

Imperial College London
Department of Mechanical Engineering

**A finite element method for GPU-based
elastic wave modelling in solid-fluid media
using displacement fields**

Yiannis Simillides

2023

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

Ultrasonic wave propagation and scattering involving both solids and fluids underpin many key configurations in non-destructive testing and underwater acoustics. The resulting interactions are highly dependent on both material parameters and geometries and are often difficult and expensive to investigate experimentally. Modelling capabilities are often used to overcome this, but these are also complex and computationally expensive due to the complexity of the fluid-solid interactions.

In this thesis, work is presented on simulating ultrasonic wave phenomena involving both solids and fluids, using the finite element method implemented on graphical processing units. The thesis begins with a literature review existing analytical and numerical methods as well as certain necessary building blocks and useful background information.

A novel explicit time-domain finite element method for simulating ultrasonic waves interacting with fluid-solid interfaces is then introduced, allowing investigation of complex, industrially relevant configurations at scale and speed. The method is displacement-based, and relies on classical hourglassing control, in addition to a modified time-stepping scheme to damp out any shear motion in an inviscid fluid. One of the key benefits of the displacement-based approach is that nodes in the fluid have the same number of degrees of freedom as those in the solid. Therefore, defining a fluid-solid model is as easy as defining an all-fluid or all-solid model, avoiding the need for any special treatments at the interfaces common in other approaches. It is thus compatible with typical elastodynamic finite element formulations and ready for implementation on a graphical processing unit.

The method has been verified across a range of problems that involve millions of degrees of freedom from different fields, such as non-destructive testing and underwater acoustics, and across different scales, with frequencies of interest ranging from kHz to MHz. There is also a brief investigation on geometrical model refinement and some initial research and simulations conducted to extend this work in three dimensions.

Acknowledgements

As my journey as a PhD student is coming to an end, there is a significant amount of people I would like to acknowledge; some of them started this journey with me, and some joined in the process.

To all my colleagues in the NDT group at Imperial, it has been a pleasure working, attending conferences and seminars and taking part in various social events with you. Specific thanks to Andy, Euan, Evripides, Filip and Georgios.

To the Graduate School at Imperial and Dr. Katerina Michalickova for allowing me to design and teach a Julia programming course from scratch; I definitely learned a lot in the process and thoroughly enjoyed the experience.

To Dr. Oliver Sanford, Dr. Stewart Haslinger and Prof. Richard Craster, for introducing me to the underwater acoustics field, and thus enhancing the research and applicability of the work done in this thesis.

To my supervisors, Dr. Michał Kalkowski, Dr Peter Huthwaite and Prof Mike Lowe: for helping guide my research over the years, providing clarification to points in the literature and for their patience as I conducted my research. Special thanks individually to Michal for reading through first drafts of my writing, laboriously helping me edit and strengthen the writing. To Mike, for leading the department, particularly over the Covid pandemic, and making it such an enjoyable and productive place to work. To Peter, for always providing support and help in all aspects and for assisting me in tracking down various opportunities related to my PhD.

And finally to my family, for always believing in me and my abilities, with special gratitude to my partner Antigoni for her unconditional love and support, and for buying me food and drinks.

Copyright Declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Statement of Originality

All research in this thesis has been carried out independently by myself, under the aid and guidance of my supervisors, Dr. Peter Huthwaite, Dr. Michał Kalkowski and Prof. Mike Lowe. Wherever other work has been used, I have acknowledged it.

Yiannis Simillides

Contents

Abstract	3
Acknowledgements	5
1 Introduction	21
1.1 Motivation	21
1.2 State of the art	24
1.3 Aims of this thesis	29
1.4 Contributions	30
1.5 Thesis Outline	30
1.6 List of Publications	31
2 Investigating the modelling of elastic wave phenomena	32
2.1 The wave equation and elastic waves	33
2.1.1 Scalar wave equation	33
2.1.2 Vector wave equation	34
2.1.3 Linearity of the wave equation	36

2.2	Review of methods for modelling wave propagation	37
2.2.1	Transfer matrix method	37
2.2.2	Distributed point source method	38
2.2.3	Boundary element method	38
2.2.4	Finite volume method	39
2.2.5	Spectral element method	39
2.2.6	Finite-difference method	40
2.2.7	Finite element method	44
2.2.8	Combined methods for modelling wave propagation	47
2.3	Boundary conditions	48
2.3.1	Physical boundary conditions	48
2.3.2	Numerical boundary conditions: modelling non reflecting boundaries	49
2.4	Summary	56
3	Development & Implementation of a novel finite element method for elastic wave phenomena	57
3.1	Introduction	57
3.2	Overview of the finite element method literature	59
3.3	The finite element formulation	62
3.3.1	Finite element formulation for solid materials	62
3.3.2	A novel finite element formulation for fluid materials	64

3.4	Algorithms and software implementation	68
3.4.1	Computational Domain	69
3.4.2	Matrix assembly	71
3.4.3	Matrix storage	74
3.5	Computational performance of the method	75
3.6	Summary	77
4	Solving Problems I: Application to Underwater Acoustics	79
4.1	Scattering from cylindrical structures	80
4.1.1	Solid cylinder	86
4.1.2	Layered cylindrical media	90
4.1.3	Cylindrically layered media, arbitrary filling	97
4.2	Objects lying and buried in the sediment	103
4.3	Setting up accurate simulations	104
4.3.1	Thin solid coating	104
4.3.2	Courant number accuracy	108
4.4	Solid sphere in the free field	110
4.5	Concluding remarks	113
5	Solving Problems II: Application to Non-Destructive Testing	114
5.1	Guided wave theory	115
5.2	Numerical simulations of guided waves	117

5.2.1	Existence of wavemodes	118
5.2.2	Phase velocity	120
5.2.3	Attenuation	123
5.3	Concluding Remarks	126
6	Concluding Remarks and Future Research Avenues	127
6.1	Review of Thesis	127
6.2	Key Contributions	128
6.3	Future Work	129
	Bibliography	130

List of Tables

1.1	Underwater acoustics benchmarks investigated at the NATO Undersea Research Centre and reported by [1].	24
4.1	Material properties used in numerical examples.	87
5.1	Comparison of phase velocities between Disperse and FE simulations.	123
5.2	Comparison of attenuation between Disperse and FE simulations.	126

List of Figures

1.1	A photograph of the immersion tank in the non-destructive testing lab at Imperial College.	23
1.2	Unexploded ordnance, spherical shape, in the free-field, taken by Sergey Dubrov [2].	25
1.3	Unexploded ordnance, cylindrical shape, proud on the seabed, taken by Sergey Dubrov [2].	25
2.1	Hourglass modes for a quadrilateral element.	46
2.2	Hourglass deformation taken from [3].	46
2.3	Absorbing boundary (striped) surrounding the area of study, which is a subset of an infinite domain.	51
3.1	Structured, unstructured, and combined meshes.	70
3.2	True object and computational approximation in pink and blue respectively.	72
3.3	Mapping of local (left) to reference element (right).	73
3.4	Node linking from the i^{th} node.	75
3.5	Ratio of solid-fluid vs fluid executable in Pogo run with (a) compression on and (b) compression off.	78

4.1	Radial displacement.	81
4.2	A schematic showing a (a) solid cylinder and (b) annulus in water insonified by a plane wave. The red dashed lines show the receiver points.	87
4.3	Scattering results a for steel cylinder submerged in water compared to analytical results.	88
4.4	Scattering patterns from (a) steel cylinder and (b) perspex cylinder submerged in water at 1 kHz.	91
4.5	Analytical and frequency-shifted scattering patterns from (a) steel cylinder and (b) perspex cylinder submerged in water at 1 kHz and 0.99 kHz respectively.	92
4.6	Scattering patterns from a steel annulus submerged in water with a (a) water interior and (b) air interior.	94
4.7	As in Figure 4.6 but for a perspex annulus.	95
4.8	Scattering patterns for normalized displacement from a cylindrical glycerine inclusion.	96
4.9	Setup for partially filled annulus $\frac{2}{3}$ full submerged in water with an incoming plane wave at (a) 0, (b) 90 and (c) 270 degrees respectively.	98
4.10	Scattering patterns from a steel annulus $\frac{2}{3}$ full submerged in water with an incoming plane wave at (a) 0, (b) 90 and (c) 270 degrees respectively.	99
4.11	Same as in Figure 4.10 but for a perspex annulus.	100
4.12	Scattering patterns from an annulus submerged in water with different levels of filling comprised of (a) steel and (b) perspex. Plane wave is travelling from the right.	102
4.13	Comparison between scattering pattern in Pogo and COMSOL for a steel annulus with a water interior, as described in Section 4.1.3.	103

4.14	Object proud or buried in sediment, with levels of burial at (a) 0, (b) 50 and (c) 100% levels of burial. Receivers are placed 6m above the water - sediment interface.	105
4.15	Results for the directivity patterns for a steel cylinder insonified with a plane wave at (a) 0, (b) 50 and (c) 100% levels of burial.	106
4.16	As in Figure 4.15 but for perspex.	107
4.17	Scattering patterns from a thin steel annulus, with (a) 10 elements through-thickness and (b) 5 elements through-thickness.	109
4.18	Scattering patterns for a steel sphere submerged in water compared to analytical results.	112
5.1	Schematic of an immersed steel plate alongside cartesian axes.	117
5.2	Dispersion curves for a 1 mm thick steel plate submerged in water (a) shown with no processing and amplitude visible (b) calculated by processing results from an FE simulation, (c) overlaid with Disperse curves (dashed-lines) at 1 MHz.	119
5.3	As in Figure 5.2 but for 5 MHz.	121
5.4	(a) Combined signal traces of the A_0 propagating wave measured at two points as described in Section 5.2.2 alongside (b) their amplitude spectrum. Red stars in (b) indicated local minima on the graph, found through use of Matlab routines.	122
5.5	A_0 mode shape data from Disperse for a 1 mm thick steel plate at 1 MHz submerged in water with (a) displacement (b) and phase.	124
5.6	FFT of first and second signal.	125

Chapter 1

Introduction

1.1 Motivation

Fast numerical methods for simulating elastic wave phenomena in coupled solid-fluid media are needed to model a range of complex geometries in real-world applications. Ultrasonic wave propagation and scattering in domains composed of both solids and fluids is a key research interest in diverse subject areas such as medical science [4], geophysics [5], non-destructive testing [6], or underwater acoustics [7]. My research, as detailed in this thesis, focuses on the last two fields: non-destructive testing (NDT) and underwater acoustics (UA).

Non-destructive testing is primarily concerned with the identification and sizing of defects in high-value, safety-critical engineering components without damaging or destroying them. Techniques can be broadly classified as surface, identifying surface or near-surface level defects, or volumetric, identifying defects or the lack thereof, in the internals of the tested object. There is a wide range of methods and techniques which fall into these categories. While this thesis focuses on ultrasonic propagation, it is helpful to mention some of these techniques in what follows. However, this is by no means an exhaustive list and the reader is referred to papers such as [8, 9, 10], which cover them in more detail, alongside their respective advantages, disadvantages and use cases. There are also technical standards one must

follow when using these techniques in industrial settings which vary from region to region. Details of an applicable standard which covers the usage of guided-wave testing in the United Kingdom is found in [11].

Common and broadly used surface methods are eddy-current testing, fluorescent penetrant inspection and thermal methods. Eddy-current testing involves electric currents and magnetic fields to identify defects, fluorescent penetrant inspection involves fluorescent penetrant applied to the surface of the object to aid visual inspection of defects, and thermal methods involve heating the object and monitoring the heat flow across the surface using an infrared camera [8].

One of the most commonly used volumetric testing methods is radiography. It has been heavily studied, formulated and researched. Ionising x-rays, gamma-rays and the object's absorption properties are used to inspect the component of interest [12]. This can be extended to create three-dimensional cross-sectional images by rotating the source or the object and is the basis of computer tomography [13].

Ultrasonic testing methods, which rely on ultrasonic wave phenomena, are also frequently used for both surface and volumetric inspection. These are varied from the conceptually relatively simple to more complex techniques, which rely on information from many sources and receivers. Classical setups include the pulse-echo technique in which ultrasound is transmitted to the surface (pulse), and the response (echo) is measured and then used to determine the size or location of a defect [14], or the through-transmission technique, in which two transducers, the source and the receiver are placed on opposite sides of the component. More advanced techniques exist, such as the Time-of-Flight-Diffraction (TOFD) technique, mainly used to size cracks, relying on the differing arrival times from the diffracted waves on the crack extremities [14], or phased-array testing in which multiple single-element transducers are used to create a focused beam to sweep the component of interest [15]. Another method is the full matrix capture technique (FMC) which allows one to capture all possible data from a phased array setup. This is then often used in conjunction with the total focusing method (TFM) algorithm which allows us to generate an image of the object at every depth and in




Figure 1.1: A photograph of the immersion tank in the non-destructive testing lab at Imperial College.

focus [16]. This technique can be applied both in contact but also in an immersion setup [17]. Guided-wave testing is another such method, in which guided-waves and their dispersion characteristics are used to identify defects in structures with one dimension significantly larger than the cross-section. Such configurations support the propagation of guided-waves which allow for long-range inspection [18, 19].

One of the key industrial use cases of ultrasonic propagation involving domains of both solid and fluid materials is immersion testing, where the examined component is immersed in a fluid and ultrasound travels to the object and then to the receiver, so the object can be investigated for any defects. The fluid acts as a couplant, allowing for greater transmission of the acoustic wave. It has been successfully used to characterise materials [20], or to retrieve a signal from some defect [21]. A photograph of an immersion tank is found in Figure 1.1.

Geometry	Type	Environment	Filling
Cylinder	Perfectly rigid	free-field	fully-filled (rigid)
Cylinder	Shell	free-field	void (air)
Cylinder	Shell	free-field	partially-filled
Cylinder	Shell	proud	void (air)
Cylinder	Shell	proud	partially-filled
Sphere	Shell	half-buried	fully-filled
Sphere	Shell	buried	fully-filled

Table 1.1: Underwater acoustics benchmarks investigated at the NATO Undersea Research Centre and reported by [1].

Another use-case is the inspection of pipes filled with, or surrounded by fluids. While quite general, specific applications where ultrasonic testing has been used are those of determining the composition of fluid [22], or in guided wave inspection of pipe such as in [23].

Underwater acoustics is also a broad application field considering the interaction of sound waves with both fluid and solid media. The applications of these range from mapping the ocean floor [24] to conducting biodiversity assessments of marine habitats [25]. In [1], a series of benchmark problems for underwater acoustics relating to unexploded ordnance was produced and is given in Table 1.1. These benchmark problems are closely related to the research conducted as part of this thesis as it looks at similar configurations for the numerical simulations presented in subsequent chapters. Other motivating examples are identifying and characterising sub-sea objects using acoustic scattering from targets located on a sandy sea floor [26] or unexploded ordnance in the free field [27]. Photographs of these cases are found in Figure 1.2 and Figure 1.3 respectively.

1.2 State of the art

While the use-cases and motivating examples in the two application fields detailed above can differ in materials, size of the object or domain, both would greatly benefit from fast, versatile modelling capabilities for coupled systems. Accurate simulation of ultrasound propagation and scattering effects in coupled solid-fluid interactions for a range of complex geometries and



Figure 1.2: Unexploded ordnance, spherical shape, in the free-field, taken by Sergey Dubrov [2].



Figure 1.3: Unexploded ordnance, cylindrical shape, proud on the seabed, taken by Sergey Dubrov [2].

materials in real-world applications is thus needed. What follows is a higher-level overview of different modelling methods. These are also discussed in more detail later on in this thesis.

As a result of the complexities and prominence of the solid-fluid interaction problem, research has been carried out to improve modelling capabilities for ultrasonic wave propagation, using analytical, semi-analytical, and numerical methods. The governing physics behind the wave equation is well understood, but analytical results are generally limited to simple geometries, which are not able to represent realistic practical scenarios often involving complex geometrical configurations [28, 29]. Semi-analytical methods rely on approximating the underlying physical phenomena.

One example of a semi-analytical method is the Kirchhoff approximation, which in this case is used to describe acoustic wave scattering in both underwater acoustics [30, 31] and non-destructive testing [32, 33] contexts from (rough) surfaces. The Kirchhoff approximation is based on an underlying assumption that each point on the surface of the scatterer can be treated as if it was a point on an infinite tangential plane to the scatterer [33]. The approximation thus assumes the full insonification of each point by the incident wavefield and that there is only a single interaction of the wavefield at that point. This, neglects among other effects, multiple scattering and surface waves, and is primarily valid for high frequencies.

The Born approximation is a semi-analytical method [34, 35] for low frequency approximations, which works by approximating the total field in the object of interest by the incident field [36]; a direct corollary of this is that the scattered field must be significantly weaker than the incident field. Both the Kirchhoff and Born approximation have been implemented in NDT simulation software such as CIVA [37].

Another semi-analytical method is the distributed point source method [38], where the total ultrasonic field is obtained via a summation of point sources distributed over a finite transducer. These semi-analytical methods allow us to solve a greater range of problems and configurations than analytical methods and at higher speed than numerical methods,

although these also suffer from the drawback of being unable to fully represent all possible practical scenarios.

Despite invaluable physical insights offered by the analytical and semi-analytical models, only numerical methods enable a full range of complex, industry-relevant configurations to be simulated. The benefit to these numerical simulations are that researchers and practitioners can simulate the experiment directly and adjust the design process using the model, optimizing the use of expensive experimental tests. Furthermore, one can develop new signal processing, inspection and imaging methods to capture specific phenomena, (eg. chosen defect scenarios), retaining complete control over the parameters of the defect or the structure [39]. Another more recent use of these numerical methods is to allow us to build large, diverse training sets for use in machine-learning applications [40].

Several numerical techniques have been used to accurately model ultrasound propagation in solid and fluid materials. Some of the most widely used methods found in the literature are the finite-difference (FD) [41, 7], boundary-element (BEM) [42] and the finite-element (FE) methods [43]. There are two specific challenges when using numerical methods to model physical phenomena

- i. numerical models of realistic three-dimensional setups can quickly grow to extremely large computational sizes,
- ii. the coupling between the solid and the fluid can add additional numerical complexities depending on the formulation used.

One solution to the problem of computational size is to implement the numerical methods on graphical processing units (GPUs). Cases involving only solid media can currently be solved using highly-efficient displacement-based elastodynamic FE solvers benefiting from GPUs [44]. As described in [44], GPUs are exceptionally well-suited to solving highly parallelisable problems involving large matrix calculations compared to central processing units (CPUs). This is thanks to the GPU architecture, where a large number of processors allow many

threads to run in parallel. They have significantly higher processing power and memory bandwidth compared to CPUs. These characteristics make GPUs an ideal architecture for executing explicit time-domain simulations which model acoustic wave phenomena where each time step only depends on field quantities at previous steps and can be solved independently of each other.

In the Finite Element (FE) literature for the fluid-structure interaction (FSI), various formulations have been proposed to model the acoustic wave propagation. Most common are the pressure formulation and the velocity potential formulation [45], which have the benefit of only requiring one degree of freedom per node. The coupling of solid and fluid domains using a pressure-based formulation using the finite element method can result in non-symmetric matrices. This is a result of off-diagonal terms resulting from the coupling between the solid and fluid interfaces [46, 47, 48]. These non-symmetric matrices are computationally expensive to use in numerical calculations [49]. Therefore, the pressure-based formulation cannot be easily used in calculations involving graphical-processing-units because of the aforementioned matrix inversion and the required special treatment at the solid-fluid boundary.

A displacement-based formulation allows modelling both domains using the same field variable. The resulting matrix is sparse and symmetric, speeding up numerical calculations. This also minimises data loading on the GPU, providing further speedups. However, displacement-based formulations for fluids have been troubled by their own challenges in the past. Most commonly reported problems include added computational requirements when compared to pressure-based formulations. This is because displacement is a vector-based quantity compared to the scalar pressure, and the presence of spurious circulation modes [50].

This is where the research presented in this thesis contributes to the existing literature: there is no efficient way to model fluid-solid systems that would permit simulating realistic inspection setups directly for the reasons mentioned above. The aims of this thesis in addressing these challenges are detailed in the following section.

1.3 Aims of this thesis

To improve the versatility of time-domain finite element modelling in coupled fluid-solid media, displacement-based FE formulation designed for use on a GPU is proposed. The method is explicit-in-time and relies on a modified damping matrix to minimise any non-physical shear motion in the fluid. In addressing the concerns mentioned above, the formulation has relied on hourglass control developed by Flanagan and Belytschko [51]. The development also draws from the work by Chen [50], on reduced integration techniques for four-noded quadrilateral elements to reduce hourglassing, and Everstine [47], who configures an isotropic elastic solid to represent a compressible inviscid fluid.

The proposed formulation enables the time domain finite element method in fluid-solid coupled media to be applied to industrially relevant problems. The method's versatility for time-domain problems, which allows us to extract data for multiple frequencies from one simulation and its ability to handle arbitrary and complex geometries, are crucial for the use-cases. Furthermore, the method can handle diverse domains, which often have very different scales of interest.

Thus, one of the central aims of this research is to formulate an explicit-in-time displacement-based finite element method allowing for implementation on GPUs. I can thus leverage the GPUs' increased performance to solve large-scale simulations at speed. This allows for executing simulations involving hundreds of millions of degrees of freedom in a relatively short time, enabling fine mesh refinement or performing parametric studies efficiently. A displacement-based formulation ensures compatibility with existing solid solvers, alongside a non-diagonal damping matrix that is suitable for GPU calculations and minimises hourglassing.

1.4 Contributions

In this thesis I have made the following contributions; I have formulated, implemented and tested a displacement-based time-domain finite element method for ultrasonic propagation in domains involving both solid and fluid. The implementation was initially prototyped in Python with the use of sparse matrix functionality from NumPy [52] and SciPy [53], to allow for greater interactivity and interpretability compared to complex, highly optimised GPU code found in implementations such as Pogo [44]. I subsequently verified and validated this implementation across a range of different domains, by solving canonical examples found in the non-destructive testing and underwater acoustics literature and comparing their results to known, verified and validated analytical and semi-analytical solutions, such as those found in papers by Faran [54], Doolittle [55], experiments by Tamarkin [56] or in guided-wave software such as Disperse [19]. I have also provided results for problems without a known-analytical solution, demonstrating the versatility of my method to handle complex geometries and multiple materials.

1.5 Thesis Outline

Chapter 2 reviews the governing equations underlying ultrasound propagation, provides a literature review of analytical and numerical methods to model this propagation and covers the boundary conditions and ways to simulate these in numerical software,

Chapter 3 looks at the finite element method in more detail, covering its implementation in software, and specifically focuses on my proposed formulation to simulate ultrasound involving solid-fluid domains on a GPU,

In Chapter 4, numerical simulations related to the field of underwater acoustics are introduced and simulated. I show results for scattering patterns from canonical setups in both the free-field and proud or buried in the surface.

In Chapter 5, I look at numerical simulations regarding non-destructive-testing, particularly at simulating guided-waves. I demonstrate that I can accurately predict their existence, capture their velocity and in the leaky-case, the attenuation of the A_0 wave mode.

In Chapter 6, I summarise the work conducted as part of this thesis and provide future research avenues.

1.6 List of Publications

Research conducted in this thesis appears in the following articles

- R. V. Craster, S. G. Haslinger, P. Huthwaite, and **Y. Simillides**, “Low frequency minehunting sonar modelling, processing and design - preliminary modelling results,” Confidential technical report, Imperial College London, Imperial Consultants, 2022.
- R. V. Craster, S. G. Haslinger, P. Huthwaite, and **Y. Simillides**, “Low frequency minehunting sonar modelling, processing and design - Further modelling results,” Confidential technical report, Imperial College London, Imperial Consultants, 2022.
- **Y. Simillides**, P. Huthwaite, Michał K. Kalkowski, Michael J. S. Lowe, “A displacement-based finite element formulation for solving elastic wave problems in coupled fluid-solid media on a GPU, submitted to “Computers and Structures”.

Chapter 2

Investigating the modelling of elastic wave phenomena

This chapter lays out the theoretical foundations of modelling ultrasound propagation in fluids and solids which is used in non-destructive evaluation and in underwater acoustics. As numerically modelling ultrasound in domains composed of both solids and fluids is the core part of the research conducted as part of this thesis, an understanding of the underlying physics and equations that govern this propagation is necessary. This chapter comprises of two parts; in the first part, the theoretical treatment of wave propagation in elastic and acoustic media including governing equations, is discussed. The second part reviews analytical and numerical methods for modelling wave propagation in contexts of varying complexity, while also highlighting any issues which may be lacking in the aforementioned methods. As part of this review of numerical methods, techniques to implement the necessary boundary conditions to accurately model wave phenomena are also described.

2.1 The wave equation and elastic waves

2.1.1 Scalar wave equation

In one dimension, the scalar wave equation is given as [57]

$$c^2 \frac{\partial^2 u}{\partial x^2} = \frac{\partial^2 u}{\partial t^2}. \quad (2.1)$$

where c is the wave speed and the variable $u(x, t)$, without loss of generality, represents the field quantity, such as pressure, displacement (which is of particular interest to us), and in some cases, the velocity potential. In two or three dimensions, the scalar wave equation for acoustics is given as [57]

$$c^2 \nabla^2 p = \frac{\partial^2 p}{\partial t^2}, \quad (2.2)$$

where $u(x, t)$ has been replaced with pressure $p(x, t)$ to be consistent with the difference between a scalar and a vector in higher dimensions.

By looking for a time-harmonic solution of the wave equation of the form

$$u(x, t) = v(x) \exp^{-i\omega t}, \quad (2.3)$$

one can arrive at the following

$$\nabla^2 v + k^2 v = 0. \quad (2.4)$$

This is the Helmholtz equation, which is sometimes referred to as the reduced wave (Helmholtz) equation [58] and is used to model the equations of motion in the frequency domain [59]. In Equation (2.4), $k = \frac{\omega}{c}$ is the wavenumber, a physical parameter of the wave.

In [60], the authors give a more general form of a common linear problem as

$$\nabla^2 u + g = a\ddot{u} + b\dot{u}, \quad (2.5)$$

where $g = g(x)$, $a = a(x)$ and $b = b(x)$ are position dependent. By specifying the appropriate parameters, one can go from a general case in Equation (2.5) to the wave equation given in Equation (2.2). Using Equation (2.5), one can also see how damping could be added to the specific equation by varying the parameter $b(x)$. Other well known partial differential equations of this form, although not detailed in this thesis, are the Heat Equation [61], Laplace's Equation [62] and Poisson's Equation [63].

2.1.2 Vector wave equation

In the case of (solid) elastic media, the isotropic wave equation (Naviers Equation [64]) where the material is inhomogeneous is given by

$$\rho \ddot{\mathbf{u}} = \nabla \lambda (\nabla \cdot \mathbf{u}) + \nabla \mu \cdot [\nabla \mathbf{u} + (\nabla \mathbf{u})^T] + (\lambda + 2\mu) \nabla (\nabla \cdot \mathbf{u}) - \mu \nabla \times \nabla \times \mathbf{u}, \quad (2.6)$$

where λ and μ are the Lamé parameters or constants, which are material dependent. For isotropic media, where the material is homogeneous, i.e. when the gradient of the material is constant and thus $\nabla \lambda = \nabla \mu = 0$, the wave equation is given as

$$\rho \ddot{\mathbf{u}} = (\lambda + 2\mu) \nabla (\nabla \cdot \mathbf{u}) - \mu \nabla \times \nabla \times \mathbf{u}. \quad (2.7)$$

This equation admits two oscillatory solutions, compressional (P) and shear (S) waves [65]. In seismology, these are also often referred to as (P)rimary and (S)econdary waves due to their differences in speed and therefore arrival time [66]. To obtain the solution of these two separate waves, the Helmholtz decomposition is used [67]

$$\mathbf{u} = \nabla \Phi + \nabla \times \Psi, \quad (2.8)$$

$$\nabla \cdot \Psi = F(\mathbf{r}, t), \quad (2.9)$$

where $\Phi(\mathbf{x}, t)$ is a scalar function, $\Psi(\mathbf{x}, t)$ is a vector valued function and F is a function of \mathbf{r} , the position vector, and t the time. Inspecting Equation (2.8), the equation has n degrees of freedom on the left-hand-side and $n + 1$ degrees of freedom on the right-hand-side. To determine \mathbf{u} uniquely, gauge fixing from gauge theory is relied upon [68], which allows F to be chosen arbitrarily. Choosing F to be zero [69], makes Ψ a zero-divergence vector. This allows \mathbf{u} to be uniquely defined.

Substituting Equation (2.8) into Equation (2.7) gives

$$\rho \frac{\partial^2}{\partial t^2} (\nabla \Phi + \nabla \times \Psi) = (\lambda + 2\mu) \nabla (\nabla^2 \Phi) - \mu \nabla \times \nabla \times (\nabla \times \Psi), \quad (2.10)$$

which simplifies to

$$\rho \frac{\partial^2}{\partial t^2} (\nabla \Phi + \nabla \times \Psi) = (\lambda + 2\mu) \nabla (\nabla^2 \Phi) + \mu \nabla^2 (\nabla \times \Psi), \quad (2.11)$$

and by separating variables one arrives at the following

$$\nabla \left[(\lambda + 2\mu) \nabla^2 \Phi - \rho \frac{\partial^2 \Phi}{\partial t^2} \right] = -\nabla \times \left[\mu \nabla^2 \Psi - \rho \frac{\partial^2 \Psi}{\partial t^2} \right]. \quad (2.12)$$

Setting both the left-hand-side and right-hand-side to zero, gives two wave equations

$$\frac{\partial^2 \Phi}{\partial t^2} = \frac{(\lambda + 2\mu)}{\rho} \nabla^2 \Phi, \quad (2.13)$$

$$\frac{\partial^2 \Psi}{\partial t^2} = \frac{\mu}{\rho} \nabla^2 \Psi. \quad (2.14)$$

The wave speed of the primary wave given in Equation (2.13) is $\sqrt{\frac{\lambda + 2\mu}{\rho}}$ and is greater than the wave speed of the secondary wave given in Equation (2.14) which is $\sqrt{\frac{\mu}{\rho}}$ as the Lamé parameters are positive.

So far elastodynamics have been discussed. The approach taken in the rest of this thesis is to treat fluid materials as solid materials with analogous choice of material parameters. This is often used in the literature, such as for example in Disperse [19] or [47] by Everstine for

a fluid-displacement finite element method. Thus, the same formality should hold for fluids, with the exception of inviscid fluids, i.e. when $\mu = 0$, that cannot support shear stress. Consequently shear waves cannot propagate in such media, which will need to be accounted for in the numerical treatment of the equation in subsequent sections.

2.1.3 Linearity of the wave equation

The wave equations presented throughout this chapter are linear, i.e. they exhibit the following two properties: additivity and homogeneity. The additivity property means that wave equations are subject to constructive and destructive interference [70]. In one dimension, the addition of two plane waves (which are not necessarily in phase)

$$\begin{aligned} w_1 &= \sin(kx - \omega t), \\ w_2 &= \sin(kx - \omega t + \phi), \end{aligned} \tag{2.15}$$

gives

$$w_1 + w_2 = w_t = 2 \cos\left(\frac{\phi}{2}\right) \sin\left(\frac{\phi}{2} + kx - \omega t\right). \tag{2.16}$$

Thus, at $\phi = 2\pi\mu$, $\forall \mu \in \mathbb{Z}$, there is constructive interference, i.e.

$$w_t = 2 \sin(kx - \omega t), \tag{2.17}$$

and at $\phi = (2\pi + 1)\mu$, $\forall \mu \in \mathbb{Z}$, there is destructive interference, i.e.

$$w_t = 0. \tag{2.18}$$

The homogeneity principle implies that results can be scaled. Practically, this means that solutions and results are often shown with respect to some combination of parameters, such as frequency and size. Furthermore, due to linearity, combinations of solutions are also a solution to the wave equation.

2.2 Review of methods for modelling wave propagation

Let us now consider methods for modelling various configurations and solving the wave equation, accounting for different application contexts. Some of the earliest work on the problem of ultrasound propagation in coupled solid-fluids domains was done by Faran [54], who investigated sound scattering from solid cylinders and spheres in liquid. Interestingly, at the time this research was carried out, limitations in computing power meant the calculation of the Bessel functions proved difficult, although this poses no problems with modern computing capabilities. The formulations provided by Faran were later expanded upon by Doolittle [55], who provided scattering results for elastic cylindrical shells. By taking appropriate material parameters, they derived a general case of the solution described earlier by Faran [54], and confirmed earlier experimental results from Tamarkin [56], on scattering from a fluid inclusion inside a fluid. Skelton and James [71] developed that concept further and provided results for arbitrary layered media in cylindrical and spherical domains.

One can gain physical insights from these analytical models which allow one to investigate different properties of wave propagation and scattering. Despite their good accuracy, they are limited to simple axisymmetrical geometries, such as spheres or cylinders, which are mainly of academic interest. There is growing interest in modelling complex setups, closely related to realistic experimental scenarios, which requires developing numerical models, which are capable of overcoming the limitations of analytical models. These are detailed and reviewed in subsequent sections.

2.2.1 Transfer matrix method

The transfer matrix method, originally developed for electromagnetic waves [72], works by setting up matrix equations for the reflected and transmitted plane waves at each interface, repeated for all subsequent interfaces on the propagation path. When combined, these equations enable basic wave characteristics to be evaluated. One of the main disadvantages of

this method is that it needs to be set up for each separate geometry one wishes to model, and only certain geometries can be fully represented [73]. This drawback renders the method unsuitable for the complex geometries one encounters in domains such as non-destructive testing.

2.2.2 Distributed point source method

The use of Huygens' principle applied to ultrasonic field modelling can also be frequently found in the literature [74], where the total ultrasonic field is obtained from a linear summation of point sources distributed over a finite transducer. Various studies have been conducted using techniques based on Huygens' principle such as multi-Gaussian paraxial models detailed in [75], which work by superimposing Gaussian beam solutions. However, they cannot handle problems involving transmission at an interface near grazing incidence, or problems where the interface has different curvature. An improvement to this technique is the distributed-point-source-method (DPSM) [38], where the point sources are distributed on both sides of the problem boundary and are placed at a constant distance from it. Using two layers of point sources on two sides of the interface avoids singularity and the need of reflection and transmission coefficient expressions that are required by other point source superposition based methods. Once again, one of the main disadvantages with this class of methods is that one needs to derive an implementation of this method for each geometry one wants to investigate, and only certain geometries have been represented using this method.

2.2.3 Boundary element method

In the boundary element method, which is based on the integral representation of the acoustic field, the surface of the object is discretised into "boundary elements" [42]. Matrices are then assembled based on these boundary elements and the resulting linear equations must be solved. This method is often considered a "direct competitor" to the finite element method in the literature [76]. There is no clear agreement on which method is advantageous, with

benefits and disadvantages to both methods. One often reported benefit of the boundary element method is that only the surface needs discretisation, significantly reducing the number of degrees of freedom needed when assembling the matrices compared to the finite element method. Despite the reduction in the number of degrees of freedom, the resultant matrix is not sparse; this has its own computational challenges.

2.2.4 Finite volume method

In the finite volume method the governing equations are discretised over each control volume and conservation laws are upheld. The finite volume method (FVM) is frequently used in computational fluid dynamics, although it not as popular for structural mechanics. In [77] states that “it is commonly accepted in the literature that the finite element method is the more convenient option for problems involving solid materials”, although in [78] it is stated that they are equivalent in some cases. There is very limited work done on linear ultrasonic wave propagation and scattering using the finite volume method, but maybe there is some unexplored potential which could benefit ultrasound simulation, particularly when dealing with problems involving fluid flow as the FVM is well suited for such problems [79].

2.2.5 Spectral element method

The spectral element method (SEM) is a special case of the finite element method, which relies on high-order basis function [80]. It has been extensively used in geophysics [81], a field which has many similarities with the non-destructive testing and underwater acoustics space. However, the spectral element method is not as well-known in these two spaces. It has the large advantage compared to the “general” finite element method, that it does not need much refinement per wave-length; in [81] it is stated that a single spectral element per wavelength is sufficient. The drawbacks of this method is that we need to use higher-order polynomials in the SEM method [82], otherwise the results are inaccurate [80] unlike the conventional finite element method. Furthermore, the resulting matrices in the spectral element method are

dense. The limitation of dense matrices is particularly challenging for GPU-based software and is thus not suitable for our use-case. In [83], it is shown that the two methods produce very similar results, which is to be expected due to the inherent similarities of the methods.

2.2.6 Finite-difference method

The finite difference method obtains solutions to the wave equation by discretising the domain and replacing the continuous differential equations with a set of discrete equations (finite-differences) that can be solved to obtain the solution at discrete points. Finite-difference (FD) techniques have frequently been used throughout the literature, due to their comparative ease of use and computational power required as well as the fact that they are conceptually simple, a direct discretisation of the wave equation. Their main disadvantage is that they cannot accurately model complex domains. Numerical modelling of the wave equation in the NDT context using the FD method can be found in [84] where they modelled the Scholte wave [85] or in [86] dealing with elastic wave propagation in submerged solids. In [7], the authors use the finite difference method in polar coordinate form to calculate the response from submerged cylinders with an arbitrary amount of filling. While such calculations are not possible to do using conventional axisymmetric techniques, the finite difference method is still limited to simple geometries.

As the finite difference method is still frequently used to solve the wave-equation [7], the derivation is described from first principles so one can refer back to certain aspects of it later on. Consider the Taylor series expansion of $u(x + \Delta x)$ and $u(x - \Delta x)$ for some $\Delta x > 0$

$$u(x + \Delta x) = u(x) + \frac{\dot{u}(x)\Delta x}{1!} + \frac{\ddot{u}(x)(\Delta x)^2}{2!} + R_k(x), \quad (2.19)$$

$$u(x - \Delta x) = u(x) - \frac{\dot{u}(x)\Delta x}{1!} + \frac{\ddot{u}(x)(\Delta x)^2}{2!} + R_k(x). \quad (2.20)$$

In equations (2.19)-(2.20), terms of $O(\Delta x^2)$ or higher (remainder function $R_k(x)$) are ignored

as these are assumed to be infinitesimally small. Rearranging, gives

$$\dot{u}(x) \approx \frac{u(x + \Delta x) - u(x)}{\Delta x}, \quad (2.21)$$

and similarly, from equation (2.20),

$$\dot{u}(x) \approx \frac{u(x) - u(x - \Delta x)}{\Delta x}. \quad (2.22)$$

By combining equations (2.21) and (2.22), one arrives at the following approximation

$$\dot{u}(x) \approx \frac{u(x + \Delta x) - u(x - \Delta x)}{2\Delta x}. \quad (2.23)$$

Equation (2.21) is called the forward difference, equation (2.22) is called the backward difference and equation (2.23) is called the central difference discretisation of the first derivative of u with respect to x . One can perform similar approximations for higher-order derivatives. For instance, the second-order central difference would read

$$\ddot{u}(x) \approx \frac{u(x + \Delta x) - 2u(x) + u(x - \Delta x)}{(\Delta x)^2}. \quad (2.24)$$

Second-order forward and backward differences follow the same paradigm.

The finite difference method works by approximating (partial) derivatives with difference quotients calculated above in equations (2.21)-(2.24). Consider the following initial value problem

$$\dot{u}(x) = f(x, u(x)), \quad (2.25)$$

$$u(0) = u_0. \quad (2.26)$$

By replacing equation (2.25) with equation (2.21) one obtains

$$u(x + \Delta x) = u(x) + \Delta x f(x, u(x)), \quad (2.27)$$

This is the forward Euler method, which is an explicit method, as $u(x + \Delta x)$ can be computed explicitly, without solving a nonlinear system of equations. Similarly, by replacing equation (2.25) with (2.22) one gets

$$u(x) = u(x - \Delta x) + \Delta x f(x, u(x)). \quad (2.28)$$

This is the backward Euler method, an implicit method, as $u(x)$ can only be computed implicitly, necessitating the solution of nonlinear equations. Solving these nonlinear equations is computationally inefficient, particularly for large simulations, despite often giving a more stable solution. Another alternative solution is combining the two into the so-called implicit-explicit methods (IMEX) [87] where both an explicit and an implicit method are used to numerically model the wave equation. Although such methods have been applied to model the solution of the wave equation in coupled solid-fluid domains, such as in [88, 89], these are currently unsuitable for GPU calculations due to their additional complexity, as one would need to handle the explicit and implicit parts separately.

Explicit methods are therefore preferred to implicit methods for GPU calculations, as the variables at one time-step can be computed from their known values at previous time-steps. The reason for this preference is the large number of processors which allow many threads to run in parallel as well as the GPUs high-memory bandwidth for loading in data. This is well known in the literature, as can be seen in [90], where it is stated that "explicit methods map well to the GPU architecture as the output elements can be computed independently of each other, thus making full use of the high computing throughput of a GPU" or in [44], where an in-depth explanation of the GPU architecture and how this maps to explicit-in-time finite elements methods is given and the reader is referred to these for further information about the GPU architecture.

Let us now apply the central difference scheme to the wave equation originally introduced in Equation (2.1). The field variable $u(x, t)$ is a function of space and time, so we apply the central difference approximation to both dimensions. Without loss of generality, denote the

time steps by the superscript k and the spatial steps by the subscript n to arrive at

$$\frac{\partial^2 u}{\partial x^2} = \frac{u_{n+1}^k - 2u_n^k + u_{n-1}^k}{\Delta x^2}, \quad \frac{\partial^2 u}{\partial t^2} = \frac{u_n^{k+1} - 2u_n^k + u_n^{k-1}}{\Delta t^2}, \quad (2.29)$$

and by rewriting the equations using equation (2.1), one arrives at the following update step

$$u_n^{k+1} = r^2[u_{n+1}^k - 2u_n^k + u_{n-1}^k] + 2u_n^k - u_n^{k-1}, \quad (2.30)$$

where $r = c \frac{\Delta t}{\Delta x}$. The update step shown in Equation (2.30) is explicit in time and thus, well-suited for implementation on a GPU[91].

Von Neumann stability analysis

We have seen why explicit based numerical methods are the way forward for numerically modelling the wave equation using a GPU. However, explicit numerical methods based on discretisation may lose stability if the discretisation is not sized correctly. A common approach to this is to perform Von Neumann stability analysis (which Trefethen defines as “the analysis of stability by Fourier methods” [41]) to arrive at the Courant-Friedrichs-Lewy (CFL) condition [92] necessary for convergence of the numerical method. Let us assume a solution has the form

$$u_n^k = \lambda^k \exp(inx), \quad (2.31)$$

where λ is the growth factor. If $|\lambda| > 1$, the solutions will be unbounded as time increases (i.e. as $k \rightarrow \infty$) which has no physical meaning in the context of the wave equation. Hence the necessary condition is that $|\lambda| < 1$. Substituting equation (2.31) into equation (2.30), gives

$$\lambda = r^2[\exp(ix) - 2 + \exp(-ix)] + 2 - \frac{1}{\lambda}, \quad (2.32)$$

which can be rearranged as

$$\lambda^2 - \lambda(r^2[\exp(ix) - 2 + \exp(ix)] + 2) + 1 = 0. \quad (2.33)$$

This then gives us the following two solutions for λ

$$\lambda = \mu \pm \sqrt{\mu^2 - 1}, \text{ where } \mu = 1 - 2r^2 \sin^2\left(\frac{ix}{2}\right). \quad (2.34)$$

The condition $|\lambda| < 1$ implies that $|\mu| < 1$ and thus

$$\left|1 - 2r^2 \sin^2\left(\frac{ix}{2}\right)\right| < 1, \quad (2.35)$$

$$|2r^2| < 2, \quad (2.36)$$

$$|r| < 1. \quad (2.37)$$

Given that $r = c \frac{\Delta t}{\Delta x}$ is positive, this leads to the stability condition $c \frac{\Delta t}{\Delta x} < 1$. The above stability condition means simulations must be setup according to the wavenumber of the propagating wave and this is often one of the main limitations when modelling numerical methods, as for very high frequencies the mesh refinement needs to be very small.

2.2.7 Finite element method

The finite element method calculates the solution of the wave equation by discretising the domain into “finite elements” and assembling matrices over these elements. An in-depth mathematical explanation will be given in subsequent chapters. In the finite element literature for the fluid-structure interaction (FSI), various formulations have been proposed to model the acoustic wave propagation. Most common is the pressure-based or velocity potential formulation [45], which have the benefit of only requiring one degree of freedom (DoF) per node. There have also been attempts with non-standard elements for the fluid domain, such as the Raviart-Thomas elements which are edge based and do not suffer from

hourglassing [93]. Despite much promise, they are not generally compatible with conventional FE software used for solid materials, and therefore associated functionality such as meshing software is not currently suited for such elements. Hourglassing is the phenomenon where an element provides no resistance to a particular mode shape; this is a non-physical phenomenon but tends to pollute the result rapidly - within a mesh such modes can quickly propagate through the domain and render the solution unusable [3]. Hourglass modes can be found in Figure 2.1, and the deformation that hourglass causes to the solution can be found in Figure 2.2. In the literature, these can often be termed zero-energy modes as they produce zero strain or stress [3]; an alternate terminology is spurious as they have no physical basis. Initially, the presence of these modes was not fully understood in the numerical modelling community. Nowadays theoretical explanations describing the phenomenon exist, and relate to the rank of the assembled finite element stiffness matrix. Without loss of generality, for the two-dimensional displacement case, the 8×8 stiffness matrix has rank 5, which implies it has an infinite number of solutions, not a unique solution. To obtain a unique solution, rigid modes are removed according to prescribed boundary conditions. We refer the reader to works such as [94] or [95] for a detailed explanation.

Coupling solid and fluid domains using a pressure-based formulation in FE can result in non-symmetric matrices [46, 47, 48] or the need for matrix inversion [49]. A demonstration of this non-symmetry will be shown in the next chapter. Non-symmetric matrices, or matrices which require inversion are not immediately suitable for large-scale problems due to their computational cost. Such industry-relevant cases involving only solid media can be solved using highly-efficient displacement-based elastodynamic FE solvers benefiting from graphical computing cards [44]. Unfortunately, the pressure-based formulation cannot be used in that framework because of the aforementioned matrix inversion and the required special treatment at the solid-fluid boundary.

Displacement-based formulations have also appeared in the literature, that is where both the fluid and solid are modelled using a displacement field. One of the key benefits of the displacement-based approach is that nodes in the fluid have the same number of degrees of




Figure 2.1: Hourglass modes for a quadrilateral element.




Figure 2.2: Hourglass deformation taken from [3].

freedom as those in the solid. Therefore, defining a fluid-solid model is as easy as defining an all-fluid or all-solid model, avoiding the need for any special treatments at the interfaces. It is thus compatible with typical elastodynamic finite element formulations and ready for implementation on a graphical processing unit. Owing to the added computational requirements when compared to pressure-based formulations alongside the presence of spurious circulation (hourglass) modes, these have not been widely used in the past [50]. These spurious circulation modes can be found throughout the numerical method literature and extensive research has been carried out to control these instabilities, such as by Hamdi [96] who applied a penalty method, by Belytschko who used mesh stabilization techniques [3] or Bathe [97] who used a mixed method. Unfortunately, these methods described above to handle the instabilities are often not suitable for GPU calculations due to the added mathematical terms, which result in non-symmetric or non-sparse matrices. In subsequent chapters, formulations of the finite element method related to the wave equation will be shown and expanded upon.

2.2.8 Combined methods for modelling wave propagation

Apart from the methods detailed above to model wave propagation in coupled solid-fluid domains, it is also possible to use a combination of different methods, often termed coupled or combined approaches in the literature. These can often alleviate some of the disadvantages of individual methods, at the cost of increased theoretical complexity as well as the need for two separate implementations in software. Due to these inherent drawbacks, it has been decided not to attempt to use these to model the solution of the wave equation on a GPU, although for completeness some of the more common combinations are mentioned, particularly those involving a finite element method alongside another method. A common approach is to combine a boundary element method to model the fluid with a finite element method to model the solid material. Numerous examples of such approaches can be found in the literature such as in the following [98, 99, 100]. Other approaches can be found in [101] where a coupled finite-volume-finite-element method is presented or in [102] where the analytical solution for the fluid is combined with a finite element method for the solid.

2.3 Boundary conditions

Different methods to model the wave equation and obtain its solution have been detailed. To fully describe the wave equation, one must also consider the boundary of the domain. Boundary conditions describe the behaviour of a system at its boundaries, and they play a critical role in determining the accuracy of the numerical solution, and its computational complexity.

2.3.1 Physical boundary conditions

The general form of boundary conditions for a second-order partial differential equation (such as the wave equation in Equation (2.2)) are specified as

$$\alpha \frac{\partial u}{\partial n} + \beta u + \gamma \dot{u} + \delta \ddot{u} + \epsilon = 0. \quad (2.38)$$

Commonly used boundary conditions are of the form

$$u = u_D, \quad (2.39)$$

where the condition is imposed on the variable itself, called Dirichlet conditions, and

$$\frac{\partial u}{\partial n} = u_N, \quad (2.40)$$

where the condition is imposed on the normal derivative of the variable, called Neumann conditions. Special (termed homogeneous) cases occur when the Dirichlet boundary condition is $u = 0$ in Equation (2.39) and when the Neumann condition is $u_N = 0$ in Equation (2.40), and therefore $\frac{\partial u}{\partial n} = 0$. These change the sign of the wave and perfectly reflect the wave respectively and are termed reflecting boundary conditions [103]. In the latter case, this represents a free surface. Furthermore, if one fixes the boundary to zero for displacement, one has a “sound hard” boundary, whereas if one fixes the pressure to zero one has a “sound

soft” boundary. In the finite element literature [104], the Dirichlet condition in (2.39) is termed “essential or strongly imposed”, owing to its explicit appearance in the definition of the variational form, whereas the Neumann condition in (2.40) is termed “natural or weakly imposed”, due to its implicit appearance in the variational form.

2.3.2 Numerical boundary conditions: modelling non reflecting boundaries

As the goal is to model physical setups, which are infinite and continuous, using numerical methods, one needs to impose numerical boundary conditions to describe the approximation of the infinite domain with the computational domain. In this section, different types of boundary conditions will be described, focusing on numerical boundary conditions to handle the transition described above. The Sommerfeld radiation condition [105], states that the asymptotic behaviour of the solution should correspond to a wave radiating outwards to infinity and this must be met in the numerical model. While some problems decay rapidly in space, solutions to the wave equation are generally oscillatory, thus only decay with respect to the following formula

$$\text{decay} \propto \frac{1}{r^{\frac{d-1}{2}}}, \quad (2.41)$$

where r is the distance and d is the dimension of the problem, as given in [106]. These problems arise when solving PDEs by volume discretisation methods (such as finite difference or finite element methods mentioned in Section 2.2).

Standard boundary conditions, such as Neumann or Dirichlet can lead to significant deviations in simulations from the real physical scenarios of interest as a result of reflections from the computational boundary. Therefore, the aim should be to minimise reflections from incident waves at the boundary of the discrete domain. Difficulties arise, as mentioned in [41], due to the fact that the boundary of the computational domain is artificial, not representing the physical reality. While theoretically a significantly larger domain could be modelled, until a significant decay has arisen, as in equation (2.41), practically this is not possible, as the

models grow too large, and thus cannot be solved computationally, even with the advances in numerical performance. One must also be aware that proposed methods must be applicable to common FE packages, without needing to modify the underlying source code which is not practical or often feasible.

In the numerical methods literature, boundary conditions can be split into three categories, those of “infinite element methods”, where elements simulating infinite space are placed around the discrete domain, those of “non-reflecting boundary conditions”, where additional conditions are added into the equations to prevent reflections, and those of “absorbing layer methods” detailed below. In this thesis the focus is on the case of absorbing layer methods in particular, as “infinite element methods” have been shown in [107, 108] to perform poorly for elasticity, and non-reflecting boundary conditions are not feasible to implement in standard FE codes. Absorbing layer methods for the problems of interest can be further split into two distinct categories, those of “perfectly matched layers” and “absorbing layers by increased damping”. In Figure 2.3, the absorbing boundary layer surrounding the “area of study” is shown.

Perfectly matched layer

The perfectly matched layer (PML) method is an artificial absorbing layer which surrounds the boundary of the domain (see Figure 2.3), absorbing incident waves, attenuating them and thus the wave decays exponentially. The method was first introduced by Berenger in [109] for Maxwell’s equations, splitting the solutions to the equation, into two “split sums”. Since then, frequent improvements have appeared in the literature, and for different wave equations or conditions, such as for elastic wave propagation [110], for acoustic waves [111] or for the seismic wave equation [112].

For reference, let us recall the basic equations underpinning the PML formulation. The essence of the method is the coordinate transformation as follows (shown here for the x-




Figure 2.3: Absorbing boundary (striped) surrounding the area of study, which is a subset of an infinite domain.

direction)

$$\frac{\partial}{\partial x} \rightarrow \frac{1}{1 + i \frac{\sigma_x(x)}{\omega}} \frac{\partial}{\partial x}. \quad (2.42)$$

Considering the coupled acoustic wave equation

$$\frac{\partial u}{\partial t} = - \frac{1}{\rho} \frac{\partial p}{\partial x}, \quad (2.43)$$

$$\frac{\partial v}{\partial t} = - \frac{1}{\rho} \frac{\partial p}{\partial y}, \quad (2.44)$$

$$\frac{\partial p}{\partial t} = - \rho c^2 \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right), \quad (2.45)$$

one can define the following attenuation coefficients

$$\sigma_x = \frac{x_i}{W_{PML}} \sigma_{max}, \quad (2.46)$$

$$\sigma_y = \frac{y_i}{W_{PML}} \sigma_{max}, \quad (2.47)$$

where W is the width, x_i and y_i are the distances to the inner boundary of the PML Layers,

and σ_{max} is the largest value of the attenuation coefficient. Performing the transformation as in [113] leads to

$$\frac{\partial u}{\partial t} + c\sigma_x u = -\frac{1}{\rho} \frac{\partial p}{\partial x}, \quad (2.48)$$

$$\frac{\partial v}{\partial t} + c\sigma_y v = -\frac{1}{\rho} \frac{\partial p}{\partial y}, \quad (2.49)$$

$$\frac{\partial p}{\partial t} + c(\sigma_x + \sigma_y)p = -\rho c^2 \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) - \rho c^2 \sigma_x \frac{\partial u}{\partial x} - \rho c^2 \sigma_y \frac{\partial v}{\partial y}. \quad (2.50)$$

Despite the performance of the PML, significant downsides are still inherent in the method. One major difficulty is that the effectiveness analytical coordinate stretching responsible for attenuating the waves decreases after discretisation [106]. The absorption is also angle-dependent, as one only takes into account the vertical distance from the boundary. Thus, it is ineffective in the absorption of evanescent waves [114]. Despite this inherent limitation, practically it is possible to circumvent the aforementioned issues by increasing the “thickness” of the PML. The most important downside is given in [115], and is the difficulty of implementation in existing finite element codes, due to it being a split-method, and is the reason it is not relied upon in this thesis.

Convolution-perfectly matched layer

The Convolution-PML method was first introduced in [116], as a method to handle the case of evanescent waves, by modifying the coordinate transform to introduce extra (frequency-dependent) terms. This requires additional computational memory to store these terms, when compared to the conventional PML method.

As in [117], by using the coordinate transformation for the “classical” PML method given in Equation (2.42), and by making s_x variable dependent on $a_x \geq 0$ and $k_x \geq 1$ one gets

$$s_x = k_x + \frac{d_x}{a_x + \omega}. \quad (2.51)$$

By setting $k_x = 1$ and $a_x = 0$, one returns to equation (2.42). Following the derivation in the paper [117], let us define

$$\hat{S}_x(t) = \mathcal{F}^{-1} \left(\frac{1}{S_x} \right), \quad (2.52)$$

and replace d_x by

$$d_{\hat{x}} = \hat{S}_x(t) * d_x. \quad (2.53)$$

By changing the form of equation (2.51) to the following

$$\frac{1}{S_x} = \frac{1}{k_x} - \frac{d_x}{k_x^2} \frac{1}{(d_x/k_x + a_x) + i\omega}, \quad (2.54)$$

and using the fact that the Fourier transform of the following functions $\delta(t)$, $H(t)$, the delta and Heaviside functions are

$$\mathcal{F}(\delta(t)) = 1, \quad \mathcal{F}(\exp^{-at} H(t)) = \frac{1}{a + i\omega}, \quad (2.55)$$

one can calculate the inverse Fourier transform similarly. Thus, when using equations (2.52) and (2.55) one can obtain the following value for \hat{S}_x

$$\hat{S}_x(t) = \frac{\delta(t)}{k_x} - \frac{d_x}{k_x^2} H(t) \exp \left(- \left(\frac{d_x}{k_x} + a_x \right) t \right), \quad (2.56)$$

and denoting the following substitution

$$\zeta_x(t) = - \frac{d_x}{k_x^2} H(t) \exp \left(- \left(\frac{d_x}{k_x} + a_x \right) t \right), \quad (2.57)$$

gives the following transformation for $\partial_{\hat{x}}$

$$\partial_{\hat{x}} = \frac{1}{k_x} \partial_x + \zeta_x(t) * \partial_x. \quad (2.58)$$

The first term in equation (2.58) is simply the existing spatial derivative modulo a constant. For the second term, the exact derivation is not reproduced here, although some advantages of its implementation will be described. Calculating the second term (introducing the time

step n) as

$$\psi_x^n = b_x \psi_x^{n-1} + a_x (\partial_x)^{n+1/2}, \quad (2.59)$$

where

$$b_x = \exp(-(d_x/k_x + a_x) \Delta t), \quad a_x = \frac{d_x}{k_x(d_x + k_x a_x)} (b_x - 1). \quad (2.60)$$

To compute each timestep in equation (2.58) one only needs to keep track of the previous iteration value and update it, which although requires additional memory, significantly reduces the computational power when computing the convolution, as we no longer need to iterate over all timesteps each time (one can also simply recompute the value over storing it, and depending on the specific problem constraints, this might be another alternative option).

Stiffness reduction method

Let us now discuss a stiffness reduction method (SRM), introduced in [118], which has comparable results to the PML method introduced in Section 2.3.2, while maintaining ease-of-implementation in finite element software.

The wave equation has the following spatial FE formulation [43]

$$[K]\mathbf{u} + [C]\dot{\mathbf{u}} + [M]\ddot{\mathbf{u}} = \mathbf{F}, \quad (2.61)$$

where $[K]$, $[C]$, $[M]$ are the stiffness, damping and mass matrices respectively, \mathbf{F} is the forcing term/source and \mathbf{u} , $\dot{\mathbf{u}}$, $\ddot{\mathbf{u}}$ are the displacement, velocity and acceleration vectors respectively.

The time-harmonic solution of equation (2.61) is

$$u(x, t) = \exp(i(kx - \omega t)). \quad (2.62)$$

Replacing the time derivatives in equation (2.61) with

$$\frac{\partial u}{\partial t} = -i\omega u, \quad \frac{\partial^2 u}{\partial t^2} = -\omega^2 u, \quad (2.63)$$

and inserting equation (2.63) into (2.61) gives

$$-[M]\omega^2 u - [C]i\omega u + [K]u = [F]. \quad (2.64)$$

Let us then describe the damping matrix, as a summation of the mass and stiffness matrix as follows

$$[C] = C_M[M] + C_K[K], \quad (2.65)$$

and now combine equations (2.63) and (2.64) to obtain

$$-[M] \left(1 + i\frac{C_M}{\omega}\right) \omega^2 + [K] (1 - i\omega C_K) u = [F]. \quad (2.66)$$

The damping introduced in (2.65) is termed Rayleigh damping [119]. The complex-valued density, ρ , and Young's modulus, E , is now defined as (with some minor abuse of notation)

$$\rho = \rho \left(1 + \frac{iC_M}{\omega}\right), \quad E = E (1 - i\omega C_K). \quad (2.67)$$

Furthermore, in the absorbing layers, the wavenumber is $k = k_r + ik_i$, for some real part k_r and some imaginary part k_i . When inserting into (2.62), this results in

$$u(x, t) = \exp(i(k_r x - \omega t)) \exp(-ik_i x). \quad (2.68)$$

Conventional FE software does not allow the alteration of the wavenumber, so this is achieved by altering the values of the FE matrices directly, increasing the value of k_i inside the absorbing region. As this is effectively altering the wavenumber, the CFL condition changes, decreasing the stable time increment. In [118], it is shown that a high value C_M causes only a small perturbation in the time increment as opposed to C_K . Thus, one can control damping purely by modifying C_M and letting $C_K = 0$. Let us now define

$$C_M(x) = C_{M_{max}} X(x)^p. \quad (2.69)$$

There is currently no known method of optimally choosing these parameters, so they are generally adjusted until the model performs as expected, based on a certain tolerance. The wavenumber can thus be increased by altering equation (2.69) or by decreasing the value of Young's modulus.

Caughey absorbing layer method

Another variation on the “classical” PML scheme encountered in literature is the Caughey PML scheme [120], which generalises the damping introduced in equation (2.65), as follows

$$[C] = [M] \sum_{j=0}^{m-1} a_j ([M]^{-1} [K])^j. \quad (2.70)$$

When $m = 2$, Caughey damping is equivalent to Rayleigh damping. This formulation is also easily computed in most FEM software, and has shown promise in types of problems involving elastic waves [120].

2.4 Summary

In this chapter the author has presented the theory of wave propagation in elastic media and subsequently detailed various analytical, numerical and combinations of these methods to solve the wave propagation problem as well as any of their disadvantages. In the following chapter, work done on deriving and implementing a finite element formulation to solve the wave propagation and scattering problem in domains involving both solids and fluids is detailed, while in subsequent chapters canonical numerical simulations carried out to verify this implementation are described.

Chapter 3

Development & Implementation of a novel finite element method for elastic wave phenomena

3.1 Introduction

Partial differential equations are ubiquitous throughout the scientific and engineering domains and are used to describe and therefore model key physical interactions and phenomena. Some important and well-known partial differential equations are the wave, the heat, the Laplace, and the Navier-Stokes equations. As mentioned in Chapter 1, analytical results are generally limited to simple geometries, which are not able to represent realistic practical scenarios often involving complex configurations. Despite invaluable physical insights offered by the analytical models, only numerical methods enable a full range of complex, industry-relevant configurations to be simulated. Among the most commonly used methods are the finite-difference method, described in detail in Chapter 2 and the finite element method, although depending on the specific problems and spatial domain one wishes to model, other numerical methods such as the finite volume method or the boundary element method also exist. For

my research and the remainder of this thesis, the focus is on the finite element method for a particular equation, the linear wave equation.

From a historical viewpoint, it is of no coincidence that the advent of the finite element method coincided with the advent of computing resources, as it is generally not possible to perform such calculations by hand. In [121], Wilson details some of the first software used to automate the finite element analysis of plane stress structures. Some of the first software packages can be traced back to NASTRAN (NASA Structure Analysis) [122], and SAP [123]. Nowadays, many software packages exist, both commercial and free. Some of the more well-known packages are Deal.ii [124], FEniCS [125], FreeFem++ [126], which are open-source, and Pogo [44], Abaqus [127] and COMSOL [128], which are commercial and proprietary. Various programming languages, such as Mathematica [129] or Matlab [130] also provide toolboxes to handle finite element problems.

Although many such packages exist, few are optimized and allow the user to make use of GPU acceleration. One such package is Pogo [44]. This also necessitates the development of new algorithms and implementations suitable for such GPU acceleration, which is the central research objective of this thesis.

In the first part of this chapter, the author will explore the literature on the finite element method and associated software packages. The second part will focus on the specific formulation and implementation that the author has developed for the problem of ultrasonic propagation and scattering in domains involving both solids and fluids. A full historical overview of the development of the finite element method will not be given in this thesis, although readers interested are referred to such works as [131].

3.2 Overview of the finite element method literature

As mentioned already, the finite element method is a numerical method well-suited for the solution of problems arising in complex, large-scale industry-relevant configurations. It is a classical and well established numerical method, with a rigorous theoretical mathematical basis, and the reader is referred to works such as [132, 133, 104, 134, 125] for further details. Rigorous mathematical definitions of a mesh, a finite element and a local-to-global mapping can be found in the aforementioned texts, although a full understanding of them is not required to use and implement the formulations and algorithms derived in this chapter.

The finite element method, and thus the associated software implementations, generally have the following structure. First, the area of interest is discretised into smaller finite elements, using some meshing algorithm implemented in software (such as the Delaunay triangulation [135] implemented in software such as Gmsh [136] or Pogo [44]). Then, one assembles local matrices on each element according to the underlying constitutive equations and associated material properties and uses a local-to-global mapping to assemble the global matrices. Finally, the resulting system of equations is solved and the solution is interpolated across the whole domain.

The research presented in this chapter largely focuses on the assembly and solution of the matrices detailed above, as one can rely on meshing functionality already present in software packages. Meshio [137], a software tool for mesh files details many of the common mesh formats for both structured and unstructured grids, as well as their efficiency, in both speed and storage, and the reader is referred to this for further details.

Much work has been done on the finite element method for the ultrasonic propagation problem involving domains comprised of both solid and fluid, often termed the fluid-structure interaction (FSI) problem. By investigating the solid-fluid interaction problem, the aim of this thesis is to understand how this can be implemented efficiently in software such as Pogo, without sacrificing on performance unnecessarily and taking into account restrictions due to the GPU-based nature of the package. Difficulties arise when imposing the solid-fluid

boundary conditions, and when different non-diagonal forms of damping are introduced.

In Chapter 1, some common formulations for the FSI problem have been detailed. For completeness, they will be complemented with additional formulations in the literature. Force-based formulations have also attracted some interest for structural analysis, first by Robinson [138], whose method can be found as the standard force method, then by Patnaik *et al.* [139] and then by Kaljević *et al.* [140] who improved upon the method and termed in the integrated force method. Despite some apparent advantages, it requires specialist software and a different mathematical formulation, and to the author's knowledge, no extensive studies have been carried out for the FSI problem using this method. Nevertheless, it may be a future avenue of research.

In Chapter 1, it was stated that displacement-based formulations for fluids may be the way forward although there are additional computational requirements compared to pressure-based formulations. While pressure is a scalar quantity, and thus only requires one degree of freedom per node, displacement is a vector-field corresponding to the dimension of the problem. As the size of the matrices in the finite element method have dimensions corresponding to $N \times N$, where N is the number of nodes multiplied by the degrees of freedom, one can see that the size of the problem may grow rapidly for large models. Despite this, by using new computing hardware, such as the GPU, which have increased performance for parallel mathematical applications, one can alleviate this. Furthermore, by storing only non-trivial nodal values in sparse matrix formats, which is a common practice in FE software, further savings can be made. Thus, the size of the model is no longer the main constraint for methods deployed to GPUs, given their excellent ability to handle large models when properly implemented [44, 141].

At the same time though, pressure-based formulations for the fluid generally result in non-symmetric matrices. The coupled fluid-solid finite element equation is given in [142] as

$$\begin{bmatrix} [M] & [0] \\ -\rho[A]^T & [Q] \end{bmatrix} \begin{bmatrix} \ddot{\mathbf{u}} \\ \ddot{\mathbf{p}} \end{bmatrix} + \begin{bmatrix} [B] & 0 \\ -\rho 0 & [C] \end{bmatrix} \begin{bmatrix} \dot{\mathbf{u}} \\ \dot{\mathbf{p}} \end{bmatrix} + \begin{bmatrix} [K] & [A] \\ 0 & [H] \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix} \quad (3.1)$$

where \mathbf{u} is the displacement vector in the solid, \mathbf{p} is vector of pressure values in the fluid. The matrices $[M]$ and $[Q]$ are the solid and fluid mass matrices respectively, $[K]$ is the solid stiffness matrix, $[H]$ is the fluid stiffness matrix, and $[B]$ and $[C]$ are the solid and fluid damping matrices. The vectors \mathbf{f}_1 and \mathbf{f}_2 are the forces in the solid and fluid regions respectively and ρ is the fluid density. The matrix $[A]$ is the area matrix, which is responsible for converting the fluid pressure at the interface to structural displacements, and its off-diagonal position is the reason for the non-symmetry in the final assembled matrices. It should be noted that this does not violate the Betti-Maxwell principle of reciprocity [143], as while the individual matrices $[M]$, $[Q]$, $[K]$, $[H]$, $[B]$, $[C]$, $[A]$ are symmetric,

$$\begin{bmatrix} [M] & 0 \\ -\rho[A]^T & [Q] \end{bmatrix} \text{ and } \begin{bmatrix} [K] & [A] \\ 0 & [H] \end{bmatrix}, \quad (3.2)$$

are not. It was shown in [60], that by reformulating the pressure-based approach given in (3.1) as a velocity-potential-based approach, one can regain symmetric matrices. The non-symmetry of the coupled pressure-displacement based formulation is in contrast to displacement-based formulations in which the matrices are symmetric. The non-symmetry and need for matrix inversion also pose significant challenges to GPU implementation. We are no longer guaranteed a sparse matrix; this results in a significant amount of increased memory and storage, and as we are often bandwidth limited, as discussed in [44], this can drastically slow down calculations, introducing additional bottlenecks. Furthermore, methods must be applicable to common FE software, which in general is written for displacement-based approaches. Nevertheless, one can always go from an approach with higher number of degrees of freedom to one of lower number of degrees of freedom by setting certain variables to be zero. Thus, the symmetric nature of the problem, as well as desirability to couple with existing displacement-based solid FE software, make modelling the fluid using a displacement field a very beneficial proposition, despite the fact that displacement-based approaches can suffer from hourglassing [144, 3].

3.3 The finite element formulation

3.3.1 Finite element formulation for solid materials

FE formulations for displacement-based solid elements are well-known [131] and are briefly discussed below so that the necessary changes made for fluid materials can be introduced. The local stiffness matrix $[K]^e$, where the superscript is to denote “local element”, of an element is given as

$$[K]^e = \int [B]^T [E] [B] dV. \quad (3.3)$$

The matrix $[B]$ in Equations (3.3) is the strain-displacement matrix which comes from the strain-displacement formula as

$$\boldsymbol{\epsilon} = [B]\mathbf{u}. \quad (3.4)$$

These are well known in the literature and depend on the choice of element, for which further information will be provided in Section 3.3.2. The matrix $[E]$ is the stress strain matrix which depends on the material parameters and for an isotropic material is given as [131]

$$[E] = \frac{E}{(1+v)(1-2v)} \begin{bmatrix} 1-v & v & v & 0 & 0 & 0 \\ v & 1-v & v & 0 & 0 & 0 \\ v & v & 1-v & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1-2v}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1-2v}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1-2v}{2} \end{bmatrix}, \quad (3.5)$$

where E is the Young’s modulus and v is Poisson’s ratio.

For the mass matrix $[M]$, a diagonal lumped mass matrix [145] is considered, of the form

$$[M] = \begin{bmatrix} m_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & m_N \end{bmatrix}, \quad (3.6)$$

which is purely diagonal and where m_i correspond to the element mass (area in 2D, volume in 3D).

For the damping matrix, mass proportional damping is assumed (this is equivalent to Rayleigh damping introduced in Equation (2.65) with the appropriate choice of constants) of the form

$$[C] = \alpha[M]. \quad (3.7)$$

Using the diagonal forms of the mass and damping matrices has the computational advantages of being easy to compute, store and invert.

For the time derivative, one can use the central finite difference method [92] to arrive at

$$[M] \frac{U_{t+1} - 2U_t + U_{t-1}}{\delta t^2} + [C] \frac{U_{t+1} - U_{t-1}}{2\delta t} + [K]U_t = F, \quad (3.8)$$

which has an error of $\mathcal{O}(\delta t^2)$. Here, U_{t-1} , U_t and U_{t+1} are the displacement vector values at the previous, current and next time steps respectively. By re-arranging equation (3.8) one arrives at

$$\left([M] \frac{1}{\delta t^2} + [C] \frac{1}{2\delta t} \right) U_{t+1} = F + \left([M] \frac{2}{\delta t^2} - [K] \right) U_t + \left([C] \frac{1}{2\delta t} - [M] \frac{1}{\delta t^2} \right) U_{t-1}. \quad (3.9)$$

Due to the choice of the diagonal damping matrix $[C]$ and the diagonal mass matrix $[M]$, the expression $\left([M] \frac{1}{\delta t^2} + [C] \frac{1}{2\delta t} \right)$ that appears in equation (3.9) can easily be inverted as it is purely diagonal. The matrices $[K]$, $[M]$, $[C]$ and \mathbf{F} can be calculated and stored before the time-marching begins and do not need to be recomputed at each time step.

3.3.2 A novel finite element formulation for fluid materials

The derivation of the finite element matrices for fluid materials, which is based on the constitutive equation for a compressible fluid, is presented. The formulation and assembly routines described in this section, and in Section 3.4 are comparable for both two and three-dimensional physical and computational domains. Therefore, they will be demonstrated in the simpler two-dimensional case, while making note of any differences where necessary. Computational examples and some further discussion on the three-dimensional case will be shown in subsequent chapters.

The constitutive equation for a compressible Stokesian fluid with Newtonian viscosity alongside a full derivation can be found in [19]. Although the full derivation is out of scope for this thesis, some necessary equations and their context are provided in what follows. The constitutive equation for a compressible fluid is given as [19]

$$\sigma_{ij} = -p\delta_{ij} - \frac{2}{3}\dot{\mu}\dot{\epsilon}_{kk}\delta_{ij} + 2\dot{\mu}\dot{\epsilon}_{ij}, \quad (3.10)$$

where σ_{ij} is the stress tensor, p is the equilibrium pressure, $\dot{\mu}$ is the dynamic viscosity of the fluid, ϵ is the strain, $\dot{\epsilon}$ is the strain rate and δ_{ij} the Kronecker delta. Considering the static part of the equation one arrives at (this also follows from the isotropy of the fluid)

$$\sigma_{xx} = \sigma_{yy} = \sigma_{zz} = -p = \lambda(\epsilon_{xx} + \epsilon_{yy} + \epsilon_{zz}). \quad (3.11)$$

Here, λ is the bulk modulus of the fluid, which is related to the acoustic speed of the fluid, c , and its density, ρ , as

$$\lambda = \rho c^2. \quad (3.12)$$

Thus, without loss of generality, over a 2-dimensional $x - y$ plane, the stress-strain and

stress-strain rate relationships are

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} \lambda & \lambda & 0 \\ \lambda & \lambda & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{xy} \end{bmatrix} + \begin{bmatrix} -\frac{4}{3}\dot{\mu} & \frac{2}{3}\dot{\mu} & 0 \\ \frac{2}{3}\dot{\mu} & -\frac{4}{3}\dot{\mu} & 0 \\ 0 & 0 & -\mu \end{bmatrix} \begin{bmatrix} \dot{\epsilon}_{xx} \\ \dot{\epsilon}_{yy} \\ \dot{\epsilon}_{xy} \end{bmatrix}. \quad (3.13)$$

The matrices defined in Equation (3.13) are then used to calculate D and \dot{D} which are needed in the stiffness and element damping matrices used in the finite element assembly for the fluid material. These are

$$D = \begin{bmatrix} \lambda & \lambda & 0 \\ \lambda & \lambda & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \dot{D} = \begin{bmatrix} -\frac{4}{3}\dot{\mu} & \frac{2}{3}\dot{\mu} & 0 \\ \frac{2}{3}\dot{\mu} & -\frac{4}{3}\dot{\mu} & 0 \\ 0 & 0 & -\mu \end{bmatrix}. \quad (3.14)$$

The local element stiffness and local element damping matrices can then be assembled (in parallel, without needing to recalculate the matrix $[B]$ for both $[K]$ and $[C]$) as follows

$$[K]^e = [B]^T [D] [B], \quad (3.15)$$

$$[C]^e = [B]^T [\dot{D}] [B]. \quad (3.16)$$

The mass matrix retains its form from the solid case, i.e. it is proportional to element mass on the diagonal. Matrix $[B]$ in Equation (3.15) is the strain-displacement matrix introduced in Equation (3.4). These are well-known in the literature [146] and depend on the choice of element and the numerical quadrature scheme.

Quadrature rules allow us to numerically approximate integrals of a function using a sum of weight and function values at specific points. Gauss-Legendre quadrature rules (commonly referred to as Gaussian quadrature) allow us to approximate this integral using polynomials and then integrating that polynomial exactly [147]. Thus the integral of $f(x)$ can be approximated using

$$\int_{-1}^1 f(x) dx \approx \sum_{i=1}^n w_i f(x_i), \quad (3.17)$$

where w_i are the weights and x_i are the specified points. It remains to calculate these weights and corresponding points. Much research has been carried out to efficiently calculate these and improve upon them, such as in [148, 149, 150]. It should be noted that in Equation (3.17) the domain of the integral is given as $[-1, 1]$, but one can always perform a change of interval from $[a, b]$ to $[-1, 1]$ to match this. For the commonly used Gaussian quadrature rules in the finite-element literature for the elastodynamic problem, these are $x_i = 0, w_i = 2$ for the first order Gaussian quadrature scheme, and $x_i = \pm \frac{1}{\sqrt{3}}, w_i = 2$ for the second order Gaussian quadrature scheme.

The formulation for the under-integrated or reduced integration quadrilateral element is now shown, i.e. the element with one point of integration in the centre. This has been demonstrated in [50] to reduce, alongside irrotational constraints or damping, spurious modes which are a particular issue for displacement-based fluid formulation [144]. For an isoparametric quadrilateral with diagonal nodes at points $[-1, -1]$ and $[1, 1]$ with shape functions

$$\mathbf{N}^e = \frac{1}{4} \begin{bmatrix} (1+x)(1-y) \\ (1-x)(1-y) \\ (1+x)(1+y) \\ (1-x)(1+y) \end{bmatrix}, \quad (3.18)$$

the strain-displacement matrix $[B]$ is given as

$$[B] = \begin{bmatrix} \frac{\partial N_1^e}{\partial x} & 0 & \frac{\partial N_2^e}{\partial x} & 0 & \frac{\partial N_3^e}{\partial x} & 0 & \frac{\partial N_4^e}{\partial x} & 0 \\ 0 & \frac{\partial N_1^e}{\partial y} & 0 & \frac{\partial N_2^e}{\partial y} & 0 & \frac{\partial N_3^e}{\partial y} & 0 & \frac{\partial N_4^e}{\partial y} \\ \frac{\partial N_1^e}{\partial y} & \frac{\partial N_1^e}{\partial x} & \frac{\partial N_2^e}{\partial y} & \frac{\partial N_2^e}{\partial x} & \frac{\partial N_3^e}{\partial y} & \frac{\partial N_3^e}{\partial x} & \frac{\partial N_4^e}{\partial y} & \frac{\partial N_4^e}{\partial x} \end{bmatrix}. \quad (3.19)$$

If one considers the single integration point to be from the first order Gaussian quadrature scheme at point $[0, 0]$ defined above, which is commonly used for bilinear quadrilateral

elements with reduced integration [148], then

$$[B] = \frac{1}{4} \begin{bmatrix} -1 & 0 & 1 & 0 & 1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 & 0 & 1 & 0 & 1 \\ -1 & -1 & -1 & 1 & 1 & 1 & 1 & -1 \end{bmatrix}. \quad (3.20)$$

Having derived the formulae for assembling the necessary matrices, we now recall the spatial FE formulation from Equation (2.61) and its finite difference in time formulation from Equation (3.9). One challenging step becomes immediately apparent - Equation (3.9) requires inverting $\left([M]\frac{1}{\delta t^2} + [C]\frac{1}{2\delta t}\right)$. It is no longer possible to execute this efficiently owing to the non-diagonal structure of the damping matrix $[C]$ used for the fluid elements. Since the GPU FE method works best for explicit time domain methods [91], a new approach must be developed. A modification of the central finite difference scheme is proposed and implemented:

$$[M]\frac{U_{t+1} - 2U_t + U_{t-1}}{\delta t^2} + [C]\frac{U_t - U_{t-1}}{\delta t} + [K]U_t = F. \quad (3.21)$$

The difference in Equation (3.21) compared to Equation (3.8) is that the first order time derivative multiplied by C is changed to depend purely on values at U_t and U_{t-1} which are known. Thus, \mathbf{U}_{t+1} can be rearranged to

$$\mathbf{U}_{t+1} = \delta t^2 [M]^{-1} (\mathbf{F} - [K]\mathbf{U}_t) + 2\mathbf{U}_t - \mathbf{U}_{t-1} + \delta t [M]^{-1} [C] (\mathbf{U}_{t-1} - \mathbf{U}_t). \quad (3.22)$$

As in [44], this is computed as

$$\mathbf{U}_{t+1} = \mathbf{F}' + [C']\mathbf{U}_{t-1} + [K']\mathbf{U}_t, \quad (3.23)$$

where the modified matrices have been introduced

$$[K'] = -\delta t^2 [M]^{-1} [K] + 2[I] - \delta t [M]^{-1} [C], \quad (3.24)$$

$$[C'] = -[I] + \delta t [M]^{-1} [C], \quad (3.25)$$

and the vector

$$\mathbf{F}' = \delta t^2 [M]^{-1} \mathbf{F}. \quad (3.26)$$

Now, as in [44], $[K']$, $[C']$ and \mathbf{F}' remain constant, they can be assembled, stored and used for all time steps. The modified numerical scheme introduced and described above in Equation (3.21) has the advantage that it is explicit in time, without relying on inversion of the matrix $[C]$, and is suitable for GPU calculations. The stability of the method remains the same as it occupies the same stencil in time [41], requiring the same number of nodal points. Thus, it has the same Courant-Friedrichs-Lewy condition [92] as the standard central difference scheme i.e.

$$c \frac{\Delta t}{\Delta x} < 1. \quad (3.27)$$

The order of the error of this method related to the modified numerical calculation of the velocity increases from $\mathcal{O}(\delta t^2)$ to $\mathcal{O}(\delta t)$ due to the use of the current time step instead of the future one.

3.4 Algorithms and software implementation


To summarise, so far the author has introduced and derived a mathematical finite element formulation for fluid materials which damps out the shearing motion through the use of a non-diagonal damping matrix alongside the use of a reduced-integration Gaussian quadrature scheme and is displacement based. It is thus compatible with existing software designed for solid materials and explicit in time and suitable for GPU calculations. For completeness, algorithms and some notes of their efficient implementation will be given for the common operations needed to assemble and solve the equations given above. It is appreciated that different software packages such as Pogo or Comsol may make different decisions in the algorithms than those presented here, but they are sufficiently generic to be useful to any practitioners wishing to implement finite element algorithms in general and more specifically, for the formulations derived in this thesis.

3.4.1 Computational Domain


The first step of using a finite element method software is defining and importing the computational domain, i.e. the mesh, into our software. This can be done with a variety of tools detailed earlier in this chapter, such as Gmsh [136], meshio [137] or Pogomesh [44]. While different software implementations provide additional information, necessary information needed are the position of the nodes and their corresponding numbering, a connectivity matrix detailing which nodes form each finite element, the material parameters (and therefore material type of each element, whether solid or fluid in our case) and the associated quadrature rules for each element, which are usually referred to by the element type (reduced integration or fully integrated for the four-noded quadrilateral). An extensive list of different element types can be found in [151], where Arnold and Logg introduce a “periodic table of finite elements”, although not all of them are applicable to the wave equation.

Grids can be further separated into structured and unstructured grids, examples of which are in Figure 3.1a and Figure 3.1b respectively. These have associated advantages and disadvantages, which are briefly detailed here. Unstructured grids conform to the object geometry, which is particularly important for complex real-world objects; they allow further refinement in areas of interest and thus can generally provide us with more accurate results. Structured and equispaced/uniform grids, do not necessarily conform to the object geometry, and thus may suffer from staircasing effects depending on the object geometry. One might also have structured and non-equispaced/uniform grids; these are usually used in cylindrical coordinates, when discretising an object such as a pipe, and these can also suffer from such staircasing issues.


If one wants to represent the circle in Figure 3.2a as a different material, without loss of generality assume this is solid and the remaining mesh is of fluid, the result of this can be found in Figure 3.2b. As seen, the purple shading is correct, the blue shading is simulated as a solid when it represents a fluid material and vice-versa, the pink shading represents a fluid material when it should be solid, thus both overestimating and underestimating the



(a) Structured rectangular mesh



(b) Unstructured mesh, conforming to a cylindrical shell



(c) Unstructured mesh, conforming to a cylindrical shell, using structured mesh elements for the majority of the domain


Figure 3.1: Structured, unstructured, and combined meshes.

true domain. Despite this, they are significantly faster to create, simpler to understand and cheaper to store in memory.


Furthermore, as will be discussed in the assembly routine section, the total number of elements and nodes is directly related to the computational time to solve a problem, and overall for the same discretization, structured grids produce fewer nodes and elements than the comparable unstructured grid. For this reason, the majority of simulations conducted in this thesis will use structured grids unless otherwise noted. It should be noted that one can combine these two methods of structured and unstructured grids, by defining the majority of the computational domain using structured grids and only using unstructured grids around the complex object. Finite element software and their associated meshing functionality such as Pogo provides this. An example of such a mesh can be found in Figure 3.1c.

3.4.2 Matrix assembly

In Equation (2.61), the matrices $[K]$, $[C]$, $[M]$ were defined as the stiffness, damping and mass matrices respectively, and Equations (3.15) and (3.16) gave the equations for $[K]^e$, $[C]^e$ the local stiffness and damping matrices for a fluid element. For the local stiffness and damping matrices, one needs to calculate the strain-displacement matrix $[B]$, which in our derivation relied on the shape function with diagonal nodes at points $[-1,-1]$ and $[1,1]$ and a single integration point, based of a first order Gaussian quadrature scheme at point $[0,0]$. This is called the reference element, and the mesh element is related to this by a map or function G [152]. Figure 3.3 shows this mapping from the local or mesh element to the reference element. One can construct the shape function on this reference element and then through a coordinate change map to the original element. The reason we use them is to keep all the formulas constant, irrespective of the size and position of each element, reducing computational complexity. It should be noted that a coordinate transformation with the associated Jacobian must be calculated, but this poses minimal computational challenges nowadays.



(a) Circular domain (pink) I wish to represent using a structured grid.



(b) Staircasing results, blue represents the computational domain, both over-fitting and under-fitting

Figure 3.2: True object and computational approximation in pink and blue respectively.




Figure 3.3: Mapping of local (left) to reference element (right).

Formulations for the local matrices have been derived and use a “local-to-global” mapping, which provides us with information about how each local degree of freedom relates to the global degree of freedom. Mathematically, this can be described as

$$[K] = \sum_{e=1}^N G([K]^e), \quad (3.28)$$

where G is the “local-to-global” mapping function and the summation is over the elements of the computational domain. It is noted that each matrix position in the global matrix can have a contribution from more than one local matrix.

For efficient assembly of the matrices, the numbering of the mesh nodes must be chosen in such a way that it minimizes the bandwidth of the assembled matrices, but this is well-known in the literature and the reader is referred to papers such as [153] by Rosen, or [154] by Cuthill and Mckee for further information.

3.4.3 Matrix storage

The formulations for the finite element matrices, assembly routines and how one can efficiently assemble the system in parallel using a “local-to-global” mapping have been described above. The dimensions of the local matrices $[K]^e, [M]^e, [C]^e$ are theoretically $(nk)^2$ where n is the number of nodes on the reference finite element ($n = 4$ for tetrahedrons, $n = 3$ for triangles, $n = 8$ for the common three-dimensional case of a brick element) and k is the number of degrees of freedom associated with the formulation. As we are now in a displacement-based setting, this is $k = 2$ for two-dimensional simulations and $k = 3$ for three-dimensional simulations. For a pressure-based formulation $k = 1$ in all dimensions, although this introduces other complexities previously discussed. The dimensions of the global matrices $[K], [M], [C]$ are $(Nk)^2$, where N is the total number of nodes in the discretised domain. These matrices must also be stored efficiently, as for large systems, such as those simulated in [155], one would require enormous amounts of storage, which is not feasible.

One thus relies on sparse matrix functionality, where only non-zero elements of the matrices are stored, compared to dense matrix formats, where all the information is stored. This functionality is well suited to finite element algorithms. For the mass matrix defined in Equation (3.6), this can be stored in a diagonal form, keeping only the values on the diagonal. Due to the finite elements algorithms used here, it is known that the resulting global stiffness and damping matrices will be sparse, as non-trivial values exist only in locations with adjacent nodes. Without loss of generality, for the i^{th} row values are non-trivial at columns associated with nodes which are part of the same elements as i [44]. An illustration of this is in Figure 3.4. The symmetry of the matrix, which is a corollary of the Betti-Maxwell principle of reciprocity [143] can be seen here. If node i is in the same element as node j , then clearly node j is also in the same element as node i . Thus the storage complexity for the stiffness, damping and mass matrix reduces from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$, allowing us to more readily solve large-scale industrially relevant simulations.

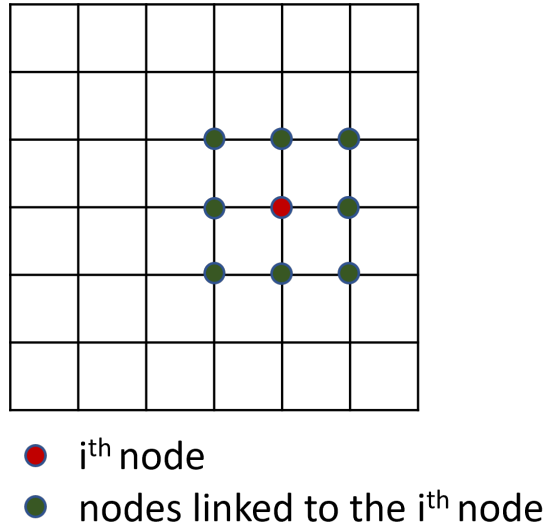


Figure 3.4: Node linking from the i^{th} node.

3.5 Computational performance of the method

Direct timing comparisons between the formulation and implementation of the method on a GPU described in this thesis and existing finite element (FE) software for fluid-solid interaction are not straightforward. This is due to FE software packages being often heavily dependent on not just the formulation used, but also the hardware used. The choice of programming language and compiler optimisations [156] also significantly impact the performance of the algorithms. Furthermore, there is frequent improvement in all aspects of computational mathematics [157] and hardware, which can complicate direct timing comparisons. The formulation in this thesis will benefit from such computational improvements.

Despite the above, it is appreciated that timing comparisons can help validate the proposed performance improvements claimed by this method. Therefore, as we know the timing ratio of Pogo Solid to conventional CPU-based FE software, if we calculate the ratio of the fluid-solid implementation to the solid implementation in Pogo (the method presented in this thesis), we can indirectly compare the GPU-based method presented in this thesis to CPU-based methods. A mathematical description of this can be as follows,

$$\frac{\text{Pogo Fluid-Solid GPU}}{\text{Conventional CPU}} = \frac{\text{Pogo Fluid-Solid GPU}}{\text{Pogo Solid GPU}} \times \frac{\text{Pogo Solid GPU}}{\text{Conventional CPU}}. \quad (3.29)$$

This way of calculating the performance difference also has the benefit of being run on the same underlying hardware for a more accurate comparison.

We setup a pair of simulations as follows: We define a square grid of rectangular elements. For the first simulation, this is fully solid and for the second simulation we set the right-half to be fluid. We progressively increase the size of the square grid to obtain measurements across a wide-range of total number of degrees of freedom. We run the simulation at grid size three times and take the average, to mitigate any anomalous results inherent in timing computing processes.

The simulations in this study are executed using two modes offered by Pogo [44]. These modes are denoted as “Compression on (default)” and “Compression off.” Compression, as defined in the Pogo documentation, refers to the optimization technique where large sections of uniform elements within a model exhibit redundancy in the stiffness matrix, K . In such scenarios, many coefficients in the matrix are identical. To reduce the storage required, Pogo has the capability to detect these uniform elements and optimize the problem by storing these coefficients once and pointing to them when needed. This approach results in a reduction in memory usage and often increased computational speed.

The results of these two measurements can be found in Figure 3.5, where we plot the ratio of Pogo-Fluid-Solid to Pogo-Solid. We can immediately see a stark contrast between the two images. In Figure 3.5a, where we use compression, we see only a marginal slow down between using the solid-fluid executable and the solid executable of approximately 2%. Whereas in Figure 3.5b, where we do not use compression, we can see a significant slow down, which also increases as the number of degrees of freedom increase. There are a few observations we can make here, although a full understanding of this difference is likely to be hardware and implementation specific, and would require significantly more analysis which is outside the scope of this work. The first observation we can make, is that despite the additional computation required by the method described in this thesis, this does not introduce significant additional slowdown compared to the GPU FE method for solid materials introduced in [44] and has been shown to be significantly more performant than comparable CPU based


methods. The second observation relates to when compression is turned off; as results are no longer cached and data must be transferred and stored for both the stiffness and damping matrices, the performance diverges and the solid-fluid case is much more expensive. As the number of degrees of freedom increases, this becomes more pronounced.

Despite the differences between the two cases we described above, we can indirectly state that the method described in this thesis will be between 70 and 100 times faster than an equivalent commercial CPU FE software, based on the solid GPU FE method in [44] being 100 times faster, using (3.29).

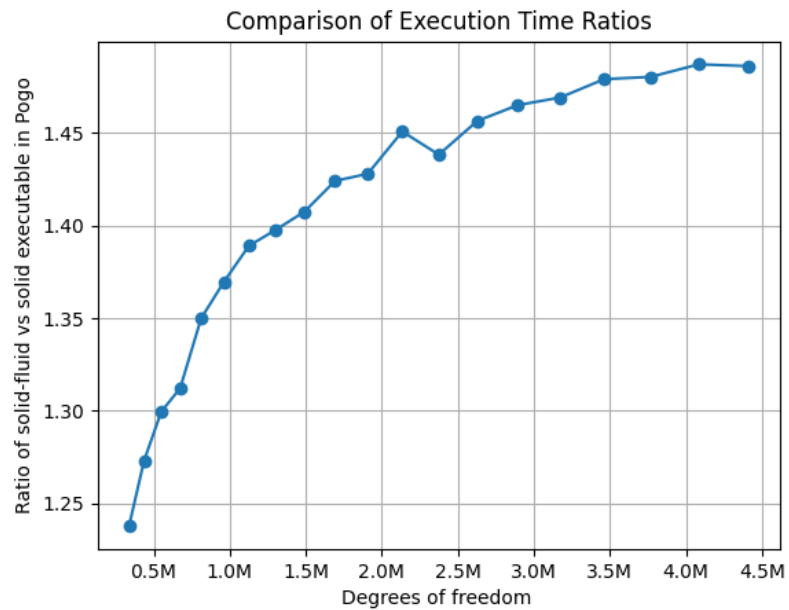
3.6 Summary

The author has provided formulations for the matrices, described how to calculate these formulations, and also stated how they can be stored. The author has not provided an explicit implementation of these, as they are heavily software-dependent and many different implementations exist, from meshing software to sparse matrix functionality and their associated linear system solvers.

Running GPU FE software has been demonstrated to offer significant improvements, even when accounting for the change in hardware, and the author refers the reader to such implementations for further reference [44]. For the examples that follow, the author uses the formulation as described in this chapter, implemented in Pogo [44], although the proposed method is sufficiently generic to be implemented in any other FE software.



(a) Compression on (default)



(b) Compression off

Figure 3.5: Ratio of solid-fluid vs fluid executable in Pogo run with (a) compression on and (b) compression off.

Chapter 4

Solving Problems I: Application to Underwater Acoustics

The scattering from rigid and elastic objects submerged in fluid has long been of interest to the underwater acoustics community, who use this scattering response to identify objects proud or buried in the seabed [158], to calculate target-echo strength measurements of such objects [159], or for defence related reasons [1].

The objective of the simulations in this chapter is twofold; verify and validate the theoretical finite element formulation derived in Chapter 3 and show that this formulation can be applied to problems in the underwater acoustics area. Canonical setups in the underwater acoustics community will be looked at, such as those mentioned in Chapter 1 by [54, 56, 55, 158]. The objects of interest can be represented by cylinders, spheres, and cylindrical shells, with the latter further being empty, partially-filled or fully-filled. All of these objects can then be present in the free-field in the fluid, proud on a sediment interface or partially or fully buried. In this chapter, results will be shown for scattering from some of these objects in some configurations in 2D and 3D, which can be found in Table 1.1. The results will be compared to existing analytical or semi-analytical solutions where these exist and in the cases they do not, informed observations and comments regarding the validity of the results will be made.

All the simulations presented in this thesis have relied on the (elastic) orthogonal hourglass control method described in [3] which relies on using an additional artificial stiffness to reduce this phenomenon. This method has been implemented in FE software such as [127] or [44]. As we are using the displacement-fluid based formulation formulated in Pogo for these simulations, we have access to this hourglass control for "free" i.e. it is not needed to re-implement it, demonstrating the versatility of the method and formulation described in this thesis.

An hourglass control parameter of 0.05 is used. Interestingly, we observe that, unlike solid materials, fluid materials demonstrate less hourglass effects when using reduced-integration as opposed to full integration. A full investigation of why this happens is out of the scope of this thesis and could be an interesting avenue for future research. The works of Chen and Taylor [50] who use a reduced integration formulation with a projected mass approach could be a good starting point.

In what follows, the damping parameter formulated in Chapter 3 has been set to zero, i.e. it is not strictly necessary to achieve accurate simulations. The reason for this is not currently known, but is possibly related to the elastic damping present when we use hourglass control from [3]. It has been kept in the formulation to more easily enable extension of this work to viscous fluids and damping, and is a possible avenue for future research.

4.1 Scattering from cylindrical structures

We now proceed to compute FE simulations and compare with analytical scattering results across three separate related physical cases. The results in this chapter will be given in terms of radial displacement, defined as

$$u_r = u_x \cos \theta + u_y \sin \theta. \quad (4.1)$$

An illustration of this can be found in Figure 4.1. To achieve scattered results numerically, the following relation is used

$$u_r^{tot} = u_r^{inc} + u_r^{scatt}, \quad (4.2)$$

where u_r^{tot} is the total displacement, u_r^{inc} is the displacement of a homogeneous field (i.e. where no scatterer is present) and u_r^{scatt} is the scattered displacement. To obtain this numerically, two FE simulations are computed; one with the scatterer, and one with a homogeneous field, giving us u_r^{tot} and u_r^{inc} , allowing us to calculate u_r^{scatt} from (4.2) As our mathematical formulation and our software implementation are displacement-based, the response for the u_x and u_y are readily available.




Figure 4.1: Radial displacement.

In all simulations in this section, an excited plane wave travels across the domain simulated and response nodes are placed at a distance from the centre of the scatterer, details of which are given when describing the particular simulation. Due to the nature of the FE method, the response in the far field can be computed (arbitrary distances can be measured given

appropriate computational power and storage). Alternatively, one can calculate the response in the near field and then propagate the result to the far field, although this is not done in this thesis. In all cases, a Courant-Friedrichs-Lewy [92] number of 0.8 with respect to the highest velocity in the simulation (usually the solid longitudinal wave speed c_L) is used to determine the time step

$$dt = \frac{dx}{c_L} \times \text{Courant}, \quad (4.3)$$

and the number of time steps is set to allow enough time for the wave to pass across the domain and back with reference to the fluid wavespeed. For a quantifiable comparison, the relative error defined as

$$\frac{\|u_{true} - u_{FE}\|_1}{\|u_{true}\|_1}, \quad (4.4)$$

is used to compare the measured FE scattering versus the analytical solution, where $\|u\|_1$ is the ℓ_1 vector norm of u given by

$$\|u\|_1 = \sum_{i=1}^n |u_i|, \quad (4.5)$$

i.e. the sum of the absolute vector values.

Although the simulations were carried out in the time-domain, to allow for a clear comparison between the analytical case, which is generally given in the frequency domain for single-frequencies, the frequency result closest to the central frequency excitation is extracted with the use of fast-Fourier-transforms. The analytical expression for the pressure in the scattered field for infinitely long cylindrical structures can be found in [55], and is given as

$$p_r^{scatt} = \sum_{n=0}^{\infty} i^n \epsilon_n b_n H_n^{(1)}(k_{ext}r) \cos n\theta. \quad (4.6)$$

From this, one can calculate the radial derivative to give us the scattered displacement as

$$u_r^{scatt} = \sum_{n=0}^{\infty} i^n \epsilon_n b_n k_{ext} H_n^{(1)'}(k_{ext}r) \cos n\theta, \quad (4.7)$$

where ϵ_n is the Neumann factor defined as $\epsilon_0 = 1$, $\epsilon_n = 2$, for $n > 0$. The coefficients b_n are used to define the scattered fields for each scatterer and are well-known in the literature for

specific configurations, details of which will be given below. The functions $H_n^{(1)}$, $H_n'^{(1)}$ are the Hankel function of the first kind and its respective derivative, k_{ext} is the radial wavenumber in the exterior fluid, r is the distance from the centre of the cylindrical scatterer and θ is the angular position.

For an elastic cylindrical shell the general form of b_n is given as [55]

$$b_n = \frac{1}{D} \begin{vmatrix} \beta_1 & \alpha_{12} & \alpha_{13} & \alpha_{14} & \alpha_{15} & 0 \\ \beta_2 & \alpha_{22} & \alpha_{23} & \alpha_{24} & \alpha_{25} & 0 \\ 0 & \alpha_{32} & \alpha_{33} & \alpha_{34} & \alpha_{35} & 0 \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} & \alpha_{45} & \alpha_{46} \\ 0 & \alpha_{52} & \alpha_{53} & \alpha_{54} & \alpha_{55} & \alpha_{56} \\ 0 & \alpha_{62} & \alpha_{63} & \alpha_{64} & \alpha_{65} & 0 \end{vmatrix}, \quad (4.8)$$

where D is given by

$$D = \begin{vmatrix} \alpha_{11} & \alpha_{12} & \alpha_{13} & \alpha_{14} & \alpha_{15} & 0 \\ \alpha_{21} & \alpha_{22} & \alpha_{23} & \alpha_{24} & \alpha_{25} & 0 \\ 0 & \alpha_{32} & \alpha_{33} & \alpha_{34} & \alpha_{35} & 0 \\ 0 & \alpha_{42} & \alpha_{43} & \alpha_{44} & \alpha_{45} & \alpha_{46} \\ 0 & \alpha_{52} & \alpha_{53} & \alpha_{54} & \alpha_{55} & \alpha_{56} \\ 0 & \alpha_{62} & \alpha_{63} & \alpha_{64} & \alpha_{65} & 0 \end{vmatrix}, \quad (4.9)$$

and the coefficients α_{ij} and β_i (for convenience of notation, the subscript n which refers to the series summation found in Equations (4.6)-(4.7) has been dropped) can be calculated from

$$\begin{aligned} \beta_1 &= \left(\frac{x_{int}}{k_{int}} \right)^2 J_n(x_{int}), \\ \beta_2 &= x_{int} J_n'(x_{int}), \end{aligned} \quad (4.10)$$

and

$$\begin{aligned}
\alpha_{11} &= - \left(\frac{x_{int}}{k_{int}} \right)^2 H_n^1(x_{int}), \\
\alpha_{12} &= 2\mu x_L^2 J_n''(x_L) - \lambda x_L^2 J_n(x_L), \\
\alpha_{13} &= 2\mu x_L^2 N_n''(x_L) - \lambda x_L^2 N_n(x_L), \\
\alpha_{14} &= 2\mu n [J_n(x_{Sh}) - x_{Sh} J_n'(x_{Sh})], \\
\alpha_{15} &= 2\mu n [N_n(x_{Sh}) - x_{Sh} N_n'(x_{Sh})], \\
\alpha_{21} &= -x_{int} H_n'^{(1)}(x_{int}), \\
\alpha_{22} &= -\rho_{ext} \omega^2 x_L J_n'(x_L), \\
\alpha_{23} &= -\rho_{ext} \omega^2 x_L N_n'(x_L), \\
\alpha_{24} &= \rho_{ext} \omega^2 n J_n(x_{Sh}), \\
\alpha_{25} &= \rho_{ext} \omega^2 n N_n(x_{Sh}), \\
\alpha_{32} &= 2n [x_L J_n'(x_L) - J_n(x_L)], \\
\alpha_{33} &= 2n [x_L N_n'(x_L) - N_n(x_L)], \\
\alpha_{34} &= -x_{Sh}^2 J_n''(x_{Sh}) + x_{Sh} J_n'(x_{Sh}) - n^2 J_n(x_{Sh}), \\
\alpha_{35} &= -x_{Sh}^2 N_n''(x_{Sh}) + x_{Sh} N_n'(x_{Sh}) - n^2 N_n(x_{Sh}), \\
\alpha_{42} &= 2\mu y_L^2 J_n''(y_L) - \lambda y_L^2 J_n(y_L), \\
\alpha_{43} &= 2\mu y_L^2 N_n''(y_L) - \lambda y_L^2 N_n(y_L), \\
\alpha_{44} &= 2\mu n [J_n(y_{Sh}) - y_{Sh} J_n'(y_{Sh})], \\
\alpha_{45} &= 2\mu n [N_n(y_{Sh}) - y_{Sh} N_n'(y_{Sh})],
\end{aligned} \tag{4.11}$$

$$\begin{aligned}
\alpha_{46} &= - \left(\frac{y_{ext}}{k_{ext}} \right)^2 J_n(y_{ext}), \\
\alpha_{52} &= -\rho_{int}\omega^2 y_L J'_n(y_L), \\
\alpha_{53} &= -\rho_{int}\omega^2 y_L N'_n(y_L), \\
\alpha_{54} &= \rho_{int}\omega^2 n J_n(y_{Sh}), \\
\alpha_{55} &= \rho_{int}\omega^2 n N_n(y_{Sh}), \\
\alpha_{56} &= -y_{ext} J'_n(y_{ext}), \\
\alpha_{62} &= 2n[y_L J'_n(y_L) - J_n(y_L)], \\
\alpha_{63} &= 2n[y_L N'_n(y_L) - N_n(y_L)], \\
\alpha_{64} &= -y_{Sh}^2 J''_n(y_{Sh}) + y_{Sh} J'_n(y_{Sh}) - n^2 J_n(y_{Sh}), \\
\alpha_{65} &= -y_{Sh}^2 N''_n(y_{Sh}) + y_{Sh} N'_n(y_{Sh}) - n^2 N_n(y_{Sh}).
\end{aligned} \tag{4.12}$$

Here, ω is the frequency of the incoming plane wave, k_{ext} is the radial wavenumber in the exterior fluid, k_{int} is the radial wavenumber in the interior fluid, k_L is the longitudinal radial wavenumber in the solid, k_{Sh} is the shear radial wavenumber in the solid, λ and μ are the Lamé elastic constants of the solid, ρ_{ext} is the density of the exterior fluid and ρ_{int} is the density of the interior fluid. The functions J_n and N_n are Bessel functions of the first and second kind respectively and J'_n and N'_n , J''_n and N''_n are their respective first and second derivatives, and x_i and y_i are calculated as $x_i = ak_i$, $y_i = bk_i$, where a is the outer radius of the cylindrical shell and b is the inner radius of the cylindrical shell (equivalently b is the radius of the internal fluid). It should be noted that Hankel functions are related to Bessel functions of the first and second kind by

$$H_n^{(1)} = J_n + iN_n. \tag{4.13}$$

In what follows, different configurations of cylindrical structures are considered. Their setup is described and results from the simulations are presented, together with comparisons with analytical results. All diagrams are logarithmically scaled on the y-axis and their amplitudes

normalised between 0 and 1 based on the maximum amplitude present.

4.1.1 Solid cylinder

The analytical expression for the radial displacement for a solid-fluid scattering problem for an infinitely long solid cylinder submerged in fluid can be found in [54] but can also be calculated from Equation (4.7) with b_n given as

$$b_n = \frac{1}{D} \begin{vmatrix} \beta_1 & \alpha_{12} & \alpha_{14} \\ \beta_2 & \alpha_{22} & \alpha_{24} \\ 0 & \alpha_{32} & \alpha_{34} \end{vmatrix}, \quad (4.14)$$

and D by

$$D = \begin{vmatrix} \alpha_{11} & \alpha_{12} & \alpha_{14} \\ \alpha_{21} & \alpha_{22} & \alpha_{24} \\ 0 & \alpha_{32} & \alpha_{34} \end{vmatrix}. \quad (4.15)$$

The coefficients α_{ij} and β_i are given in Equations (4.10)-(4.12).

The considered setup of this is schematically depicted in Figure 4.2a. The domain is

$$-10 \text{ m} < x < 10 \text{ m},$$

$$-10 \text{ m} < y < 10 \text{ m},$$

with a grid spacing of 12.5 mm corresponding to 1601×1601 grid points. The forcing vector is a three-cycle Hann-windowed toneburst. The frequencies investigated range from 200 Hz to 1 kHz. Absorbing layers are placed on the boundaries of the numerical model to satisfy the Sommerfeld radiation condition [160], as an infinite continuous space is represented using a finite discrete space in the simulations.

Different materials for the solid have been considered, such as steel and perspex, to demonstrate the method's capabilities in low and high impedance situations. Their parameters




Figure 4.2: A schematic showing a (a) solid cylinder and (b) annulus in water insonified by a plane wave. The red dashed lines show the receiver points.

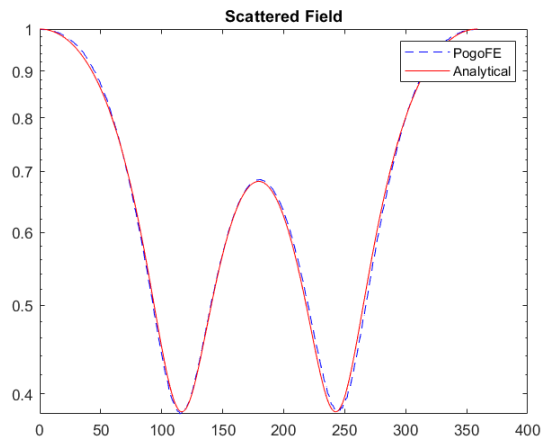
Material	Longitudinal wave speed (m/s)	Shear wave speed (m/s)	Density (kg/m ³)
Steel	5960	3260	7932
Perspex	2730	1430	1180
Water	1500	-	1000
Air	350	-	1.125
Glycerine	1860	-	1258

Table 4.1: Material properties used in numerical examples.

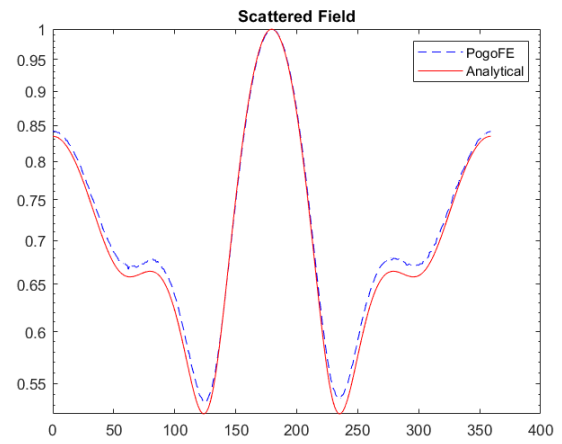
can be found in Table 4.1. Furthermore, perspex has the interesting property that it has a shear wave speed less than the acoustic wavespeed of the surrounding fluid, confirming the accuracy of the model in different scenarios.

The first set of simulations to consider is scattering from a steel cylinder for $k\alpha$ numbers between 0.84 and 4.19. These correspond to frequencies ranging from 200 Hz to 1000 Hz with a fixed radius of 1 m. To compute these simulations, the radius is fixed and the frequency modified so the geometry can be fixed and thus ease computations by eliminating the need for remeshing.


The results are shown in Figure 4.3. There is good agreement throughout the considered range, with the exception of scattered angles perpendicular to the propagation direction,




(a) Scattering at 200 Hz



(b) Scattering at 250 Hz



(c) Scattering at 500 Hz



(d) Scattering at 1 kHz

Figure 4.3: Scattering results a for steel cylinder submerged in water compared to analytical results.

which show some slight deviation. One probable cause for this is due to the direction of the plane wave, with only small amounts of energy being transmitted at grazing incidence and a greater effect of structured mesh staircasing. This might be improved by the use of a conforming mesh, although this does introduce additional complexities during pre-processing as one must carefully define the object boundaries, the sources and receivers and accurately mesh around these. In [141], it is shown that both approaches are equally good for waves in polycrystals, which suffer from the same error regarding staircasing, and thus structured grids are normally used due to the comparative ease of use.

The relative errors of the scattering field at the frequencies shown computed with Equation (4.4) are 1.05%, 1.03%, 1.51%, and 2.17% respectively. At a lower mesh refinement (25 mm), the errors increased to 1.12%, 1.16%, 2.68%, and 4.37%. This clearly shows a reduction in error with a finer mesh, and therefore confirming the convergence of the method. It should be noted that there will always be a trade-off in the additional computational power and storage required to run calculations on a finer mesh for higher accuracy, and it is left to the reader to decide what error-rate is suitable for their simulations. The discrepancy at higher frequencies is to be expected as there is comparatively less resolution per wavelength compared to lower frequencies, as the element sizes are fixed for each set of simulations.

The simulations are then repeated for a perspex cylinder. The geometrical setup remains the same as described above, and a central frequency of 1 kHz with a cylinder radius of 1 m is considered. Results for this can be found in Figure 4.4, with the equivalent-sized steel cylinder present for comparison. This has a relative error of 3.36% for the same mesh discretisation as the steel cylinder. While this may appear initially surprising, the explanation for this is two-fold. Perspex has a lower longitudinal and shear wavespeed than steel, and thus needs further refinement to maintain the same number of elements per wavelength. Furthermore, as visible in Figure 4.5, the perspex scattering pattern is much more susceptible to small changes in frequency compared to the steel pattern. For a $\pm 1\%$ shift in frequency at 1 kHz, the relative error of the analytical solution of the steel cylinder scattering is 1.3% (both downwards and upwards), whereas for a $\pm 1\%$ shift in frequency at 1 kHz, the relative error

of the analytical solution of the perspex cylinder scattering is 3.26% for the downward shift and 4.20% for the upward shift, indicating that errors are not linear in frequency change and are also material dependent.

4.1.2 Layered cylindrical media

The second scattering example considers scattering by solid cylindrically layered media submerged in fluid and enclosing another fluid. Analytical results for this case can be calculated from Equations (4.7) and (4.8) - (4.9).


A schematic diagram of the setup can be found in Figure 4.2b. The same solid materials as those in Section 4.1.1 are considered to demonstrate capabilities in low and high impedance situations. The exterior and interior layers are of fluid, not necessarily of the same material, which surround a solid layer. The interior fluid layer and solid layer both have a radius of 1m for a total of 2 m. The domain is

$$-40 \text{ m} < x < 40 \text{ m},$$

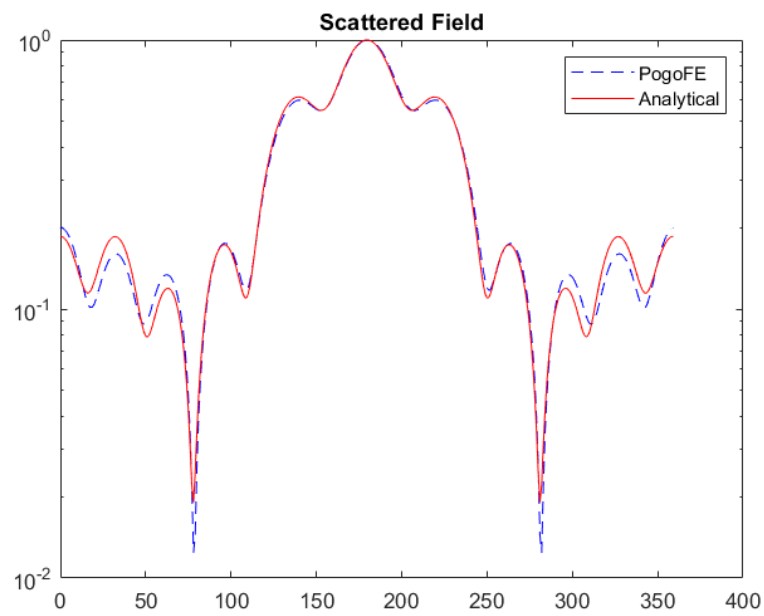
$$-40 \text{ m} < y < 40 \text{ m},$$

with a grid spacing of 50 mm corresponding to 1601×1601 grid points. The forcing vector was a three-cycle Hann-windowed toneburst with a central frequency of 500 Hz. Absorbing boundaries are also placed on the boundaries of the setup as detailed in Section 4.1.1.

In this configuration, steel and perspex is used for the solid layers, water as the exterior fluid and both water and air as the interior layer, giving a total of four computed cases. In Figure 4.6, which shows the scattering pattern for the steel layer, has very little change in the scattering pattern between the two cases of an air and water interior. This is due to a comparatively high density of the solid material compared to the interior fluid, so much of the energy is reflected regardless of the filling material. For the perspex layer, the scattering pattern between the two different materials in the interior layer is clearly different as can be

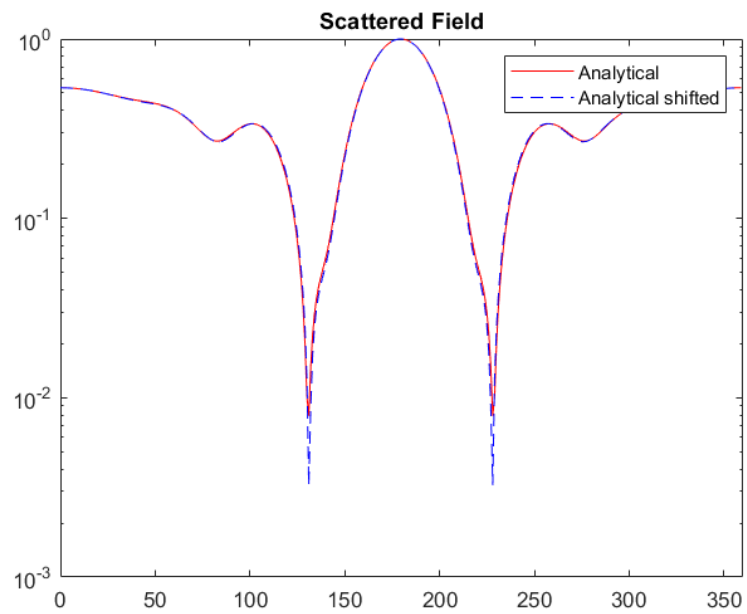


(a) Steel cylinder




(b) Perspex cylinder

Figure 4.4: Scattering patterns from (a) steel cylinder and (b) perspex cylinder submerged in water at 1 kHz.



(a) Steel cylinder



(b) Perspex cylinder

Figure 4.5: Analytical and frequency-shifted scattering patterns from (a) steel cylinder and (b) perspex cylinder submerged in water at 1 kHz and 0.99 kHz respectively.

seen in Figure 4.7. Good agreement between the analytical and numerical solution can be observed in all cases when comparing with analytical results from [55], with a relative error of 2.68%, 3.00% for the steel layer with water and air interiors respectively and 3.73%, 3.19% for the perspex layer with water and air interiors respectively.


The next set of scattering simulations presented are that of ultrasonic scattering from liquid cylindrical obstacles. An experimental investigation of this setup was first performed by Tamarkin in [56]. Analytical results for this case can be calculated from Equation (4.7) with b_n given as

$$b_n = \frac{\begin{vmatrix} \beta_1 & \alpha_{12} \\ \beta_2 & \alpha_{22} \end{vmatrix}}{\begin{vmatrix} \alpha_{11} & \alpha_{12} \\ \alpha_{21} & \alpha_{22} \end{vmatrix}}. \quad (4.16)$$


For full details of the coefficients found in Equation (4.16) the reader is referred to [55]. A setup of the experiment can be found in Figure 4.2a, by replacing the solid cylinder with a fluid one.

The circular fluid inclusion has a radius of 2 m. The incident plane wave has a central frequency of 500 Hz and measurements are taken at a distance of 10 m from the centre of the cylindrical inclusion. The geometrical aspect of the domain is identical to that above. The circular inclusion material in this case is glycerine. The velocity and density material parameters for glycerine found in Table 4.1 are used it is treated as an inviscid fluid for the purpose of this study. The glycerine inclusion is then surrounded by an infinite water layer and absorbing layers are also placed on the boundaries of the setup as detailed in Section 4.1.1 so it can be computationally solved.

The comparison with analytical results shows excellent agreement at all angles as seen in Figure 4.8. Results for this simulation are expected to be more accurate due to the comparable wavespeeds of both glycerine and water which allow the Courant-Friedrichs-Lewy number to be closer to 1, as described in [161]. The relative error for this example is 1.12%. This




(a) Water interior

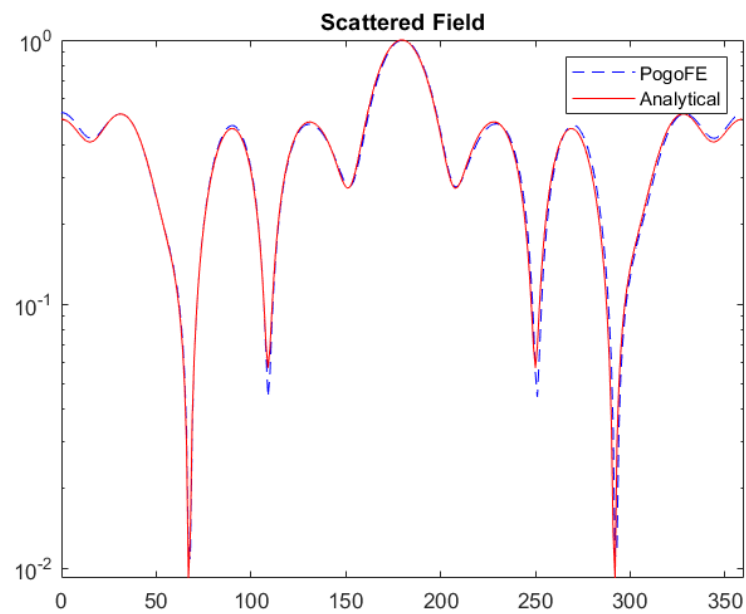


(b) Air interior

Figure 4.6: Scattering patterns from a steel annulus submerged in water with a (a) water interior and (b) air interior.



(a) Water interior



(b) Air interior

Figure 4.7: As in Figure 4.6 but for a perspex annulus.

scenario was considered, as the ability to accurately model the interaction between two fluids is also necessary to model objects on the sediment interface or buried below it, which is extremely useful in underwater acoustics. The current state of the art research frequently models the sediment as a fluid, as in [158, 162].





Figure 4.8: Scattering patterns for normalized displacement from a cylindrical glycerine inclusion.

4.1.3 Cylindrically layered media, arbitrary filling


Now that the method has good results for solid cylinders in Section 4.1.1 and for filled or empty layered cylindrical media, as well as for interactions between fluids in Section 4.1.2, results for arbitrarily filled objects such as those in [7] are computed. Geometrical objects of the same dimensions and materials as in Section 4.1.2 are considered to allow for direct comparison between the scattering patterns. This also provides greater certainty that the simulations are not running into any issues with convergence or refinement. These objects will be partially filled with fluid and the remaining filling will be air. The excitation has a central frequency of 500 Hz as previously. In the majority of the simulations, the filling level is two-thirds water and one-third air. As the objects are no longer axisymmetric, the responses are calculated with excitation from three separate directions, with the angle of incidence of the plane wave being at 0, 90 and 270 degrees. These correspond to a plane wave travelling from the side, above and below respectively. For the plane wave excitation at 0 degrees we will also look at different levels of filling across the two solid materials. These will be filled with one-third water or half-filled with water. The setup for the two-third filling with plane wave excitation at 0, 90 and 270 degrees can be found in Figure 4.9.

There exist no analytical results for this setup although in [7] a finite difference method is used to acquire results. Despite this, some observations and comparisons with existing results from Section 4.1.2 are made. For the steel annulus, with scattering pattern shown in Figure 4.6, no observable difference between water and air filling is obvious. Thus, one can expect the partially filled water-air steel annulus to match the previous scattering pattern, which indeed it does as seen in Figure 4.10. The scattering patterns for the partially filled water-air perspex annulus are found in Figure 4.11. In this case, one can clearly see the difference that the partially filled interior has on the scattering pattern, because of the reduced impedance mismatch.


Another reason we can expect accurate results is due to the nature of the FE method, which works for arbitrary geometries and was one of the motivating factors for using this specific



(a)




(b)




(c)


Figure 4.9: Setup for partially filled annulus $\frac{2}{3}$ full submerged in water with an incoming plane wave at (a) 0, (b) 90 and (c) 270 degrees respectively.



(a)




(b)




(c)


Figure 4.10: Scattering patterns from a steel annulus $\frac{2}{3}$ full submerged in water with an incoming plane wave at (a) 0, (b) 90 and (c) 270 degrees respectively.



(a)



(b)




(c)


Figure 4.11: Same as in Figure 4.10 but for a perspex annulus.

method. While results for a subsection of simulations are shown, often of academic interest, it should be noted that it is straightforward to change materials, dimensions and filling level to better suit the investigations of interest for more practical applications. To demonstrate the ability to handle different filling levels, which is not possible to do with analytical or semi-analytical models, scattering patterns for different filling levels in both steel and perspex are shown in Figure 4.12. For the steel annulus, as previously, no significant changes can be seen in the scattering patterns due to the different filling levels, with minor differences, some which may arise due to staircasing mentioned in Chapter 3. For the perspex annulus, the changes in scattering pattern due to the different filling levels are clearly evident.

Another way to verify and validate the results in the non-axisymmetrical case, where no analytical results exist, is to compare with results other researchers have obtained using existing FE software. In Figure 4.13, one can see the scattering patterns using the FE method implemented in Pogo, compared to that of COMSOL, for the setup shown in Figure 4.9, a two-thirds filled steel annulus with a plane wave incoming at 0 degrees. Excellent agreement between the two simulations can be seen, with only slight deviation at certain angles, most likely due to staircasing or differences in mesh refinement, as the exact meshes and refinement used are different between different software implementations. It should be noted that, although this comparison shows that it is possible to solve non-axisymmetrical problems with existing FE software, only the method developed and presented in this thesis is in the time-domain and easily usable on GPUs. The fact that this is in the time-domain allows users to extract data for multiple frequencies from one simulation as opposed to frequency-domain software where one would have to solve for each individual frequency.



(a) Steel annulus



(b) Perspex annulus

Figure 4.12: Scattering patterns from an annulus submerged in water with different levels of filling comprised of (a) steel and (b) perspex. Plane wave is travelling from the right.




Figure 4.13: Comparison between scattering pattern in Pogo and COMSOL for a steel annulus with a water interior, as described in Section 4.1.3.

4.2 Objects lying and buried in the sediment

Now that the method has been verified for ultrasonic waves in a wide variety of cases, such as those involving fluids, fluids and solids and fluids interacting with fluids, we consider more complex setups, such as objects lying proud on the sediment, buried or similar. As previously, objects of a similar $k\alpha$ number are considered. This is to ensure the simulations have been sufficiently refined to reach convergence, as no analytical solutions exist to compare to.

In this section, the simulation consists of a solid cylinder of 1 m with a plane-wave toneburst signal with central frequency of 500 Hz. The sediment is modelled as a fluid with density equal to 1700 kg m^{-3} and a sound speed equal to 1700 m/s, following [158]. These are very similar parameters to glycerine used to calculate the scattering in Figure 4.8. As the sediment is modelled as a fluid, and not as a solid with a very low shear speed, there is no need to refine the grid to handle the low shear speed and thus no additional complexities are introduced for the meshing stage of the solver.

The cylinder is in three different configurations, as shown in Figure 4.14. These correspond to objects lying on the sediment, partially covered and fully covered, or equivalently with level of burial equal to 0, 50 and 100% respectively. A semi-circular ring of receivers is placed

above the cylinder and sediment, in the water. These are 6m above the water - sediment interface. Modelling the sediment as a fluid is a well accepted practice as in [158, 162]. Results for the directivity patterns can be seen in Figures 4.15 and 4.16 for steel and perspex cylinders respectively, insonified with a plane wave at different levels of burial. Due to the absence of analytical results for these cases, a comparison to COMSOL Multiphysics [128] or similar may be useful in the future and is an avenue for extending this work.


4.3 Setting up accurate simulations

Given the validity of the method as demonstrated in the previous sections for a range of scatterers and configurations, the factors that can affect the accuracy of these simulations is discussed. Namely, the number of elements investigated in examples of thin solid coating and the Courant number accuracy.


4.3.1 Thin solid coating

Another aspect of interest, is the possibility to accurately model thin solid coatings on various objects of interest. It is known from extensive modelling done for the NDE industry that this is possible to do and that the limiting factor is usually a requirement for a certain number of elements per wavelength. It should be noted that a certain number of elements is also required to capture the geometry of the object accurately. For this section, the scatterer is a steel annulus with a different number of elements used to describe the solid part, and thus its thickness, while keeping the rest of the simulation constant, using the same refinement and total simulation size as described in Section 4.1.2.


Figure 4.17 shows results for 10 and 5 elements per wavelength thick solid part. For the simulation with 10 elements per wavelength, the FE simulation approximately matches the analytical result although some noise can be seen, while for the simulation with 5 elements



(a)




(b)




(c)


Figure 4.14: Object proud or buried in sediment, with levels of burial at (a) 0, (b) 50 and (c) 100% levels of burial. Receivers are placed 6m above the water - sediment interface.



(a)




(b)




(c)


Figure 4.15: Results for the directivity patterns for a steel cylinder insonified with a plane wave at (a) 0, (b) 50 and (c) 100% levels of burial.



(a)



(b)




(c)

Figure 4.16: As in Figure 4.15 but for perspex.


per wavelength, the simulation is no longer completely accurate, particularly in the lower amplitude regions. Thus, when considering geometries on a scale approaching the wavelength of the wave, the assumption of accurate modelling centered on elements per wavelength breaks down and convergence studies should be undertaken to ensure the method still generates accurate results. Looking through the literature, there does not seem to be any mathematical theory behind this, unlike the Courant-Friedrichs-Lewy condition. However, this is a phenomenon that is well known in simulations regarding static analysis in which one effectively has infinite wavelength, but still needs to sufficiently refine the mesh to accurately capture the geometry. Thus in simulations, particularly those involving “thin layers” or complex geometries, we are often “geometry-limited” as opposed to “wavelength-limited” and care must be taken to sufficiently refine both aspects.

4.3.2 Courant number accuracy

Another closely related issue that one must be aware of when setting up numerical simulations involving both solids and fluids, is the often large differences in the longitudinal velocity (c_L) in the solid and the acoustic wave speed in the fluid. In [161], it is shown that for solid materials the closer our CFL condition is to 1, the more accurate the simulations become. However, exceeding this results in instability, and if one assumes a much lower number, significantly finer refinement is needed to achieve similar results. This has the implication that, for a simulation involving steel and water, two of the most common materials one might encounter, a Courant number of 1 for the steel immediately implies a Courant number of approximately 0.25 for the water, significantly reducing the accuracy of the simulations at low level of refinements. Therefore, one way to alleviate this (and what is done in this thesis), is to simply refine the mesh further than one might for an equivalent simulation involving only solids.



(a)



(b)

Figure 4.17: Scattering patterns from a thin steel annulus, with (a) 10 elements through-thickness and (b) 5 elements through-thickness.

4.4 Solid sphere in the free field

In the final section of this chapter, initial results for a simulation of a 3-dimensional object are shown, specifically for a solid steel sphere surrounded by water, i.e. in the free-field. The analytical expression for the pressure in the scattered field for a solid sphere can be found in [54], and is given as

$$p_r^{scatt} = \sum_{n=0}^{\infty} b_n [j_n(k_{ext}r) - in_n(k_{ext}r)] P_n(\cos n\theta). \quad (4.17)$$

From this, one can calculate the radial derivative to give the scattered displacement as

$$u_r^{scatt} = \sum_{n=0}^{\infty} b_n k_{ext} [j'_n(k_{ext}r) - in'_n(k_{ext}r)] P_n(\cos n\theta). \quad (4.18)$$

These formulations closely follow those given in Equations (4.6)-(4.7) with the addition of j_n , n_n , which are spherical Bessel functions of the first kind and second kind respectively and j'_n , n'_n are their respective derivatives, while P_n is the Legendre polynomial. The scattering coefficient b_n is given as

$$b_n = -(2n + 1)\epsilon_n(-i)^{n+1} \sin(\eta_n) \exp(i\eta_n) \quad (4.19)$$

where η_n is the phase-shift angle of the n^{th} scattered wave given as

$$\eta_n = \arctan \left(\tan(\delta_n(x_{ext})) \times \frac{(\tan(\Phi_n) + \tan(\alpha_n(x_{ext})))}{\tan(\Phi_n) + \tan(\beta_n(x_{ext}))} \right). \quad (4.20)$$

Here α_n , β_n and δ_n are the intermediate scattering phase-angles given by

$$\alpha_n(x) = \arctan \left(\frac{-xj'_n(x)}{n_n(x)} \right), \quad (4.21)$$

$$\beta_n(x) = \arctan \left(\frac{-xn'_n(x)}{n_n(x)} \right), \quad (4.22)$$

$$\delta_n(x) = \arctan \left(\frac{-xj_n(x)}{n_n(x)} \right), \quad (4.23)$$

Φ_n is the measure of the boundary impedance at the surface of the solid cylinder given as

$$\Phi_n = \left(-\frac{\rho_{ext}}{\rho_{int}} \right) \tan(\zeta_n(x_L, \sigma)), \quad (4.24)$$

and $\zeta_n(x_L, \sigma)$ is the new scattering phase-angle and is given by

$$\zeta_n(x_L, \sigma) = \arctan \left[\frac{\frac{\tan(\alpha_n(x_L))}{\tan(\alpha_n(x_L)) + 1} - \frac{n^2}{\tan(\alpha_n(x_{Sh})) + n^2 - \frac{1}{2}x_{Sh}^2}}{\frac{\frac{x_{Sh}^2}{2} \left(\frac{\tan(\alpha_n(x_{Sh})) + n^2 - \frac{1}{2}x_{Sh}^2}{\tan(\alpha_n(x_L)) + 1} - \frac{n^2(\tan(\alpha_n(x_{Sh})) + 1)}{\tan(\alpha_n(x_{Sh})) + n^2 - \frac{1}{2}x_{Sh}^2} \right)}{\tan(\alpha_n(x_L)) + 1}}}{\tan(\alpha_n(x_L)) + 1} - \frac{n^2(\tan(\alpha_n(x_{Sh})) + 1)}{\tan(\alpha_n(x_{Sh})) + n^2 - \frac{1}{2}x_{Sh}^2}} \right]. \quad (4.25)$$

The setup considered is similar to the one given in Section 4.1.1, extended to three dimensions.


The domain is

$$\begin{aligned} -10 \text{ m} < x < 10 \text{ m}, \\ -10 \text{ m} < y < 10 \text{ m}, \\ -10 \text{ m} < z < 10 \text{ m}, \end{aligned}$$


with a grid spacing of 50 mm corresponding to $401 \times 401 \times 401$ grid points. The forcing vector is a three-cycle Hann-windowed toneburst, and absorbing boundaries are placed on the boundaries of the setup as detailed in Section 4.1.1.

In Figure 4.18, scattering comparisons are shown between the FE simulation and those calculated using the analytical result in the x-y plane with $z = 0$. Results are presented for 100 Hz and 200 Hz, and excellent agreement can be observed, with relative errors of 1.36% and 2.10% respectively. Although results are shown in a single plane, the simulation carried out has been a full three-dimensional reconstruction and does not rely on any axisymmetrical reductions.

The results presented in this section consider lower frequencies compared to earlier setups



(a) Scattering at 100 Hz



(b) Scattering at 200 Hz

Figure 4.18: Scattering patterns for a steel sphere submerged in water compared to analytical results.

due to computational limitations inherent in three-dimensional simulations. Even though the three-dimensional simulation has less refinement per wavelength than the equivalent two-dimensional simulation (by a factor of 4), we have 25 times more nodes in our mesh for the sphere than that of the cylinder. Ways of expanding on three-dimensional simulations are discussed in subsequent chapters.

4.5 Concluding remarks

This chapter is concluded by recalling what was set out to be shown with this set of simulations. The goal was twofold; verify and validate the theoretical finite element formulation derived in Chapter 3 and show that this formulation could be applied to problems in the underwater acoustics area. This has been achieved by computing results from canonical underwater acoustics problems, such as scattering by solid cylinders, empty or filled solid annuli which have analytical solutions in the literature [54, 55], but also those of arbitrarily filled annuli and those of proud or buried cylinders in the sediment. Investigations have been conducted on the thickness refinement, which has shown to be important, in addition to the classical wavelength based refinement in finite element calculations. Finally, initial results for three-dimensional simulations have been shown, namely for a solid sphere in the free-field. As with all problem areas of this type, there is an infinite number of combinations one might be interested in, such as different materials, frequencies, orientations and filling levels, but the author has shown that they have developed capability to solve these problems, and the exact setup is left to the reader.

In the next chapter, we will look at simulations related to the non-destructive testing area of interest.

Chapter 5

Solving Problems II: Application to Non-Destructive Testing

Non-destructive testing techniques are used to identify and size defects in high-value, safety critical components, without damaging or destroying them [9]. A key technique which is used is guided-wave testing. Significant research has been carried out to demonstrate, improve and document the capabilities of guided-wave testing in non-destructive testing, such as in [163, 18]. There are also software packages, such as Disperse [19], which allow us to calculate guided-wave characteristics and their dispersion curves, which are used as reference values for the simulations shown in this chapter.

The objective of the simulations in this chapter is once again twofold; verify and validate the theoretical finite element formulation derived in Chapter 3 and show that this formulation can be applied to problems in the non-destructive field. While the aim of this thesis is to develop capabilities for modelling ultrasonic interactions, as opposed to performing or researching non-destructive testing itself, it is useful to give a brief overview of guided-waves specifically in the case when fluid is present in the simulated domain, so it can guide our modelling work.

It is also helpful to mention that objects of interest in the two communities of underwater acoustics (Chapter 4) and non-destructive testing (this chapter) vary greatly in scale. In

the NDE area, objects are often sized in millimetres and investigated with frequencies in the megahertz range [164], whereas in the underwater acoustics community, objects are often sized in metres with frequencies in the hertz range [7]. Due to the linearity of the wave equation, details of which can be found in Section 2.1.3, there is an equivalence between the two; in dimensionless units (i.e. scaling against the wavenumber) objects and their computational complexity become similar in size between the two areas.

5.1 Guided wave theory

Guided waves are elastodynamic waves which are guided along the boundary of the object and their propagation characteristics depend on both the elastic properties of the object, but also of the surrounding material. This is unlike bulk waves which propagate freely in the surrounding space. Although a full overview of guided-waves is out of the scope of this thesis, for completeness, the key characteristics of them are given and the reader is referred to work such as [165] for further information.

One of the key features of guided waves and therefore their usage in non-destructive testing, is their ability to propagate over long distances with limited amplitude drop, because the ability of the wave (and therefore energy) to spread in space is limited. Another key feature is that they typically become dispersive, i.e. their (group) velocity varies with frequency, as opposed to bulk waves which do not for the majority of materials of interest. This velocity-frequency relationship is captured by a dispersion curve. These dispersion-curves depend not only on the material characteristics but also on the individual guided-waves modes; there are also multiple modes in any given setup. These dispersion curves are numerical solutions of the following characteristic function

$$f(\omega, c) = 0, \tag{5.1}$$

where ω is the frequency and c is the velocity of the guided wave.

Of great importance to NDT, but also to the specific problem of ultrasonic wave propagation in domains involving both solid and fluid are the guided-wave modes known as (leaky) Lamb wave modes and the Scholte wave mode. It should be noted that although theoretically Lamb waves travel indefinitely for elastic plates in a vacuum, the equations and approximations found in the literature also hold for the case of elastic plates in air or water, due to the large impedance mismatch between the solid elastic material and the surrounding material [163].

The characteristic function (from Equation (5.1)) for Lamb waves [165] can be split into two separate classes of solutions, those of symmetric and antisymmetric modes. They have shapes which are perfectly symmetric or antisymmetric to the midplane of the plate respectively. These modes are labelled S_n and A_n , for $n = 0, 1, 2, \dots$, for symmetric and antisymmetric modes respectively, with S_0 and A_0 being termed the fundamental modes. The characteristic function for the symmetric and antisymmetric case are given respectively by

$$(B^2 - k_x^2)^2 \tan(Ah) + (4k_x^2 AB) \tan(Bh) = 0, \quad (5.2)$$

$$(B^2 - k_x^2)^2 \cot(Ah) + (4k_x^2 AB) \cot(Bh) = 0, \quad (5.3)$$

with A and B given as

$$A = \sqrt{\left(\frac{\omega}{c_L}\right)^2 - k_x^2}, \quad (5.4)$$

$$B = \sqrt{\left(\frac{\omega}{c_{Sh}}\right)^2 - k_x^2}. \quad (5.5)$$

Here, c_L and c_{Sh} are the longitudinal and shear wave speed in the solid respectively, and k_x is the wavenumber in the direction along the structure, given by $k_x = \frac{\omega}{c}$ and the thickness of the plate is $2h$. It should be noted that Equations (5.2)-(5.3) have an infinite number of solutions; as the frequency increases, more modes are present in the steel plate and as stated in [163] “all of the modes can extend upwards to infinite phase velocity”. While this may seem paradoxical at first, this does obey conservation laws as energy is transported at group velocity, not phase velocity [163].

5.2 Numerical simulations of guided waves

In this section, the numerical simulation of guided waves involving both solid and fluid materials is considered. This is essential to numerous NDE applications. Three characteristics of guided waves are investigated, with further detail given in each corresponding section. These are: the existence of wave modes, phase velocity, and attenuation of the A_0 mode. The considered setup can be found in Figure 5.1. This corresponds to us considering the 2D representation of the cross-section through the solid steel plate, which is 1 mm thick and surrounded by an infinite fluid (water in our case) on either side. The wave propagates in the direction of the plane of the cross-section, left to right in Figure 5.1. The domain is

$$0 \text{ mm} < x < 150 \text{ mm}, \quad (5.6)$$

$$0 \text{ mm} < y < 22 \text{ mm}, \quad (5.7)$$

with a grid spacing of 0.05 mm corresponding to 3001×441 grid points. A Courant number of 0.8 as detailed in Equation (4.3) is used and enough time for the wave to travel the domain in the horizontal direction and back is assumed. The stiffness reduction method by [118] is once again used and absorbing layers are placed on the vertical boundaries of the setup as detailed in Section 4.1.1. This is to be able to model the infinite fluid in finite space in our finite-element simulations.

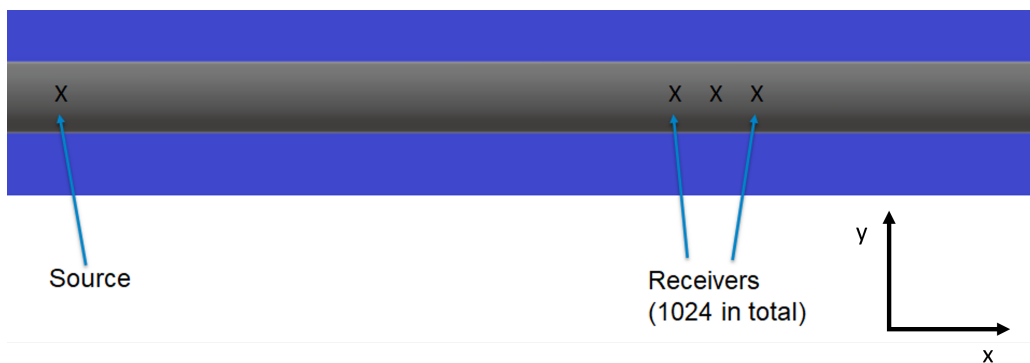


Figure 5.1: Schematic of an immersed steel plate alongside cartesian axes.

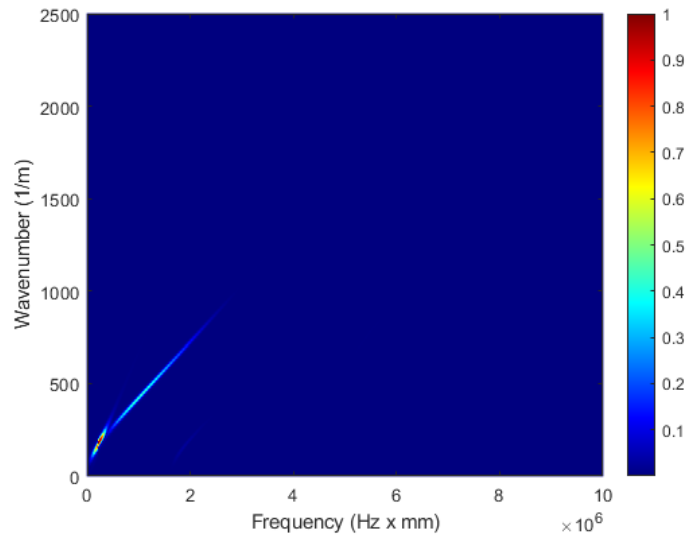
5.2.1 Existence of wavemodes

For the first set of simulations, the goal is to investigate the existence of the antisymmetric guided-wavemodes in a steel plate immersed in an infinite inviscid fluid. A Hann-windowed three-cycled sinusoidal source of 1 MHz excited the 2nd DoF of the node corresponding to the y -direction along the midpoint of the steel plate at $y = 11$ mm. This is expected to excite the antisymmetric modes and not excite the symmetric ones. A series of receivers were placed along the middle of the plate's thickness in the far field. The distance was far enough to allow the wave to propagate and form, so that the guided wave response can be recorded at specific intervals.


To demonstrate the existence of the wave modes, we have relied on the two-dimensional Fourier transform method described in [166]. Signal traces are taken from the simulated receivers and converted to the frequency-wavenumber domain using a zero-padded 2D fast Fourier transform.

The result from the 2D fast Fourier transform is then overlaid with theoretical dispersion curves from Disperse [19]. The results are processed as follows: the bottom 0.1% values are set to 0 and any other values are set to 1. This allows for greater visual clarity of the curves, with the drawback of loss of amplitude information. Therefore, it has been decided to present both pre-processed and processed results. Figure 5.2 clearly shows that both A_0 , A_1 , A_2 and the Scholte wave are captured in the FE solution. The presence of the Scholte wave is particularly encouraging as it is an interface wave, governed by the interaction between the two media at the interface. It also confirms the expectation that the source excitation has excited the antisymmetric modes but not the symmetric ones.


To demonstrate that the model can handle higher frequencies and excite other antisymmetric modes, the simulation is repeated with a central frequency of 5 MHz. The rest of the setup is unchanged. Dispersion curves calculated with the 2D fast Fourier transform method (left) are again overlaid with dispersion curves from Disperse (right) and presented in Figure 5.3. In this case, the higher-order A_2 , A_3 and A_4 antisymmetric modes are also present, and at



(a)



(b)



(c)

Figure 5.2: Dispersion curves for a 1 mm thick steel plate submerged in water (a) shown with no processing and amplitude visible (b) calculated by processing results from an FE simulation, (c) overlaid with Disperse curves (dashed-lines) at 1 MHz.


higher frequencies, away from 5 MHz which was the central frequency excitation, one can see deviation of the finite element results and the theoretical results. This is to be expected, due to the comparatively coarser discretisation per wavelength at higher frequencies.

5.2.2 Phase velocity


For the evaluation of the phase velocity, a Hann-windowed three-cycled sinusoidal source of 2 MHz is applied to the 2nd DoF of the node at the midpoint of the plate. The amplitude spectrum method developed in [167] is used, where the two signal traces containing the A_0 mode in the far field are added together. The amplitude spectrum of the result is then obtained via a Fourier transform and the minima from the resulting graph. Phase velocity can be calculated using the following formula

$$c = \frac{2lf}{m}, \quad (5.8)$$


where c is the guided-wave velocity, l is the distance between the two signals, f is the frequency and m is the index of the minima. The distance between the two response points was 0.01 m. Minima were calculated numerically via Matlab routines and are highlighted on Figure 5.4. Figure 5.4 also shows the combined signal traces and the FFT of the response. Some ringing and low frequency components are visible in the combined signal traces, but as this is far away from the central frequency it does not affect the results. Due to the nature of the method, the results are given for discrete frequencies, and changing the receiver placement can change these frequencies. In Table 5.1, the velocity matches that of Disperse with high accuracy, with relative errors of less than 1%. For the frequency closest to the central frequency, the results are even better, with a near exact match to those calculated using semi-analytical results.



(a)




(b)




(c)

Figure 5.3: As in Figure 5.2 but for 5 MHz.



(a) Combined signal traces.



(b) Amplitude spectrum of combined signals.

Figure 5.4: (a) Combined signal traces of the A0 propagating wave measured at two points as described in Section 5.2.2 alongside (b) their amplitude spectrum. Red stars in (b) indicated local minima on the graph, found through use of Matlab routines.

Simulation frequency (MHz)	Disperse Frequency (MHz)	Phase velocity from FE (m/ms)	Phase velocity from Disperse (m/ms)	Relative error difference (%)
1.266	1.267	2.532	2.519	0.515
1.589	1.589	2.648	2.643	0.189
1.912	1.912	2.731	2.732	0.0366
2.235	2.235	2.794	2.797	0.108
2.558	2.558	2.842	2.845	0.106

Table 5.1: Comparison of phase velocities between Disperse and FE simulations.


5.2.3 Attenuation

Guided waves attenuate due to absorption, scattering and leakage. Absorption and scattering do not contribute to attenuation in this case; in the simulation considered, we do not have attenuation due to absorption as elastic materials with no damping are considered, and we have no scattering attenuation as there is no roughness or inhomogeneities in the simulated steel plate.


Leakage, in particular, occurs in guided waves when the waveguide (in this case the steel plate) is in contact with another medium (the surrounding fluid). The correct prediction of this leakage and therefore the attenuation characteristics are of great importance as they are part of the modal properties of the plate-water system [165].

Modelling attenuation is very sensitive to the contributions of multiple modes present in a signal (it is rarely possible to achieve single mode excitation, in both numerical and experimental settings). Hence, it requires careful generation and windowing of signals to achieve accurate results. For excitation “mode matching” has been used, where the source perfectly matches the displacement and phase characteristics of the desired A_0 wave mode. Displacement and phase information for this is readily available from Disperse and can be seen in Figure 5.5 and ensures that the simulated wave does not excite any additional modes. This information is used as an input to the forcing vector in the simulations.

Measurements are then taken arbitrarily at two points where the receivers are placed, as can be seen in Figure 5.1. In Figure 5.6, one can clearly see that the second signal has less amplitude than the first signal, giving a qualitative view of the attenuation. Attenuation is



(a) Displacement data.



(b) Phase data.

Figure 5.5: A_0 mode shape data from Disperse for a 1 mm thick steel plate at 1 MHz submerged in water with (a) displacement (b) and phase.

then calculated by computing a Fourier transform of the received signals as shown in Figure 5.6, dividing these and then scale this based on the distance travelled.




Figure 5.6: FFT of first and second signal.

The most accurate results for attenuation are achieved around the central frequency of the source. FFT spectra of the signals used in the 1 MHz simulation are shown in Figure 5.6. The attenuation simulation is repeated for frequencies of 100 kHz and 200 kHz either side of the original 1 MHz excitation for a total of 5 simulations. In Table 5.2, the attenuation calculated from FE results is compared to those coming from Disperse. The values computed from the FE simulations agree well with those extracted from Disperse [19], with relative errors of less than 1%. This indicates the fluid formulation can accurately model the leakage phenomenon in guided wave propagation.

Frequency (kHz)	Attenuation from FE (dB/m)	Attenuation from Disperse (dB/m)	Relative error difference (%)
800	556.1	555.1	0.18%
900	530.9	530.7	0.2%
1000	511.2	514.2	0.583%
1100	498.2	502.6	0.875%
1200	496.6	494.6	0.404%

Table 5.2: Comparison of attenuation between Disperse and FE simulations.

5.3 Concluding Remarks

This chapter is once again concluded by recalling what was set out to show with this set of simulations. The goal was twofold; as in Chapter 4, verify and validate the theoretical finite element formulation derived in Chapter 3 and show that this formulation could be applied to problems in a different research area, in particular that of non-destructive testing. This has been done by computing guided wave properties, such as dispersion curves, velocity and attenuation. As mentioned, of great importance was capturing the existence of the Scholte wave-mode and the leaky characteristics of the $A0$ wave mode, which have both been successfully simulated.

In the following chapter, this thesis is concluded by providing an overview of the work presented so far as well as possible ways to extend this further.

Chapter 6

Concluding Remarks and Future Research Avenues

6.1 Review of Thesis

The aim of this thesis was to investigate and implement a method that would allow us to accurately simulate ultrasonic wave phenomena in domains involving both solid and fluid materials. This method would need to be able to handle a range of complex geometries and materials found in real-world engineering applications.

Chapter 2 reviewed the governing equations underlying ultrasound propagation, providing a literature review of analytical and numerical methods to model this propagation. It also included a review and description of the necessary boundary conditions and ways to simulate these in numerical software with the use of the stiffness reduction method.

Chapter 3 looked at the finite element method in more detail, covering its implementation in software, and specifically focused on the proposed displacement-based formulation to simulate ultrasonic wave phenomena in solid-fluid domains ready for implementation on a GPU.

In Chapter 4, numerical simulations related to the field of underwater acoustics were intro-

duced and simulated. It began by introducing analytical results found in the literature from [54, 55]. It proceeded to compare results for scattering patterns from canonical setups of different materials in both the free-field and proud or buried in the surface for objects such as cylinders, annuli and spheres. For the annuli, different levels of filing and incident directions of the excited plane wave were also investigated.

Chapter 5, the final results chapter of this thesis, looked at numerical simulations regarding non-destructive-testing, particularly at simulating guided-waves. It was demonstrated that this numerical method can accurately predict their existence, capture their velocity and in the leaky-case, the attenuation of the $A0$ wave mode. The calculation of these used various methods found in the literature, such as the amplitude spectrum method [167] or Fourier-transform methods described in [166].

6.2 Key Contributions

There is an ever increasing demand in academia and industry on acquiring a significant number of high-fidelity results from large-scale computational simulations of ultrasonic wave phenomena, at speed. However, this has always been challenging, due to the computational resources required, as well as the complexities of implementing the underlying formulation in software. While significant improvements have been made in recent years on both fronts, there is currently no one clear performant solution to modelling these ultrasonic interactions in complex, industry-relevant configurations. Therefore, the author has formulated an explicit finite-element numerical method in space, coupled with a finite-difference method in time. The method is displacement-based and is thus compatible with typical elastodynamic finite element formulations and ready for implementation on a graphical processing unit. The method has been subsequently verified and validated this implementation across a range of problems involving millions of degrees of freedom.

Specifically, it was shown that the formulation can compute results from canonical underwater

acoustics problems, both in the free-field and proud or buried on the sediment, and for arbitrary amounts of filling. Two-dimensional objects looked at are cylinders and cylindrically-layered media (annuli) of different solid and fluid materials. The arbitrarily filled objects, and cylinders proud or buried in the sediment do not have analytical solutions, and thus only numerical methods such as those presented here can accurately simulate them. Preliminary work has been completed showing that this formulation also works for three-dimensional objects, by computing scattering results from a solid steel sphere surrounded by water.

Another aspect of this work which has been important is the demonstration of the need to account for “geometrical refinement” as well as the commonly accepted “wavelength refinement” when setting up numerical simulations, particularly those involving complex geometries.

Furthermore, it is shown that one can simulate the propagation and characteristics of guided waves, such as their existence, velocity and attenuation. Errors were calculated for guided wave phase velocity and attenuation and these were less than 1%, while capturing the Scholte wave correctly gives the readers of this thesis confidence in the implementation of the solid-fluid coupling.

For the different cases considered above, good results are achieved for both low frequencies common in underwater acoustics and high frequencies common in NDT settings, confirming the accuracy of the method for capturing fluid-solid wave interactions.

6.3 Future Work

There are various avenues that one might explore to extend this work in the future. Below, three such avenues are detailed, which would be a direct continuation of the work presented in this thesis.

The first potential investigation is regarding absorbing boundary conditions, specifically for fluid materials. In this thesis, the stiffness reduction method described in [118] has been relied

upon, which was designed for solid materials with both longitudinal and shear waves. As inviscid fluid materials do not support shear wave propagation, there should be a more efficient way to damp out these waves, thus reducing the width of the absorbing layer. This would reduce the computational storage and power needed to perform simulations. The reduction would allow one to further refine the computational grid; or simply perform simulations faster with the same hardware.

Another avenue of investigation one might wish to pursue is verifying and validating that the proposed formulation also accurately works for viscous fluids. While the formulation has been designed to allow for viscous damping, there may be some additional complexities. Thus this needs to be properly investigated, with comparisons to existing analytical or numerical results across a range of scenarios. This would allow the formulation to be used in additional investigations than just the ones involving inviscid fluids it can currently solve.

The third possible extension is to use the formulation for larger, industry-relevant complex simulations. It has been shown we can accurately simulate such objects, although due to the size of them, it was limited in the frequency ranges it could investigate. It may be possible to refine the computational mesh in such a way that we have different levels of refinement depending on the “importance” each feature has, or for different materials with different wavespeeds. One immediate downside which would need to be accounted for is the possibility of additional numerical reflections due to the change. There are also constant advances being made in software, such as improvements to matrix calculations, and hardware that will help alleviate this problem.

Bibliography

- [1] M. Zampolli, F. B. Jensen, and A. Tesei, “Benchmark problems for acoustic scattering from elastic objects in the free field and near the seafloor,” *The Journal of the Acoustical Society of America*, vol. 125, no. 1, pp. 89–98, 2009.
- [2] Sergey Dubrov (Photographer), “Russian diving,” 2014. [Online; <https://www.shutterstock.com/g/sdubrov/sets/1748716> accessed 15 September 2022].
- [3] T. Belytschko, J. S.-J. Ong, W. K. Liu, and J. M. Kennedy, “Hourglass control in linear and nonlinear problems,” *Computer Methods in Applied Mechanics and Engineering*, vol. 43, no. 3, pp. 251–276, 1984.
- [4] F. M. Abu-Zidan, A. F. Hefny, and P. Corr, “Clinical ultrasound physics,” *Journal of Emergencies, Trauma and Shock*, vol. 4, no. 4, p. 501, 2011.
- [5] D. Aggelis, E. Kordatos, M. Strantza, D. Soulioti, and T. Matikas, “NDT approach for characterization of subsurface cracks in concrete,” *Construction and Building Materials*, vol. 25, no. 7, pp. 3089–3097, 2011.
- [6] B. N. Pavlakovic, *Leaky guided ultrasonic waves in NDT*. PhD thesis, 1998.
- [7] S. D. Adams, R. V. Craster, and D. P. Williams, “Two-dimensional frequency domain scattering and point excitation of partially filled elastic cylinders,” *Journal of sound and vibration*, vol. 329, no. 11, pp. 2147–2157, 2010.

- [8] M. R. Jolly, A. Prabhakar, B. Sturzu, K. Hollstein, R. Singh, S. Thomas, P. Foote, and A. Shaw, “Review of non-destructive testing (ndt) techniques and their applicability to thick walled composites,” *Procedia CIRP*, vol. 38, pp. 129–136, 2015.
- [9] A. Lopez, R. Bacelar, I. Pires, T. G. Santos, J. P. Sousa, and L. Quintino, “Non-destructive testing application of radiography and ultrasound for wire and arc additive manufacturing,” *Additive Manufacturing*, vol. 21, pp. 298–306, 2018.
- [10] B. Wang, S. Zhong, T.-L. Lee, K. S. Fancey, and J. Mi, “Non-destructive testing and evaluation of composite materials/structures: A state-of-the-art review,” *Advances in mechanical engineering*, vol. 12, no. 4, p. 1687814020913761, 2020.
- [11] “Non-destructive testing. Guided wave testing. General guidance and principles.,” *British Standard 9690-1*, 2011.
- [12] S. Eckel, P. Huthwaite, M. Lowe, A. Schumm, and P. Gu erin, “Establishment and validation of the channelized hotelling model observer for image assessment in industrial radiography,” *NDT & E International*, vol. 98, pp. 1–7, 2018.
- [13] E. Deprins, “Computed radiography in ndt applications,” *Insight-Non-Destructive Testing and Condition Monitoring*, vol. 46, no. 10, pp. 590–593, 2004.
- [14] H. Gebru, “Assessment of ultrasonic inspection techniques and models for nuclear power plant components,” 2015.
- [15] J. Russell, “The development and implementation of advanced ultrasonic phased array technology,” 2010.
- [16] M. Li and G. Hayward, “Ultrasound nondestructive evaluation (nde) imaging with transducer arrays and adaptive processing,” *Sensors*, vol. 12, no. 1, pp. 42–54, 2011.
- [17] L. Le Jeune, S. Robert, P. Dumas, A. Membre, and C. Prada, “Adaptive ultrasonic imaging with the total focusing method for inspection of complex components immersed in water,” in *AIP Conference proceedings*, vol. 1650, pp. 1037–1046, American Institute of Physics, 2015.

- [18] D. N. Alleyne, *The nondestructive testing of plates using ultrasonic Lamb waves*. PhD thesis, Department of Mechanical Engineering, Imperial College, 1991.
- [19] 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.
- [20] A. Castellano, P. Foti, A. Fraddosio, S. Marzano, and M. D. Piccioni, “Mechanical characterization of CFRP composites by ultrasonic immersion tests: Experimental and numerical approaches,” *Composites Part B: Engineering*, vol. 66, pp. 299–310, 2014.
- [21] H.-J. Kim, S.-J. Song, and L. W. Schmerr, “Modeling ultrasonic pulse-echo signals from a flat-bottom hole in immersion testing using a multi-Gaussian beam,” *Journal of Nondestructive Evaluation*, vol. 23, no. 1, pp. 11–19, 2004.
- [22] W. Cobb, “Non-Intrusive, Ultrasonic Measurement of Fluid Composition,” in *Review of Progress in Quantitative Nondestructive Evaluation*, pp. 2177–2183, Springer, 1999.
- [23] P. Petcher, S. Burrows, and S. Dixon, “Shear horizontal (SH) ultrasound wave propagation around smooth corners,” *Ultrasonics*, vol. 54, no. 4, pp. 997–1004, 2014.
- [24] R. E. Hansen, “Mapping the ocean floor in extreme resolution using interferometric synthetic aperture sonar,” in *Proceedings of Meetings on Acoustics ICU*, vol. 38, p. 055003, Acoustical Society of America, 2019.
- [25] T. A. Mooney, L. Di Iorio, M. Lammers, T.-H. Lin, S. L. Nedelec, M. Parsons, C. Radford, E. Urban, and J. Stanley, “Listening forward: approaching marine biodiversity assessments using acoustic methods,” *Royal Society open science*, vol. 7, no. 8, p. 201287, 2020.
- [26] M. Zampolli, A. L. Espana, K. L. Williams, S. G. Kargl, E. I. Thorsos, J. L. Lopes, J. L. Kennedy, and P. L. Marston, “Low-to mid-frequency scattering from elastic objects on a sand sea floor: Simulation of frequency and aspect dependent structural echoes,” *Journal of Computational Acoustics*, vol. 20, no. 02, p. 1240007, 2012.

- [27] K. Williams, T. McGinnis, V. Miller, B. Brand, and R. Light, “Limited scope design study for multi-sensor towbody,” tech. rep., University of Washington Seattle United States, 2016.
- [28] R. P. Feynman, R. B. Leighton, and M. Sands, “The feynman lectures on physics; vol. i,” *American Journal of Physics*, vol. 33, no. 9, pp. 750–752, 1965.
- [29] J. Achenbach, *Wave propagation in elastic solids*. Elsevier, 2012.
- [30] E. Pouliquen, A. P. Lyons, and N. G. Pace, “Penetration of acoustic waves into sandy seafloors at low grazing angles: The helmholtz-kirchhoff approach,” 1998.
- [31] N. J. Pignier, C. J. O’Reilly, and S. Boij, “A kirchhoff approximation-based numerical method to compute multiple acoustic scattering of a moving source,” *Applied Acoustics*, vol. 96, pp. 108–117, 2015.
- [32] E. I. Thorsos, “The validity of the kirchhoff approximation for rough surface scattering using a gaussian roughness spectrum,” *The Journal of the Acoustical Society of America*, vol. 83, no. 1, pp. 78–92, 1988.
- [33] F. Shi, *Elastic wave scattering from randomly rough surfaces*. PhD thesis, Imperial College London, 2015.
- [34] J. Gubernatis, E. Domany, J. Krumhansl, and M. Huberman, “The born approximation in the theory of the scattering of elastic waves by flaws,” *Journal of Applied Physics*, vol. 48, no. 7, pp. 2812–2819, 1977.
- [35] M. Born and E. Wolf, *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- [36] R. Phillips, “Efficient numerical modelling of the ultrasonic scattering from complex surface-breaking defects,” 2018.
- [37] P. Calmon, S. Mahaut, S. Chatillon, and R. Raillon, “Civa: An expertise platform for simulation and processing ndt data,” *Ultrasonics*, vol. 44, pp. e975–e979, 2006.

- [38] S. Banerjee and T. Kundu, “Ultrasonic field modeling in plates immersed in fluid,” *International Journal of Solids and Structures*, vol. 44, no. 18-19, pp. 6013–6029, 2007.
- [39] E. Ginzel, “NDT modelling: an overview,” in *Proc. Conference on Modelling in Non-Destructive Testing, Pretoria, South Africa*, 2007.
- [40] R. J. Pyle, R. L. Bevan, R. R. Hughes, R. K. Rachev, A. A. S. Ali, and P. D. Wilcox, “Deep learning for ultrasonic crack characterization in nde,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 68, no. 5, pp. 1854–1865, 2020.
- [41] L. N. Trefethen, “Finite difference and spectral methods for ordinary and partial differential equations,” 1996.
- [42] S. M. Kirkup, *The boundary element method in acoustics*. Integrated sound software, 2007.
- [43] K.-J. Bathe, *Finite element procedures*. Klaus-Jurgen Bathe, 2006.
- [44] P. Huthwaite, “Accelerated finite element elastodynamic simulations using the GPU,” *Journal of Computational Physics*, vol. 257, pp. 687–707, 2014.
- [45] M. Gennaretti, G. Bernardini, C. Poggi, and C. Testa, “Velocity-potential boundary-field integral formulation for sound scattered by moving bodies,” *AIAA Journal*, vol. 56, no. 9, pp. 3547–3557, 2018.
- [46] O. Omid and V. Lotfi, “A symmetric implementation of pressure-based fluid–structure interaction for nonlinear dynamic analysis of arch dams,” *Journal of Fluids and Structures*, vol. 69, pp. 34–55, 2017.
- [47] G. Everstine, “Finite element formulations of structural acoustics problems,” *Computers & Structures*, vol. 65, no. 3, pp. 307–321, 1997.
- [48] O. C. Zienkiewicz, “Coupled vibrations of a structure submerged in a compressible fluid,” in *Proc. of Symposium on Finite Element Techniques Held at the University of Stuttgart*, 1969.

- [49] S. Rugonyi and K.-J. Bathe, “On finite element analysis of fluid flows fully coupled with structural interactions,” *CMES- Computer Modeling in Engineering and Sciences*, vol. 2, no. 2, pp. 195–212, 2001.
- [50] H. C. Chen and R. L. Taylor, “Vibration analysis of fluid–solid systems using a finite element displacement formulation,” *International Journal for Numerical Methods in Engineering*, vol. 29, no. 4, pp. 683–698, 1990.
- [51] D. Flanagan and T. Belytschko, “A uniform strain hexahedron and quadrilateral with orthogonal hourglass control,” *International journal for numerical methods in engineering*, vol. 17, no. 5, pp. 679–706, 1981.
- [52] C. R. Harris, K. J. Millman, S. J. Van Der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, *et al.*, “Array programming with numpy,” *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.
- [53] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, *et al.*, “Scipy 1.0: fundamental algorithms for scientific computing in python,” *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.
- [54] J. J. Faran Jr, “Sound scattering by solid cylinders and spheres,” *The Journal of the acoustical society of America*, vol. 23, no. 4, pp. 405–418, 1951.
- [55] R. Doolittle and H. Überall, “Sound scattering by elastic cylindrical shells,” *The Journal of the Acoustical Society of America*, vol. 39, no. 2, pp. 272–275, 1966.
- [56] P. Tamarkin, “Scattering of an underwater ultrasonic beam from liquid cylindrical obstacles,” *The Journal of the Acoustical Society of America*, vol. 21, no. 6, pp. 612–616, 1949.
- [57] P. A. Tipler and G. Mosca, *Physics for Scientists and Engineers, Volume 1: Mechanics, Oscillations and Waves; Thermodynamics*. Macmillan, 2003.

- [58] F. Ihlenburg and I. Babuska, “Finite element solution of the Helmholtz equation with high wave number part II: the hp version of the FEM,” *SIAM Journal on Numerical Analysis*, vol. 34, no. 1, pp. 315–358, 1997.
- [59] H. Gan, P. Levin, and R. Ludwig, “Finite element formulation of acoustic scattering phenomena with absorbing boundary condition in the frequency domain,” *The Journal of the Acoustical Society of America*, vol. 94, no. 3, pp. 1651–1662, 1993.
- [60] G. Everstine, F. Henderson, and R. Lipman, “Finite Element Prediction of Acoustic Scattering and Radiation from Submerged Elastic Structures,” 1984.
- [61] J. R. Cannon, *The one-dimensional heat equation*. No. 23, Cambridge University Press, 1984.
- [62] D. Medková, “The laplace equation,” *Boundary value problems on bounded and unbounded Lipschitz domains*. Springer, Cham, 2018.
- [63] F. W. Dorr, “The direct solution of the discrete poisson equation on a rectangle,” *SIAM review*, vol. 12, no. 2, pp. 248–263, 1970.
- [64] M. Marion and R. Temam, “Navier-stokes equations: Theory and approximation,” *Handbook of numerical analysis*, vol. 6, pp. 503–689, 1998.
- [65] R. L. Higdon, “Radiation boundary conditions for elastic wave propagation,” *SIAM Journal on Numerical Analysis*, vol. 27, no. 4, pp. 831–869, 1990.
- [66] K. Aki and P. G. Richards, *Quantitative seismology*. 2002.
- [67] L. Koenigsberger, “Hermann von Helmholtz,” 1902.
- [68] K. Moriyasu, *An elementary primer for gauge theory*. World Scientific, 1983.
- [69] D. C. Gazis, “Three-dimensional investigation of the propagation of waves in hollow circular cylinders. i. analytical foundation,” *The journal of the Acoustical Society of America*, vol. 31, no. 5, pp. 568–573, 1959.

- [70] D. A. Russell, “Acoustics and vibration animations,”
- [71] E. A. Skelton and J. H. James, *Theoretical acoustics of underwater structures*. World Scientific, 1997.
- [72] T. G. Mackay and A. Lakhtakia, “The transfer-matrix method in electromagnetics and optics,” *Synthesis lectures on electromagnetics*, vol. 1, no. 1, pp. 1–126, 2020.
- [73] Y. Xu, D. N. Johnston, Z. Jiao, and A. R. Plummer, “Frequency modelling and solution of fluid–structure interaction in complex pipelines,” *Journal of sound and vibration*, vol. 333, no. 10, pp. 2800–2822, 2014.
- [74] P. Wilcox, R. Monkhouse, M. Lowe, and P. Cawley, “The use of Huygens’ principle to model the acoustic field from interdigital Lamb wave transducers,” in *Review of Progress in Quantitative Nondestructive Evaluation*, pp. 915–922, Springer, 1998.
- [75] L. W. Schmerr, H.-J. Kim, R. Huang, and A. Sedov, “Multi-gaussian ultrasonic beam modeling,” *coordinates*, vol. 1, no. 2, p. 3, 2003.
- [76] S. N. Chandler-Wilde and S. Langdon, “A galerkin boundary element method for high frequency scattering by convex polygons,” *SIAM Journal on Numerical Analysis*, vol. 45, no. 2, pp. 610–640, 2007.
- [77] D. Lopes, R. Agujetas, H. Puga, J. Teixeira, R. Lima, J. Alejo, and C. Ferrera, “Analysis of finite element and finite volume methods for fluid-structure interaction simulation of blood flow in a real stenosed artery,” *International Journal of Mechanical Sciences*, vol. 207, p. 106650, 2021.
- [78] L. Ma, G. N. Barakos, and Q. Zhao, “A 3d implicit structured multi-block grid finite volume method for computational structural dynamics,” *Aerospace Science and Technology*, vol. 117, p. 106980, 2021.
- [79] E. Stein, R. De Borst, and T. J. Hughes, “Encyclopedia of computational mechanics,” 2004.

- [80] J. D. De Basabe and M. K. Sen, “A comparison of finite-difference and spectral-element methods for elastic wave propagation in media with a fluid-solid interface,” *Geophysical Journal International*, vol. 200, no. 1, pp. 278–298, 2015.
- [81] E. Rosenkrantz, A. Bottero, D. Komatitsch, and V. Monteiller, “A flexible numerical approach for non-destructive ultrasonic testing based on a time-domain spectral-element method: Ultrasonic modeling of lamb waves in immersed defective structures and of bulk waves in damaged anisotropic materials,” *NDT & E International*, vol. 101, pp. 72–86, 2019.
- [82] Y. Liu, J. Teng, H. Lan, X. Si, and X. Ma, “A comparative study of finite element and spectral element methods in seismic wavefield modeling,” *Geophysics*, vol. 79, no. 2, pp. T91–T104, 2014.
- [83] J. O. Robertsson, R. Laws, C. Chapman, J.-P. Vilotte, and E. Delavaud, “Modelling of scattering of seismic waves from a corrugated rough sea surface: a comparison of three methods,” *Geophysical Journal International*, vol. 167, no. 1, pp. 70–76, 2006.
- [84] R. van Vossen, J. O. Robertsson, and C. H. Chapman, “Finite-difference modeling of wave propagation in a fluid–solid configuration,” *Geophysics*, vol. 67, no. 2, pp. 618–624, 2002.
- [85] F. Cegla, P. Cawley, and M. Lowe, “Material property measurement using the quasi-scholte mode—a waveguide sensor,” *The Journal of the Acoustical Society of America*, vol. 117, no. 3, pp. 1098–1107, 2005.
- [86] G. Dahake and S. Gracewski, “Finite difference predictions of p-sv wave propagation inside submerged solids. i. liquid–solid interface conditions,” *The Journal of the Acoustical Society of America*, vol. 102, no. 4, pp. 2125–2137, 1997.
- [87] E. Hairer, S. P. Nørsett, and G. Wanner, “Solving ordinary differential equations i,” 1991.

- [88] E. Burman, M. A. Fernández, and P. Hansbo, “Continuous interior penalty finite element method for oseen’s equations,” *SIAM journal on numerical analysis*, vol. 44, no. 3, pp. 1248–1274, 2006.
- [89] E. Burman, R. Durst, M. A. Fernández, and J. Guzmán, “Fully discrete loosely coupled robin-robin scheme for incompressible fluid–structure interaction: stability and error analysis,” *Numerische Mathematik*, vol. 151, no. 4, pp. 807–840, 2022.
- [90] F. J. Buwalda, E. De Goede, M. Knepflé, and C. Vuik, “Comparison of an explicit and implicit time integration method on gpus for shallow water flows on structured grids,” *Water*, vol. 15, no. 6, p. 1165, 2023.
- [91] Z. Fu, T. J. Lewis, R. M. Kirby, and R. T. Whitaker, “Architecting the finite element method pipeline for the gpu,” *Journal of Computational and Applied Mathematics*, vol. 257, pp. 195–211, 2014.
- [92] R. Courant, K. Friedrichs, and H. Lewy, “Über die partiellen differenzgleichungen der mathematischen physik,” *Mathematische annalen*, vol. 100, no. 1, pp. 32–74, 1928.
- [93] A. Bermúdez and R. Rodríguez, “Finite element computation of the vibration modes of a fluid—solid system,” *Computer Methods in Applied Mechanics and Engineering*, vol. 119, no. 3-4, pp. 355–370, 1994.
- [94] T. J. Hughes, *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation, 2012.
- [95] T. Belytschko, W. K. Liu, B. Moran, and K. Elkhodary, *Nonlinear finite elements for continua and structures*. John wiley & sons, 2014.
- [96] M. A. Hamdi, Y. Ousset, and G. Verchery, “A displacement method for the analysis of vibrations of coupled fluid-structure systems,” *International Journal for Numerical Methods in Engineering*, vol. 13, no. 1, pp. 139–150, 1978.

- [97] K. Bathe, C. Nitikitpaiboon, and X. Wang, “A mixed displacement-based finite element formulation for acoustic fluid-structure interaction,” *Computers & Structures*, vol. 56, no. 2-3, pp. 225–237, 1995.
- [98] G. C. Everstine and F. M. Henderson, “Coupled finite element/boundary element approach for fluid–structure interaction,” *The Journal of the Acoustical Society of America*, vol. 87, no. 5, pp. 1938–1947, 1990.
- [99] R. A. Jeans and I. C. Mathews, “Solution of fluid–structure interaction problems using a coupled finite element and variational boundary element technique,” *The Journal of the Acoustical Society of America*, vol. 88, no. 5, pp. 2459–2466, 1990.
- [100] J.-H. Park, H. Koh, and J. Kim, “Fluid-structure interaction analysis by a coupled boundary element-finite element method in time domain,” *Boundary element technology VII*, pp. 227–243, 1992.
- [101] R. Suliman, O. Oxtoby, A. Malan, and S. Kok, “A matrix free, partitioned solution of fluid–structure interaction problems using finite volume and finite element methods,” *European Journal of Mechanics - B/Fluids*, vol. 49, pp. 272–286, 2015.
- [102] Y. Yang, B. R. Mace, and M. J. Kingan, “Wave and finite element method for predicting sound transmission through finite multi-layered structures with fluid layers,” *Computers & Structures*, vol. 204, pp. 20–30, 2018.
- [103] H. P. Langtangen, *Computational partial differential equations: numerical methods and diffpack programming*, vol. 2. Springer Science & Business Media, 2013.
- [104] S. Brenner and R. Scott, *The mathematical theory of finite element methods*, vol. 15. Springer Science & Business Media, 2007.
- [105] A. Bayliss and E. Turkel, “Radiation boundary conditions for wave-like equations,” *Communications on Pure and applied Mathematics*, vol. 33, no. 6, pp. 707–725, 1980.
- [106] S. G. Johnson, “Notes on perfectly matched layers (PMLs),” *Lecture notes, Massachusetts Institute of Technology, Massachusetts*, vol. 29, 2008.

- [107] S. Mahaut, N. Leymarie, C. Poidevin, T. Fouquet, and O. Dupond, “Study of complex ultrasonic NDT cases using hybrid simulation method and experimental validations,” *Insight-Non-Destructive Testing and Condition Monitoring*, vol. 53, no. 12, pp. 664–667, 2011.
- [108] D. Givoli, “High-order local non-reflecting boundary conditions: a review,” *Wave motion*, vol. 39, no. 4, pp. 319–326, 2004.
- [109] J.-P. Berenger, “A perfectly matched layer for the absorption of electromagnetic waves,” *Journal of computational physics*, vol. 114, no. 2, pp. 185–200, 1994.
- [110] F. D. Hastings, J. B. Schneider, and S. L. Broschat, “Application of the perfectly matched layer (PML) absorbing boundary condition to elastic wave propagation,” *The Journal of the Acoustical Society of America*, vol. 100, no. 5, pp. 3061–3069, 1996.
- [111] Q.-H. Liu and J. Tao, “The perfectly matched layer for acoustic waves in absorptive media,” *The Journal of the Acoustical Society of America*, vol. 102, no. 4, pp. 2072–2082, 1997.
- [112] D. Komatitsch and J. Tromp, “A perfectly matched layer absorbing boundary condition for the second-order seismic wave equation,” *Geophysical Journal International*, vol. 154, no. 1, pp. 146–153, 2003.
- [113] E. Turkel and A. Yefet, “Absorbing PML boundary layers for wave-like equations,” *Applied Numerical Mathematics*, vol. 27, no. 4, pp. 533–557, 1998.
- [114] J.-P. Bérenger, “An effective PML for the absorption of evanescent waves in waveguides,” *IEEE Microwave and Guided Wave Letters*, vol. 8, no. 5, pp. 188–190, 1998.
- [115] P. Rajagopal, M. Drozd, E. A. Skelton, M. J. Lowe, and R. V. Craster, “On the use of absorbing layers to simulate the propagation of elastic waves in unbounded isotropic media using commercially available finite element packages,” *NDT & E International*, vol. 51, pp. 30–40, 2012.

- [116] J. A. Roden and S. D. Gedney, “Convolution PML (CPML): An efficient FDTD implementation of the CFS–PML for arbitrary media,” *Microwave and optical technology letters*, vol. 27, no. 5, pp. 334–339, 2000.
- [117] D. Komatitsch and R. Martin, “An unsplit convolutional perfectly matched layer improved at grazing incidence for the seismic wave equation,” *Geophysics*, vol. 72, no. 5, pp. SM155–SM167, 2007.
- [118] J. Pettit, A. Walker, P. Cawley, and M. Lowe, “A stiffness reduction method for efficient absorption of waves at boundaries for use in commercial finite element codes,” *Ultrasonics*, vol. 54, no. 7, pp. 1868–1879, 2014.
- [119] J. W. S. B. Rayleigh, *The theory of sound*, vol. 2. Macmillan, 1896.
- [120] J.-F. Semblat, L. Lenti, and A. Gandomzadeh, “A simple multi-directional absorbing layer method to simulate elastic wave propagation in unbounded domains,” *International Journal for Numerical Methods in Engineering*, vol. 85, no. 12, pp. 1543–1563, 2011.
- [121] E. L. Wilson, “Automation of the finite element method—A personal historical view,” *Finite Elements in Analysis and Design*, vol. 13, no. 2-3, pp. 91–104, 1993.
- [122] R. H. MacNeal, *The NASTRAN theoretical manual*, vol. 221. Scientific and Technical Information Office, National Aeronautics and Space Administration, 1970.
- [123] E. L. Wilson, K. Bathe, F. Peterson, and H. Dovey, “SAP—A structural analysis program for linear systems,” *Nuclear Engineering and Design*, vol. 25, no. 2, pp. 257–274, 1973.
- [124] W. Bangerth, R. Hartmann, and G. Kanschat, “deal. II—a general-purpose object-oriented finite element library,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 33, no. 4, p. 24, 2007.

- [125] A. Logg, K.-A. Mardal, and G. Wells, *Automated solution of differential equations by the finite element method: The FEniCS book*, vol. 84. Springer Science & Business Media, 2012.
- [126] O. Pironneau, F. Hecht, A. Hyaric, and K. Ohtsuka, “FreeFEM,” *URL: <http://www.freefem.org>*, 2006.
- [127] G. Abaqus, “Abaqus 6.11,” 2011.
- [128] C. Multiphysics, “Introduction to COMSOL Multiphysics®,” *COMSOL Multiphysics, Burlington, MA, accessed Feb*, vol. 9, p. 2018, 1998.
- [129] S. Wolfram, *The mathematica book*, vol. 1. Wolfram Research, Inc., 2003.
- [130] C. B. Moler, *Numerical computing with MATLAB*. SIAM, 2004.
- [131] C. A. Felippa, “Introduction to finite element methods,” *University of Colorado*, vol. 885, 2004.
- [132] G. Strang and G. J. Fix, *An analysis of the finite element method*, vol. 212. Prentice-hall Englewood Cliffs, NJ, 1973.
- [133] K.-J. Bathe and E. L. Wilson, *Numerical methods in finite element analysis*. Prentice Hall, 1976.
- [134] R. D. Cook *et al.*, *Concepts and applications of finite element analysis*. John wiley & sons, 2007.
- [135] B. Delaunay *et al.*, “Sur la sphere vide,” *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, vol. 7, no. 793-800, pp. 1–2, 1934.
- [136] C. Geuzaine and J.-F. Remacle, “Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities,” *International journal for numerical methods in engineering*, vol. 79, no. 11, pp. 1309–1331, 2009.

- [137] N. Schlömer, “meshio: Tools for mesh files,” Mar. 2022. If you use this software, please cite it as below.
- [138] J. Robinson, “Automatic selection of redundancies in the matrix force method-rerank technique,” in *AIAA and Canadian Aeronautics and Space Institute, Joint Meeting*, p. 801, 1965.
- [139] S. N. Patnaik, L. Berke, and R. H. Gallagher, “Integrated force method versus displacement method for finite element analysis,” *Computers & Structures*, vol. 38, no. 4, pp. 377–407, 1991.
- [140] I. Kaljević, S. N. Patnaik, and D. A. Hopkins, “Development of finite elements for two-dimensional structural analysis using the integrated force method,” *Computers & Structures*, vol. 59, no. 4, pp. 691–706, 1996.
- [141] A. Van Pamel, G. Sha, S. I. Rokhlin, and M. J. Lowe, “Finite-element modelling of elastic wave propagation and scattering within heterogeneous media,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 473, no. 2197, p. 20160738, 2017.
- [142] O. Zienkiewicz and P. Bettess, “Fluid-structure dynamic interaction and wave forces. an introduction to numerical treatment,” *International Journal for Numerical Methods in Engineering*, vol. 13, no. 1, pp. 1–16, 1978.
- [143] A. Ghali, A. M. Neville, and T. G. Brown, *Structural analysis: a unified classical and matrix approach*. Crc Press, 2017.
- [144] T. Belytschko, D. Flanagan, and J. Kennedy, “Finite element methods with user-controlled meshes for fluid-structure interaction,” *Computer Methods in Applied Mechanics and Engineering*, vol. 33, no. 1-3, pp. 669–688, 1982.
- [145] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu, *The finite element method: its basis and fundamentals*. Elsevier, 2005.

- [146] L. Chen, “iFEM: an integrated finite element methods package in MATLAB,” tech. rep., 2009.
- [147] G. H. Golub and J. H. Welsch, “Calculation of gauss quadrature rules,” *Mathematics of computation*, vol. 23, no. 106, pp. 221–230, 1969.
- [148] G. Cowper, “Gaussian quadrature formulas for triangles,” *International Journal for Numerical Methods in Engineering*, vol. 7, no. 3, pp. 405–408, 1973.
- [149] M. Abramowitz, I. A. Stegun, and R. H. Romer, “Handbook of mathematical functions with formulas, graphs, and mathematical tables,” 1988.
- [150] L. N. Trefethen, “Is gauss quadrature better than clenshaw–curtis?,” *SIAM review*, vol. 50, no. 1, pp. 67–87, 2008.
- [151] D. N. Arnold and A. Logg, “Periodic table of the finite elements,” *Siam News*, vol. 47, no. 9, p. 212, 2014.
- [152] R. C. Kirby, “A general approach to transforming finite elements,” *The SMAI journal of computational mathematics*, vol. 4, pp. 197–224, 2018.
- [153] R. Rosen, “Matrix bandwidth minimization,” in *Proceedings of the 1968 23rd ACM national conference*, pp. 585–595, 1968.
- [154] E. Cuthill and J. McKee, “Reducing the bandwidth of sparse symmetric matrices,” in *Proceedings of the 1969 24th national conference*, pp. 157–172, 1969.
- [155] E. Van Brummelen, “Partitioned iterative solution methods for fluid–structure interaction,” *International Journal for Numerical Methods in Fluids*, vol. 65, no. 1-3, pp. 3–27, 2011.
- [156] L. S. Blackford, A. Petitet, R. Pozo, K. Remington, R. C. Whaley, J. Demmel, J. Dongarra, I. Duff, S. Hammarling, G. Henry, *et al.*, “An updated set of basic linear algebra subprograms (blas),” *ACM Transactions on Mathematical Software*, vol. 28, no. 2, pp. 135–151, 2002.

- [157] A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatin, A. Novikov, F. J. R Ruiz, J. Schrittwieser, G. Swirszcz, *et al.*, “Discovering faster matrix multiplication algorithms with reinforcement learning,” *Nature*, vol. 610, no. 7930, pp. 47–53, 2022.
- [158] P. Cristini, Y. Pailhas, R. Hamon, A. Xenaki, and G. Urso, “Influence of the sediment characteristics and of the level of burial on the acoustic response of a hollow cylinder in shallow water,” in *OCEANS 2019-Marseille*, pp. 1–6, IEEE, 2019.
- [159] J. Bucaro, B. Houston, H. Simpson, Z. Waters, M. Saniga, S. Dey, A. Sarkissian, D. Calvo, L. Kraus, and T. Yoder, “Wide area detection and identification of underwater uxo using structural acoustic sensors,” tech. rep., NAVAL RESEARCH LAB WASHINGTON DC PHYSICAL ACOUSTICS BRANCH, 2011.
- [160] O. Zienkiewicz, D. Kelly, and P. Bettess, “The sommerfeld (radiation) condition on infinite domains and its modelling in numerical procedures,” in *Computing Methods in Applied Sciences and Engineering, 1977, I*, pp. 169–203, Springer, 1979.
- [161] M. B. Drozd, *Efficient finite element modelling of ultrasound waves in elastic media*. PhD thesis, Imperial College London, 2008.
- [162] D. Burnett and M. Zampolli, “Festa: a 3-d finite element program for acoustic scattering from undersea targets,” *NURC Report SR*, vol. 394, 2004.
- [163] M. J. S. Lowe, *Plate waves for the NDT of diffusion bonded titanium*. PhD thesis, Imperial College London (University of London), 1992.
- [164] M. K. Kalkowski, J. M. Muggleton, and E. Rustighi, “Axisymmetric semi-analytical finite elements for modelling waves in buried/submerged fluid-filled waveguides,” *Computers Structures*, vol. 196, pp. 327–340, 2018.
- [165] S. G. Braun, D. J. Ewins, and S. S. Rao, *Encyclopedia of Vibration: FP*, vol. 2. Academic press, 2002.

- [166] D. Alleyne and P. Cawley, "A two-dimensional fourier transform method for the measurement of propagating multimode signals," *The Journal of the Acoustical Society of America*, vol. 89, no. 3, pp. 1159–1168, 1991.
- [167] T. Pialucha, C. Guyott, and P. Cawley, "Amplitude spectrum method for the measurement of phase velocity," *Ultrasonics*, vol. 27, no. 5, pp. 270–279, 1989.