAM specimens where by, in situations where rolling is not applied the original β -grains simply re-grow the orientations of the substrate during epitaxial solidification of the next layer deposited. Therefore, in situations where rolling is not applied, knowledge of the original orientation of the substrate may be crucial for control of the large prior- β structures.

5.1.5 Crystallography

The captured velocity maps describe the acoustic response of the α -phase, thus it is interesting that prior- β grains can still be discerned. As discussed, the fine lathe size found in the α -phase



Figure 5.5: SRAS is able to map the texture of the substrate as well as the deposition. (a) Shows the full SRAS velocity map from the 50 kN rolled specimen, captured at 0°, the indicated area represented the interface between the deposition and substrate. Magnified velocity maps of this interface are given in (b)-(d) at propagation angles of 0, 45 and 90°, respectively. There is a clear grain refinement in the deposition. Strong anisotropic response from the substrate. (e) Shows the corresponding vector map for the substrate area, this map indicates the dominant alignment of hexagonal crystals with respect to the c-axis. Whilst no clear preference can be seen in the deposition, there is a clear preference towards the prism plane (ϕ_1 close to 0°), due to insensitivity to ϕ_2 this plane may be {1010} or {1120}. Scale bar indicates 2 mm.

prevents direct recovery of the crystalline orientation, but several conclusions may still be drawn on the crystallography of the WAAM depositions. Fig. 5.6 defines the hexagonal α -phase crystal, and Euler-Bunge rotation notation that will be referenced throughout this section.



Figure 5.6: (a) Definition of axis and planes in reference to α -phase hexagonal crystal, as given by Euler-Bunge notation. Vector map plot shown in fig. 5.5 is relative to the c-axis direction. Linear acoustic technques are insensitive to variations in ϕ_2 [284], due to elastic transverse isotropy. In practice, this limits the ability to differentiate between $(11\bar{2}0)$ and $(10\bar{1}0)$ planes [284]. (b) and (c) Measured hexagonal phase and corresponding reconstructed cubic β -phase inverse pole figures, captured using EBSD by Donoghue et al. courtesy of Materials Characterization ©2016. Scale bars indicate 1 mm. (b)(i) Microstructrual segregation at the top of the specimen, corresponding to the final layer of deposition. (ii) and (iii) Repeating pattern in prior- β orientation. (iv) Clear columnar growth can be seen at a similar angle to that shown in fig. 5.2.

In previous work by Donoghue et al. [126], recovery of the prior- β macrostructure from EBSD imaging, revealed a repeating orientation of columnar grains along the x-direction. This is demonstrated in fig. 5.6, (b) hexagonal phase inverse pole figure and (c) reconstructed cubic phase inverse pole figure. In fig. 5.6b clear columnar growth at an angle similar to that presented in acoustic map

in fig. 5.2. Furthermore, the orientation of grains can be seen to alternate between an orientation of [001] and one closer to [101], demonstrated by grains indicated by (i) and (ii). Observation of the SRAS data suggests a similar repeating pattern in the columnar growth, along the x-direction. This patterns appears more obviously in the SRAS due to the stark velocity contrast in the 90° scan between prior- β grains. This suggests that the acoustic variation between prior- β grains, seen in fig. 5.3, is due to changes in orientation of the prior- β grains.



Figure 5.7: Block diagram of methods utilised in this section to understand the crystallography and texture development in the probed WAAM specimens, with respect to the SAW velocity information captured by SRAS in §5.1.2.

The general approach taken in this section has been summarised by the block diagram shown in fig. 5.7. EBSD data has been captured for all prepared side-wall specimens imaged in this study. Due to the orientation variants possible in the α -phase, direct interpretation of the inverse pole figures is challenging. The texture distribution is more clearly seen by considering the respective pole figures, fig. 5.8, for each specimen. The undeformed specimen fig. 5.8a shows a strong texture in the (0001) plane, comparatively the 50 kN specimen shows a weaker texture, primarily in the (0001) plane but also observable in the (11 $\overline{2}0$) plane. This is consistent with a greater variety of α -phase orientations, due to the break-up of prior- β grains.



Figure 5.8: Pole figure orientation maps of EBSD measured hexagonal α -phase in (a) undeformed and (b) 50 kN rolled specimens in x - z side-walls, confirming a stronger crystallographic texture in the undeformed specimen. Pole figures are generated from small area EBSD maps.

Results from §5.1.2 suggested the angular dependence on SAW velocity varied between the three specimens. To compare this quantitatively, fig. 5.9 plots the change in velocity at each measurement point between the 0 and 90° propagation direction scans. In the undeformed specimen fig. 5.9a prior- β grains are clearly seen, with a variation of up to 100 ms⁻¹. Within the prior- β grains themselves there is little acoustic contrast, less than 40 ms⁻¹. This is a useful property for the identification of unrefined microstructure acoustically. Whilst angular variations are seen in both the fig. 5.9b 50 kN and fig. 5.9c 75 kN rolled specimens, there are much more in the 75 kN sample and on a smaller scale. Given that 75 kN specimen has the finest prior- β size, this suggests the contrast feature size shown in fig. 5.9b and fig. 5.9c are representative of the size of prior- β grains. Additionally, the frequency of angular acoustic velocity contrast appears to be related to the efficacy of grain refinement.

From comparison with the work of Donoghue et al. it has been established that the acoustic contrast between prior- β grains is due to a change in orientation of the prior- β grains. However, this does not explain the limited acoustic variation seen within a prior- β grain or why some exhibit a greater angular dependency than others, as observed in both the undeformed and rolled specimens as in fig. 5.9. To look at this further the measured EBSD data has been used to predict the macroscopic elastic tensor of the measured specimen. It should be noted that the EBSD from



Figure 5.9: Contrast mapping of the change in velocity (Δv) at each measurement point between the 0 and 90° propagation direction scans. (a) Undeformed (b) 50 kN rolled and (c) 75 kN rolled. Scale bar shown in (a) indicates 2 mm.

the undeformed specimen was captured fully within a prior- β grain showing significant acoustic variation with propagation direction. Given the small feature size in the rolled specimens it is not possible to know how many prior- β grains are captured within these EBSD datasets. Firstly, having captured the EBSD data it was then possible to combine this with the previously described Voigt-Reuss-Hill method, for calculation of the average elastic tensor. The initial hexagonal tensor is defined in Table 5.2. The resulting tensor has triclinic crystal symmetry, and given the dependency of SAW velocity on stiffness it is of interest to plot the Young's modulus as a function of direction. Fig. 5.10 shows the averaged Young's modulus for the three prepared (x - z) side-wall specimens, calculated using the Hill method from the EBSD data shown in fig. 5.8. The Hill method is an arithmetic average of the upper and lower bounds given by the Reuss and Voigt methods.

| | C_{11} | C_{12} | C_{13} | C ₃₃ | C_{44} | ρ |
|---|----------|----------|----------|-----------------|----------|----------------|
| | GPa | GPa | GPa | GPa | GPa | $\rm kgm^{-3}$ |
| α | 170 | 92 | 70 | 192 | 52 | 4430 |
| β | 135 | 113 | - | - | 54.9 | 4430 |

Table 5.2: Elastic constants for titanium, used to predict velocity distribution [287, 288].

This plot exhibits the anisotropic behaviour of the undeformed specimen, as seen in the angular

acoustic response, fig. 5.3. The shift of $\sim 4 \,\mathrm{GPa}$ in the undeformed specimen would correspond to a shift of over $50 \,\mathrm{ms}^{-1}$ in SAW velocity, this compares well with the change in velocity seen in fig. 5.9a. In contrast, the rolled specimens exhibit a much smaller stiffness variation and thus smaller angular velocity change, this again matches well with the results of fig. 5.3, to within $25 \,\mathrm{ms}^{-1}$, which compares well with the accuracy suggested by Kube et al. for the prediction of acoustic wave velocity from the Hill method [285]. Both rolled specimens have a maximum change in Young's modulus of 1 GPa, corresponding to a variation in SAW velocity of around $15 \,\mathrm{ms}^{-1}$. This compares with the response from the 50 kN rolled specimen (fig. 5.9b); however, the 75 kN rolled specimen appears to exhibit large velocity contrasts in small regions, more notable towards the top of the specimen. These small regions are likely to correlate to the refined prior- β grain size, the EBSD data from the 75 kN rolled specimen is captured from a region with low velocity variation, which agrees with the suggested velocity response in fig. 5.10. These results confirm the change in bulk stiffness expected in the undeformed specimen and suggest the EBSD data in the rolled specimens have primarily been captured on prior- β grains with small angular variations. This succinctly demonstrates the acoustic variations observed with varying propagation angle is due to changes in stiffness anisotropy between prior- β grains.



Figure 5.10: Young's modulus calculated from EBSD data, using Hill method for tensor averaging. The undeformed specimen shows a significant variance in elasticity with angle, compared to the two rolled specimens. This agrees well with the results presented in fig. 5.3. The shift of ~ 4 GPa in the undeformed specimen would correspond to a shift of over 50 ms⁻¹ in SAW velocity.

By taking the elastic constants of titanium from literature for the α - and β -phase, given in Table 5.2, it is possible to calculate the expected velocity distribution for both phases, shown in fig. 5.11b. Comparing the measured distributions 5.11a to the modelled velocities 5.11c for each phase suggest that measured velocities match well with the expected distribution for pure α , with all specimens exhibiting a peak value within $100 \,\mathrm{ms}^{-1}$ of the mean velocity of the forward model. If there was a significant influence from the presence of the retained β -phase a velocity shift towards $2500 \,\mathrm{ms}^{-1}$ (the peak β -phase value), would be expected, as the retained β -phase is known to agglomerate at grain boundaries, the shift in the peak of $50 \,\mathrm{ms}^{-1}$ value seen from the undeformed to the 75 kN rolled specimen, is the effect of a greater number of grain boundary crossings, causing the β -phase to have a greater impact of the measured SAW velocity. In fig. 5.11b the full range of possible α -phase velocity measurements is compared to the expected velocity measurement based on the EBSD-measured α -orientation in the undeformed specimen. The expected velocity values match well with the measured range and demonstrate that velocities towards $2000 \,\mathrm{ms}^{-1}$ are not measured as the orientations present in this specimen have characteristic acoustic slowness surfaces with velocities in the upper range $(2800 - 3300 \,\mathrm{ms}^{-1})$ of the velocity spectrum for Ti-6Al-4V. The expected velocity distribution suggests measured values should occur around $3200 \,\mathrm{ms}^{-1}$ in greater abundance than measured. This is possibly due to the limited number of angles (3) where measurements have been captured compared to the calculated angles in the theoretical data (180). fig. 5.11c is a bivariate plot of Euler angles ϕ_1 and Φ , from the measured EBSD in the undeformed specimen. As expected, a limited number of α -phase variants are seen, typical of a $\beta \rightarrow \alpha$ phase transition.

From fig. 5.6c it can be seen that all grains have the [001] direction aligned or close to the zdirection axis, but the orientation in the y-direction is unknown. Looking at the prior- β structure, the top face is fixed to [001] for the fibre, however, the crystalline direction may still rotate freely in the x-y plane. This is of particular interest as this is the plane probed by SRAS when investigating the side-wall specimens (the crystal face exposed in the x-z plane). This suggests the variation in crystallographic response is due to a change in the orientation of the prior- β grain in the y-direction, inducing the formation of α -phase variants with a different characteristic slowness surface. Using Burger's orientation relationship, equation 2.1, the twelve unique α -phase variants were calculated for prior- β orientations of [0 0 1] and [0.6 0 1] (aligned with the z-direction). The resulting slowness surfaces for each variant are plotted as a function of propagation angle in fig. 5.11d. Whilst the twelve variants are calculated, fewer can be seen in the plotted slowness surface; this is due to several orientations being symmetric acoustically. Furthermore, several orientations pairs produce slowness surfaces that are exactly 180 degrees out of phase.

The actual measured velocity will vary depending upon the exact composition of variants found

in the prior- β grain; variant selection is known not to occur evenly [289]. Notably, the [0.6 0 1] orientation will likely show a measurable shift in velocity between 0 and 45 degrees as each variant creates a unique velocity surface. This matches well with the results of fig. 5.3. However, For the [0 0 1] orientation, little to no SAW velocity variation is likely to be measured due to the lack of unique velocity surfaces with only three surfaces possible. One of these is the isotopic basal plane, and the other two are precisely 180° out of phase with each other, and therefore likely cancel out any angular variation. Again, the modelled slowness surfaces for (001) match well with the results of fig. 5.3, where the larger grains showed little to no SAW velocity variance with propagation angle in the undeformed specimen. Thus, it can be concluded that the prior- β grains can be detected using SRAS as a limited number of possible α -phase variants for a given prior- β orientation lead to large microtextured regions with distinct slowness surfaces. This is an important finding because it explains the contrast in the SAW velocity maps between macro-regions, despite the β phase transforming to α . Furthermore, this suggests that if a single propagation direction is to be used to assess the efficacy of rolling, the propagation direction must be chosen carefully. The strong fibre textures in this material means angles towards 90° should always generate the greatest contrast in SAW velocity. With knowledge of the exact orientation distribution in the region, it would be possible to predict the measured slowness surfaces.

5.1.6 Minimum stiffness direction

One commonly cited issue with strong columnar growth is the directional mechanical properties it bestows upon components. Thus far this has been demonstrated by the change in Young's modulus with angle (fig. 5.10). In addition to this it is useful to discuss the extent to which this anisotropy can be detected and characterised by SRAS.

As the material stiffness, E, is proportional to the SAW velocity, SRAS velocity maps captured at multiple propagation directions can be used to indicate the local material anisotropy and stiffness mismatches. Bearing this in mind, regions of measurable velocity gradients, and deviations away from the bulk mean velocity are of interest. Fig. 5.12a, plots the maximum change in SAW velocity for the undeformed cross-section (y - z) specimen. The region in the lower left discussed in fig. 5.4 exhibits a maximum change of $300 \,\mathrm{ms}^{-1}$ in comparison to the bulk, which varies by around $200 \,\mathrm{ms}^{-1}$. Interestingly, layer interfaces are clearly indicated in this plot, and suggest a change in velocity of up to $450 \,\mathrm{ms}^{-1}$. A similar change is seen in the crown of the deposition; as there has not been the deposition of subsequent layers, recrystallization has not been triggered, indicating initiation of recrystallization is crucial in achieving components with isotropic properties.

The anisotropy information contained within SRAS datasets has also been explored by Mark et
al. [254], to indicate the E_{min} direction, where SRAS velocity maps have been captured at multiple propagation angles. This result for the prepared cross-section (y-z) is shown in fig. 5.12b. For the majority of the specimen there appears to be no clear preferential direction, however for the region in the lower left, there is a clear dominance of angle at or close to 90°. This suggests this region will be highly deformable under loading in the y-direction, i.e. though thickness of deposition.

5.1.7 Extracting build process metrics from acoustic data

It is clear from the acoustic maps that the rolling process is successful in homogenising the specimen microstructure, and this change can be sensed using the SRAS technique. In order to quantify the measured SRAS data, a robust metric must be developed, to relate rolling force effect with acoustic response.

In the unrolled specimen, a clear repetitive banding can be seen in the acoustic map, this can be seen in fig. 5.2. These bands are appropriately 1 mm in height and can be correlated to each deposited layer. Martina et al. have previously reported that the α -lamella thickness is known to increase towards the top of a band, due to thermal cycling [26]. This work has recently been expanded by Ho et al. who found this heat affected zone corresponded to coarsening of the lamellar structure, due to transformation coarsening [139]. This variation can be discerned in the acoustic image. Fig. 5.13 plots the instantaneous the velocity gradient in the z-direction. In future, this approach may be used to extract the layer height of deposition for quality control, however this data has been extracted from prepared specimens and further investigations are required to determine if this is feasible on as-deposited specimens.

One of the key aims of this work was to understand how the rolling process affects the acoustic response of the material, and determine if this could be used as a quality control method. In the results presented thus far it is clear that the undeformed specimen can be easily distinguished from the rolled specimens. However, differentiating between the rolled specimens, is not elementary. The large prior- β grains can clearly been seen and sized (Fig. 5.2), however as the prior- β grain size becomes small and the acoustic variations between grains less distinct sizing of these features becomes difficult. Additionally, for this approach to be viable in-line it would be desirable for the differentiation to be feasibly automated.

A simple automated algorithm was developed in Matlab[®], using standard image processing function, as a tool for differentiating between the effects of rolling force. The algorithm is summarised as follows:

• To remove high frequency noise, a Gaussian filter is applied to velocity maps. This has the effect of making only larger features visible. Representative outputs from this stage, for the

(a) undeformed and (b) 75 kN rolled specimen, are shown in fig. 5.14.

- Taking a nominal location for a seed pixel, the algorithm counts the number of connected pixels which have a measured SAW velocity within 20 ms⁻¹.
- The location of the seed pixel is moved across the specimen.

Fig. 5.14, shows the distribution of continuous regions within the specimens, as calculated from the algorithm. As expected, the undeformed specimen was characterised by fewer regions in total but typically with a larger size compared to the rolled specimen. Significantly, the algorithm was able to differentiate between the 50 kN and 75 kN rolled specimen. The developed metric is simple and not computationally intensive, potentially making real time processing in-line viable, subject to further development.



Figure 5.11: Investigation of origin of acoustic response due to texture. (a) Measured velocity distribution for the three side-wall specimens. (b) Bi-variate plot of Euler angles ϕ_1 and Φ , showing a limited number of orientations, as predicted by the Burgers orientation relationship for one prior- β orientation. (c) Comparison of simulated velocity calculated form elastic constants and the predicted velocity distribution, calculated from EBSD measured for both phases. (d) Theoretical slowness surfaces for 6 α -phase orientation from prior- β grain with [0.6 0 1] (upper) and [0 0 1] (lower) orientation in y-direction, similar to the orientations seen in fig. 5.6.



Figure 5.12: Prepared cross-section (y-z) of undeformed specimen, (a) plot of maximum change in SAW velocity with varying propagation angle, larger variations indicate more anisotropic behaviour. Scale bar indicates 2 mm. (b) Scatter plot indicating direction of minimum stiffness, E_{min} , relative to the y-direction. Areas with clusters of similar E_{min} directions will be highly deformable under loading.

| hah | -Yran | hanh | m | a Wasaa | v-n/n |
|-----|-------|------|---|---------|-------|
| | | | | | |

Figure 5.13: Extract of effective layer height in deposition from velocity gradient. Example of the response of the layer height algorithm in the undeformed specimen, which calculates the instantaneous velocity gradient in the z-direction (image has been rotated through 90° CW to aid interpretation).





Figure 5.14: Representative outputs from algorithm after Gaussian filter has been applied to remove high frequency noise, making detection of larger features such as prior $-\beta$ feasible for the (a) undeformed and (b) 75 kN rolled specimen. The contrasting feature size is clear to see. (c) RGB Histograms from the three prepared side-wall (x - z) specimens, showing the distribution of continuous regions areal size, as calculated by the developed algorithm.

5.1.8 Preliminary as-deposited results

All results presented thus far have been captured on prepared specimens, however for the proposed approach of taking measurements in-line and non-destructively for quality control, measurements must be taken from as-deposited surfaces. As-deposited surfaces are optically rough and therefore pose three primary challenges to the use of laser ultrasonics: roughness (out of phase reflections interfere to create a speckle pattern), waviness (drastic changes in the surface normal can mean little light is returned to the detector), and attenuation of the acoustic wave within the material. As SRAS does not rely on time-of-flight measurements, the propagation distance between the generation patch and detection beam can be small to minimise attenuation, allowing the later point to be neglected in this work. The use of the normal KED necessitates an optically smooth surface, (generally an arithmetic mean roughness, R_a less than 20 nm). Rough surfaces lead to a diffuse reflection of the detection beam and create a speckle pattern from the interference of offset reflections at varying phase. Several techniques exist which can adapt to the speckle patterns that are concomitant with rough surfaces, such as random quadrature demodulation [290], two-wave mixing [291] and the speckle knife edge detector [269]. Other interferometry techniques such as Fabry-Pérot interferometry, are inherently insensitive to speckle [292]. Use of one such technique, in place of the KED, in the SRAS system would allow the speckle pattern from the rough surface to be efficiently detected.

The waviness of the specimens causes two issues, illustrated in fig. 5.15, firstly as the height of the deposition varies the detection beam will move out of focus meaning the spot imaged on to the specimen is of larger diameter - detection efficiency is maximised when $D_{detection} \approx \lambda_g/2$ [56]. Furthermore, the curvature means the beam is no longer perpendicular to the surface, causing the reflected beam to miss the detector.

Within the discussion of applying SRAS during the deposition process, it is also necessary to conciser what part of the specimen should be measured. In the section measurements have been presented from the x - z and y - z faces, but clearly, the cross-section (y - z) is not available for inspection during deposition, thus the side-wall (x - z) face is more appropriate for discussing as a realistic on-line solution. Fig. 5.16 shows preliminary SRAS velocity maps captured on the asdeposited side-wall in the undeformed specimen, using two rough-surface capable SRAS systems, which makes use of an interferometer based on random quadrature demodulation. The velocity contrast in these maps correspond to the prior- β structure and despite the lower resolution this is sufficient to indicate the presence of the prior- β structure and hence the failure of the rolling process. It is also interesting to note that again the contrast between the prior- β grains is more discernible in the 90° deg propagation direction, as seen in fig. 5.3.



Figure 5.15: Schematic optical detection on a surface that is both rough and wavy, as found in the side-wall of WAAM depositions. As the detection beam is scanned along the z-axis the beam moves out of focus due to the change in height of the specimen. This reduces detection efficiency. The rough surface creates a diffuse reflection, visualised as a cone of light, this necessitates the use of a 'rough surface detector'. Furthermore, the curvature of the specimen also changes the surface normal as the beam is scanned, in extreme cases this can mean almost no light is returned to the detector.

Micrographs of the as-deposited surface are shown in fig. 5.17b undeformed and d 75 kN rolled. A protective oxide layer of varying thickness is distinguished by the varying colour across the undeformed specimen. Furthermore, a prior- β grain boundary can be seen to run across the specimen. Extracted surface profiles for both specimens are shown in fig. 5.17b. Both specimens exhibit a regular periodicity of approximately 1 mm, that can be correlated to the layer height of deposition, with a change in height up to 20 μ m. Whilst the undeformed specimen presents little roughness within layers, the rolled specimen repeating striations were seen to develop within the layers from the rolling process. This effect is observed in the surface profile map as the increase in high frequency roughness.

Comparing the velocity maps in fig. 5.16 to the surface profiles in fig. 5.17a, suggests the rough surface detectors are able to cope well with the high frequency roughness but some drop out is observed between layers due to the low frequency waviness, corresponding to build layer periodicity. This topographical variation causes deviations in the position of the returned detection beam which are orders of magnitude larger than that caused by surface roughness.

The results presented in fig. 5.17a show the current ability to capture velocity information on as-deposited surfaces. As expected, the signal degrades due to the deflection of the detection beam from the surface profile, as shown in fig. 5.15. In practice it would be possible to employ some form



Figure 5.16: SRAS velocity map captured from the side wall of an as-deposited undeformed specimen. A interferometer based on random quadrature demodulation has been used to over come the 'rough' surface and capture these SAW velocity maps. Grey areas in these figures indicate signal drop-out, the scale bar in (b) indicates 2 mm. In keeping with the results from the prepared surfaces, the presence of the prior- β structure can be discerned most clearly in the 90° propagation direction. The lack of contrast in the 0° direction map is due to the unique velocity surfaces which arise form the limited number of α variants in a parent prior- β grain.

of contour following, allowing the detector system position to be adjusted to maintain a normal to the inspected surface and maintain the correct focal length. However, despite the significant loss of signal in the as-deposited datasets, the information captured is sufficient to still clearly distinguish the unrolled specimen from the rolled specimen. Whilst there is a significant drop-out due to specimen waviness, the result shown in fig. 5.17a suggests that the macro features can still be seen without the need to scan the full area.

Additionally, measurements may not be feasible during the period where the deposited layer cools, due to both the on-going phase change from a primarily β microstructure to primarily α and the change in temperature altering the elastic modulus and subsequent velocity measurement from the specimen. With the benefit of accurate temperature measurement in the deposit, these effects could be decoupled from the SAW velocity; however, in order to make meaningful measurements capturing the effect of the rolling process measurement shall be made at room temperature (rolling



Figure 5.17: (a) Line plots of surface topology in as-deposited undeformed and 75 kN rolled specimens. Data plotted in (a) is extracted from high resolution optical microscopy images shown in (b) undeformed and (c) 75 kN rolled, captured using focus variation microscopy. (b) Changes in colour due to varying oxide thickness. A prior- β grain can be seen to run across the specimen. The scale bar shown in (b) indicates 2 mm.

is applied after each layer of deposition has cooled to ambient). In practice, this avoids the complexities of measuring specimens with a temperature gradient or undergoing phase-change.

Finally, for SRAS to be considered a realistic solution for in-line monitoring the speed of measurement must be considered with respect to the build rate. The WAAM equipment used in this study deposits at a horizontal travel rate of $270 \text{ mm} \cdot \text{min}^{-1}$. If the primary aim of the in-line inspection system is to decide if rolling has enacted grain refinement, then the step size between measurement points can be large given the width of prior- β grains in the undeformed specimen, 2-6 mm in width. The WAAM system takes around 0.45 s to deposit 2 mm therefore a SRAS system should be able to capture acoustic data in an area (horizontal and vertical) sufficiently large as to identify undeformed prior- β grains in this time. The current SRAS instrumentation can capture ~ 2000 points per second, limited by the generation laser repetition rate, thus in this time it would be possible to capture 900 unique acoustic data points. At an acoustic resolution of 100 µm in both axes, an area of 20×45 points or $2 \times 4.5 \text{ mm}$ could be interrogated; this would sufficiently capture the acoustic response to identify large prior- β grains, whilst keeping pace with the deposition process. Thus, from a practical standpoint there is clearly a good opportunity to integrate SRAS in-line without affecting the fabrication process.

5.2 Laser powder bed fusion

This section will go through the L-PBF results in titanium, nickel and silicon steel. This section has been structured in this order in a bid to aid the narrative by commencing with more elementary results before progressing to more complex results towards the end of the chapter. This section will not discuss the time implications and technical challenges to a SRAS inspection in-situ for powder bed fusion, this topic requires a greater breadth of discussion and will be dealt with expressly in chapter 7.

5.2.1 Ti-6Al-4V

SAW velocity measurements

This section will discuss the correlation of results between the microstructure and porosity of Ti-6Al-4V parts fabricated with L-PBF and ultrasonic measurements; fig. 5.18 presents initial results.

Looking first at the SAW velocity map, fig. 5.18a, the first observation is the relative uninteresting nature of these velocity maps, which stems from two sources. Firstly, the grain size is much smaller than the spatial resolution of the system again, such that the stark contrast observed in large grain materials is absent. Secondly, titanium is relatively elastically isotropic, thus variation in the orientation induce comparatively small changes in the SAW velocity. This has meant no spatial inhomogeneities can be discerned from this velocity map. Comparing the measured velocities to the range of SAW velocities as calculated by the forward model, fig. 5.18c, it was seen that the measured velocity is within the normal range for Ti-6Al-4V and the average measured velocity is close the dominant velocity from the model ($\sim 2950 \,\mathrm{ms}^{-1}$).

One of the problems which plague powder bed fusion is the formation of porosity, and therefore, it is worthwhile to examine the optical dataset for discontinuities. Fig. 5.18b presents the optical level across the specimen, for the same area as the velocity map, indications of several surfacebreaking defects can be determined. These are marked in white in the velocity map.

Comparing the SAW velocity and optical datasets, several low velocity locations were seen, many of these surround optical indications not in the optical map, suggesting these are not surface breaking defects. Comparing the velocity of these region low velocity regions ($\leq 2800 \,\mathrm{ms}^{-1}$) to the output of forward model for Ti-6Al-4V, these velocities are towards the lower limit predicted by the forward model. However, from a crystallographic perspective, velocities $\leq 2800 \,\mathrm{ms}^{-1}$ correspond to the region where $\phi_1, \Phi \approx 90^\circ$ only, given this is incompatible with the fibre texture of L-PBF specimens, this suggests these low velocity regions in the SAW velocity map are not a crystallographic effect.

Smith et al. previously observed this slow SAW velocity effect is in similar Ti-6Al-4V specimens,



Figure 5.18: Results in L-PBF Ti-6Al-4V. The main feature from the SAW velocity map (a) is the presence of surface breaking defects (these are also seen in the optical map (b)) and anomalous low velocity regions. The forward model (c) suggests this velocity is not a crystallographic effect. Instead, the presence of shallow subsurface defects induces mode-conversion to the disperse A_0 Lamb wave, where measured velocity corresponds thickness of the 'plate'.

and suggested defects located within a wavelength of the free surface were found to cause mode conversion of the customary Rayleigh wave to a Lamb plate wave, specifically the fundamental antisymmetric mode, A_0 . Fig. 5.18d plots the dispersion curves of the fundamental modes in titanium. In specimens, such as titanium with a small velocity range, subsurface features are distinguished from variations in the crystalline texture by the large shifts in velocity. Dispersion curves have been calculated using the Disperse computation package [293]. this is a potentially powerful result as it allows subsurface defects to be detected by SRAS, furthermore by comparing the measured velocity to the thickness on the dispersion curve it would - in theory - be possible to also extract the depth of these features. This topic will be briefly revisited when considering nickel specimens in the next section.

5.2.2 Nickel superalloy

SAW velocity measurements

A sample has been manufactured using the common aerospace nickel super alloy, Inconel 718 often favoured in gas turbine applications for its temperature, corrosion and creep resistance. Samples were manufactured using a Relaizer SLM50 build platform. An island scan pattern was used, in effect this means the specimens were formed of an array of smaller squares.



Figure 5.19: (a) SRAS velocity map of Inconel 718 triangle sample, which was fabricated using an island scan strategy, as shown in (b).

Fig. shows the (a) SAW velocity map and (b) a diagram of the hatching strategy used in the build. It is clear that the islands have created a strong microstructral texture discontinuity that can be clearly observed in the velocity map. These texture boundaries correspond to the interfaces between hatched areas, the acoustic image bares a remarkable similarity to the photograph of the as-deposited sample. This is a clear exemplification of the impact varying hatching patterns can have on component microstructure. Additionally, the islands located on the perimeter of the triangle have shorter scan vectors than the bulk of the material which may be the cause of the

repeated high velocity clustering within these areas. A smaller vector length will reduce the time between the melting of successive vectors, which will promote a different solidification to that seen in the bulk.

Elastic constants for Inconel 718 are not widely available, thus simulations were performed for a number of similar nickel superalloys and pure nickel. For brevity, only pure nickel is shown in fig. 5.20c, as this was found to be a good approximation for all nickel alloys (elastic constants from Salama and Alers [294]). In all simulated nickels, faster velocities suggested an orientation towards [001], and slower moving towards [111]. Thus from comparing the output of the forward model for nickel to the SAW velocity map (fig. 5.19a) it can be said that the interface zones are closer to a [001] orientation, whilst the island areas are closer to a [111] orientation. This is in good agreement with the findings of Carter et al. [295], who used EBSD to study the influence of scanning interface regions on the formation of nickel superalloys.

EBSD comparison

To further investigate the crystallographic information contained within the SAW velocity map, EBSD imaging was carried on the specimen as a comparator. Fig. 5.20 compares the SRAS velocity map and EBSD orientation information for the tip of the triangle. The EBSD data demonstrated the tip of the triangle consists primarily of colonies at [001] and [101] orientation. Remembering the spatial resolution of the SRAS velocity map in question is ~ 100 μ m, compared to the nanometre scale resolution of EBSD, the datasets compare relatively well. Most notably, both suggested a strong [001] orientation at the very tip of the triangle. Moving further up, both suggest a similar formation of [101] towards the perimeter before transitioning to [001], this may be an effect of a different cooling profile due to be being at the edge of the specimen.

Surface and sub-surface porosity extraction

Interestingly, the SAW velocity map presented in fig. 5.20a does not immediately suggest the presence of subsurface defects. As the presence of several surface breaking features can be seen in the velocity map (white regions), it is therefore unlikely that the volume was free from such defects. Instead, measurement of subsurface features is obscured due to the crystalline anisotropy of the nickel. Previous samples were fabricated in Ti-6Al-4V, which is characterised by small variations in velocity due to texture. In contrast to the small SAW variations in texture in titanium, a single grain in nickel can have variations of $\sim 1000 \,\mathrm{ms}^{-1}$, across a single plane. As the depth of the media approaches twice the acoustic wavelength, the velocity tends towards the normal Rayleigh wave velocity. In the case of nickel, even large shifts in velocity will often fall within the normal Rayleigh wave velocity range, meaning subsurface defects are obscured by normal microstructural variations



Figure 5.20: Crystallography of the triangle tip. (a) and (b) compare the SRAS velocity map and EBSD captured IPF from the tip of the triangle specimen. Whilst the SRAS velocity maps does not directly reveal the crystallographic orientation, the general texture can be inferred by comparing the measured velocity to the average across the orientation space (c).

in the velocity map. This suggests that merely looking for large shifts in the SAW velocity is an unreliable method for the detection of subsurface defects in AM.

5.2.3 Fe-Si

High silicon steel has provided a final fascinating material to study in this chapter, thanks to its strong - and unusual - microstructural texturing. As detailed in §2.2.3, within the domain of soft magnetic materials lies silicon steel, iron alloys with a silicon content between 2-7%. Within the family of silicon steels, Fe-6.9Si alloys have a high potential in magnetic device applications due to their high electrical resistivity, near-zero magnetostriction, and low magnetocrystalline anisotropy. Magnetic anisotropy is the dependence of the magnetic properties on the direction of the applied field with respect to the crystal lattice. Depending on the orientation of the magnetic field, one would need a lower or higher magnetic field to reach the saturation magnetisation. [001] directions are the easy axes of magnetisation in this material as this is the least densely packed direction. Previous research in the manufacture of Fe-Si materials by powder bed fusion has shown the build parameters can be used to control the formation of texture relative to the build direction, allowing the easy axes of magnetisation to be aligned, making the component suitable for magnetic applications [148]. The formation of this texture is essential to ensuring the function of potential Fe-6.9Si components; however, this texture has only been observed in a small region in the bulk of the material by EBSD previously. Therefore, the challenge is to detect the presence of a cube texture formation across the specimen using SRAS.

As disused through this chapter, left unperturbed additive components will tend to form a fibre texture, aligning the [001] axes with the build direction. In reference to fig. 5.21, this results in a predominance of the (001) plane (red plane) when examining the x - y plane. However, the rotation of the crystal lattice around the z-axis is not usually controlled. The ingenious feature of texture control in Fe-6.9Si is also to control this rotation, such that the (101) plane (green plane) is aligned with the x-axis of the build - this is usually called a cube texture.



Figure 5.21: Definition of textures and relation to SAW velocities for silicon steel.

Three specimens were fabricated for study in this section, using the build parameters listed in table 5.3. It was expected that the texture control would not be initiated in cube 1, meaning only the typical fibre texture would be observed. Whereas in specimen 3, substantial texture control would have been generated, causing the formation of cube textures, with the alignment of the (001) plane to the x - y plane and the (101) plane to the x - z and y - z planes (subsequent references to these planes will only mention the x - z plane for brevity, but the analysis equally applicable to the y - z plane). It was unclear what microstructure would result in specimen 2, where intermediate processing parameters have been used.

Again, the forward model was utilised to comment on the crystallographic texture of the specimens. Unfortunately, measured elastic constants for Fe-6.9Si are not available and were interpolated from measurements in Fe-6.3Si and Fe-8.5Si, as reported by Machova et al. [296]. The SAW

| Specimen | Hatch spacing | Scan speed | Exposure time | Point distance | Power | Energy |
|----------|---------------|-----------------|---------------|----------------|-------|---------------|
| | μm | ${\rm ms}^{-1}$ | μs | μm | W | $\rm Jm^{-1}$ |
| 1 | | 0.50 | 40 | | 70 | 140 |
| 2 | 60 | 0.25 | 60 | 20 | 70 | 280 |
| 3 | | 0.50 | 40 | | 90 | 180 |

Table 5.3: Build parameters for Fe-6.9Si specimens used in this study. It is expected these parameters will generate a cube texture in specimen 1 and [001] fibre-texture in specimen 3.

velocity on the three principal planes for the three elastic constant sets of FeSi along with pure iron are compared in fig. 5.22, the pure silicon model is not shown as this is significantly faster $(\sim 1000 \,\mathrm{ms^{-1}}$ on all planes). It is seen that the difference between the two known constant sets (Fe-6.3Si and Fe-8.5Si) is $\sim 50 \,\mathrm{ms}^{-1}$, and as the Fe-Si walls between these two, the error is expected to be no greater than $\sim 50 \,\mathrm{ms}^{-1}$. It is also noted that as all materials have an anisotropic ratio, η , of > 1 then all velocity surfaces have similar shapes (see §4.2.3) for more discussion on this), meaning some comment on the change of SAW velocity with propagation angle is still possible without knowing the precise elastic constants. Therefore, these interpolated values are expected to accurately represent the Fe-6.9Si material and are sufficient to provide insight into the material crystallography. As an aside, it is interesting to note the increase in silicon content slows the velocity of the RSAW, but on the (001) plane, the velocity of the PSAW increases with the silicon content. This is due to the value of C_{11} and C_{12} decreasing, but the shear parameter C_{44} increasing.

| Material | C ₁₁ | C ₁₂ | C ₄₄ | ρ | Ref |
|--------------------------|-----------------|-----------------|-----------------|----------------|-------|
| | | GPa | | $\rm kgm^{-3}$ | |
| Iron | 230.4 | 137.1 | 115.9 | 7874 | [297] |
| Silicon | 165.0 | 64.0 | 79.2 | 8960 | [298] |
| Iron-Silicon $(6.9\%)^*$ | 217.0 | 129.0 | 137.0 | 7600 | [296] |

Table 5.4: Elastic constants of pure iron, silicon and Fe-Si.

* interpolated from measurements in 6.3 and 8.6% silicon content specimens.



Figure 5.22: Comparison of the velocity surfaces on the three principal planes, in the alloy 6.9Si and the two constituent elements iron and silicon.

x - y plane measurements

The x - y plane is first examined in fig. 5.23, comparing specimens 1 and 3. The first noticeable feature is the contrast in measured SAW velocity observed at the edge of the specimen. This stems from the contour scan strategy employed in the fabrication of the specimens, where the bulk of the material is hatched at an angle of, and then the outer border area hatched with parallel vectors. In both figures, the thickness of these anomalous regions is measured to be approximately 450 μ m, which correlates with the thickness of the build contour.



Figure 5.23: SAW velocity maps from the x - y plane of (a) specimen 2 and (b) specimen 3.



Figure 5.24: y-direction inverse pole figure from edge of cube 3, captured by EBSD.

Moving to the consideration of the bulk region, both cubes exhibit a dominance of a fast SAW velocity ($\sim 3300 \,\mathrm{ms}^{-1}$). To extract further information, the velocities may be compared to the velocity model for Fe-6.9Si, fig. 5.25. Again, velocities on the (001) plane are on average the fastest across the orientation space and exhibit peaks at $\sim 2800 \,\mathrm{ms}^{-1}$ and $\sim 3300 \,\mathrm{ms}^{-1}$, the latter of which is in good agreement with the measured velocity of these specimen. The range of velocities observed in the bulk is commensurate with an (001) fibre orientation. In contrast, a velocity closer to (111) is slower, suggesting the different scan strategy in the border region has prevented the formation of a fibre texture.

For comparison, EBSD captured inverse pole figure from specimen 3 is shown in fig. 5.24. In this figure, the normal direction is in the z-direction. i.e. a red pixel indicates where [001] crystal direction is aligned with z-direction (build direction), this is the hallmark of a fibre texture. This is in good agreement with the SRAS velocity measurements from the bulk. Similarly, an anomalous region can be seen on the edge of the pole figure, with orientations closer to (101) and primarily (111). This is the border region shown in the velocity maps, and its presence in the EBSD dataset is evidential for the lower SAW velocity measurement indicating the lack of fibre texture.

This first analysis has shown the presence of a fibre texture may be confirmed by the SAW velocity measurements; we may now consider the additional constraints of a cube texture. The velocity distribution from the bulk regions shows a clear dominance of velocities around 3200 ms^{-1} , from fig. 5.25 we see around in the 001 plane a second peak around 2800 ms^{-1} which does not appear to be represented in the experimental data. The 3200 ms^{-1} peak is due to the PSAW centred at 45° , whereas 2800 ms^{-1} is centred at $0^{\circ}/90^{\circ}$. For a cube texture, we should see the 45° direction of the (001) plane perpendicular to the x and y axes of the build cube. Therefore, to test this hypothesis, two further phenomena would be expected. The 45° prorogation direction scan should align with 0° of the (001) plane, generating a peak at 2800 ms^{-1} . This data is presented in fig. 5.26 and is seen to agree well with the hypothesis, confirming the alignment of the rotation of the crystal lattice to 45° , in turn, confirming the presence of a cube texture. The exact velocity of the peaks do not match exactly (discrepancy of 80 ms^{-1}), which is most likely due to the uncertainty in the



Figure 5.25: SAW velocity distributions in Fe-6.9Si. The velocity increases towards the (001) plane and has a velocity distribution quite distinct from that near the (111) plane.

elastic constants and density input to the forward model to predict the velocities.

x-z plane measurements

The vertical planes of the specimens were also sectioned and exposed for analysis. Fig. 5.27 compares the SAW velocity maps from the three specimens. Taking the specimens in turn, cube one exhibits few features of note, indeed the specimen appears to have little spatial coherence in the SAW velocity, suggesting there is not a dominant crystal plane under the generation patch. In specimen two, the majority of the map remains disordered, however, there are initial indications of texturing (regions of consistent velocity, as marked in the figure). This is followed by specimen 3, where significant velocity clustering is observed, suggestive of a predominant orientation (i.e. always seeing the same crystallographic plane). Evident elongation of grains and an increase in the velocity can also be seen from specimen 1 to specimen 3.

The data presented in fig. 5.27 was captured at the 0° propagation direction. By capturing at multiple directions, the averaged velocity surface can be interrogated. This is useful as where a strong cube texture exists (specimen 3), a velocity curve similar to the (101) line shown in fig. 5.21 would be expected. Where there is a fibre texture (specimen 1), then the lattice is free to rotate relative to the x - z plane, the result is likely to be random, and thus no real relationship with varying propagation direction is expected. The velocity surface between 0° and 180° for specimens 1 (a) and 3 (b) are presented in fig. 5.28. It is seen that the SAW velocity of specimen



Figure 5.26: (a) Velocity distributions from bulk area of specimen 3 at 0° (green) and 45° (purple) propagation directions. (b) Comparing the velocity distributions to the velocity surface of the (101) plane shows the 0° of the specimen is aligned with the 45° rotation of plane, and vice-versa.



Figure 5.27: SRAS velocity maps from three Si-Fe L-PBF cubes fabricated with different build powers. The formation of cube texture is expected to increase from specimen 1 (a) to 3(c), with specimen 2(b), expected to fall somewhere between. The large features in specimen 3 (c) are indicative of many grains have similar crystallographic orientations. The scale bar in (c) indicates 1 mm.

193

3 is highly sensitive to the propagation angle. The measured velocity surface is seen to follow the calculated (001) surface well, moving from a high velocity at 0° to a low velocity at 90° , with a range of $3200 \,\mathrm{ms}^{-1}$. In contrast, the median velocity in specimen 1 varies between just $2600 \,\mathrm{ms}^{-1}$ and $2760 \,\mathrm{ms}^{-1}$ over the same angular range, suggesting this specimen is not as sensitive to propagation angle. Furthermore, the measured average velocity surface bears little resemblance to the theoretical (101) velocity surface.



Figure 5.28: Average velocity surfaces in the x - z plane for specimen 1 (a) and specimen 3 (b). When a cube texture exists, the measured SAW velocity would be expected to resemble the velocity surface of the (101) plane.

Simulated measurements

For a final study, it is possible to compare experimentally measured velocity maps with simulated datasets, following the procedure defined in §4.4.3. Firstly, inverse pole figures of these orientation distribution functions are given in fig 5.29. These represent strong and weak cube textures and are in good agreement with experimental data published in the existing literature. The experimental and simulated velocities of the three measurement planes are then given in fig. 5.30 for specimen 3. Considering first the x - y plane, in both specimens, there is seen to be a dominance of velocities $\sim 3300 \text{ ms}^{-1}$, representing the (001) plane. The experimental data has a greater number of lower velocity regions ($\sim 2600 \text{ ms}^{-1}$), suggesting the actual texture is weaker than the simulated case.



Figure 5.29: Pole figures taken from ODF generated from a (a) strong cube texture (b) weak cube texture.

Two points are worthy of note in the x-z plane, the grain morphology and the spatial resolution of the instrument. Firstly, the simulated data has followed the Voronoi tessellation method, as detailed in §4.4.3, which can be used to create an equiaxed microstructure. This is a sensible approach for the x - y plane but does not capture the elongated grains seen in the x - z planes. Secondly, the simulated dataset displays the velocity of each individual orientation, however as the grain size in the actual specimen is beneath the spatial resolution of the instrument, the measured SAW velocity is an averaged response of the grains under the generation patch. As such, the data in the x - z plane of the simulated cube has been elongated and blurred to better replicate measurement. This plane presents large regions of velocity ~ 3100 ms⁻¹, as seen in the experimental data, with the intermediate velocity (green regions) now seen, which are not present in the x - z plane, in better agreement with the experimental dataset.

5.3 Summary

This chapter has presented a series of results in additive manufacturing specimens, by firstly the wire-arc method, then discussing results in samples preprepared by laser powder bed fusion in three interesting and industrially relevant materials. Across materials and techniques, attention has been paid to understanding the microstructural origin of the measured SAW velocity, the purposing being such measurements can be applied as quality control and to inform on the additive



Figure 5.30: Isometric cubes of SAW velocity in specimen 3 of Fe-6.9Si specimens. The border region has been excluded from experimental datasets. The x - z plane of the simulated data has been spatially averaged in order to simulate a spatial resolution similar to the EMDA system.

build process. In all of the materials studied in this chapter, the grain size is below the spatial resolution of the SRAS instrument, and therefore the orientation cannot be recorded through the inversion method outlined in chapter 4. The primary focus of this chapter has been to overcome this limitation by leveraging the forward model in a multitude of new ways to reveal microstructure information still. While the exact orientation cannot be determined, in the case of the two extended studies in this chapter (WAAM Ti-6Al-4V and Fe-6.9Si), the formation of texture is of more significant concern to the performance of the component than the orientation of the individual crystal.

In wire-arc additive manufacture, the process of inter-pass rolling has previously been shown to drastically improve the functional mechanical properties of components. By mapping the texture of WAAM components through the generation and detection of surface acoustic waves, SRAS scan can be used to probe the changes in texture. This work constitutes a primary step towards a viable quality control tool requisite to the development of WAAM as an industrial solution.

Specifically, the work in wire-arc specimens has demonstrated seven notable outcomes. Firstly, the ability of SRAS to detect grain refinement in both side-wall and cross-section WAAM specimens, allowing the differentiation between undeformed and rolled specimens. For the first time, SRAS has been utilised on a dual-phase material. It has been shown inherited texturing from the parent β -phase to the child α phase leads to regions with distinct slowness surfaces, which can be

correlated to prior- β grain growth.

Variations in measured acoustic wave velocity with varying propagation angle have been shown to compare well with Young's modulus extracted from averaged elasticity tensors and are indicative of material anisotropy. Further useful build information, such as the minimum stiffness direction and demonstration of strong crystallographic alignment in the substrate, can be obtained using SRAS. The development of a metric that allows a direct correlation between rolling force and the acoustic response by measuring the size of distinct regions. This was shown to differentiate between the effect of varying rolling force, as well as undeformed and rolled more generally. This should prove to be a fundamental tool for allowing closed-loop control from an in-line system.

The rates of deposition seen in WAAM pair well with the acquisition speed of a theoretical in-line SRAS system. Identification of large prior- β grains would be possible without impact on the rate of manufacture. A preliminary demonstration of the capability of a 'rough' surface detector system to unravel the speckle pattern, allowing acoustic images to be captured on an as-deposited WAAM surface, was presented. This addresses many of the practical difficulties of taking such measurements in-situ.

Laser powder bed fusion has provided a more complex process to investigate. The 'low-hanging fruit' of using optical measurements to observe surface-breaking porosity was used in nickel and titanium specimens. In line with the prior art, subsurface defects were observed in the titanium specimen due to mode conversion. However, this effect was not immediately apparent in the nickel specimen. The dispersion curves for nickel suggested that the variations in SAW velocity due to subsurface defects are likely to be swamped by the significant variations in SAW velocity due to the high anisotropy of nickel. This suggests relying on the waves of slower mode converted waves for the detection of subsurface defects is not a robust approach and that, more generally, SRAS is not best suited to the detection of void-like defects.

The SAW velocity captured from a single velocity map was shown to highlight spatial inhomogeneities within the build which formed due to the build process. It was shown these could be related to the orientation by comparing the measured velocity to the output from the forward model. The detection of anomalous texture at the interface between scan areas (IN718) or the edge of the specimen (Fe-6.9Si) are examples of this. It has been shown that examining the could indicate the formation of a fibre texture in all specimens. Multiple propagation directions were then used to confirm the presence of a cube texture in specimen 3 and a lack thereof in specimen 1.

The perhaps natural approach is to study and inspect the x - y plane as this presents the most recently deposited layer; in cases where the build strategy is selected to generate crystallographic interfaces within the x - y plane, this can offer useful information. This has been demonstrated, such as the hatching in the nickel or contour border in the silicon specimens. However, due to the tendency for strong fibre texturing in additive manufacture, the x - y plane can exhibit very little velocity contrast, as shown in the Ti-4Al-4V specimens and the bulk of the Fe-6.9Si samples. This lack of contrast is not necessarily an issue, as it can indicate the formation of a strong texture, but it is worth remarking upon. Prior to this work, the study of the x - z plane has been neglected; through the confirmation of the cube texture in the Fe-6.9Si specimen, it has been demonstrated that in certain materials, these planes can also reveal important microstructure information.

Chapter 6

Finding the single-crystal elasticity and the orientation

6.1 Introduction

The previous chapter investigated the wealth of information contained in spatially resolved acoustic spectroscopy (SRAS) measurements of additively manufactured specimens, using the interaction between the generated surface waves and the crystalline sample to inform on the microstructure and, in turn, the build process. The size of crystallographic features in the specimens studied thus far are relatively small compared to the spatial resolution of the EMDA system; therefore, Chapter 4 focused on extracting crystallographic information without the use of the inversion method introduced in §4.3. However, the power of SRAS is traditionally derived from its ability to determine the crystallographic orientation, which will be an important capability of a viable quality control tool within additive manufacturing. This then naturally poses the question of how accurately and reliably the orientation can be determined. This chapter briefly sets the study of additive manufacturing specimens to one side, as larger grain specimens of greater availability are studied as a proof-of-concept. However, the work is eminently applicable to additive manufacturing.

The first half of this chapter ($\S6.2$) is concerned with understanding the accuracy of crystallographic orientation determination in hexagonal crystal systems, focussing on the ubiquitous titanium. Hexagonal materials are particularly worthy of discussion as they exhibit transverse isotropy - the elastic properties are direction-independent in-plane but vary through the thickness. The modelling techniques detailed in \$4.4 are utilised to study the impact of crystal anisotropy (\$6.2.2), velocity resolution (\$6.2.3) and the number of scanned propagations directions (\$6.2.4) on orientation determination to be investigated for a wide range of experimental parameters. Experimental results are then presented in §6.2.5 and compared to EBSD, with the determined errors compared to expected errors for the given experimental parameters. This section concludes with the importance of known elastic constants to the accurate determination of crystallographic orientation.

Observing the impact of elastic constants leads into the latter half of this chapter, which will describe how the measurement of velocity surface spectra may be used to simultaneously determine the crystallographic orientation and elastic constants in polycrystalline specimens. A brute-force method is prosed in §6.3.1, based on the forward and inversion models previously used for orientation determination. The influence of perturbations in the elastic constants on the surface acoustic wave (SAW) velocity is then studied in §6.3.2 before experimental results in nickel, a nickel superalloy and titanium are presented in §6.3.3, with orientation results compared to EBSD and elastic constants to existing literature values. The experimental results include a previously problematic specimen, where the orientation could not be accurately determined using elastic constants from previously literature. Finally, the solution shape is visualised for both cubic and hexagonal crystal systems, allowing errors in the determination of elastic constants to be correlated with the precision of SAW velocity measurement in §6.3.4.

6.2 Accuracy of orientation determination in titanium

SRAS measurements have previously been used to recover the exact grain orientations in materials of cubic crystal structure [252, 255]. This section extends this approach to hexagonal materials, exploring the limitations and accuracy of this methodology. The prevalence of titanium throughout safety-critical sectors such as aerospace and healthcare gives an industrial context to the need to understand the ability to map crystallographic texture in hexagonal materials. Furthermore, hexagonal crystal structures exhibit unique acoustic properties, making them particularly worthy of investigation.

Common engineering metals typically have one of five crystal structures: cubic, hexagonal, tetragonal, rhombohedral and orthorhombic. The resulting shape of the elasticity tensor for an exemplary material of each crystal class is shown in fig. 6.1. As can be seen in fig. 6.1b hexagonal tensors exhibit mechanical transverse isotropy; rotations about the c-axis do not change the directional mechanical properties of the crystal. In contrast, all other crystal class systems, including the previously reported cubic, show two axes where the elasticity varies. Given the relationship between material elasticity and SAW wave velocity, this phenomenon will undoubtedly impact the ability to determine the orientation of hexagonal crystal structures. Furthermore, in addition to

the Rayleigh SAW, a supersonic PSAW wave mode is known to propagate on certain planes in hexagonal crystals, as defined by Ting and Barnett [276]. With these unique properties in mind, it is of interest to understand the ability of SAW measurements to map the orientation of hexagonal crystals.



Figure 6.1: Stereographic projections of directional Young's modulus for different crystal classes on irreducible area of unit sphere, these are the keys used for inverse pole figure representation of the given crystal.

The impact of the symmetry of the hexagonal crystal is seen most clearly by visualising the velocity surfaces of titanium. Fig. 6.2a and b plot the acoustic velocity surfaces seen on the principal planes using three sets of elastic constants which have previously been reported for hcp titanium, given in table 6.1. The constants sets from the work of Fisher et al. and Hearmon gave similar results on both principal planes (difference of 3 ms^{-1} , on average), the difference between the models become more clear closer to the prism plane ($\sim 10 \text{ ms}^{-1}$). As the mode transition occurs at difference location then difference between the dominant velocity of each elastic constant set can exceed 40 ms^{-1} for certain orientations. The Ogi constants set produced a notably faster SAW velocity. It should be noted these three elastic constant sets vary significantly from those measured by Kim and Rokhlin for Ti-alloys [240]. The elastic constants sets for the Ti-alloys produce velocity surfaces slower than those shown in fig. 6.2. This demonstrate SAW velocity models for pure-Ti is not representative of Ti-based alloys, elastic constants must be specific to the material under evaluation.

| Reference | c_{11} | c_{12} | c_{13} | c_{33} | c_{44} | ρ (kg m ⁻³) |
|------------------------|-----------------|----------|----------|-----------------|----------|------------------------------|
| Ori et al [200] | (GI a) 162.6 | (01 a) | (GI a) | (GI a) 195.0 | (61 a) | (Kg III) 4420 |
| Ugi et al. [299] | 163.6 | 92.3 | 67.9 | 185.2 | 47.1 | 4429 |
| Hearmon et al. $[275]$ | 160.0 | 90.0 | 66.0 | 181.0 | 46.5 | 4510 |
| Fisher et al. $[300]$ | 162.4 | 92.0 | 69.0 | 180.7 | 46.7 | 4506 |

Table 6.1: Elastic constants sets of titanium from literature, these are named after the author of the publication in which these values are first reported. This table is a truncated version of table B.1, in appendix B.

As seen in Fig. 6.1, hexagonal tensors exhibit transverse isotropy. Consequently this leads to a few cases where elastic properties, and subsequently SAW velocity, do not vary with the crystal orientation. For example, the velocity surface seen on the basal plane is isotropic - rotation of the velocity surface does not changed the SAW velocity (fig. 6.2b - c), the characteristic acoustic velocity surface of the principal prism planes are identical - the velocity surface is insensitive to changes in ϕ_2 . This is to say, linear acoustic techniques cannot determine ϕ_2 . A final special case occurs when the Euler angle Φ of two planes are the same and both plane normals are in the same plane with axis Z, the Euler angle ϕ_1 of these planes can be narrowed down to either 0° or 180°.

6.2.1 Representing orientation data

Inverse pole figures (IPF) are the most common method of presenting crystalline orientation. Usually, for EBSD captured datasets, the IPF shows the combination of Euler angles Φ and ϕ_2 ; however, fig. 6.2 has demonstrated that ϕ_2 cannot be determined from the velocity surface. Therefore, an alternative representation is needed for Euler angles ϕ_1 and Φ as resolvable by linear acoustic techniques. Fig. 6.3, compares a traditional IPF (Euler angles Φ and ϕ_2) and the proposed c-axis map (Euler angles ϕ_1 and Φ). Both images have been generated from the same dataset as captured by EBSD. The c-axis map scales linearly for both ϕ_1 and Φ . Changes in hue represent variations in ϕ_1 . Similarly, changes in saturation represent variations in Φ . The key in the c-axis map shows the orientation of the hexagonal crystal.

There are several experimental factors that may affect the accuracy of crystallographic determination: lattice anisotropy, velocity resolution, systematic measurement error, signal-to-noise ratio (SNR), number of measured SAW propagation directions and the interval between scan directions. The following sections will consider each of these in isolation to understand their impact on orientation determination accuracy.



planes, showing mode transition from supersonic to Rayleigh wave, and isotropic nature near basal plane.

Figure 6.2: SAW properties of titanium.



Figure 6.3: Titanium orientation results by EBSD: IPF (left, Euler angles Φ and ϕ_2) and *c*-axis map (right, Euler angles ϕ_1 and Φ). These figures are plotted from the same dataset. The Φ orientation can be determined in both plots, for example the basal plane is seen as red in the IPF and white in the c-axis map.

6.2.2 Orientation dependent accuracy

It was firstly of interest to consider the fundamental limitations imposed by the acoustic properties of the hexagonal crystal. Fig. 6.4a, shows the orientation accuracy as a function of SNR and Φ angle - the accuracy of orientation recovery is dependent on the underlying orientation. When $\Phi \ge 28^{\circ}$ the ϕ_1 error is < 1°, for all simulated SNRs. When $\Phi \le 28^{\circ}$ the ϕ_1 error can increase to > 5°, due to the isotropy of the velocity surface before the mode transition. Increasing the SNR of measurement was found to increase the range of Φ for which ϕ_1 can be determined. For example, at 0.1 dB large errors are seen $\Phi \le 28^{\circ}$, but this reduces to $\Phi \le 18^{\circ}$ for 9.7 dB.

Errors in Φ , fig. 6.4b, were found to be lowest in the region of $28^{\circ} < \Phi < 60^{\circ}$. The $0^{\circ} < \Phi \leq 28^{\circ}$ results appear to match very well with the results from ϕ_1 error, suggesting the quasi-isotropy of the velocity surfaces (compared to the velocity sensitivity) in this region is again the cause of the determination errors. It was unclear at this stage why determination deteriorated for $60^{\circ} < \Phi \leq 90^{\circ}$.

In fig. 6.4, the orientation determination accuracy is dependent upon the position within the orientation space. The uniqueness of the velocity surface is not consistent - highly unique surfaces increase the likelihood of the overlap function finding the true orientation. The velocity gradient, fig. 6.5 is a good representation of this; where velocity gradients are large, velocity surfaces are more unique. Fig. 6.5a and fig. 6.5b show the velocity gradients across ϕ_1 and Φ respectively.

Fig. 6.5b explains the improvement of determination in Φ in the region $28^{\circ} \leq \Phi \leq 85^{\circ}$, as seen in fig. 6.4b. The presence of two SAW modes in this region, and thus a mode transition gives large velocity gradients (> $10 \,\mathrm{ms}^{-1}$). Plotting the mean gradient shows the change in velocity



Figure 6.4: Simulation results which show the error of angles ϕ_1 and Φ with respect to SNR and the crystallographic plane given by Φ . 30-fringe (150 ms⁻¹ FWHM) generation pattern has been used in this simulation. The simulation process used here is defined in §4.4. The SNR is originally defined by the perpetuation in the acoustic waveform, then converted to dB.

between planes is small outside of this range, making the differentiation between neighbouring planes velocity surfaces increasingly difficult. Similarly, fig. 6.5a again highlights this isotropy before the mode transition. After this increasingly strong variations occur around $\phi_1 = 45^{\circ}$ and $\phi_1 = 135^{\circ}$. Fig. 6.5c is the product of fig. 6.5a and fig. 6.5b, giving an effective sensitivity map regions with high velocity gradients are easier to determine. When sampling a velocity surface it is important to capture ϕ_1 angles with large velocity gradients.

It should now be clear that orientation accuracy is highly dependent on the anisotropy of the crystal. Hence, the data thus far applies only to hcp titanium and other hexagonal materials such as zinc will have a different characteristic accuracy response, which could be found through the methodology presented in this work. To show this, fig. 6.6 plots the velocity variation across each plane for $0 \leq \Phi \leq 90$. The variation was calculated by Equation 6.1.

$$\Delta v(\Phi) = \frac{\max(v(\phi_1)) - \min(v(\phi_1))}{\frac{1}{N} \sum_{i=1}^{N} v(\phi_1)}$$
(6.1)



(a) Velocity gradient across ϕ_1 . (b) Velocity gradient across Φ . (c) Combined velocity gradient effective sensitivity map.

Figure 6.5: Velocity gradients across the orientation space. These are effectively sensitivity maps, showing the ease of determining the orientation by SAW measurement across the orientation space.

Fig. 6.6, the distinct mode transition in titanium can be seen at 28° and later around 45° in both Magnesium and Ruthenium. In Zinc however, no dominant supersonic mode is seen to propagate, instead the pure Rayleigh mode is seen across all planes. The large variations across the plane means orientation determination in zinc is relatively easier than titanium, and magnesium would be particularly difficult outside of the range $40^{\circ} \leq \Phi \leq 50^{\circ}$. Note, all materials show an isotropic basal plane as this is a shared property of all hexagonal crystals, imposed by the symmetry of the elasticity tensor. These curves inform on the velocity resolution that would be required experimentally to determine the crystalline orientation in each material.

6.2.3 Velocity resolution

If we consider next the most elementary experimental parameter, how well the SAW velocity can be measured. Two issues generally affect the velocity resolution in SAW measurements: 1) the measurement precision and 2) measurement accuracy. Considering first the measurement precision, this can be equated to the signal bandwidth. Fig. 6.7a shows the relationship between angular error in determination of Φ and the FWHM, for six levels of noise. The result shows an exponential increase error with increasing FWHM. For all non-zero values of noise, there was error (> 1°) when the FWHM is greater than 400 ms⁻¹. It can be seen that reducing signal bandwidth below 20 ms⁻¹ had a marginal impact on the determination precision and for most levels of noise 150 ms⁻¹ is sufficient to match the resolution of the forward model. For non-zero



Figure 6.6: Comparison of the SAW velocity variation across planes in four common hexagonal materials. As discussed, the distinct mode transition in titanium can be seen at around 28°. In Zinc however no dominant supersonic mode is seen to propagate, instead the pure Rayleigh mode is seen across all planes. Elastic constants for zinc, magnesium and ruthenium from [301, 302, 303].

levels of noise, errors converged with increasing FWHM. For FWHM greater than 1000 ms^{-1} all levels of SNR gave errors > 2°, rapidly increasing to 10° by 2000 ms^{-1} FWHM. For context, typical SRAS experimental measurements have a velocity bandwidth of 400 to 800 ms^{-1} [253] as indicated in the figure, but SNR can be controlled by changing the number of fringes in the generation patch and/or using temporal averaging.

Fig. 6.7b shows the impact of a systematic error in the velocity measurement. Velocity errors of $\geq \pm 1 \%$ have a catastrophic effect on the angle determination, inducing errors of $\geq 5^{\circ}$ in Φ determination. Interestingly, the overlap function appears more tolerant to positive velocity errors. Signals of increasing bandwidth show a improved robustness to systematic velocity errors as expected, but as discussed in fig. 6.7a, increasing bandwidth also reduces overall accuracy. This trade-off is a crucial consideration of an experimental system. In summary, the highest orientation accuracy can be achieved with very narrowband signals, however such measurements can be corrupted by even small systematic velocity errors.

6.2.4 Number of measured propagation directions and interval

The velocity surface of hexagonal crystals has both mirror and inversion symmetry, as discussed in the authors' previous paper [304]. According to the Nyquist theory, the minimum number of



(a) Effect of varying signal bandwidth - representative to velocity measurement precision, on determination of Φ .

(b) Effect of a velocity measurement error on determination of Φ .

Figure 6.7: Effect of systematic error in velocity measurement on the determination of Φ .
the measured directions to determine orientation is two, for hexagonal structure crystals. Fig. 6.8 shows the simulation results, also using 30 fringes (150 ms⁻¹ FWHM) in the generation pattern: the angle Φ 's accuracy varies to both the number of the directions scanned and the data SNR as expected, because the Φ is determined by the shape of the slowness surface.



Figure 6.8: Simulation results which show the error of angles ϕ_1 and Φ with respect to SNR and the number of scanned directions. 30-fringe (150 ms⁻¹ FWHM) generation pattern has been used in this simulation.

The error of Euler angle ϕ_1 increased dramatically when only two directions were used, due to the inversion and reflection symmetry across all planes in hexagonal crystals. Hence at a certain rotation of the plane, velocities from two orthogonal directions are not enough to reliably identify Euler angle ϕ_1 . In this case, the error may be close to 0° but is equally likely to be close to 180°. This poses the question, if only two-directions are scanned, which combination of the two directions is most efficient? The impact of varying the scan angle interval on the determination of ϕ_1 and Φ is illustrated in Fig. 6.9. It is seen that scan interval has little impact on the determination of the both ϕ_1 and Φ , so long as scans are not orthogonal.



Figure 6.9: Error bar of Euler angles against the intervals when using two directions. Simulation at 30 fringes $(150 \text{ ms}^{-1} \text{ FWHM})$.

6.2.5 Orientation accuracy - experimental

In this section, a series of SRAS datasets have been captured experimentally in a grade 1 commercially pure titanium (> 99% Ti) specimen of α -phase with hexagonal close packed structure. Specimens were ground and polished on the surface to be probed according to standard preparation for titanium [305], followed by a vibratory polish for 12 h in colloidal silica to prepare the samples for EBSD mapping, giving a final surface roughness of $R_a < 100 \text{ nm}$ - expressed by the arithmetic mean. Further information on the specific material studied, its preparation and EBSD scanning procedure can be found in [306].

Full orientation map at standard scan parameters

The specimen was raster scanned to build up a velocity map, the process was then repeated a varying propagation directions so as to build-up a velocity surface for each point. Detailed scan parameters are given in table 6.2. Based on these parameters, and using the models above, the predicted errors of this experimental study are expected to be $\phi_1 = 2.35^{\circ}$ and $\Phi = 3.94^{\circ}$.

| Scanned directions | 18 | Wavelength | $24\mu\mathrm{m}$ |
|-------------------------|---------------------|-----------------------|-------------------|
| Scan direction spacing | 10° | Averages | 1 |
| Noise value | $15 \mathrm{~dB}$ | Elasticity model | Hearmon |
| Signal bandwidth | $480~{\rm ms}^{-1}$ | | |
| Expected Error ϕ_1 | 2.35° | Expected Error Φ | 3.94° |

Table 6.2: Experimental parameters for SRAS scan of titanium sample.

It should be noted that the large scale EBSD map took 18 hours to capture with multiple IPFs having to be stitched together, in comparison to the 4 hours it took to acquire all eighteen of the SRAS velocity maps. It would be possible to further reduce this by using a single velocity map to identify grains and capturing only single pixel velocity surfaces within these.

Fig. 6.10 compares the *c*-axis maps captured by EBSD and SRAS on a Ti specimen. In all datasets, non-measured points have been shown in grey. Due to the isotropy at and near the basal plane, for grains which are close to basal ($\Phi \sim 0^{\circ}$, grains that appear white) angle ϕ_1 gave inconsistent results. Beside this special situation, the other measurements are comparable. It is worth noting that the spatial resolution of the RAS set-up in these experiments is 100 µm, which is approximately half of the generation patch diameter [253]. It is also possible to quantitatively compare the individual Euler angles, fig. 6.11a,b and fig. 6.11c,d compare ϕ_1 and Φ respectively, plotting the absolute angular difference on a pixel-by-pixel basis. In the comparative maps grain

boundaries have been ignored. Average errors are summarised in table 6.3.

The expected errors based on the settings in table 6.2 were smaller than those seen experimentally. Comparison of ϕ_1 between experimental measurements is difficult as this angle is taken in reference to the alignment of the sample. Slight misalignment of the specimen between experiments can cause errors of a few degrees. This is not the case for Φ , however the standard deviation of the acoustic wavelength in this experimental deployment is ~ 70 nm which corresponds to a velocity error of ~ 0.33%. This was an instrumental error due to the EMDA system where the divergent imaging beam passed through a tilted dichroic mirror; when the patch rotated to a different direction for velocity measurements, the projection is slightly distorted by the mirror and caused the wavelength error. From fig. 6.7b, this would be sufficient to cause a 1° - 2° error, and is a reasonable cause of the under-estimate of the error in Φ determination. Additionally, discrepancies between the simulated and experimental dataset can also be attributed to: variance between the elastic constants used to generate the forward model and the actual elasticity of the specimen, such as the variances shown in fig. 6.6 between the Ogi, Fisher and Hearmon elastic constant sets.



Figure 6.10: Comparison of SRAS and EBSD c-axis maps in α -titanium specimen.

Single pixel orientation at optimal experimental settings

The experimental settings of the previous section are commonly used as they offer sufficient velocity resolution without compromising spatial resolution or significantly increasing the scan time. However, when these factors are not of importance it is feasible to capture more detailed velocity surfaces. In this section the bandwidth of the signal and the number of scanned directions were



Figure 6.11: Comparison of orientation determination by SRAS and EBSD in α -titanium specimen for ϕ_1 (a,b) and Φ (c,d). Note the ϕ_1 scale wraps around (0° = 180°), whilst the Φ scale does not, hence why two different colourmaps have been used for these angles. In (b) and (d) two grains with a large error have been highlighted.

| Tab | le | 6.3: | Average | error | over | the | whole | e specimen: | SRAS | vs. | EBSD. |
|-----|----|------|---------|------------------------|------|-----|-------|-------------|------|-----|-------|
|-----|----|------|---------|------------------------|------|-----|-------|-------------|------|-----|-------|

| | Simulated Error | | | Experimental Error | | |
|--------------|-----------------|-------------------|---------------------------|--------------------|-------------------|---------------------------|
| Euler angles | Overall | $\Phi < 28^\circ$ | $\Phi \geqslant 28^\circ$ | Overall | $\Phi < 28^\circ$ | $\Phi \geqslant 28^\circ$ |
| ϕ_1 | 2.35° | 2.71° | 2.19° | 5.14° | 9.24° | 4.81° |
| Φ | 3.94° | 10.43° | 1.10° | 6.99° | 12.37° | 6.48° |

pushed to their practical maximum.

Experimental data is captured at two points in the specimen. Detailed scan parameters are given in table 6.4. Based on these parameters the predicted errors of this experimental study are expected to be $< 1^{\circ}$ on both grains - given the forward model has been calculated with 1 degree between planes this is the maximum orientation resolution. As expected, these errors were significantly smaller than those suggested in §6.2.5 thanks to the improved velocity resolution. Orientation determination is again compared to EBSD.

Scanned directions 360 Wavelength $24 \, \mu m$ 1° Scan direction spacing 128Averages $0\,\mathrm{ms}^{-1}$ SNR value 19 dBVelocity error Signal bandwidth $190 \, {\rm m s^{-1}}$ Elasticity model Fisher Expected Error $< 1^{\circ}$ Grain A $\Phi = 65^{\circ}$ $< 1^{\circ}$ Grain B $\Phi = 40^{\circ}$

Table 6.4: High velocity resolution experimental parameters for SRAS scan of titanium sample.

Fig. 6.12 shows the measured velocity surfaces from the two grains, with the fitted acoustic surfaces superimposed. In both grain A and B the resultant planes are within the imagined 1° of the EBSD measurements. Fig. 6.12c and fig. 6.12d show the FoM plots for grain A and B respectively.

One further important point to note is the influence of elastic constants. The low errors in orientation determination, compared to EBSD, are achieved using the Fisher model for α -titanium. When the values given by Hearmon are used, the Φ of the two velocity surface spectra are 58° and 42°, giving rise to errors of 7° and 1°, respectively. The values of both elastic constant sets are very similar but sufficient to lead to slightly different velocity surfaces (fig. 6.2). The discrepancy in the inversion result between these two elastic constant sets is quite surprising, particularly in the case of grain A. In both grains, the correlation value was higher using the Fisher model compared to the Hearmon model, suggesting the velocity surfaces generated by Fisher are a better reflection of the experimentally measured dataset. For a final comparison, the Ogi model, which produced notably fast velocity surfaces, led to errors in the Φ determination of 25° and 18° in grains A and B, respectively. The correlation values of this model were lower than both the Fisher and Hearmon sets., suggesting the inverse solver has the power to discriminate on both orientation and elastic constants simultaneously.





(a) Grain A velocity spectrum with velocities from fitted plane superimposed.





Figure 6.12: Inversion of two high resolution velocity surface spectrum for two grains in titanium sample.

Number of experimentally scanned directions and interval

Finally, the effect of number of scanned directions and the interval has been investigated experimentally (simulated results shown in §6.2.4). From fig. 6.13a, the experimental results show good agreement with the simulated data sets, within 2° for both ϕ_1 and Φ , following the same trend, with errors increasing with fewer scanned directions.

Similarly, both the simulated and experimental data show little influence in the determination of either ϕ_1 and Φ , except from when scans are close to orthogonal. This is a useful result as 2 direction velocity vector maps have previously been used to illustrate approximate c-axis position from orthogonal scans [253]. In future, such maps should use a scan interval between [30°, 60°].



Figure 6.13: Experimental results, compared with simulation, for varying number of scanned directions and scan interval when using directions.

6.3 Elastic constant determination

The proceeding section has demonstrated the excellent crystallographic orientation determination capability of SRAS. However, it is clear that reasonably accurate knowledge of the material elastic constants is required. Indeed, this problem was essentially encountered in the previous chapter when studying the high silicon steel and WAAM specimens, where uncertainty in the materials elastic constants is the limiting factor in generating velocity surfaces genuinely representative of the studied material. Again, this poses the question of how well these parameters need to be known to return credible results in orientation determination. This problem could be tackled with a methodology similar to that followed in the first half of this chapter, allowing uncertainty in elastic constants to be related to crystallographic orientation determination errors.

However, it is worthwhile to contemplate the problem at hand briefly. The elastic constants for many engineering materials are either unknown or have significant uncertainty due to the established measurement techniques, as described in Chapter 2, requiring exacting preparation of a specimen, usually of singe crystal facsimile. Of course, ambiguity in the elastic constants is detrimental to the quality of information obtained acoustically. Furthermore, it impacts material science, where elastic constants are essential for modelling plasticity and predicting ductility. Finally, framing this discussion within the aims of this thesis, the ability to vary the build material and process 'on-the-fly' in additive manufacturing mean the elastic constants are known to vary across the manufactured components, and the inherently polycrystalline nature further reduces the list of viable measurement techniques. Therefore, an actionable method of determining this in materials with arbitrary grain structures (and orientations) is interesting. With this in mind, the following section considers the exciting proposition of determining the elastic constants from SRAS measurements, presenting a method for simultaneously determining the elastic constants and crystallographic orientation.

6.3.1 Alogrithm

The general method outlined in this section can be explained by providing a worked example, outlined in fig. 6.14. (a) Staring with a polycrystalline specimen of unknown elastic constants and crystalline orientation, using SRAS the SAW is captured across the specimen in multiple propagation directions (b). The measurements from three of the grains are shown as radial velocity surfaces, note the different wave modes and velocities on each plane. Fig. 6.14c The possible velocity surfaces for each elastic constant permutation and crystalline orientation is then determined using a brute-force search of the forward model. The forward model looks for solutions to the Christoffel equation by scanning the determinant. Several wave modes, including the Rayleigh (RSAW) and

pseduo (PSAW) surface waves, can be seen to propagate on this plane - each plane is a single element of the forward model solution (See Chapter 4 for further information on the calculation of SAW velocities). Fig. 6.14d Each element of the forward model is then compared to the measured velocity surfaces by the inverse solver. Each grain input will then produce an independent figure of merit for the elasticity and orientation space. By assuming the elastic constants are a global property of the specimen, the elasticity figures of merit for each grain can be combined to give a final set of elastic constants for the full specimen. Finally, the orientation of each grain is then recalculated using the determined elastic constants. Fig. 6.14e shows the figure of merit for orientation space, showing the orientation goodness of fit for grains 1, 2 and 3.

This can be formalised by equation 6.2, where: Ng is the number of grains measured; F_{EO} is the elasticity-orientation figure of merit for each of these grains; $\overline{F_E}$ is the ensemble elasticity figure of merit for the whole specimen; (hkl) and ϕ_1 denotes the modelled plane and rotation; and C_{ij} is the modelled elastic constant matrix. The orientation with the greatest correlation value is selected for each element of C_{ij} . The elastic constants derived from $\overline{F_E}$ are then substituted back in to each F_{EO} to determine the correct orientation in each grain.

$$\overline{F_E}(C_{ijkl}) = \prod_{1}^{N_g} F_{EO}(C_{ij}, W, X, Y, Z)$$

$$(W(C_{ij}), X(C_{ij}), Y(C_{ij}), Z(C_{ij})) =$$

$$\underset{(h,k,l,\phi_1)}{\operatorname{arg\,max}} F_{EO}(C_{ij}, h, k, l, \phi_1).$$
(6.2)

6.3.2 Influence of elasticity on ultrasonic velocity

The first step in extending the inversion is a sensitivity analysis relating changes of elastic constants to the SAW velocity. Fig. 6.15a-c shows the change in SAW velocity on the {111} plane in pure Ni for the three elastic constants independently perturbed by ± 20 GPa. The first point to note is that the effect of C₁₁ is the inverse of C₁₂, this is suggestive of interdependence between these constants. Furthermore, there is an anisotropy in the velocity change, for example comparing C₄₄ at 45° and 70°. Over this range, the change in velocity is in the order of 600 ms⁻¹. The repetition spaced by 60 degrees is due to the six-fold symmetry on the plane {111}, planes near {001} have four-fold symmetry and intermediate planes have two-fold symmetry.

Fig. 6.15d-h shows the change in SAW velocity on the plane $\Phi = 45^{\circ}$ in the five independent elastic constants of titanium, again across a range of ± 20 GPa. Titanium is characterised by the presence of the PSAW, which dominates close to the basal plane before giving way to the RSAW as the cut-plane tilts away, $\Phi = 45^{\circ}$ has been chosen as it exhibits both wave modes. From d and



Figure 6.14: The new inversion method described in this work, detailing the process of calculating the crystallographic orientation and elastic constants for a synthetic specimen. (a) The SAW velocity is measured in multiple grains using SRAS. By rotating the generation patch velocity surfaces can then be built up for each grain. The velocity surfaces shown are from three arbitrary planes of nickel (here three primary planes are used as examples). (c) The forward model is searched for the range of possible elastic constants (C_{11}, C_{12}, C_{44}) and crystalline orientation (hkl). Each element of the forward model determines the velocity of all SAWs which can propagate across that plane. (d) The output of the forward model search and one experimentally measured velocity surface are input to the inverse solver. The inverse solver calculates the 'goodness of fit' between the velocity surface and each element of the forward model. This is then repeated for each measured grain before taking the product of the figure of merit for all grains. The specimens elastic constants are then determined by the location of the maxima. (e) Calculated crystallographic orientation from the inverse solver, showing the three grains on one pole figure. (f) Final calculated elastic constants for this specimen.

e it is again seen that C_{11} is the inverse of C_{12} , in the RSAW regions. However, in the PSAW region varying C_{12} appears to have no impact on the SAW velocity.

This illustrates governing properties for the RSAW and PSAW are different and thus have different interactions with elastic constants. The basal plane is isotropic in hexagonal materials, this means only Euler-Bunge angles ϕ_1 and Φ can be determined by linear acoustic techniques [307]), an important limitation of this work for orientation determination.

The results for C_{13} and C_{33} are given in Fig. 6.15g-h, respectively. For the RSAW mode both constants exhibit minimal effect on the SAW velocity. In the PSAW region the change becomes more distinct (~ 150 ms⁻¹ over ±20 GPa). It is also seen that the value of C_{13} and C_{33} changes the position of the mode transition. Given this effectively causes changes in SAW velocity of ~ 500 ms⁻¹ accurately measuring this position is likely to greatly improve the accuracy to which these elastic constants can be determined.

Given the interaction shown in fig. 6.15a,b and d,e it is useful to consider the value of $C_{11} - C_{12}$, henceforth referred to as C_{Δ} to prevent confusion with the plane $C_{11} - C_{12}$, C_{11}/C_{12} is used to refer to both constants in cases where their behaviour is similar.

Fig. 6.16a shows the elastic constant which causes the greatest change in SAW velocity across the orientation space in Ni. Velocity surfaces are symmetric about the indicated line and thus C_{Δ} can be excluded in the second half to provide information on the succeeding sensitivity. The result is a simple conclusion where the order of sensitivity C_{Δ} , C_{44} , C_{11}/C_{12} is constant across all planes. The change in SAW velocity however, remains anisotropic. In short, measuring a plane near {0.5 0.5 1} makes constant determination an easier task than near {001}, as the change in SAW velocity - as a function elastic constants - is greater.

Similarly, fig. 6.16b shows the elastic constant sensitivity in Ti, across the orientation space. C_{Δ} and C_{44} are the dominant elastic constants across the orientation space, as seen in the left hand-side of the plot. SAWs in hexagonal crystal structures are symmetric about $\phi_1 = 90^\circ$; C_{44} and C_{Δ} have been excluded for $\phi_1 > 90^\circ$. The response in this region is harder to interpret, with a chaotic behaviour flipping between dominant constants. In general, directions close to $\phi_1 = 90^\circ$ are sensitive to C_{11}/C_{12} , then give way to C_{13} as $\phi_1 \to 0^\circ$.

In summary, these results demonstrate the underlying anisotropy of the elastic constant influence in SAW velocity. The variation in the SAW velocity, within a useful range of elastic constants, are well within the bounds of typical SAW measurement accuracy at certain angles (> 100ms^{-1}), but necessitate challenging levels of accuracy at other angles (< 1ms^{-1}). By sampling various crystalline orientations the sensitivity to each elastic constant can be maximised, as elastic anisotropy means some crystalline planes yield little sensitivity to certain elastic constants. Therefore, the inversion algorithm proposed in this work is most effective when using velocity surfaces from multiple



Figure 6.15: (a)-(c) present results in cubic Ni. The change in SAW velocity across the {111} plane for perturbations in the elastic constants (a) C_{11} , (b) C_{12} , (c) C_{44} over a range of 40 GPa. It is important to note the dependence between C_{11} and C_{12} , simply perturbing C_{11} and C_{12} in equal and opposite directions (i.e. $C_{11} \uparrow, C_{12} \downarrow$ by 10 GPa) results in a near identical SAW velocity. (d) - (h) present results for hexagonal Ti on the plane $\phi = 45^{\circ}$, for perturbations in each of the five independent elastic constants over a range of 40 GPa. $\phi = 45^{\circ}$ is selected as it has both RSAW and PSAW components. (d) C_{11} , (e) C_{12} , (f) C_{44} , (g) C_{13} , (h) C_{33} .



Figure 6.16: (a)The elastic constant that causes the greatest change across the orientation space in Ni. C_{Δ} dominates across all planes. The change in SAW velocity at the four markers positions are 1) 70 ms⁻¹, 2) 140 ms⁻¹, 3) 55 ms⁻¹ and 4) 22 ms⁻¹. The change from perturbations C_{11}/C_{12} is not shown but varies between 6 ms⁻¹ near {001} and 2 ms⁻¹ near {111}. (b) The elastic constant that causes the greatest change across the orientation space in Ti. C_{44} and C_{Δ} have been excluded for $\phi_1 > 90^{\circ}$ (line of symmetry). The change in SAW velocity at the four markers positions are 1) 89 ms^{-1} , 2) 76 ms⁻¹, 3) 15 ms⁻¹ and 4) 1 ms⁻¹.

grain orientations to ensure some sensitivity to all elastic constants is included.

6.3.3 Experimental results

Nickel and nickel superalloy

The following section presents the results of elastic constant and crystallographic orientation determination in three real-world specimens. Firstly, the inversion of a velocity surface (fig. 6.17a) of a Ni single crystal specimen on the (001) plane is presented, the R-value between the determined orientation and the SRAS inversion is 2° , due to slight misalignment of the rotation - collecting velocity measurements at increments smaller than every 5° would likely improve this). Elastic constant values for pure Ni are well established in published literature and thus provides a useful basis for validating the inversion method. The determined crystallographic plane is {001}, and elastic constants calculated by inversion are in good agreement with the prior literature, fig. 6.17b - results from this specimen are indicated by \checkmark markers.

This has demonstrated the technique with a single high quality velocity surface in a specimen with well characterised elastic properties. A more complex Ni alloy is now considered, using a greater number of grains but lower quality velocity surfaces. Two CMSX-4 polycrystalline specimens were fabricated (with one annealed and one left in the 'as-built' state) and then scanned by SRAS, with the SAW velocity measured in 18 directions at an interval of 10°. The orientation result of the 'as-built' specimen has previously been reported [308], and showed good agreement to EBSD, using literature elastic constants. The annealed specimen (studied below) is of particular of interest to the present work as the original orientation result showed poor agreement to EBSD, suggesting an error in elastic constants.

The seven grains indicated in fig. 6.18a were selected for the inversion process as these had unique orientations and a large number of pixels with good SNR. The waveforms from each pixel in these grains were then averaged (each grain had at least 1000 pixels once boundary regions and locations with poor SNR ratio were excluded). From fig. 6.17b (results from this specimen are indicated by \star markers), the calculated elastic constants are within the reported range for CMSX-4; C₁₁ and C₁₂ are particularly in good agreement with the measurements of Amulele et al. for CMSX-4 [309].

However, the calculated value of C_{44} is ~ 8 GPa lower than reported values. Mendik et al. found that alongside the impact of surface condition and residual stress, annealing a sample will change the measured SAW velocity and lower the values of C_{44} measured [310]. Annealing of the specimen is the only difference between this specimen and that previously reported with good orientation agreement using literature elastic constants, and therefore the probable source of discrepancy in $C_{44}.$

Example results of the inversion process for grain 5 is shown in fig. 6.17c, directions with high SNR (such as 90°) had a smaller velocity standard deviation, this was due to the anisotropy in the amplitude of the SAW [311]. The mean difference between the measured data and inversion is 19 ms^{-1} , compared to 78 ms^{-1} for the literature elastic constants in this grain.

Bearing in mind the variation in SAW velocity due to elastic constant perturbation is anisotropic across a plane (see fig. 6.15), it is useful to compare the measured propagation directions to the sensitive directions, as when measuring only 18 directions it may be possible to miss propogation directions that cause large changes in the SAW velocity, for a given elastic constant. Fig. 6.17d plots the angle measured velocity surfaces against the change in velocity for a perturbation in C₄₄ of 10 GPa for seven planes, corresponding to those used in the inversion. Aside from grain 3, the velocity surface was measured in a region with a change in velocity greater than $50 \,\mathrm{ms}^{-1}$. This suggests the inversion should be sensitive to C₄₄ and the calculated value of 119 GPa is accurate.

Fig. 6.18 shows the orientation result of the SRAS inversion compared to the EBSD dataset and the orientation calculated from literature values for CMSX-4 elastic constants. In general, the EBSD data is in much better agreement with the orientation calculated in this work. Considering the EBSD data has not been used as an input to the inversion and is generally considered the benchmark for crystallographic orientation calculation this helps confirm the validity of the new elastic constant set for this specimen. For Φ , the rotation between (001) and (111), the disagreement between inversion elastic constants and the orientation measured by EBSD was 1.3°, compared to 5.9° for the literature values. For ϕ_2 the discrepancies were 6.9° and 13.1°, respectively. Given the values of C_{Δ} , C_{11} and C_{12} match those found in literature, the updated value of C_{44} has clearly been crucial in improving the orientation determination.

α -Titanium

Determination of the elastic constants in hexagonal materials is more complex as the search space becomes significantly larger due to five independent elastic constants. Furthermore, from fig. 6.15gh and fig. 6.20d the SAW velocity has a weak sensitivity to the two additional constants C_{13} and C_{33} across much of the orientation space. To demonstrate the proposed method, a polycrystalline titanium specimen was scanned using SRAS, fig. 6.19c shows the single direction SRAS velocity map and indicates the two arbitrary grains targeted for the inversion. The inversion result and calculated orientation are shown in fig. 6.19a,b for grains 1 and 2, respectively. The crystalline plane, Φ , of the two grains was calculated as 66° and 39°, respectively, within 1° of the mean EBSD result from the grain in both cases. The determined elastic constants show good agreement with existing literature values, fig. 6.19d. In all cases, the elastic constants fall within the range found in prior literature, and agree particularly well with oft-quoted values determined by Hearmon [226].

In the case of C_{33} there is a discrepancy of ~ 6%. In Grain 1 (fig. 6.19a), the sensitivity in C_{33} comes from the RSAW velocity around $\phi_1 = 0/180$, which in the measured data is relatively lower amplitude than at $\phi_1 \approx 45^{\circ}/135^{\circ}$. In Grain 2 (fig. 6.19b), the position of the mode transition provides the sensitivity to C_{33} , however data around this position this was omitted from the inversion, due to the position of the mode-transition being in correctly predicted by the forward model, as discussed in §4.2.5.



Figure 6.17: (a) Measured data for single crystal Ni specimen, with fitted data superimposed. The measured data-point $\sim \phi_1 = 25^{\circ}$ has been excluded from inversion as this occurs at the transition between wave modes as shown. (b) The elastic constants of Ni and Ni-alloys. Both experimental datasets show good agreement to existing values. Red bars indicate the expected standard deviation in elastic constants determined in this work. The SAW inversion is particularity sensitive to the value of C_{Δ} . (c) Velocity surface from grain 5 compared to the velocity surfaces (red-line) from SRAS inversion and literature value of elastic constants. The colourscale of the measured data indicates the relative amplitude of the SAW in the given direction. Error bars indicate the standard deviation of the measurement across the pixels used for the inversion. (d) Plots the velocity variation as a function of 10 GPa perturbation in C_{44} , for the 7 grains used as inputs for the inversion and the position of the 18 ultrasonic measurements on each grain. In all grains except 3, several measurements have been taken at angles with high sensitivity.



Figure 6.18: Compares the inverse pole figures for the CMSX-4 specimen, from (a) EBSD, (b) SRAS using inversion elastic constants and (c) SRAS using existing elastic from literature. Numeric labels in the EBSD data indicate the grains used in the inversion process. The elastic constants calculated by the SRAS inversion in this work have significantly improved the orientation result, relative to the EBSD dataset.



Figure 6.19: (a) Measured velocity surface in Grain 1, with fitted result from inversion superimposed, calculated to be $\Phi = 66^{\circ}$. (b) Measured velocity surface in Grain 1, with fitted result from inversion superimposed, calculated to be $\Phi = 39^{\circ}$. The orientation was determined within 1° of EBSD in both cases, areas shown in white indicate where data was omitted from inversion, due to the uncertainty of the mode transition. (c) Single direction SRAS velocity map of whole specimen surface, G1 and G2 indicate the two grains used for inverse analysis. (d) Comparison of elastic constants determined by SRAS and literature values, for Ti. Red bars indicate the expected standard deviation on each elastic constant from SRAS inversion. Black bars indicate the range of values found in prior literature. All constants fall within the range found in literature and all except C₃₃ show good agreement with the values of Hearmon [226].

6.3.4 Analysis of error and uncertainty

Shape of inversion solution shape

Having demonstrated experimental results, the accuracy and sources of uncertainly in the inversion shall now be investigated. The first step is to investigate the solution space, understanding this is an invaluable tool for tailoring the search, informing on the accuracy of elastic constant determination and dictating the search range. In general, it is beneficial for the solution space to be concave with a single global minima with large variations in F_E from perturbation in a member of C_{ijkl} . Fig. 6.20(a) shows this to be true in the plane $C_{11} - C_{44}$, the position of the minima in this plane could be well determined by any simple optimisation algorithm. In contrast, the minima is poorly defined in the plane $C_{11} - C_{12}$; values along the same diagonal trajectory (where C_{Δ} remains constant) produce similar values of F_E . This is the impact of the interdependence between C_{11} and C_{12} , as shown in fig. 6.15.

As shown by fig. 6.20(a), the orthogonal $C_{11} + C_{12}$ is therefore the direction with poor sensitivity. This phenomena can be visualised by looking at two typical velocity surfaces from elastic constant perturbation; fig. 6.20(c) where SAWs on the plane (111) for a perturbation in the plane $C_{11} - C_{44}$ (blue line) and the plane $C_{11} - C_{12}$ (red line), of equal perturbation from a nominal value (black line). Clearly, the perturbation in the C_{44} direction is much larger, 102 ms^{-1} on average, than the $C_{11} - C_{12}$ direction where the change is only 2 ms^{-1} on average.

The range over which to search can be informed by plotting the elastic constants of several cubic materials on a scale normalised by the value of C_{44} . Most materials have several similar literature vales however, it appears most are tightly spread in the C_{Δ} direction but can be extremely spread out in the $C_{11} + C_{12}$ direction. β – Ti (red diamonds) and Al (blue circles) are notable examples. Whilst not conclusive, this would be typical of measurements which can determine C_{Δ} but lack the accuracy to determine the specific values of C_{11} and C_{12} along the minima. Applying the familiar Born stability criterion $C_{11} > C_{12}$ significantly reduces the search space (E < 0 region in black). Given this, and the importance of searching close to the plane $C_{11} + C_{12}$, searching in a square grid will generate many solutions far way from the minima and in some cases solutions that are unstable (inset square grid). Instead, a grid tiled around values of C_{Δ} allows greater refinement of search (inset tilted grid). The Zener anisotropy ratios, α , 0.5, 1 and 2 are indicated by dashed white lines. It is worth remarking on the difference between anisotropy where $\alpha > 1$ and $\alpha < 1$. Aside the from the absolute difference in velocity, the shapes of the velocity surface in KCL ($\alpha > 1$) and Ni ($\alpha < 1$) are noticeably different, as the directions of high velocity are 45° out of phase. Care should be taken not to perform the inversion of measured velocity surfaces with modelled surfaces from the other side of $\alpha = 1$ as the incorrect rotation on the plane may be determined.

Further information of this style of plot can be found in [312].

In titanium, the solution space of the planes $C_{11} - C_{44}$ and $C_{11} - C_{12}$ are similar to (a) and (b), such that they are not worth reproducing here. Instead, two particularly interesting planes are shown. Firstly, fig. 6.20(e) plots the value of F_E for $\Phi = 90^\circ$ in the plane $C_{13} - C_{33}$, note the scale has changed from fig. 6.20(a)-(b). Whilst we see a minima does exist, the function is weakly concave. Finally, fig. 6.20(g) plots the value of F_E for $\Phi = 0^\circ$ in the plane $C_{11} - C_{33}$, this region has a greater sensitivity to C_{33} (fig. 6.15(h)) but has an interdependence on the value of C_{11} , making the minima poorly defined. To the authors knowledge, this is the first reporting of the solution space in hexagonal crystal systems.

Relating instrument performance to errors in elastic constant determination

The ability to accurately determine the elastic constants of the specimen is primarily dependent on the instrument accurately measuring the SAW velocity. By simulating the velocity surfaces that would be determined by a range of δv the corresponding errors in elastic constant determination were elucidated. Fig. 6.21 shows the resulting standard deviation in elastic constant determination as a function of the velocity error standard deviation for (a) nickel and (b) titanium. The experimental availability of these velocity errors is contextualised by fig 6.21(c), showing the expected velocity standard deviation as a function of number of signal averages and fringes in the generation patch. In both materials we see C₄₄ and C_{Δ} is determined most accurately, whilst the interdependency between C₁₁ and C₁₂ causes significantly more uncertainty. As hypothesised from fig. 6.15(f), C₁₃ is slightly easier to determine than C₃₃. Reducing the standard deviation in the determination of all elastic constants to below 1 GPa is practically achievable if >100 averages can be used, allowing the elastic constants to be determined with meaningful accuracy.

The standard deviation of elastic constants given in fig. 6.21 should be considered a lower bound of the errors achievable without further modification to the method. Firstly, these errors are dependent on capturing the velocity surface at 180 uniformly incremented propagation directions; reducing the number of captured propagation directions will reduce accuracy. Elastic constant standard deviations expected when using 18 measurement angles, as in the present work, are calculated in this fashion and included with experimental results above.

Furthermore, given the acoustic generation mechanism relies on thermoelastic expansion, some temperature change from the pulsed laser would be expected. Simple finite-element modelling estimates the resultant temperature rise in Ti is $< 1 \,^{\circ}$ C. Experimental data presented in this work is spatially over-sampled leading to DC heating, this is estimated to be $2 - 3 \,^{\circ}$ C. To give these values context, an increase in specimen temperature of $2 \,^{\circ}$ C reduces the three elastic constants [313] in turn slowing the SAW velocity by $\sim 10 \,\mathrm{ms}^{-1}$ in CMSX-4 on average [314].

Finally, in most crystalline materials, certainly in the nickel and titanium specimens studied in this work, several surface wave modes can propagate. In the single crystal Ni and Ti results presented, velocity measurements near the mode transition have been excluded as the current method of discriminating between the modes is somewhat imprecise and leads to uncertainty in the location of these transitions on the plane (see §4.2.5). In future, coupling the forward model with a finite-element model to calculate the wave mode displacements will further improve the inversion accuracy - further consideration of this idea is given as part of a discussion future work in §8.5.

6.3.5 Implementation

Experimental considerations

The current implementation of SRAS allows rapid high resolution velocity maps to be acquired by prioritising raster scanning. This has limited the present work to either several grains with few angular measurements or detailed angular measurements in few grains. Making adjustments to the instrumentation to prioritise angular scanning will allow high velocity resolution velocity surfaces to be captured in many grains, with the expectation this would further improve the accuracy in elastic constant determination. Experimental data in this work has been captured at 18 propagation directions with single-shot measurements, optimal parameters for minimising scan time but increases the errors in elastic constant and orientation determination. By using temporal averaging and scanning a greater number of directions, the accuracy of determination can further be improved. The expected temperature variation from the generation process is small ($< 3^{\circ}$) but can be further reduced by enacting rigorous temperature control of the specimen during the SRAS experiment.

The standard deviation of elastic constants given in fig. 6.21 should be considered a lowerbound of the errors achievable without further modification to the method. Firstly, these errors are dependent on capturing the velocity surface at 180 uniformly incremented propagation directions, reducing this will reduce accuracy. Elastic constant standard deviations expected when using 18 measurement angles, as in the present work, were calculated in this fashion and included with experimental results above. Furthermore, given that the acoustic generation mechanism relies on thermoelastic expansion, we would expect to see some temperature change from the pulsed laser. Simple finite-element modelling estimates the resultant temperature rise in Ti is < 1 °C. Experimental data presented in this work are spatially over-sampled - a given location on the specimen is exposed to multiple laser pulses, therefore the build up in temperature from successive pulses of the generation laser must be considered, this is estimated to be 2 - 3 °C, again by finiteelement modelling. To give these values context, an increase in specimen temperature of 2° C reduces the three cubic elastic constants [313] in turn slowing the SAW velocity by ~ 10ms^{-1} in CMSX-4 on average [314]. Enacting rigorous temperature control of the specimen during the SRAS experiment will be required to realise improved accuracy.

The grain size in all of specimens studied in this section are relatively large (mm-scale). For this proposed method to operate grains must be spatially resolved by the experimental system. Presently, the spatial resolution of the experimental system is $\sim 50 \,\mu\text{m}$. This covers a wide range of relevant materials and processing routes. However, it is envisaged that the spatial resolution can be further improved by a factor of ~ 10 through minor changes to the optical imaging system, making the proposed method applicable to a wider range of materials

Numerical considerations

Combined improvements in hardware and algorithm efficiency have allowed a brute-force approach, as used in this work, to become a viable method for the first time. The calculation of a single velocity plane now takes ~ 5s at a velocity resolution of 2 ms^{-1} , quoted speeds are for an Intel[®] i9-10900X @ 3.70 GHz. The generation of forward models and the inversion are both suited to parallel computation, allowing modern multi-core hardware to be leveraged. As a demonstration of the computational demand, table 6.5 tabulates the search parameters and computation time for the inversion of the CMSX-4 specimen studied in this work (using the pc described above). The brute-force method has proved attractive given the local minimum and weakly concave shape in the C₁₁ - C₁₂ plane, mitigating previous issues with optimisation solvers [239]. The generation of the forward models is assumed to be completed de-novo, however after completing the inversion process for a material the forward models are saved. For example, having studied nickel and the nickel alloy CMSX-4 in this chapter, the search area spanning most nickel-like materials is now pre-stored, further reducing the inversion time for future nickel specimens.

Where possible, points close to the mode transition were omitted so as to prevent this artefact from dominating the inversions reported in this work. In most crystalline materials, certainly in the nickel and titanium specimens studied in this work, several surface wave modes can propagate. In the single crystal Ni and Ti results presented, velocity measurements near the mode transition have been excluded as the current method of discriminating between the modes is somewhat imprecise and leads to uncertainty in the location of these transitions on the plane. In future, coupling the forward model with a finite-element model to calculate the wave mode displacements will further improve the inversion accuracy.

| Parameter | Resolution | | |
|---------------------|---|--|--|
| Coarse search | | | |
| (hkl) | 0.1 to (011) to (111) | | |
| C_{44} | $1\mathrm{GPa}$ over $30\mathrm{GPa}$ | | |
| C_Δ | $0.5\mathrm{GPa}$ over $10\mathrm{GPa}$ | | |
| Computation time | 6 hours | | |
| fine search | | | |
| (hkl) | 0.05 to (011) to (111) | | |
| C_{44} | $0.5\mathrm{GPa}$ over $8\mathrm{GPa}$ | | |
| C_{11} | $0.5\mathrm{GPa}$ over $\mathrm{C}_{\Delta}\pm 4\mathrm{GPa}$ | | |
| C_{12} | $0.5{\rm GPa}$ over $C_{\Delta}\pm 4{\rm GPa}$ | | |
| Computation time | 16 hours | | |

Table 6.5: Forward model search resolution for inversion of SAW velocities in CMSX-4 specimen.

6.4 Summary

This chapter has investigated the concept of determining the crystallographic orientation and elastic constants from SRAS velocity surface spectra. The chapter first focuses on the accuracy of crystallographic orientation determination before considering the simultaneous determination of orientation and elastic constants. This layout followed a logical narrative as the importance of elastic constants was emphasised in the earlier work of this chapter, and the simultaneous determination of both orientation and elastic constants is of greater complexity.

The main thrust of this first section has been to investigate the sensitivity of orientation accuracy to experimental parameters when measuring SAW velocity in hexagonal materials. The results presented have focussed on titanium, given its industrial relevance and the results presented in Chapter 4, but this method can be extended to any hexagonal material. For example, Zirconium is another common hexagonal material in nuclear energy and would be an obvious candidate to be similarly studied.

The calculated SAW velocity model of titanium has shown that its velocity surface is isotropic on the basal plane ($\Phi = 0^{\circ}$), this prevents the determination of ϕ_2 . As the cut plane tilts away from the basal plane, the velocity surfaces become elliptical, before the occurrence of a wave mode hop around 28°, where the dominant wave mode transitions from a supersonic wave to the normal Rayleigh wave. Accuracy in the determination of ϕ_1 was shown to vary across the orientation space, with errors of $< 2^{\circ}$ degrees easily achievable in the Rayleigh wave range, for almost any level of SNR. Similar precision can be achieved in the supersonic zone (apart from the special case on the basal plane) but relies on higher SNR (10 dB) or increased velocity resolution. The orientation determination accuracy was shown to be primarily dependent on the accuracy of the SAW velocity measurement. A FWHM of $> 400 \,\mathrm{ms^{-1}}$ is required to achieve average errors of 1° across the orientation space. Determination close to the basal plane becomes increasing difficult as velocity resolution must approach sub $- \,\mathrm{ms^{-1}}$ to differentiate between planes (e.g. $\Phi = 0^{\circ}$ and $\Phi = 1^{\circ}$). Conversely, determination in the region $40^{\circ} \leq \Phi \leq 60^{\circ}$ requires less precise velocity measurement as there is a greater change in velocity between planes in this region.

Full specimen comparison with EBSD yielded an error of 6.48° in Φ determination. High SNR velocity surfaces, representing the best possible performance of the current instrument, reduced the error in orientation determination to < 1°, with the caveat that acquiring such high-quality data remains a significant challenge in and of itself. In both cases, it was notable that the choice of published elastic constants used to generate the forward model could have a significant impact on the accuracy of orientation determination, particularly given the elastic constants used were both stated to be for α -titanium.

It was found that for both ϕ_1 and Φ capturing > 3 propagation directions significantly reduces the error in orientation determination and is relatively insensitive to varying SNR levels. Increasing the number of directions to 90 reduces the associated error by a further ~ 0.5°.

To conclude this section, two experimental studies have been compared to the simulation. Firstly a polycrystalline α -phase hcp titanium specimen has been scanned using SRAS using standard parameters (FWHM > 480 ms⁻¹ over 18 propagation directions) and the orientation output compared to measurements by EBSD. The calculated orientation was found to compare well with EBSD with an average error of 4.81° and 6.48°, for ϕ_1 and Φ in the Rayleigh wave region, errors approximately double in the supersonic region ($\Phi < 28^{\circ}$). Pertinently, the experimental errors are in good agreement with the expected errors from simulation (overall errors are 2.80° for ϕ_1 and 3.05° for Φ). Discrepancies are suggested to be from variation in the acoustic wavelength, inaccurate elastic constants and variation in the FWHM across the specimen. A second measurement with higher velocity resolution and increased scanning angles (FWHM of > 190 ms⁻¹ over 360 propagation directions) was made in two grains. With these scanning parameters, orientation errors were expected to be reduced to less than 1°, which was found when comparing the calculated orientation to EBSD. The second experimental setting represents the highest resolution that is currently practically achievable.

Whilst the sensitivity to the elastic constants could be considered a limitation of the technique, it has provided an opportunity to determine the elastic constants in polycrystalline specimens with unknown crystallographic orientation. The measurement of elastic constants has traditionally been an onerous pursuit, requiring exacting preparation and knowledge of the crystallographic orientation and dimensions of the specimen. A demonstration of how SRAS may be used to simultaneously determine the elastic constants and crystallographic orientation has been provided. Experimental measurements are made in pure Ni (cubic), Ti (hexagonal) and the Ni-alloy CMSX-4. In all cases, the determined crystallographic orientation showed good agreement with EBSD ($< 6^{\circ}$ error for 18 propagation direction scans and $< 1^{\circ}$ for 180 propagation direction scans), including the previously problematic CMSX-4 specimen. Therefore, aside from determining the elastic constants, the complete inversion algorithm has the benefit of making the established orientation inversion process more robust by removing the dependency on literature values for elastic constants. The determined elastic constants are compared with existing literature values are found to agree within 1

Solutions spaces in both nickel and titanium have been visualised, illustrating the convace nature of C_{44} and C_{Δ} , contrasting the C_{11} and C_{12} . This has been explored further to demonstrate that the accuracy to which the constants may determined was vary for each constant. For example, the elastic constants in CMSX-4 are estimated to have associated standard deviations of 2 GPa for C_{44} and 10 GPa for C_{11} and C_{12} (although the value of C_{Δ} is found with a standard deviation of 1 GPa). Similarly, C_{44} and C_{Δ} are easier to determine in hexagonal materials, with C_{11} , C_{12} , C_{13} and C_{33} all showing an interdependence, in agreement with prior art.

A series of simulations have allowed errors in the elastic constant determination to be correlated to the precision of SAW velocity measurement. Whilst not fully comprehensive, notably not accounting for variations in temperature or systematic errors, this suggests current accuracy. Furthermore, this study has indicated that sub-GPa accuracy is possible by minimising errors in the velocity measurement, $< 1 \,\mathrm{ms}^{-1}$. Potential sources of error and future improvements to the methodology were briefly discussed to bring the chapter to a close. Experimental considerations focussed on careful temperature control of the specimen. Discussed improvements to the algorithm focused on developing a higher fidelity forward model that captures the thermoelastic generation process.



Figure 6.20: Examples of elastic constants solution spaces for nickel (a)-(b). (a) The plane $C_{11} - C_{44}$. The minima is well defined at $(C_{44} = 130 \text{ GPa}, C_{11} = 250 \text{ GPa})$. (b) The plane $C_{11} - C_{12}$. The minima is poorly defined at $(C_{12} = 150 \text{ GPa}, C_{11} = 250 \text{ GPa})$. (c) Compares SAWs on the plane (111) for a perturbation in the plane $C_{11} - C_{44}$ (blue line) and the plane $C_{11} - C_{12}$ (red line), of equal perturbation (10 GPa) from the minima (black line). The relative position of the three elastic constants sets are shown in the inset scatter plot. Perturbations in C_{44} are easier to measure as they causes larger changes in the SAW velocity. (d) Elastic constants of several cubic materials on a normalised scale. The background colorscale is the normalised Young's modulus (E/C_{44}) in the [001] direction and α is the Zener ratio [312]. Several materials appear spread along the diagonal of anisotropy, this is the direction of least sensitivity. Thus, when determining elastic constants a wide range of values along this diagonal should be considered. Examples of the solution space in hexagonal titanium, note the scales have been reduced compared to Ni (d) - (e). (d) The plane $C_{13} - C_{33}$ in the RSAW region. (e) The plane $C_{11} - C_{33}$ the PSAW region, this region has a greater sensitivity to C_{33} , but has a co-dependency on the value of C_{11} .



Figure 6.21: Simulations of errors in the determination of elastic constants, as a function of velocity standard deviation for pure (a) Ni and (b) pure Ti. (c) provides context to (a) and (b) by illustrating the typically available velocity standard deviation in the SRAS experiment, as a function of signal averages. From (c) it can be seen that a velocity standard deviation of $< 10 \text{ ms}^{-1}$ is readily achievable even in the 8 fringe case (more broadband signal) with few averages (~ 10), this corresponds to a standard deviation of 10 GPa for C₁₁ and C₁₂ (least sensitive constants) in Ni (approximately 4% and 6.65%, respectively).

Chapter 7

The case for SRAS in additive manufacturing

7.1 Introduction

In the previous two chapters, experimental results on additive manufacturing specimens and the development of elastic constant determination have been reported. It is now worthwhile to consider the impact of a spatially resolved acoustic spectroscopy (SRAS) instrument in a production environment, focussing primarily on the monetary and temporal costs. The motivation for developing this model is that Renishaw wishes to understand the case for the integration of a SRAS system within a commercial additive manufacturing build chamber. As a build system manufacturer Renishaw are in a position to add inspection modules to their laser powder bed fusion (L-PBF) machines, and have already begun integrating passive monitoring systems. As demonstrated in Chapters 4 and 5, SRAS potentially offers significant benefits within the manufacturing cycle. This work has focussed on post-build inspection; at the outside of this EngD, the proposition of integrating SRAS within a build system was also considered, but the economic case must also be credible.

This chapter constructs a cost-time model for powder-bed additive processes ($\S7.5$) based on existing literature (summarised in \$7.2) to assess the economic impact of in-situ inspection. The model firstly calculates the time to build a given component, along with the additional time to complete a SRAS scan. Renishaw has provided commercial context and detail, such as the cost of additive equipment for this chapter. The characteristics of the EMDA system and alternative rough-surface-capable systems are used to define a window of hypothetical scan speed (\$7.5.2). This scan speed and the SRAS system costs are then used as an input to the cost-time model. The model first considers the temporal impact of in-situ SRAS scanning, describing the increase in total build time and reducing total output in manufacturing a typical aerospace AM component.

The component's embodied cost is then considered a function of the build failure rate in §7.6.1, leading to a formulation of the problem as a function of failure rate and inspection performance. Results are compared to a no-inspection and x-ray computed tomography (XCT) inspection (as is industry standard) scenarios. One potentially powerful driver of in-situ inspection is the ability to repair defective builds. The cost model is extended in §7.6.3 to assess the impact of in-situ repair on the additive process. Unfortunately, despite being clearly desirable, the current generation of additive build systems do not offer the capability to repair defective components as standard. As such, §7.6.4 describes a preliminary investigation of the capacity to repair defects observed in additive components. Observed surface-breaking defects are identified ex-situ and returned to the build system to trial repair strategies.

The chapter concludes with a discussion around the possible implementation and technical considerations of a SRAS system within an additive build system (§7.7). This section includes the demonstration of a rough surface capable SRAS system within the carcass of a Renishaw AM250 build system. The final discussions desire to challenge the underlying assumptions used in the model and discuss possible avenues to reduce the temporal and economic impact of in-situ inspection.

7.2 Existing cost-time models

7.2.1 Early work

Additive manufacturing is a complex process with many user-controlled build variables. This results in a build time (and quality) highly dependent on the parameters used. Significant work has been undertaken to build models that can evaluate the cost and time implications of the laser powder bed fusion process. The build process can be broken down into three key components: scanning, recoating and setting. The scanning phase is the time spent irradiating the powder bed. The recoating stage refers to the time spent lowering the bed after a deposited layer and depositing a fresh layer of powder. Finally, setting is the time spent by the operator both setting up the build and fulfilling tasks to complete the build, such that a subsequent job can begin.

General cost models for subtractive manufacturing are well established, having been available for several decades [315]. The remainder of this section is dedicated to a brief overview of pertinent cost-time models in additive manufacturing. Works discussed here provide a basis for the development of the cost-time model for a hypothetical SRAS system, which is presented in §7.5. The development of such models has progressed in lockstep with the technology, providing a tool to study the process, for example, the contribution of activities within the process to the total cost or the sensitivity to part geometry. This field was initiated by the seminal work of Alexander et al. [316], where the authors proposed an activity-based costing approach that categorises costs into direct and indirect costs. In cost models of this type, indirect costs are generally allocated based on the time requirement of various processes occurring in the AM workflow. This approach has become standard in most subsequent modelling.

Later the focus of cost-time models moved from the study of single builds to the scalability of multiple builds. In essence, most of these models were built to address long-standing questions such as the economies of scale within additive manufacturing. A significant body of work is associated with demonstrating that economies of scale are available in AM, with all authors agreeing that the cost of additive systems is primarily dependent on maximising utilisation of the build system, achieved by filling the build volume and running builds back-to-back with minimal downtime. Baumers et al. built the first cost-time model for additive manufacturing which accounted for the indirect costs of the process and has been used to provide insight into the unsuitability of additive manufacturing for high volume applications [317]. This work is extremely comprehensive as it attempts to account for every ancillary cost associated with the AM process, for example, estimating the power consumption of the build machine.

7.2.2 Modelling failure and defects

More recently, there has been a desire to include effects such as failure and inspection within these models. Ding et al. noted that existing models did an excellent job of capturing the embodied cost for an AM component due to the equipment cost and associated activities but ignored the potentially significant impact of build failure [318]. This additional costing is based on the idea that a failed build must be replaced by a conforming component, increasing the cost of production. Experimental findings concluded the cost impact of build failure rate can be significant in AM and produces a complex relationship between average unit cost and quantity at the level of the build. The authors also importantly find that the longer a build runs for then, the greater the likelihood of failure [319]. This is an essential point as it suggests that the failure rates may become significant when considering large and complex components.

Colosimo et al. developed a complementary model for the economic impact of in-situ monitoring tools in AM [320]. The authors conducted a series of studies to estimate the failure rate of typical AM components. They then calculated the break-even point of a hypothetical inspection system as a function of this failure rate. Unlike prior literature, the authors acknowledged no inspection tool is perfect and therefore accounted for type I and type II errors in the detection. One interesting point is the model developed by the authors assumes that defective parts which have not been scrapped during processes due to a type II error are still flagged as defective post-build. Whilst not explicitly stated, this assumes that post-build inspection is still required. The conclusion was an estimation of $\sim 9\%$ saving in the cost of additive components could be achieved by the use of in-situ monitoring. All of the tools the authors considered were passive, observing the emission spectra of the build process. These are essentially 'free' measurements in that there is no temporal impact on the build; therefore, only the cost of the equipment needs to be considered. However, a prospective SRAS system will need to scan the sample, causing some addition to the time, hence the importance of the time element of the model compared to existing literature.

7.3 System costing

7.3.1 Additive build system

In order to model the accurate costs of SRAS based inspection, the first step is to calculate the capital equipment costs, including the build system. The calculations in this chapter have been conducted for a RenAM 500Q L-PBF build system, fig. 7.1. Costs and associated return period of the build system are given in table 7.1. Values such as the system lifetime are open to debate, but five years is taken as an approximation in agreement with costing literature discussed in §7.2.

Multiple laser build systems are becoming increasingly common. In the future, this may open the door to more exotic processing routes, for example, using the lasers in tandem to pre-heat and build. However, the current driver is simply to realise faster build times. As such, the lasers work independently, dividing the powder bed up into sections, and the processing time scales linearly.

 Table 7.1: Specification of the build system and supporting costs.

| Machine Cost 1 laser | £k | 430 |
|------------------------|---------------------------------|----------------------|
| Machine Cost 4 laser | £k | 650 |
| Return Period | Years | 5 |
| Setting time per build | Hours | 2 (1 pre + 1 post) |
| Staffing rate | $\pounds\cdot \mathrm{hr}^{-1}$ | 15 |



Figure 7.1: Photograph of Renishaw RenAM 500Q multi-laser additive system.

7.3.2 SRAS

Table 7.2 lists the components required for SRAS based inspection and basic costs associated. Most of these and the set-up can be seen in fig. 7.2. This assumes that the system will utilise the standard galvanometer steering mirrors and voice-coil assembly for focal plane control and the build PC for data processing. It can be seen that over 50% of the equipment costs are due to the lasers. The costs discussed in this section refer only to the cost of purchasing the equipment. It is beyond the scope of this analysis to account for the additional time required to build the system or for a profit margin on the SRAS instrumentation. §3.3 described the practical differences between the EMDA system and a rough surface capable SRAS instrument, but a few of the key requirements of rough surface equipment are emphasised below.

| Item | Function | Cost $[\pounds\ \mathbf{k}]$ |
|-----------------|--|------------------------------|
| Infra-red laser | Generate surface acoustic waves | 20 |
| Detection laser | Detect surface perturbation | 20 |
| Oscilloscope | Capture waveforms | 10 |
| Detector | Receive light from detector beam | see below |
| General optics | Fringe pattern imaging | 5 |
| Electronics | Amplifier chain and other signal processing features | 5 |
| Mask | Optical grating pattern for imaging | 5 |
| Total | - | 65 |

 Table 7.2:
 Basic costs of SRAS instrumentation

It should be noted that one crucial component not valued in the above table is the detector. Whilst the components listed above are integral to any deployment of SRAS, the detector system and concomitant costs vary based on the specific instrument selected. §1.7.3 contains a detailed discussion of the challenges posed by optically rough surfaces. The classic and 'cheap' two photodiode knife edge detector is unsuitable for as-deposited L-PBF applications. Costs of a two-wave mixing or Fabry-Pérot interferometry system can vary between $\pounds 60 - 90$ k. These turnkey systems include the detection laser, and therefore $\pounds 20$ k can be removed from the total cost, giving a total system cost of approximately $\pounds 100 - 130$ k. Aside from the cost, one primary issue with photo-refractive techniques is the rate of adaptation, which tends to be in the order of tens of Hz, limiting the scan rate.

The speckle knife edge detector (SKED) is currently developed and licensed by The University of Nottingham under patent, GB201120774D0. The cost of this device is currently difficult to estimate; the cost of fabricating a SKED is nominal compared to the commercially available systems, but at present, no information is available to capture the cost of licensing fees accurately. As such, calculations in this report will assume a nominal SRAS system cost of £100k, but to fairly judge the implications of a detector choice, the analysis will illustrate the effects of increasing SRAS system cost, where appropriate.



Figure 7.2: SRAS rough surface set-up, showing primary components. Physical arrangements of a SRAS system suitable for integration within an additive build system are discussed in §7.7.

Aside from the lasers, the majority of components are non-specialist equipment that can easily be acquired. Due to the specialist nature of the lasers, it is worthwhile noting the specifications of the lasers used in the current SRAS rough surface system, as used to acquire the velocity map in fig. 7.6. Firstly, the operational properties of the generation laser can be found in table 7.3.

| Specification | Unit | Value |
|-----------------|------|-------|
| Wavelength | nm | 1064 |
| Pulse width | ns | 10 |
| Average power | W | 4 |
| Repetition rate | kHz | 5 |
| Pulse Energy | μJ | 800 |

 Table 7.3: Specification of current generation laser, Elforlight FQ-600-5-Y-1064

Important features worth noting include the pulse width, which determines the frequency content contained within the pulse. Too long a pulse length, > 20 ns, will mean there is no suitably high-frequency content for the SRAS inspection. Additionally, increasing the pulse energy increases the amplitude of the SAW packet, thus improving Signal-to-Noise Ratio (SNR). Although, a balance must be struck between large amplitude waves and ablation. Finally, the repetition rate is the critical factor in determining the scan rate of the inspection.

The properties of the detection laser are given in table 7.4. The previously described EMDA system uses significantly less power, < 500 mW, however, the presence of an optically rough surface means that only a small fraction of the incident laser irradiance is returned to the detector. By working at a lower acoustic frequency (controlled by selecting a longer wavelength from the mask), a larger detection spot can be used, allowing a greater power of detection laser to be used without damaging the surface.

Table 7.4: Specification of current detection laser, Cobolt Samba

| Specification | Unit | Value |
|-----------------|------|-------|
| Wavelength | nm | 5064 |
| Average power | W | 1.5 |
| Repetition rate | kHz | CW |
| Noise, rms | % | < 0.1 |

7.3.3 X-ray computed tomography

The existing literature compares inspection against a non-inspected build but does not address the point that post-build inspection still has an associated cost. The integrity of AM components is currently assessed by XCT. Therefore, to truly reflect current practice, the cost of post-build XCT must also be accounted for. Generally, XCT is used as a service rendered and thus, a
flat rate of cost per part is used to estimate the cost of XCT inspection. Based on available industry information, the cost of XCT based inspection is taken as \pounds 200 per component. Based on information from Renishaw, this represents a system capable of inspecting one component of the sample build (§7.5.1), with the ability to detect defects on the order of \ge 100µm with a focus on detecting 'critical' defects of \ge 250µm.

7.4 Definition of defects and the impact on processing routes

Before proceeding to the presentation of the cost-time model, it is worth defining what is meant by a defect in the context of this chapter. Unfortunately, the following is necessarily laborious in the semantics of the manufacturing chain but introduces a pertinent argument on the place of SRAS within this chain.

As demonstrated in chapter 5, the optical datasets which accompany SRAS measurements can be used to detect surface breaking porosity and the acoustic datasets have some capacity to sense subsurface porosity. However, it is clear that SRAS is a more suitable tool for probing for microstructural anomalies. Therefore, a 'defect' in the parlance of this chapter refers to any anomaly, including microstructural effects, deviating from the intended design.

We have outlined above the three systems which will be considered in the following cost model: SRAS, the build system and an XCT inspection system. Several arrangements of these three systems can be imagined within the overall build workflow. For clarity, fig .7.3 presents the two schemes which will be considered in this chapter, where **A** represents the current practice and **B** a hypothetical SRAS in-situ inspection. The SRAS process denoted here may be replaced by any suitable in-situ inspection tool.

Presently, the process is straightforward, a part is designed, the build machine prepared, and the entire component fabricated without further intervention. On completion of the build, the part is removed from the build system and then interrogated. If it is found to fall outside of acceptable conformance, then the part is scrapped a replacement must be made *de novo*. How a failed build is 'corrected' is an exercise that generally relies on operator skill and accompanying heuristics [319].

A less linear process is observed when an in-situ inspection is enacted. As the build progresses, the integrity is continually assessed, the built is terminated, and the partial component scrapped. The reader should note the hatched XCT process in scenario **B**, as this raises a key point. The literature covered in $\S7.2$ has not sought to replace XCT in the process; instead, the authors have focussed on using 'free' passive measurements to detect erroneous builds, allowing these builds to be scrapped earlier. All completed builds are then subjected to post-build inspection, as in **A**, where further defective parts may be found.

At the outset of this EngD, the envisaged concept was to replace XCT with in-situ inspection, and whilst this may be plausible in-future, the current capability of SRAS, in line with the work presented in this thesis, demonstrates the power of SRAS is found in the microstructure and elastic property assessment, rather than the detection of porosity, where XCT excels. Thus, for completeness, SRAS values will be compared to costs with and without XCT inspection. However, the reader should be aware that whilst possible to compare the costs associated with SRAS and XCT, it is not a like-for-like replacement. The likelihood being some form of separate volumetric porosity detection will be necessary in addition to a SRAS system.

However, this does not necessarily mean there is no case for SRAS within the build cycle. The aforementioned work of Colosimo et al., §7.2.2, demonstrated that the embodied cost of a final AM component could be reduced by in-situ inspection, even when the inspection system significantly increases the capital cost of the system (20%). This proposition is insensitive to the cause of build failure or type of defects, and therefore we may consider microstructure related issues, in addition to surface-breaking defects. It is this use case that shall be primarily explored in the following work, therefore.

Of course, in the case of a defect-free component with the incorrect microstructure, no amount of XCT will reveal the inferior mechanical properties, but as microstructure is not currently assessed as part of the normal build procedure, there is no 'like-for-like' for which to compare. Therefore, to caveat the conclusions presented in this chapter, values calculated using XCT represent a useful benchmark for the current cost of inspection, and putting the cost of a potential SRAS system in perspective - while noting the primary functions of these two inspection methods are different.

We may argue that a fairer comparison would be against current practice to assess the microstructure in production components, but as discussed at length in chapters 1 and 2, this is not readily achieved currently, due to a lack of suitable NDE methods - this is in essence one of the fundamental motivations of this thesis. Whilst it is true that the use of SRAS would provide a unique dataset, proving a wealth of information on the microstructure, which would hopeful enable a greater number of safety-critical parts to be created by AM; the aim of this chapter is to describe the monetary impact as far as possible. Therefore, as an 'acid-test' we can envisage a cost-free microstructure assessment tool that can be applied to the production components post-manufacture, rather than test facsimiles - this then has the same embodied cost as the no in-situ inspection case. This makes the cost comparison between the no in-situ inspection scenario and SRAS inspection the most relevant. If a there is a situation where the opportunities provided by in-situ inspection allow the SRAS to realise a lower embodied cost then there would be an extremely strong argument to pursue this avenue.



Scenario A: post-build inspection

Scenario B: in-situ inspection

Figure 7.3: The two charts presented here outline the workflow for **A** AM process with post-build inspection and **B** AM process with in-situ inspection.

7.5 Cost and time model

7.5.1 Sample build

Having estimated the capital cost of equipment for SRAS inspection, the temporal impact of inspection must now be parametrised - relative to the normal build time. A typical aerospace component manufactured using L-PBF is shown in fig. 7.4. A normal build has three of these components per build plate and takes approximately sixty hours to manufacture. This example build shall be used as a basis for estimating the temporal impact of SRAS inspection. The processing parameters used for this sample build are given in table 7.5. This component is not dissimilar to the Airbus A320 nacelle hinge bracket discussed in the work of Tomlin and Meyer [321], which is later used by Colosimo et al. as a case study example [320]. Note that the size of this component is at the upper-limit of what can be analysed by XCT without sectioning the sample.

The build time and cost of any build described by the parameters given in table 7.5 may be evaluated by this model; for example, the benchmark build described by Baumers et al. [317] was also used as an input to compare the predicted build times to validate the model. Whilst not necessary for the work at hand, the logical progression of this model would be to take the input from the build file used by the additive system directly.



Figure 7.4: Typical aerospace component manufactured using L-PBF.

 Table 7.5: Build parameters for sample component.

| Volume | m^3 | 9e - 4 | Layer height | μm | 30 |
|-----------------|-------------------------|--------|--------------|----------------|-----------|
| Deposition rate | ${\rm m}^3{\rm h}^{-1}$ | 5e-6 | Density | $\rm kgm^{-3}$ | 4430 |
| Feedstock cost | $\pounds kg^{-1}$ | 160 | Material | - | Ti-6Al-4V |
| Parts per build | - | 3 | Build lasers | - | 3 |

7.5.2 Time

To estimate the time spent inspecting the sample build, the 'speed' of the inspection system must be defined. The 'speed' of the SRAS system can be essentially broken down to the rate of data acquisition and the number of data points to be captured. The number of data points to be captured can be considered as a function of the generation patch size (which defines the spatial resolution) and, in turn, the acoustic wavelength. Fig. 7.5, plots the impact on build time as a function of points acquired per second and the wavelength of inspection, as defined by equation 7.1.



$$R = \frac{T_{build} + T_{scan}}{T_{build}} \tag{7.1}$$

Figure 7.5: Ratio of build+scan time to build time, as a function of inspection wavelength and data points captured per second. The white marker indicates the input subsequently used for the cost-time model in this chapter. The discontinuous nature of this figure is due to the assumption that the acoustic wavelength equates the depth of inspection and therefore the number of build layers probed in one scan can be increased by increasing the acoustic wavelength.

Where T is the time spent on that activity. The closer to one the ratio is, the smaller the impact of inspection on the total build time. With increasing wavelength, more build layers can be inspected per scan, hence the discontinuous nature of the contours. This assumes the

mode conversion phenomena can be used for void sensing, to depths of approximately 1λ . On first inspection, it would seem desirable to work as close to the lower right-hand side as possible in order to minimise the temporal penalty of inspection; however, increasing the wavelength of inspection leads to a decreased resolution and decreased measurement fidelity. Referring to existing defects classification tables previously published in the literature, it can be seen that defects occur over a wide length scale range, and a 300 µm presents a sensible lower bound of resolution in order to capture the majority of defects.

Clearly, the acquisition rate of the SRAS system is a vital parameter for the model, and hence some context is required for the speed of rough surface SRAS systems. Fig. 7.6 presents SRAS data from two as-deposited powder bed specimens, captured using rough surface systems. To obtain this data, 200 averages of each acquired waveform were required in both scans. The parameter points per second shown in fig. 7.5 is then simply $\frac{\text{rep rate}}{\text{averages}}$. In both cases, the final acquisition rate is 25 points \cdot s⁻¹, in comparison, the previously described EMDA system has an acquisition rate of 2000 points \cdot s⁻¹. Several factors feed into both the points per second variable, which need to be explained. On a smooth surface, the primary influence on the points per second is the repetition rate of the generation laser. In the 'on-the-fly' scanning approach, as used by the EMDA system, one pulse is equivalent to one measurement point - thus, a 5 kHz repetition rate laser allows 5000 measurements per second. However, on rough surfaces, the signal-to-noise ratio (SNR) is significantly lower and necessitates averaging of the acoustic waveform to improve SNR. This can be achieved in two manners, temporal or spatial averaging. In temporal averaging multiple traces are captured at a specific point and averaged to give the measurement for that point before proceeding to the next point. In spatial averaging, a succession of single waveforms are captured across a small region and averaged together to give the measurement for that area. Both approaches are valid, but from a practical viewpoint, spatially averaging is significantly faster. By spatially averaging, the system of motion, i.e. galvanometer mirrors or scanning stages, can be driven at constant velocity. However, only temporal averaging is used in the case of the roughsurface data presented in fig. 7.6. the actual scan rate would be further reduced thanks to the time to move between scan points when averaging like this (for more information §3.2.4 has discussed the differences between point-to-point scanning and 'on-the-fly'); however, for the calculations at hand we will assume spatial averaging could be implemented and neglect the time to reposition the generation image and detection spot between measurements.

It would be preferable to correlate the number of averages needed with surface roughness and integrate this within the model, however, the relationship is extremely complex. Simply describing surface roughness with a statistical parameter like R_a does not capture the complexity of the surface. The specific acquisition rate chosen for the analysis presented here is somewhat open to



(a) Optical microscopy image of asdeposited titanium L-PBF specimen surface. (b) SRAS velocity map of surface shown in (a), captured using SKED system.



(c) SRAS velocity map of as-deposited aluminium L-PBF specimen surface, captured using SRAS-QT system. The top right and lower left edges of the specimen move out of focus, hence the greater loss of data in these areas.

Figure 7.6: Examples of high quality SRAS data captured on as-deposited L-PBF surfaces. Compared to the velocity maps presented in previous chapters, the velocity maps (b) and (c) clearly have a greater drop-out and lower spatial resolution. The differences between the systems used to captured these maps and EMDA has been discussed in §3.3.

debate, but as the existing rough-surface systems are not optimised for speed, it is reasonable that a production system might achieve higher scanning rates - although still slower than the EMDA system. The assumed value is marked in fig. 7.5, along with the two rough surfaces systems of fig. 7.6 and the EMDA system described in Chapter 3.



Fig. 7.5 may be more clearly understood by fixing the points captured per second and extracting a single line plot, as shown in fig. 7.7a. It is clear to see that at high resolutions, small changes in the wavelength can have a major impact on the scan to build ratio, whilst at lower resolutions this effect diminishes. The exponential rise in scan time at resolutions better than 190 μ m make this area infeasible to operate in. Given the usefulness of the data drops off sharply after around 350 μ m, this actually provides a relatively small operating window, approximately 190 μ m – 350 μ m. Thus, rather than increasing inspection wavelength, the primary route for temporal savings is by moving down on the vertical axis, increasing points captured per second.

Fig. 7.7a assumes a fixed number of six grating fringes are present in the generation patch. The effect of the number of fringes in a patch on the effective spatial resolution can be seen in fig. 7.7b. As the number of fringes in the patch increase, the excitation becomes more narrowband, improving the velocity resolution. Taking a six fringe patch, this gives a maximum wavelength of $\sim 100 \,\mu\text{m}$ for 300 μm resolution. In the performance used for calculations in this report, a 5 kHz repetition rate laser is used, with 10 averages per point, and a wavelength of $\sim 100 \,\mu\text{m}$, as indicated by the white mark in fig. 7.5.

Fig. 7.5 shows the ratio of time spent on each major part of the L-PBF process, showing that nearly twice as long is spent building the components as inspecting for the sample build. The time to set up the system and recoating the powder bed between layers is of minor importance for this component - although recoating may play a role in tall parts with a narrow cross-section, where the ratio of time spent recoating to irradiating the powder bed increases. The cost of setting up and ending a build by the operator is taken as a fixed time, and thus, given the relatively long build time, this is also of little importance in this component. The model predicts the normal build operation to take 66 hours which is in good agreement with reality, suggesting the component modelling the build process is valid.



Figure 7.7: Minimising the time of SRAS scan is a trade off with resolution, which in turn is a function of the number of the acoustic wavelength and the number of fringes in the generation image. The line shown in (a) is extracted from fig. 7.5 for a single acquisition rate; hence the discontinuous nature is again due to the number of layers inspected increasing with wavelength assumption.



Figure 7.8: Percentage of total build time spent on each of the major activities in the L-PBF process, for a single build of the example component.

7.5.3 Cost based on annual production

The knowledge of the time to complete a single build of the sample component, combined with an understanding of the costs of a SRAS system, the L-PBF process, and XCT analysis, allows the cost comparison between fabrication without inspection and with inspection by either XCT or SRAS. The net cost of an additive part may be found by the sum cost of the component activities, equation 7.2.

$$C_{\rm net} = C_{\rm material} + C_{\rm build} + C_{\rm setting} + C_{\rm inspect} \tag{7.2}$$

Where the C_{material} is the cost of the raw feedstock, C_{build} is the amortised capital cost of the AM system, C_{setting} is the operator cost to set up the machine for the next build and C_{inspect} is the amortised capital cost of the inspection system.

Given the build time is relatively long, ~ 66 hours, the difference in machine output with and without inspection needs to be assessed over a suitably long period. For this, a year has been selected, with the output from a single machine over this time period compared. Upon the completion of a build cycle, an operator is required to reset the machine, removing the build and preparing for the next cycle. Based on industrial guidance from Renishaw, two situations are of interest with respect to the operator working: firstly, a shift attendance where an operator is only available during the hours 0900 to 1700 and secondly, a full-attended scenario, where an operator is always available to prepare the next build. In this scenario, a new build is initiated immediately after the completion and machine resetting of the previous build. In contrast, in the shift-attendance scenario, builds completing outside of the attended window wait until the next shift period before manufacture is recommenced. Whilst this analysis is quite specific, it is worth discussion as it reveals two operating regimes, from fig. 7.10 it is seen that despite the relatively long build time, the build system is significantly underutilised in the shift attendance case (a), compared to the full attendance case (b), resulting in 74 fewer builds completed over the course of a year. Therefore, the important point is that these two scenarios represent a build system fully utilised and under-utilised, with the industrial drivers of this utilisation less important to the conclusions of the current analysis. Calculations are first performed on a single machine basis for full-time and shift attendance. Fig. 7.9 shows components output and the costs associated on a per-kg basis. The costing categories are summarised as: feedstock, raw powder cost; machine, the capital cost of build system; staffing, cost of staff to attend machine and set up new builds; inspection, the capital cost of inspection equipment.

Returning to the discussion of inspection, we see for shift attendance, SRAS manages the same build throughput over the year as a build with no inspection. The downtime in shift inspection is such that the extra-build time due to scanning is not penalised. This suggests that whilst SRAS inspection can increase the build time by $\sim 50\%$; this additional time may be absorbed by the typical underutilisation seen in industrial systems, resulting in the same number of parts being produced. This interesting result challenges the initial preconception that the scanning speed of a SRAS system is the greatest obstacle to in-situ inspection.

Unsurprisingly, both forms of inspection lead to parts with a higher specific cost compared to no inspection - in both cases, the specific cost increases by over 100%. Compared to the no inspection cost case, there is an increase of $\pounds 4750 \,\mathrm{kg}^{-1}$ in the embodied cost for the SRAS inspection. Comparing XCT and SRAS, it can be seen that SRAS has a higher associated cost of inspection than XCT by around $\pounds 170 \,\mathrm{kg}^{-1}$ for shift attendance. One might conclude from this that whilst SRAS is the most expensive option, it is on the same order as the cost of XCT and depending on the actual cost of the detector it might be slightly cheaper. If we assume a SRAS system cost of $\pounds 50 \,\mathrm{k}$ then XCT inspection becomes the most expensive option by $\pounds 70 \,\mathrm{kg}^{-1}$.

However, as the duty cycle approaches 1 for full-time attendance SRAS based inspection has a throughput of around 50 builds less per annum. A throughput of 75 builds per annum has reduced the SRAS inspection costs by around $\pounds 300 \,\mathrm{kg}^{-1}$. However, a direct comparison with other scenarios is unfair due to the reduced output. The previous section compared the costs based on a single machine output, but, depending on the use case, it may be more realistic to fix the output per year, at 127 builds, and compare the costs. Where SRAS based inspection yields a reduced build rate, production will be balanced by the addition of another build system(s), with costs on a pro-rata basis. Results are shown in fig. 7.10. Comparing XCT and SRAS, it can be seen that inspection costs between SRAS and XCT are comparable, but as the build rate is reduced in order to fabricate 127 builds per annum an extra portion of a build system is required to balance output for comparison, hence the difference in machine costs per kilogram, $\pounds 670$ compared to £220. In total, the calculations show XCT to be cheaper by around $\pounds 700 \, \text{kg}^{-1}$ compared to SRAS, $\pounds 870 \,\mathrm{kg}^{-1}$ and $\pounds 1580 \,\mathrm{kg}^{-1}$, respectively. Over the five-year life span of the build system, this would extrapolate to an additional cost of around $\pounds 1.8M$ (this is driven by the assumption additional build systems are required to make up the short fall). This suggests that in the case of maximum utilisation, the time impact of SRAS becomes a critical factor and is therefore not a viable method.



Figure 7.9: Annual output and costs per machine. The left-hand side column shows output for shift based attendance, whilst full-time attendance is given in the right-hand column. The costs shown are broken down by activities, the total costs per kg are (a) £890 (b) £560 (c) £1370 (d) £1060 (e) £1190 (f) £870.



Figure 7.10: Comparison of elemental costs on a per-kilogram basis for (a) SRAS-based inspection and (b) XCT-based inspection, based on fixed per-annum output. The costs shown are broken down by activities, the total costs per kg are £1580 for SRAS inspection, approaching double the cost of XCT based inspection at £870.

7.6 Build failure

L-PBF is well suited to the manufacture of complex components, such as the sample build §7.5.1. However, the manufacture of such components is plagued by relatively high failure rates. The work presented in this chapter thus far has not leveraged any of the benefits of in-situ inspection. Intuitively, the presence of in-situ inspection facilitates early termination of a flawed build - as opposed to post-build evaluation. In theory, this should save the cost of the material and processing to complete the component. The following section presents modifications to the cost-time model to account for build failure. This body of works assumes operation in the under-utilised domain, where the extra build time due to inspection does not reduce the function of the system and that there is sufficient machine availability to fabricate replacement components.

7.6.1 Scrap rate

For the sample build the failure rate, $\gamma \in [0, 1]$, is estimated to fall in the range 0.2 - 0.5, in line with values suggested in literature. Assuming that a defective part is scrapped and replaced *de novo*, the final cost of the component without inspection is found by equation 7.3.

$$C_{\text{part, as-is}} = \frac{C_{\text{net}}}{1 - \gamma} \tag{7.3}$$

The advantage of in-situ monitoring is that a build may be terminated as soon as a defect is observed, minimising the wasted build time and material. The cost of a suppressed build can be found by equation 7.4, where $\eta \in [0, 1]$ is the completion of a build before termination due to a defect.

$$C_{\text{scrap}} = \eta \left(C_{\text{material}} + C_{\text{proc}} + C_{\text{inspection}} \right) + C_{\text{setting}}$$
(7.4)

Equation 7.5 then formalises the relationship between failure rate and scrap cost to determine the total cost of a component.

$$C_{\text{part, inspect}} = \frac{(1-\gamma)C_{\text{net}} + \gamma C_{\text{scrap}}}{1-\gamma}$$
(7.5)

From this equation, fig. 7.11a compares the cost of one aerospace bracket as a function of build failure rate, γ , for three scenarios (i) no inspection cost, (ii) XCT inspection and (iii) SRAS inspection. It can be seen that for failure rates below 0.1, typical of simple builds, then SRAS does not prove a cost-effective tool in this scenario. On the other hand, as the failure rate increases, an in-situ inspection tool capable of informing early termination becomes cost-effective for failure rates ≥ 0.15 , as may be expected in complex and long-run builds. The additional £100k for the inspection system is offset by the reduced waste in materials and build time.

The cost of the inspection system, the failure rate (γ) and the completion rate (η) before build suppression are the three inputs to equation 7.5 which are not well defined. Therefore, fig. 7.11b calculates the break-even cost of the in-situ inspection system as a function of η and γ . As would be expected, the viable cost of an inspection system increases with higher rates of failure rate and falling completion rates. The upper left-hand corner raises an interesting, if likely to be strictly hypothetical, point; here, the failure rate is so high (> 0.35) that an inspection system equal to the cost of the original build system could be justified. This emphasises that whilst the initial cost of the build system is significant, the value of lost feedstock to scrapped builds can also run into the hundreds of thousands, although this is likely to be harder to appreciate as it gradually accumulates rather than being a one-off penalty.

It is important to note that the completion rate is influential on the break-even cost due to the relatively long build time; if the sample build used in this analysis took only one hour, for example, then the completion rate would have significantly less impact on the cost-saving. Nevertheless, it is important to note this effect as in prior literature authors have attempted to qualify the failure rates of builds - but less detail has been captured on the progress of the build when the failure occurs. Further investigations to shed light on this would be informative - hypothetically, if failure occurs early in the build, $\eta = 0.1$ for example, then the case for in-situ inspection is certainly strengthened.

The analysis presented thus far in this section assumes SRAS is 100% reliable in the detection of defects of interest. Of course, we may use the standard type I and type II errors terminology to further parametrise the problem, fig. 7.12. α is the probability of type I error and likewise β the probability of type II error, fig 7.12.

The cost of a part in the inspection scenario, taking into account the possibility of false alarms and false build suppression, may then be found by equation 7.6. Given the number of terms in this equation, the influence of each parameter is not immediately apparent. An increasing value of α classifies a greater number of parts as defective, leading to the undesirable scrapping of conforming builds and inducing a replacement cost. In contrast, an increasing value of β classifies a greater number of defective parts as conforming.

$$C_{\text{part}} = \frac{(1-\gamma)\left(\frac{(1-\alpha)C_{\text{net}} + \alpha C_{\text{scrap}}}{1-\alpha}\right) + \gamma\left(\beta C_{\text{net}} + (1-\beta)C_{\text{scrap}}\right)}{1-\gamma}$$
(7.6)

The embodied cost of the part for two high and low error scenarios are presented in fig. 7.13, the values used for α and β are in-line with previous work [320]. In fig. 7.13(a) when the associated



(a) Breakeven cost, as a function of failure rate, for inspection in-situ and post-build, costs are shown for XCT and SRAs along with illustrative 'free' scenarios.

(b) Breakeven cost for SRAS in-situ inspection, as a function of failure rate and completion rate.

Figure 7.11: Comparison of break even points between SRAS in-situ inspection and the conventional post-build inspection. The merit of SRAS in-situ inspection is well correlated to the failure rate of the component, making clear that this would only be viable in complx parts with high failure rates.

error rates of the inspection are low SRAS becomes a cost-effective solution when $\gamma \sim 0.14$, similar to the no error case presented in fig. 7.11. When both error errors increase by a factor of 10, the break even failure rate moves considerably - $\gamma \sim 0.25$ in the case of the no inspection cost scenario.

It is again true that the values of α and β are not well known, and it, therefore, makes sense to repose the problem and find the 'break-even' failure rate as a function of these two parameters. From fig. 7.14, it can be seen that the failure rate window spanned by the break even failure rate is relatively narrow (0.15 to 0.25), and therefore, for the purpose of this analysis, the impact of inspection errors can be neglected - given that there is already considerable uncertainty in the failure rate of the component, it is not possible to draw conclusions of greater detail on the error

| | Defect present | No defect present |
|-----------------|---------------------------------|----------------------------------|
| Defect detected | Correct outcome True postive | Type I error False positive |
| Defect free | Type II error False negative | Correct outcome True negative |

Figure 7.12: Chart defining the two error types of false positive (type I error) and false negative (type I error). Accordingly, the false positive rate is $\alpha \in [0,1]$ and the false negative rate is $\beta \in [0,1]$.

window of inspection.

In reference back to fig. 7.3, if SRAS replaces XCT then the probability of false-negative (i.e. a missed defect) must be zero ($\beta = 0$), for the purpose of cost calculation as the part will not later be flagged as defective and needs to be replaced. Similarly, if the defect in question is one not usually probed as part of the workflow, such as microstructural anomalies, then again $\beta = 0$. §7.4 has discussed the definition of defects, for the purposes of this chapter, in greater depth.



Figure 7.13: Comparison of break even costs for two error scenarios. Increasing errors associated with the inspection system requires a failure rate of the component to justify there use. These errors are manifest as correctly build parts being incorrectly assessed as defective (α error), and defective parts being classified as conforming (β error).



Figure 7.14: The values of α and β are not well defined, it makes sense to repose the problem as a function of α and β , such that the impact of these to parameters on the break even failure rate may be discerned. The break even point is the inherent failure rate of a part at which in-situ inspection becomes cost-effective, compared to build then scrap.

7.6.2 Opportunity to repair

The results presented in §7.6.1 have assumed that following the detection of a defective build, the only available course of action is to scrap and begin anew. A logical approach is then to consider the possibility of repairing the defective part in-situ.

A brief schematic description of a repair, as it is envisaged, is given in fig. 7.15. A layer of the additive manufacture process is illustrated in fig. 7.15a, followed by component measurement fig. 7.15b, a decision on cancelling the build is made in fig. 7.15c, rework path definition in fig. 7.15d, applying the rework and verify in fig. 7.15e, with the final part completed at the end in fig. 7.15f. This build-characterise-repair strategy allows for the integration of a measurement technique and repair strategies within an AM build environment to allow targeted repair on a layer-by-layer basis.

The benefits of this approach are that the number of scrapped components, along with wasted feedstock and processing time, can be further reduced beyond that associated with early termination, as discussed in §7.6.1. Existing literature considering the cost implications have explicitly not considered the possibility of repair as this is not a feature of current generation machines. However, as in-situ repair is a logical progression and the level of research into this is on the rise (as will be outlined in §7.6.4), it is worthy of evaluation.



Figure 7.15: Flowchart of proposed targeted rework approach simulated in this work. (a) outlines the completed processing of an L-PBF layer; (b) shows a measurement step, where defects are identified and it is determined to either cancel the build or to rework the defects; (c) if the defects are above a predefined threshold, the build may be cancelled; (d) if below threshold, through data processing, rework paths can be defined for (e) targeted re-melting with the option of remeasurement to, again, cancel, or continue manufacture; This process may repeat on a layer per layer basis with the finished part shown in (f).

7.6.3 Cost of repair

Following a similar methodology to that of §7.5, the cost of a repaired component maybe found by equation 7.7.

$$C_{\text{part}} = C_{\text{net}} + nC_{\text{repair}} \tag{7.7}$$

where n is the number of repair events per build, C_{repair} is the cost to repair the defect, a function of any additional processing time and material required to enact the repair. Based on the work thus far, equations 7.7 and 7.5 may be combined to give equation 7.8, describing the final embodied cost of a part.

$$C_{\text{part}} = C_{\text{net}} + \frac{\xi C_{\text{scrap}}}{1-\xi} + (\gamma - \xi)nC_{\text{repair}}$$
(7.8)

where ξ is the failure rate of parts that cannot be repaired - hence, $\gamma - \xi$ is the failure rate of repairable parts. Therefore, to conclude the cost analysis of this chapter, fig.7.16 compares embodied cost of an inspected and repaired part against the cost of post-build inspection as a function of repairable rate, $\gamma - \xi$, for a fixed failure rate of $\gamma = 0.2$, thought to be conservative for the sample build.

Fig. 7.16 plots the cost of a component as a function of repair rate and number of repair operations. This suggests that even for relatively low levels of repair $(0.2 > \gamma - \xi)$, it becomes the cost-effective choice.

We could then extend this by modelling the rate of defect formation by a Poisson distribution, well suited to modelling the number of distinct events in a fixed period - in this case we care about the number of defects forming within the build time. However, the model has now moved far beyond contemporary, and the wisest course of action would be to stage a series of contemporary experimental studies to validate these results as far as possible.

7.6.4 Preliminary demonstration of repair by reworking surface defects

Having established there is a cost basis for enacting repair, the reality of achieving this will be briefly investigated. Previously, laser re-melting by applying additional skin scans has been employed to improve surface roughness [322, 117]. It was shown, however, that compressive residual stresses were created as a consequence of this re-melting process. Further to this, Yasa et al. [323] showed that complete layer re-melting for defect rework on steel L-PBF components could reduce near surface porosity, but at the expense of a change in microstructure akin to a recast layer. This approach has been reported to yield L-PBF parts with unpredictable and anisotropic microstructures [324]. Mireles et al. [325] used IR imaging for in-situ monitoring in a selective electron beam



(a) Embodied cost per part as a function of repairable rate and failure rate,
★ indicates the break even point between inspection with repair and no inspection (i.e. build then scrap).

(b) Comparison of the embodied cost per part as a function of the repair rate and cost, for the cases of repair, inspection with early build termination to scrap and no inspection.

Figure 7.16: Embodied cost of a part when it is possible to repair. The breakeven point is now dependent on the failure rate (γ) and the repairable rate ($\gamma - \xi$). As the repairable rate tends towards zero ($\xi \rightarrow \gamma$) then results approach those presented in fig. 7.11.

melting process, reworking artificial defects in Ti-6Al-4V components via large area re-melting. It was shown that re-melting successfully restored component integrity, however the authors used two inspection methods, thermography and X-ray computed tomography, for defect sizing - this yielded significantly different results, thus the efficacy of re-melting the artificial defects remains unclear. Furthermore, the defects seeded in the component were an order of magnitude greater than those seen to commonly occur in L-PBF in practice.

The focus of this study is to show defects, identified using NDE techniques at the surface of a given layer, can be reworked using defined scan strategies in a bid to repair them. SRAS is utilised in this study to locate and characterise surface defects, and a number of targeted rework strategies are investigated for their ability to rework the defect. The reduction in the size of an array of defects within an L-PBF build subjected to these various strategies is then presented. Information on the defect type, location and depth obtained pre- and post-rework is shown, along with a texture map and a methodology in which the proposed technique can be integrated into L-PBF builds to guarantee part integrity is discussed.

Methodology

L-PBF test specimens were manufactured using a Realiser SLM50 with the nickel superalloy, Inconel 718, the composition of which can be found in the literature [326]. A $10 \times 10 \times 10$ mm cube was built using a meander scan strategy with a hatch spacing of 60 µm and a 67° rotation between each layer. A laser power of 100 W and scanning speed of 560 mms⁻¹ were used to build layer thickness of 40 µm. Gas atomised powder sized $15 - 45 \mu$ m was used as the feedstock.

After manufacture, the specimen was polished to a mirror-like surface to facilitate inspection by SRAS and optical microscopy. The prepared surfaces were also imaged using a focus variation microscope (Alicona InfiniteFocus G5) with a $20 \times$ lens and in-built image stitching tool. The lateral and vertical resolution of the 3D micrographs was $2.94 \,\mu\text{m} \cdot \text{px}^{-1}$ and $0.1 \,\mu\text{m} \cdot \text{px}^{-1}$, respectively. The original build parameters were reused for the rework operation. In order to solely investigate the effect of re-melting, no additional powder was used in the rework process.

Three pore rework strategies were investigated, as shown in the schematic presented in fig. 7.17: A hatch melting pattern illustrated in fig. 7.17a; a spiral melting pattern in fig. 7.17b; and singleshot melting for pores similar in size to the laser spot size (40 μ m) in fig. 7.17c. The hatch melting pattern was selected since the samples were manufactured identically. The spiral melting pattern was proposed by Jhabvala et al., showing optimised thermal gradients in both the x and y direction designed to avoid overheating [327]. Since re-melting hatch patterns are limited in physical size (< 300 μ m) by the size of the defect, more complex melting patterns such as fractal, chessboard and 'paintbrush' scanning are not feasible as these are designed for bulk part processes.

In order to generate rework paths, an automated algorithm was developed, based on the flow chart in fig. 7.18. The original optical SRAS data is shown in fig. 7.18a. A size thresholding step was then used to disregard defects and outliers smaller than 20 μ m from the SRAS optical scan data, shown in fig. 7.18b, given the image resolution of 10 μ m (pixel size) and the defined minimum cluster size of 2×2 pixels [328]. This step can be adjusted to ensure that only those defects that take the part out of specification are reworked. A filtering step, based on eroding and dilating, closes non-uniform surface defects, fig. 7.18c. This step is necessary for avoiding partial or uncontrolled splitting of hatch patterns on a targeted defect. The resultant binary image from these steps is then processed to segment defects based on their aspect ratio; the centroid and dimensions are obtained, fig. 7.18d, to only target pores. The final processing step is the generation of the rework



Figure 7.17: Scan strategies for pore rework trialled in this work: (a) a simple hatch pattern as used in the manufacture of a part; (b) a spiral scan pattern where the gap between the scan lines is equal to a single hatch spacing; (c) discrete shot exposures for defects smaller than the laser spot diameter, $40 \,\mu\text{m}$.

paths (fig. 7.18e). In this step the rework area is defined by dilating the defect area by 150%. The laser path vector spacing for both hatch and spiral rework was defined as 40 μ m. All scan strategies were defined from the outer edge of the defect inwards, in order to draw material into the defect. The 150% repair area was selected in order to process sufficient re-melt material to sufficiently reduce the defect size. An example rework path as created by the algorithm is shown in fig. 7.18f, with a zoom of the highlighted sector in fig. 7.18e. The resulting vector rework paths were then verified through simulation in the ReaLizer RDesigner software to ensure data integrity.



Figure 7.18: Micrograph analysis steps outlining the order of operations for preparation of pores for rework: (a) the original optical data set; (b) removing outliers (defects $< 20 \,\mu$ m); (c) morphological closing based on eroding and dilating the binary image; (d) segmentation based on their aspect ratio and centroid for determination of dimensions and location; (e) the rework paths are created and overlaid; (f) magnified image of rework area showing individual hatch lines. The defect highlighted has been denoted d1 and its respective rework path has been denoted r1.

Results

To show the efficacy of the rework, focus variation micrographs of the samples were obtained before and after the rework was conducted to extract defect depth profiles. In order to define a successful rework, the height difference between the surface and defect depth should be significantly reduced. It is expected that complete removal of the defect would be impossible without additional feedstock; however, as this would be a difficult extra step, the aim of this work is to reduce defect size to within tolerance rather than to completely eliminate it.

Examples of defects before rework are shown in fig. 7.19a – d. These are all non-uniform acicular pores, as defined by shape morphology. The resulting surfaces after applying the rework strategies are shown in fig. 7.19e – h. A hatch pattern was applied to fig. 7.19e and fig. 7.19f for defects of $\sim 350 \,\mu\text{m}$ and $\sim 120 \,\mu\text{m}$ in diameter, respectively. The hatch rework areas show good coverage, completely encompassing the underlying defect. The individual melt tracks are visible, akin to a conventional hatch pattern in bulk material addition. The result of a spiral rework is shown in fig. 7.19g on a defect 150 μ m in size. Here, the resulting surface exhibits a smooth appearance with no individual melt tracks visible. A single shot rework was applied to the defect ($\sim 40 \,\mu\text{m}$ in size) highlighted in fig. 7.19d and the reworked region is shown in (h), alongside a larger spiral rework. The efficacy of rework of the single shots cannot be clearly identified from the micrographs alone, and quantification of rework is necessary in order to classify these.

In order to quantitatively evaluate the effectiveness of the rework, a 3D representation of the focus variation data obtained from the defect presented in fig. 7.19c is shown in fig. 3(i). The corresponding 3D representation of the region after the rework is shown in fig. 7.19k. Through plotting a 2D profile through the lowest point of the dataset, a cross-sectional view of the pre-rework (fig. 7.19j) and post-rework (fig. 7.19l can be extracted. The minimum depth of the defect pre-rework and post-rework are highlighted as d2 and r2, respectively. The pre-rework profile, d2, clearly shows a very steep trench, making it difficult for subsequent powder layers to backfill this defect in the L-PBF manufacture. Contrary to this, it can be seen that the slopes of the reworked area, r2, are more normalised compared to the initial defect. This will allow easier infill of powder particles in the next layer in the L-PBF manufacture. It must be noted, that through the rework operation, material has moved above the level surface to $+6.3 \,\mu\text{m}$. However, since the layer thickness of the L-PBF manufacture has been defined as $40 \,\mu\text{m}$ and can commonly be larger than this (> 50 $\,\mu\text{m}$), the material above the level surface should not impact production quality. Furthermore, the powder feedstock deposited for the next layer is still able to flow in and around the treated area.

Comparing d2 and r2, it is seen that the rework process has also reduced the depth of the defect



Figure 7.19: Optical micrographs of spherical pores of Inconel 718 samples (a) – (d) and their respective reworks (e) – (h); (e) and (f) are hatching rework strategies, (g) is a spiral rework strategy and (h) is a single shot rework. 3D views and line profiles of an exemplary pore are shown for pre rework d2 and post rework r2; (i) 3D view of existing surface pore; (j) line profile through the lowest point below the surface of the pore; (k) 3D view of the reworked area; (l) line profile through the lowest point below the surface of the rework. The line profile location is marked in red and blue for the pre rework and post rework, respectively.

relative to the nominal surface. This data is presented again in fig. 7.20, showing a comparison of the defect depths before rework and post rework for all three rework methodologies as extracted from line profiles through the lowest available points on all targeted rework areas. The minimum points from the profiles in fig. 7.19j and l labelled d2 and r2 have been transposed onto the data presented in fig. 7.20 to highlight the effectiveness of reworking the targeted areas.

fig. 7.20 clearly demonstrates the ability of selective re-melting to reduce the local variations in surface height caused by pores, opening the door to repairing surface breaking defects. A reduction in the depth of the defect along with a profile more favourable for infill with powder in the next build



Figure 7.20: Plot of defect depths for all three rework methodologies; hatch, spiral and single shot (for defects smaller than $40 \,\mu\text{m}$) with data shown for pre rework (red) and post rework (blue) for each. A reduction in variance can be observed for hatch and spiral rework patterns. The single shot rework patterns show no statistically relevant improvement. The minima of r^2 and d^2 are highlighted.

layer means that if such reworks are performed in-situ as defects develop, conditions are favourable for powder filling the defect in the next layer and the avoidance of the defect propagating through multiple layers.

Statistics of the data presented in fig. 7.20 are shown in table 7.6. The hatch rework methodology is shown to be more successful than the spiral method in terms of height reduction, with a higher reduction average of $\Delta 12.1 \,\mu m$ (50.0%), compared to $\Delta 9.5 \,\mu m$ (31.0%). The standard deviation of the data sets improved by 76.7% and 54.0% for the hatch and spiral strategies, respectively. Since Inconel 718 exhibits a relatively low thermal conductivity, the hatch pattern allows for heat dissipation outwards of the rework area, whereas the spiral rework traps the heat, which can result in over-melting at the centre, similarly reported by Jhabvala et al [327].

The single-shot method appears to have a minimal effect on the defect depth, with an average reduction of 17.1% and an increase in the standard deviation of 22.4%. The low probability of success of the single shot strategy can firstly be explained by the smaller initial pore depth leaving less room for improvement and secondly by the difficulty in realignment of the specimen for the rework. The small size of the single shot reworks resulted in many of the single shots being only partially reworked, thus limiting the effectiveness of the technique, as can be seen in fig. 7.19h. The latter point can be attributed to the angular relocation accuracy of $\pm 0.13^{\circ}$ in the build chamber

since the measurements were conducted ex-situ and the samples needed to be relocated. This problem will be eliminated when the method is applied in-situ.

Table 7.6: Mean and standard deviation of defect depths as presented in fig. 7.20 for both prerework and post rework.

| Repair | Pre Depth | Pre Depth | Post Depth | Post Depth | Δ Mean | Δσ |
|---------------|---------------------------------|---------------------------|----------------------------------|---------------------------|---------------|-------|
| Strategy | $\mathrm{Mean}~[\mu\mathrm{m}]$ | $\sigma \; [\mu {\rm m}]$ | $\mathrm{Mean}\;[\mu\mathrm{m}]$ | $\sigma \; [\mu {\rm m}]$ | [%] | [%] |
| Hatch | 24.2 | 14.1 | 12.1 | 3.3 | 50.0 | 76.7 |
| Spiral | 30.7 | 16.3 | 21.2 | 7.5 | 31.0 | 54.0 |
| \mathbf{SS} | 12.7 | 4.7 | 10.5 | 5.7 | 17.1 | -22.4 |

7.7 Implementation and improvements

The final consideration of this chapter is to look at how a SRAS system could be physically integrated within an AM build system, and the potential to address some of the underlying assumptions built into the cost-time model presented in this chapter which may lead to significant improvements.

7.7.1 Scan rate and time

§7.5.2 defines two domains of operation, full-utilisation and under-utilisation. The full-utilisation scenario has been neglected almost immediately due to the significant (53% increase in build time). To reduce the increase in build time to just 20% (a more tolerable figure as this would only reduce the per annum output by 19 components) would require a ~ 250% increase in scan rate. As shown in the above calculations, if a SRAS based inspection could achieve a throughput close to the nominal amount 'free' in-situ inspection, then there would be a strong case for the use of SRAS in both domains of operation. Therefore, prior to concluding this chapter is worth discussing if the SRAS scan really need be as slow as defined in §7.5.2.

Firstly, the simplest and most effective method for increasing the scan rate is increasing the repetition rate of the generation laser. In the current implementation of SRAS, one pulse is equivalent to one measurement - thus, increasing the number of pulses per second would linearly increase the effective scan rate. This can be continued until the point that the subsequent pulse begins to interfere with the previous measurement. A rough theoretical upper limit of the pulse rate can be calculated, as per equation 7.9, where the minimum time between pulses is the time for a SAW to propagate from the far edge of the generation patch to the detection position. This

distance is assumed to be no greater than twice the maximum patch size of 1 mm. A limit of 1.5 MHz is around three orders of magnitude greater than any of the current generation lasers used in SRAS systems. Unfortunately, whilst extremely high repetition rate lasers are available, these begin to trade-off other parameters increase the number of averages required. Therefore an improvement on the order of $\times 10$ is the maximum possible through this avenue.

$$t_{\text{propogation}} = \frac{2 \,\text{mm}}{3000 \,\,\text{ms}^{-1}} \approx 0.67 \,\mu\text{s} \qquad 0.67 \,\mu\text{s} \rightarrow 1.5 \,\,\text{MHz} \quad \text{rep rate}$$
(7.9)

Further changes to the hardware offer little scope to greatly increase the scan rate, but reconfiguring the nature of how the SRAS scan is conducted may provide a more to reduce the temporal impact. For example, calculations in this report have assumed that build time and scan time are asynchronous, but it may be possible to scan and build simultaneously. This will minimise the temporal and, in turn, the monetary impact of inspection.

Whilst L-PBF is characterised by extremely steep temperature gradients, an increase of 1° C in the material temperature is known to slow the SAW velocity by 5 ms^{-1} , on average (the reader is referred to chapter 8 for a more complete discussion of the effects of temperature on the SAW velocity). In the presence of large temperature gradients, there is, therefore, the potential for the SRAS measurement to sense primarily the change in temperature across the specimen, rather than microstructure. Scanning during deposition, of course, comes with its own challenges, temperature, the 'dirty' environment due to loose powder and spatter particle ejection and acoustic emission from the build laser. Points two and three can both likely be negated by ensuring a suitably large offset between the build laser system and scanning.

Several additional points could be beneficial to increasing scan-rates, for example it is anecdotally claimed that typical surface finishes are improving, leading to reduced surface roughness, thanks to the ongoing improvement of build systems. One may extrapolate fewer averages will be required to obtain a SRAS measurement of acceptable SNR, and thus the scan rate will be improved. However, by lacking anything greater than anecdotal evidence such assertions stray in to the realm of 'hope', rather than reality. Recent research demonstrated that even when optimising for surface finish in current-generation build systems the roughness is still significant, $R_a > 5 \,\mu m$ [329].

Finally, calculations included thus far have assumed 100% area scanning. However, one approach is to use a variable resolution scan, where a coarse mesh is used for the majority of the part, and in areas of key features such as internal cooling channels or where an initial coarse scan has highlighted potential issues. However, putting this into practice requires a high degree of build specific knowledge on locations of likely failure and interest.

In summary, variable scanning is the most promising route to increasing the scan rate as it

does not compromise the acquired data in the key regions of the component, rely upon significant improvements to the hardware (further increasing the capital cost of the system), or greatly increasing the scan complexity (requiring the microstructure to be unravelled from the effects of temperature in the measured SAW velocity). However, to realise this would require a deal of understanding and *a priori* knowledge of the component under inspection that would have to designer. Increasing the scan rate by simply increasing the repetition rate of the generation laser can be easily archived, increased cost being the main impact, and does not require any novel research like variable resolution scanning will. However, realistically an order of magnitude improvement in scan speed is probably the maximum that could be found by only increasing the repetition rate.

7.7.2 Example system

Finally, the following section details progress towards the integration of a SRAS system within an L-PBF build chamber. The chamber in question was acquired as a stripped carcass and thus is not suitable for fabrication testing in future; however, it does serve as a useful proof-of-concept to demonstrate the ability to fit the system within the constraints of the chamber and drive the scan using galvanometers rather than linear scanning stages. This work was carried out in collaboration with Rikesh Patel and Don Pieris. An optical diagram of the system is shown in fig. 7.21, giving access to the x - y plane of the built component.

A basic software package has been developed in step with the system to enable scanning and data processing, using the event-driven programming language C#. The use of non-proprietary control and processing software highlights the simple signal processing required of the SRAS technique. At first glance, multi-build laser systems appear to offer a simple route in, where one build laser could be 'swapped' for a SRAS system, given the all-optical nature.

7.8 Summary

This chapter has dealt with the case for SRAS as an in-situ inspection tool within the laser powder bed manufacture. As of interest to Renishaw, the primary focus has been the development of a cost-time model capable of evaluating the performance and evaluating the total embodied cost of additive components when a SRAS system is integrated within the build system - with comparisons drawn against the current approach. Necessarily, the specific results presented in this chapter have been limited to a few example inputs to demonstrate the cost-time model. Of course, having developed this framework, it is trivial to calculate further specific cases, and it is hoped this model will prove useful in future.



Figure 7.21: Build chamber SRAS system, green and red lines are used to indicate corresponding laser paths. The area marked build chamber is the area of the system where additive manufacturing would occur within a working system.

system

The scanning time of SRAS has often been cited as the biggest stumbling block to the integration within an AM build system. Indeed, the cost-time model developed in this chapter agrees that the speed of scanning is the limiting factor in the case of full utilisation of the build system. The total output over a year would be reduced from 127 to 75 components per annum in this scenario. Based on the assumption that the short fall would need to be made up by additional build systems, the additional cost was calculated to be on the order of \pounds 700kg⁻¹ - making SRAS inspection based around full utilisation unviable, based around current scan rates and system costs. However, in the case of underutilisation, the additional build times due to SRAS scanning are absorbed by the idle time, without a reduction in the number of manufactured components. In this scenario, the embodied cost of SRAS inspection was found to be similar to XCT but around \pounds 470kg⁻¹ more expensive than than the no inspection cost scenario - assuming a SRAS system would cost \pounds 100k.

Of course, the fact that the additional cost of the inspection system results in a higher embodied final cost of the component compared to the no inspection scenario is no great surprise. Accessing the SRAS datasets presented in previous chapters may alone be justification for the integration of a SRAS system within a build system. Nevertheless, the power of in-situ inspection comes from the capability to enact early termination of failed builds. When this is considered it was shown for relatively low failure rates ($\gamma \approx 0.15$), thought to be realistic for the sample build based on prior art, then SRAS would become a more cost effective option than the post-build no inspection cost scenario.

Moreover, the potential financial savings of repair have been evaluated for the first time. This methodology was then used to demonstrate that if the repair of components in-situ could be realised, then even low failure rates would (0.15) merit the use of SRAS. This chapter has demonstrated the possibility of repairing L-PBF defects during the build process, adding weight to the argument for in-situ inspection. The methodology of rework shown in this section is highly suitable for integration into an entirely in-situ system, as measurements can be made in-situ using a modified L-PBF laser/optics system, rework scan paths can be computed on the fly, and the rework paths can be applied within the same manufacturing process with a slight increase in overall build time and hence insignificant increase in cost. This section demonstrates the possibility of repairing L-PBF defects during the build process, adding weight to the argument for in-situ inspection. The methodology of rework demonstrated in this section is highly suitable for integration into an entirely in-situ system, as measurements can be made in-situ using a modified L-PBF laser/optics system, rework scan paths can be computed on the fly, and the rework paths can be applied within the same manufacturing process with a non-significant increase in overall build time. Whilst this study has focussed on targeted surface rework, the advent of in-situ inspection would allow for remelting to occur throughout the build process in order to rework defects within the build volume of the component. The results presented in this section show that repair will be a fundamental building block towards realising utility from in-situ measurements. Further studies should focus on understanding the physics of this process and generating a model describing the process.

Evidently, the high capital cost of a laser ultrasound system results in a higher base cost per component in powder bed fusion when SRAS is used for inspection when ignoring the potential for repair and failure detection. Therefore, it is essential to clearly state that in the case of non-complex components, with short build times and correspondingly low failure rates, SRAS (and likely many in-situ tools) are less attractive propositions due to the additional capital cost not being recovered by the need to replacement few failed components. Instead, the inputs used for the model in this chapter have described a complex and long build time component with a concomitant high failure rate and using relatively expensive material (Ti-6Al-4V). When any of the build time, build complexity (taken as a proxy for failure rate) or feedstock cost increase, the case for in-situ inspection becomes stronger. The question of the ultrasonic detector requires clarification if the costs outlined in this chapter are to be made more concrete. Commercially available detectors are technically capable but significantly increase the system cost, costing approximately $\pounds 100k$.

Potential improvements to the assumptions used for the cost-time model are presented. It was noted that whilst improving scan rates could be most simply achieved by using generation lasers with higher repetition rates, targeted scanning utilising variable resolution to focus on areas of interest within the build may be a more promising avenue. This led to a brief discussion on the potential physical implementation of SRAS within a build system to conclude the chapter. The presentation of a demonstration build system with SRAS instrumentation installed is the most notable outcome of this section.

Based on the points noted thus far, there is a reasonably strong case for the integration of SRAS within a build system. Equally, this chapter highlights that much work remains before this is realised. It is not unreasonable to conclude the case for SRAS is at a tipping point. The value offered by rough surface imaging is the crucial point. If this avenue is to be explored further, then a further detailed investigation of the defect detection capability of SRAS is required.

Chapter 8

Future work

8.1 Introduction

This work has investigated the ultrasonic technique spatially resolved acoustic spectroscopy (SRAS) as a tool for materials characterisation in additive manufacturing. The ability to determine the elastic constants and crystallographic orientation simultaneously, by measurement of the surface acoustic wave (SAW) velocity, in novel polycrystalline materials was determined to be a requisite for meaningful characterisation and illustrates the most exciting aspect of this work. However, there remains much work to be done. Quite obviously, the work of chapter 6 needs applying to additive manufacturing materials. More generally, this opens several new potential avenues of research. As such, this chapter outlines sensible developments and investigations to follow this EngD, before the thesis is concluded in the following chapter. Close attention is paid to the impact these developments will facilitate beyond mere improvements to the experiment. The ordering in this chapter is loosely sorted by the ease of implementation.

8.2 Instrumentation development

Progress towards the next generation SRAS instrument is shown in fig .8.1, specific developments and steps towards this are outlined below.

8.2.1 Random access scanning

A typical SRAS velocity map data is captured in a sequential grid in the x and y directions. When determining the elastic constants or orientation map in relatively large grain (100s of μ m) materials, there is a high degree of data redundancy. If only one velocity surface could be captured



Figure 8.1: On-going development of new SRAS instrument, dove prism is shown for controlling the orienation of the generation patch and detection position (§8.2.2).

per grain, the redundant data could be eliminated and the scan time substantially reduced, without a meaningful loss of information. There is a precedent for applying a similar approach in EBSD, where the quickly captured forescatter electron image is used to define gain boundaries and a sparse EBSD dataset captured for each region, providing a $100 \times$ speed up for microstructure mapping [187].

A similar speed up could be expected by applying this technique with SRAS. A single velocity map would be used to highlight grains, the centres of which can then be targeted for velocity surface scanning, illustrated in fig. 8.2, so-called random-access scanning. This time saving can offset the use of temporal averaging, rather than single-shot acquisition, in each waveform to ensure high signal-to-noise ratio measurement.

The EMDA instrument is designed to acquire velocity maps rapidly but is slow to capture velocity surfaces due to the ad-hoc rotation of the generation patch and detection beam. Therefore, an automated method of controlling the propagation direction is now required to realise this potential time-saving.

8.2.2 Multi-angle scanning

It should be clear to the reader by this stage that the ability to propagate SAWs in multiple directions is intrinsic to the function of SRAS. Whilst the generation patch may be automatically rotated, both the placement of the detection spot and the rotation of the KED are done manually (see §3.4). Subsequently, the operator must adjust the light balance on the knife-edge detector prior to each scan. Clearly, this is not an appropriate solution for acquiring velocity surfaces with fine angular spacing (i.e. 1 degree between SRAS measurements) or in a production system. Practical methods of controlling the propagation direction is an active area of research within the Optics and Photonics research group at Nottingham. Based on this work, two possible approaches


Figure 8.2: Schematic of a polycrystalline specimen, showing one SAW velocity surface spectrum captured per grain. The grain contrast highlighted by a SRAS velocity map can be used to target grains for velocity surface scanning, significantly reducing the number of pixels to be scanned per direction, in turn reducing acquisition time.

are discussed below.

Single-pixel

Presently, there this a growing interest in the field of compressed sensing. For instance, the development of a camera that could produce 2-D images using only a single photodiode rather than the megapixels found now in an everyday camera.

It was hypothesised that a coded grating pattern could be used for SAW excitation in order to reduce data acquisition time and improve the signal to noise ratio. There were a few issues with this approach; the rapid modulation of the phases in the generation pattern necessitates the use of a spatial light modulator (SLM), adding a not insignificant further cost to the optical system. More critically, it was found through experimental testing that this approach has very little tolerance to a displacement of the detection point - the detection beam must be located at precisely the centre of the patch, else the correct phase modulations are not known [256]. Further to this point, the rough surface found in AM specimens would cause the path length from generation point to detection to vary randomly, corrupting the control of the phase of each segment.

Dove prism

A dove prism is a truncated right-angle prism, typically used to invert or rotate an image without dispersion, fig 8.3. By using this optical element, the generation patch and detection spot can be

rotated in unison on the sample, whilst the returned detection beam stays static on the detector thanks to the double pass through the prism. This allows a single optical element to be rotated to achieve multi-direction SAW propagation.

A simple experiment has been used to test a dove prism in a simple low optical power system that mimics SRAS. Fig. 8.3 shows the image formed on a camera at the sample position during the rotation of the dove prism in a SRAS system allowing the image form at the specimen plane to be observed on a camera.

A problem with dove prisms is the resulting substantial astigmatism from the refraction at the entrance and exit surfaces when used in convergent light. Thus, images passed through the dove prism must usually be focussed to infinity. Thus, the prism must be placed between the dichroic mirror and the objective lens - at this point the detection beam and the zero-order generation image are focussed to infinity. When working with smooth surface specimens, the reflected beam will be refocussed to infinity by the objective lens, meaning the second pass through the prism returns the detection beam to the original position.

Having successfully tested the performance of the prism in the LED system, the next step is to integrate it into a SRAS experimental system that is capable of capturing data. The addition of a high power q-switched laser in place of the LED poses a further problem. In the LED set-up, the Fourier image of the grating is formed at the midpoint of the dove prism; however, given the high pulse energy of the generation laser, forming the Fourier image here is likely to cause terminal damage to the prism.

This approach appears viable in any future additive manufacturing SRAS system, with a caveat. The detection spot moves in and out of focus with the undulations of the rough surface found in additive components. The result is the second pass of the dove prism introduces astigmatism - the repercussions of this on SAW detection efficiency remains untested and requires further investigation.

8.2.3 Spatial resolution

Currently, the primary impediment to applying techniques developed in Chapter 6 to additive materials is the system's spatial resolution. With this in mind, the new system will operate with a minimum acoustic wavelength of 6 μ m, assuming an eight fringe patch this will improve the standard spatial resolution from ~ 100 μ m to ~ 24 μ m. For a typical metal of SAW velocity 3000 ms⁻¹, this corresponds to a working acoustic frequency of 500 MHz.

Practically, when attempting to make measurements outside of the current operating range of EMDA, several factors come into play. For example, the generation laser in the EMDA system



(a) Schematic of a dove prism, showing the face cut from a rightangle prism. The collimated image is rotated through total internal reflection along the long face of the prism. The image experiences double rotation, thus the prism need only be rotated through 90° to rotate the generation image through 180° .



(b) Generation mask and detection spot imaged through dove prism at 0° position.

(c) Generation mask and detection spot imaged through dove prism at 120° position.

Figure 8.3: The SRAS generation mask and detection spot may be imaged through a dove prism, allowing control of the propagation direction of generated SAWs, without the need to realign other optical components.

has a pulse width of 1.6 ns, at the typical operating frequency of near 125 MHz there is ample energy in the laser pulse. However, moving to 500 MHz, there is negligible content - the pulses are sufficiently long that SAWs will be destroyed. Therefore, a new generation laser with a pulse width of 800 ps has been selected for this system. The spot size of the detection beam must scale with the acoustic wavelength, with diameter $\approx \frac{\lambda_g}{2}$ [56]. In this case, a focussed detection spot diameter of $\sim 3\mu$ m would be required, which can be achieved by the use of the new 50× objective lens. The corresponding grating spacing on the chrome mask is then 150µm, assuming a f100 mm tube lens is again used. Finally, 500 MHz also exceeds the upper-frequency limit of the current knife-edge detector design, although the actual Hamamatsu silicon photodiodes have a cut off-frequency of 500 MHz and would therefore be worth testing before embarking on the arduous journey to source new high-frequency photodiodes. Instead, the current limiting factor is the RF transformer fitted to the printed circuit board of the detector, which has a working range of 50 – 300 MHz.

8.2.4 Velocity resolution

As detailed in Chapter 6, accurate measurement of the SAW velocity is the cardinal factor in determining the elastic constants and crystallographic orientation to a meaningful accuracy. To realise sub-GPa errors in elastic constants, the standard deviation of the measured velocity needs to be less than 1 ms^{-1} . This can be improved using averaged temporally averaged waveforms and using a greater number of fringes where possible. A greater number of generation fringes can be contained in a fixed size generation patch by working at a higher acoustic frequency.



Figure 8.4: SAW velocity surfaces in tungsten. As tungsten is highly isotropic there is little varation in the velocity surfaces.

The refractory and highly isotropic material tungsten has potential uses in future nuclear power

generation. With this in mind, tungsten structures are now regularly being fabricated by wire-arc additive manufacturing [330]. Acoustic characterisation of tungsten is extremely challenging due to its isotropy; velocity surfaces of tungsten (first shown in §4.2.3) are shown in fig. 8.4, note the velocity range of 5 ms^{-1}). Therefore, determining the crystallographic orientation in tungsten will prove a valuable benchmark for the velocity resolution of the new system. Furthermore, it shall be worthwhile to consider the difference in elastic response when treating tungsten as an isotropic material, defined by E and μ , compared to an anisotropic material, defined by C_{ij}.

8.2.5 Temperature control

Continuing the theme of accurate measurement of the velocity. As discussed in chapter 6, the temperature of the specimen has a tangible impact on the acoustic wave velocity. As a demonstration, the SAW velocity on the (001) plane of nickel at temperatures between 100 and 760 K is shown in fig. 8.5a. The temperature range shown displays a near-linear trend, which can be approximated as a rate of change of the surface wave velocity of $-0.4 \text{ ms}^{-1} \text{ K}^{-1}$, fig. 8.5b. From chapter 6, we have seen a measurement error better than 1 ms^{-1} is needed for determination. Thus, a temperature instability of a few degrees can become the limiting factor in the accuracy.

Experimental results presented in this work have been at nominal room temperature but was not controlled and probably spans a range of between 5 - 10°C. Therefore, to measure the velocity with the best possible resolution, the temperature of the experiment needs to be carefully controlled in future. The work of Optics and Photonics group alumnus Robert Ellwood was concerned with non-linear measurements for fatigue, where the requisite velocity sensitivity is on the order of mms⁻¹. In Dr Ellwood's work, a Peltier heat pump controlled chamber was shown to achieve long term temperature stability of 1 - 2 K. A similar approach can be implemented in the new SRAS system to ensure a stable environment. Specialised microscopy linear stages are then available for more precise temperature control; however, these limit the size of specimens that can be interrogated and will only be used when absolutely necessary, therefore.

8.3 High temperature measurements

Having concluded a description of the new SRAS instrument, the remainder of this chapter will explore requisite ancillary developments and areas of interest meriting further investigation.

As discussed in the previous section, whilst it is desirable to isolate the effects of temperature from a measurement, the idea of manipulating the specimen temperature also opens the door to measuring the elastic constants and microstructure at elevated temperatures. Bear in mind, many challenging applications, such as gas turbine engines, require materials that present con-



Figure 8.5: (a) SAW velocity on the {001} plane of nickel, as a function of temperature. (b) Over this range there is a linear decrease in velocity with increasing temperature.

sistent mechanical properties across a wide temperature range - making temperature dependent measurements a vital research field [135]. In fig. 8.5, the linear temperature dependence of the SAW velocity in pure nickel is shown. Mapping the temperature-dependent elastic constants for the next generation of nickel superalloys is an active area of research [331].

A more interesting case is the measurement of the elastic constants of titanium. Prior to melting at 1,940 K, the crystal structure of titanium transforms from the hexagonal α -phase to the cubic β -phase at 1160 K, as discussed in Chapter 6. From Burger's orientation relationship, the (0001)_{α} plane transforms to (101)_{β} plane. Fig 8.6 shows the SAW velocity on the (0001)_{α} plane prior to the phase transformation and on the (101)_{β} plane at 1273 K.

Measurements of the β -phase elastic constants have proved difficult for established methods due to transformation-induced twinning, causing a polycrystalline to form [299]. Rather, ultrasonic measurements assuming the anisotropy of the crystal and using approximate relationships to engineering moduli have been used [332]. SRAS would be well suited to making such measurements, allowing the elastic constants through the phase transition to be tracked for the first time, and this is not limited to titanium; many materials have temperature allotropes. Clearly, achieving



Figure 8.6: SAW velocity surfaces, as a function of temperature, in titanium. According to Burger's orientation relationship, inset schematic, at 1160°C the $(0001)_{\alpha}$ plane transforms to $(101)_{\beta}$ plane. Elastic constants values from [299], for hexagonal structure, and [332], for cubic structure.

such high temperatures would require apparatus beyond temperature-controlled linear stages but would provide a meaningful application of SRAS going forward.

8.4 Verification of elastic constant measurements

Elastic constants determined in this work have been indirectly validated by crystallographic orientation comparison and literature values. For validation, direct measurement of the elastic constants by alternative techniques, such as resonant ultrasound spectroscopy, is an essential next step. The method outlined in this work benefits, particularly in hexagonal materials, from using multiple unique velocity surfaces in the inversion. However, most elastic constant determination techniques require single crystal materials, making the comparison of experimental measurements on the same specimen challenging. Large grain polycrystalline materials may prove viable, where individual grains can be isolated and extracted. Otherwise, a single crystal specimen can be cut in such a way that multiple planes are exposed, as shown in fig. 8.7, allowing characterisation by volumetric approaches and SRAS.



Figure 8.7: A schemutic of a cut single crystal in titaniium, showing the three unique velocity surfaces which could be exposed. This would facilitate comparison of elastic constants determined by SRAS, resonant ultrasound spectroscopy and through transmission measurements.

8.5 Additional information for inversion

8.5.1 Multiple wave modes

Currently, the inversion of the velocity surface to calculate the elastic constants and/or orientation uses only the 'dominant' mode, as discriminated by the forward model. However, modes are regularly measured in one propagation direction - as an example, fig. 8.8 shows two planes in nickel where two surface wave modes can be seen.

Including the velocity of this second mode in the inversion offers a distinct advantage for the determination of elastic constants. In Chapter 6, we have seen the primary challenge in the accurate determination of the elastic constants is resolving the subtle changes along the C_{11} - C_{12} diagonal. Fig. 8.9 shows the shape of the solution space in the plane C_{11} - C_{12} for the dominant and secondary mode, as previously the minima contour is located along the diagonal; however, the gradient of these contours is different - making the true solution more unique.

The relationship between C_{11} and C_{12} for the two modes is seen more easily in fig. 8.10, where the gradient of the minima contours is extracted. Using the velocity of this secondary mode for inversion makes the solution more deterministic, making the process of determining the values of C_{11} and C_{12} more robust.



Figure 8.8: Velocity surfaces from single crystal nickel, the dominant mode is overlaid in black. In both planes an additional mode is seen measured over certain angles, the secondary mode, from the forward model, is overlaid in red in these region. This information can be used for the inversion to make increase the uniqueness of each plane-elastic constant set, making the problem more deterministic.

8.5.2 Amplitude

As shown in §4.3.1, the forward model data used for the inversion is a binary matrix, effectively discarding any amplitude information. However, the amplitude information offers the potential of an additional fitting parameter, in which case the forward model matrix can be construed by equation 8.1.

$$I_{\rm hkl}(v,\phi_1) = \begin{cases} 0, & v \neq v_c(h,k,l,\phi_1) \\ \overrightarrow{d}, & v = v_c(h,k,l,\phi_1) \end{cases}$$

$$\tag{8.1}$$

where \overrightarrow{d} is the displacement vector of the *i*-th wave mode, as calculated by the forward model.

Of course, to realise this, the displacement calculated by the forward model must reflect experimental reality. Here we encounter the problem of the ad-hoc method by which the amplitudes are currently derived, which only offers relative rather than absolute amplitude. This is compounded by the thermoelastic generation mechanism not being well captured.



Figure 8.9: Shape of solution space in the plane C_{11} - C_{12} for (a) dominant surface wave mode and (b) secondary surface wave mode.



Figure 8.10: Diagonal of the minima in C_{11} - C_{12} for the two wave modes, extracted from fig. 8.9. The differing gradients make the solution more deterministic.

There is a wealth of literature concerning the calculation of acoustic amplitudes from thermoelastic laser generation. Primarily this deals with the similar case of isotropic media, however generalised analytical approaches in anisotropic materials have proved elusive. The majority of the techniques developed for modelling nano-second pulse width sources have substituted appropriate boundary stress conditions for the laser source, creating a shear stress dipole. However, contrary to the isotropic case, for an arbitrary sample orientation, the equivalent surface normal stress was non-zero, and the equivalent surface shear stress was not a simple dipole.

Numerical modelling is a more appropriate tool to tackle this problem in the anisotropic case. The work of Taheri et al. provides a solid foundation in this domain, here the authors used a multiphysics finite element model to simulate the generation process of a thermoelastic line source [333]. The multiphysics element stems from using thermal equations to calculate the expansion of the illuminated region, then stress-stress relations to find the resultant particle displacements representing the acoustic wave.

8.6 Spatially resolved properties

In this thesis, it has been assumed that elastic constants are a global property of the specimen to which all measurements should conform. Fig .8.11 shows an extreme example that demonstrates this is not always true. The additively manufactured titanium specimen has an increasing chromium content in the y-axis. After manufacture, a spatially varying heat treatment is applied with a thermal gradient along the x-axis [334]. The varying chromium content causes the formation of the phases (from top to bottom) $\alpha \rightarrow \alpha + \beta \rightarrow \beta$, the final section also contains significant TiCr₂ [335].

Mapping spatial heterogeneities in additive materials, resulting from either deliberate functional grading as shown here or as an unintentional consequence of the build process, is highly desirable. In such specimens, the elastic constants are no longer a global property of the material. A first estimation of this effect could be obtained by mapping the deviation from the globally determined value across a specimen - giving a qualitative indication. It would be particularly interesting to observe spatial coherence in this variance across materials expected to present a consistent elastic response.

The desire to quantitatively measure these elastic constants presents two problems. Firstly, the algorithm proposed generates only a single solution for the material - assuming a greater sensitivity to the global solution can be found by using the measurements of multiple unique planes. It is possible to use only one velocity surface for the inversion; however, depending on the crystallographic orientation, reduced or even no sensitivity is given to certain constants - for example, the supersonic region of hexagonal materials is insensitive to C_{12} . More work is required to relate the number and orientation of velocity surfaces required for the inversion and the ensuing errors in determination.



Figure 8.11: Additively manufactured titanium sample with varying chromium content in the y-axis, after manufacture a spatially varying heat treatment is applied, causing grain refinement along the x axis.

Secondly, the symmetry and approximate chemical composition are unknown for a given measurement pixel in such materials. The result is a vast solution space (in this example, this would require searching both cubic and hexagonal symmetries across the range of titanium to chromium elastic constants), which, despite the computation improvements given in Chapter 4, is intractable. The inverse solver must be adapted to pre-filter the measured velocity surfaces to address this possible approaches to this are outlined below.

8.7 Hybrid solvers

The determination of elastic constants and crystallographic orientation has relied on a brute-force style solver, essentially evaluating the objective function for every point in the solution space. The brute force approach has been made viable by performance improvements in the forward model and is a good strategy for dealing with local minima in the values of C_{11} and C_{12} . However, when working with more challenging materials where the symmetry may not be known at a given location, such as dual-phase titanium, we can see how the possible solution space can explode, again rendering the problem intractable.

8.7.1 Artificial neural network classifier

Initially developed as a mathematical model to mimic the function of the human brain, artificial neural networks have found uses in pattern recognition and classification, amongst others. Pertinently, neural networks have seen a great deal of interest in ultrasonics [336, 337], including inverse problems in materials characterisation [338].

The low-hanging fruit would be to simply classify the crystal symmetry of an input velocity surface spectrum. In practice, this would primarily be differentiating between hexagonal and cubic symmetries. However, the neural network could also be trained to indicate the likelihood of a range of hexagonal materials (an example of a confidence index like this is shown in fig. 8.12), again reducing the possible search space. To give some context, the range of C_{11} captured by the matrix below is 60 – 720 GPa, a solution space too large to solve directly, identifying the measured velocity surface as a form of α -titanium (lower right square) reduces the range to 160 – 165 GPa. Neural networks require large databases of classified images for training and validation, to the extent that it would be impossible to fully sample with experimental data, but the simulation method described in §4.4 can be used to generate large training datasets quickly.

Du et al. have taken this idea to its logical conclusion and used a convolution neural network to predict the elastic constants from the measured velocity surfaces directly. The authors remarked their analytic optimisation approach agreed well with literature values before going on to build an artificial neural network to allow instantaneous prediction of the elastic constants from SAW velocities. The authors reported that the network showed promising results but did not provide numerical qualification of this [225]. In any case, this is a worthwhile theme of future investigation. If a robust and accurate neural network is viable, it will offer near-instantaneous calculation of the material properties.

8.7.2 Optimisation algorithms

Consider the case of cubic materials, whilst finding the true values of C_{11} and C_{12} are challenging, we have seen the C_{Δ} and C_{44} solution space is convex and relatively easy to determine. Cherry et al. demonstrated the use of Nelder–Mead optimisation to determine the elastic constants from simulated SAW velocity measurements, whilst the values of C_{11} and C_{12} were underestimated by ~ 8 GPa, the values of C_{Δ} and C_{44} were found to within 1 GPa. Using an optimisation search could give a good indication of the robust parameters C_{Δ} and C_{44} in short order, reducing the space to tediously brute-force.



Figure 8.12: Example of similarity index between different hexagonal materials, an artificial neural network can be trained to produce similar maps for measured velocity surfaces, which can then be used as a guide for narrowing the search space for determination.

8.8 Textured materials

One of the challenges additive manufacturing presents, powder bed fusion, in particular, is the small grain size, meaning individual grains are not fully resolved by current SRAS instrumentation. Measuring velocity surfaces which are some function of the multiple grains under the generation patch, prevents the direct calculation of the crystallographic orientation. As discussed in \$8.2.3, the intention is to improve the spatial resolution to around $20 - 30 \,\mu\text{m}$ - this will allow the grains to be fully resolved in a much more comprehensive range of materials.

Despite these resolution improvements, many materials, such as Ti-6Al-4V, present an ultra-fine structure that makes resolving the grains a significant challenge even for electron microscopy - one which incremental improvements to the SRAS instrumentation cannot tackle. However, it is the presence (or lack-there-of) of microtextured regions, such as those found in WAAM (§5.1), which is of greater importance than the specific orientation of any single grain. These heterogeneities are thought to be the governing factor in the premature failure of these alloys [339]. Measuring the feature size [339, 340], crystallographic texture [204], elastic constants [246] and predicting mechanical performance in these alloys are all important contemporary problems in ultrasonics [341]. From an industrial viewpoint, these alloys are commonly found in aerospace and play an important role in dwell fatigue failure [342]. A growing understanding of the importance of these regions, combined with the catastrophic consequences of rotor failure, has necessitated forthcoming regulations to specify limits on the presence of these textured regions [343].

For simplicity when framing the problem, we can consider three domains of operation, as shown in fig. 8.13. Domain ' \mathbf{A} ', Where grains are well resolved, is the primary use case of SRAS, with crystallographic orientation and elastic constants able to be determined (as demonstrated in Chapter 6). 'B', the grains are now micro-regions of a dominant texture, with a notable concentration of divergent orientations within each region. Finally, in \mathbf{C} , there is little to no spatial structure, with each pixel exhibiting a random orientation compared to its neighbour. To some extent, the work of Chapter 5 is concerned with elucidating application-specific material information where the grains are not resolved. The heuristic is that the crystallographic orientation cannot be returned in specimens where the grains are not well resolved.



increasing disorder



high texture



low texture

Figure 8.13: As the level of texture decreases there is a greater variance of orientations within a region through until the disorder is at a maximum each point displays a random orientation. This assumes a constant spatial resolution. SRAS is typically used with ' \mathbf{A} ' type specimens to measure the orientation. Understanding the breadth of applications of SRAS measurements to 'B' type materials is an important area of future research. Finally, the longer length-scale ordering breaks down in 'C' type materials, with each grain below the spatial resolution of the system.

Consider then the results of fig. 8.14, which compares the *c*-axis orientation map as captured

by SRAS (fig. 8.14a) and EBSD (fig. 8.14b), in textured Ti-6Al-4V. The EBSD image is split in half, the left-hand section plots the raw data as experimentally measured before spatially averaging is applied in the right-hand image. Averaging highlights micro-textured regions and reduces the spatial resolution to a similar order as SRAS to compare the datasets. The original EBSD dataset looks to suffer from significant noise. On closer inspection, this is due to small areas of an orientation distinct from the region in which it lies (a glance at the EBSD results of similar alloys in published literature show the same 'noisy' structure [339, 340]). Of course, these fine regions ($\leq 20\mu$ m) are not resolved by the SRAS measurement.



(a) SRAS *c*-axis orientation map.

(b) Hybrid EBSD *c*-axis orientation map.

Figure 8.14: Comparison of SRAS and EBSD *c*-axis maps in textured Ti-6Al-4V. In the EBSD dataset the left-hand side plots the originally sampled data, the right-hand side is spatially averaged, allowing easier observation of MTRs. The maps indicate the position of the hexagonal *c*-axis relative to the observer, the plane, Φ , is mapped by saturation and the rotation, ϕ_1 , by hue.

The c-axis maps show relatively good agreement between the measurement techniques, and this is to some extent an unexpected result given the fine structure contained within the aforementioned micro-textured regions. The ϕ_1 angle is in particularly good agreement, with an error. Still, a larger discrepancy exists between the Φ angle measurements of the two techniques - preliminary results suggest this is an effect of the incorrect elastic constants being used to calculate the forward model. This satisfactory result perhaps belies the complexity of the problem and poses the question to what level of orientation disorder can single-crystal information be determined? Can the misorientated grains be effectively treated as a noise source in the case of highly ordered grains such as this specimen? If so, can the texture strength be directly related to the accuracy of orientation determination (in the vain of results presented in the first half of Chapter 6)? Addressing these questions will require further investigations of the generation and propagation process in textured materials.

Furthermore, as the order breakdowns, there remains the question of how much information can be deduced from the SRAS measurements without the *a priori* process-specific knowledge which has been leveraged in chapter 5. Extracting the texture coefficients and asymmetric elastic constants, as others have done using bulk waves and resonance, would present a leap forward in the quantitative information that can be deduced from these difficult materials.

Again, it is important to remember that whilst the study of textured materials poses interesting problems from an acoustics perspective, there are implications for the characterisation of industrially relevant materials, such as the titanium alloys found in aerospace [340]. Furthermore, there is a great deal of synergy with additive manufacturing problems. Indeed, it is this increased disorder that is measured in additive high silicon steel specimens, §2.2.3, to describe the alignment of the easy axes of magnetisation.

8.9 Closed loop-feed back for microstructure control in additive manufacturing

To conclude, it is worth a final contemplation of the outlook for SRAS within additive manufacturing. As shown in §, the integration of a SRAS system to capture measurements during the wire-arc build process appears a relatively achievable prospect within the near future.

Finding a sensible deployment of the technique which is robust to the industrial nature of the build environment appears to be the greatest challenge. Autofocus, vibration and electromagnetic interference from the build process, and robust 'packaging' of the optical system all need to be addressed to enable the as-deposited images captured in the lab (see §5.1.8) to be replicated on the shop floor. This is not to overstate the problem, as all of these challenges have been tackled in other topics and should provide a solid basis to realise an industrial instrument for WAAM.

Powder bed fusion presents a more challenging prospect. Presently, work is ongoing to integrate a SRAS system within a custom L-PBF system based at the University of Nottingham. The ability to take SRAS measurements during the build in-effect overcomes the limitation of surface wave imaging by allowing a volumetric dataset of the microstructure to be acquired. Mapping the formation of the microstructure during the build process should allow processing parameters to be dynamically controlled to deliver bespoke microstructures. There are few immediate examples of such closed-loop control being enacted. The work of Zhong et al., who recently reported on a closedloop model which uses the thermal history of the most recently deposited layer to compensate for differences in accumulated heat, is one of the few illustrations of this idea [344]. However, Zhong's work was focused on ensuring high-density components, a more straightforward challenge than dynamic control of the microstructure. Indeed, this often prompts comment that such closed-loop control of the microstructure is 'pie in the sky', pointing to the current absence of any credible method to provide this as evidence. The principal reason this field is undeveloped is the full process-microstructure relationships within powder bed fusion are not yet well understood. Whilst schemes have demonstrated the ability to control the grain orientation and morphology (see §2.2.1), these are limited to single material builds of simple geometry. Of course, the issue is to derive relations that guide how user-controlled process variables should be adjusted in repose to varying build conditions. The work of this thesis, combined with the integration of a future SRAS system within a build chamber, will provide a tool ideally suited to addressing this problem. It is hoped that the additive manufacturing community will 'pick up the ball and run with it'.

More specifically, integrating a SRAS system within the build system will allow the microstructure to be observed over a wide area, as it forms on a layer-by-layer basis. Therefore, direct measurements over a wide area will be possible for the first time, rather than relying on indirect observations, such as melt-pool temperature, to infer the likely grains structure. This will allow errors or deviations in the expected microstructure to be corrected by dynamically correcting the build parameters based on the SRAS measurements, providing a platform for progressing this field towards dynamic control of the microstructure in 'real' components. The tools developed in this thesis will be particularly applicable for characterising materials alloyed in-situ and controlling the build process to obtain novel microstructures.

Chapter 9

Conclusions

The work presented in this thesis has concerned the development of the laser ultrasonic technique spatially resolved acoustic spectroscopy (SRAS) as a materials characterisation tool, primarily within additive manufacturing although the outcomes of this work have been shown to have impact in the wider materials science community. Specifically, the work towards the goals of this thesis have focused on the characterisation of texture and microstructure within specimens fabricated by wire-arc manufacture and laser powder bed fusion and the development of a method for elasticity determination.

Through a thorough introduction of the underlying methods and a survey of the existing techniques in chapters 1 and 2, the place of SRAS within the materials characterisation community at large has been comprehensively discussed. Particular emphasis has been given to established techniques for imaging the microstructure and determination of the elasticity matrix and the current applications of non-destructive evaluation within AM.

The instrument used throughout this thesis, known as the EMDA system, has been presented in detail in Chapter 3. The system uses a broadband Q-switched laser and chrome-coated optical mask to create the generation fringes, which subsequently generate ultrasound through thermoelastic absorption. This allows the frequency of the acoustic wave packet to be encoded, importantly this is a function of only the properties under the generation patch. The surface perturbation from the acoustic wave is detected by the angular deviation of a second laser on a knife-edge detector, thus allowing the surface wave velocity to be measured. The results presented in this thesis have generally used 8 generation fringes at an acoustic wavelength of 24 μ m, resulting in a spatial resolution of ~ 100 μ m. Two SAW velocity maps were presented as examples of SRAS measurements in large and fine-grained materials.

The two available methods of capturing velocity surfaces by rotating either the optical mask

or the specimen are presented, and the subsequent difficulty in regularly capturing high-resolution velocity surfaces was noted. Rotating the specimen requires meticulous alignment of the sample to the rotation centre, whilst the more commonly used optical mask rotation leads to variances on the order of 0.1 μ m in the acoustic wavelength and displaces the generation image by ~ 10 μ m due to the centre of rotation being misaligned, inducing both velocity and registration errors between propagation directions.

The numerical methods used in this work have been presented in detail in chapter 4. A great deal of the chapter is given over to the explanation of the forward model used to calculate the SAW velocity as a function of material properties and propagation direction. Previously, the full power of the forward model could not be realised due to the protracted computation times. Subsequent improvements detailed in this chapter have allowed the solution space to be visualised - an essential step in the development of a method for the determination of the elastic constants. As far as we are aware, this is the most efficient published method for calculation of the velocities and displacements of SAW modes. Reducing the forward model computation time by a factor of up to 100 (from a week to an hour), has been a essential step in realising the subsequent outcomes of this thesis. Most obviously, the determination of elastic constants has previously not been possible because the was too slow. Beyond this obvious example, the forward model is leveraged throughout the thesis. Whilst not explicitly stated, one must also consider the impact this has had on the experimental procedure. Previously, scanning a sample to capture SRAS velocity maps and velocity surface spectra was fast whilst computing the forward model of the material was slow. Now, the opposite is true; it is feasible to complete several models for a prospective experimental material beforehand, allowing the user to make informed choices on the requisite velocity resolution and bandwidth of the measurement before experimental investigation.

The velocities as calculated by the forward model were used to generate synthetic measurements, which has provided a basis for studying the impact of experimental parameters. The accompanying inverse model, which calculated the overlap between the experimental data and the range of possible velocity surfaces from the forward model, has been presented. This algorithm allows the calculation of both crystallographic orientation and elastic constants; examples have been provided to this effect.

SAW velocity measurements of several additive manufacturing specimens have been presented in Chapter 5. Materials studied have included high silicon steel, Ti-6Al-4V and nickel. These materials were selected for their industrial ubiquity and fascinating microstructural characteristics, making a worthwhile proving ground for the measurement technique. Wire-arc results have considered the dual-phase material, Ti-6Al-4V, which presents an intruding set of challenges for acoustic measurements. The transformation from the cubic β -phase to the hexagonal α phase as the material cools to room temperature presents a potential problem. It is the size of the β structure that must be controlled. Again, the forward model was utilised to explain the source of this acoustic contrast. The SRAS velocity maps proved to be sufficient to determine the refinement of this β structure and early as-deposited results provide hope that an in-situ tool could be realised in short order. This will be an essential step in the industrialisation of WAAM.

Results in laser powder bed fusion demonstrated the capacity to detect surface-breaking defects using the optical datasets complementary to SRAS measurements. In line with prior literature, subsurface defects have been observed in titanium specimens through mode-conversion of the Rayleigh wave to a Lamb-type wave. However, this effect was not observed in nickel-based specimens, despite the specimens containing many surface-breaking defects. As nickel is highly anisotropic, the greater range of SAW velocities (compared to titanium) was found to potentially mask the presence of subsurface defects. That is to say, the variations due to crystalline texture are greater than those generated by mode conversion. Therefore, it was concluded this was a potentially unreliable mechanism for the detection of subsurface defects.

A further detailed study of the fascinating material high silicon (6.9%) steel, was presented. The formation of a cube crystalline texture over the more typical fibre texture was observed through SRAS measurements by comparing measured 'averaged' velocity surfaces to the forward model. An anomalous microstructure was also observed in the border of all studied components, which corresponded to a different scan strategy being used in this area when fabricating. The important issue in this is that despite the exact crystallographic orientation of each grain not being calculated (due to the small grain size), the more important crystallographic texture could still be inferred from the measured SAW velocity. This result is significant for two main reasons. Considering the material firstly, study by SRAS has allowed comment on the texture of the full surface for the first time, including the previously unreported border area in this material, demonstrating how the texture forms over a wide scale. Furthermore, this materials was selected due to the unusual texturing, so as to demonstrate the ability of SRAS - through the forward model - to act as a microstructure qualification tool.

As a materials characterisation tool, the power of SRAS has stemmed primarily from the ability to determine the crystallographic orientation. With a view to the use in novel manufacture, the question of how accurate this may be determined was investigated for hexagonal materials in Chapter 6. The study demonstrated the influence of experimental parameters, including the number of fringes in the generation patch and the signal-to-noise ratio of the acoustic measurement. The use of 'typical' experimental settings was shown to lead to a disagreement of $< 7^{\circ}$ between EBSD and SRAS results. However, this could be improved to $< 1^{\circ}$ through diligently capturing high-resolution velocity surface spectra. This is the greatest orientation accuracy which has been

reported by SRAS thus far, and is as good as x-ray diffraction - the primary alternative to EBSD. A surprising outcome was the importance of the material elastic constants on the orientation result, with a disagreement in ϕ of $< 7^{\circ}$ between two sets ubiquitous throughout the literature, despite both being quoted for pure titanium - this demanded further investigation and ideally the development of a to make SRAS more robust to uncertainty in elastic constants.

Chapters 4, 5 and the first half of 6 emphasise the importance of elastic constant knowledge for accessing the richer information contained within SRAS measurements. To address this, the latter part of chapter 6 has proposed a new method for simultaneously determining the elastic constants and crystallographic orientation in polycrystalline specimens. It has long been understood that knowledge of the SAW velocity surface and one of the crystallographic orientation or the elastic constants allows calculation of the unknown. By assuming the material elasticity is a global property, constant across all measurement positions, the full inverse problem has been tackled by combining the inversion results across both the elasticity and orientation spaces from multiple grains. Aside from effectively averaging measurement errors, this approach has been demonstrated to be beneficial as the sensitivity to each elastic constant varies significantly across the orientation space. Therefore, by this approach, the polycrystalline nature of materials, which has previously made them difficult to study by existing techniques, is actually exploited in the pursuit of accurate determination of the elasticity matrix. It has been shown reliable results can be determined with just a few unique grains, a surprising but powerful result. As discussed in Chapter 8, the few grains required to determine the orientation opens the possibility of tracking spatial variations in the elasticity - an exiting prospective given the ease and regularly with which chemically heterogeneous materials are created in additive manufacturing.

Calculated elastic constants in pure materials showed good agreement with literature values (within 1 GPa, on average). The improved agreement with EBSD in the troublesome CMSX-4 specimen, error reduced from $> 14^{\circ}$ to $< 7^{\circ}$, has evidential value for the enhanced robustness in orientation determination garnered by finding the elastic properties of the material at hand. The development of a technique to determine the elasticity matrix in polycrystalline materials without knowledge of or conditions upon their grain orientations is, without doubt, the most exciting outcome of this thesis. Development of techniques for the measurement of elastic constants had stalled prior this work, and the measurement of elasticity in polycrystalline materials remained a real difficulty. The realisation of a practical method to deliver both crystallographic orientation and elasticity through SRAS will become a significant new tool for materials sciences.

Having considered and developed the technical basis for the use of SRAS within additive manufacturing as a tool, chapter 7 presents the case for in-situ inspection within laser powder fusion specifically, with particular focus paid to the monetary and temporal impact of a hypothetical SRAS system. Whilst the scan time of a SRAS system is significant, increasing the build time by 53% in the sample build, this was shown to only be a key factor in scenarios where the utilisation of the build system is exceptionally high. Indeed, for typical utilisation rates, the additional build time due to SRAS scanning would be absorbed by the 'downtime' - the end result being no reduction in the number of manufactured components. This has allowed a minimum scan 'speed' to be determined, which will prove to be a key criterion for shaping future SRAS system development.

The monetary impact was then considered separately. The addition of a SRAS system would increase the cost of a build system by a minimum of 10% and lead to a rise in the embodied cost of $\pounds 470 \text{kg}^{-1}$. This additional cost may be acceptable to realise the microstructure information detailed in the previous two chapters. However, this cost can be further justified by considering the propensity of additive manufacturing builds to fail and using the in-situ inspection to allow early termination of failed builds. To this end, a cost model has been developed in line with existing literature. Furthermore, for the first time, the cost of potential repair of defects was presented and paired with a rudimentary demonstration of repair in additive components. The realisation of even relatively low repair rates shall go a long way to making the presence of SRAS within a build system the cost-effective choice.

Many potential avenues of future research have been opened by the work contained within this thesis. The possibility of integrating a SRAS system within the wire-arc build appears achievable in short order and offers a convenient method to assure quality control of the grain refinement. Yet, it is the measurement of the elastic constants that has the broadest range of applications. For example, it provides the ability to measure the elastic properties in 'real' specimens through thermal processing. This measurement capability comes at an opportune moment for additive manufacture as new alloys and functionally graded materials become an ever greater focus within the field.

To conclude, it is clear from the work presented in this thesis that SRAS is positioned to offer a unique perspective on the microstructure within additive manufacturing materials. The capability to rapidly capture wide-area maps of the microstructure and leverage the forward model to elucidate a wealth of information vital to materials characterisation. Additive manufacturing offers a paradigm shift in capability, including the possibility of creating designer microstructures, as discussed in the early part of this thesis. Despite this potential, AM is not yet used to create 'mission critical' components, but the ability to rapidly image the microstructure along with measuring the crystalline texture and elasticity in these components using SRAS significantly increases the likelihood of harnessing the potential power of additive manufacturing.

Appendix A

Elastic constants in cubic materials

Table B.1 tabulates elastic constants for the cubic materials: pure nickel, the nickel superalloys Inconel 718, silicon, iron, tungsten and aluminium. The two forms of nickel are discussed at length through the course of this thesis. Aluminium has been used for generating illustrative figures, such as dispersion curve plotting, in chapter 1, for simplicity when introducing basic concepts as it can be treated as an isotropic material. Pure nickel and tungsten have been used as illustrative materials for acoustic wave calculations in chapter 4, whilst chapter 6 presents the elastic constant (values calculated in this work are highlighted below) and crystallographic orientation determination in both nickel and Inconel 718. Discussion concerning silicon and iron are presented in chapter 5. Furthermore, as these are all common engineering materials there have been multiple attempts to measure the elastic constants of these materials - highlighting the range of disagreement across prior literature.

| Material | C ₁₁ | C ₁₂ GPa | C ₄₄ | C ₄₄ Notes on errors | |
|-----------|----------------------|------------------------|--|---|---------------|
| Aluminium | 108.2 | 62.3 | 28.4 calculated from compliance | | [275] $[345]$ |
| | 107.3 | 60.08 | 28.3 | > 0.5%, particularly C ₁₂ | [346] |
| Iron | 230.37 ± 0.53 | 137.1 ± 0.24 | 115.87 ± 0.17 | fitted RUS data, errors propagated | [297] |
| | 226 ± 2 | 140 ± 8 | 116 ± 1 | | [347] |
| | 228 ± 2 | 132 ± 4 | 116.5 ± 1 | C_{12} propagated from direct measurement of other constants | [348] |
| Nickel | 246 ± 3.5 | 153 ± 3.5 | 121 + 1 | - | * |
| | 247 | 153 | <u> </u> | | [275] |
| | 251.6 | 154.4 | 122 | > 0.5%, particularly C ₁₂ | [294] |
| | | | | propagated from | |
| | 252.8 ± 2 | 152 ± 3 | 123.8 ± 1 | measurement | [237] |
| | | | | errors | |
| CMSX-4 | 251.5 ± 5 | 163 ± 5 | 119 ± 1 | - | * |
| | 251 | 159 | 132 | - | $[349]^{A}$ |
| | 248 | 155 | 132 | - | $[349]^{B}$ |
| | 258 | 159 | 129 | likely $\geq 2\%$ | [309] |
| | | | | errors are FWHM | |
| | 243 ± 2 | 153 ± 2 | 128 ± 1 | of the minima | [349] |
| | | | | function | |
| Silicon | 168 | 66 | 84 | $\geq 5\%$ | [350] |
| | $165.779 \pm 0.05\%$ | $63.9365 \pm 0.2\%$ | $\begin{array}{c} 79.6346 \pm \\ 0.05\% \end{array}$ | - | [351] |
| | 165.64 | 63.94 | 79.51 | errors expected to be similar to [351] | [352] |
| Tungsten | 512.57 | 205.82 | 152.67 | calculated from compliance in | [353] |

 Table A.1: Experimentally derived elastic constants (GPa) of various cubic materials from existing literature.

*Values determined in this work are shown in **bold**.

^A Reported to the author in personal communication by W. Hermann.

^B Reported to the author in personal communication by V. Alberts.

 $^{\rm C}$ Likely to be misprint and should be 152.67 GPa.

Appendix B

Elastic constants in hexagonal materials

Table B.1 tabulates elastic constants for the hexagonal materials: pure titanium, the titanium alloy Ti-6Al-4V, magnesium and zinc. The two forms of titanium are discussed at length through the course of this thesis; Ti-6Al-4V in chapter 5, whilst results in pure titanium are discussed in chapters 4 and 6, specifically. Magnesium and zinc have been presented as they are included as part of the discussion of Chapter 6, and as they are common engineering materials, there have also been several attempts to measure their elasticity.

| Material | C ₁₁ | C ₁₂ | C ₁₃ | C ₃₃ | C ₄₄ | Notes on errors | Ref |
|-----------|--------------------|-------------------|---|---|-------------------|--|--------|
| Hexagonal | | | | | | | |
| Magnesium | $59.4\pm0.7\%$ | $25.61 \pm 1.8\%$ | $\begin{array}{c} 21.44 \pm \\ 1.8\% \end{array}$ | $61.6\pm0.7\%$ | $16.4 \pm 0.7\%$ | propagated from measurement errors | [302] |
| Titanium | 161 ± 0.6 | 90.5 ± 0.5 | 66 ± 1.8 | 174 ± 3.6 | 47 ± 0.3 | - | * |
| | $154.0 \pm 0.5\%$ | 86.0 | 67.3 | $183.0 \pm 0.5\%$ | $46.7\pm0.5\%$ | larger errors for $C_{12} \& C_{13}$ standard | [236] |
| | 160 ± 5 | 90 ± 4 | 66 ± 3 | 181 ± 2 | 46.5 ± 0.4 | deviation from | [226] |
| | 163.6 | 92.3 | 67.92 | 185.2 | 47.05 | three repeats in tabulated from graph in [354] | [299] |
| | $162.4 \pm 0.2\%$ | 92 | 69.7 ± 0.8 | $\begin{array}{c} 180.7 \pm \\ 0.2\% \end{array}$ | $46.7\pm0.2\%$ | $C_{12} \geq 0.2\%$ | [300] |
| | | | | | | fitted to | |
| Ti-6Al-4V | 136.0 | 78 | 68.0 | 163 | 40 | experimental | [355] |
| | 141 + 907 | | | 100 1 007 | | mechanical testing | [0.40] |
| | $141 \pm 3\%$ | $70.9 \pm 3\%$ | $57 \pm 3\%$ | $103 \pm 3\%$ | $48.70 \pm 3\%$ | - | [240] |
| | 170.0 | 92.0 | 70.0 | 192.0 | 52.0 | diffraction elastic moduli | [287] |
| Zinc | 177.6 | 24.0 | 60.2 | 74.0 | 39.8 | Calculated from compliance in [301] Calculated from | [356] |
| | 163.5 | 17.1 | 41.5 | 55.3 | 39.8 | compliance in [301] | [356] |
| | $163.68 \pm 0.5\%$ | $36.4\pm2\%$ | $53.0\pm5\%$ | $63.47 \pm 0.5\%$ | $38.79 \pm 0.5\%$ | - | [357] |

Table B.1: Experimentally derived elastic constants (GPa) of various hexagonal materials from existing literature.

*Values determined in this work are shown in ${\bf bold}.$

Appendix C

Generalised anisotropic elastic modulus relations

C.1 Cubic materials

General forms of engineering elastic moduli in cubic materials are given as a function of direction cosines l_i below

$$\frac{1}{E} = S_{11} - 2 \left(S_{11} - S_{12} - \frac{1}{2} S_{44} \right) \left(l_1^2 l_2^2 + l_1^2 l_3^2 + l_2^2 l_3^2 \right)
\frac{1}{K} = \frac{1}{3} \left(S_{11} + 2 S_{12} \right)
\frac{1}{G} = S_{44} + 4 \left(S_{11} - S_{12} - \frac{1}{2} S_{44} \right) \left(l_1^2 l_2^2 + l_1^2 l_3^2 + l_2^2 l_3^2 \right)
v_{lm} = -\frac{S_{12} + \left(S_{11} - S_{12} - \frac{1}{2} S_{44} \right) \left(l_1^2 m_1^2 + l_2^2 m_2^2 + l_3^2 m_3^2 \right)}{S_{11} - \left(S_{11} - S_{12} - \frac{1}{2} S_{44} \right) \left(l_1^2 l_2^2 + l_1^2 l_3^2 + l_2^2 l_3^2 \right)}$$
(C.1)

For completeness, the bulk modulus K is given but this term is isotropic, hence no direction cosines appear in its formula. In the case of Poisson's ratio, the vector m_i perpendicular to l_i is also required to define the polarisation.

C.2 Hexagonal materials

Where θ is the angle between the c-axis and the direction of loading.

$$\frac{1}{E} = \cos^{2}\theta \left(S_{11} + S_{33} - 2S_{13} + S_{44}\right) + \cos\theta \left(2S_{13} + S_{44} - 2S_{11}\right) + S_{11}
K = S_{11} + S_{12} + S_{13} - \cos\theta \left(S_{11} - S_{33} + S_{12} - S_{13}\right)
\frac{1}{G} = \S_{44} + \left(S_{11} - S_{12} - \frac{1}{2}S_{44}\right) \left(1 - \cos^{2}\theta\right) + 2\cos^{2}\theta \left(S_{11} + S_{33} - 2S_{13} + S_{44}\right) \left(1 - \cos^{2}\theta\right)
v_{=} - \frac{\cos^{2}\theta S_{13} + (1 - \cos^{2}\theta) S_{12}}{\cos^{2}\theta \left(S_{11} + S_{33} - 2S_{13} + S_{44}\right) + \cos\theta \left(2S_{13} + S_{44} - 2S_{11}\right) + S_{11}}$$
(C.2)

In this instance we see the bulk modulus, K, has a direction term as the crystal is not isometric (as it is in the cubic case).

Appendix D

Generalised anisotropic bulk wave velocity and elasticity relations

The simple relations between the bulk wave acoustic velocities and the elastic constants in high symmetry directions have been given in §2.4.1. It is possible to generalise these forms for any plane and direction, however, significant complexity is introduced. The non-zero off-diagonal terms in the characteristic equation mean that the modes are coupled, and that it is not possible to get exact expressions for the velocities in a generalised form. Neighbours et al. used a perturbation method to find the generalised roots of the secular equation; the equations derived from this method are given below.

D.1 Cubic materials

As the full form of the equations is rather-long winded, it is first useful to define some contractions based on the primed elastic constants

$$C'_{11} = C_{11} - (C_{11} - C_{12} - 2C_{44}) \left((l_1^2 l_2^2 + l_1^2 l_2^2 + l_1^2 l_2^2) \right)$$

$$C'_{15} = (C_{11} - C_{12} - 2C_{44}) - \frac{l_3(l_1^4 + l_2^4 - l_3^2(l_1^2 + l_2^2))}{l_1^2 + l_2^2}$$

$$C'_{16} = (C_{11} - C_{12} - 2C_{44}) \frac{1}{l_1 l_2 (l_2^2 - l_1^2)(l_1^2 + m_1^2)}{l_1^2 + l_2^2}$$

$$C'_{55} = C_{44} + (C_{11} - C_{12} - 2C_{44}) \frac{2l_3^2(l_1^4 l_1^2 l_2^2 + l_2^4)}{l_1^2 + l_2^2}$$

$$C'_{56} = (C_{11} - C_{12} - 2C_{44}) \frac{l_1 l_2 l_3(l_1^2 l_2^2)}{l_1^2 + l_2^2}$$

$$C'_{66} = C_{44} + (C_{11} - C_{12} - 2C_{44}) \frac{2l_1^2 l_2^2}{l_1^2 + l_2^2}$$

$$C'_{66} = C_{44} + (C_{11} - C_{12} - 2C_{44}) \frac{2l_1^2 l_2^2}{l_1^2 + l_2^2}$$

The generalised relations between the three bulk wave velocities and elastic constants in cubic materials, in a direction defined by the vector l, is then given by

$$\rho v_L^2 = C_{11}' + \left(\frac{C_{15}'^2}{C_{11}' - C_{55}'} + \frac{C_{16}'^2}{C_{11}' - C_{66}'} + \frac{2C_{15}'C_{16}'C_{56}'}{(C_{11}' - C_{66}')(C_{11}' - C_{55}')} \right)
\rho v_L^2 = C_{66}' + \left(\frac{C_{16}^2}{C_{66} - C_{11}} + \frac{C_{56}'^2}{C_{66}' - C_{55}'} + \frac{2C_{15}'C_{16}'C_{56}}{(C_{66}' - C_{11}')(C_{66}' - C_{55})} \right)
\rho v_L^2 = C_{55}' + \left(\frac{C_{15}'^2}{C_{55}' - C_{11}'} + \frac{C_{56}'^2}{C_{55}' - C_{66}'} + \frac{2C_{15}'C_{16}'C_{56}'}{(C_{55}' - C_{11}')(C_{55}' - C_{66}')} \right)$$
(D.2)

Simplified forms for the planes (001) and (101) are given in table 2.2.

D.2 Hexagonal materials

Again, the full form of the bulk velocity relations in hexagonal materials is rather long-winded, and thus some contractions based on the primed elastic constants are useful, as define by Eros and Smith [358]. The angle θ is given by the angle between the *c*-axis and the axis of wave propagation.

$$C'_{11} = C_{11} \sin^4 \theta + C_{33} \cos^4 \theta + 2(Cele_{13} + 2C_{44}) \sin^2 \theta \cos^2 \theta$$

$$C'_{15} = (C_{13} - C_{11} + 2C_{44}) \sin^3 \theta \cos \theta + (C_{33} - C_{13} - 2C_{44}) \cos^3 \theta \sin \theta \qquad (D.3)$$

$$C'_{55} = (C_{13} - C_{11} + 2C_{44}) \sin^2 \theta \cos^2 \theta$$

We may then define the three bulk wave velocities as

$$\rho v_L^2 = C'_{11} + \frac{C'_{15}{}^2}{C'_{11} - C'_{55}}
\rho v_{T,1}^2 = \frac{1}{2} (C_{11} - C_{12}) \sin^2 \theta + C_{44} \cos^2 \theta
\rho v_{T,2}^2 = C'_{55} + \frac{C'_{15}{}^2}{C'_{11} - C'_{55}}$$
(D.4)

In the case of $\theta = 0^{\circ}$ or $\theta = 90^{\circ}$, the equations above reduce to the forms given in table 2.2.

Appendix E

R-value calculation

The *R*-value has been used in chapter 4 and 6 to describe the disagreement between crystallographic orientations of nickel specimens. From [359] there are five common definitions of the *R*-value; the definition used in this work is the two-angle cosine. A graphical definition of this method and he two requisite angles (ϕ and τ) is shown in fig. E.1.

| Name | Formula |
|--------------------|--|
| One-angle cosine | $R = \min(\theta_{1,\dots,24})$ |
| Two-angle cosine | $R = \cos^{-1}(\cos\phi\cos\tau)$ |
| Two-angle RMS | $R = \sqrt{\phi^2 + \tau^2}$ |
| Three-angle cosine | $R = \cos^{-1}(\cos(\alpha_1 - \alpha_2)\cos(\beta_1 - \beta_2)\cos(\gamma_1 - \gamma_2))$ |
| Three-angle RMS | $R = \sqrt{(\alpha_1 - \alpha_2)^2 + (\beta_1 - \beta_2)^2 + (\gamma_1 - \gamma_2)^2}$ |

 Table E.1: Definitions of R-value.



Figure E.1: Definition of two-angle cosine *R*-value.

Bibliography

- D. Hoang, D. Perrault, M. Stevanovic, and A. Ghiassi, "Surgical applications of threedimensional printing: a review of the current literature & how to get started," Annals of translational medicine, vol. 4, pp. 456–456, Dec. 2016.
- [2] M. Gorelik, "Additive manufacturing in the context of structural integrity," International Journal of Fatigue, vol. 94, pp. 168–177, 2017. Fatigue and Fracture Behavior of Additive Manufactured Parts.
- [3] J. Mazumder, "Design for metallic additive manufacturing machine with capability for "certify as you build"," *Proceedia Cirp*, vol. 36, pp. 187–192, 2015.
- [4] S. Sharples, All-Optical Scanning Acoustic Microscope. University of Nottingham, 2003.
- [5] W. Li, Laser Ultrasonic Method for Determination of Crystallographic Orientation of Large Grain Metals by Spatially Resolved Acoustic Spectroscopy (SRAS). PhD thesis, University of Nottingham, Electrical and Electronic Engineering, September 2012.
- [6] F. Abe, K. Osakada, M. Shiomi, K. Uematsu, and M. Matsumuto, "The manufacturing of hard tools from metallic powders by selective laser melting," *Journal of Materials Processing Technology*, vol. 111, no. 1-3, pp. 210–213, 2001.
- [7] A. Voet, J. Mingneau, J. Dehaes, J.-P. Kruth, and v. J. Vaerenbergh, "Study of the wear behaviour of conventional and rapid tooling mould materials," in *International Conference Polymers and Moulds Innovations PMI*, 2005.
- [8] F. P. W. Melchels, M. A. N. Domingos, T. J. Klein, J. Malda, P. J. Bartolo, and D. W. Hutmacher, "Additive manufacturing of tissues and organs," *Progress in Polymer Science*, vol. 37, no. 8, pp. 1079–1104, 2012.
- [9] D. Brackett, I. Ashcroft, and R. Hague, "Topology optimization for additive manufacturing," in *Proceedings of the solid freeform fabrication symposium*, Austin, TX, vol. 1, pp. 348–362, 2011.

- [10] A. Großmann, P. Weis, C. Clemen, and C. Mittelstedt, "Optimization and re-design of a metallic riveting tool for additive manufacturing—a case study," *Additive Manufacturing*, vol. 31, p. 100892, 2020.
- [11] M. Leary, L. Merli, F. Torti, M. Mazur, and M. Brandt, "Optimal topology for additive manufacture: A method for enabling additive manufacture of support-free optimal structures," *Materials & Design*, vol. 63, pp. 678–690, 2014.
- [12] T. J. Horn and O. L. A. Harrysson, "Overview of current additive manufacturing technologies and selected applications," *Science Progress*, vol. 95, no. 3, 2012.
- [13] M. Simonelli, N. T. Aboulkhair, P. Cohen, J. W. Murray, A. T. Clare, C. Tuck, and R. J. M. Hague, "A comparison of Ti-6Al-4V in-situ alloying in selective laser melting using simply-mixed and satellited powder blend feedstocks," *Materials Characterization*, vol. 143, pp. 118 126, 2018. Metal Additive Manufacturing: Microstructures and Properties.
- [14] K. Moussaoui, W. Rubio, M. Mousseigne, T. Sultan, and F. Rezai, "Effects of selective laser melting additive manufacturing parameters of inconel 718 on porosity, microstructure and mechanical properties," *Materials Science and Engineering: A*, vol. 735, pp. 182–190, 2018.
- [15] B. Wu, Z. Pan, D. Ding, D. Cuiuri, H. Li, J. Xu, and J. Norrish, "A review of the wire arc additive manufacturing of metals: properties, defects and quality improvement," *Journal of Manufacturing Processes*, vol. 35, pp. 127–139, 2018.
- [16] Y. Kok, X. Tan, P. Wang, M. Nai, N. Loh, E. Liu, and S. Tor, "Anisotropy and heterogeneity of microstructure and mechanical properties in metal additive manufacturing: A critical review," *Materials & Design*, vol. 139, pp. 565–586, 2018.
- [17] M. J. Kenney, K. O'Donnell, M. J. Quintana, and P. C. Collins, "Spherical pores as 'microstructural informants': Understanding compositional, thermal, and mechanical gyrations in additively manufactured ti-6al-4v," *Scripta Materialia*, vol. 198, p. 113827, 2021.
- [18] J. H. Martin, B. D. Yahata, J. M. Hundley, J. A. Mayer, T. A. Schaedler, and T. M. Pollock, "3d printing of high-strength aluminium alloys," *Nature*, vol. 549, pp. 365–369, Sept. 2017.
- [19] W. E. Frazier, D. Polakovics, and W. Koegel, "Qualifying of metallic materials and structures for aerospace applications," *Jom*, vol. 53, no. 3, pp. 16–18, 2001.
- [20] C. A. Brice, Unintended consequences: How qualification constrains innovation, vol. 241, pp. 241–256. TMS (The Minerals, Metals & Materials Society), 2011.

- [21] M. O'Brien, "Existing standards as the framework to qualify additive manufacturing of metals," in 2018 IEEE Aerospace Conference, pp. 1–10, IEEE, 2018.
- [22] Federal Aviation Administration, "Material strength properties and design values. (a). us," tech. rep., US Government, 2010.
- [23] T. Pereira, J. V. Kennedy, and J. Potgieter, "A comparison of traditional manufacturing vs additive manufacturing, the best method for the job," *Procedia manufacturing*, vol. 30, pp. 11–18, 2019.
- [24] M. Seifi, A. Salem, J. Beuth, O. Harrysson, and J. J. Lewandowski, "Overview of materials qualification needs for metal additive manufacturing," *Jom*, vol. 68, no. 3, pp. 747–764, 2016.
- [25] A. G. Demir and B. Previtali, "Investigation of remelting and preheating in slm of 18ni300 maraging steel as corrective and preventive measures for porosity reduction," *The International Journal of Advanced Manufacturing Technology*, vol. 93, no. 5, pp. 2697–2709, 2017.
- [26] F. Martina, P. A. Colegrove, S. W. Williams, and J. Meyer, "Microstructure of interpass rolled wire+ arc additive manufacturing ti-6al-4v components," *Metallurgical and Materials Transactions A*, vol. 46, no. 12, pp. 6103–6118, 2015.
- [27] S. C. Cowin, Continuum Mechanics of Anisotropic Materials. Springer, 2013.
- [28] H. M. Ledbetter and R. P. Reed, "Elastic properties of metals and alloys, i. iron, nickel, and iron-nickel alloys," *Journal of Physical and Chemical Reference Data*, vol. 2, no. 3, pp. 531– 618, 1973.
- [29] M. H. Sadd, *Elasticity: theory, applications, and numerics*. Academic Press, 2009.
- [30] M. Jamal, S. Jalali Asadabadi, I. Ahmad, and H. A. Rahnamaye Aliabad, "Elastic constants of cubic crystals," *Computational Materials Science*, vol. 95, pp. 592 – 599, 2014.
- [31] A. G. Every, "Determination of the elastic constants of anisotropic solids," NDT & E International, vol. 27, no. 1, pp. 3–10, 1994.
- [32] A. G. Every and A. K. McCurdy, Landolt-B[']ornstein group III condensed matter, ch. Volume 29A - Low Frequency Properties of Dielectric Crystals, Second and Higher Order Elastic Constants. Springer, 1992.
- [33] M. Farley, "40 years of progress in NDT history as a guide to the future," in AIP Conference Proceedings, vol. 1581, pp. 5–12, American Institute of Physics, 2014.
- [34] P. Mix, Introduction to Nondestructive Testing. Wiley and Sons., 2005.
- [35] P. Cawley, "Non-destructive testing—current capabilities and future directions," Proceedings of the Institution of Mechanical Engineers, Part L: Journal of Materials: Design and Applications, vol. 215, no. 4, pp. 213–223, 2001.
- [36] T. D. et al., "A landscape for the future of NDT in the UK economy," tech. rep., Materials: Knowledge Transfer Network, 2014.
- [37] M. Turnbow and J. Thompson, "Reliability of nde the criticality of procedures and personnel," 4th European-American Workshop on Reliability of NDE, 2009.
- [38] B. Raj, V. Moorthy, T. Jayakumar, and K. B. S. Rao, "Assessment of microstructures and mechanical behaviour of metallic materials through non-destructive characterisation," *International Materials Reviews*, vol. 48, no. 5, pp. 273–325, 2003.
- [39] C. B. Scruby, "An introduction to acoustic emission," Journal of Physics E: Scientific Instruments, vol. 20, pp. 946–953, aug 1987.
- [40] B. Ribic, P. Burgardt, and T. DebRoy, "Optical emission spectroscopy of metal vapor dominated laser-arc hybrid welding plasma," *Journal of Applied Physics*, vol. 109, no. 8, p. 083301, 2011.
- [41] J. C. Limited, "Probability of detection (pod) curves: Derivation, applications and limitations," tech. rep., Health and Safety Executive, 2006.
- [42] R. Halmshaw, Industrial radiogrpahy: theory and practice, ch. 8, pp. 182–205. Springer science, second ed., 1995.
- [43] W. A. Kalender, "X-ray computed tomography," *Physics in Medicine and Biology*, vol. 51, pp. R29–R43, jun 2006.
- [44] S. K. Dinda, J. M. Warnett, M. A. Williams, G. G. Roy, and P. Srirangam, "3d imaging and quantification of porosity in electron beam welded dissimilar steel to fe-al alloy joints by x-ray tomography," *Materials & Design*, vol. 96, pp. 224–231, 2016.
- [45] A. Speidel, R. Su, J. Mitchell-Smith, P. Dryburgh, I. Bisterov, D. Pieris, W. Li, R. Patel, M. Clark, and A. T. Clare, "Crystallographic texture can be rapidly determined by electrochemical surface analytics," Acta Materialia, vol. 159, pp. 89 – 101, 2018.
- [46] V. S. Tong and T. Ben Britton, "Trueebsd: Correcting spatial distortions in electron backscatter diffraction maps," *Ultramicroscopy*, vol. 221, p. 113130, 2021.

- [47] J. Cheeke and N. David, Fundementals and Applications of Ultrasonic Waves. CRC Press, 2002.
- [48] H. Lamb, "On waves in elastic plate," Proceedings of the Royal Society of London, vol. 93, no. 648, 1917.
- [49] I. A. Viktorov, Rayleigh and Lamb waves Physical theory and applications. Plenum Press, english ed., 1967.
- [50] D. N. Alleyne and P. Cawley, "Optimization of Lamb wave inspection techniques," Ndt & E International, vol. 25, no. 1, pp. 11–22, 1992.
- [51] A. E. H. Love, "Some problems of geodynamics," Nature, vol. 89, no. 2228, pp. 471–472, 1912.
- [52] J. L. Rose, Ultrasonic Waves in Solid Media. Cambridge Press, 1999.
- [53] A. Migliori, J. Sarrao, W. M. Visscher, T. Bell, M. Lei, Z. Fisk, and R. Leisure, "Resonant ultrasound spectroscopic techniques for measurement of the elastic moduli of solids," *Physica B: Condensed Matter*, vol. 183, no. 1, pp. 1 – 24, 1993.
- [54] F. Lépine, G. Sansone, and M. J. Vrakking, "Molecular applications of attosecond laser pulses," *Chemical Physics Letters*, vol. 578, pp. 1–14, 2013.
- [55] A. V. Korzhimanov, A. A. Gonoskov, E. A. Khazanov, and A. M. Sergeev, "Horizons of petawatt laser technology," *Physics-Uspekhi*, vol. 54, no. 1, p. 9, 2011.
- [56] C. Scruby and L. Drain, Laser Ultrasonics: techniques and applications. Taylor and Francis, 1990.
- [57] C. Scruby, R. Dewhurst, D. Hutchins, and S. Palmer, "Quantitative studies of thermally generated elastic waves in laser-irradiated metals," *Journal of Applied Physics*, vol. 51, no. 12, pp. 6210–6216, 1980.
- [58] S.-L. Chen and C. Tian, "Recent developments in photoacoustic imaging and sensing for nondestructive testing and evaluation," Visual Computing for Industry, Biomedicine, and Art, vol. 4, no. 1, p. 6, 2021.
- [59] S. G. Pierce, B. Culshaw, and Q. Shan, "Laser generation of ultrasound using a modulated continuous wave laser diode," *Applied Physics Letters*, vol. 72, no. 9, pp. 1030–1032, 1998.

- [60] G. Wissmeyer, M. A. Pleitez, A. Rosenthal, and V. Ntziachristos, "Looking at sound: optoacoustics with all-optical ultrasound detection," *Light: Science & Applications*, vol. 7, no. 1, p. 53, 2018.
- [61] J.-P. Monchalin, "Optical detection of ultrasound," IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 33, no. 5, pp. 485–499, 1986.
- [62] B. F. Pouet, R. K. Ing, S. Krishnaswamy, and D. Royer, "Heterodyne interferometer with two-wave mixing in photorefractive crystals for ultrasound detection on rough surfaces," *Applied Physics Letters*, vol. 69, no. 25, pp. 3782–3784, 1996.
- [63] D. Cook and Y. Berthelot, "Detection of small surface-breaking fatigue cracks in steel using scattering of rayleigh waves," NDT & E International, vol. 34, no. 7, pp. 483–492, 2001.
- [64] P. D. Warren, C. Pecorari, O. V. Kolosov, S. G. Roberts, and G. A. D. Briggs, "Characterization of surface damage via surface acoustic waves," *Nanotechnology*, vol. 7, no. 3, p. 295, 1996.
- [65] V. Jaya Rao, E. Kannan, R. V. Prakash, and K. Balasubramaniam, "Fatigue damage characterization using surface acoustic wave nonlinearity in aluminum alloy aa7175-t7351," *Journal* of Applied Physics, vol. 104, no. 12, p. 123508, 2008.
- [66] S. D. Sharples, M. Clark, and M. G. Somekh, "Fast noncontact imaging of material microstructure using local surface acoustic wave velocity mapping," in *IEEE Ultrasonics Symposium*, 2005., vol. 2, pp. 886–889, IEEE, 2005.
- [67] A. Briggs, "Acoustic microscopy-a summary," *Reports on Progress in Physics*, vol. 55, no. 7, p. 851, 1992.
- [68] M. Clark, S. Sharples, and M. Somekh, "Fast, all-optical rayleigh wave microscope: Imaging on isotropic and anisotropic materials," *IEEE Transactions on Ultrasonics, Ferroelectrics,* and Frequency Control, vol. 47, no. 1, pp. 65–74, 2000.
- [69] M. Clark, S. D. Sharples, and M. Somekh, "Optimisation using measured green's function for improving spatial coherence in acoustic measurements," *Ultrasonics*, vol. 42, no. 1, pp. 205– 212, 2004. Proceedings of Ultrasonics International 2003.
- [70] S. D. Sharples, M. Clark, W. Li, and M. G. Somekh, "Rapid imaging of microstructure using spatially resolved acoustic spectroscopy," in 1st International Symposium on Laser Ultrasonics: Montreal, Canada, Citeseer, 2008.

- [71] S. D. Sharples, M. Clark, and M. G. Somekh, "Spatially resolved acoustic spectroscopy for fast noncontact imaging of material microstructure," *Opt. Express*, vol. 14, pp. 10435–10440, Oct 2006.
- [72] W. E. Frazier, "Metal additive manufacturing: A review," Journal of Materials Engineering and Performance, vol. 23, 2014.
- [73] W. E. King, A. T. Anderson, R. M. Ferencz, N. E. Hodge, C. Kamath, S. A. Khairallah, and A. M. Rubenchik, "Laser powder bed fusion additive manufacturing of metals; physics, computational, and materials challenges," *Applied Physics Reviews*, vol. 2, no. 4, p. 041304, 2015.
- [74] H. K. Rafi, N. V. Karthik, H. Gong, T. L. Starr, and B. E. Stucker, "Microstructures and mechanical properties of ti6al4v parts fabricated by selective laser melting and electron beam melting," *Journal of Materials Engineering and Performance*, vol. 22, no. 12, pp. 3872–3883, 2013.
- [75] I. Yadroitsava, A. du Plessis, and I. Yadroitsev, "Chapter 12 bone regeneration on implants of titanium alloys produced by laser powder bed fusion: A review," in *Titanium for Consumer Applications* (F. Froes, M. Qian, and M. Niinomi, eds.), pp. 197–233, Elsevier, 2019.
- [76] R. Liu, Z. Wang, T. Sparks, F. Liou, and J. Newkirk, "13 aerospace applications of laser additive manufacturing," in *Laser Additive Manufacturing* (M. Brandt, ed.), Woodhead Publishing Series in Electronic and Optical Materials, pp. 351–371, Woodhead Publishing, 2017.
- [77] S. Sanchez, P. Smith, Z. Xu, G. Gaspard, C. J. Hyde, W. W. Wits, I. A. Ashcroft, H. Chen, and A. T. Clare, "Powder bed fusion of nickel-based superalloys: A review," *International Journal of Machine Tools and Manufacture*, vol. 165, p. 103729, 2021.
- [78] M. Van Elsen, Complexity of selective laser melting: a new optimisation approach. PhD thesis, Katholieke Universiteit Leuven, 2007.
- [79] D. Gu, H. Zhang, D. Dai, C. Ma, H. Zhang, Y. Li, and S. Li, "Anisotropic corrosion behavior of sc and zr modified al-mg alloy produced by selective laser melting," *Corrosion Science*, vol. 170, p. 108657, 2020.
- [80] C. L. Frederick, A. Plotkowski, M. M. Kirka, M. Haines, A. Staub, E. J. Schwalbach, D. Cullen, and S. S. Babu, "Geometry-induced spatial variation of microstructure evolution during selective electron beam melting of rene-n5," *Metallurgical and Materials Transactions* A, vol. 49, no. 10, pp. 5080–5096, 2018.

- [81] M. Hirsch, S. Catchpole-Smith, R. Patel, P. Marrow, W. Li, C. Tuck, S. D. Sharples, and A. T. Clare, "Meso-scale defect evaluation of selective laser melting using spatially resolved acoustic spectroscopy," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 473, no. 2205, 2017.
- [82] T. Niendorf, S. Leuders, A. Riemer, F. Brenne, T. Tröster, H. A. Richard, and D. Schwarze, "Functionally graded alloys obtained by additive manufacturing," *Advanced engineering materials*, vol. 16, no. 7, pp. 857–861, 2014.
- [83] Y. Lei, K. Aoyagi, and A. Chiba, "A method to manipulate non-steady-state columnar-toequiaxed transition in powder bed fusion additive manufacturing using an electron beam," *Acta Materialia*, vol. 227, p. 117717, 2022.
- [84] R. R. Dehoff, M. M. Kirka, W. J. Sames, H. Bilheux, A. S. Tremsin, L. E. Lowe, and S. S. Babu, "Site specific control of crystallographic grain orientation through electron beam additive manufacturing," *Materials Science and Technology*, vol. 31, no. 8, pp. 931–938, 2015.
- [85] S.-H. Sun, T. Ishimoto, K. Hagihara, Y. Tsutsumi, T. Hanawa, and T. Nakano, "Excellent mechanical and corrosion properties of austenitic stainless steel with a unique crystallographic lamellar microstructure via selective laser melting," *Scripta Materialia*, vol. 159, pp. 89–93, 2019.
- [86] K. Sofinowski, M. Wittwer, and M. Seita, "Encoding data into metal alloys using laser powder bed fusion," *Additive Manufacturing*, p. 102683, 2022.
- [87] A. Plotkowski, J. Ferguson, B. Stump, W. Halsey, V. Paquit, C. Joslin, S. Babu, A. Marquez Rossy, M. Kirka, and R. Dehoff, "A stochastic scan strategy for grain structure control in complex geometries using electron beam powder bed fusion," *Additive Manufacturing*, vol. 46, p. 102092, 2021.
- [88] C. Sly, "Counterfeit parts: The scourge of industry," *Practical Failure Analysis*, vol. 1, no. 4, pp. 8–12, 2001.
- [89] M. Sandborn, C. Olea, J. White, C. Williams, P. A. Tarazaga, L. Sturm, M. Albakri, and C. Tenney, "Towards secure cyber-physical information association for parts," *Journal of Manufacturing Systems*, vol. 59, pp. 27–41, 2021.
- [90] M. Yampolskiy, W. E. King, J. Gatlin, S. Belikovetsky, A. Brown, A. Skjellum, and Y. Elovici, "Security of additive manufacturing: Attack taxonomy and survey," *Additive Manufacturing*, vol. 21, pp. 431–457, 2018.

- [91] M. Yampolskiy, J. Gatlin, and M. Yung, "Myths and misconceptions in additive manufacturing security: Deficiencies of the cia triad," in *Proceedings of the 2021 Workshop on Additive Manufacturing (3D Printing) Security*, pp. 3–9, 2021.
- [92] L. Graves, W. King, P. Carrion, S. Shao, N. Shamsaei, and M. Yampolskiy, "Sabotaging metal additive manufacturing: Powder delivery system manipulation and material-dependent effects," *Additive Manufacturing*, vol. 46, p. 102029, 2021.
- [93] A. Dachowicz, S. C. Chaduvula, M. Atallah, and J. H. Panchal, "Microstructure-based counterfeit detection in metal part manufacturing," *JOM*, vol. 69, no. 11, pp. 2390–2396, 2017.
- [94] C. Wei, Z. Sun, Y. Huang, and L. Li, "Embedding anti-counterfeiting features in metallic components via multiple material additive manufacturing," *Additive Manufacturing*, vol. 24, pp. 1–12, 2018.
- [95] A. Sola, Y. Sai, A. Trinchi, C. Chu, S. Shen, and S. Chen, "How can we provide additively manufactured parts with a fingerprint? a review of tagging strategies in additive manufacturing," *Materials*, vol. 15, no. 1, p. 85, 2022.
- [96] M. H. Mosallanejad, B. Niroumand, A. Aversa, and A. Saboori, "In-situ alloying in laserbased additive manufacturing processes: A critical review," *Journal of Alloys and Compounds*, vol. 872, p. 159567, 2021.
- [97] P. Vora, K. Mumtaz, I. Todd, and N. Hopkinson, "Alsi12 in-situ alloy formation and residual stress reduction using anchorless selective laser melting," *Additive Manufacturing*, vol. 7, pp. 12–19, 2015.
- [98] Z. Liu, D. Zhang, S. Sing, C. Chua, and L. Loh, "Interfacial characterization of slm parts in multi-material processing: Metallurgical diffusion between 316l stainless steel and c18400 copper alloy," *Materials Characterization*, vol. 94, pp. 116–125, 2014.
- [99] R. Duan, S. Li, B. Cai, W. Zhu, F. Ren, and M. M. Attallah, "A high strength and low modulus metastable β ti-12mo-6zr-2fe alloy fabricated by laser powder bed fusion in-situ alloying," *Additive Manufacturing*, vol. 37, p. 101708, 2021.
- [100] A. Pfaff, M. Jäcklein, M. Schlager, W. Harwick, K. Hoschke, and F. Balle, "An empirical approach for the development of process parameters for laser powder bed fusion," *Materials*, vol. 13, no. 23, 2020.

- [101] C. Druzgalski, A. Ashby, G. Guss, W. King, T. Roehling, and M. Matthews, "Process optimization of complex geometries using feed forward control for laser powder bed fusion additive manufacturing," *Additive Manufacturing*, vol. 34, p. 101169, 2020.
- [102] J. C. Najmon, S. Raeisi, and A. Tovar, "2 review of additive manufacturing technologies and applications in the aerospace industry," in *Additive Manufacturing for the Aerospace Industry* (F. Froes and R. Boyer, eds.), pp. 7–31, Elsevier, 2019.
- [103] P. Edwards, A. O'Conner, and M. Ramulu, "Electron beam additive manufacturing of titanium components: Properties and performance," *Journal of Manufacturing Science and Engineering*, vol. 135, no. 6, p. 061016, 2013.
- [104] M. Fousová, D. Vojtěch, J. Kubásek, E. Jablonská, and J. Fojt, "Promising characteristics of gradient porosity ti-6al-4v alloy prepared by slm process," *Journal of the mechanical behavior* of biomedical materials, vol. 69, pp. 368–376, 2017.
- [105] S. Leuders, M. Thöne, A. Riemer, T. Niendorf, T. Tröster, H. A. Richard, and H. J. Maier, "On the mechanical behaviour of titanium alloy tial6v4 manufactured by selective laser melting: Fatigue resistance and crack growth performance," *International Journal of Fatigue*, vol. 48, pp. 300–307, 2013.
- [106] S. K. Everton, M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare, "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing," *Materials & Design*, 2016.
- [107] S. Tammas-Williams, H. Zhao, F. Léonard, F. Derguti, I. Todd, and P. B. Prangnell, "Xct analysis of the influence of melt strategies on defect population in ti-6al-4v components manufactured by selective electron beam melting," *Materials Characterization*, vol. 102, pp. 47– 61, 2015.
- [108] A. du Plessis, I. Yadroitsava, and I. Yadroitsev, "Effects of defects on mechanical properties in metal additive manufacturing: A review focusing on x-ray tomography insights," *Materials & Design*, vol. 187, p. 108385, 2020.
- [109] K. Darvish, Z. W. Chen, and T. Pasang, "Reducing lack of fusion during selective laser melting of cocrmo alloy: Effect of laser power on geometrical features of tracks," *Materials* and Design, vol. 112, pp. 357–366, 2016.
- [110] C. L. Qiu, N. J. E. Adkins, and M. M. Attallah, "Microstructure and tensile properties of selectively laser-melted and of hiped laser-melted ti-6al-4v," *Materials Science and Engineer*-

ing a-Structural Materials Properties Microstructure and Processing, vol. 578, pp. 230–239, 2013.

- [111] M. J. Matthews, G. Guss, S. A. Khairallah, A. M. Rubenchik, P. J. Depond, and W. E. King, "Denudation of metal powder layers in laser powder bed fusion processes," *Acta Materialia*, vol. 114, pp. 33–42, 2016.
- [112] R. Li, J. Liu, Y. Shi, L. Wang, and W. Jiang, "Balling behavior of stainless steel and nickel powder during selective laser melting process," *The International Journal of Advanced Manufacturing Technology*, vol. 59, no. 9-12, pp. 1025–1035, 2011.
- [113] D. Gu and Y. Shen, "Balling phenomena in direct laser sintering of stainless steel powder: Metallurgical mechanisms and control methods," *Materials and Design*, vol. 30, no. 8, pp. 2903–2910, 2009.
- [114] J. P. Kruth, L. Froyen, J. Van Vaerenbergh, P. Mercelis, M. Rombouts, and B. Lauwers, "Selective laser melting of iron-based powder," *Journal of Materials Processing Technology*, vol. 149, no. 1–3, pp. 616–622, 2004.
- [115] L. Thijs, F. Verhaeghe, T. Craeghs, J. V. Humbeeck, and J.-P. Kruth, "A study of the microstructural evolution during selective laser melting of ti–6al–4v," Acta Materialia, vol. 58, no. 9, pp. 3303–3312, 2010.
- [116] N. J. Harrison, I. Todd, and K. Mumtaz, "Reduction of micro-cracking in nickel superalloys processed by selective laser melting: A fundamental alloy design approach," Acta Materialia, vol. 94, pp. 59–68, 2015.
- [117] P. Mercelis and J.-P. Kruth, "Residual stresses in selective laser sintering and selective laser melting," *Rapid Prototyping Journal*, vol. 12, no. 5, 2006.
- [118] M. F. Zaeh and G. Branner, "Investigations on residual stresses and deformations in selective laser melting," *Production Engineering*, vol. 4, no. 1, pp. 35–45, 2009.
- [119] K. Murali, A. N. Chatterjee, P. Saha, R. Palai, S. Kumar, S. K. Roy, P. K. Mishra, and A. R. Choudhury, "Direct selective laser sintering of iron-graphite powder mixture," *Journal* of Materials Processing Technology, vol. 136, no. 1-3, pp. 179–185, 2003.
- [120] H. J. Niu and I. T. H. Chang, "Instability of scan tracks of selective laser sintering of high speed steel powder," *Scripta Materialia*, vol. 41, no. 11, pp. 1229–1234, 1999.

- [121] S. P. Moylan, J. Drescher, and M. A. Donmez, "Powder bed fusion machine performance testing," in Proceedings of the 2014 ASPE Spring Topical Meeting–Dimensional Accuracy and Surface Finish in Additive Manufacturing, vol. 57, 2014.
- [122] P. Almeida and S. Williams, "Innovative process model of ti-6al-4v additive layer manufacturing using cold metal transfer (cmt)," 21st solid Freeform Fabrication Symposium, 2010.
- [123] F. Martina, J. Mehnen, S. W. Williams, P. Colegrove, and F. Wang, "Investigation of the benefits of plasma deposition for the additive layer manufacture of ti-6al-4v," *Journal of Materials Processing Technology*, vol. 212, no. 6, pp. 1377–1386, 2012.
- [124] F. Martina, J. Ding, S. Williams, A. Caballero, G. Pardal, and L. Quintino, "Tandem metal inert gas process for high productivity wire arc additive manufacturing in stainless steel," *Additive Manufacturing*, vol. 25, pp. 545 – 550, 2019.
- [125] S. W. Williams, F. Martina, A. C. Addison, J. Ding, G. Pardal, and P. Colegrove, "Wire + arc additive manufacturing," *Materials Science and Technology*, vol. 32, pp. 641–647, May 2016.
- [126] J. Donoghue, A. Antonysamy, F. Martina, P. A. Colegrove, S. Williams, and P. Prangnell, "The effectiveness of combining rolling deformation with wire-arc additive manufacture on beta-grain refinement and texture modification in ti-6al-4v," *Materials Characterization*, vol. 114, pp. 103–114, 2016.
- [127] M. Simonelli, Y. Y. Tse, and C. Tuck, "On the texture formation of selective laser melted ti-6al-4v," *Metallurgical and Materials Transactions A*, vol. 45, no. 6, pp. 2863–2872, 2014.
- [128] I. Bantounas, D. Dye, and T. C. Lindley, "The role of microtexture on the faceted fracture morphology in ti-6al-4v subjected to high-cycle fatigue," Acta Materialia, vol. 58, no. 11, pp. 3908–3918, 2010.
- [129] B. E. Carroll, T. A. Palmer, and A. M. Beese, "Anisotropic tensile behavior of ti-6al-4v components fabricated with directed energy deposition additive manufacturing," Acta Materialia, vol. 87, pp. 309–320, 2015.
- [130] P. A. Colegrove, H. E. Coules, J. Fairman, F. Martina, T. Kashoob, H. Mamash, and L. D. Cozzolino, "Microstructure and residual stress improvement in wire and arc additively manufactured parts through high-pressure rolling," *Journal of Materials Processing Technology*, vol. 213, no. 10, pp. 1782–1791, 2013.

- [131] P. A. Colegrove, F. Martina, M. J. Roy, B. A. Szost, S. Terzi, S. W. Williams, P. J. Withers, and D. Jarvis, "High pressure interpass rolling of wire+ arc additively manufactured titanium components," in *Advanced Materials Research*, vol. 996, pp. 694–700, Trans Tech Publ, 2014.
- [132] A. R. McAndrew, M. Alvarez Rosales, P. A. Colegrove, J. R. Honnige, A. Ho, R. Fayolle, K. Eyitayo, I. Stan, P. Sukrongpang, A. Crochemore, and Z. Pinter, "Interpass rolling of ti-6al-4v wire + arc additively manufactured features for microstructural refinement," *Additive Manufacturing*, vol. 21, pp. 340–349, 2018.
- [133] S. Singh, S. kumar Sharma, and D. W. Rathod, "A review on process planning strategies and challenges of waam," *Materials Today: Proceedings*, vol. 47, pp. 6564–6575, 2021. International Conference on Advances in Design, Materials and Manufacturing.
- [134] F. Wang, S. Williams, P. Colegrove, and A. A. Antonysamy, "Microstructure and mechanical properties of wire and arc additive manufactured ti-6al-4v," *Metallurgical and Materials Transactions A*, vol. 44, no. 2, pp. 968–977, 2013.
- [135] T. M. Pollock and S. Tin, "Nickel-based superalloys for advanced turbine engines: chemistry, microstructure and properties," *Journal of propulsion and power*, vol. 22, no. 2, pp. 361–374, 2006.
- [136] S. Catchpole-Smith, N. Aboulkhair, L. Parry, C. Tuck, I. Ashcroft, and A. Clare, "Fractal scan strategies for selective laser melting of 'unweldable' nickel superalloys," *Additive Manufacture*, vol. In Press, 2017.
- [137] M. M. Attallah, R. Jennings, X. Wang, and L. N. Carter, "Additive manufacturing of nibased superalloys: The outstanding issues," *MRS bulletin*, vol. 41, no. 10, pp. 758–764, 2016.
- [138] M. J. Donachie, Titanium: A Technical Guide. ASM International, second edition ed., 2000.
- [139] A. Ho, H. Zhao, J. W. Fellowes, F. Martina, A. E. Davis, and P. B. Prangnell, "On the origin of microstructural banding in ti-6al4v wire-arc based high deposition rate additive manufacturing," Acta Materialia, vol. 166, pp. 306 – 323, 2019.
- [140] L. Germain, S. R. Dey, M. Humbert, and N. Gey, "Determination of parent orientation maps in advanced titanium-based alloys," *Journal of microscopy*, vol. 227, no. 3, pp. 284–291, 2007.
- [141] P. Davies, An Investigation of Microstructure and Texture Evolution in the Near-alpha Titanium alloy Timetal 834. PhD thesis, University of Sheffield, 2009.

- [142] F. Dunne, A. Walker, and D. Rugg, "A systematic study of hcp crystal orientation and morphology effects in polycrystal deformation and fatigue," *Proceedings of the Royal Society* of London A: Mathematical, Physical and Engineering Sciences, vol. 463, no. 2082, pp. 1467– 1489, 2007.
- [143] V. Sinha, M. Mills, J. C. Williams, and J. Spowart, "Observations on the faceted initiation site in the dwell-fatigue tested ti-6242 alloy: crystallographic orientation and size effects," *Metallurgical and Materials Transactions A*, vol. 37, no. 5, pp. 1507–1518, 2006.
- [144] G. Ouyang, X. Chen, Y. Liang, C. Macziewski, and J. Cui, "Review of fe-6.5 wt % si high silicon steel—a promising soft magnetic material for sub-khz application," *Journal of Magnetism and Magnetic Materials*, vol. 481, pp. 234–250, 2019.
- [145] R. M. Bozorth, Ferromagnetism. Wiley-IEEE Press, 1993.
- [146] C. Gheorghies and A. Doniga, "Evolution of texture in grain oriented silicon steels," Journal of Iron and Steel Research, International, vol. 16, no. 4, pp. 78–83, 2009.
- [147] P. Beckley, *Electrical steels for rotating machines*. No. 37 in IEE Energy Engineering Series, IET, 2002.
- [148] M. Garibaldi, I. Ashcroft, M. Simonelli, and R. Hague, "Metallurgy of high-silicon steel parts produced using selective laser melting," Acta Materialia, vol. 110, pp. 207–216, 2016.
- [149] A. Du Plessis, E. MacDonald, J. M. Waller, and F. Berto, "10 non-destructive testing of parts produced by laser powder bed fusion," in *Fundamentals of Laser Powder Bed Fusion* of Metals (I. Yadroitsev, I. Yadroitsava, A. du Plessis, and E. MacDonald, eds.), Additive Manufacturing Materials and Technologies, pp. 277–300, Elsevier, 2021.
- [150] B. Foster, E. Reutzel, A. Nassar, B. Hall, S. Brown, and C. Dickman, "Optical, layerwise monitoring of powder bed fusion," in 2014 International Solid Freeform Fabrication Symposium, University of Texas at Austin, 2015.
- [151] P. Avrampos and G.-C. Vosniakos, "A review of powder deposition in additive manufacturing by powder bed fusion," *Journal of Manufacturing Processes*, vol. 74, pp. 332–352, 2022.
- [152] M. Kalms, R. Narita, C. Thomy, F. Vollertsen, and R. B. Bergmann, "New approach to evaluate 3d laser printed parts in powder bed fusion-based additive manufacturing in-line within closed space," *Additive Manufacturing*, vol. 26, pp. 161–165, 2019.

- [153] S. Berumen, F. Bechmann, S. Lindner, J. P. Kruth, and T. Craeghs, "Quality control of laser- and powder bed-based additive manufacturing (am) technologies," *Laser Assisted Net Shape Engineering 6, Proceedings of the Lane 2010, Part 2*, vol. 5, pp. 617–622, 2010.
- [154] P. Bidare, I. Bitharas, R. Ward, M. Attallah, and A. Moore, "Fluid and particle dynamics in laser powder bed fusion," Acta Materialia, vol. 142, pp. 107 – 120, 2018.
- [155] P. Bidare, R. Maier, R. Beck, J. Shephard, and A. Moore, "An open-architecture metal powder bed fusion system for in-situ process measurements," *Additive Manufacturing*, vol. 16, pp. 177 – 185, 2017.
- [156] M. Wittwer and M. Seita, "A machine learning approach to map crystal orientation by optical microscopy," *npj Computational Materials*, vol. 8, no. 1, p. 8, 2022.
- [157] B. Gaskey, L. Hendl, X. Wang, and M. Seita, "Optical characterization of grain orientation in crystalline materials," *Acta Materialia*, vol. 194, pp. 558–564, 2020.
- [158] F. Wang, H. Mao, D. Zhang, X. Zhao, and Y. Shen, "Online study of cracks during laser cladding process based on acoustic emission technique and finite element analysis," *Applied Surface Science*, vol. 255, no. 5, Part 2, pp. 3267–3275, 2008.
- [159] M. Strantza, D. G. Aggelis, D. De Baere, P. Guillaume, and D. Van Hemelrijck, "Evaluation of shm system produced by additive manufacturing via acoustic emission and other ndt methods," *Sensors*, vol. 15, no. 10, pp. 26709–26725, 2015.
- [160] B. Pouet, A. Wartelle, and S. Breugnot, "Multi-detector receiver for laser ultrasonic measurement on the run," *Nondestructive Testing and Evaluation*, vol. 26, no. 3-4, pp. 253–266, 2011.
- [161] L. Koester, H. Taheri, T. Bigelow, and L. Bond, "Nondestructive testing for metal parts fabricated using powder based additive manufacturing," *Materials evaluation*, vol. 76, 2018.
- [162] S. Shevchik, C. Kenel, C. Leinenbach, and K. Wasmer, "Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks," *Additive Manufacturing*, 2017.
- [163] N. Eschner, L. Weiser, B. Häfner, and G. Lanza, "Classification of specimen density in laser powder bed fusion (l-pbf) using in-process structure-borne acoustic process emissions," *Additive Manufacturing*, vol. 34, p. 101324, 2020.

- [164] I. Eriksson, P. Gren, J. Powell, and A. F. Kaplan, "New high-speed photography technique for observation of fluid flow in laser welding," *Optical Engineering*, vol. 49, no. 10, p. 100503, 2010.
- [165] F. Herzog, F. Bechmann, S. Berumen, J. Kruth, and T. Craeghs, "Inventors method for producing a three-dimensional component," 2013.
- [166] D. S. Egan and D. P. Dowling, "Influence of process parameters on the correlation between in-situ process monitoring data and the mechanical properties of ti-6al-4v non-stochastic cellular structures," *Additive Manufacturing*, vol. 30, p. 100890, 2019.
- [167] S. Clijsters, T. Craeghs, S. Buls, K. Kempen, and J. P. Kruth, "In situ quality control of the selective laser melting process using a high-speed, real-time melt pool monitoring system," *The International Journal of Advanced Manufacturing Technology*, vol. 75, no. 5-8, pp. 1089– 1101, 2014.
- [168] M. Grasso and B. M. Colosimo, "Process defects and in situ monitoring methods in metal powder bed fusion: a review," *Measurement Science and Technology*, vol. 28, 2017.
- [169] C. Mandache, "Overview of non-destructive evaluation techniques for metal-based additive manufacturing," *Materials Science and Technology*, vol. 35, pp. 1007–1015, June 2019.
- [170] P. A. Hooper, "Melt pool temperature and cooling rates in laser powder bed fusion," Additive Manufacturing, vol. 22, pp. 548–559, 2018.
- [171] Y. Chivel, "Optical in-process temperature monitoring of selective laser melting," *Physics Proceedia*, vol. 41, pp. 904–910, 2013. Lasers in Manufacturing (LiM 2013).
- [172] T. Furumoto, T. Ueda, M. R. Alkahari, and A. Hosokawa, "Investigation of laser consolidation process for metal powder by two-color pyrometer and high-speed video camera," *CIRP Annals*, vol. 62, no. 1, pp. 223–226, 2013.
- [173] J. Raplee, A. Plotkowski, M. M. Kirka, R. Dinwiddie, A. Okello, R. R. Dehoff, and S. S. Babu, "Thermographic microstructure monitoring in electron beam additive manufacturing," *Scientific Reports*, vol. 7, pp. 43554–, Mar. 2017.
- [174] H. Krauss, C. Eschey, and M. Zaeh, "Thermography for monitoring the selective laser melting process," in *Proceedings of the Solid Freeform Fabrication Symposium*, 2012.
- [175] D. Cerniglia and N. Montinaro, "Defect detection in additively manufactured components: laser ultrasound and laser thermography comparison," *Procedia Structural Integrity*, vol. 8, pp. 154–162, 2018.

- [176] A. Thompson, I. Maskery, and R. K. Leach, "X-ray computed tomography for additive manufacturing: a review," *Measurement Science and Technology*, vol. 27, p. 072001, jun 2016.
- [177] A. B. Spierings, M. Schneider, and R. Eggenberger, "Comparison of density measurement techniques for additive manufactured metallic parts," *Rapid Prototyping Journal*, vol. 17, pp. 380–386, Jan. 2011.
- [178] H. D. Carlton, A. Haboub, G. F. Gallegos, D. Y. Parkinson, and A. A. MacDowell, "Damage evolution and failure mechanisms in additively manufactured stainless steel," *Materials Science and Engineering: A*, vol. 651, pp. 406–414, 2016.
- [179] F. Léonard, S. Tammas-Williams, P. B. Prangnell, I. Todd, and P. J. Withers, "Assessment by x-ray ct of the effects of geometry and build direction on defects in titanium alm parts," in *Conference on industrial computed tomography (iCT)*, pp. 19–21, 2012.
- [180] M. Montazeri, R. Yavari, P. Rao, and P. Boulware, "In-Process Monitoring of Material Cross-Contamination Defects in Laser Powder Bed Fusion," *Journal of Manufacturing Science and Engineering*, vol. 140, 07 2018. 111001.
- [181] S. Siddique, M. Imran, M. Rauer, M. Kaloudis, E. Wycisk, C. Emmelmann, and F. Walther, "Computed tomography for characterization of fatigue performance of selective laser melted parts," *Materials & Design*, vol. 83, pp. 661–669, 2015.
- [182] J. Kastner, B. Harrer, and H. P. Degischer, "High resolution cone beam x-ray computed tomography of 3d-microstructures of cast al-alloys," *Materials Characterization*, vol. 62, no. 1, pp. 99–107, 2011.
- [183] S. A. McDonald, T. L. Burnett, J. Donoghue, N. Gueninchault, H. Bale, C. Holzner, E. M. Lauridsen, and P. J. Withers, "Tracking polycrystal evolution non-destructively in 3d by laboratory x-ray diffraction contrast tomography," *Materials Characterization*, vol. 172, p. 110814, 2021.
- [184] J. Waller, B. Parker, K. Hodges, J. Walker, and E. Burke, "Nondestructive evaluation of additive manufacturing state-of-the-discipline report," tech. rep., NASA, 11 2014.
- [185] E. Todorov, P. Boulware, and K. Gaah, "Demonstration of array eddy current technology for real-time monitoring of laser powder bed fusion additive manufacturing process," in Nondestructive Characterization and Monitoring of Advanced Materials, Aerospace, Civil Infrastructure, and Transportation XII (P. J. Shull, ed.), vol. 10599, pp. 190 – 201, International Society for Optics and Photonics, SPIE, 2018.

- [186] D. Eisenbarth, P. Stoll, C. Klahn, T. B. Heinis, M. Meboldt, and K. Wegener, "Unique coding for authentication and anti-counterfeiting by controlled and random process variation in l-pbf and l-ded," *Additive Manufacturing*, vol. 35, p. 101298, 2020.
- [187] V. S. Tong, A. J. Knowles, D. Dye, and T. B. Britton, "Rapid electron backscatter diffraction mapping: Painting by numbers," *Materials Characterization*, vol. 147, pp. 271–279, 2019.
- [188] B. Winiarski, A. Gholinia, K. Mingard, M. Gee, G. Thompson, and P. Withers, "Correction of artefacts associated with large area ebsd," *Ultramicroscopy*, vol. 226, p. 113315, 2021.
- [189] D. Rowenhorst, "Removing imaging distortions through automatic stitching of ebsd mosaics," *Microscopy and Microanalysis*, vol. 19, no. S2, p. 840–841, 2013.
- [190] M. Glavicic, P. Kobryn, T. Bieler, and S. Semiatin, "A method to determine the orientation of the high-temperature beta phase from measured ebsd data for the low-temperature alpha phase in ti-6al-4v," *Materials Science and Engineering: A*, vol. 346, no. 1, pp. 50 – 59, 2003.
- [191] M. Kral, M. Mangan, G. Spanos, and R. Rosenberg, "Three-dimensional analysis of microstructures," *Materials Characterization*, vol. 45, no. 1, pp. 17–23, 2000.
- [192] M. P. Echlin, A. Mottura, C. J. Torbet, and T. M. Pollock, "A new tribeam system for three-dimensional multimodal materials analysis," *Review of Scientific Instruments*, vol. 83, no. 2, p. 023701, 2012.
- [193] A. T. Polonsky, N. Raghavan, M. P. Echlin, M. M. Kirka, R. R. Dehoff, and T. M. Pollock, "3d characterization of the columnar-to-equiaxed transition in additively manufactured inconel 718," in *Superalloys 2020*, pp. 990–1002, Springer, 2020.
- [194] Y. Chen, S. J. Clark, C. L. A. Leung, L. Sinclair, S. Marussi, M. P. Olbinado, E. Boller, A. Rack, I. Todd, and P. D. Lee, "In-situ synchrotron imaging of keyhole mode multi-layer laser powder bed fusion additive manufacturing," *Applied Materials Today*, vol. 20, p. 100650, 2020.
- [195] C. Zhao, K. Fezzaa, R. W. Cunningham, H. Wen, F. De Carlo, L. Chen, A. D. Rollett, and T. Sun, "Real-time monitoring of laser powder bed fusion process using high-speed x-ray imaging and diffraction," *Scientific reports*, vol. 7, no. 1, pp. 1–11, 2017.
- [196] H. Rieder, M. Spies, J. Bamberg, and B. Henkel, "On and offline ultrasonic inspection of additively manufactured components," 19th World Conference on Non-Destructive Testing, 2016.

- [197] L. Bai, A. Velichko, A. T. Clare, P. Dryburgh, D. Pieris, and B. W. Drinkwater, "The effect of distortion models on characterisation of real defects using ultrasonic arrays," NDT & E International, vol. 113, p. 102263, 2020.
- [198] Y. Javadi, C. N. MacLeod, S. G. Pierce, A. Gachagan, D. Lines, C. Mineo, J. Ding, S. Williams, M. Vasilev, E. Mohseni, and R. Su, "Ultrasonic phased array inspection of a wire + arc additive manufactured (waam) sample with intentionally embedded defects," *Additive Manufacturing*, vol. 29, p. 100806, 2019.
- [199] Y. Javadi, C. N. Macleod, S. G. Pierce, A. Gachagan, W. Kerr, J. Ding, S. Williams, M. Vasilev, R. Su, and J. Mineo, C. andDziewierz, "Ultrasonic phased array inspection of wire + arc additive manufacture samples using conventional and total focusing method imaging approaches," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 61, no. 3, pp. 144–148, 2019.
- [200] R. Zimermann, E. Mohseni, D. Lines, R. K. Vithanage, C. N. MacLeod, S. G. Pierce, A. Gachagan, Y. Javadi, S. Williams, and J. Ding, "Multi-layer ultrasonic imaging of as-built wire + arc additive manufactured components," *Additive Manufacturing*, vol. 48, p. 102398, 2021.
- [201] N. Huang, O. J. Cook, R. L. Smithson, C. M. Kube, A. P. Argüelles, and A. M. Beese, "Use of ultrasound to identify microstructure-property relationships in 316 stainless steel fabricated with binder jet additive manufacturing," *Additive Manufacturing*, vol. 51, p. 102591, 2022.
- [202] F. E. Stanke and G. S. Kino, "A unified theory for elastic wave propagation in polycrystalline materials," *The Journal of the Acoustical Society of America*, vol. 75, no. 3, pp. 665–681, 1984.
- [203] C. M. Kube, "Iterative solution to bulk wave propagation in polycrystalline materials," The Journal of the Acoustical Society of America, vol. 141, pp. 1804–1811, Mar. 2017.
- [204] B. Lan, T. Ben Britton, T.-S. Jun, W. Gan, M. Hofmann, F. P. E. Dunne, and M. J. S. Lowe, "Direct volumetric measurement of crystallographic texture using acoustic waves," *Acta Materialia*, vol. 159, pp. 384 – 394, 2018.
- [205] K. M. M. Tant, E. Galetti, A. J. Mulholland, A. Curtis, and A. Gachagan, "Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing," *Inverse Problems in Science and Engineering*, vol. 0, no. 0, pp. 1–25, 2020.

- [206] M. Klein and J. Sears, "Laser ultrasonic inspection of laser cladded 316ss and ti-6-4," in International Congress on Applications of Lasers & Electro-Optics, p. 1006, Laser Institute of America, 2004.
- [207] J. Yu, D. Zhang, H. Li, C. Song, X. Zhou, S. Shen, G. Zhang, Y. Yang, and H. Wang, "Detection of internal holes in additive manufactured ti-6al-4v part using laser ultrasonic testing," *Applied Sciences*, vol. 10, no. 1, p. 365, 2020.
- [208] G. Davis, R. Nagarajah, S. Palanisamy, R. A. R. Rashid, P. Rajagopal, and K. Balasubramaniam, "Laser ultrasonic inspection of additive manufactured components," *The International Journal of Advanced Manufacturing Technology*, vol. 102, no. 5, pp. 2571–2579, 2019.
- [209] S. Everton, P. Dickens, C. Tuck, and B. Dutton, "Using laser ultrasound to detect subsurface defects in metal laser powder bed fusion components," *The Journal of The Minerals, Metals & Materials Society*, vol. 70, pp. 378–383, Mar 2018.
- [210] Y. Fan, S. Dixon, R. Edwards, and X. Jian, "Ultrasonic surface wave propagation and interaction with surface defects on rail track head," NDT & E International, vol. 40, no. 6, pp. 471–477, 2007.
- [211] J. Cooper, R. Crosbie, R. Dewhurst, A. McKie, and S. 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.
- [212] J. Zhang, J. Wu, X. Zhao, S. Yuan, G. Ma, J. Li, T. Dai, H. Chen, B. Yang, and H. Ding, "Laser ultrasonic imaging for defect detection on metal additive manufacturing components with rough surfaces," *Appl. Opt.*, vol. 59, pp. 10380–10388, Nov 2020.
- [213] T. Murray, J. Deaton, and J. Wagner, "Experimental evaluation of enhanced generation of ultrasonic waves using an array of laser sources," *Ultrasonics*, vol. 34, no. 1, pp. 69–77, 1996.
- [214] C. Pei, K. Demachi, T. Fukuchi, K. Koyama, and M. Uesaka, "Cracks measurement using fiber-phased array laser ultrasound generation," *Journal of Applied Physics*, vol. 113, p. 163101, Apr. 2013.
- [215] W. P. Stratoudaki T, Clark M, "Laser induced ultrasonic phased array using full matrix capture data acquisition and total focusing method," *Optics Express*, vol. 24, no. 19, 2016.
- [216] D. Pieris, T. Stratoudaki, Y. Javadi, P. Lukacs, S. Catchpole-Smith, P. D. Wilcox, A. Clare, and M. Clark, "Laser induced phased arrays (lipa) to detect nested features in additively manufactured components," *Materials & Design*, vol. 187, p. 108412, 2020.

- [217] Adaptive Data Acquisition for Fast Ultrasonic Imaging Using Laser Induced Phased Arrays, vol. 2021 48th Annual Review of Progress in Quantitative Nondestructive Evaluation of Quantitative Nondestructive Evaluation, 07 2021. V001T18A003.
- [218] F. Dong, X. Wang, Q. Yang, A. Yin, and X. Xu, "Directional dependence of aluminum grain size measurement by laser-ultrasonic technique," *Materials Characterization*, vol. 129, pp. 114–120, 2017.
- [219] C. Scruby, R. Smith, and B. Moss, "Microstructural monitoring by laser-ultrasonic attenuation and forward scattering," NDT international, vol. 19, no. 5, pp. 307–313, 1986.
- M. Malmström, A. Jansson, B. Hutchinson, J. Lönnqvist, L. Gillgren, L. Bäcke, H. Sollander,
 M. Bärwald, S. Hochhard, and P. Lundin, "Laser-ultrasound-based grain size gauge for the hot strip mill," *Applied Sciences*, vol. 12, no. 19, 2022.
- [221] Y. Zhan, C. Liu, J. Zhang, G. Mo, and C. Liu, "Measurement of residual stress in laser additive manufacturing tc4 titanium alloy with the laser ultrasonic technique," *Materials Science and Engineering: A*, vol. 762, p. 138093, 2019.
- [222] Y. Ma, X. Hu, Z. Hu, Z. Sheng, S. Ma, Y. Chu, Q. Wan, W. Luo, and L. Guo, "Simultaneous compositional and grain size measurements using laser opto-ultrasonic dual detection for additive manufacturing," *Materials*, vol. 13, no. 10, 2020.
- [223] Y. Ma, Z. Hu, Y. Tang, S. Ma, Y. Chu, X. Li, W. Luo, L. Guo, X. Zeng, and Y. Lu, "Laser opto-ultrasonic dual detection for simultaneous compositional, structural, and stress analyses for wire + arc additive manufacturing," *Additive Manufacturing*, vol. 31, p. 100956, 2020.
- [224] T. Thréard, E. de Lima Savi, S. Avanesyan, N. Chigarev, Z. Hua, V. Tournat, V. E. Gusev, D. H. Hurley, and S. Raetz, "Photoacoustic 3-d imaging of polycrystalline microstructure improved with transverse acoustic waves," *Photoacoustics*, vol. 23, p. 100286, 2021.
- [225] X. Du and J.-C. Zhao, "Facile measurement of single-crystal elastic constants from polycrystalline samples," *npj Computational Materials*, vol. 3, no. 1, p. 17, 2017.
- [226] R. F. S. Hearmon, "The elastic constants of crystals and other anisotropic materials," Landolt-Bornstein Tables, III/18, vol. 1154, 1984.
- [227] G. R. Speich, A. J. Schwoeble, and W. C. Leslie, "Elastic constants of binary iron-base alloys," *Metallurgical Transactions*, vol. 3, no. 8, pp. 2031–2037, 1972.
- [228] M. Born, "On the stability of crystal lattices. i," Mathematical Proceedings of the Cambridge Philosophical Society, vol. 36, no. 2, p. 160–172, 1940.

- [229] G. A. Alers and J. R. Neighbours, "Crystal stability and elastic constants," Journal of Applied Physics, vol. 28, no. 12, pp. 1514–1514, 1957.
- [230] A. Roy, D. W. Branch, D. S. Jensen, and C. M. Kube, "propSym: a tool to establish relationships between property constants for material property tensors of any order," Journal of Applied Crystallography, vol. 54, pp. 1894–1901, Dec 2021.
- [231] R. W. Fenn, W. R. Hibbard, and H. A. Lepper, "The elastic coefficients of single crystals of alpha brass," *JOM*, vol. 2, no. 1, pp. 175–181, 1950.
- [232] Y. Shirakawa and K. Numakura, "Elasticity constants of singlecrystals of nickel-copper alloys," Science reports of the Research Institutes, 1957.
- [233] K. Honda and Y. Shirakawa, "On young's modulus of elasticity of single crystals of nickel and cobalt," Science reports of the Research Institutes, Tohoku University, 1949.
- [234] S. Pathak and S. R. Kalidindi, "Spherical nanoindentation stress-strain curves," Materials Science and Engineering: R: Reports, vol. 91, pp. 1 – 36, 2015.
- [235] A. Vinckier and G. Semenza, "Measuring elasticity of biological materials by atomic force microscopy," FEBS Letters, vol. 430, no. 1, pp. 12 – 16, 1998.
- [236] J. W. Flowers, K. C. O'Brien, and P. C. McEleney, "Elastic constants of alpha-titanium single crystals at 25 °C," *Journal of the Less Common Metals*, vol. 7, no. 5, pp. 393 – 395, 1964.
- [237] J. R. Neighbours, F. W. Bratten, and C. S. Smith, "The elastic constants of nickel," *Journal of Applied Physics*, vol. 23, no. 4, pp. 389–393, 1952.
- [238] T. Grabec, P. Sedlák, K. Zoubková, M. Ševčík, M. Janovská, P. Stoklasová, and H. Seiner, "Evolution of elastic constants of the niti shape memory alloy during a stress-induced martensitic transformation," Acta Materialia, vol. 208, p. 116718, 2021.
- [239] Y. Xu, T. Aizawa, and J. Kihara, "Simultaneous determination of elastic constants and crystallographic orientation in coarse-grained nickel by acoustic spectro-microscopy," *Materials Transactions*, JIM, vol. 38, no. 6, pp. 536–545, 1997.
- [240] J.-Y. Kim and S. I. Rokhlin, "Determination of elastic constants of generally anisotropic inclined lamellar structure using line-focus acoustic microscopy," *The Journal of the Acoustical Society of America*, vol. 126, no. 6, pp. 2998–3007, 2009.

- [241] K. Zoubková, P. Stoklasová, T. Grabec, P. Sedlák, and H. Seiner, "Transient grating spectroscopy for complete elastic anisotropy: beyond the measurement of surface acoustic waves," in 2021 IEEE International Ultrasonics Symposium (IUS), pp. 1–3, IEEE, 2021.
- [242] D. Gasteau, N. Chigarev, L. Ducousso-Ganjehi, V. E. Gusev, F. Jenson, P. Calmon, and V. Tournat, "Single crystal elastic constants evaluated with surface acoustic waves generated and detected by lasers within polycrystalline steel samples," *Journal of Applied Physics*, vol. 119, no. 4, p. 043103, 2016.
- [243] M. Tane, K. Yamori, T. Sekino, and T. Mayama, "Impact of grain shape on the micromechanics-based extraction of single-crystalline elastic constants from polycrystalline samples with crystallographic texture," Acta Materialia, vol. 122, pp. 236–251, 2017.
- [244] R. G. Leisure and F. A. Willis, "Resonant ultrasound spectroscopy," Journal of Physics: Condensed Matter, vol. 9, pp. 6001–6029, jul 1997.
- [245] B. Bales, L. Petzold, B. R. Goodlet, W. C. Lenthe, and T. M. Pollock, "Bayesian inference of elastic properties with resonant ultrasound spectroscopy," *The Journal of the Acoustical Society of America*, vol. 143, no. 1, pp. 71–83, 2018.
- [246] J. Rossin, P. Leser, K. Pusch, C. Frey, S. P. Murray, C. J. Torbet, S. Smith, S. Daly, and T. M. Pollock, "Bayesian inference of elastic constants and texture coefficients in additively manufactured cobalt-nickel superalloys using resonant ultrasound spectroscopy," Acta Materialia, vol. 220, p. 117287, 2021.
- [247] B. R. Goodlet, S. P. Murray, B. Bales, J. Rossin, C. J. Torbet, and T. M. Pollock, "Temperature dependence of single crystal elastic constants in a coni-base alloy: A new methodology," *Materials Science and Engineering: A*, vol. 803, p. 140507, 2021.
- [248] S. D. Sharples, M. Clark, I. Collison, and M. G. Somekh, "Adaptive acoustic imaging using aberration correction in difficult materials," *Insight-Non-Destructive Testing and Condition Monitoring*, vol. 47, no. 2, pp. 78–80, 2005.
- [249] 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.
- [250] M. Clark and S. Sharples, 2009.

- [251] W. Li, S. D. Sharples, M. Clark, and M. G. Somekh, "Frequency Spectrum Spatially Resolved Acoustic Spectroscopy for Microstructure Imaging," *Journal of Physics: Conference Series*, vol. 278, no. 012026, pp. 1 – 4, 2011.
- [252] W. Li, S. D. Sharples, R. J. Smith, M. Clark, and M. G. Somekh, "Determination of crystallographic orientation of large grain metals with surface acoustic waves," *Journal of the Acoustical Society of America*, vol. 132, pp. 738–745, 2012.
- [253] R. J. Smith, W. Li, J. Coulson, M. Clark, S. D. Sharples, and M. G. Somekh, "Spatially resolved acoustic spectroscopy for rapid imaging of material microstructure and grain orientation," *Measurement Science and Technology*, vol. 25, no. 055902, pp. 1 – 11, 2014.
- [254] W. Mark, A.F.and Li, S. D. Sharples, and P. J. Withers, "Comparison of grain to grain orientation and stiffness mapping by spatially resolved acoustic spectroscopy and ebsd," *Journal of Microscopy*, vol. 627, no. 1, pp. 89–97, 2017.
- [255] R. Patel, W. Li, R. J. Smith, S. D. Sharples, and M. Clark, "Orientation imaging of macrosized polysilicon grains on wafers using spatially resolved acoustic spectroscopy," *Scripta Materialia*, vol. 140, pp. 67 – 70, 2017.
- [256] R. Patel, S. D. Sharples, M. Clark, M. G. Somekh, and W. Li, "Single pixel camera methodologies for spatially resolved acoustic spectroscopy," *Applied Physics Letters*, vol. 118, no. 5, p. 051102, 2021.
- [257] D. Clark, S. D. Sharples, and D. C. Wright, "Development of online inspection for additive manufacturing products," *Insight - Non-Destructive Testing and Condition Monitoring*, vol. 53, no. 11, pp. 610–613(4), 2011.
- [258] R. Patel, M. Hirsch, P. Dryburgh, D. Pieris, S. Achamfuo-Yeboah, R. Smith, R. Light, S. Sharples, A. Clare, and M. Clark, "Imaging material texture of as-deposited selective laser melted parts using spatially resolved acoustic spectroscopy," *Applied Sciences*, vol. 8, no. 10, p. 1991, 2018.
- [259] L. Rayleigh, "On waves propogated along the plance surface of an elastic solid," Proceedings of the London Mathematical Society, vol. 1, no. 1, 1885.
- [260] R. Stoneley, "The Propagation of Surface Waves in an Elastic Medium with Orthorhombic Symmetry," *Geophysical Journal International*, vol. 8, pp. 176–186, 12 1963.

- [261] R. Stoneley, "The propagation of surface elastic waves in a cubic crystal," Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, vol. 232, no. 1191, pp. 447–458, 1955.
- [262] B. Lüthi, Physical acoustics in the solid state, ch. 3, pp. 39–44. Springer, 2005.
- [263] M. Destrade, "The explicit secular equation for surface acoustic waves in monoclinic elastic crystals," *The Journal of the Acoustical Society of America*, vol. 109, no. 4, pp. 1398–1402, 2001.
- [264] T. C. T. Ting, "An Explicit Secular Equation for Surface Waves in an Elastic Material of General Anisotropy," *The Quarterly Journal of Mechanics and Applied Mathematics*, vol. 55, pp. 297–311, 05 2002.
- [265] P. Stoklasová, P. Sedlák, H. S. Landa, and Michal, "Forward and inverse problems for surface acoustic waves in anisotropic media: A Ritz-Rayleigh method based approach," *Ultrasonics*, vol. 56, pp. 381–389, 2015.
- [266] G. W. Farnell, Physical Acoustics Principles and Methods, vol. VI, ch. 3, pp. 109–166. New York and London: Academic Press, 1970.
- [267] M. R. Cherry, S. Sathish, and R. Grandhi, "A numerical method for predicting rayleigh surface wave velocity in anisotropic crystals," *Journal of Computational Physics*, vol. 351, pp. 108 – 120, 2017.
- [268] S. Sharples, M. Clark, R. Smith, R. Ellwood, W. Li, and M. G. Somekh, "Laser ultrasonic microscopy," *Nondestructive Testing and Evaluation*, vol. 26, no. 3-4, pp. 367–384, 2011.
- [269] S. D. Sharples, R. A. Light, S. O. Achamfuo-Yeboah, M. Clark, and M. G. Somekh, "The SKED: speckle knife edge detector," *Journal of Physics: Conference Series*, vol. 520, p. 012004, jun 2014.
- [270] T. Blum, B. Pouet, S. Breugnot, and P. Clémenceau, "Non-destructive testing using multi-channel random-quadrature interferometer," in *AIP Conference Proceedings*, vol. 975, pp. 239–246, American Institute of Physics, 2008.
- [271] A. Maradudin and D. Mills, "The attenuation of rayleigh surface waves by surface roughness," Annals of Physics, vol. 100, no. 1, pp. 262–309, 1976.
- [272] A. Ruiz and P. B. Nagy, "Laser-ultrasonic surface wave dispersion measurements on surfacetreated metals," *Ultrasonics*, vol. 42, no. 1, pp. 665 – 669, 2004.

- [273] J. Bercoff, M. Tanter, and M. Fink, "Supersonic shear imaging: a new technique for soft tissue elasticity mapping," *IEEE transactions on ultrasonics, ferroelectrics, and frequency* control, vol. 51, no. 4, pp. 396–409, 2004.
- [274] J. Matas, J. Bass, Y. Ricard, E. Mattern, and M. Bukowinski, "On the bulk composition of the lower mantle: predictions and limitations from generalized inversion of radial seismic profiles," *Geophysical Journal International*, vol. 170, no. 2, pp. 764–780, 2007.
- [275] R. F. S. Hearmon, "The elastic constants of anisotropic materials II," Advances in Physics, vol. 5, no. 19, pp. 323–382, 1956.
- [276] T. C. T. Ting and D. M. Barnett, "Classifications of surface waves in anisotropic elastic materials," *Wave Motion*, vol. 26, no. 3, pp. 207 – 218, 1997.
- [277] W. Li, J. Coulson, J. W. Aveson, R. J. Smith, M. Clark, M. G. Somekh, and S. D. Sharples, "Orientation characterisation of aerospace materials by spatially resolved acoustic spectroscopy," 3rd International Symposium on Laser Ultrasonics and Advanced Sensing, vol. 520, 2014.
- [278] J. Sermeus, Thermoelastic characterization of thin functional films on anisotropic substrates. PhD thesis, KU Leuven, 2014.
- [279] W. Li, S. D. Sharples, R. J. Smith, M. Clark, and M. G. Somekh, "Determination of crystallographic orientation of large grain metals with surface acoustic waves," J Acoust Soc Am, vol. 132, no. 2, pp. 738–45, 2012.
- [280] S. Kumar and S. K. Kurtz, "Simulation of material microstructure using a 3d voronoi tesselation: Calculation of effective thermal expansion coefficient of polycrystalline materials," *Acta metallurgica et materialia*, vol. 42, no. 12, pp. 3917–3927, 1994.
- [281] A. Okabe, B. Boots, and K. Sugihara, Spatial Tessellations: concepts and applications of Voronoi diagrams. John Wiley and Sons, second ed., 1998.
- [282] F. Fritzen, T. Böhlke, and E. Schnack, "Periodic three-dimensional mesh generation for crystalline aggregates based on voronoi tessellations," *Computational Mechanics*, vol. 43, no. 5, pp. 701–713, 2009.
- [283] T. van Nuland, J. van Dommelen, and M. Geers, "An anisotropic voronoi algorithm for generating polycrystalline microstructures with preferred growth directions," *Computational Materials Science*, vol. 186, p. 109947, 2021.

- [284] W. Li, J. Coulson, P. Marrow, R. J. Smith, S. J. Lainé, M. Clark, and S. D. Sharples, "Crystallographic orientation determination of hexagonal structure crystals by laser ultrasonic technique," *Journal of Physics: Conference Series*, vol. 684, no. 1, pp. 012001–, 2016.
- [285] C. M. Kube and J. A. Turner, "Voigt, reuss, hill, and self-consistent techniques for modeling ultrasonic scattering," AIP Conference Proceedings, vol. 1650, pp. 926–934, Mar. 2015.
- [286] R. Acharya and S. Das, "Additive manufacturing of in100 superalloy through scanning laser epitaxy for turbine engine hot-section component repair: Process development, modeling, microstructural characterization, and process control," *Metallurgical and Materials Transactions A*, vol. 46, pp. 3864–3875, Sep 2015.
- [287] J. L. W. Warwick, J. Coakley, S. L. Raghunathan, R. J. Talling, and D. Dye, "Effect of texture on load partitioning in Ti-6Al-4V," *Acta Materialia*, vol. 60, no. 10, pp. 4117–4127, 2012.
- [288] J. Bernier, J.-S. Park, A. Pilchak, M. Glavicic, and M. Miller, "Measuring stress distributions in ti-6al-4v using synchrotron x-ray diffraction," *Metallurgical and Materials Transactions* A, vol. 39, no. 13, pp. 3120–3133, 2008.
- [289] D. Qiu, R. Shi, D. Zhang, W. Lu, and Y. Wang, "Variant selection by dislocations during α precipitation in α/β titanium alloys," Acta Materialia, vol. 88, pp. 218–231, 2015.
- [290] B. Pouet, S. Breugnot, and P. Clémenceau, "Robust laser ultrasonic interferometer based on random quadrature demodulation," in *AIP Conference Proceedings*, vol. 820, pp. 233–239, AIP, 2006.
- [291] J. Ke, C. Duan, W. Yi, and C. Yan, "Application of an adaptive two-wave mixing interferometer for detection of surface defects," in 2016 Progress in Electromagnetic Research Symposium (PIERS), pp. 2142–2146, 2016.
- [292] I. Bogatyreva, V. Dan'ilchenko, and L. Pospelov, "Theory of fabry-perot interferometer with statistically rough mirrors," *Radiophysics and Quantum Electronics*, vol. 26, no. 1, pp. 49–54, 1983.
- [293] B. Pavlakovic, M. Lowe, D. Alleyne, and P. Cawley, "Disperse: A general purpose program for creating dispersion curves," in *Review of progress in quantitative nondestructive evaluation*, pp. 185–192, Springer, 1997.
- [294] K. Salama and G. A. Alers, "The composition dependence of the third-order elastic constants of the Cu–Ni system," *physica status solidi* (a), vol. 41, no. 1, pp. 241 – 247, 1977.

- [295] L. N. Carter, C. Martin, P. J. Withers, and M. M. Attallah, "The influence of the laser scan strategy on grain structure and cracking behaviour in slm powder-bed fabricated nickel superalloy," *Journal of Alloys and Compounds*, vol. 615, pp. 338–347, 2014.
- [296] A. Machová and S. Kadečková, "Elastic constants of iron-silicon alloy single crystals," *Czechoslovak Journal of Physics B*, vol. 27, no. 5, pp. 555–563, 1977.
- [297] J. J. Adams, D. Agosta, R. Leisure, and H. Ledbetter, "Elastic constants of monocrystal iron from 3 to 500 k," *Journal of applied physics*, vol. 100, no. 11, p. 113530, 2006.
- [298] R. Hearmon, "The elastic constants of anisotropic materials," *Reviews of Modern Physics*, vol. 18, no. 3, p. 409, 1946.
- [299] H. Ogi, S. Kai, H. M. Ledbetter, R. Tarumi, M. Hirao, and K. Takashima, "Titanium's hightemperature elastic constants through the hcp-bcc phase transformation," Acta Materialia, vol. 52, pp. 2075–2080, 2004.
- [300] E. S. Fisher and C. J. Renken, "Single-Crystal Elastic Moduli and the hcp → bcc Transformation in Ti, Zr, and Hf," *Physical Review*, vol. 135, pp. A482–A494, Jul 1964.
- [301] H. M. Ledbetter, "Elastic properties of zinc: A compilation and a review," Journal of Physical and Chemical Reference Data, vol. 6, no. 4, pp. 1181–1203, 1977.
- [302] L. J. Slutsky and C. W. Garland, "Elastic constants of magnesium from 4.2 K to 300 K," *Physical Review*, vol. 107, no. 4, p. 972, 1957.
- [303] B. D. Pandey, D. Singh, and P. Yadawa, "Ultrasonic study of osmium and ruthenium," *Platinum Metals Review*, vol. 53, no. 2, pp. 91–97, 2009.
- [304] W. Li, J. Coulson, R. J. Smith, S. J. Lainé, and S. D. Sharples, "Crystallographic orientation determination of hexagonal structure crystals by laser ultrasonic technique," *Journal of Physics: Conference Series*, vol. 684, no. 1, p. 012001, 2016.
- [305] W. L. Finlay, J. Resketo, and M. B. Vordahl, "Optical metallography of titanium," Industrial & Engineering Chemistry, vol. 42, no. 2, pp. 218–222, 1950.
- [306] S. J. Laine and K. M. Knowles, "1124 deformation twinning in commercial purity titanium at room temperature," *Philosophical Magazine*, vol. 95, no. 20, pp. 2153–2166, 2015.
- [307] P. Dryburgh, R. J. Smith, P. Marrow, S. J. Lainé, S. D. Sharples, M. Clark, and W. Li, "Determining the crystallographic orientation of hexagonal crystal structure materials with surface acoustic wave velocity measurements," *Ultrasonics*, vol. 108, p. 106171, 2020.

- [308] W. Li, J. Coulson, J. W. Aveson, R. J. Smith, M. Clark, M. G. Somekh, and S. D. Sharples, "Orientation characterisation of aerospace materials by spatially resolved acoustic spectroscopy," *Journal of Physics: Conference Series*, vol. 520, no. 012017, pp. 1 – 4, 2014.
- [309] G. M. Amulele, A. G. Every, and S. C. Yates, "Acoustic microscopy wall thickness measurements on nickel based superalloy gas turbine blades," in *Review of Progress in Quantitative Nondestructive Evaluation*, pp. 2001–2007, Springer, 1999.
- [310] M. Mendik, S. Sathish, A. Kulik, G. Gremaud, and P. Wachter, "Surface acoustic wave studies on single-crystal nickel using Brillouin scattering and scanning acoustic microscope," *Journal of Applied Physics*, vol. 71, no. 6, pp. 2830–2834, 1992.
- [311] A. A. Maznev, A. M. Lomonosov, P. Hess, and A. A. Kolomenskii, "Anisotropic effects in surface acoustic wave propagation from a point source in a crystal," *The European Physical Journal B - Condensed Matter and Complex Systems*, vol. 35, no. 3, pp. 429–439, 2003.
- [312] D. C. Gazis, R. Herman, and R. F. Wallis, "Surface elastic waves in cubic crystals," *Physical Review*, vol. 119, no. 2, p. 533, 1960.
- [313] Y. P. Varshni, "Temperature dependence of the elastic constants," *Physical Review B*, vol. 2, no. 10, p. 3952, 1970.
- [314] A. Epishin, B. Fedelich, M. Finn, G. Künecke, B. Rehmer, G. Nolze, C. Leistner, N. Petrushin, and I. Svetlov, "Investigation of elastic properties of the single-crystal nickelbase superalloy CMSX-4 in the temperature interval between room temperature and 1300 °C," Crystals, vol. 11, no. 2, 2021.
- [315] M. Ashby, Material selection in mechanical design. Butterworth-Heinemann, third edition ed., 1999.
- [316] P. Alexander, S. Allen, and D. Dutta, "Part orientation and build cost determination in layered manufacturing," *Computer-Aided Design*, vol. 30, no. 5, pp. 343–356, 1998.
- [317] M. Baumers, P. Dickens, C. Tuck, and R. Hague, "The cost of additive manufacturing: machine productivity, economies of scale and technology-push," *Technological Forecasting* and Social Change, vol. 102, pp. 193 – 201, 2016.
- [318] J. Ding, M. Baumers, E. A. Clark, and R. D. Wildman, "The economics of additive manufacturing: Towards a general cost model including process failure," *International Journal of Production Economics*, vol. 237, p. 108087, 2021.

- [319] M. Baumers and M. Holweg, "On the economics of additive manufacturing: Experimental findings," J Oper Manag, vol. 65, no. 8, pp. 794–809, 2019.
- [320] B. M. Colosimo, S. Cavalli, and M. Grasso, "A cost model for the economic evaluation of insitu monitoring tools in metal additive manufacturing," *International Journal of Production Economics*, vol. 223, p. 107532, 2020.
- [321] M. Tomlin and J. Meyer, "Topology optimization of an additive layer manufactured (alm) aerospace part," in *Proceeding of the 7th Altair CAE technology conference*, pp. 1–9, 2011.
- [322] J. Grum and J. M. Slabe, "Possibility of introducing laser surfacing into maintenance service of die-casting dies," *Surface and Coatings Technology*, vol. 180-181, pp. 596–602, 2004.
- [323] E. Yasa, J. Deckers, and J. Kruth, "The investigation of the influence of laser re-melting on density, surface quality and microstructure of selective laser melting parts," *Rapid Prototyp*ing Journal, vol. 17, no. 5, pp. 312–327, 2011.
- [324] T. R. Newton, S. N. Melkote, T. R. Watkins, R. M. Trejo, and L. Reister, "Investigation of the effect of process parameters on the formation and characteristics of recast layer in wireedm of inconel 718," *Materials Science and Engineering a-Structural Materials Properties Microstructure and Processing*, vol. 513-14, pp. 208–215, 2009.
- [325] J. Mireles, S. Ridwan, P. A. Morton, A. Hinojos, and R. B. Wicker, "Analysis and correction of defects within parts fabricated using powder bed fusion technology," *Surface Topography: Metrology and Properties*, vol. 3, no. 3, p. 034002, 2015.
- [326] M. Rahman, W. K. H. Seah, and T. T. Teo, "The machinability of inconel 718," Journal of Materials Processing Technology, vol. 63, no. 1-3, pp. 199–204, 1997.
- [327] J. Jhabvala, E. Boillat, T. Antignac, and R. Glardon, "On the effect of scanning strategies in the selective laser melting process," *Virtual and Physical Prototyping*, vol. 5, no. 2, pp. 99– 109, 2010.
- [328] M. Hirsch, R. Patel, W. Li, G. Guan, R. K. Leach, S. D. Sharples, and A. T. Clare, "Assessing the capability of in-situ nondestructive analysis during layer based additive manufacture," *Additive Manufacturing*, vol. 13, pp. 135–142, 2017.
- [329] S. Rott, A. Ladewig, K. Friedberger, J. Casper, M. Full, and J. H. Schleifenbaum, "Surface roughness in laser powder bed fusion – interdependency of surface orientation and laser incidence," *Additive Manufacturing*, vol. 36, p. 101437, 2020.

- [330] G. Marinelli, F. Martina, S. Ganguly, and S. Williams, "Development of wire + arc additive manufacture for the production of large-scale unalloyed tungsten components," *International Journal of Refractory Metals and Hard Materials*, vol. 82, pp. 329–335, 2019.
- [331] P. Aba-Perea, T. Pirling, P. Withers, J. Kelleher, S. Kabra, and M. Preuss, "Determination of the high temperature elastic properties and diffraction elastic constants of ni-base superalloys," *Materials & Design*, vol. 89, pp. 856–863, 2016.
- [332] H. Ledbetter, H. Ogi, S. Kai, S. Kim, and M. Hirao, "Elastic constants of body-centered-cubic titanium monocrystals," *Journal of applied physics*, vol. 95, no. 9, pp. 4642–4644, 2004.
- [333] H. Taheri, L. W. Koester, T. A. Bigelow, and L. J. Bond, "Thermoelastic finite element modeling of laser generated ultrasound in additive manufacturing materials," in ASNT Annual Conference 2017, pp. 188–198, 2017.
- [334] B. Martin, P. Samimi, and P. Collins, "Engineered, spatially varying isothermal holds: Enabling combinatorial studies of temperature effects, as applied to metastable titanium alloy β-21s," Metallography, Microstructure, and Analysis, vol. 6, no. 3, pp. 216–220, 2017.
- [335] B. W. Martin, Development of a novel, bicombinatorial approach to alloy development, and application to rapid screening of creep resistant titanium alloys. PhD thesis, Iowa State University, 2017.
- [336] F. Margrave, K. Rigas, D. Bradley, and P. Barrowcliffe, "The use of neural networks in ultrasonic flaw detection," *Measurement*, vol. 25, no. 2, pp. 143–154, 1999.
- [337] M. Kesharaju and R. Nagarajah, "Feature selection for neural network based defect classification of ceramic components using high frequency ultrasound," *Ultrasonics*, vol. 62, pp. 271– 277, 2015.
- [338] G. Trtnik, F. Kavčič, and G. Turk, "Prediction of concrete strength using ultrasonic pulse velocity and artificial neural networks," *Ultrasonics*, vol. 49, no. 1, pp. 53–60, 2009.
- [339] A. Moreau, L. Toubal, P. Bocher, M. Humbert, E. Uta, and N. Gey, "Evaluation of macrozone dimensions by ultrasound and ebsd techniques," *Materials Characterization*, vol. 75, pp. 115– 128, 2013.
- [340] S. Rokhlin, G. Sha, J. Li, and A. Pilchak, "Inversion methodology for ultrasonic characterization of polycrystals with clusters of preferentially oriented grains," *Ultrasonics*, vol. 115, p. 106433, 2021.

- [341] R. Ma, Z. Li, W. Gu, L. Wei, Y. Wang, Y. Wang, J. Zuo, and A. Yin, "Measurement of micro-region elastic anisotropy in textured metal by laser-ultrasonics," *AIP Advances*, vol. 10, no. 10, p. 105301, 2020.
- [342] S. Gosha, M. Mills, S. Rokhlin, V. Sinha, W. Soboyejo, and J. Williams, "The evaluation of cold dwell fatigue in ti-6242," tech. rep., Federal Aviation Administration, 2007.
- [343] D. Dye. private communication, Dec. 2021.
- [344] Q. Zhong, X. Tian, X. Huang, C. Huo, and D. Li, "Using feedback control of thermal history to improve quality consistency of parts fabricated via large-scale powder bed fusion," *Additive Manufacturing*, vol. 42, p. 101986, 2021.
- [345] E. Goens, "Elastizitätskonstanten des aluminiumeinkristalls," Annalen der Physik, vol. 409, no. 3, pp. 233–242, 1933.
- [346] J. Vallin, M. Mongy, K. Salama, and O. Beckman, "Elastic constants of aluminum," Journal of Applied Physics, vol. 35, no. 6, pp. 1825–1826, 1964.
- [347] J. Leese and A. E. Lord, "Elastic stiffness coefficients of single-crystal iron from room temperature to 500°c," *Journal of Applied Physics*, vol. 39, pp. 3986–3988, July 1968.
- [348] A. E. Lord and D. N. Beshers, "Elastic stiffness coefficients of iron from 77° to 673°k," Journal of Applied Physics, vol. 36, pp. 1620–1623, May 1965.
- [349] X. Zhang, P. R. Stoddart, J. D. Comins, and A. G. Every, "High-temperature elastic properties of a nickel-based superalloy studied by surface Brillouin scattering," *Journal of Physics: Condensed Matter*, vol. 13, pp. 2281–2294, feb 2001.
- [350] S. C. Prasad and W. A. Wooster, "The determination of the elastic constants of silicon by diffuse X-ray reflexions," Acta Crystallographica, vol. 8, p. 361, Jun 1955.
- [351] H. J. McSkimin and P. Andreatch Jr, "Elastic moduli of silicon vs hydrostatic pressure at 25.0 C and 195.8 C," *Journal of Applied Physics*, vol. 35, no. 7, pp. 2161–2165, 1964.
- [352] J. J. Hall, "Electronic effects in the elastic constants of n-type silicon," *Physical Review*, vol. 161, no. 3, p. 756, 1967.
- [353] P. W. Bridgman, "Some properties of single metal crystals," Proceedings of the National Academy of Sciences, vol. 10, no. 10, pp. 411–415, 1924.

- [354] N. Dumontet, D. Connétable, B. Malard, and B. Viguier, "Elastic properties of the alpha' martensitic phase in the ti-6al-4v alloy obtained by additive manufacturing," *Scripta Materialia*, vol. 167, pp. 115–119, 2019.
- [355] V. Hasija, S. Ghosh, M. J. Mills, and D. S. Joseph, "Deformation and creep modeling in polycrystalline Ti–6Al alloys," Acta Materialia, vol. 51, no. 15, pp. 4533–4549, 2003.
- [356] A. W. Hanson, "Elastic behavior and elastic constants of zinc single crystals," *Physics Reviews*, vol. 45, pp. 324–331, Mar 1934.
- [357] G. A. Alers and J. R. Neighbours, "The elastic constants of zinc between 4.2° and 670°K," Journal of Physics and Chemistry of Solids, vol. 7, no. 1, pp. 58–64, 1958.
- [358] S. Eros and C. Smith, "Low-temperature elastic constants of magnesium alloys," Acta Metallurgica, vol. 9, no. 1, pp. 14–22, 1961.
- [359] K. Clay, J. Jackson, P. Quested, and R. Morrell, "Improving single-crystal orientation determination for advanced nickel-based alloys," tech. rep., NPL, 2009.