# Exploring the Sparsity of Seismic Wave Propagation in Complex Media

Claudia Maria Haindl

University College
University of Oxford

*A thesis submitted for the degree of*
*Doctor of Philosophy*

Michaelmas 2021

## Abstract

Owing to continual improvements in computation power, we are getting closer to modelling seismic waves in ways that reflect the realistic complexity of the Earth. Nevertheless, modelling the true multi-scale structure of the Earth at large scale and/or high frequencies is currently still beyond our reach from a computational cost perspective. Hence, we rely on efficient modelling software to increase the scope of seismic simulations. Azimuthal Complexity Adaptation (ACA) is an innovative way of reducing the computational cost of seismic simulations by exploiting inherent sparsity of the wavefield. In this thesis, AxiSEM3D, an ACA-based modelling solver, is expanded to model wave propagation in local-scale settings. I find that the sparsity of wavefields persists in the presence of complex local-scale structures such as thrust faults and salt bodies and in models with full anisotropy. Further, I develop 3D wavefield scanning to locate complex areas of the wavefield and to investigate how they relate to structures in the subsurface model. The results suggest that high wavefield complexity is typically localised near abrupt wavespeed contrasts, on the low-wavespeed side thereof. I attempt to develop a new modelling solver which overcomes some of the limitations of AxiSEM3D. The resulting method is unstable, but it demonstrates the algorithmic complexity required to perform ACA in settings like coast lines. I conclude that a hybrid method which combines classic 3D modelling and ACA may be a simpler solution for increasing the scope of AxiSEM3D. Finally, I implement fault ruptures to simulate an earthquake scenario in the San Francisco Bay Area. Using ACA, I can include low-wavespeed bay muds which are usually neglected to save computational cost. Neglecting these muds leads to an underestimate of the shaking in the San Francisco Bay, even at low frequencies.

# Exploring the Sparsity of Seismic Wave Propagation in Complex Media




Claudia Maria Haindl

University College

University of Oxford

A thesis submitted for the degree of

*Doctor of Philosophy*

Michaelmas 2021

# Acknowledgements

I owe my deepest thanks to my supervisor, Prof. Tarje Nissen-Meyer, for his continued support, both on a technical and emotional level, and for his seemingly infinite enthusiasm and patience. Throughout these past four years, Prof. Nissen-Meyer has encouraged me to take my research into the direction that I am truly interested in, empowering me to make my own decisions all the while reassuring me that there is someone in my corner when the going gets tough. It is an immeasurable gift for a young scientist to be treated as an equal by their mentor.

Dr Kuangdai Leng did not only develop the software which is the basis of all the work presented here, but he also suffered though a myriad of emails and help requests from me. Effectively, he became my unofficial secondary supervisor. I want to express my sincerest gratitude for the many hours he spent helping me in his free time, despite having no obligation to do so and, moreover, despite also having a full-time job and a family with two children. This work would absolutely not have been possible without him.

I thank Prof. Gregory Beroza and Dr Arthur Rodgers for their advice and valuable suggestions regarding the San Francisco Bay Area project. Further, I thank Dr James Rickett for the interesting discussions and for setting up a promising collaboration with Schlumberger, which was unfortunately cancelled due to the COVID-19 pandemic.

I am very grateful to my colleagues at the Department of Earth Sciences for the patient explanations, engaging discussions and many technical troubleshoots. Dr Kasra Hosseini helped me finding my feet with Linux and GitHub which was a vital step on my way to becoming productive. I also particularly want to thank Dr Benjamin Fernando, who pioneered the use of AxiSEM3D. He helped me through many starting problems and he told me that I am not an idiot when I needed to hear it the most. He, Harrison Bowers and Dr Harriet Godwin further supported me by proof-reading this thesis. Ben, Harriet, Eoghan, Maria, Laura, Sophie, Matthew, Edward, Cornelius, Clancy and Lauren - your friendly faces made life in Oxford joyful and rewarding. Thank you for being there.

Also, I want to thank the faculty and administrative staff of the Department of Earth Sciences and University College for keeping everything running smoothly, for giving me a welcoming place to live and to work. And thank you for organising

so many great events. I will always fondly remember the Department Christmas parties and the amazing college formals.

I owe my thanks to the NERC CDT in Oil & Gas and the Radcliffe Scholarship who sponsored this thesis. Moreover, the CDT provided many invaluable courses and field trips throughout the years. In particular, I want to thank the CDT organisers Lorna Morrow and Anna Clark who are the heart and soul of the program.

To my friends from the Oil & Gas CDT I want to extend my gratitude for the many fun memories of the courses and field trips we did together. It is a very enriching experience to be connected to so many wonderful people who are at the same stage of their career, and who share the fears, hopes and dreams of this stage of life.

Finally, I want to thank my family for their unwavering faith and their support at every turn of the way. My partner, Achim Ditscher, has kept me sane during the lock-down and he has provided me with food and coffee whenever I was too absorbed to get these things myself. I would not have gotten this far without you.

# Abstract

Owing to continual improvements in computation power, we are getting closer to modelling seismic waves in ways that reflect the realistic complexity of the Earth. Nevertheless, modelling the true multi-scale structure of the Earth at large scale and/or high frequencies is currently still beyond our reach from a computational cost perspective. Hence, we rely on efficient modelling software to increase the scope of seismic simulations. Azimuthal Complexity Adaptation (ACA) is an innovative way of reducing the computational cost of seismic simulations by exploiting inherent sparsity of the wavefield. In this thesis, AxiSEM3D, an ACA-based modelling solver, is expanded to model wave propagation in local-scale settings. I find that the sparsity of wavefields persists in the presence of complex local-scale structures such as thrust faults and salt bodies and in models with full anisotropy. Further, I develop 3D wavefield scanning to locate complex areas of the wavefield and to investigate how they relate to structures in the subsurface model. The results suggest that high wavefield complexity is typically localised near abrupt wavespeed contrasts, on the low-wavespeed side thereof. I attempt to develop a new modelling solver which overcomes some of the limitations of AxiSEM3D. The resulting method is unstable, but it demonstrates the algorithmic complexity required to perform ACA in settings like coast lines. I conclude that a hybrid method which combines classic 3D modelling and ACA may be a simpler solution for increasing the scope of AxiSEM3D. Finally, I implement fault ruptures to simulate an earthquake scenario in the San Francisco Bay Area. Using ACA, I can include low-wavespeed bay muds which are usually neglected to save computational cost. Neglecting these muds leads to an underestimate of the shaking in the San Francisco Bay, even at low frequencies.

# Contents

**List of Abbreviations** **viii**

**1 Introduction** **1**

**2 Literature Review: A History of Wave Propagation Modelling** **4**

2.1 Introduction . . . . . 4

2.2 The Beginnings of Seismic Modelling . . . . . 5

2.3 Major Developments in Finite Difference Modelling . . . . . 6

2.4 Major Developments in Finite Element Modelling . . . . . 8

2.5 Alternative Global Wavefield Methods . . . . . 11

2.5.1 Normal Mode Summation . . . . . 11

2.5.2 Direct Solution Method . . . . . 12

2.5.3 Green's Functions Evaluation Method . . . . . 12

2.6 Alternative Grid Methods . . . . . 12

2.6.1 Pseudo-Spectral Method . . . . . 12

2.6.2 Lattice Method . . . . . 13

2.6.3 Discontinuous Galerkin Method . . . . . 14

2.7 Method Comparisons . . . . . 14

2.8 Remarks . . . . . 16

**3 Theoretical Background** **18**

3.1 Introduction . . . . . 18

3.2 Azimuthal Smoothness in Wavefields in 3D Media . . . . . 19

3.3 Equations Governing the ACA-SEM . . . . . 19

3.4 Wavefield Learning - Creating Azimuthal Complexity Maps . . . . . 25

3.5 Method Architecture . . . . . 26

3.5.1 Efficient Time-Stepping . . . . . 26

3.5.2 3D Models and Particle-Relabelling . . . . . 28

**4 Absorbing Boundaries** **32**
4.1 Introduction . . . 32
4.2 The Mixed Absorbing Boundary Method . . . 33
4.2.1 Absorbing Boundary Condition . . . 35
4.2.2 Sponge Boundary . . . 35
4.2.3 A Note on Stability . . . 36
4.3 Optimising the Sponge Parameters for the Mixed Boundary . . . 36
4.3.1 Test Set-Up . . . 37
4.3.2 Results for Elastic Media . . . 39
4.3.3 Results for Acoustic Media . . . 41
4.3.4 How does the Error Decrease with Sponge Width? . . . 44
4.4 Discussion . . . 46

**5 Azimuthal Complexity Adaptation at Local Scales** **47**
5.1 Introduction . . . 47
5.2 Exploiting Wavefield Smoothness - Verification and Benchmark . . . 48
5.2.1 The Benchmark Model . . . 48
5.2.2 Azimuthal Complexity Adaptation: Verification, wavefield learning, complexity maps . . . 49
5.2.3 Comparison with an Established Modeling Method . . . 55
5.3 A Complex 3D Setting - Limitations of Azimuthal Smoothness? . . . 59
5.4 Discussion . . . 65
5.5 Conclusions . . . 68

**6 Wavefield Scanning in 3D** **70**
6.1 Introduction . . . 70
6.2 Theory . . . 71
6.2.1 Wavefield Scanning in 2D . . . 71
6.2.2 Wavefield Scanning in 3D . . . 73
6.3 A First Look at 3D Scanning Results . . . 75
6.4 Comparing the Effect of an On-Axis and an Off-Axis Structure on Wavefield Complexity . . . 77
6.5 3D Wavefield Scanning of the Homogenised Overthrust Model . . . 80
6.6 Conclusions . . . 84

**7 Moving Beyond the Limitation of Azimuthal Continuity** **86**
7.1 Introduction . . . 86
7.2 The Short-Time Fourier Transform . . . 87
7.2.1 Window Functions . . . 89
7.2.2 Proof of Concept . . . 90

7.3 Handling Non-Periodic Fourier Derivatives . . . . . 91
7.4 Interpolation . . . . . 93
7.5 Changes to the Method Architecture of AxiSEM3D . . . . . 95
7.6 Numerical Instability . . . . . 96
7.7 A Way Forward . . . . . 97

**8 Using ACA for Hazard Prediction: A Case Study of the San Francisco Bay Area 98**
8.1 Introduction . . . . . 98
8.1.1 Local Geology . . . . . 100
8.1.2 Previous Simulations of Hayward Fault Earthquakes . . . . . 101
8.2 Simulation Set-Up . . . . . 102
8.3 Comparison of a Magnitude 6.5 Point Source and a Finite Rupture Earthquake Scenario . . . . . 104
8.3.1 Ground Motion . . . . . 104
8.3.2 Cost . . . . . 105
8.4 Testing the Effect of Extremely Low S-Wave Velocities . . . . . 108
8.5 Analysis of the Berkeley Earthquake on 4 Jan 2018 . . . . . 110
8.6 Conclusions . . . . . 118

**9 Discussion 120**
9.1 The Potential of ACA for Exploration-Type Settings . . . . . 120
9.2 Cost Scaling of ACA for Global-Scale and Local-Scale Models . . . 121
9.3 Other Lessons from 2D and 3D Wavefield Scanning . . . . . 122
9.4 Overcoming the Limitations of ACA . . . . . 123
9.5 Making Expensive Simulations Possible Through ACA . . . . . 124

**10 Conclusions 127**

**References 129**

# List of Abbreviations

**ABC** . . . . . Absorbing boundary condition.
**ABT** . . . . . Absorbing boundary term.
**ACA** . . . . . Azimuthal complexity adaptation.
**BFSM** . . . . . Buffered Fourier spectral method.
**CWT** . . . . . Continuous Wavelet transform.
**DOF** . . . . . Degrees of freedom.
**DWT** . . . . . Discrete Wavelet transform.
**FDM** . . . . . Finite difference method.
**FEM** . . . . . Finite element method.
**FFT** . . . . . Fast Fourier transform.
**FT** . . . . . Fourier transform.
**FWI** . . . . . Full waveform inversion.
**GLL** . . . . . Gauss-Lobatto-Legendre.
**GLC** . . . . . Gauss-Lobatto-Chebyshev.
**MPI** . . . . . Message passing interface.
**NFFT** . . . . . Non-uniform Fourier transform.
**OLA** . . . . . Overlap-and-add.
**PGV** . . . . . Peak ground velocity.
**PML** . . . . . Perfectly matched layers.
**PSM** . . . . . Pseudo-spectral method.
**RMS** . . . . . Root mean squared.
**SEG** . . . . . Society of exploration geophysicists.
**SEM** . . . . . Spectral element method.
**SFBA** . . . . . San Francisco Bay Area.
**STFT** . . . . . Short-term Fourier transform.
**USGS** . . . . . United States Geological Survey.
**WT** . . . . . Wavelet transform.

*Somewhere, something incredible is waiting to be known.*

— Carl Sagan

# 1 Introduction

Wave propagation modelling plays a major role in modern seismology. It is used for ground motion predictions, but also for tracing back the sources of earthquakes and for imaging the subsurface. When constructing a subsurface model (or, equivalently, a source model), the accuracy of this model can be tested by running a seismic simulation. The synthetic data obtained in this manner is compared to real data, and the misfit between the two datasets helps to identify errors in the model. Such a process can even be formalised into a workflow where the model is updated iteratively, minimising the misfit between real and synthetic data on each turn. This is called seismic inversion and it is one of the main state-of-the-art imaging methods for the subsurface on global, regional and local scales.

Immense progress has been made in the last decades when it comes to numerical methods for modelling seismic waves, making them more accurate, versatile and efficient. Nevertheless, the high computational cost of numerical modelling at high frequencies and in 3D remains a challenge (e.g. Virieux et al., 2009; Vigh et al., 2019). This challenge will eventually be solved if the development of computational resources continues at its current exponential rate but, until then, we rely on innovative solutions to approximate the true complexity of wave propagation in the subsurface of the Earth.

Azimuthal Complexity Adaptation (ACA) is one such innovative solution. It offers a way to adapt the resolution of a seismic simulation to the observed complexity of the wavefield, thus eliminating oversampling and making the simulation more cost efficient. This method was introduced by Leng et al. (2016) as part of the seismic modelling solver, AxiSEM3D, which was originally developed for global scale problems. Leng et al. show that, for tomographic models, ACA can reduce the way simulation cost scales with frequency from power 4 to power 3, which is a game-changing development.

In global seismology, the multi-scale nature of the Earth means that there is always something to be gained from modelling higher frequencies and higher resolutions. But there are more fields where ACA could be beneficial. In this thesis, I put the focus on two other cost-sensitive applications for seismic modelling: resource exploration settings and seismic hazard prediction. In both of these areas the need for more efficient modelling methods cannot be overstated.

In exploration, wave propagation modelling is used for full waveform inversion and, more recently, for migration velocity analysis. Both are expensive iterative methods which require many seismic simulations. In some settings cheaper imaging methods which neglect or approximate aspects of wave physics (e.g. Kirchhoff migration or Gaussian beam migration) are viable. But such approximations lead to significant errors in the presence of steep wavespeed contrasts (Leveille et al., 2011; Warner et al., 2013; Chauris and Cocher, 2017; Alkhalifah and Wu, 2017). Even in studies where the full range of inversion methods is applied, not all wave physics can be accounted for: with the exception of some experimental studies (e.g. Raknes et al., 2015), the exploration industry is still limited to using acoustic methods, since doing inversions with full elasticity is prohibitively expensive.

In seismic hazard prediction, hypothetical sources are modelled in a region of interest to estimate the shaking and thus the damage that may be caused by a future event. Here, the challenge lies in the uncertainty inherent to the seismic source used to approximate a future event. Proper treatment of this uncertainty necessitates probabilistic modelling which, again, consists in countless seismic simulations (e.g.

Graves et al., 2011). Technically, a high level of uncertainty also exists for the subsurface model used in these simulations, but accounting for this lies well beyond our current computational capabilities.

Inspired by these needs and the promising results from Leng et al. (2016), I investigate if the improved cost scaling of ACA also applies to typical exploration and earthquake hazard settings. I expand and adapt the functionality of AxiSEM3D for local-scale models, and I develop methods to overcome current limitations of ACA and to answer open questions about the nature of azimuthal wavefield complexity.

To start off, I conduct an extensive literature review looking into the history of seismic modelling, and into the divergence of approaches currently taken by global seismologists and exploration seismologists. Then, I provide a theoretical background explaining the concept of azimuthal wavefield complexity and the wave physics behind AxiSEM3D. This leads into my first major addition to the code: the absorbing boundaries. This new boundary condition (plus a number of minor changes to AxiSEM3D) enabled us to conduct and publish an investigation of the performance of ACA in exploration-type settings. The resulting paper, published in Geophysics in 2021, is presented as a chapter in this thesis. To answer questions raised but not yet addressed by the paper, I develop a method to obtain a 3D representation of azimuthal wavefield complexity. I go on to use this new method to investigate the azimuthal wavefield complexity in a fully anisotropic model of another exploration-type setting. Having taken this first step towards disentangling the azimuthal dimension in ACA, I attempt to develop a new solver which overcomes many of the current limitations of AxiSEM3D. Such limitations include that sampling rate is fixed along the azimuth, that the modelling domain must take the shape of a cylinder or sphere and that fluid and solid media must be azimuthally continuous. Finally, I move on to modelling earthquake scenarios with ACA. I investigate the azimuthal complexity of earthquake ruptures and I employ ACA to model the notoriously difficult (from a computational cost perspective) San Francisco Bay Area.

*Work is a way of bringing order to chaos, and there's a basic satisfaction in seeing that we are able to make something a little more coherent by the end of the day.*

— Alain de Botton

# 2

# Literature Review: A History of Wave Propagation Modelling

## Contents


**2.1 Introduction . . . . . 4**
**2.2 The Beginnings of Seismic Modelling . . . . . 5**
**2.3 Major Developments in Finite Difference Modelling . 6**
**2.4 Major Developments in Finite Element Modelling . . 8**
**2.5 Alternative Global Wavefield Methods . . . . . 11**
2.5.1 Normal Mode Summation . . . . . 11
2.5.2 Direct Solution Method . . . . . 12
2.5.3 Green's Functions Evaluation Method . . . . . 12
**2.6 Alternative Grid Methods . . . . . 12**
2.6.1 Pseudo-Spectral Method . . . . . 12
2.6.2 Lattice Method . . . . . 13
2.6.3 Discontinuous Galerkin Method . . . . . 14
**2.7 Method Comparisons . . . . . 14**
**2.8 Remarks . . . . . 16**


## 2.1 Introduction

The method around which this thesis is built, AxiSEM3D, is an efficient 3D re-invention of the global 2.5D spectral element method, AxiSEM. This new 3D method has the additional capability to run local-scale simulations which is the focus of

this work. When testing and optimising AxiSEM3D for exploration-type models, I find myself confronted with the gap between the more academically oriented global modelling community and seismic modelling in industry. It turns out that this gap is indeed immense, starting with the fact that spectral element modelling, the basis of AxiSEM3D, is barely used in exploration. The stark difference in approaches to wave propagation taken by seismologists in industry and in different research groups has likely evolved over time. It may be caused by a combination of well-supported reasons, in terms of which method is appropriate for a certain type of problem, and reasons of a more psychological nature: how easily a method can be understood, availability and even group dynamics. In this chapter, I investigate how the different approaches emerged and separated, giving a historical overview over the development of seismic modelling.

The finite difference method and the spectral element method have emerged as "the big two" modelling methods in seismology, and the development of each is considered separately. I also provide an overview over some alternative methods, some of which were competitive for a period of time and then faded back into obscurity, and some of which are still relevant today. This part is separated into two sections: one for a set of methods that deal with seismic modes in a way such that they can only be used for global wavefields and one for methods using grids which can be applied in both local and global settings. Finally, I discuss model comparisons and benchmarks. The goal of this review is to answer the questions: How did the field of seismic modelling get to where it is today? And how does a local-scale expansion of AxiSEM3D fit within the state-of-the-art?

## 2.2 The Beginnings of Seismic Modelling

The research of numerical wave propagation has always been, and still is, closely tied to the development of computing technology. While the basis of finite difference methods was known since the $18^{\text{th}}$ century, the first published attempts at using them to model seismic waves were made in the 1960s which coincides with the wider availability of transistor computers. Notably, Maenchen (1963) wrote a generalised

open source code capable of modelling visco-elastic waves in heterogeneous media in 2D. Finite element methods emerged around the same time (Turner et al., 1956; Clough, 1960) and they were initially applied to seismic waves by Costantino (1967). The development of finite difference methods was pushed by a research group based in Israel who successfully modelled layered half-spaces (Alterman and Karal, 1968), acoustic homogeneous spheres (Alterman and Kornfeld, 1968), acoustic heterogeneous spheres (Alterman and Aboudi, 1968) and elastic heterogeneous spheres (Alterman et al., 1970). First comparisons of synthetics created by numerical modelling with more cheaply obtainable seismograms created using ray theory proved the importance of the more rigorous approach (Hilterman, 1970). Moreover, the new methods allowed modelling of surface waves (Boore, 1970).

Arguably, more widespread interest in the numerical modelling of seismic waves was sparked with the release of the book "Methods in Computational Physics: Advances in Research and Applications" in 1972, where comprehensive introductions to both finite difference and finite element methods, respectively, were made available to the general research community (Boore, 1972b; Lysmer and Drake, 1972).

## 2.3 Major Developments in Finite Difference Modelling

In the 70s and early 80s, FDMs were adapted and used to model topography (Boore, 1972a) and the effects of dynamic faulting (Madariaga, 1976; Virieux and Madariaga, 1982). Kelly et al. (1976) states that FD modelling was already being used in the exploration industry to aid interpretations of seismic data. However, there were still issues associated with the approach. As discussed by Alford et al. (1974), low order FDMs suffered from poor dispersion whereas high order FDMs are more memory demanding and less flexible. This problem was addressed by Virieux (1984) and Virieux (1986) who developed staggered-grid velocity-stress formulations of the FDM (shown in Figure 2.1), thus significantly reducing dispersion. Progress to this end continued as Mora (1989) and Crase (1990) introduced an alternative way of handling the spatial derivatives, by performing band-limited convolutions,

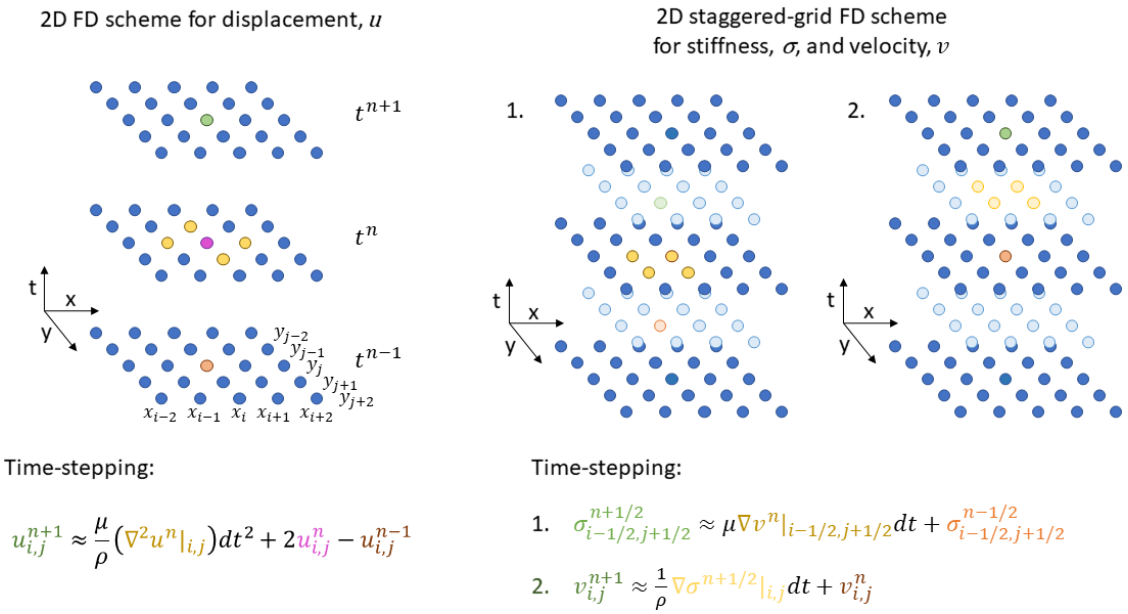



**Figure 2.1:** Schematic explanation of FDMs comparing a non-staggered grid and a staggered grid approach. The shown examples are simple low order schemes applicable to homogeneous models. The equations for the spatial derivatives (yellow) were omitted for brevity but would be central difference schemes using the grid points marked in yellow and purple. $\mu$ and $\rho$ stand for the shear modulus and density of the medium, respectively.

which again reduced dispersion errors. Around the same time, the first successful full-waveform inversions were performed using FD frameworks (Gauthier et al., 1985; Gauthier et al., 1986; Crase et al., 1990). In the 90s, ability of state-of-the-art FDMs to handle strong anisotropy was demonstrated by Igel et al. (1995) and Igel and Weber (1995), and Chaljub and Tarantola (1997) made first attempts at modelling global wave propagation with 1D models. A final major step in improving the accuracy of FDMs was made when Geller and Takeuchi (1998) and Takeuchi and Geller (2000) introduced a systematic method to find optimally accurate spacial and temporal operators for any FD scheme of a given order. The representation of free surfaces and topography remains a challenge with FDMs, although progress was made by Komatitsch et al. (1996) who implemented topography with undulating meshes, Robertsson (1996) who developed an improved free surface boundary condition and Hestholm (2003) who increased the stability of the free surface in elastic simulations using non-equidistant grids.

## 2.4 Major Developments in Finite Element Modelling

The main difference between element-based modelling methods (namely finite element methods, FEMs, and spectral element methods, SEMs) and FDMs is that, in element-based modelling, the wave equation is solved in its integral form. This requires integration and therefore interpolation of the wavefield between nodes by means of shape functions. A basic explanation of element-based methods is shown in Figure 2.2. While more complex to program, element-based methods implicitly handle free surfaces and allow for irregular meshes. The advantages of this were demonstrated early on by Smith (1975), who incorporated topography as well as mesh refinements in his simulation. Schlue (1979) expanded Lysmer and Drake's FE formulation to handle surface waves in the presence of 3D topography. Despite this, the application of FEMs remained limited because low-order schemes

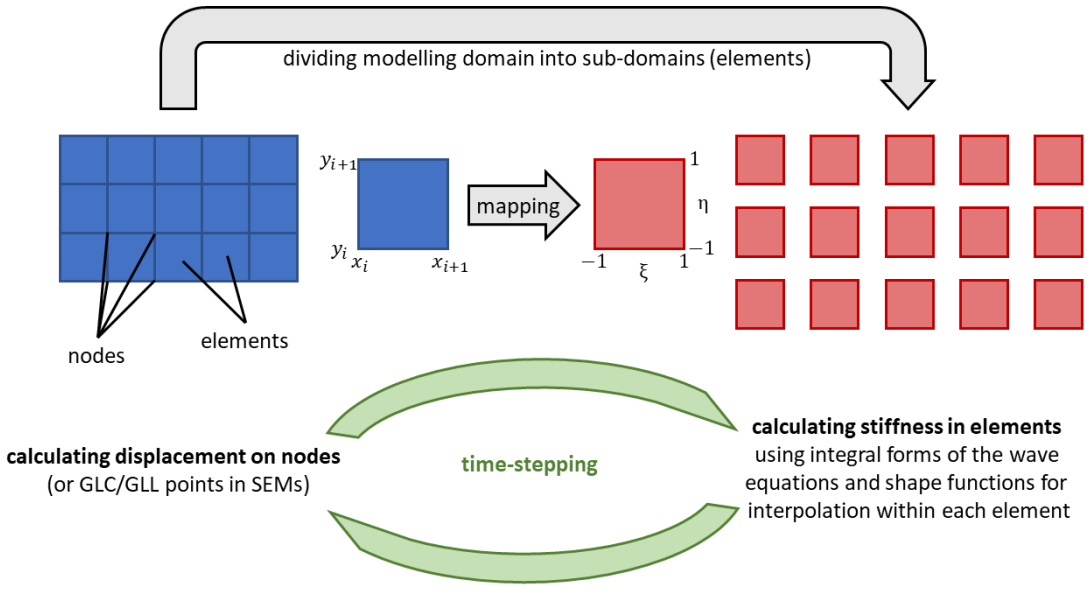



**Figure 2.2:** Schematic explanation of element-based seismic modelling methods.

suffer from poor dispersion properties, as shown in a direct comparison of FEM and FDM by Marfurt (1984), while high-order schemes produce parasitic modes (Mullen and Belytschko, 1982). Serón et al. (1990) continued to work with low-order

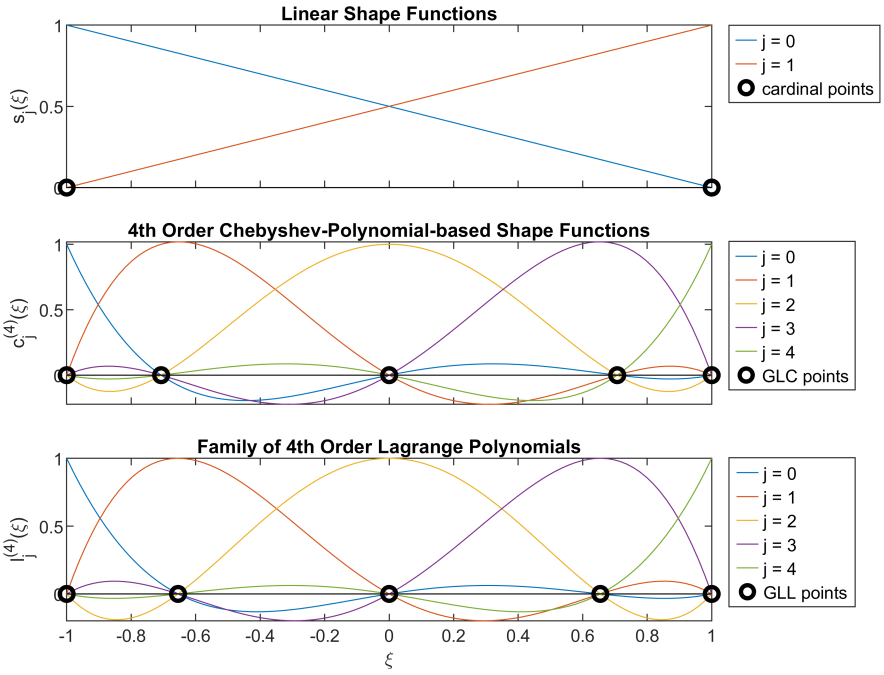



**Figure 2.3:** Shape functions commonly used for element-based modelling methods. For the polynomial shape functions, order $N = 4$ is shown. $\xi$ denotes the mapped coordinate as shown in Fig. 2.2. **Top:** Linear shape functions as used by FEMs. **Middle:** Shape functions which are a linear combination of Chebyshev polynomials, as used by early SEMs. The wavefield is known exactly on $N + 1$ co-location points between the nodes, called Gauss-Lobatto-Chebyshev (GLC) points. **Bottom:** Lagrange polynomials which are used as shape functions by later SEMs. Here, the co-location points between the nodes are called Gauss-Lobatto-Legendre (GLL) points.

FEMs, finding the most accurate and efficient spatial and temporal operators among commonly used ones and later producing an open source code (Serón et al., 1996).

The breakthrough of this family of methods occurred in the mid-to-late 90s, when the transition to spectral elements was made. The SEM had been developed in the field of fluid dynamics (Patera, 1984) and is commonly advertised as having the set-up and flexibility of an FEM but the accuracy and low dispersion of a pseudo-spectral method. The increase in accuracy is achieved by replacing the shape functions, which are linear in the traditional FEM, with specific polynomials and introducing non-uniformly spaced co-location points. Some commonly used shape functions are shown in Figure 2.3. This approach was first applied to acoustic wave propagation by Padovani et al. (1994) and Seriani and Priolo (1994) and soon extended to elastic waves by Carcione (1996). These early SEMs used Chebyshev

polynomials as basis functions (Fig. 2.3, middle). When these polynomials were replaced by Lagrange polynomials (Fig. 2.3, bottom) the method became more efficient as this made the mass matrix diagonal (Faccioli et al., 1996; Faccioli et al., 1997; Komatitsch and Vilotte, 1998). The mass matrix contains the density of the medium, and it being diagonal reduces the number of non-zero coefficients and thus memory needs and computation steps required for the division by density which occurs in each time step of an SEM. By this point, applications of the traditional FEM had become rare (e.g. Bao et al., 1998) while the Lagrange-based SEM was taken up by many research groups and developed into the range of open source codes available today. Komatitsch and Tromp (2002a) and Komatitsch and Tromp (2002b) developed a global modelling code, including gravity and rotation of the Earth, which Capdeville et al. (2003) expanded by adding solid-fluid interactions where the wavefield in the acoustic region is calculated in terms of normal modes. Rigorous modelling of the fluid as well as self-gravitation was added by Chaljub and Valette (2004). Following advances in Chebyshev-based SEMs, Seriani and Oliveira (2007) introduced the idea of lumped and blended operators to seismic modelling. These operators offer a kind of workaround which makes spectral element modelling with Chebyshev polynomials equally as efficient as it is with Lagrange polynomials. This idea did not catch on, however, and Seriani and Oliveira (2008) showed that the Chebyshev-based method is indeed slightly more dispersive than the Lagrange-based method. Around the same time began the development of one of the first open source SEM codes, AxiSEM, which was designed for modelling wavefield in an axisymmetric earth. Nissen-Meyer et al. (2007a) laid the groundwork by introducing a source decomposition to represent general moment tensor sources in the axisymmetric framework, Nissen-Meyer et al. (2007b) introduced the algorithm and Nissen-Meyer et al. (2008) added solid-fluid interaction. The final release which further added visco-elasticity and anisotropy was given in Nissen-Meyer et al. (2014). By this point, another open source code, SPECFEM3D Version 2.0 (Peter et al., 2011), which handles general heterogeneous 3D models in Cartesian and Spherical meshes, had also been published. While full 3D modelling is of course desirable,

applications on a global scale are limited by the availability of computational power and memory. This was addressed by Leng et al. (2016), who developed the open source code AxiSEM3D, an algorithm which is based on AxiSEM but which can transition gradually from 2D to 3D modelling using azimuthal pseudo-spectra. Thanks to the azimuthal smoothness inherent to wavefields, full 3D resolution is often not required and the cost of the simulation can be reduced accordingly without loss of accuracy. The mixing of a 2D mesh with azimuthal pseudo-spectra, however, limits the use of irregular meshes in AxiSEM3D, an issue which was partly alleviated using particle relabelling (Leng et al., 2019). van Driel et al. (2020) also recognised the potential of wavefield smoothness and incorporated wide azimuthal spacing in the global 3D meshes of their commercial software package, Salvus.

## 2.5 Alternative Global Wavefield Methods

### 2.5.1 Normal Mode Summation

As first attempts were made at modelling wave propagation in spheres with finite difference methods, another approach emerged which allowed the calculation of global wavefields more cheaply: a displacement formulation of normal mode summation was suggested by Gilbert (1971). Normal modes had been used prior to this but in a more general descriptive manner of Earth seismicity, studying the effects of ellipticity and rotation (e.g. Dahlen, 1968). 3D heterogeneity was successfully included by Woodhouse and Dahlen (1978). Tromp and Dahlen (1990) calculated more accurate seismograms by moving to higher order Born approximations, an effort repeated by Friederich et al. (1993) who only considered forward scattering which allows for cheaper calculations. Improved formulations and algorithms which take into account attenuation, self-gravitation and anisotropy have since been proposed (e.g. Yang et al., 2017; Liu and Zhang, 2017) and many open source codes are in circulation which allow the efficient calculation of low-frequency wavefields. The cost of normal mode summation scales exponentially with the range of frequencies considered due to mode coupling, so for periods below 50s other methods need to be considered.

### 2.5.2 Direct Solution Method

Geller and Ohminato (1994) introduced their direct solution method (DSM) as an alternative to normal mode summation which is capable of dealing with high frequencies. As the name suggests the method consists in directly solving the momentum equation by means of a matrix inversion. Consequently, an inversion is necessary for every source-receiver pair, but the frequency limitations of normal mode summation methods no longer apply. The operators are improved by Geller and Takeuchi (1995) and extended to spherical coordinates by Cummins et al. (1997). Takeuchi et al. (2000) managed to speed up the method by treating heterogeneities as perturbations, i.e. using the Born approximation, and anisotropy was added by Kawai et al. (2006).

### 2.5.3 Green's Functions Evaluation Method

A third approach to computing seismograms was released with the GEMINI code, based on the Green's functions evaluation method (Friederich and Dalkolmo, 1995). This is another highly efficient method which, however, is currently limited to spherically symmetric models.

## 2.6 Alternative Grid Methods

### 2.6.1 Pseudo-Spectral Method

Like many other numerical modelling methods, the pseudo-spectral method (formerly also Fourier method), sometimes considered a variant of the finite difference method, was first used in the field of fluid dynamics (e.g. Orszag, 1971; Kreiss and Oliger, 1972). It was applied to acoustic waves for the first time by Gazdag (1981) and extended to elastic waves by Kosloff et al. (1984). The PSM partially operates in the spatial Fourier domain, where the otherwise complex step of calculating gradients in a discrete grid becomes a simple multiplication. This reduces the number of required grid points and leads to smaller dispersive errors than FEMs and FDMs had at the time. The method held up well in comparisons and it was received with great

enthusiasm (e.g. Kosloff and Baysal, 1982; Fornberg, 1987; Fornberg, 1988). Kosloff et al. (1990) developed the Chebyshev Spectral Method, where the PSM is combined with an SEM in the vertical direction to keep some of the benefit of the PSM while being able to implement an accurate free-surface boundary condition. Global simulations were performed successfully in 2D (Furumura et al., 1998) and 3D (Igel, 1999), where the latter employed the Chebyshev Spectral Method. As the FDM and SEM continued to be optimised, however, the advantages of the PSM diminished in comparison (e.g. Mizutani et al., 2000). Moreover, a "pure" PSM cannot be parallelised which rendered the method unattractive with the rise of supercomputers.

### 2.6.2 Lattice Method

The basis of lattice methods is the idea of treating seismic waves as the movements and collisions of a myriad of small particles, much like in an ideal gas. Rothman (1987) first demonstrated the feasibility of this approach for a homogeneous acoustic case using lattice gas automata. Limited inhomogeneities were introduced by Huang et al. (1988) and full heterogeneity was enabled when Mora and Maillot (1990) replaced the lattice gas automata with the phononic lattice solid, i.e. they replaced free moving gas particles with particle representations of lattice vibrations which are naturally tied to medium properties. The modelling of scattering effects was further improved by implementing the Boltzmann lattice solid (Mora, 1992). When Toomey and Bean (2000) extended the method to elastic media, they noticed large dispersion errors, an issue which limits the adaptation of the method to this day. Nevertheless, the method was taken up again based on the advantage of small communication overheads in parallel implementations (O'Brien and Bean, 2004; O'Brien, 2014; Zergani et al., 2016). Xia et al. (2017) use a version of the lattice method called lattice spring method and argue that, while the method requires a very fine mesh to mitigate the dispersion error, it is more flexible in its application than FDMs are, and it may be of use for modelling crack propagation and similar problems.

### 2.6.3 Discontinuous Galerkin Method

Another method particularly suited to crack propagations modelling is the Discontinuous Galerkin Method, which stems from the field of neutron transport (Reed and Hill, 1973). The method has similarities with SEMs, but here, the wavefield is not required to be fully continuous between elements. Instead, there is a term which punishes discontinuities. This makes for a more expensive but also more flexible simulation which can easily account for cracks and similar phenomena. An implementation of this method for seismic applications, called ADER-DG, was developed quite recently and published in a long series of release papers accounting for visco-elasticity, anisotropy and local time-stepping (Käser and Dumbser, 2006; Dumbser and Käser, 2006; Käser et al., 2007; de la Puente et al., 2007; Dumbser et al., 2007). Duru et al. (2017) further improved the boundary conditions of the proposed scheme.

## 2.7 Method Comparisons

When it comes to choosing a numerical method for a given seismic modelling problem, the variety is overwhelming. Hence, many authors have benchmarked and compared a selection of methods in an attempt to give a recommendation. Some of the more recent model comparisons are introduced here.

The COSY project by Igel et al. (2000) introduced a benchmark for global modelling methods and compared normal mode summation, the DSM, GEMINI and a global FD implementation. For a spherically symmetric model, they managed to get a good agreement between the results of all the methods with the exception of the FDM which struggled with dispersion errors in the surface waves. However, Igel et al. (2000) conceded that they did not use the most optimal available operators in their FD implementation. When using a 3D model, they find considerable differences between the solutions which cannot be fully explained. They also find that low-order (up to order 3) Born approximations in both the DSM and NMS leads to erroneous seismograms for a typical 3D Earth model.

Perhaps infamously, the Grenoble benchmark was a comparison of grid methods in an alpine valley setting (Chaljub et al., 2010). The tested methods were a staggered-grid FDM, the ADER-DGM and two implementations of the SEM. A good agreement was found at lower frequencies (up to 1 Hz). The increasing divergence between the synthetic waveforms observed at higher frequencies was put down to the differing implementations of physical effects like attenuation and grid set-up etc. Chaljub et al. (2010) note that the FDM simulations ran 35% faster than the SEM simulations and that the ADER-DG was significantly more costly to run than both. On the other hand, the FDM implementation used in their test was not able to handle topography, whereas the other two methods account for it implicitly.

Taking a more theoretical approach, Moczo et al. (2011) had an in-depth look at the dispersion errors in different implementations of the FDM, FEM, SEM and DGM. They used isotropic homogeneous media with different ratios of P- and S-wave velocity and tested how densely spaced the grid must be in order to achieve a given accuracy. The best performing method, i.e. the method which consistently gave accurate results with the largest grid spacing, was the staggered grid FDM. The DGM was the most sensitive to the $v_p/v_s$ ratio, requiring significant up-sampling in high ratio models.

Another benchmark based on a real setting was the Euroseistest Verification and Validation Project (E2VP) which looked at the Mygdonian basin in Greece (Maufroy et al., 2015). The tested methods were two FDMs, an SEM, a PSM and a DGM. After the issues faced when fitting the different synthetics in the Grenoble case (i.e. significant mismatch at higher frequencies with no way of identifying the correct solution), the submitted datasets were tested against real data this time. Similar problems persisted, however, and this was investigated in detail by Chaljub et al. (2015). They found that a significant source of errors was the improper treatment of material discontinuities. The proper handling of such discontinuities could easily be achieved in an SEM framework whereas the other methods required workarounds. Eventually, all methods were brought into better agreement but the

workarounds proved to be an additional source of error requiring several iterations of altered input models until the methods gave the desired results.

## 2.8 Remarks

Considering the developmental history of all the methods described above, it appears on first sight that methods like the PSM were given up on for good reasons. However, the methods that prevailed did so because of continuing development and contributions from several research groups which lead to break-throughs like the staggered grid formulation for the FDM. So the question remains, did some methods lose traction simply because their development stopped or did the development stop because there was no way to keep the methods competitive?

The accessibility of modelling methods appears to play a huge role in their continued use and development. A positive example of this is the FDM which is very simple to implement. It is completely feasible to do the necessary research and write a working 3D FDM code within a day, without previous experience in the subject. This opens the door for developers from many different backgrounds. Moreover, understanding the underlying algorithm is important when it comes to trusting the output of a solver. So, it is not surprising that many research groups are committed to keeping this accessible method competitive with newer, more complex methods.

The different approaches of exploration and global seismologists are, for the large part, due to the different strengths of FDMs and SEMs: As stated by Chaljub et al. (2010) FDMs are about 35% faster than SEMs, but SEMs are better at handling topography. In reflection seismology (commonly used in industry), the bulk of the useful information is contained in body waves, which are not sensitive to topography. Surface waves, in the other hand, are very sensitive to topography but they only become relevant at larger scales i.e. in regional and global seismology. Thus, in global seismology, the additional accuracy of SEMs is considered worth the additional cost whereas, in exploration seismology, the benefit of using an SEM is considered minimal. From personal communications (J. Rickett (Schlumberger), June 2019, and M. Ravasi (Equinor), May 2020) I found that the complexity of SEMs

poses an additional hindrance to their incorporation in exploration workflows, as does corporate routine. FDMs have been in use for a long time (20 years longer than SEMs) and the various in-house codes are tried and tested and highly optimised.

Leaving aside all technical detail and looking at AxiSEM3D from a purely historical perspective, there are some things to be learned. On one hand, the complexity of the underlying algorithm makes AxiSEM3D rather inaccessible, and this impedes continual development outside of the core research group. On the other hand, the development of new and innovative methods like AxiSEM3D is vital for the progress of the field of computational seismology. It has been concluded repeatedly (e.g. Virieux et al., 2011; Chaljub et al., 2015) that having a large repertoire of modelling methods is important because each use case has different requirements and because the agreement between the synthetics from different methods is still not very good. Relying only on the method which is easiest to implement would make us blind to the errors which evidently still exist.

*The practice of science happens at the border between the known and the unknown. Standing on the shoulders of giants, we peer into the darkness with eyes opened not in fear but in wonder.*

— Brian Cox

# 3 Theoretical Background

## Contents


**3.1 Introduction . . . . . 18**
**3.2 Azimuthal Smoothness in Wavefields in 3D Media . . 19**
**3.3 Equations Governing the ACA-SEM . . . . . 19**
**3.4 Wavefield Learning - Creating Azimuthal Complexity Maps . . . . . 25**
**3.5 Method Architecture . . . . . 26**
3.5.1 Efficient Time-Stepping . . . . . 26
3.5.2 3D Models and Particle-Relabelling . . . . . 28


## 3.1 Introduction

The seismic modelling software AxiSEM3D (developed by Leng et al., 2016; Leng, 2017) serves as the basis of all work conducted as a part of this thesis. This chapter explains the theoretical foundation of the underlying method and gives an overview over the method architecture of AxiSEM3D and the implementation of 3D models. I use the nomenclature ACA-SEM for azimuthal complexity adapting spectral element method (introduced in Haindl et al., 2021) to distinguish the specific software implementation, i.e. AxiSEM3D, from the general method.

## 3.2 Azimuthal Smoothness in Wavefields in 3D Media

The basis for the efficiency of the ACA approach is the inherent azimuthal smoothness of wavefields. Wavefronts spread spherically from the source, and only strong heterogeneities can significantly disrupt this shape through a combination of off-plane scattering and refraction. Such strong variations are mainly found along the depth axis and thus the azimuthal symmetry is often largely left intact. By choosing a cylindrical coordinate system centred around the source, this can be exploited to reduce computational effort and memory requirements of seismic simulations. Figure 3.1 illustrates the concept of azimuthal smoothness and, in Figure 3.2, a wavefield snapshot with apparent azimuthal smoothness is shown.

The degree of azimuthal smoothness in a given simulation depends on the 3D model and the seismic source. Hence, it should be expected that the efficacy of the ACA approach is somewhat situational. The limits of the ACA approach are explored and tested in later parts of this thesis.

## 3.3 Equations Governing the ACA-SEM

The ACA-SEM is a hybrid of an SEM in the radial and vertical directions of a cylindrical coordinate system and a PSM in the azimuthal direction. In the following, the mathematical basis of this method is derived for solid and fluid media.

Consider a 3D velocity model with fluid regions, such as an ocean or the outer core, and solid regions. Within the solid, waves are described by the momentum equations for elastodynamics:

$$\rho\ddot{\mathbf{u}} - \nabla \cdot \mathbf{T} = \mathbf{f}_\mathrm{s}, \tag{3.1}$$

$$\mathbf{T} = \mathbf{C} \colon \nabla\mathbf{u}, \tag{3.2}$$

where $\rho$ is the medium density, $\mathbf{u}(\mathbf{r}, t)$ is the displacement and $\ddot{\mathbf{u}}$ its second time derivative, $\mathbf{f}_\mathrm{s}$ is a moment tensor source, $\mathbf{T}$ is the second-order stress tensor, $\mathbf{C}$ denotes the fourth-order elastic stiffness tensor and $:$ stands for a tensor

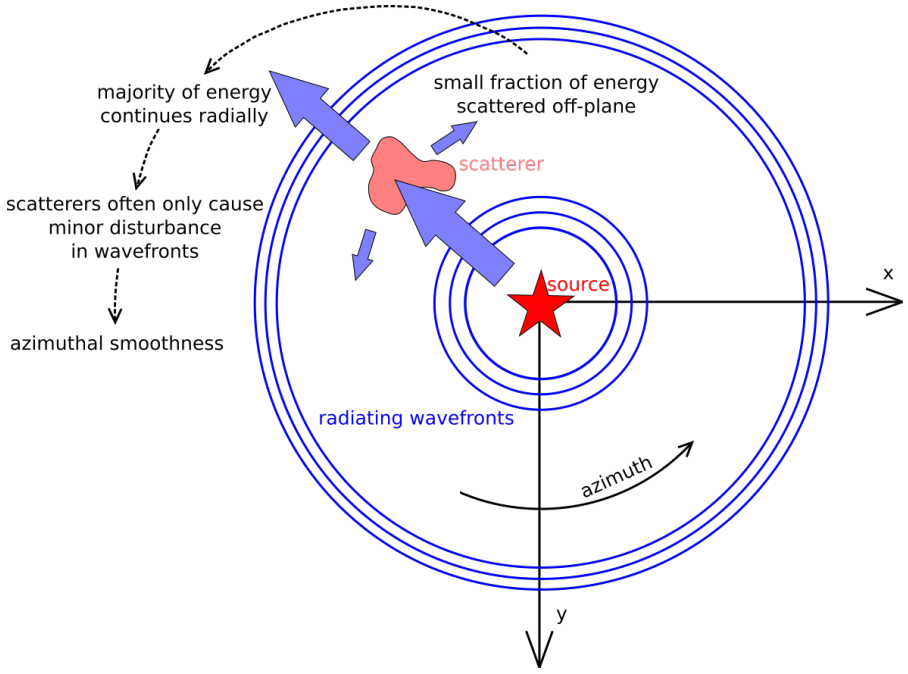



**Figure 3.1:** Sketch illustrating wavefield smoothness along the azimuth of a coordinate system placed with the depth axis on the seismic source. Radial propagation of the wavefield outward from the source is referred to as in-plane, whereas propagation deviating from this direction is referred to as off-plane.

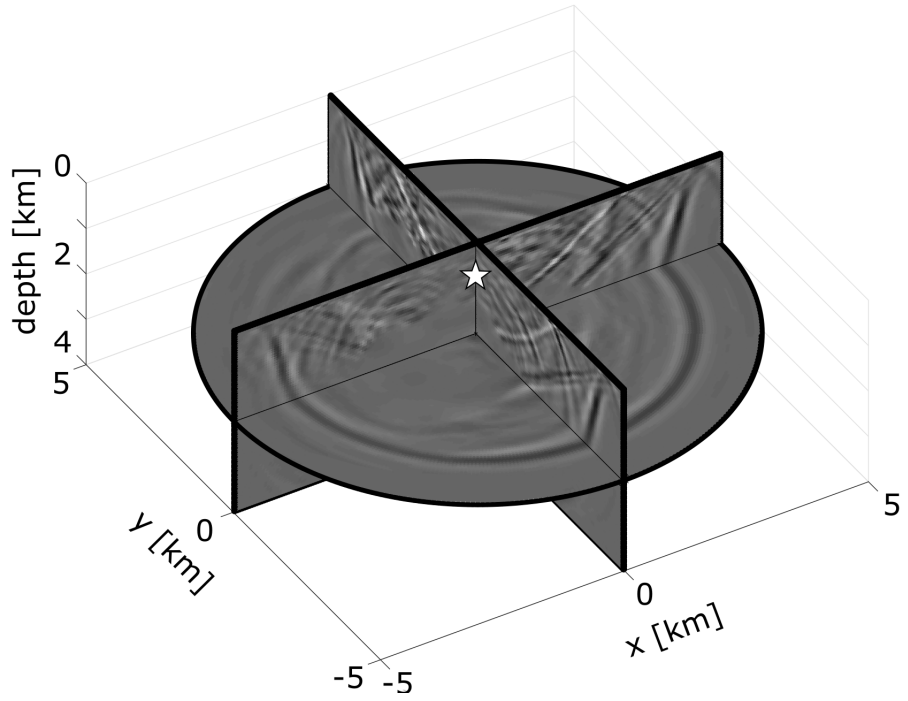



**Figure 3.2:** Cross sections of a wavefield (vertical displacement) snapshot in a complex medium (SEG/EAGE salt model) 2.4 s after an explosion was set off at the position marked by the star.

contraction. In the fluid, the acoustic wave equation is commonly expressed in terms of pressure $p(\mathbf{r}, t)$:

$$c^{-2}\ddot{p} - \nabla^2 p = \nabla^2 f_{\mathrm{f}}, \tag{3.3}$$

where $c$ is the acoustic wave speed and $f_{\mathrm{f}}$ is a pressure source term (typically an explosion or pressure caused by a directional force, such as from wind). This form has the advantage of being cheaper than displacement or velocity formulations, since pressure has only one degree of freedom. However, solid-fluid interactions at an interface between an acoustic and an elastic medium can be calculated much more elegantly using velocity. A way around this, which was first suggested by Everstine (1981) and which is now standard practice, is to use a formulation in terms of a fluid potential $\chi(\mathbf{r}, t)$. There are several valid fluid potential formulations, and we use a displacement potential introduced in Nissen-Meyer et al. (2008):

$$\kappa^{-1}\ddot{\chi} - \nabla \cdot (\rho^{-1}\nabla\chi) = \kappa^{-1} f_{\mathrm{f}}, \tag{3.4}$$

$$\mathbf{u} = \rho^{-1}\nabla\chi. \tag{3.5}$$

Fluid potentials have only one degree of freedom and, at the same time, the potential-based equations allow for an easy integration of solid-fluid interfaces. The formulation used here also holds for inhomogeneous fluids.

Waves travelling through these solid and fluid regions are coupled through the continuity of normal traction and normal displacement across the solid-fluid boundaries. Such coupling can be expressed as

$$\hat{\mathbf{n}} \cdot \mathbf{T} = \ddot{\chi}\hat{\mathbf{n}}, \quad \text{and} \tag{3.6}$$

$$\hat{\mathbf{n}} \cdot \mathbf{u} = \rho^{-1}\hat{\mathbf{n}} \cdot \nabla\chi, \tag{3.7}$$

respectively, where $\hat{\mathbf{n}}$ denotes the normal on the solid-fluid boundary.

Some numerical modelling methods, most notably the FDM, are based on the above equations which are also referred to as strong forms of the momentum equations. The SEM, on the other hand, works with the integral or weak forms of the momentum equations. These forms facilitate a more natural handling of

boundary conditions - one of the major attractions of both SEMs and FEMs. In order to derive the integral forms of Equations 3.1 and 3.4, they are multiplied by arbitrary time-independent test functions $\mathbf{w}(\mathbf{r})$ and $w(\mathbf{r})$, respectively, and then integrated by parts over volume. The volume where the velocity model is defined is denoted $\Omega$, with fluid and solid regions distinguished as $\Omega_\mathrm{f}$ and $\Omega_\mathrm{s}$, respectively. The weak forms of the elastic and acoustic wave equations are given by:

$$\int_{\Omega_\mathrm{s}} \rho\ddot{\mathbf{u}} \cdot \mathbf{w}\, \mathrm{d}^3\mathbf{r} + \int_{\Omega_\mathrm{s}} \nabla\mathbf{u} \colon \mathbf{C} \colon \nabla\mathbf{w}\, \mathrm{d}^3\mathbf{r} + \int_{\Gamma_\mathrm{sf}} \mathbf{w} \cdot \ddot{\chi}\hat{\mathbf{n}}\, \mathrm{d}^2\mathbf{r} = \int_{\Omega_\mathrm{s}} \mathbf{f}_\mathrm{s} \cdot \mathbf{w}\, \mathrm{d}^3\mathbf{r} \tag{3.8}$$

$$\int_{\Omega_\mathrm{f}} \kappa^{-1} w\ddot{\chi}\, \mathrm{d}^3\mathbf{r} + \int_{\Omega_\mathrm{f}} \rho^{-1}\nabla\chi \cdot \nabla w\, \mathrm{d}^3\mathbf{r} - \int_{\Gamma_\mathrm{sf}} w\mathbf{u} \cdot \hat{\mathbf{n}}\, \mathrm{d}^2\mathbf{r} = \int_{\Omega_\mathrm{f}} \kappa^{-1} w f_\mathrm{f}\, \mathrm{d}^3\mathbf{r}, \tag{3.9}$$

with $\hat{\mathbf{n}}$ denoting the upward-pointing normal vector on the solid-fluid boundary, $\Gamma_\mathrm{sf}$.

At this point, the ACA-SEM deviates from the classic SEM approach. Instead of subjecting the above weak forms to 3D Legendre spectral element discretisation (e.g. Komatitsch and Vilotte, 1998), we move to a source-centred cylindrical coordinate system and apply the spectral element discretisation only in 2D: along the in-plane coordinates, $s$ and $z$. Meanwhile, a spatial Fourier transform is applied to the azimuthal $\phi$-direction. As a result, the solver operates on a 2D mesh which can be pictured as a slice through the cylindrical velocity model (see Fig. 3.3). To be more precise, this mesh can be described as a 2D field where each point contains a Fourier series expansion of the wavefield as recorded in a circle around the axis. The advantage of this set-up is that it avoids meshing the azimuth and hence it does

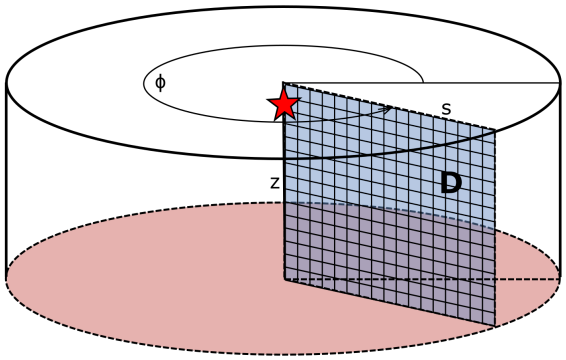



**Figure 3.3:** Sketch illustrating the geometry of the mesh used in an ACA method. The $s$ and $z$ coordinates are kept in the physical domain while along the $\phi$ coordinate the Fourier series expansions are considered. The star shows the source position. The modelling domain is labelled as $D$.

not impose a fixed resolution in this dimension. This gives us flexibility to adapt to any wavefield smoothness observed along $\phi$: the resolution can be decreased by simply truncating the Fourier series where the coefficients become negligibly small.

Let us start with the 2D spectral element discretisation. An issue that comes with using the weak form of the wave equation, is that fields like the displacement are in an integral expression and thus must be considered globally. This means that, without a workaround, a huge system of equations must be solved at every time step. The answer to this problem is the use of local shape functions which allow the modelling domain $D$ to be broken up into smaller segments, $D_e$, which we call elements (see Fig. 2.2). The first step to introducing our shape functions is to map the coordinates of a given 2D element, $\mathbf{x} = (s, z)$, onto a reference coordinate frame, $\xi = (\xi, \eta)$, where $-1 \leq \xi \leq 1$ and $-1 \leq \eta \leq 1$. In this new coordinate frame, the wavefield is expressed in terms of a polynomial expansion using high-order Lagrange polynomials, $l_j(\xi)$. The subscript $j = \{0..n_l\}$ is the degree of the Lagrange polynomial and $n_l$ is the order of the set of polynomials (typically 4 in AxiSEM3D). These polynomials are defined such that for $n_l + 1$ pre-determined points in the interval between -1 and 1 one polynomial of the set takes the value 1 and all other polynomials take the value 0 (see Fig. 2.3). The points where this happens are called Gauss-Lobatto-Legendre (GLL) points. Outside the element, the polynomials are defined to be 0. The spatial discretisation of the wavefield inside the element is performed along the GLL points:

$$\mathbf{u}(\mathbf{x}(\xi, \eta), \phi, t)|_{\mathbf{x}\in D_e} \approx \sum_{j=0}^{n_l} \sum_{k=0}^{n_l} \mathbf{u}^{j,k}(\phi, t) l_j(\xi) l_k(\eta), \tag{3.10}$$

where $\mathbf{u}^{j,k}(\phi, t) = \mathbf{u}(\mathbf{x}(\xi_j, \eta_k), \phi, t)|_{\mathbf{x}\in D_e}$ and $\xi_j$ and $\eta_k$ are positions of GLL points. Using our set of Lagrange polynomials in place of the test functions, $w$ and $\mathbf{w}$, seen in the weak forms, we are now able to perform cheaper local integration. Moreover, weight functions and spatial derivatives of the Lagrange polynomials can be pre-computed analytically to make both spatial integration and differentiation cheaper. An important property of the SEM is that the mass matrix becomes exactly diagonal. To obtain the acceleration, $\ddot{\mathbf{u}}$, the stiffness must be divided by

density (i.e. multiplied by the inverse mass matrix) at every time-step. If the mass matrix is diagonal the inverse mass matrix is also diagonal and thus contains fewer coefficients. This reduces the number of calculations required and, further, a diagonal inverse mass matrix makes these calculations easily parallelisable.

Let us now consider the azimuthal dimension and derive the partially pseudo-spectral weak form of the elastic wave equation. To not complicate the equations further, I perform the following steps on the displacement, $\mathbf{u}(\mathbf{r}, t)$, in its original form without the spectral element discretisation. The displacement is approximated in terms of a Fourier series along $\phi$ which is truncated after $n_u$ coefficients. Using the unit cylindrical basis $(\hat{\mathbf{s}}, \hat{\phi}, \hat{\mathbf{z}}) = (\hat{\mathbf{x}}\cos\phi + \hat{\mathbf{y}}\sin\phi, -\hat{\mathbf{x}}\sin\phi + \hat{\mathbf{y}}\cos\phi, \hat{\mathbf{z}})$, expressed for the Einstein summation convention as $(\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \hat{\mathbf{g}}_3)$, one may write

$$\begin{aligned}\mathbf{u} &= \hat{\mathbf{g}}_i(\phi) u_i(s, \phi, z; t) \\ &= \hat{\mathbf{g}}_i(\phi) \sum_{\alpha=0}^{\infty} u_i^\alpha(s, z; t) e^{\mathbb{i}\alpha\phi} \\ &\approx \hat{\mathbf{g}}_i(\phi) \sum_{\alpha=0}^{n_u} u_i^\alpha(s, z; t) e^{\mathbb{i}\alpha\phi}.\end{aligned} \tag{3.11}$$

$\mathbb{i}$ stands for the imaginary unit and $u_i^\alpha(s, z; t)$ is a 2D field containing the $\alpha$-th Fourier coefficients of $u_i(s, \phi, z; t)$. These coefficient fields are defined in $D$, which can be visualised as a 2D slice through the 3D modelling domain $\Omega$ (Fig. 3.3). If the test functions, material properties and source terms are reformulated in an equivalent manner, Equation 3.8 can be rewritten as

$$\begin{aligned}\sum_{|\alpha|\leq n_u} \Big( \int_D \rho^{-(\alpha+\beta)} \ddot{u}_i^\alpha w_i^\beta s \,\mathrm{d}s\mathrm{d}z + \int_D u_{i;j}^\alpha C_{ijkl}^{-(\alpha+\beta)} w_{k;l}^\beta s \,\mathrm{d}s\mathrm{d}z \\ + \int_{L_{\mathrm{sf}}} \ddot{\chi}^\alpha n_i^{-(\alpha+\beta)} w_i^\beta s \,\mathrm{d}s \Big) = \int_D f_{s,i}^{-\beta} w_i^\beta s \,\mathrm{d}s\mathrm{d}z, \quad \forall \beta \in \{-n_u, ..., n_u\},\end{aligned} \tag{3.12}$$

where $L_{\mathrm{sf}}$ is a line where the solid-fluid interface intersects $D$. $u_{i;j}^\alpha$ is a shorthand for the Fourier coefficients of the deformation gradient $\nabla\mathbf{u}$. In full index notation, these coefficients are expressed as

$$u_{i;j}^\alpha = \partial_{x_i} u_j^\alpha + \frac{1}{s}\delta_{i2}(\mathbb{i}\alpha\delta_{jk} - \epsilon_{3jk}) u_k^\alpha, \tag{3.13}$$

with $\delta_{ij}$ being the Kronecker delta and $\epsilon_{ijk}$ the Levi-Civita symbol.

Equation 3.9 may be reformulated in the same way:

$$\int_{D_\mathrm{f}} \kappa^{-1} \ddot{\chi}^{-\beta} w^\beta s \,\mathrm{d}s\mathrm{d}z + \int_{D_\mathrm{f}} \rho^{-1} \left( \chi_{;i}^{-\beta} w_{;i}^\beta + \beta^2 \chi^{-\beta} w^\beta s \right) \mathrm{d}s\mathrm{d}z$$
$$- \sum_{|\alpha| \leq n_u} \left( \int_{L_\mathrm{sf}} u_i^\alpha n_i^{-(\alpha+\beta)} w^\beta s \,\mathrm{d}s \right) = \int_{D_\mathrm{f}} \kappa^{-1} f_\mathrm{f}^{-\beta} w^\beta s \,\mathrm{d}s\mathrm{d}z, \quad \forall \beta \in \{-n_u, ..., n_u\}. \tag{3.14}$$

The system of coupled equations is now expressed as $n_u + 1$ two-dimensional weak forms which can be solved using a cheaper 2D SEM rather than a full 3D method. The number of equations to be solved is thereby determined by the azimuthal complexity of the 3D input model and the ensuing complexity of the wavefield.

An important factor to consider in terms of azimuthal complexity is the existence of non-isotropic seismic sources. This is taken care of by applying a source decomposition available to on-axis sources in spherical/cylindrical coordinate systems, which was introduced by Nissen-Meyer et al. (2007a): Any responses of a homogeneous medium to a point moment tensor can be represented by a superposition of a monopole, a dipole and a quadrupole source. This means that a moment tensor source can be expressed exactly using a Fourier expansion order of $|\beta| \leq 2$. In the same way, a point force can be expressed using only two terms, i.e. $|\beta| \leq 1$. Hence, the source term adds little to the complexity of the problem.

## 3.4 Wavefield Learning - Creating Azimuthal Complexity Maps

The choice of an appropriate Fourier expansion order $n_u$ is critical for the efficiency and accuracy of the ACA method. As shown by Leng et al. (2016), the solver converges for sufficiently large $n_u$. This property is used by a technique called wavefield scanning. Introduced by Leng et al. (2019), wavefield scanning utilises a test run with a large starting $n_u$ and determines the smallest $n_u$ where the error in the truncated Fourier series approximation of the wavefield is still negligibly small. This is done independently for each point on the mesh creating a 2D map of the required resolution. The "learned" $n_u$-map, which I refer to as a complexity map, is then saved to a file which can be used in future runs with the same or similar models.

The error tolerance of the wavefield learning algorithm can be specified by the user. I use a default value of 0.1%. This condition is applied to the displacement potential in fluid media, and separately to the radial, tangential and vertical components of the displacement in solid media. Further, it is necessary to carefully choose the parts of the time series where the learning algorithm is applied, such as to avoid fitting numerical noise. For this purpose, the H2-norm over the Fourier coefficients is used as a measure for amplitude in the 2D mesh. Good results can be achieved by ensuring a proper resolution of only the largest amplitude of the time series. This makes use of the fact that, usually, the highest amplitude phases are also the phases with the smallest wavelength and thus with the highest sensitivity to 3D complexity. To be specific, single peak wavefield scanning tends to resolve surface waves or, if these are not present, first arrival S-waves or, if these are not present either, first arrival P-waves. In many cases, this produces a complexity map which properly resolves first arrivals of both S- and P-waves as well as any topography. However, in models where multiple off-plane reflections form a significant part of the wavefield or where the first arrivals do not illuminate all the present structures, more peaks in the time series need to be considered. For this purpose, I expanded the wavefield scanning algorithm of Leng et al. (2019). The new algorithm sorts peaks in the time series by amplitude and outputs the minimum value for $n_u$ required to resolve the largest $k$ peaks, where $k$ can be specified by the user. A detailed description of this process is given in Chapter 6.

## 3.5 Method Architecture

### 3.5.1 Efficient Time-Stepping

The ACA-SEM largely operates in partial Fourier space where the wavefield can be represented by half as many coefficients and gradient operations are simplified and cheaper. However, interactions with the 3D medium must be calculated in physical space. These considerations lead to the complex method architecture illustrated in Figure 3.4. As can be seen in this schematic, 3D model interactions in physical space occur twice per time step: 1. when the stress tensor is calculated from strain

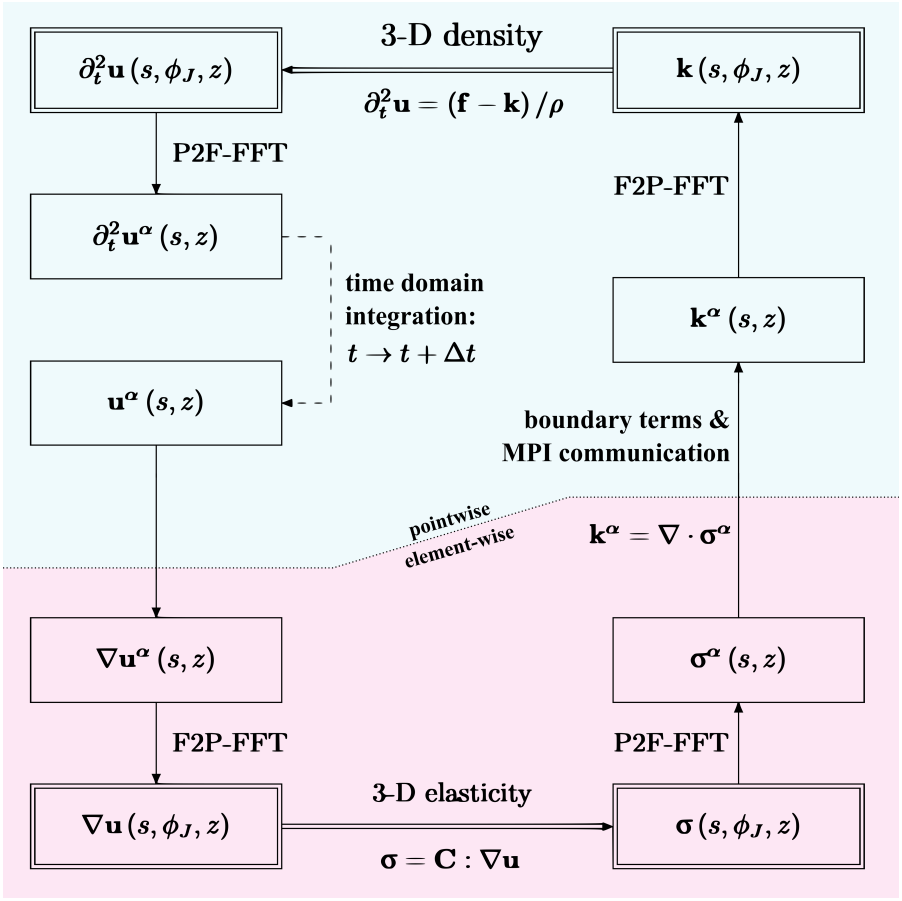



**Figure 3.4:** Workflow of one time step in AxiSEM3D (modified from Leng, 2017). $\sigma$ and $\mathbf{k}$ denote the stress tensor and the stiffness force respectively. P2F-FFT stands for the Fourier transform from physical space to partial Fourier space (only the azimuthal dimension is transformed) and F2P-FFT stands for its inverse.

and the stiffness of the medium, and 2. when acceleration is calculated from the stiffness force and the density of the medium. There may be more interactions in physical space, such as along non-horizontal solid-fluid boundaries, but these are usually limited to a small fraction of the points in the simulation.

The Fourier transforms are the most computationally expensive part of an ACA simulation. AxiSEM3D uses the FFTW3 library which handles the Fourier transforms with unparalleled efficiency. Nevertheless, the additional cost arising from the four Fourier transforms required per time step currently outweighs the cost savings from operating in partial Fourier space. This means that it is essential for the efficient use of AxiSEM3D to make proper use of the ACA approach to optimise the azimuthal resolution.

### 3.5.2 3D Models and Particle-Relabelling

ACA methods by design only require a 2D input mesh. This avoids the challenges associated with hexahedral meshing which are seen as a hurdle when it comes to the applicability of spectral element methods (discussed e.g. in Peter et al., 2011). At the same time, this limits the ways in which 3D models can be implemented. In AxiSEM3D, 3D models are divided into "volumetric" (spatially varying medium properties in a regular mesh) and "geometric" (mesh distortions) models which can be combined to approximate a desired structure.

First, the solver is set up by a 1D background velocity model and a quadrilateral

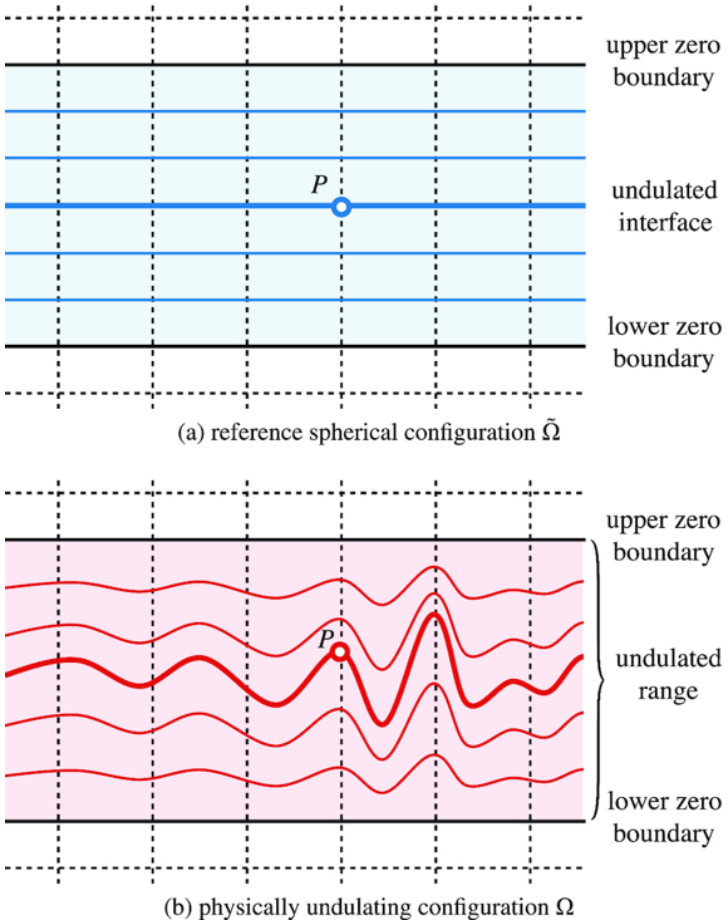



**Figure 3.5:** An illustration of geometric mapping from a selected reference spherical configuration to the physically undulating configuration (from Leng et al., 2019). The undulated range can be freely chosen by the user, who should consider the trade-off between the size of the 3D model (all points in range require larger $n_u$) and the smallest effective element size which will determine the time step based on the Courant stability criterion.

2D mesh defining the modelling domain $D$. When adding a volumetric model, the 3D medium properties are sampled around the azimuth at a rate determined by the chosen Fourier expansion order, $n_u$, and superimposed on the background model. 3D anisotropy is fully implemented (a use case is presented by Tesoniero et al., 2020), as is intrinsic attenuation (Leng et al., 2016).

Geometric model support was added by Leng et al. (2019) who implemented particle-relabelling. This method was suggested by Al-Attar and Crawford (2016) and extended to include solid-fluid interactions in Al-Attar et al. (2018). With this addition, discontinuities can be treated rigorously such as to avoid the stepping artefacts common in methods which implement models purely volumetrically.

The principle of particle-relabelling is that, given any undulating physical model (volume denoted as $\Omega$), it is possible, (by appropriately changing density and elastic/acoustic properties) to create a flat-layered anisotropic reference model (volume denoted as $\tilde{\Omega}$) which is identical to the physical model in terms of wave propagation. That is, the weak forms in Eqs. 3.8 and 3.9 are equivalent to the weak forms in the reference volume:

$$\begin{aligned}\int_{\tilde{\Omega}_\mathrm{s}} \tilde{\rho}\ddot{\tilde{\mathbf{u}}} \cdot \tilde{\mathbf{w}} \,\mathrm{d}^3\mathbf{r} + \int_{\tilde{\Omega}_\mathrm{s}} \nabla\tilde{\mathbf{u}} \colon \tilde{\mathbf{C}} \colon \nabla\tilde{\mathbf{w}} \,\mathrm{d}^3\mathbf{r} \\ + \int_{\tilde{\Gamma}_\mathrm{sf}} \tilde{\mathbf{w}} \cdot \ddot{\tilde{\chi}}\tilde{\mathbf{n}} \,\mathrm{d}^2\mathbf{r} = \int_{\tilde{\Omega}_\mathrm{s}} \tilde{\mathbf{f}}_\mathrm{s} \cdot \tilde{\mathbf{w}} \,\mathrm{d}^3\mathbf{r}, \text{and}\end{aligned} \tag{3.15}$$

$$\begin{aligned}\int_{\tilde{\Omega}_\mathrm{f}} \tilde{\kappa}^{-1}\tilde{w}\ddot{\tilde{\chi}} \,\mathrm{d}^3\mathbf{r} + \int_{\tilde{\Omega}_\mathrm{f}} \tilde{\rho}^{-1}\nabla\tilde{\chi} \cdot \tilde{\tilde{I}} \cdot \nabla\tilde{w} \,\mathrm{d}^3\mathbf{r} \\ - \int_{\tilde{\Gamma}_\mathrm{sf}} \tilde{w}\tilde{\mathbf{u}} \cdot \tilde{\mathbf{n}} \,\mathrm{d}^2\mathbf{r} = \int_{\tilde{\Omega}_\mathrm{f}} \tilde{\kappa}^{-1}\tilde{w}\tilde{f}_\mathrm{f} \,\mathrm{d}^3\mathbf{r}.\end{aligned} \tag{3.16}$$

To derive the medium properties of the reference volume, Leng et al. (2019) introduce a mapping, $\xi : \tilde{\Omega} \to \Omega$, for spherical Earth models which is restricted to the radial direction. For the purpose of local-scale models, we instead restrict it to the vertical direction:

$$\xi(\mathbf{r}) = \mathbf{z} + \tau(\mathbf{r})\,\hat{\mathbf{z}}. \tag{3.17}$$

This mapping effectively stretches a point in the reference model to its corresponding position in the original, undulating model. With this as a starting point, the undulating model can be treated as a kinematic deformation state of the

reference model. This principle is schematically illustrated in Figure 3.5. For solid media, the weak forms in the reference volume (Eq. 3.15) can be related to the weak form in the physical volume (Eq. 3.8) using the following set of equations:

$$\tilde{\mathbf{u}}(\mathbf{r},t) = \mathbf{u}(\xi(\mathbf{r}),t) \tag{3.18}$$

$$\tilde{\mathbf{w}}(\mathbf{r},t) = \mathbf{w}(\xi(\mathbf{r}),t) \tag{3.19}$$

$$\tilde{\mathbf{C}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau,\mathbf{r}) \cdot \mathbf{C}(\xi(\mathbf{r})) \cdot \mathbf{J}^{-T}(\tau,\mathbf{r}) \, |\mathbf{J}(\tau,\mathbf{r})| \tag{3.20}$$

$$\tilde{\mathbf{n}}(\mathbf{r}) = \mathbf{J}^{-\mathrm{T}}(\tau,\mathbf{r}) \cdot \hat{\mathbf{n}}(\xi(\mathbf{r})) \, |\mathbf{J}(\tau,\mathbf{r})| \tag{3.21}$$

$$\tilde{\rho} = \rho(\xi(\mathbf{r}),t) \, |\mathbf{J}(\tau,\mathbf{r})|, \tag{3.22}$$

where $|\mathbf{J}(\tau,\mathbf{r})|$ denotes the Jacobian determinant.

For fluids, the changes in medium properties are slightly different, as given by the equations:

$$\tilde{\chi}(\mathbf{r},t) = \chi(\xi(\mathbf{r}),t) \tag{3.23}$$

$$\tilde{w}(\mathbf{r},t) = w(\xi(\mathbf{r}),t) \tag{3.24}$$

$$\tilde{\mathbf{I}} = \mathbf{J}^{-1}(\tau,\mathbf{r}) \cdot \mathbf{J}^{T}(\tau,\mathbf{r}) \tag{3.25}$$

$$\tilde{\mathbf{n}}(\mathbf{r}) = \mathbf{J}^{-\mathrm{T}}(\tau,\mathbf{r}) \cdot \hat{\mathbf{n}}(\xi(\mathbf{r})) \, |\mathbf{J}(\tau,\mathbf{r})| \tag{3.26}$$

$$\tilde{\kappa} = \kappa(\xi(\mathbf{r}),t) \, |\mathbf{J}(\tau,\mathbf{r})|^{-1}, \tag{3.27}$$

$$\tilde{\rho} = \rho(\xi(\mathbf{r}),t) \, |\mathbf{J}(\tau,\mathbf{r})|^{-1}. \tag{3.28}$$

To obtain an expression for the Jacobian, we start with the mapping in Equation 3.17, from where the deformation gradient in cylindrical coordinates can be derived as:

$$\mathbf{F}(\tau,\mathbf{r}) = [\nabla(\tau\hat{\mathbf{z}})]^T = \begin{pmatrix} \hat{\mathbf{z}} & \hat{\mathbf{s}} & \hat{\phi} \end{pmatrix} \begin{bmatrix} \partial_z\tau & \partial_s\tau & \frac{\partial_\phi\tau}{s} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{pmatrix} \hat{\mathbf{z}} \\ \hat{\mathbf{s}} \\ \hat{\phi} \end{pmatrix} \tag{3.29}$$

The Jacobian of this deformation is given by

$$\mathbf{J}(\tau,\mathbf{r}) = \mathbf{I} + \mathbf{F}(\tau,\mathbf{r}), \tag{3.30}$$

and its inverse, can be explicitly written as

$$\mathbf{J}^{-1}(\tau,\mathbf{r}) = \frac{1}{1+\partial_z\tau}\begin{pmatrix}\hat{\mathbf{z}} & \hat{\mathbf{s}} & \hat{\phi}\end{pmatrix}\begin{bmatrix}1 & -\partial_s\tau & -\frac{\partial_\phi\tau}{s}\\ 0 & 1+\partial_z\tau & 0\\ 0 & 0 & 1+\partial_z\tau\end{bmatrix}\begin{pmatrix}\hat{\mathbf{z}}\\ \hat{\mathbf{s}}\\ \hat{\phi}\end{pmatrix}. \tag{3.31}$$

Note that the coordinate vectors are commuted if compared to the standard expressions used in the other equations.

There are some limitations to the types of geometries that can be implemented with particle relabelling. Most notably, the relabelled interface must be continuous. Hence, we cannot accurately model discontinuities which reach the surface or two discontinuities which merge to form a wedge between them.

*When I'm working on a problem, I never think about beauty. I think only how to solve the problem. But when I have finished, if the solution is not beautiful, I know it is wrong.*

— Freeman Dyson

# 4 Absorbing Boundaries

## Contents


**4.1 Introduction . . . . . 32**
**4.2 The Mixed Absorbing Boundary Method . . . . . 33**
4.2.1 Absorbing Boundary Condition . . . . . 35
4.2.2 Sponge Boundary . . . . . 35
4.2.3 A Note on Stability . . . . . 36
**4.3 Optimising the Sponge Parameters for the Mixed Boundary . . . . . 36**
4.3.1 Test Set-Up . . . . . 37
4.3.2 Results for Elastic Media . . . . . 39
4.3.3 Results for Acoustic Media . . . . . 41
4.3.4 How does the Error Decrease with Sponge Width? . . . 44
**4.4 Discussion . . . . . 46**


## 4.1 Introduction

At the start of this project, AxiSEM3D was predominantly a solver for global seismic simulations. Local-scale simulations with Cartesian meshes were possible, but such meshes were treated as a special case of spherical input meshes with a small lateral extent. This would induce curvature to the effective mesh, even if the input mesh had completely flat boundaries. In simple set-ups, this approximation would work. But when implementing more complex models with discontinuities,

the resulting errors in locating the model would cause simulations to fail. Many small adaptations to the code were necessary to enable the local-scale simulations shown in the following chapters. Most of these changes are trivial in theory, and are thus not documented more extensively, but they do bear mentioning. At the time of writing, AxiSEM3D treats Cartesian, curved and spherical meshes rigorously with the ability to process and even mix Cartesian, source-centred spherical and geographic user inputs. In this chapter, I elaborate on one code adaptation which does warrant a more in-depth explanation: the addition of absorbing boundaries.

In local-scale simulations, the modelling domain is truncated. The artificial edges of the truncated domain reflect the wavefield in the same way a free surface does. Handling these reflections has been a long-standing challenge in most numerical modelling applications and numerous methods have been devised for this purpose. For AxiSEM3D, a combination of the sponge layers by Kosloff and Kosloff (1986) and the 1$^{st}$ order absorbing boundary condition by Clayton and Engquist (1977) is chosen. Further, an empirical equation to optimise the input parameters of the sponge layers is determined. The theoretical basis for the absorbing boundary methods as well as the process for the automated parametrisation are described in the following sections.

## 4.2 The Mixed Absorbing Boundary Method

When it comes to the choice of an absorbing boundary method, there are many options. A comprehensive overview of absorbing boundary methods for acoustic modelling can be found in Gao et al. (2017). In short, absorbing boundaries can be categorised into absorbing layers or sponge boundaries, where the model is extended by an absorbing region within which waves are gradually attenuated (e.g. Cerjan et al., 1985; Kosloff and Kosloff, 1986; Berenger, 1994), and absorbing boundary conditions (ABCs), where a wavefield-cancelling boundary term is applied at the edge of the model (e.g. Clayton and Engquist, 1977; Reynolds, 1978; Higdon, 1986).

Perfectly matched layer methods (PMLs), a variant of the sponge layers, are currently considered the state-of-the-art of absorbing boundaries. However, PMLs often require significant restructuring of the underlying numerical modelling method

(e.g. Festa and Vilotte, 2005), and they often cause problems with stability. Gao et al. (2017) point out that, in their acoustic example, a hybrid method which combines a sponge boundary with ABCs (taken from Liu and Sen, 2010) performs well at removing unwanted reflections, and is on-par with PMLs. Given the inherent complexity of the ACA approach, it seems prudent to avoid extensive restructuring and to choose the hybrid approach instead. A sketch of such a hybrid set-up is shown in Figure 4.1.

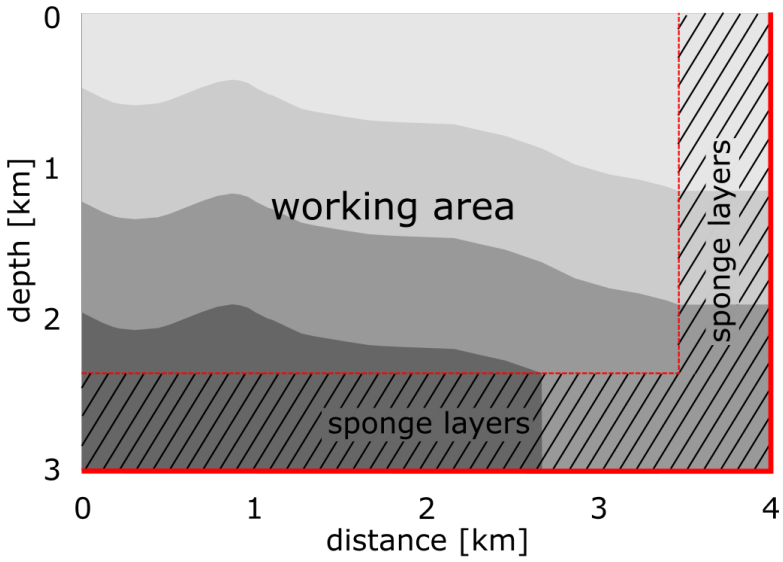



**Figure 4.1:** Sketch of the mixed absorbing boundary method in AxiSEM3D. The dashed line shows the edge of the working area. The hatched area highlights the sponge region. The bold line shows the position of elements where, additionally to the attenuation by the sponge, the ABC is applied.

I adapt an ABC introduced by Komatitsch and Vilotte (1998) based on the work of Clayton and Engquist (1977) which is specifically formulated for SEMs. This method performs best for P-waves at normal incidence. Grazing waves and S-waves, on the other hand, are better handled by a sponge region with gradually increasing attenuation. Such a method becomes more effective the longer waves travel through the attenuating region. For this purpose, I implement a sponge boundary method developed by Kosloff and Kosloff (1986).

### 4.2.1 Absorbing Boundary Condition

The absorbing boundary condition used by Komatitsch and Vilotte (1998) requires the addition of a wave-cancelling boundary term, ABT, to Equations 3.8 and 3.9 at the points along the truncated edge of the model. In the solid domain, the absorbing model boundary is denoted as $\Gamma^{\mathrm{s}}_{\mathrm{abs}}$. The boundary term can be calculated as

$$\mathrm{ABT_s} = -\int_{\Gamma^{\mathrm{s}}_{\mathrm{abs}}} \mathbf{w} \cdot \mathbf{t}\, d\Gamma, \tag{4.1}$$

where $\mathbf{t}$ is the traction which relates to velocity, $\dot{\mathbf{u}}$, as

$$\mathbf{t} = c_p \rho \dot{\mathbf{u}} \cdot \hat{\mathbf{n}}\hat{\mathbf{n}} + c_s \rho (\dot{\mathbf{u}} - \dot{\mathbf{u}} \cdot \hat{\mathbf{n}}\hat{\mathbf{n}}). \tag{4.2}$$

$c_p$ and $c_s$ represent P- and S-wave speed, respectively, and $\hat{\mathbf{n}}$ the normal of the absorbing boundary. Komatitsch et al. (2000) extend their ABCs to incorporate acoustic media. In this case, the boundary term can be calculated as

$$\mathrm{ABT_f} = \int_{\Gamma^{\mathrm{f}}_{\mathrm{abs}}} (\rho c)^{-1} w \dot{\phi}\, d\Gamma, \tag{4.3}$$

where $\Gamma^{\mathrm{f}}_{\mathrm{abs}}$ denotes the absorbing boundary in the fluid medium.

### 4.2.2 Sponge Boundary

Kosloff and Kosloff (1986) suggest an absorbing boundary method based on gradually increasing damping of the wavefield within an extended boundary (sponge) region. The damping effect is achieved by applying an absorption factor $\gamma$ to the acceleration $\ddot{\mathbf{u}}$ in Equation 3.8 (or equally to $\ddot{\chi}$ in Equation 3.9) as follows:

$$\ddot{\mathbf{u}}_{\mathrm{abs}} = \ddot{\mathbf{u}} - 2\gamma\dot{\mathbf{u}} - \gamma^2 \mathbf{u}. \tag{4.4}$$

The absorption factor, $\gamma$, is zero within the working area of the model and then smoothly increases across an extension of the model (between dashed and bold red line in Fig. 4.1) up to some maximum value $U_0$:

$$\gamma = \begin{cases} U_0 \cos^2\left(\frac{\pi d}{2w}\right), & \text{if } d < w \\ 0, & \text{otherwise,} \end{cases} \tag{4.5}$$

where $d$ denotes distance from the edge of the mesh and $w$ is the width of the absorbing boundary region. A challenge posed by this method is to find appropriate values for the width $w$ and the maximum absorption factor $U_0$.

### 4.2.3 A Note on Stability

Both of the above absorbing boundary methods work by introducing velocity terms to the wave equations. Such velocity terms are equivalent to viscous damping. As explained in Hughes (1987), viscous damping terms affect the stability of the Newmark scheme, and this can be counteracted by adding a correction term to the mass matrix, as

$$\mathbf{M}_{\text{corr}} = \mathbf{M} + \frac{1}{2}\Delta t \mathbf{C}, \tag{4.6}$$

where $\mathbf{C}$ is the damping matrix. This correction is currently not implemented in AxiSEM3D, but it is a potential avenue to explore if the performance of the absorbing boundaries in their current state proves insufficient.

## 4.3 Optimising the Sponge Parameters for the Mixed Boundary

In the introduction of their sponge boundary, Kosloff and Kosloff (1986) use a 1D acoustic example and they do not elaborate on the effect of changing $w$ and $U_0$. But for the sponge boundary to serve as a practical addition to AxiSEM3D, the relationship between these input parameters and the absorption performance must be known. The effect of the sponge boundary is non-physical and thus there is no way to predict these relationships other than by empirical systematic testing. Ideally, I seek 1. a rule of thumb for users to choose $w$ depending on their desired accuracy, and 2. an empirical equation which automates the choice of $U_0$ for optimal absorption.

Preliminary tests show that the optimal choice of $U_0$ depends on a number of input parameters: the local wave speeds of P- and S-waves, $v_p$ and $v_s$, the source period, $T$, source type (i.e., explosion, point force, rupture, etc.) and the boundary width, $w$. On the other hand, the optimum for $U_0$ is independent of the local density and the element size in the input mesh. Testing the dependency of $U_0$ across the parameter space spanned by the four relevant parameters would require a prohibitive number of simulations. However, it is possible to reduce this

parameter space based on equivalence: in the absence of attenuation and other structures in the model, a simulation with given values for $v_p$, $v_s$ and $w$ should be numerically indistinguishable from a simulation where all of these values and the mesh size are doubled. The same is true for a simulation where $T$ and $w$, the mesh size and the simulation time are doubled. So, it is really the ratios of these values that affect the optimal $U_0$. Specifically, I identify the dependency of the optimal value for $U_0$ on the sponge width relative to the S- and P-wavelengths, given by $n_p = \frac{w}{v_p T}$ and $n_s = \frac{w}{v_s T}$, respectively.

Figure 4.2 compares the seismic traces from a receiver gather with different absorbing boundary configurations for a homogeneous solid model and a homogeneous fluid model. This illustrates the necessity of including the sponge boundary if a high absorption performance is required - particularly for elastic models.

### 4.3.1 Test Set-Up

For the following tests, homogeneous models are implemented with a Cartesian mesh. The mesh size is 25 km x 25 km and the simulation time is adapted to the model such as to ensure the slowest phase (ground roll for the elastic case and P-wave for the acoustic wave, respectively) can travel at least 50 km. The reference solution is obtained by using a much larger mesh where no boundary reflections occur within the simulated time frame. A 2 Hz explosive source is placed at 1 km depth and the receivers are located on a line along the surface. For solid models and fluid models displacement and pressure are recorded respectively. The misfit is calculated as the root-mean-square (RMS) difference between the simulated traces and the reference traces for 18 evenly spaced receivers placed between 0.5 km and 9 km distance from the axis. To put the resulting values into perspective, they are expressed as a fraction of the misfit between the reference trace and a trace obtained

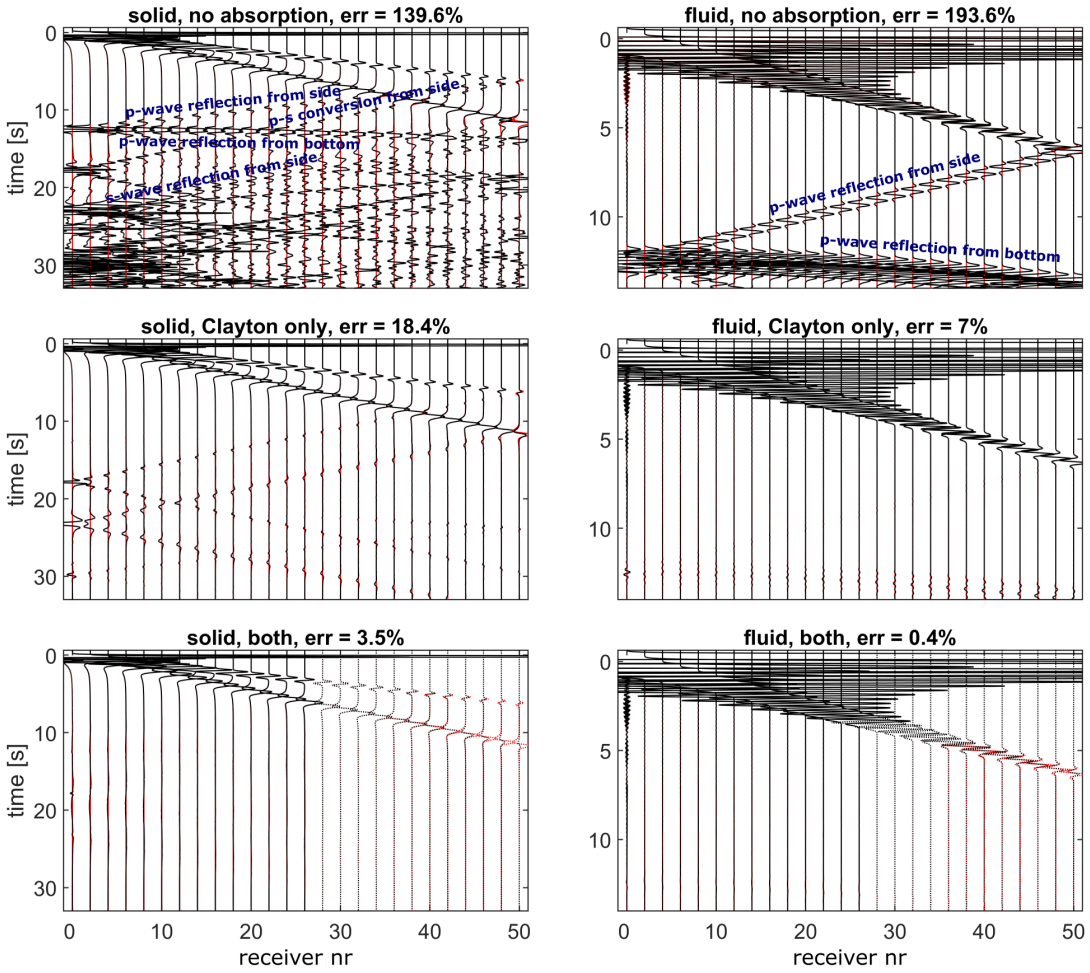



**Figure 4.2:** Example traces showing the effect of different absorbing boundary configurations in homogeneous solid and fluid media. In the given example the model parameters are $v_p = 4.0$ km/s, $v_s = 2.35$ km/s (solid only), $\rho = 2500$ kg/m$^3$ and $U_0 = 0.09$. The traces are recoded at receivers spaced regularly in 0.5 km intervals at the surface of a 25 km wide mesh. Reference traces without boundary reflections are shown in red. Dotted lines indicate that the receivers are located within the sponge region, so a mismatch between these traces and the reference is expected. The percentage error is the average RMS error relative to the RMS amplitude of the traces and does not include the receivers inside the sponge region. Some phases were labelled to help with interpretation.

by using only the Clayton-Engquist ABC without the additional sponge region:

$$\epsilon_r = \sqrt{\frac{1}{n_t}\sum_{i=0}^{n_t} \left\|\mathbf{u}(t_i) - \mathbf{u}_{\text{ref}}(t_i)\right\|^2}, \tag{4.7}$$

$$\epsilon = \sqrt{\frac{1}{n_r}\sum_{r=0}^{n_r} \epsilon_r^2}, \tag{4.8}$$

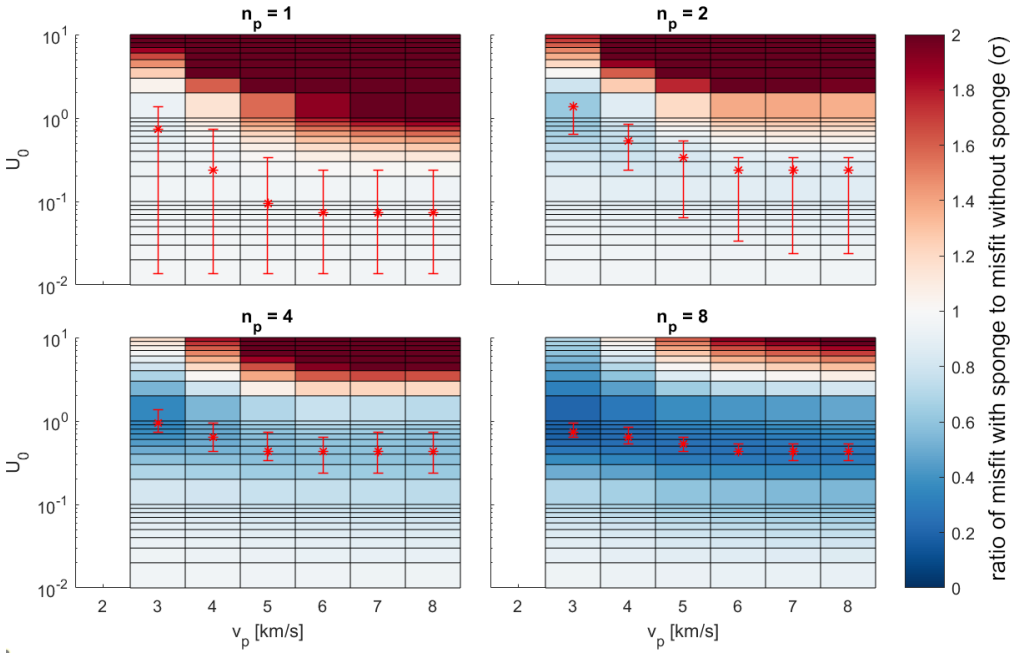
$$\sigma = \frac{\epsilon_{\text{with sponge}}}{\epsilon_{\text{without sponge}}}. \tag{4.9}$$

Here, $n_t$ denotes the number of samples in time for a given trace, and $n_r$ denotes the number of receivers considered for the misfit, respectively. $\epsilon_r$ is the error of a single trace at receiver number $r$, $\epsilon$ is the total RMS error of a given model with given absorbing boundary parameters and $\sigma$ is the error relative to the RMS error of the same model where no sponge region was used. Thus, a value of $\sigma = 1$ means that there is no improvement over using the ABC on its own, larger values mean the absorption with the given $U_0$ is worse, and smaller values mean the absorption with the given $U_0$ is better.

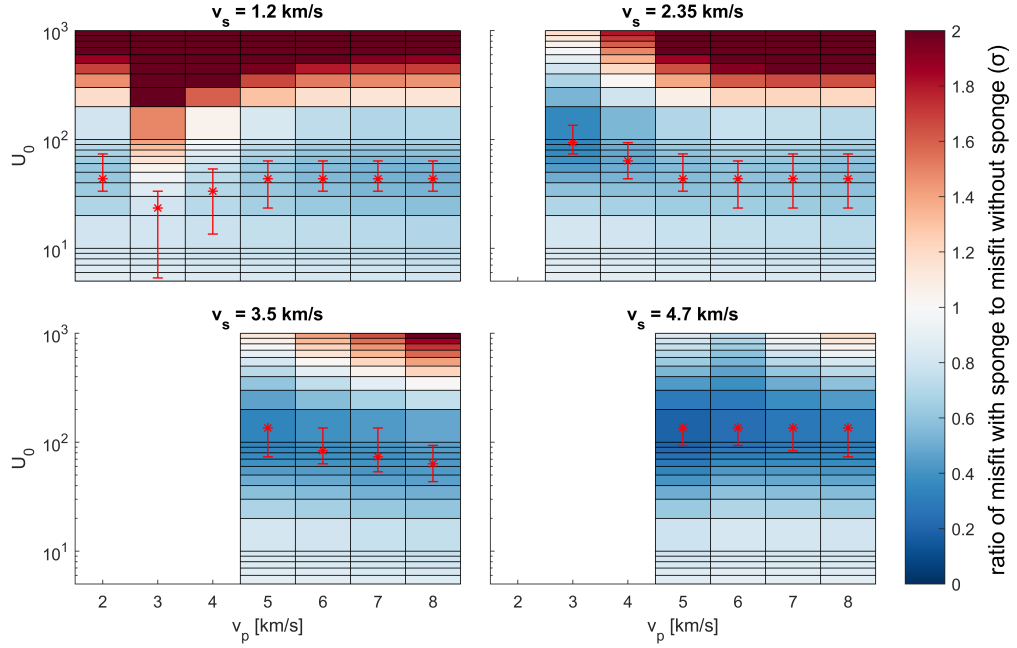
For a given model, 28 values for $U_0$ are tested: 0.01 to 0.1 in intervals of 0.01, 0.1 to 1 in intervals of 0.1 and 1 to 10 in intervals if 1. In each test series, $v_p$ is systematically increased to 8 km/s in intervals of 1 km/s. The boundary width is also increased accordingly such as to keep $n_p$ constant.

### 4.3.2 Results for Elastic Media

Figure 4.3(a) shows four test series for elastic models with a fixed $v_s$ of 2.35 km/s and with a different $n_p$ in each series. In Figure 4.3(b), $n_p$ is fixed and $v_s$ is changed in each series. It appears that for elastic models the ideal value for $U_0$ varies smoothly with both $v_p$ and the width of the sponge boundary. Further, the panels indicate that the benefit of the sponge boundary only becomes significant at boundary sizes larger than twice the P-wavelength. The misfit improves with increasing S-wave velocity. This is counter-intuitive since the number of S-wavelengths that fit into the boundary width, $n_s = \frac{w}{v_s T}$, decreases with $v_s$. The behaviour indicates that the $v_p/v_s$ ratio plays an important role when it comes to the absorption performance of the mixed boundary.




**(a)** Changing $n_p$ between series, using a constant $v_s = 2.35$ km/s.




**(b)** Changing $v_s$ between series, using a constant $n_p = 4$.

**Figure 4.3:** Results from the systematic tests of the performance of the mixed absorbing boundary using different values for $U_0$ in elastic media. Each panel displays the misfits due to remaining boundary reflections for a range of P-wave velocities and values for $U_0$. Colour scale shows the misfit relative to the misfit resulting from only using a Clayton-Engquist ABC ($\sigma$ in Eq. 4.9). The absorbing boundary width (span) is kept at a constant ratio relative to the P-wavelength, with a different ratio being shown in each panel. The red marker shows the optimal value for $U_0$ based on these misfits and the error bars indicate the 10% tolerance range.

Using the results displayed in Figure 4.3 as well as some additional test series conducted in the same manner, an empirical function can be fitted to the optimal values for $U_0$. The results of this are displayed in Figure 4.4.

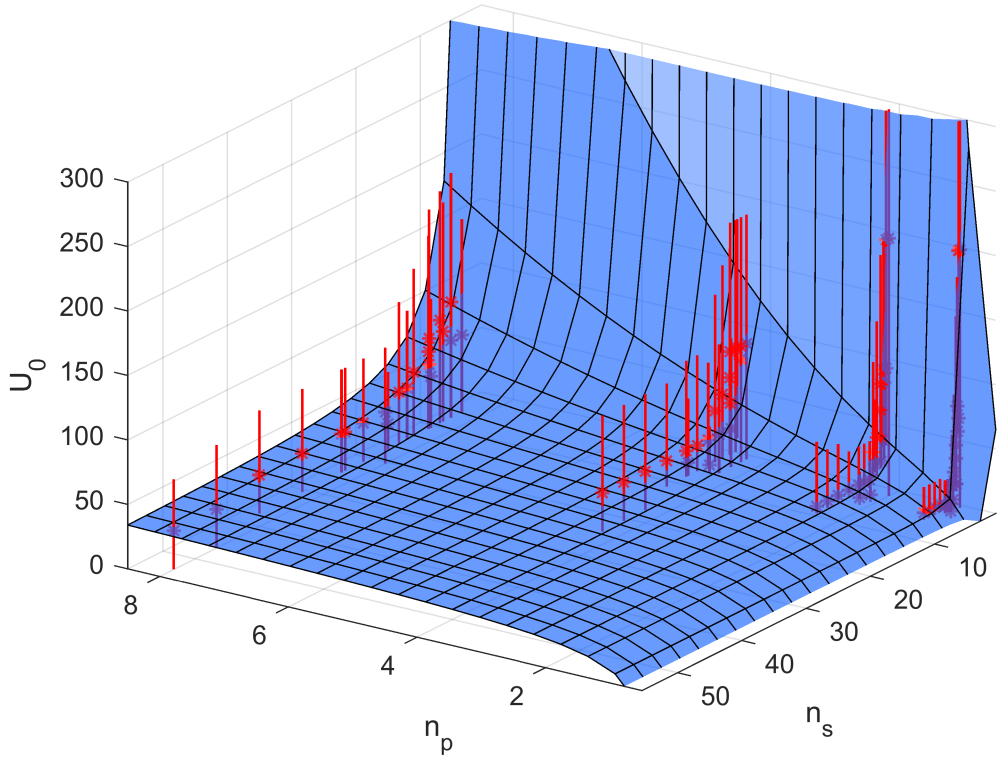



**Figure 4.4:** Comparison of optimal $U_0$ found in the systematic tests (red markers) and the $U_0$ predicted by the empirical function (blue surface). Vertical red lines show 10% error tolerance.

The equation used to create the blue surface in Figure 4.4 is:

$$U_0 = 0.8 \left( \frac{v_s}{v_p} \right)^2 - \frac{1}{2\sqrt{n_p}} + 0.35 \tag{4.10}$$

This surface does not fully fit the data at $n_p = 1$ and $n_s < 5$, and this could be fixed by using a more complex expression. However, since the sponge did not prove useful at $n_p < 2$, the simpler expression was chosen.

### 4.3.3 Results for Acoustic Media

The results of the test series for fluid models are shown in Figure 4.5. Note that the x-axes are different than in the elastic test results. Instead of increasing $v_p$

throughout each series for a few series with different $n_p$, I now keep $v_p$ constant in each series, allowing for a larger range of $n_p$ to be tested. The reason for this is that increasing $v_p$ at a constant $n_p$ in the fluid should result in an equivalent model (in the solid this was not the case, since the resulting increase in $w$ would in turn affect $n_s$).

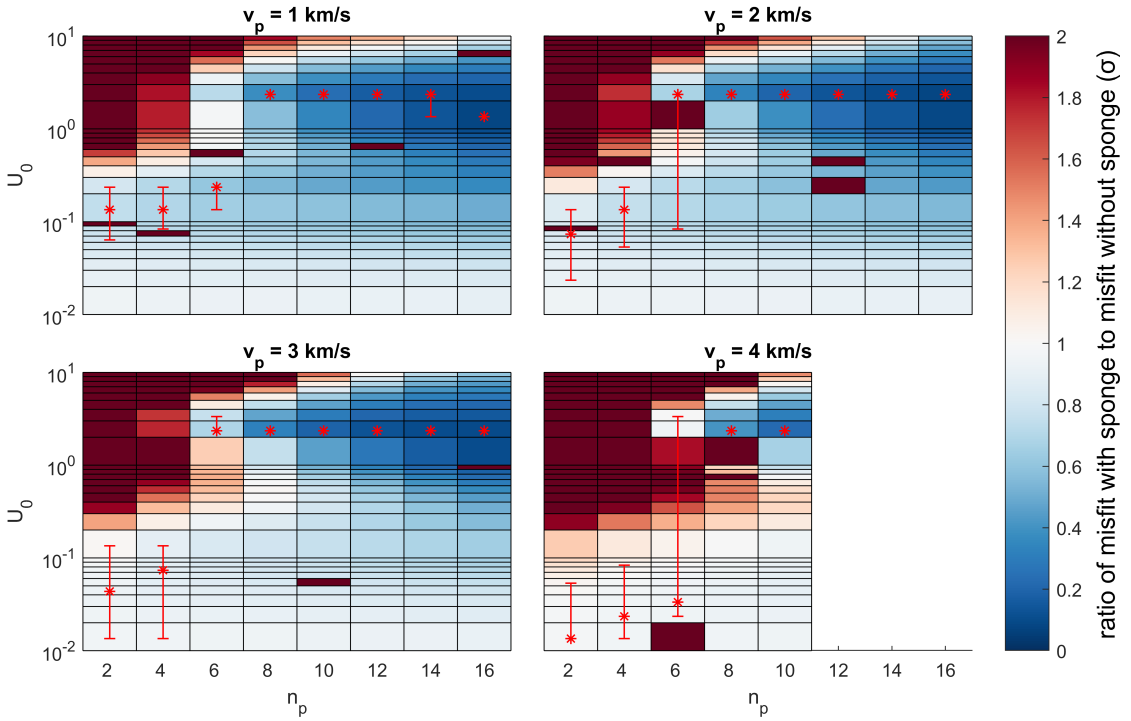



**Figure 4.5:** Results from the systematic tests of the performance of the mixed absorbing boundary using different values for $U_0$ in acoustic media. Each panel displays the misfits due to remaining boundary reflections for a range of values for $n_p$ (the ratio between wavelength and sponge width) and $U_0$. Colour scale shows the misfit relative to the misfit resulting from only using a Clayton-Engquist ABC ($\sigma$ in Eq. 4.9). The P-wave velocity, $v_p$, is kept constant for each series with a different $v_p$ being shown in each panel. The red marker shows the optimal value for $U_0$ based on these misfits and the error bars indicate the 10% tolerance range.

Compared to the elastic models, the minimum effective boundary size is significantly larger, with an $n_p$ of 8 being the smallest relative sponge width which reliably improves absorption. It should be noted, however, that the Clayton-Engquist ABC performs better in the fluid (see Figure 4.2, middle panels), making it more difficult to improve upon.

As mentioned previously, we should expect little to no difference between the panels. Yet, at small values for $n_p$ (<8) there are some clear variations. This

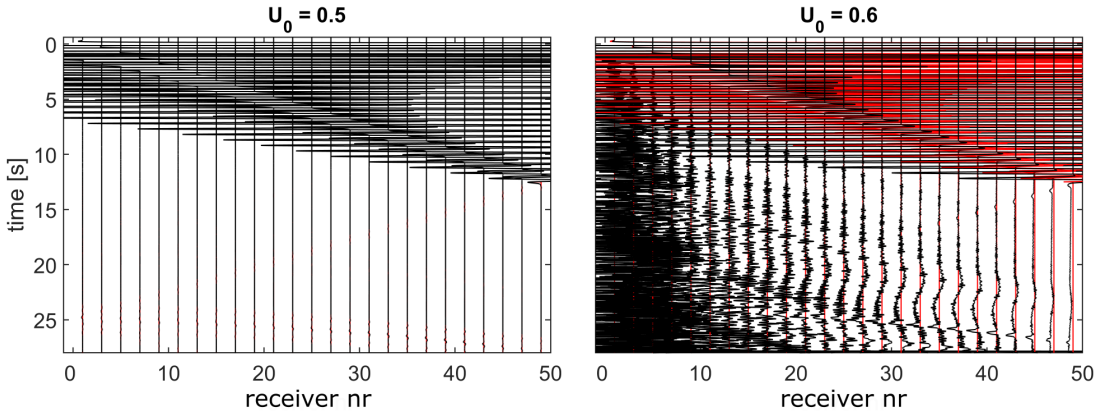



**Figure 4.6:** Example traces showing occurrence of numerical noise for seemingly random values of $U_0$ in fluid media. In the given example the model parameters are $v_p = 1.0$ km/s and $n_p = 12$. The traces are recoded at receivers spaced regularly in 0.5 km intervals at the surface of a 25 km wide mesh. Reference traces without boundary reflections are shown in red.

does not pose a problem for the empirical equation for $U_0$, however, as I am only interested in predicting $U_0$ in the parameter space which is useful for seismic simulations. Another observation from Figure 4.5 proves to be more problematic: there are some simulations with very high error which seem to be arbitrarily scattered throughout the panels. Upon closer inspection of the traces, it turns out that the random occurrences of large errors are not due to boundary reflections but to strong numerical noise which appears at certain combinations of $w$, $v_p$ and $U_0$. For some simulations, the noise can be removed by a low-pass filter but for others this is not sufficient. An example of this is shown in Figure 4.6 where, for an otherwise equal model, $U_0 = 0.5$ preforms quite well but $U_0 = 0.6$ causes very high amplitude noise and a decrease in amplitude of the direct wave.

In conclusion, the tests in fluid media show that 1. large sponge boundaries (more than 8 times the P-wavelength) are required to improve absorption, 2. $U_0 = 2$ is a good choice for the absorption factor, independently of other parameters, and 3. sponge layers in the fluid may cause erratic numerical noise. If such noise occurs, it can be fixed by slightly changing $U_0$.

### 4.3.4 How does the Error Decrease with Sponge Width?

We now have empirical equations which can be used to automate the choice of $U_0$ in the AxiSEM3D source code. But the choice of the sponge width must be left up to the user since it presents a trade-off between cost (because the mesh may need to be extended) and absorption performance. To aid future users make an educated choice, I seek another pair of empirical expressions to predict the remaining error versus sponge width. The above models have been re-run using the empirically predicted $U_0$, and the resulting errors of these simulations are shown in Figure 4.7.

Figure 4.7(a) confirms what was already apparent in the previous tests: The absorption performance in solid media depends strongly on $v_s/v_p$ ratio, with the absorption being poorest for extreme ratios at both the positive and the negative end possible of the range. The blue surface was created using the following equation:

$$\epsilon = (0.4(\frac{v_s}{v_p} - 0.4)^2 + 0.07)e^{-0.25n_p}, \tag{4.11}$$

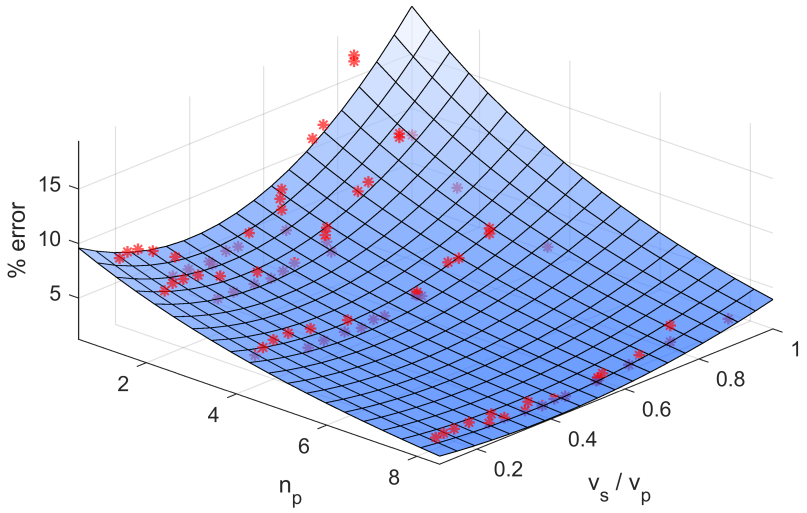
which predicts the errors with a standard deviation of $\pm 0.037$.

For the fluid, a power expression appears to match the data better than an exponential function. The best fit was achieved by

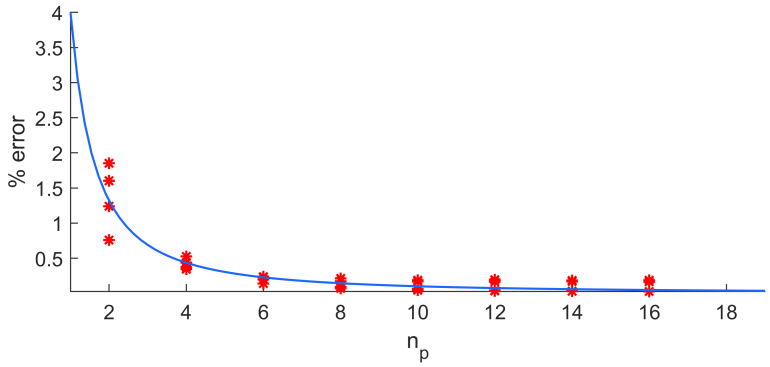
$$\epsilon = 0.05n_p^{-1.6}, \tag{4.12}$$

with a standard deviation of $\pm 0.0001$. This expression has the caveat of approaching infinity if the sponge width is reduced to 0 which does, of course, not reflect the true behaviour of the mixed boundary in fluid media. But since the use of the sponge boundary in fluids with $n_p < 8$ is deprecated, the equation can still serve users as a workable rule of thumb.

It should be noted that the above expressions are for the error, $\epsilon$, as defined in Equation 4.8, i.e. for the RMS error of 50 equally spaced stations in a 25 km wide mesh. Naturally, in a mesh with a different size and different receiver positions, the absolute values for $\epsilon$ will be different. Nevertheless, the overall trend of error versus sponge width will likely be the same.




**(a)** Error from boundary reflections versus $v_s/v_p$ ratio and relative sponge width for the entire set of solid models. Red markers show the test results having used the $U_0$ predicted by equation 4.10. Blue surface shows the empirical fit from equation 4.11.




**(b)** Error from boundary reflections versus relative sponge width for the entire set of fluid models. Red markers show the test results having used $U_0 = 2$. Blue line shows the empirical fit from equation 4.12.

**Figure 4.7:** Results from the tests of the performance of the mixed absorbing boundary using empirically predicted values for $U_0$. The errors are given as a percentage of the root-mean-square amplitude of the reference solution.

## 4.4 Discussion

This series of tests serves as the basis for the empirical equations used in AxiSEM3D to automatically select an appropriate value for $U_0$. To this end, it was necessary to find a simple solution and to find it relatively quickly. Arguably, this goal was achieved by the above test series. However, as becomes obvious from the results shown here, there are several effects at play when it comes to absorption performance of the mixed boundary which are by no means understood. The most striking example is the seemingly random appearance of noise in some configurations of the mixed boundary in fluid media. It remains to be seen if the correction term introduced in Equation 4.6 can remove this noise. There are other considerations which were knowingly neglected and may be worth revisiting: The boundary reflections consist of multiple phases and it stands to reason that each of these phases is optimally absorbed at a different value of $U_0$. So, a thorough analysis should really isolate these phases repeat the above tests for each phase individually. This may result in a min-max expression over several simple empirical equations. As it stands, the expressions given above are workable approximations of this ideal solution which would require a much larger investment of both time and computational resources to obtain.

*A good simulation, be it a religious myth or scientific theory, gives us a sense of mastery over experience. To represent something symbolically, as we do when we speak or write, is somehow to capture it, thus making it one's own.*

— Heinz Pagels

# 5

# Azimuthal Complexity Adaptation at Local Scales

## Contents


**5.1 Introduction . . . . . 47**
**5.2 Exploiting Wavefield Smoothness - Verification and Benchmark . . . . . 48**
5.2.1 The Benchmark Model . . . . . 48
5.2.2 Azimuthal Complexity Adaptation: Verification, wavefield learning, complexity maps . . . . . 49
5.2.3 Comparison with an Established Modeling Method . . . 55
**5.3 A Complex 3D Setting - Limitations of Azimuthal Smoothness? . . . . . 59**
**5.4 Discussion . . . . . 65**
**5.5 Conclusions . . . . . 68**


## 5.1 Introduction

This chapter is an excerpt of my first-author paper published in the peer-reviewed journal Geophysics (Haindl et al., 2021). The introduction and theory were not included because they are largely repeated in the overarching introduction and theory chapters of this thesis.

The purpose of this work was to test whether the ACA approach is feasible and

useful on local scales and exploration-type settings. At the point of publishing, AxiSEM3D had only been tested on global scales where the relative size of 3D structures compared to the model size is much smaller than in local-scale problems. After adapting the code for local-scale simulations as outlined in the previous chapter, I set out to analyse the azimuthal smoothness present in a variety of local-scale settings and to compare the performance of AxiSEM3D against an established SEM.

## 5.2 Exploiting Wavefield Smoothness - Verification and Benchmark

In the following section, we benchmark AxiSEM3D using a 3D model with elastic and acoustic regions. We demonstrate the use of wavefield learning to adapt the overall problem size (degrees of freedom) to the 3D complexity of the wavefield. Further, we compare synthetic data produced by AxiSEM3D with the output of SPECFEM3D Cartesian (Peter et al., 2011), an established elastic local-scale modeling solver based on a 3D SEM.

### 5.2.1 The Benchmark Model

| Medium | P-wave velocity | S-wave velocity | Density |
|---|---|---|---|
| Water | 1500 | 0 | 1000 |
| Background | 3000 | 1500 | 2000 |
| Salt | 5000 | 2800 | 2200 |

**Table 5.1:** Medium properties of the 3D benchmark model. Units are m/s and kg/m$^3$ for velocities and densities, respectively.

The goal of our benchmark design is to have a small model which uses all the 3D capacities of AxiSEM3D. To facilitate the comparison with SPECFEM3D, we further match our model design to the functionality of "xmeshfem", the automated 3D mesher which is provided as part of the SPECFEM3D software package.

The model consists of an ocean, a homogeneous background medium and an off-axis cuboid with salt-like properties (Figure 5.1). The ocean floor is undulated with a depth range between 300 m and 650 m. All medium parameters can be

found in Table 5.1. The input model for SPECFEM3D is fully honored in its mesh, adhering to all medium interfaces. In AxiSEM3D, the undulating ocean floor is treated by using particle relabeling, whereas the salt cuboid is volumetrically imposed. The whole model measures 5 km x 5 km x 5 km, but in AxiSEM3D only a cylinder of 5 km height and 5 km diameter is utilized. An explosive-type source with a Gaussian source wavelet is placed at 1 km depth. For the sake of comparability, the Clayton-Engquist type absorbing boundary condition is used on its own, since only this one is implemented in both solvers.

### 5.2.2 Azimuthal Complexity Adaptation: Verification, wavefield learning, complexity maps

#### Verification of Azimuthal Complexity Adaptation

Complexity maps are created as part of the ACA process. They reveal information about the wavefield structure and they can be used to analyze the model-specific performance of the ACA-method. For visualization and interpretation, we do

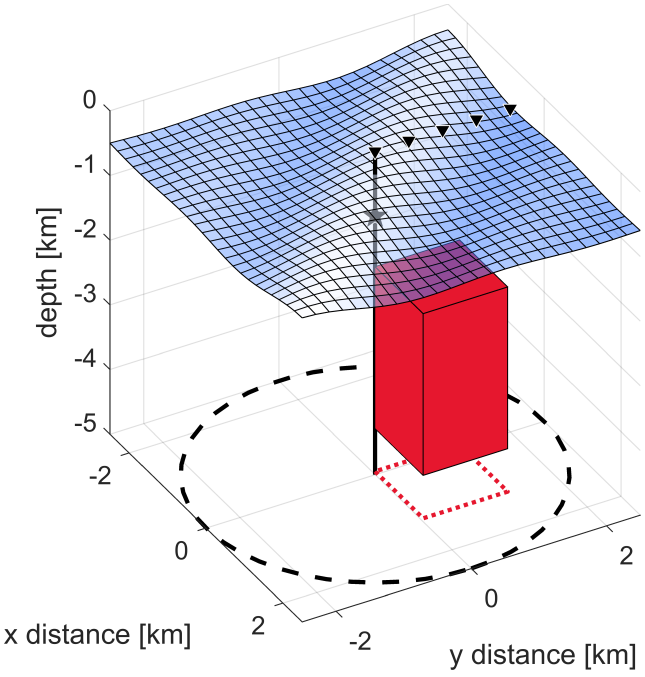



**Figure 5.1:** Image showing the 3D undulations of the ocean floor discontinuity and the position of the salt-like cuboid in the benchmark model. Solid and dashed black lines show axis and extent of the AxiSEM3D model, respectively. Black triangles indicate receiver positions and star marks source position.

not plot $n_u$ directly, as it scales with distance from the axis, $s$. We instead use the perhaps more intuitive azimuthal wavenumber $k_{\phi,max}$. This is the maximum wavenumber in azimuthal direction that needs to be resolved as suggested by the test simulation on each point in the 2D mesh:

$$k_{\phi,\max} = \frac{2n_u + 1}{4\pi s}. \tag{5.1}$$

Figure 5.2 shows two complexity maps created for the benchmark model with a 10 Hz source. They differ in the number of peaks that were considered when determining the required azimuthal resolution. Note that the resolution in the vertical and radial direction is set by the mesh spacing and is equivalent to maximum wavenumbers around 32 km$^{-1}$.

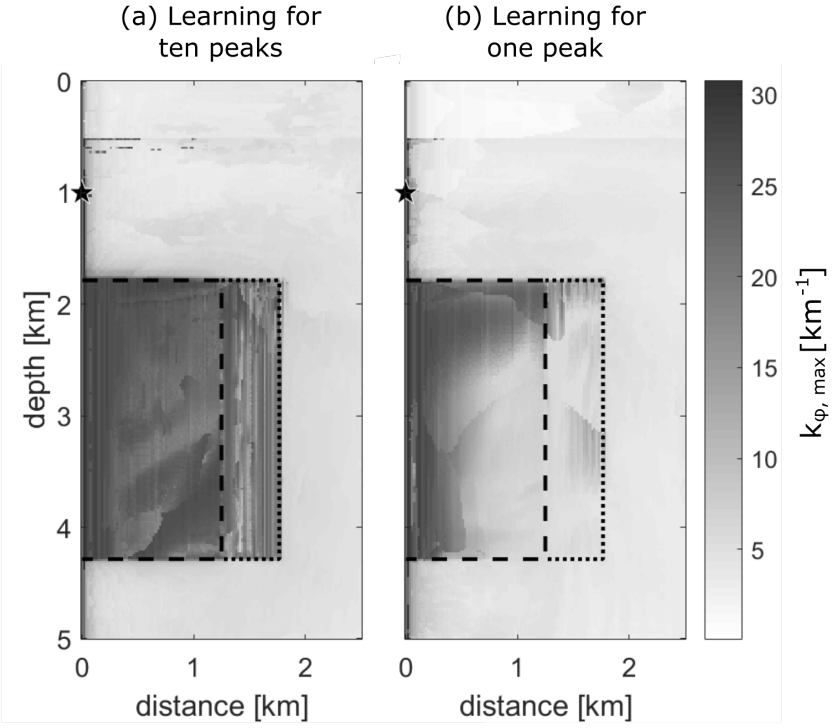



**Figure 5.2:** Comparison of complexity maps of the benchmark model created by fitting a different number of amplitude peaks in the time series with the learning algorithm. Gray scale shows maximum tangential wavenumber in km$^{-1}$. Dashed and dotted lines show the outline of the salt block: dashed lines represent the faces bordering the source-centered axis and dotted lines show the extent across the diagonal. Star indicates source position.

The reduction in degrees of freedom compared to a high resolution run (i.e. a run with azimuthal sampling equal to in-plane sampling) is 76% and 86% for the maps

shown in Figure 5.2(a) and Figure 5.2(b) respectively. In Figure 5.3, we compare synthetic seismic traces produced for each case (for the geometry see Figure 5.1). As a reference case we use traces produced at high azimuthal resolution.

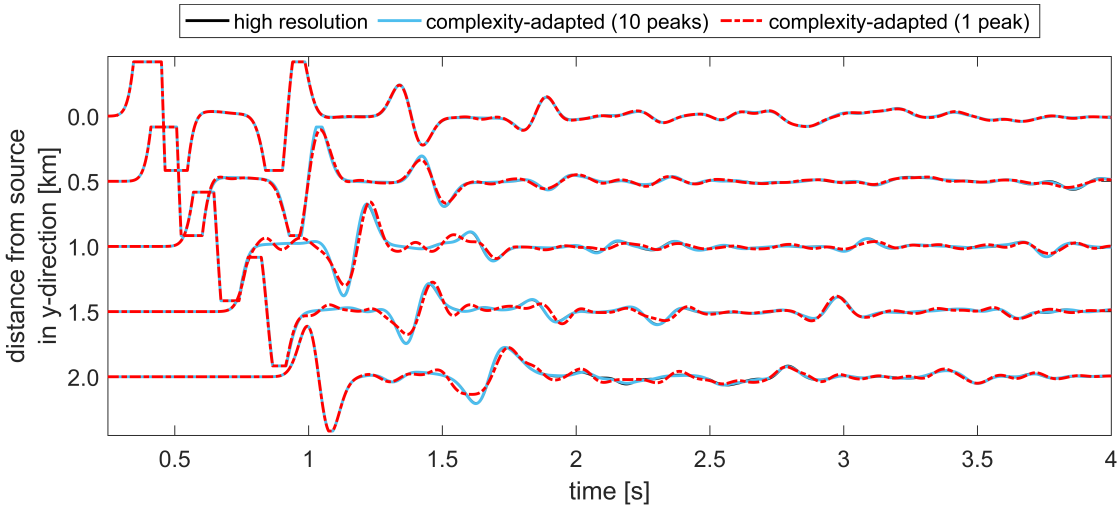



**Figure 5.3:** Comparison of synthetic traces of the benchmark model produced by AxiSEM3D using high azimuthal resolution and azimuthal complexity adaptation (respective complexity maps shown in Figure 5.2). A 10 Hz Gaussian source wavelet was used. Traces show vertical displacement clipped at 30%. The high-resolution results (black line) are barely visible as they are covered by the complexity-adapted results using ten peaks (blue line).

There is no visible difference between the traces produced by the highly resolved run (solid black) and the ACA-optimized run where ten peaks were used for learning (solid blue). The dashed line shows the result of a simulation where only the largest amplitude was considered by the wavefield learning algorithm. Using multiple peaks for wavefield learning is evidently vital for the proper resolution of all the phases in this benchmark model. Note that the method assures first arrivals to be resolved properly in either case. If only first arrivals are required for a certain problem (e.g., travel-time tomography or velocity modeling), resolving only the maximum amplitude is a viable and more computationally efficient approach.

### Azimuthal Complexity Adaptation at Different Source Frequencies

We now demonstrate how azimuthal complexity varies with source frequency, and how these findings can be used to save computational cost when performing wavefield learning.

Figure 5.4 gives an overview of how complexity maps change across a range of source frequencies. With increasing frequency, we observe a proportional increase of the highest values of $k_{\phi,max}$. The complexity in the homogeneous regions of the model does not increase at the same rate. Further, the relative effect of the seafloor on complexity becomes greater and a higher complexity feature can be seen just below the seafloor at 20 Hz and 40 Hz. There are $k_{\phi,max}$ contrasts, particularly visible inside the salt block at 10 Hz and higher, which appear to vary arbitrarily between the maps.

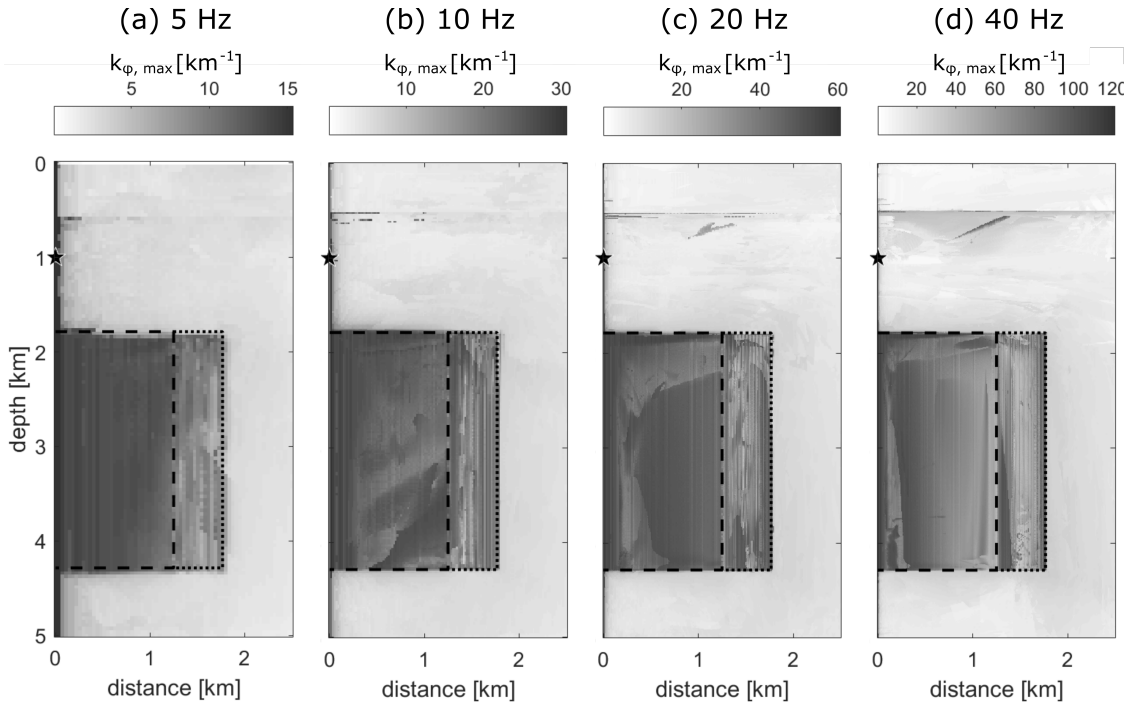



**Figure 5.4:** Comparison of complexity maps of the benchmark model at different source frequencies. Ten peaks were used for learning. Gray scale shows maximum tangential wavenumber in km$^{-1}$. Dashed and dotted lines show the outline of the salt block: dashed lines represent the faces bordering the source-centered axis and dotted lines show the extent across the diagonal. Black stars indicate source position. Note the change in scale for each frequency.

Crucially, the main features of the maps remain the same and the observed linear relationship between maximum complexity and frequency can be exploited to reduce the cost of high resolution test runs. Depending on the application, there are two ways to go about this:

Firstly, information from low-frequency test runs can be used to optimize the cost of a full test run. As a starting point for a high-$n_u$ test run, one might ensure a spatial sampling rate around the azimuth equal to the sampling rate in the radial and vertical directions, which is dictated by the spacing of the 2D mesh. However, this often results in significant oversampling and an unnecessarily high cost for the test run. As shown in Figure 5.4, the maximum value of $k_{\phi,max}$, and hence $n_u$, is proportional to the source frequency, so a cap for $n_u$ can be determined from a low resolution run. This can also be used as a first order of optimization if the full test run is to be avoided and there is little tolerance for errors.

Secondly, if a full test run is to be avoided and some amplitude error is allowed, one may simply use a smoothed (to remove frequency-specific features) complexity map from a low-frequency test run, multiplied by the ratio of target frequency to test frequency. We demonstrate this by performing a 20 Hz run with a smoothed version of the map in Figure 5.4(a) multiplied by four, a smoothed version of the map in Figure 5.4(b) multiplied by two and with the 20 Hz map shown in Figure 5.4(c) as a reference. The resulting traces are compared in Figure 5.5.

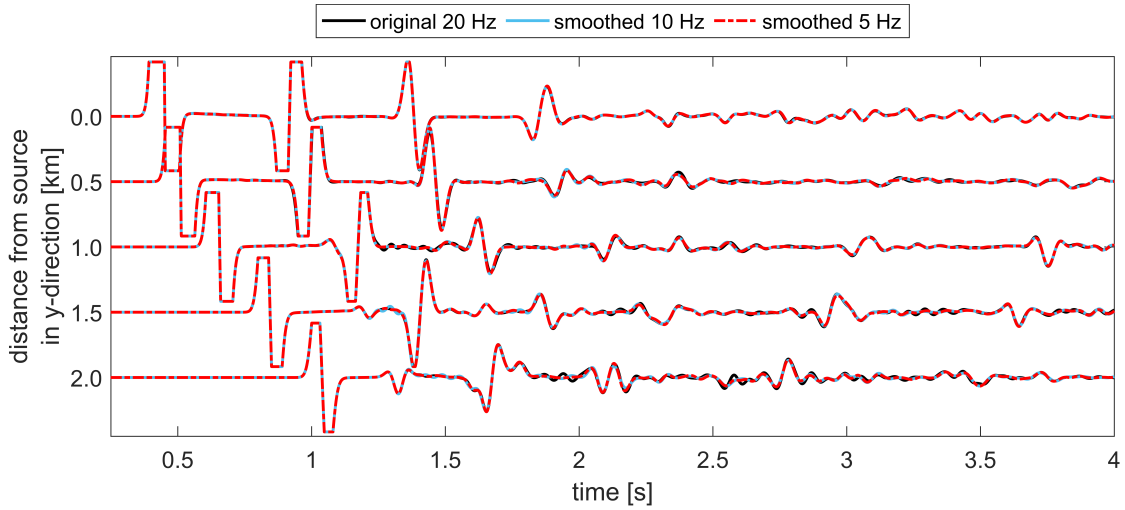



**Figure 5.5:** Comparison of synthetic seismic traces of the benchmark model produced by ACA simulations with a standard complexity map and approximate maps obtained by scaling up lower-frequency complexity maps. A 20 Hz Gaussian source wavelet was used. Traces show vertical displacement clipped at 15%.

The fit of the traces obtained using scaled lower-frequency complexity maps is close, but some deviations are visible. Most waveforms are present and arrival

times are consistent. The most noticeable deviations are seen in low-amplitude, high-frequency events. Interestingly, the 10 Hz approximation offers no improvement over the 5 Hz approximation in terms of the trace misfit. The scaled 10 Hz map does, however, contain fewer Fourier coefficients, so the final run was slightly cheaper at the expense of a more costly test run.

**Wavefield Learning as a Measure for Wave-Scattering Complexity**

We now assess whether one can gain more physical insight from wavefield learning beyond the speed-up for the ACA-SEM. The complexity maps offer a means to quantify the "3D-ness" of the wavefield. This "3D-ness" is brought about by a combination of non-axisymmetric forward scattering by refraction, as well as off-plane scattering. Analyzing these wave effects can provide quantitative information about the extent to which a wavefield may illuminate any given structures from a basic wave-physical perspective, irrespective of any forward or inverse methodologies.

An interesting observation from all the complexity maps shown so far is how the maximum tangential wavenumber instantly declines to a fairly constant background level outside the salt block. Firstly, this suggests that the highest levels of azimuthal complexity stem from regions where the Fourier series are recorded across sharp velocity contrasts. Secondly, the wavefield retains a relatively constant level of azimuthal smoothness once it passed through a "disruptive" structure that gives rise to a certain complexity.

As seen in Figure 5.2, we can also infer direct wave illumination. In Figure 5.2(a), the maximum tangential wavenumber in the vicinity of the salt block is comparatively constant. The variations that can be seen in this area may be attributed to interference effects which depend on the angle of incidence of the wavefronts on the faces of the cuboid. Similar structures are seen in Figure 5.2(b). However, in Figure 5.2(b), there is a decline of azimuthal complexity with increasing distance from source and axis. This likely results from the fact that direct waves no longer graze the edge of the salt block as they travel outwards from the source, hence do not cause the same high levels of complexity seen in the other panel.

Near the outer edge of the cuboid (dotted line), the complexity increases slightly. This is an area with an azimuthally narrow velocity perturbation, so we attribute this to asymmetric refraction. The edges closer to the axis are barely sampled by direct waves. Such observations about the wavefield complexity in a given structure may help putting results from seismic imaging and tomography into perspective: irrespective of data quality and inversion method, the basic wavefield itself does not contain the same level of complexity as the underlying structure, even above the diffraction limit. In future work, this can be helpful in better constraining true resolution limits and uncertainties for seismic images.

In Figure 5.4, we can see how the sensitivity of the wavefield with respect to certain structures changes across source frequencies. The seafloor undulations have a wavelength of around 200 m. The 5 Hz wavefield (Figure 5.4, a), with wavelengths of 300-600 m in the vicinity of the seafloor, does not pick up this complexity. The sensitivity increases for each successive panel, as the seismic wavelength is halved. Again, this can help in identifying the true resolvability of a structure by the corresponding wavefield, which becomes nontrivial in multiscale, strongly complex 3D media.

### 5.2.3 Comparison with an Established Modeling Method

Comparing different seismic modeling methods has always presented a challenge (e.g. Moczo et al., 2007). However, since AxiSEM3D converges towards the solution of a full 3D SEM with an increasing number of Fourier coefficients (Leng et al., 2016), a close fit can be expected when comparing it to an independent 3D SEM solver, such as SPECFEM3D Cartesian. SPECFEM3D is probably the most widely used SEM implementation in the seismological research community, but our comparisons are expected to fare similarly for any SEM or other fully discrete 3D implementation. They merely highlight the conceptual addition of the ACA idea. In Figure 5.6, we show synthetics produced by AxiSEM3D and SPECFEM3D. The 3D mesh for SPECFEM3D was created using the intrinsic mesher “xmeshfem”, which allows to mesh the salt block and undulating ocean floor while keeping the top of the salt flat.

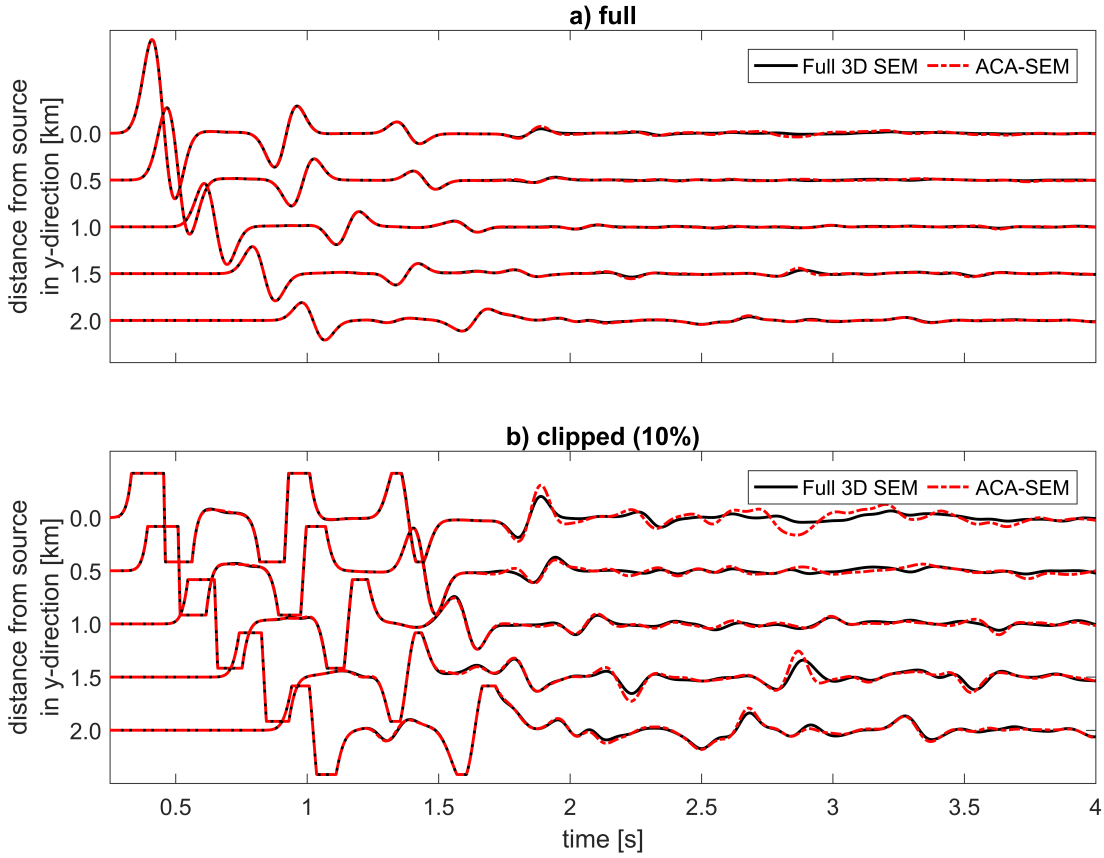



**Figure 5.6:** Comparison of synthetic traces of the benchmark model produced by the 3D ACA-SEM, AxiSEM3D, and the full 3D SEM, SPECFEM3D Cartesian. The ACA traces were calculated using ten-peak optimization. A 10 Hz Gaussian source wavelet was used. **(a)** Full range of vertical displacement. **(b)** Range clipped at 10% of maximum displacement.

The resulting fit between the two sets of synthetics is different for the large-amplitude ocean multiples (Figure 5.6, a) and the smaller reflections caused by the salt cube (Figure 5.6, b): The free-surface multiples fit well while the salt reflections do not. Regarding the small events, the fit seems to get better at some distance from the axis.

To do a quantitative comparison of the fit while accounting for the observed difference, we roughly divide the traces into two sections. The first part, up to 2.2 s for all traces, appears to be dominated by ocean multiples. For this first section of the wavefield, the fit consistently falls within the "excellent" range according to the goodness-of-fit criterion from Kristeková et al. (2009), with a maximum

envelope misfit of 0.01 and a phase misfit smaller than 0.005. For the remaining salt-dominated part of the wavefield, the fit is classified as "fair" for the two traces closer to the source (envelope misfit 0.55-0.57; phase misfit 0.3) and "good" for the three traces further away (envelope misfit 0.15-0.3; phase misfit 0.05-0.14).

In light of a better fit shown later in this paper, we interpret these results as a direct consequence of rigorous meshing versus volumetrically imposing a 3D model onto a regular mesh. This demonstrates the well-known deficiency of weak-form methods such as SEM to capture sharp boundaries within element interiors. The ocean was treated rigorously in both AxiSEM3D (using relabelling) and SPECFEM3D (using meshing). The physics behind the free-surface multiples are relatively complex: a free surface and a solid-fluid interface with 3D undulations. So, the close fit obtained here can be seen as a confirmation that ACA is in principle capable of modeling complex 3D structures accurately. The salt cuboid, on the other hand, was meshed rigorously for SPECFEM3D and imposed volumetrically in AxiSEM3D. The issues that arise from treating models differently in such a way are widely known (e.g. Chaljub et al., 2010). This comes to show that the use of 2D meshes for ACA can cause problems in media with near-vertical edges of strong contrasts if not honored explicitly by the 2D mesh. 3D Discontinuities can only be honored through relabelling, which is currently limited to azimuthally continuous interfaces.

**Cost Analysis**

One of the original motivations behind the pseudo-spectral method is the halving of degrees of freedom in each spatial dimension when operating in wavenumber domain. The ACA method with its 1D Fourier space therefore halves the overall DOF compared to a fully discrete 3D method, even before any complexity adaptation. The calculation of spatial gradients in Fourier space also becomes computationally cheaper. On the other hand, any model inhomogeneities in the dimension(s) where PSM is used (e.g., azimuthal in our case) require a forward and an inverse Fourier transform for every time step. ACA methods thus have a higher computational

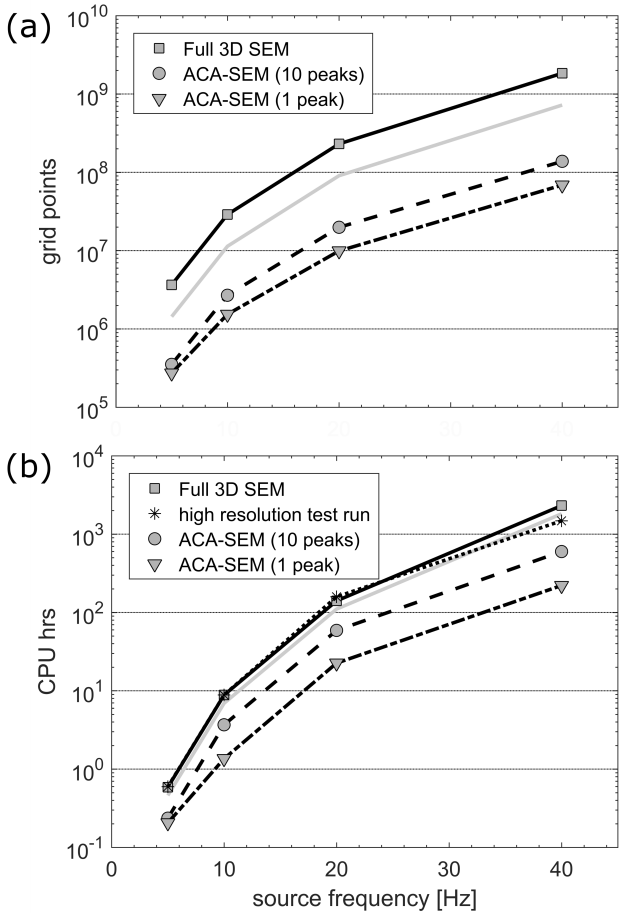



**Figure 5.7:** **(a)** Graph comparing the spatial DOF (grid pints) of a simulation of the benchmark model with a full 3D SEM and an ACA-SEM. For the 3D SEM the total number of points in the mesh is given. For the complexity maps the total summed length of the Fourier series of every 2D GLL point is given. **(b)** Graph comparing the simulation cost of running SPECFEM3D and AxiSEM3D on the Benchmark model. Gray solid lines in both graphs show the full SEM values multiplied by a factor 0.785. This is to account for cube used by the full SEM being larger than the cylinder used by the ACA-SEM.

cost per DOF, compared to a fully discrete method. This cost overhead for the ACA method needs to be offset by cost savings through complexity adaptation to render the ACA method competitive.

In Figure 5.7(a), we show how the DOF of the benchmark simulation change with source frequency for a standard 3D method and two complexity adapted AxiSEM3D runs, one of which uses ten peaks for learning, and one of which uses the largest peak only (optimized for first arrival). Using ACA while aiming for full resolution of all phases consistently reduces the DOF in ACA by a bit more than one order of magnitude. If only the first arrival is resolved by ACA, the

DOF reduce by a factor of 20-30. Note that the difference in volumes between the cylindrical ACA domain and the cuboid full 3D domain (see Figure 5.1) is accounted for in all relative cost comparisons.

A comparison in terms of DOF ignores the aforementioned, possibly compromising overhead of the ACA-SEM. We therefore show a comparison of the computational cost in CPU hours for both AxiSEM3D and SPECFEM3D Cartesian in Figure 5.7(b). Note that such a comparison is influenced by the specific implementation of each solver and hardware specifications. Both AxiSEM3D and SPECFEM3D continue to be improved and made more efficient. The simulations for both solvers were performed on the UK supercomputer ARCHER with identical CPU-cores.

As expected, the speedup for ACA is smaller than its DOF reduction: about 1/3 when using the ten-peak complexity map. In Figure 5.7(b), we also show the cost of the high-$n_u$ test runs for wavefield learning. If the high-$n_u$ cut-off is chosen carefully, the DOF can become small enough to make up for the overhead of the method, even for the test simulation.

In summary, our ACA-SEM is significantly faster for all tested examples compared to the fully discrete 3D SEM solver. At global scale, the difference is even bigger. This is because ACA methods scale better when structures become large compared to the wavelength (Leng et al., 2019). We did not reach this parameter regime of higher speedups here, but whenever high-frequency simulations become interesting, the computational advantage of ACA methods shall become even more pronounced for local-scale applications as well. Moreover, the memory requirements of AxiSEM3D are significantly smaller than the requirements of SPECFEM3D, thanks to the mesh being in 2D.

## 5.3 A Complex 3D Setting - Limitations of Azimuthal Smoothness?

Finally, we explore the performance of the ACA-SEM with the SEG/EAGE salt model, a model containing a highly irregular salt body as well as other heterogeneities. This is really the opposite of the idealized use case for ACA, but we are interested

whether such a regime of realistic complexity renders ACA pointless due to a potential breakdown of azimuthal smoothness. Further, we examine whether model smoothing and a fixed error tolerance provide some leverage for speeding up accurate ACA simulations in highly complex models.

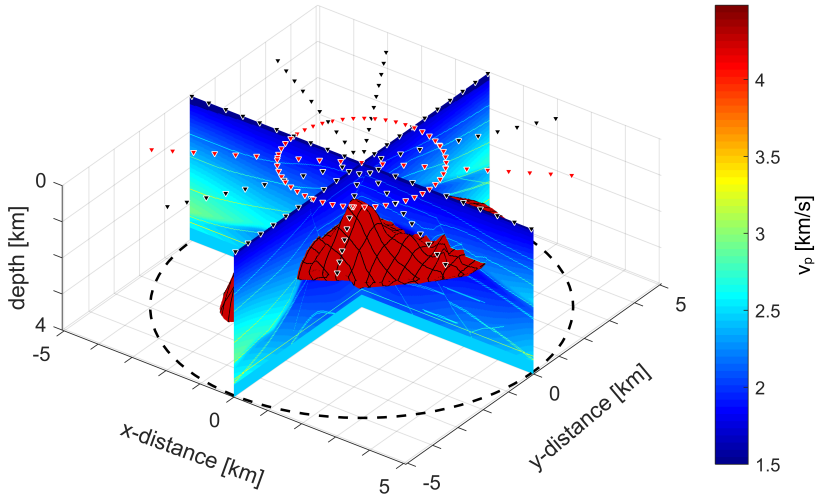



**Figure 5.8:** Cross sections of the SEG salt model with 3D outline of the salt body. Color scale shows P-wave velocity in km/s. Grid lines are placed around the salt body for better 3D visualization. Triangles indicate receiver positions. The receivers for which the synthetics are shown below are highlighted in red.

The model is shown in Figure 5.8. The ocean in the original model is replaced by a homogeneous elastic overburden and the model is treated purely volumetrically. To ensure comparability between the results from SPECFEM3D and AxiSEM3D, only the Clayton-Engquist type absorbing boundaries are used as above. An explosive source with a 10 Hz Gaussian wavelet is placed at 10 m depth.

Figure 5.9 contains the complexity maps created while resolving 16 peaks and one peak, respectively. Regions of high complexity can be found in the faulted section of the salt body, around the overburden, and near the outer edge of the model. The high complexity near the outer edge is pronounced around 3 km depth, where there is a large range of different velocities around the azimuth. The differences between the two maps are not as obvious as they were for the benchmark model. The map created while resolving the 16 largest wavefronts simply exhibits a

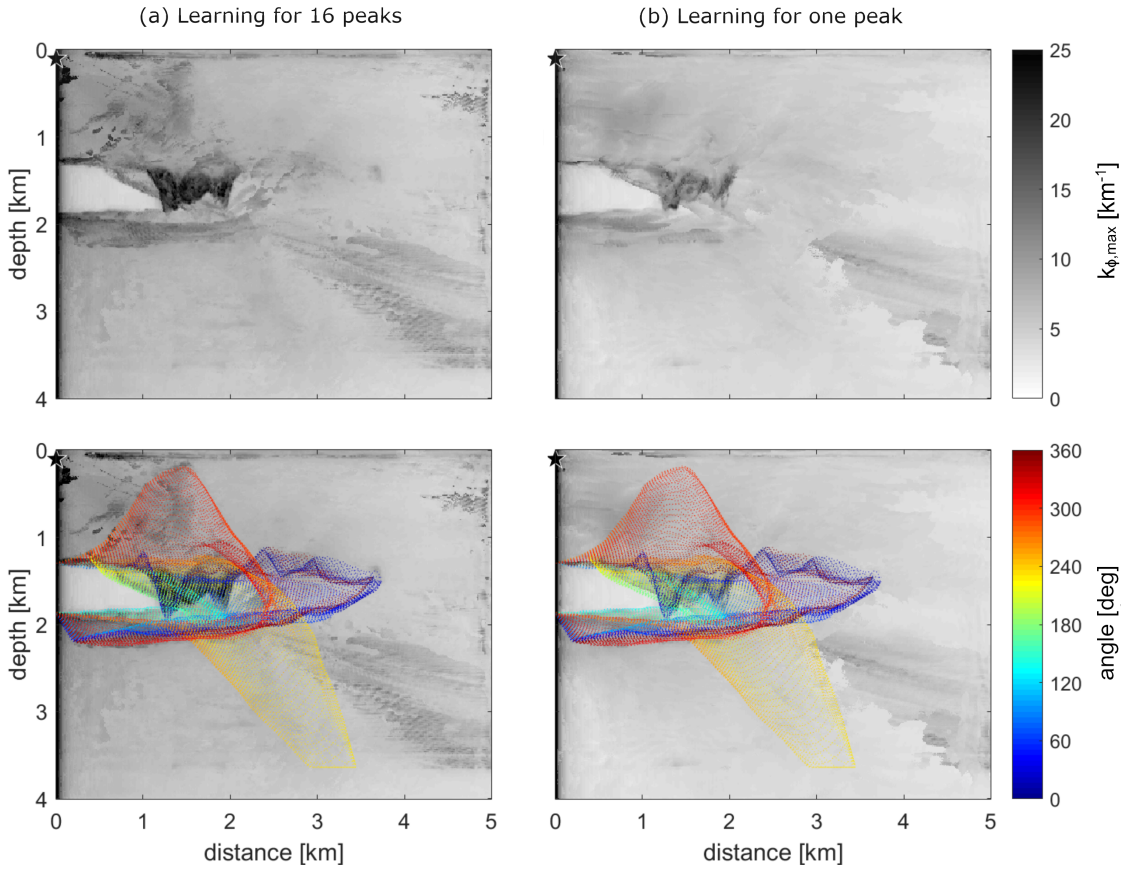



**Figure 5.9:** Comparison of complexity maps of the SEG/EAGE salt model created by fitting a different number of peaks with the wavefield learning algorithm. Top panels show the complexity maps on their own and bottom panels show the same maps superimposed by outlines of the salt body as seen in vertical cross-sections taken at different angles from the positive x-axis ($\phi$). Gray scale indicates maximum tangential wavenumber in km$^{-1}$. Color scale indicates the angle of the outline of the salt body. Black stars indicate source position.

higher level of complexity. A notable area is near the outer edge (around 0.6 km depth) where high $k_{\phi,max}$ values are seen in Figure 5.9(a), but the same feature does not exist in Figure 5.9(b). These are presumed to be boundary reflections which are not absorbed fully by the Clayton-Engquist boundary.

Figure 5.10 shows a representative subset of the calculated synthetics (positions shown as red triangles in Figure 5.8). The traces produced by AxiSEM3D closely match the ones produced by SPECFEM3D. One exception to this is the central trace in Figure 5.10(a), where a drift can be observed in the SPECFEM3D trace. This is probably an artifact due to the close proximity of the source and the central receiver. We have calculated the goodness-of-fit (Kristeková et al., 2009) for all

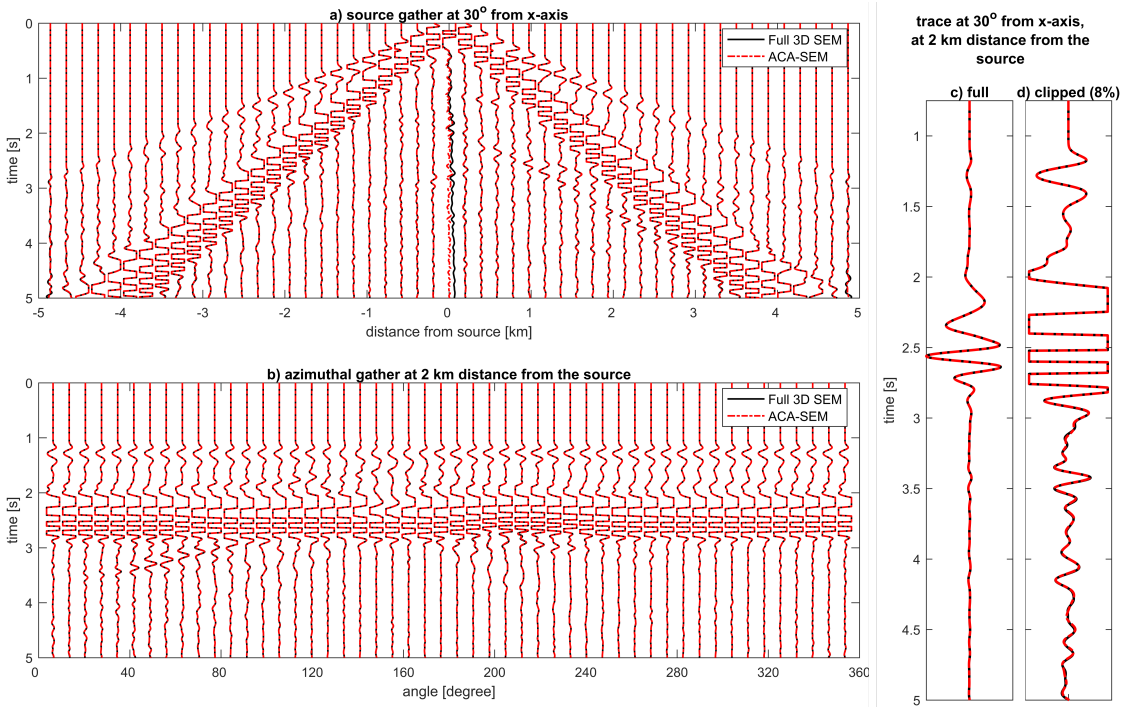



**Figure 5.10:** Comparison of synthetics (vertical displacement) produced by the ACA-SEM, AxiSEM3D, and the full 3D SEM, SPECFEM3D for the SEG salt model. The ACA simulation was optimized for 16 peaks. Receiver positions are shown on Figure 5.8. A 10 Hz Gaussian source wavelet was used. **(a)** Common azimuth shot gather. **(b)** Variable azimuth shot gather. **(c)** Single trace showing full range. **(d)** Single trace clipped at 8% of maximum displacement.

traces, including the ones not shown here, but excluding the central trace as well as the outermost traces which seem to be dominated by boundary reflections. With a maximum phase misfit below 0.005 and a maximum envelope misfit of 0.01, the fit between SPECFEM3D and AxiSEM3D can again be classified as “excellent”.

The results obtained when using ACA are consistent with the results from a high-resolution run (Figure 5.11). Notably, we can observe the effect of fitting different numbers of peaks in Figure 5.11(c) and Figure 5.11(d): the first arrival and the ground roll are fitted properly even if just one peak is used for learning. Deviations become visible later, after 3 s. Comparing the high-resolution AxiSEM3D run with the optimized runs, the required CPU hours are reduced by 77% and 85% for the 16-peak optimization and the single-peak optimization, respectively. Interestingly, an ACA simulation optimized to fit only the single largest peak is more accurate for the SEG/EAGE model than for the benchmark model. This may

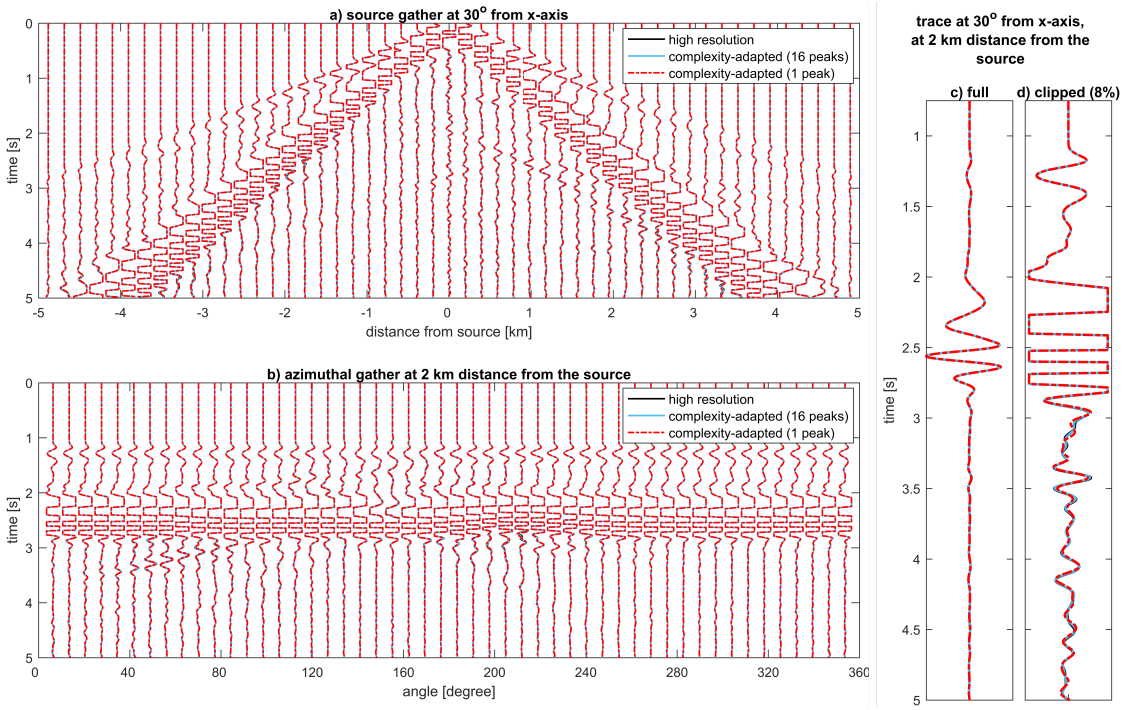



**Figure 5.11:** Comparison of synthetics (vertical displacement) produced by the ACA-SEM, AxiSEM3D, using high azimuthal resolution and azimuthal complexity adaptation (respective complexity maps shown in Figure 5.9). Receiver positions are shown on Figure 5.8. A 10 Hz Gaussian source wavelet was used. **(a)** Common azimuth shot gather. **(b)** Variable azimuth shot gather. **(c)** Single trace showing full range. **(d)** Single trace clipped at 8% of maximum displacement.

be due to the particular shape of the salt body in the benchmark model where the parallel flat sides exacerbate strong multiple reflections.

Having established the capability of the ACA method to handle complex models, we now investigate two more factors: the computational cost compared to a standard 3D SEM across a range of source frequencies, and the performance of the ACA-SEM for smoothed models. Figure 5.12 shows a version of the SEG/EAGE salt model which has been smoothed using a Gaussian filter. Similar smoothing is common for inversions, since sharp contrasts in background models may lead to local minima (Virieux and Operto, 2009). The DOF and associated costs of the simulations are summarized in Figure 5.13. All ACA runs shown in the figure are based on the largest peak only.

Despite the higher complexity of the model, the results resemble the cost comparison for the benchmark model. The computational cost of the high-resolution

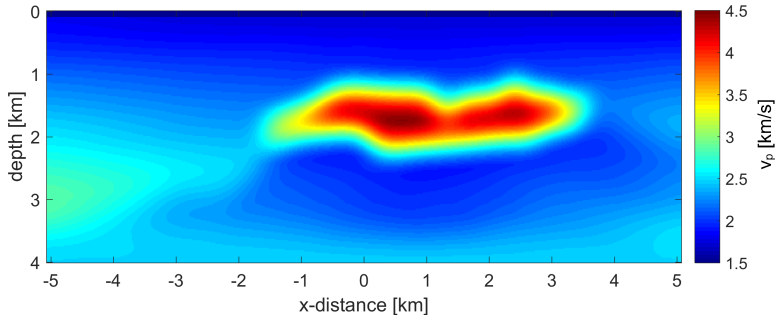



**Figure 5.12:** Cross-section through a smoothed version of the SEG/EAGE salt model. Color scale indicates P-wave velocity in km/s.

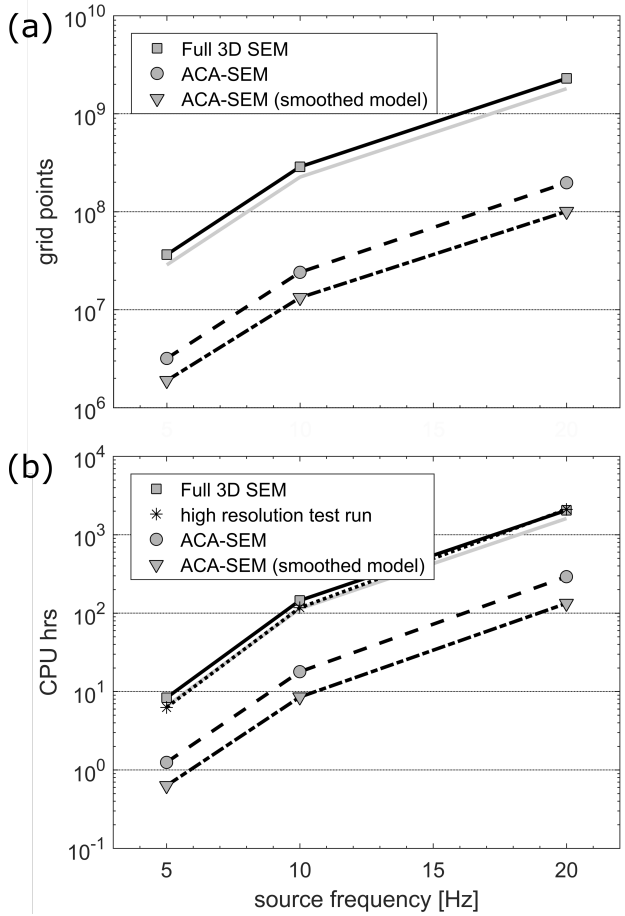



**Figure 5.13:** **(a)** Graph comparing the DOF of a simulation of the salt model with a full 3D SEM and an ACA-SEM. For the 3D SEM the total number of 3D GLL points in the mesh is given. For the complexity maps the total summed length of the Fourier series of every 2D GLL point is given. **(b)** Graph comparing the simulation cost of running SPECFEM3D and AxiSEM3D on the salt model. Gray solid lines in both graphs show the full SEM values multiplied by a factor 0.785. This is to account for cube used by the full SEM being larger than the cylinder used by the ACA-SEM.

test runs is in the order of the cost of the classic 3D simulations. The DOF reduction and associated cost savings from ACA optimization are consistent across the tested frequencies. For the original model, the DOF are reduced, on average, by a factor 11 and this results in the ACA simulation having a 7x shorter runtime than the full 3D simulation in terms of CPU hours. Model smoothing results in azimuthally smoother wavefields, manifesting itself in DOF being reduced by 20x, and a 15x reduction in computational cost.

## 5.4 Discussion

The question of how to optimally model environments of various complexities is crucial as data acquisition and hardware technologies improve. In ACAs, we see a step towards answering this question. This paper gives an introduction to the concept of ACA and its potential, showing a considerable speed-up compared to conventional 3D-discrete modeling for very complex scenarios: an order of magnitude in models with solid-fluid domains, bathymetry and discontinuous salt bodies. This is significant, as it suggests that ACA fares favorably across a large range of scales, from the deep mantle to exploration geophysics. Moreover, ACA can be further accelerated by smoothing, by considering only first arrivals and, potentially, by moving to higher frequency regimes.

ACA works best in a very particular type of geometry: when a source is in the center of a grid of receivers. This means that, for some common settings in exploration seismic, a bit of thought is required to use ACA to its fullest benefit. Here are a few examples of application strategies when using ACA:

- **Multiple sources:** It is common to run several consecutive simulations in the same model with different source locations. Usually, this is done by keeping the domain fixed and moving the source. When using an ACA, it is more efficient to adjust the modeling domain to each source position. This is trivial in AxiSEM3D, as the rotation of the domain to center on each source is done internally. Ideally, one would create a generalized complexity map

(e.g. simple layers informed by one test run) and a generalized 2D mesh to be used with all simulations. Using a fixed domain is also possible in AxiSEM3D which is fully capable of having off-axis sources and multiple sources (Szenicer et al., 2020). However, implementing sources in such a way reduces azimuthal smoothness.

- **Streamer geometry:** In most offshore seismic surveys, the source is placed at one end of a long streamer of receivers. In a source-centered ACA simulation, a large section of the cylindrical model domain would thus remain unused. The azimuthally fixed ACA method shown here is not a natural choice in such a set-up, although there are workarounds. One approach is to combine several simulations into one by arranging them around the source in a cake-slice set-up (Nissen-Meyer et al., 2019). Other ACA methods which are currently under development will solve this problem with the use of localized transforms. This will enable variable azimuthal sampling, allowing the downsampling and/or masking of unused parts of the azimuth.

- **Strong localized off-axis structures:** As demonstrated in the benchmark example, the benefit of ACA is not diminished by off-axis structures. Extreme cases (e.g. a small singular structure far from the axis which is much more complex than the background) can be dealt with using an ACA-based injection technique introduced by Leng et al. (2020).

Another aspect to keep in mind when using ACA are the benefits and limitations that come with using a 2D mesh. On the one hand, this makes meshing trivial and it saves memory. On the other hand, 3D discontinuities are honored by using particle relabeling, and this approach is limited to continuous interfaces. The effect of not honoring discontinuities was demonstrated in the Benchmark example. It must be noted, however, that meshing complex 3D structures is very challenging, and this poses a problem for standard 3D methods as well. A method to get accurate results without rigorous meshing is model homogenization (Capdeville

et al., 2013), and this should be considered when modeling 3D structures with steep discontinuities with an ACA.

As a concept, ACA is still a relatively under-explored. In the following, we outline a few avenues for future research. First, we need to gain a deeper understanding of how certain features in complexity maps arise. This is important to predict errors that may be induced by using approximate complexity maps and, ultimately, in order to predict complexity maps without running a high-resolution test simulation.

Further, we found that the degrees of freedom of the optimized simulation scaled with period to the fourth power (accounting for the change in time step) across our specific range of tested frequencies. This is much like a standard 3D modeling method. However, Leng et al. (2019) show that scaling with period decreases to the cubed period when approaching the higher-frequency range for a given structural model. Translated to our local-scale complex settings, one would then expect a substantially larger speed-up for ever higher-frequency regimes. It remains to be shown for the kinds of models presented here, but will certainly be of interest in future applications especially when computing hardware allows for routine simulations and inversions in these regimes.

Equally, the other end of the spectrum remains to be explored. We have shown that wavefields retain a degree of azimuthal smoothness even in the presence of irregular salt bodies. Nevertheless, it is expected that the ACA idea of exploiting smoothness will eventually break down for models which cause too much off-plane scattering. This could be the case for asteroids, or inside strongly topographic enclosures such as some volcanoes. Finding such counter indications is important for fully describing the ACA. It should be noted that, in such cases, one can still use the ACA solver for accurate simulations at full azimuthal resolution, with the caveat that the wavenumber-physical space transformations create a slight overhead in computational cost compared to conventional 3D meshing.

The practical use of the quantitative information contained in azimuthal complexity maps is another avenue of research. Potential applications might be: illumination studies where relative amplitudes are taken into consideration; the quantification

of uncertainties for different levels of 3D wavefield complexity; assessing intrinsic "invisibility" of certain structures; the optimization of misfit functions for full waveform inversion. Finally, a full comparison between state-of-the art FD and SEM methods with and without ACA may benefit the seismic community to choose optimal packages for any given problem.

## 5.5 Conclusions

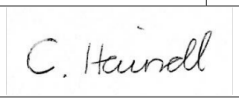
We have presented a new ACA-based seismic modeling method for complex 3D models, which can adaptively reduce the degrees of freedom and thereby computational cost by exploiting azimuthal wavefield smoothness while retaining high accuracy. The method is capable of handling models including solid and fluid domains and it has been shown to work well with a combination of absorbing boundary methods. We have benchmarked our ACA solver against an established 3D solver, obtaining an excellent fit between the outputs provided the models are implemented equally. We achieved a relative speed-up of an order of magnitude as well as significant savings in working memory with the ACA solver. We introduced azimuthal complexity maps, which are created as part of the complexity adaptation process, and we have shown how these maps may be interpreted in terms of seismic illumination and resolution. Finally, we verified the method on the SEG/EAGE salt model, quantifying the relative level of azimuthal complexity, representative of off-plane 3D scattering to be expected in highly heterogeneous models. We showed how the performance of the ACA solver increases with model smoothness. For the original model, the DOF are reduced by 92% compared to a standard 3D simulation when using ACA. In a direct comparison between an ACA-SEM and a standard SEM, this results in a cost reduction of 86%. When the model is smoothed, the cost savings increase to 95% in terms of DOF and 93% in terms of runtime. More research needs to be done to fully understand the relationship of model smoothness and ensuing wavefield smoothness.

**Statement of Authorship for joint/multi-authored papers for PGR thesis**
To appear at the end of each thesis chapter submitted as an article/paper

The statement shall describe the candidate's and co-authors' independent research contributions in the thesis publications. For each publication there should exist a complete statement that is to be filled out and signed by the candidate and supervisor **(only required where there isn't already a statement of contribution within the paper itself).**

| | |
|---|---|
| Title of Paper | A 3D Complexity-Adaptive Approach to Explore Sparsity in Elastic Wave Propagation |
| Publication Status | Published |
| Publication Details | Haindl, C.; Leng, K. & Nissen-Meyer, T. (2021). "A 3D Complexity-Adaptive Approach to Explore Sparsity in Elastic Wave Propagation". In: Geophysics, 86(5). |

Student Confirmation

| | | | |
|---|---|---|---|
| Student Name | Claudia Haindl | | |
| Contribution to the Paper | Developed absorbing boundaries and multi-peak wavefield scanning, performed simulations, analysed results, wrote manuscript | | |
| Signature |  | Date | 30/12/2021 |

Supervisor Confirmation

By signing the Statement of Authorship, you are certifying that the candidate made a substantial contribution to the publication, and that the description described above is accurate.

| | | |
|---|---|---|
| Supervisor name and title: Prof. Dr. Tarje Nissen-Meyer | | |
| Supervisor comments: Candidate made all relevant contributions to the paper | | |
| Signature  | Date | Jan 4, 2022 |

This completed form should be included in the thesis, at the end of the relevant chapter.

*The goal is to turn data into information, and information into insight.*

— Carly Fiorina

# 6 Wavefield Scanning in 3D

## Contents


**6.1 Introduction . . . . . 70**
**6.2 Theory . . . . . 71**
6.2.1 Wavefield Scanning in 2D . . . . . 71
6.2.2 Wavefield Scanning in 3D . . . . . 73
**6.3 A First Look at 3D Scanning Results . . . . . 75**
**6.4 Comparing the Effect of an On-Axis and an Off-Axis Structure on Wavefield Complexity . . . . . 77**
**6.5 3D Wavefield Scanning of the Homogenised Overthrust Model . . . . . 80**
**6.6 Conclusions . . . . . 84**


## 6.1 Introduction

Wavefield scanning is the central building block of ACA. It ensures that the efficiency of a simulation is optimised with respect to the desired accuracy. As seen in the previous chapter, the complexity maps created by wavefield scanning may offer new ways to characterise and analyse wavefields. However, these complexity maps are collapsed both in time and azimuth and, consequently, some potentially valuable information is lost. It is quite simple to resolve complexity maps on the time axis, but there are no obvious applications for this type of data. Azimuthal resolution,

on the other hand, is not as straightforward, but it would answer many questions, such as: Which parts of a given model give rise to the highest wavefield complexity? What does this tell us about the wavefield? How severe is the inefficiency induced by having a fixed azimuthal resolution for a given model? In this chapter, I tackle these questions using a wavelet transform to take a look at the azimuthal variations off-plane wavefield complexity.

## 6.2 Theory

### 6.2.1 Wavefield Scanning in 2D

2D complexity maps can be conceptualised as maps of the number of azimuthal sampling points required to resolve the most complex peak in the wavefield. But this definition does not fully encapsulate everything that goes into wavefield scanning. In the following, the wavefield scanning process is described thoroughly to create the basis from which to develop the 3D scanning method.

Recalling the Fourier expansion of the displacement given in Equation 3.11, we express the the Fourier coefficients in terms of a vector:

$$U_i(s,z;t_m) = \begin{pmatrix} u_i^0(s,z;t_m) \\ \vdots \\ u_i^{n_u}(s,z;t_m). \end{pmatrix} \tag{6.1}$$

$u_i^\alpha(s,z;t_m)$ denotes the $\alpha^{\text{th}}$ Fourier coefficient of the displacement in the $i^{\text{th}}$ dimension. $(s,z)$ are the in-plane coordinates, $t_m$ is the current time step. When scanning, an algorithm identifies whether the wavefield at the current time step qualifies as one of the largest "peaks". If so, the current time step $t_m$ is added to a list of potential candidates, $t_{\text{candidates}}$ from which the most complex peak is selected at the end of the simulation. In an elastic medium this is done separately for all dimensions, as denoted by the subscript $i$. The exact process is visualised as a flowchart in Figure 6.1.

There are three user defined parameters for the scanning process: The number of peaks which are considered, $K$, and the absolute and relative tolerances, $\epsilon_{\text{abs}}$ and $\epsilon_{\text{rel}}$, respectively. These parameters provide different ways to control the overall

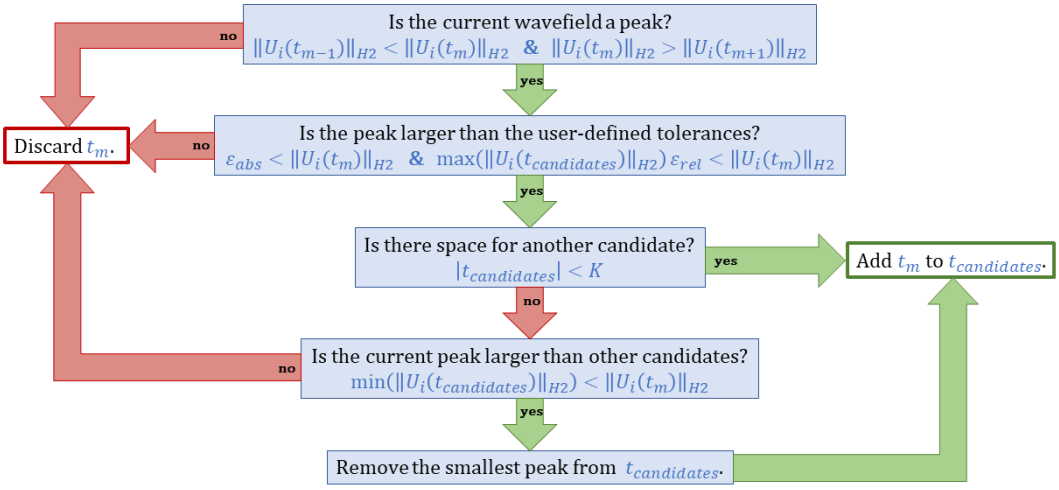



**Figure 6.1:** Flowchart explaining the process of peak selection and subsequent wavefield scanning. $U_i(t_m)$ is the azimuthal Fourier expansion of the displacement in a given dimension, $i$, and at a given in-plane coordinate $(s, z)$. Note that $(s, z)$ was dropped in the flow chart for brevity. $t_m$ indicates the current time step and $t_{\text{candidates}}$ is a collection of no more than $K$ time steps where the wavefield satisfies the user-defined criteria for a wavefield peak. $\epsilon_{\text{abs}}$ and $\epsilon_{\text{rel}}$ are user-defined tolerances.

accuracy and noise tolerance of the algorithm. Figure 6.1 shows how they play out in the scanning process.

For each of the resulting candidate peaks, the minimum number of Fourier coefficients, $n_u$ (introduced in equation 3.11), is sought with which the wavefield can still be accurately represented:

$$\min(n_u), \text{ such that } \frac{\left\|\sum_{|\alpha|>0}^{n_u} u_i^{\alpha}(s,z;t_{\text{candidates}})\right\|_2 - \|u_i(s,z;t_{\text{candidates}})\|_2}{\|u_i(s,z;t_{\text{candidates}})\|_2} < \epsilon_{\text{tol}}. \tag{6.2}$$

The error tolerance, $\epsilon_{\text{tol}}$, is also defined by the user input. At the end of the test simulation, the required number of sampling points, $n_r$, can be determined from the required number of Fourier coefficients using the Nyquist sampling criterion,

$$n_r = 2n_u + 1. \tag{6.3}$$

The user inputs allow for optimisation based on the use-case. The effect of using different values for $K$ were demonstrated in the previous chapter. $\epsilon_{\text{rel}}$ is

an attempt to tackle the same problem but in a less data-driven way. The rapid decrease in amplitude of the propagating direct wave relative to other phases makes it all but impossible to obtain a clean complexity map using this criterion instead of $K$. Nevertheless, $\epsilon_{\text{rel}}$ is useful if there is an external reason to only consider certain amplitudes. $\epsilon_{\text{abs}}$ is used to avoid over-fitting in the presence of numerical noise. Finally, $\epsilon_{\text{tol}}$ is the parameter users can change to directly manage the efficiency-accuracy trade-off.

### 6.2.2 Wavefield Scanning in 3D

There is no intrinsic spatial resolution in Fourier domain, so a different transform is needed for a 3D wavefield scanning method. The wavelet transform is an obvious choice because it covers a range of wavenumber scales, and thus it requires no prior knowledge of the wavefield complexity to optimise the scanning result. The choice of the wavelet was informed by the fact that the basis functions for the Fourier transform are sinusoids. To keep the 3D scanning results comparable to the 2D maps used for the ACA, a real-valued Morlet is chosen, i.e. a sinusoid modulated by a Gaussian.

There are two approaches to the wavelet transform: discrete (DWT) and continuous (CWT). There are upsides and downsides to both transforms: DWT is cheaper and sparse whereas the CWT is shift-invariant and has a better resolution. The Morlet does not meet the criteria for a DWT, so this choice is being made for us.

In the CWT, the wavefield is convolved by a series of wavelets which vary in one scaling parameter, $\sigma$. This effectively adds an additional dimension to the wavefield representation. In the case of the Morlet transform of wavefield along the azimuth, this can be expressed as:

$$u_i(s,z;\phi,\sigma;t) = \sqrt{\frac{2\sqrt{\pi}}{\sigma}} u(s,z;\phi;t) * e^{-\frac{\phi^2}{2\sigma^2}} \left( \cos(\frac{\kappa\phi}{\sigma}) - e^{-\frac{1}{2}\kappa^2} \right), \qquad (6.4)$$

where the star denotes a convolution and $\kappa$ is a fixed wavelet parameter which influences the trade-off between spatial and wavenumber resolution. Higher values of $\kappa$ result in a better wavenumber resolution and lower values result in a better spatial

resolution, respectively. $\kappa$ should not be smaller than 5, since several approximations (including Eq. 6.5) are not valid at smaller numbers.

The dominant frequency (or in our case more accurately the dominant wavenumber) of the Morlet can be approximated from the term inside the cosine function, i.e.

$$k_c = \frac{\kappa}{\sigma}. \tag{6.5}$$

My goal is for the results of the CWT to be as comparable to the Fourier transform as possible. Ideally, $u_i(s, z; \phi, \sigma; t)$ should be discretised into $u_i^{\beta}(s, z; \phi; t)$ as an azimuthally localised version of $u_i^{\alpha}(s, z; t)$, as seen in Equation 3.11. This informs my choice for $\kappa$ and $\sigma$: if $k_c$ is to go from 1 to $n_u$ in steps of 1, a reasonable parametrisation would be $\kappa = 2\pi$ and $\sigma = \frac{2\pi}{\beta}$, thus making $\alpha$ and $\beta$ somewhat equivalent.

The CWT produces a highly redundant representation of the wavefield. In most simulations, saving the wavefield in wavelet domain, i.e. as $u_i^{\beta}(s, z; \phi; t)$, would very quickly exhaust the available memory. So it is important to collapse this representation in a way similar to the 2D scanning algorithm, while leaving the azimuthal resolution intact. In fact, the 2D scanning algorithm explained in the previous section can be applied directly to the wavefield in wavelet domain. But rather than having one scanning operator for one given point $(s, z)$ we now have $n_r$ scanning operators for one given point $(s, z)$: one for each azimuthal sampling point $(s, z; \phi)$. This produces a 3D complexity map which azimuthally localised estimates for the required number of azimuthal sampling points.

Of course, 2D scanning and 3D scanning are not perfectly equivalent. The Fourier transform can be pictured as the zero-spatial-resolution end-member of the Morlet-based CWT. By increasing the spatial resolution, the CWT introduces uncertainty in the wavenumber domain which results in a smearing effect. Consequently, 3D scanning results likely overestimate the required resolution. Also, some ringing which stems from the convolution with the Morlet is to be expected on the 3D complexity maps.

## 6.3 A First Look at 3D Scanning Results

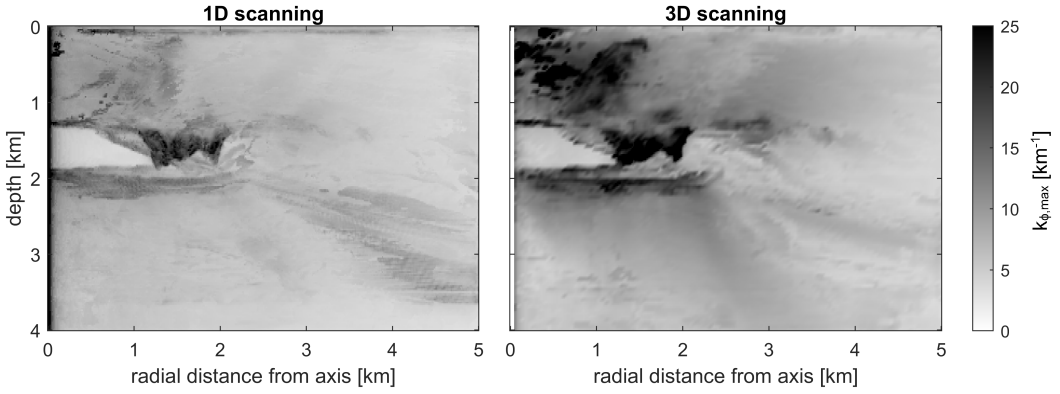



**Figure 6.2:** Comparison of 2D and 3D wavefield scanning results for the SEG salt model. Colour scale shows maximum azimuthal wavenumber in km$^{-1}$. 10 peaks were used for scanning. **Left:** Standard 2D complexity map. **Right:** Azimuthal maximum of the 3D complexity.

To test the new scanning method, I use a familiar example: Figure 6.2 shows the wavefield complexity from one of the simulations seen in the previous chapter, in Section 5.3. The set-up is identical, with an explosive source at 1 km depth and 5 s simulation time. On the left panel, we see the familiar complexity map used to optimise ACA simulations. The right panel shows the result of 3D scanning which was collapsed along the azimuth by taking the maximum value. This results in a 2D map comparable to the maps created by the standard scanning method. Note that the resolution of the 3D scanning result was lowered on purpose to reduce the size of the dataset. As expected, the 3D scanning method overestimates the complexity of the wavefield, but the order of magnitude of the two results matches as do the structures seen in the two maps. This comparison confirms that 2D scanning results and 3D scanning results indeed measure the same thing.

With the 3D scanning results, we can visualise the complexity in the whole 3D volume of the modelled domain, seeing which parts of the azimuth give rise to the features seen in the complexity map. Figure 6.3 illustrates this by showing three cross sections through the wavefield complexity.

A striking feature of these cross sections is the relatively low complexity inside the salt body, which can be explained by the high seismic velocity of salt. The

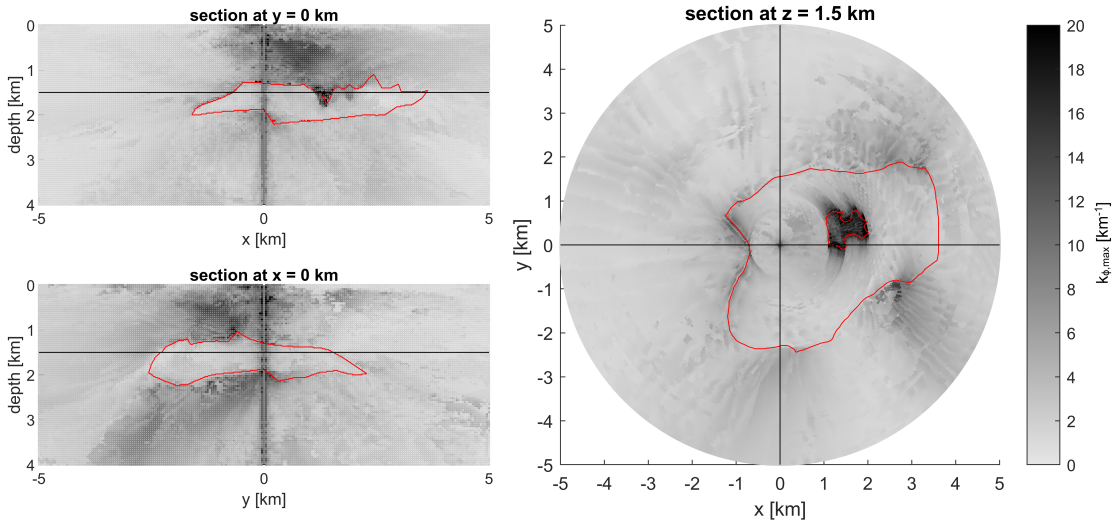



**Figure 6.3:** Cross sections of the 3D wavefield scanning result for the SEG salt model. Colour scale shows maximum azimuthal wavenumber in km$^{-1}$. Red lines outline the position of the salt body in the respective cross section. Black lines indicate positions of the other cross sections.

highest complexity, on the other hand, is found above the faulted sections of the salt body. I was able to infer this without 3D scanning by juxtaposing the salt outline on the complexity map in Figure 5.9. But the highly localised nature of the complexity only becomes clear now.

On a different note, the horizontal cross section in Figure 6.3 contains the anticipated artefacts from the wavelet transform. 1. The smearing effect of the wavelet transform is apparent around the hole in the salt body. The maximum wavenumber drops instantly at the radial boundary of this hole, but at the azimuthal boundaries the wavenumber decreases slowly. This stark contrast in behaviour around one structure indicates that this is not a real property of the wavefield. Instead, the high wavenumbers are probably actually contained to the low velocity medium, and the high wavenumbers seen in the salt are an artefact. 2. Examples of ringing are visible as near-radial stripes in several places around the horizontal cross section.

## 6.4 Comparing the Effect of an On-Axis and an Off-Axis Structure on Wavefield Complexity

The example above is the same exact set-up used for the complex example in Haindl et al. (2021). During the review process, one question which was asked repeatedly was: Having the source centred on the salt body is an ideal case for ACA - but how would it fare if the salt was off-axis? In this section, I investigate this very question by combining information available from standard wavefield scanning and 3D wavefield scanning.

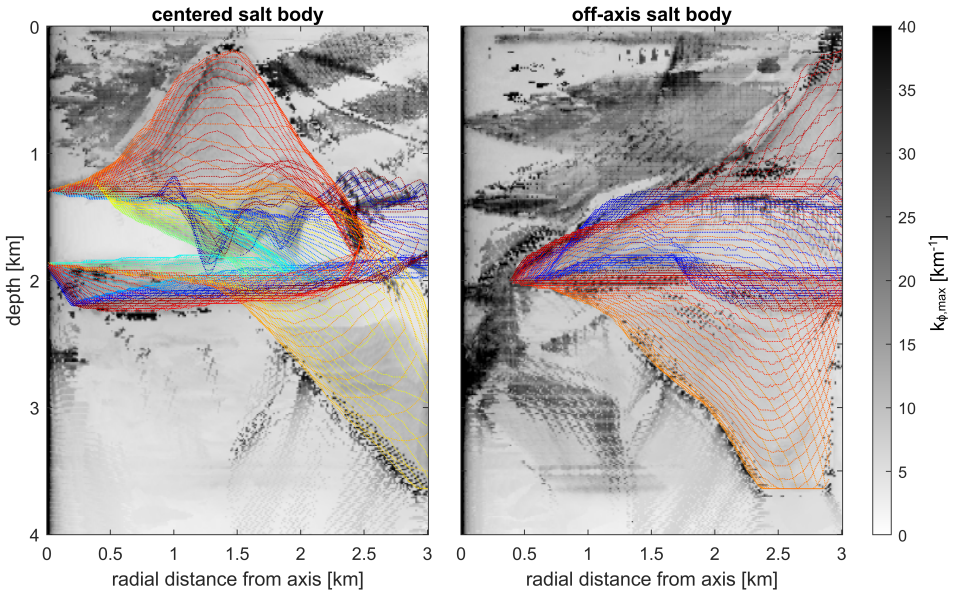



**Figure 6.4:** Complexity maps of the truncated SEG salt model before and after shifting it 2 km in the positive y-direction. Dotted lines show the outline of the salt body with colours indicating the angle from the axis.

For the following tests, the SEG salt model is truncated at 3 km distance for the axis to allow for a significant shift of the salt body. Figure 6.4 shows complexity maps of the truncated domain and the truncated domain where the model has been shifted in positive y-direction by 2 km. These maps were created with the standard learning method and the same learning parameters as the example in the previous section.

Surprisingly, despite all parameters being equal, the complexity maps of the smaller domain have more high azimuthal wavenumber regions than the complexity map shown in Figure 6.2. It has been observed in Chapter 4, that absorbing

boundaries may contribute to numerical noise, and it seems likely that this is the reason for the observed differences.

Comparing the two panels in Figure 6.2, the biggest increase in azimuthal complexity can be observed near the axis. This high complexity region likely represents the part of the wavefield which is scattered back from the salt surface towards the axis. However, the cost of the ACA is affected more strongly by high complexity if it is far away from the axis. In the region near the outer edge of the model, no significant increase in complexity is visible. The average for the maximum azimuthal wavenumber, $k_\phi$, is 11.0 km$^{-1}$ for the model with source-centred salt and 13.2 km$^{-1}$ for the model with off-axis salt. For an optimised ACA run, the shifted model would require 12% more DOF than the source-centred model.

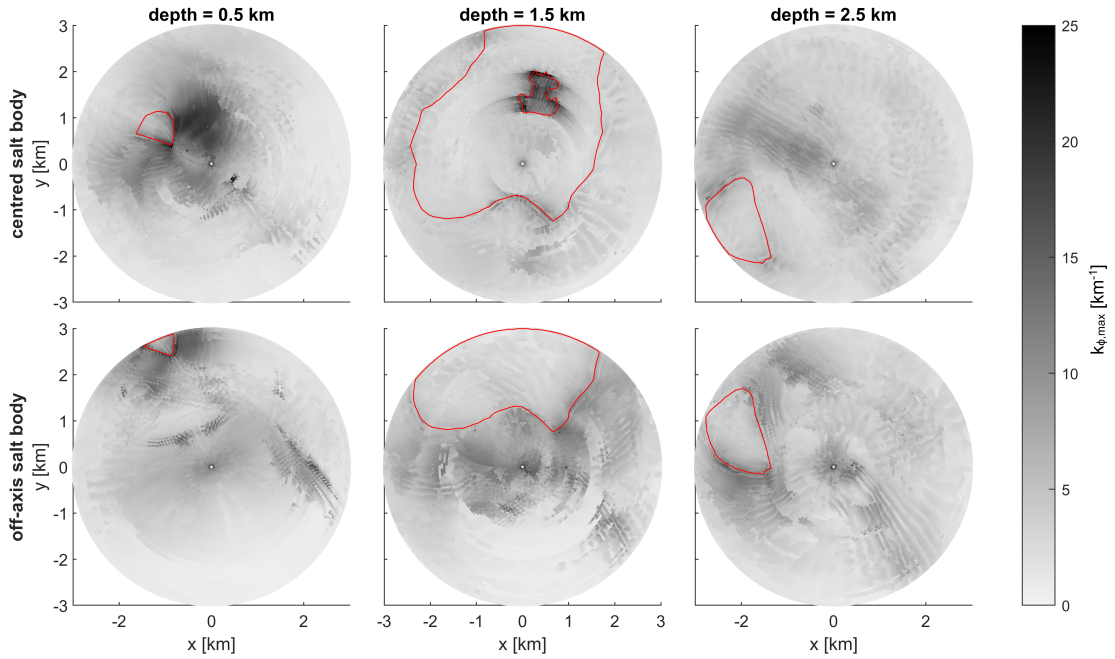



**Figure 6.5:** Comparison of horizontal cross sections of the 3D wavefield complexity in the SEG salt model. Top row shows results for the truncated domain and bottom row shows results for the shifted truncated domain. Red line indicates position of the salt body at the respective depth.

Figure 6.5 shows three horizontal cross sections of the 3D wavefield scanning results for both models. The wavefield complexity tends to be highest along the outside of the salt body. Other than the predictable shift of these regions with the shift of the salt body, there seem to be no additional effects from having an

off-plane structure in an ACA set up. It appears that, because the salt body in the SEG model is so irregular, the wavefield complexity is fairly high regardless of the position of the salt.

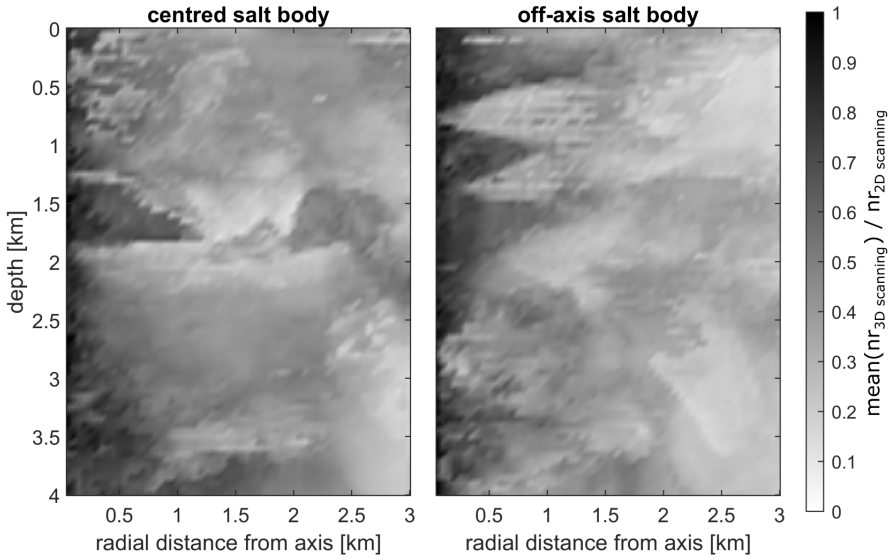



**Figure 6.6:** Map of the average ACA efficiency factor for the truncated SEG salt model and the shifted truncated SEG salt model.

With the added azimuthal resolution from 3D scanning, we can introduce another measure to evaluate the performance of ACA: We can estimate the azimuthal sampling rate which is required locally and compare it to the maximum azimuthal sampling rate which would be used in an ACA optimised simulation. This gives a ratio between 0 and 1, where 1 indicates perfect efficiency and a lower number indicates that irregular azimuthal sampling would be beneficial. It should be noted that, in a set-up with a cylindrical mesh where azimuthal sampling can be varied, limitations in 3D meshing would likely also result in an efficiency well below 1. Figure 6.6 shows maps of the average ACA efficiency for the two salt models. Both models have a relatively high efficiency near the axis. As would be expected, the efficiency for the model with the shifted salt body drops off towards the outer edge of the model. This does not happen to the same degree in the case with the centred salt. This is reflected in the overall average over the 3D efficiency factors, which is 0.40 for the centred salt model and 0.33 for the off-axis salt model, respectively.

## 6.5 3D Wavefield Scanning of the Homogenised Overthrust Model

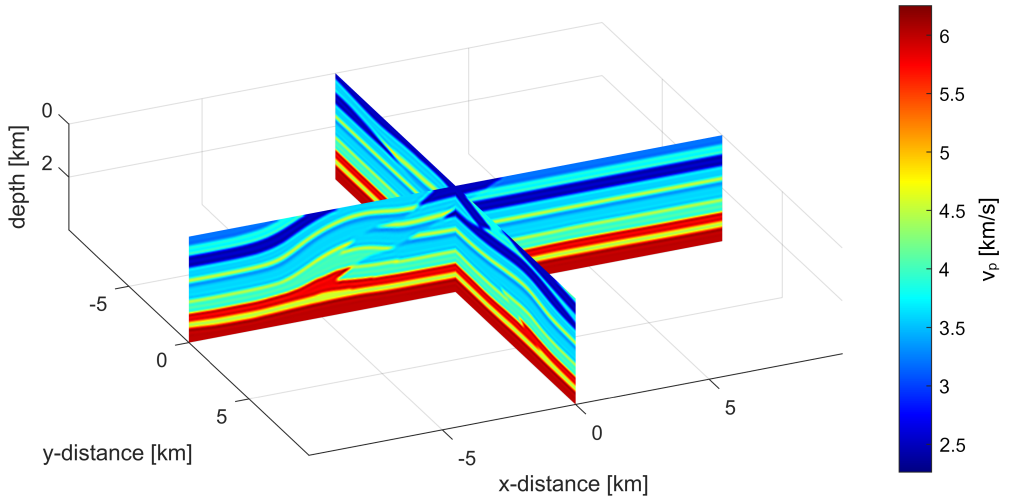



**Figure 6.7:** Vetrical cross-sections through the SEG overthrust model showing P-wave velocity in m/s.

Apart from the SEG salt model, the other open-access velocity model commonly used for testing methods on complex media is the SEG overthrust model (Figure 6.7). Yann Capdeville & Paul Cupillard graciously offered their homogenised version of the SEG overthrust model for our tests. As discussed in the Benchmark example in Section 5.2.3, implementing a complex 3D model via a volumetric grid leads to significant errors because this approach does not honour sharp discontinuities. Homogenisation is a proven method to improve results from volumetric grid models (Capdeville et al., 2010; Cupillard and Capdeville, 2018; Hedjazian et al., 2021). On the flip-side, homogenisation introduces full anisotropy, the effects of which on wavefield complexity are investigated in this section.

Figure 6.8 shows slices through the homogenised overthrust model as well as the respective Zener anisotropy ratios, $a_r$, defined as

$$a_r = \frac{2C_{44}}{C_{11} - C12}, \tag{6.6}$$

where $C_{ij}$ are elements of the elasticity tensor in Voigt notation.

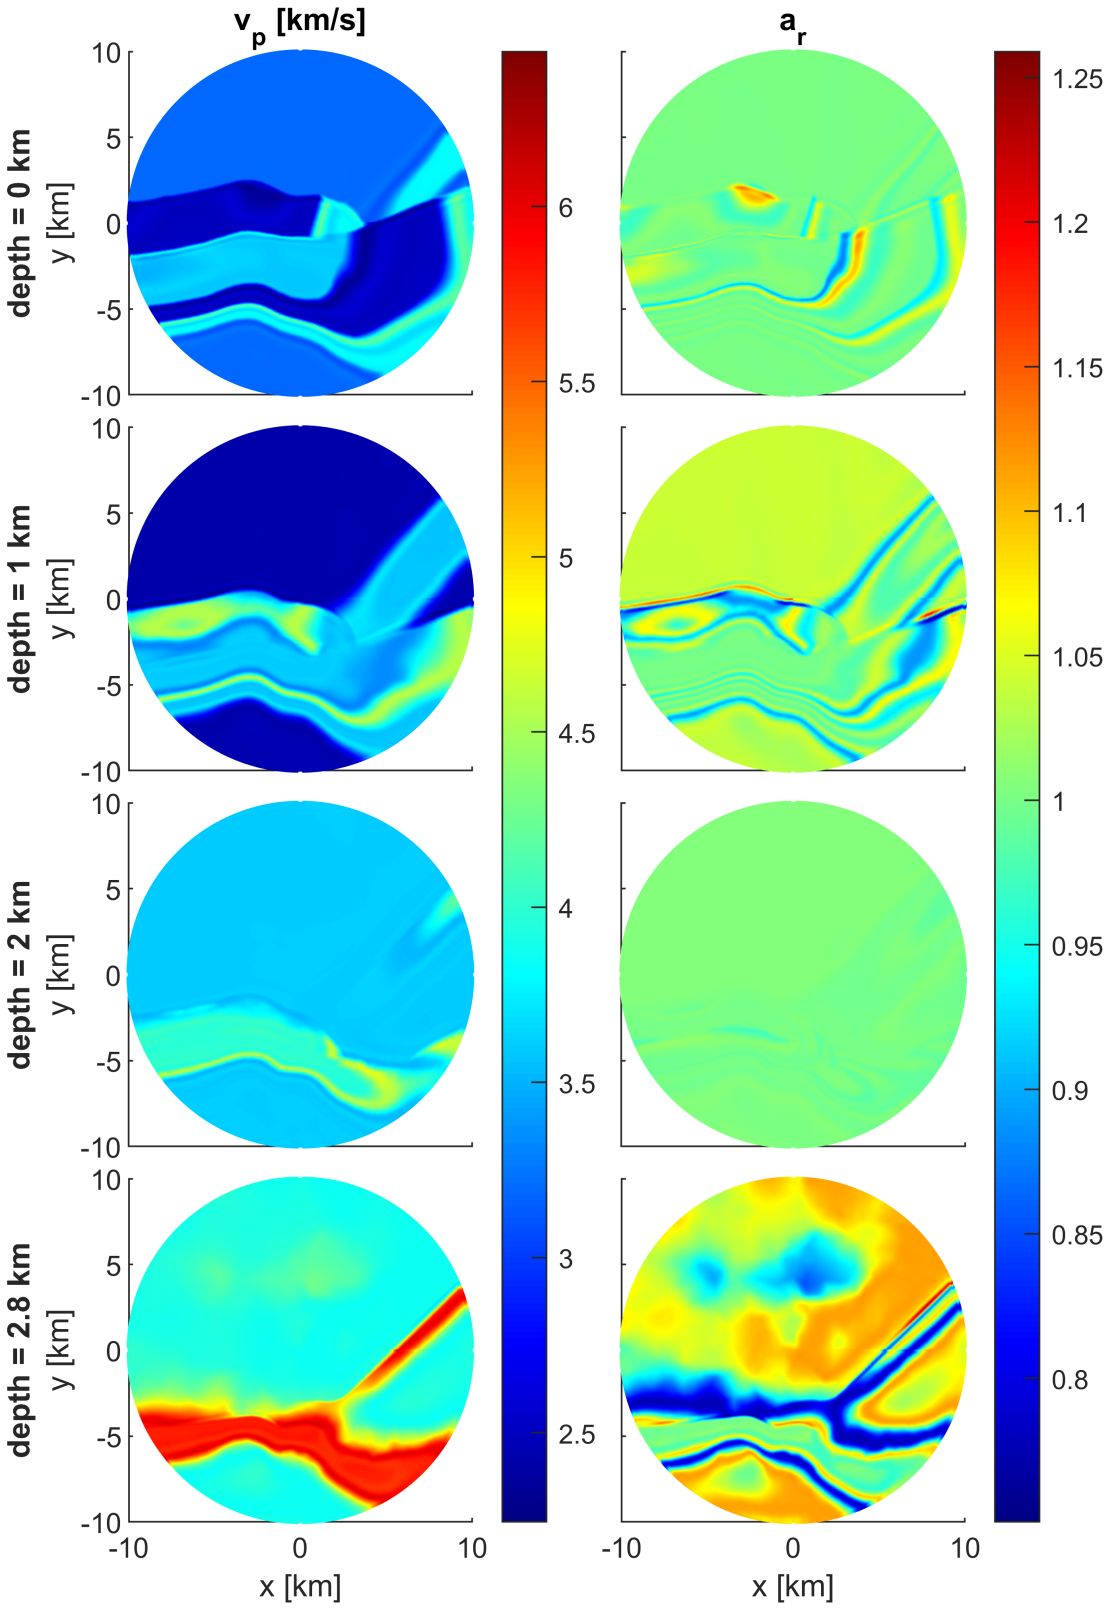



**Figure 6.8:** Horizontal cross sections through the homogenised SEG overthrust model. Panels on the left show average P-wave velocity in km/s. Right panels show Zener ratio as a measure for anisotropy. A value of 1 means the medium is isotropic.

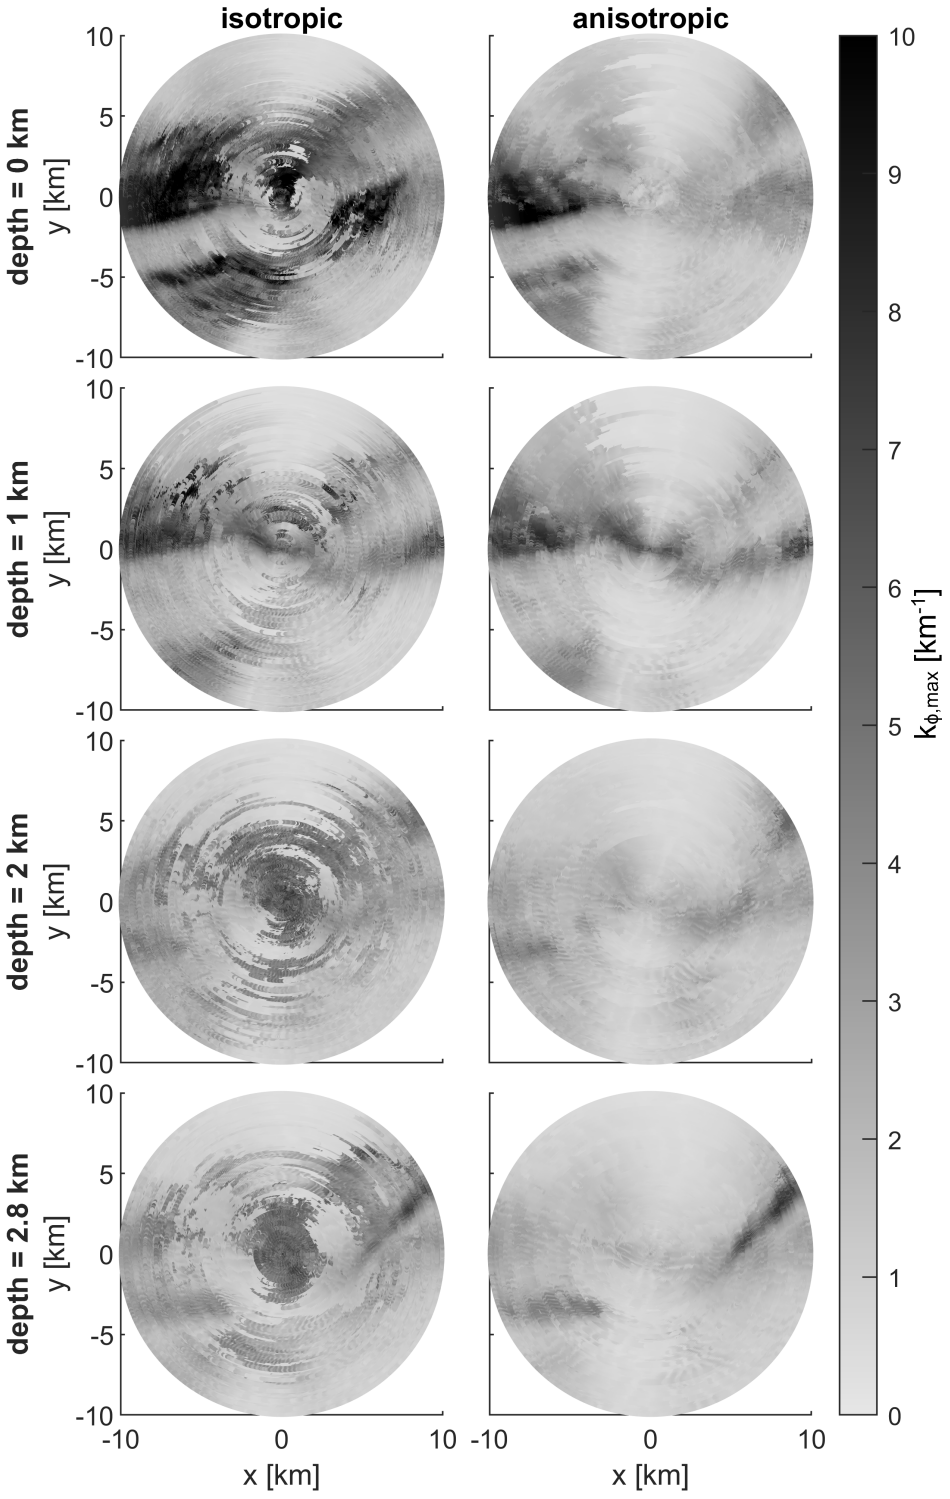



**Figure 6.9:** Horizontal cross sections through the maximum azimuthal wavenumbers (in km$^{-1}$) recorded in the homogenised overthrust model. Left panels show results for an isotropic version of the model. Right panels show results for the anisotropic version. 5 peaks were used in the scanning algorithm.

The 3D wavefield complexity for each slice can be found in Figure 6.9. On the left side of this figure, the results for an isotropic version of the homogenised model are shown. The results shown on the right side are obtained using full anisotropy.

I will start by pointing out some general observations from the anisotropic run, before going into a comparison of the two runs. Overall, there is a repetition of the trend observed in the previous examples: azimuthal complexity is relatively low in high-velocity regions, and the highest values are to be found on the low-velocity sides of strong velocity contrasts. Also, the abruptness of the change in wavespeed seems to increase the complexity more than the absolute velocity contrast. One example of this can be seen at 0 km depth: on the left side of the model, contacts between the slow and fast media are sharper than on the right side, despite the right side containing some higher velocity contrasts. On the complexity maps, we see much higher values on the left side than on the right. At 2.8 km depth is a similar example: on the left side, the contact between the fast and the slow medium is smeared whereas on the right side it is quite abrupt. This is reflected in the complexity, which is higher near the right contact than near the left one.

Clearly, the addition of anisotropy does not increase the azimuthal complexity of the wavefield. In fact, in the case shown here, it is doing quite the opposite. Apart from containing higher azimuthal wavenumbers than the anisotropic equivalent, the result from the isotropic model has a very patchy appearance. Upon closer inspection, the anisotropic model is not completely devoid of patches either, with visible examples seen in the 0 km and 1 km slices, near x = 0 km along the positive y axis. By varying the different scanning parameters, I found that this patchiness is a result of the arbitrary number of "peaks" for wavefield learning (5 in both examples shown here).

To investigate this further, a series of 3D scanning results from the isotropic model, using different peaks is compared in Figure 6.10. Apparently, there is a gradual increase in absolute complexity as more peaks are considered. But, in addition to this, there are patches with high azimuthal wavenumbers which start out small and increase in size. It stands to reason that peaks which include a

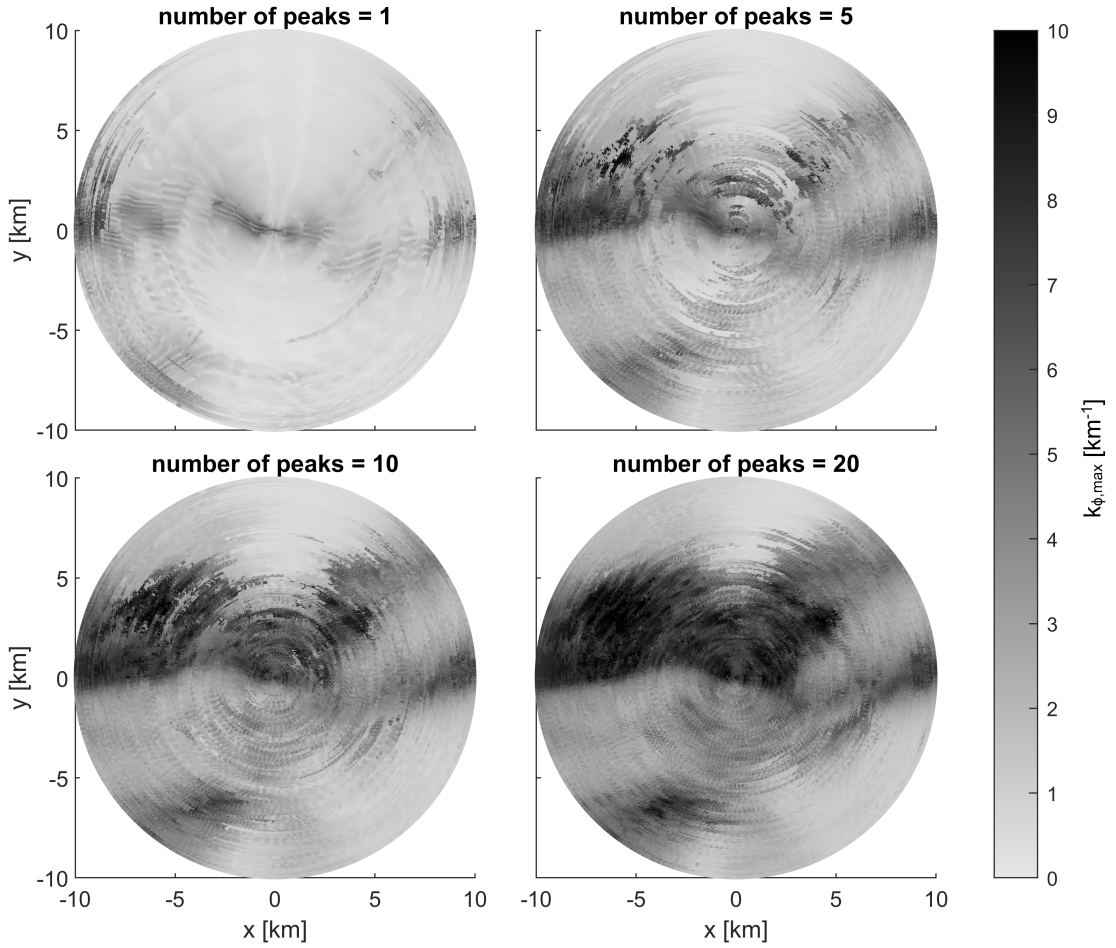



**Figure 6.10:** Comparison of horizontal cross sections through 3D scanning results of the isotropic homogenised overthrust model at 0 km depth, where different numbers of candidate peaks were used for scanning.

particularly complex phase are large enough to become candidates in some parts of the model but not in others. This creates artefacts which are purely down to the peak selection process and do not reflect the true complexity of the wavefield. Luckily for the purpose of ACA, this effect is mostly evened out by collapsing the azimuth, and thus it is very rare to observe in 2D scanning.

## 6.6 Conclusions

3D wavefield scanning is a new tool to investigate azimuthal wavefield complexity. By looking at 3D scanning results for different models, I was able to identify the most common source for high wavefield complexity: the medium with the lowest

wavespeed in the direct vicinity of a sharp velocity contrast. As seen in the SEG salt example, it is possible for such wavefield complexity to radiate outwards away from the velocity contrast, but this appears to be limited to quite extreme cases. Further, I found that neither shifting structures off-axis nor including full anisotropy significantly increases wavefield complexity.

The 3D complexity maps also offer a measure for the inefficiency caused by having a fixed azimuthal resolution. While shifting a complex structure off-axis does not increase the wavefield complexity, it does increase this inefficiency. In the case of the SEG salt model, shifting the salt 2 km to the side resulted in a 17.5% drop in efficiency compared to the centred salt.

There appears to be no obvious way to predict a complexity map of a slightly shifted model from the complexity map of the original model. Nevertheless, complexity maps clearly reflect the underlying velocity model in a systematic manner. Thus, I see potential in machine learning techniques for the prediction of complexity maps.

*Progress in science is often built on wrong theories that are later corrected. It is better to be wrong than to be vague.*

— Freeman Dyson.

# 7

# Moving Beyond the Limitation of Azimuthal Continuity

## Contents


**7.1 Introduction . . . . . 86**
**7.2 The Short-Time Fourier Transform . . . . . 87**
7.2.1 Window Functions . . . . . 89
7.2.2 Proof of Concept . . . . . 90
**7.3 Handling Non-Periodic Fourier Derivatives . . . . . 91**
**7.4 Interpolation . . . . . 93**
**7.5 Changes to the Method Architecture of AxiSEM3D . 95**
**7.6 Numerical Instability . . . . . 96**
**7.7 A Way Forward . . . . . 97**


## 7.1 Introduction

A drawback of ACA methods is the strict requirement for azimuthal continuity. 3D Models are implemented by what effectively amounts to a convolution in Fourier domain, but this approach has distinct limitations: Solid and fluid media must be continuous around the azimuth and so does the azimuthal sampling rate. Consequently, it is impossible to model realistic coastlines or fluid inclusions, and it is inefficient to use ACA in models with azimuthally localised complexity. A

quantitative measure for this inefficiency has been developed in the previous chapter, in Section 6.4. In the following sections, an attempt to overcome these limitations is made. While the proposed method is not numerically stable, the ansatz may serve as a jumping off point for future attempts to tackle the same problem.

## 7.2 The Short-Time Fourier Transform

In order to allow variations in the azimuthal sampling rate, the Fourier transform in the ACA approach has to be replaced with a different transform. Initial ideas included the non-uniform Fourier transform (NFFT) and the Wavelet transform. However, neither of these transforms proved suitable. The NFFT allows for variation in azimuthal sampling rate but its frequency resolution is limited by the total number of sampling points. Hence, it cannot serve the purpose of locally up-sampling a complex region. In Wavelet domain, there is no simple method to calculate derivatives, and this defeats the point of the transform altogether. A feasible solution was found in a combination of the Short-Time Fourier transform (STFT) and interpolation between windows with different sampling rates.

The STFT is commonly used to analyse the spectra of transient signals, such as in sound processing. This is done by dividing the signal into regular time intervals, or windows, and Fourier transforming each window separately. Usually, the windows are multiplied by a window function prior to the transform to minimize edge effects. The concept of the STFT is explained in Figure 7.1. For a function, $f(t)$, the continuous STFT is defined as:

$$\mathrm{STFT}\{f(t)\}(\tau,\omega) = \int_{-\infty}^{\infty} f(t)w(t-\tau)e^{-i\omega t}\,dt, \tag{7.1}$$

with the notable differences from the Fourier transform being the addition of the window function, $w$, and the dependency on $\tau$, i.e. the position of the window function relative to the signal.

To replace the Fourier transform in ACA with an STFT, a few adaptations to the standard STFT process are required. Firstly, instead of time, it is the azimuth which is divided into windows. Secondly, the intervals should not be regular, but adapted

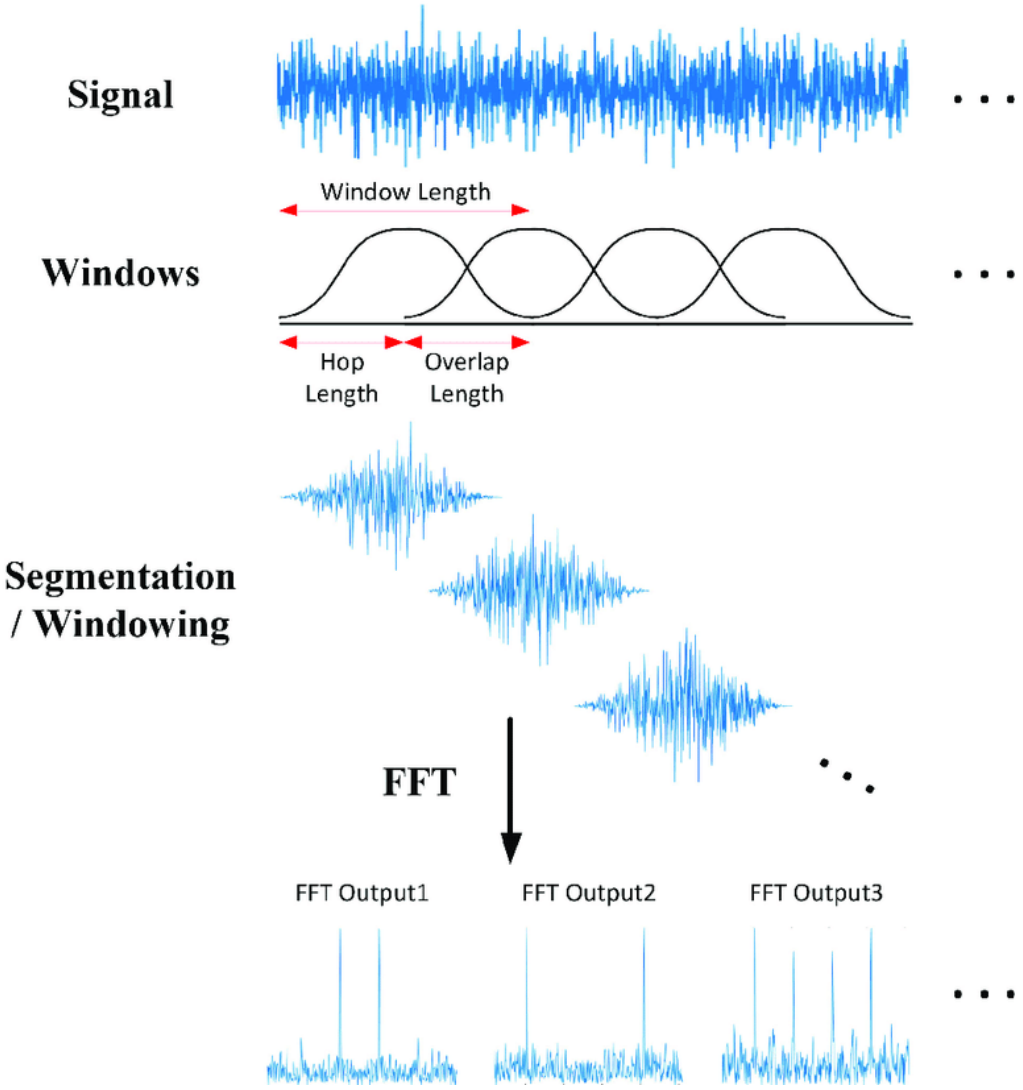



**Figure 7.1:** Graphical explanation of the STFT (from Jeon et al., 2020). The signal is multiplied by window functions and each windowed signal is Fourier transformed individually.

to the variations in azimuthal complexity. Third, an interpolation method is required to handle the overlap region of windows with different azimuthal sampling intervals.

The use of the STFT may open up other possibilities in addition to the changes in azimuthal sampling rate. Since the windows are fully decoupled, it may be possible to implement local solid/fluid regions by replacing the overlaps with solid-fluid boundaries between adjacent windows. Similarly, absorbing boundaries could be implemented on window edges and windows outside a desired azimuthal range could be masked, hence reducing the modelled volume to a wedge.

### 7.2.1 Window Functions

Multiplying a signal by a window function affects its Fourier spectrum. In fact, the spectrum of the resulting "windowed" function is simply the convolution of the spectrum of the signal and the spectrum of the window function. To keep the artefacts from this convolution as small as possible, window functions are designed to have spectra approximating a delta function. Figure 7.2 shows a few examples for window functions and their spectra.

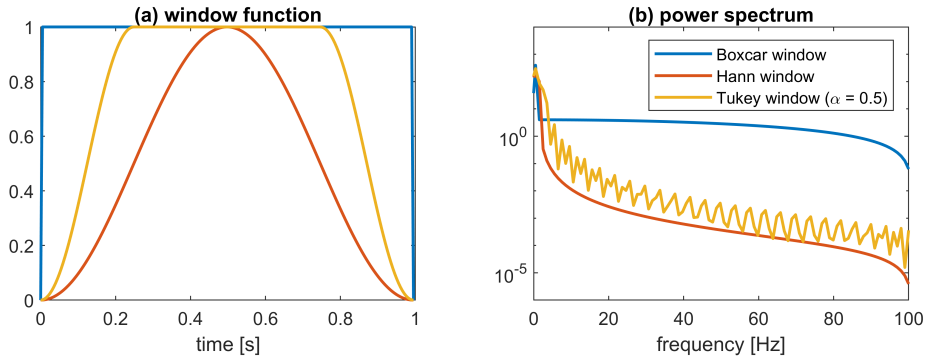



**Figure 7.2:** Examples of window functions (a) and their spectra (b). Ideally, a window function should have a spectrum similar to a delta function. Note that not using a window function has the equivalent effect of using a boxcar function.

In addition to minimising convolution artefacts, an appropriate window function for the windowed ACA must meet a number of criteria. Most importantly, it must allow an inverse STFT which is not the case for all window functions. In an inverse STFT the windows are recombined after the inverse Fourier transform, and this recombination is only accurate if the window function satisfies the overlap-and-add (OLA) criterion,

$$\sum_{m \leq M(s,z)} w_m(s, z; \phi) = 1, \tag{7.2}$$

where $M(s, z)$ is the number of windows used at a given point $(s, z)$ and $w_m(s, z; \phi)$ are the individual window functions at that point.

Most windows which satisfy the OLA criterion have a fixed overlap length equal to half or more of the window length. Such an example is shown in Figure 7.1, and the Hann window in Figure 7.2 also falls into this category. But using fixed-overlap

windows for the windowed ACA would be inefficient as the large overlaps would lead to high redundancy. An exception to this is the Tukey window, which has cosine tapered slopes of adjustable length as well as a central section without overlap. This offers a potential compromise between efficiency and mitigating the artefacts in frequency domain.

The Tukey window is defined as

$$w_m(x) = \begin{cases} \frac{1}{2}\left(1 - \cos\left(\frac{2\pi(x-x_0)}{\alpha L}\right)\right), & x_0 \leq x \leq x_0 + \frac{\alpha L}{2} \\ 1, & x_0 + \frac{\alpha L}{2} < x < (x_0 + L) - \frac{\alpha L}{2} \\ \frac{1}{2}\left(1 + \cos\left(\frac{2\pi(x-(x_0+L))}{\alpha L}\right)\right), & (x_0 + L) - \frac{\alpha L}{2} \leq x \leq x_0 + L \\ 0, & \text{otherwise,} \end{cases} \tag{7.3}$$

where $x_0$ is the starting position of the window, $L$ is its length and $\alpha$ is the fraction of the window which consists of the tapered slopes.

### 7.2.2 Proof of Concept

Replacing the Fourier transform with the STFT changes the equations governing the ACA method. To investigate the consequences of this, let us start by applying window functions $w_m(s, z; \phi)$ to separate the wavefield into $M$ overlapping segments:

$$\mathbf{u} = \sum_{m \leq M(s,z)} w_m(s, z; \phi)\mathbf{u}(s, \phi, z; t) \tag{7.4}$$

$$= \sum_{m \leq M(s,z)} \mathbf{u}_m(s, \phi, z; t). \tag{7.5}$$

This representation is exact provided the window functions satisfy OLA criterion. Just like in the standard ACA approach, these wavefield windows are then expressed in terms of a Fourier series:

$$\mathbf{u}_m(s, \phi, z; t) = \hat{\mathbf{g}}_i(\phi) \sum_{|\alpha| \leq n_u} \sum_{m \leq M(s,z)} u^{\alpha}_{m|i}(s, z; t)\Psi^{\alpha}(\phi), \tag{7.6}$$

where $u^{\alpha}_{m|i}(s, z) = w^{\alpha}_m(s, z) * u^{\alpha}_i(s, z)$ denotes the convolution of the wavenumber spectra of the window function and the wavefield over the azimuthal wavenumber $\alpha$.

Ideally, all operations previously performed on Fourier coefficients, $u^{\alpha}_i$, would now simply be performed on convolved Fourier coefficients, $u^{\alpha}_{m|i}$, instead, thus requiring minimal changes in the algorithm. However, we must ensure that artefacts

from the convolution are not propagated during time-stepping. Hence, we require that whenever the windows are transformed to physical space and recombined, all artefacts cancel. Recalling the method architecture of AxiSEM3D shown in Figure 3.4, such a recombination would occur twice in each time step (denoted F2P-FFT): once after time integration and a gradient calculation, and once after a divergence calculation.

The time integration in AxiSEM3D is done either using the Newmark or the Symplectic time scheme. In both methods, time integration consists of simple addition of fields multiplied by constants. Such operations can trivially be shown to be unaffected by the STFT. The derivatives, on the other hand, do give rise to errors in the individual wavefield windows. Conveniently, these errors cancel upon recombination of the windows. Consider the product rules for the gradient and divergence operators:

$$\nabla \mathbf{u}_m = (\nabla w_m)^{\mathbf{T}} \mathbf{u} + w_m \nabla \mathbf{u}, \tag{7.7}$$

$$\nabla \cdot \mathbf{u}_m = (\nabla w_m) \cdot \mathbf{u} + w_m \nabla \cdot \mathbf{u}. \tag{7.8}$$

In both cases an error term arises which involves the gradient of the window function, $\nabla w_m$. Since both inner and outer vector products are distributive, the OLA criterion ensures that these error terms cancel when the windows are recombined:

$$\sum_{m \leq M} (\nabla w_m) = \nabla \sum_{m \leq M} w_m = \nabla 1 = 0. \tag{7.9}$$

## 7.3 Handling Non-Periodic Fourier Derivatives

A major challenge with the STFT approach is calculating accurate Fourier derivatives. Part of the attraction of transforming the azimuthal dimension to the Fourier domain is that this simplifies the calculation of the derivative, as

$$\frac{d}{d\phi} FT(\mathbf{u}(\phi)) = i\alpha FT(\mathbf{u}(\phi)). \tag{7.10}$$

However, this Fourier derivative is only accurate for periodic functions. In standard ACA, this is a given, since the wavefield is measured along a circle. But with the

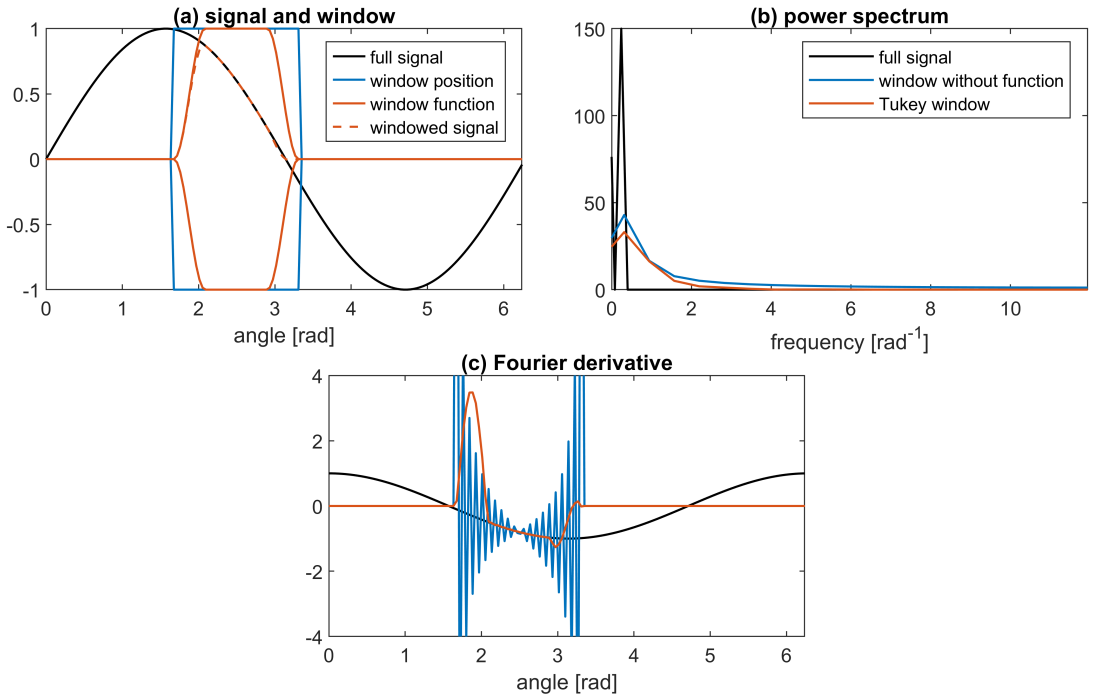



**Figure 7.3:** Graphical explanation of the problems arising when taking the Fourier derivative of a segment of the azimuth. (a) Black line: sine function representing a smooth wavefield along the whole azimuth. Blue box: window position. Red line: the respective Tukey window function. Dashed red line: signal multiplied by window function. (b) Black line: Fourier transform of whole azimuth. Blue line: Fourier transform of segment inside the blue box. Red line: Fourier transform of signal multiplied by window function. (c) Derivatives calculated from spectra shown on panel (b) according to equation 7.10.

STFT, this circle is divided into segments and the periodicity is lost rendering the Fourier derivative inaccurate (see Fig 7.3(c), blue line). Window functions can be used to restore the periodicity to a degree (see Fig 7.3(c), red line). The side-lobes in Figure 7.3(c) are caused by the tapered cosine slopes and, as shown in the previous section, they cancel when the derivatives of adjacent windows are recombined.

Nevertheless, even just considering the region between the side-lobes, the error remains significant (1% relative RMS error in the given example). When the windowed ACA method turned out to be numerically unstable, this inaccurate Fourier derivative appeared to be an obvious culprit. Thus, the windowed Fourier derivative was replaced by the buffered Fourier spectral method (BFSM). This method was proposed by Fu and Liu (2011) as an accurate way to calculate derivatives for non-periodic functions. The accuracy of the derivative is restored

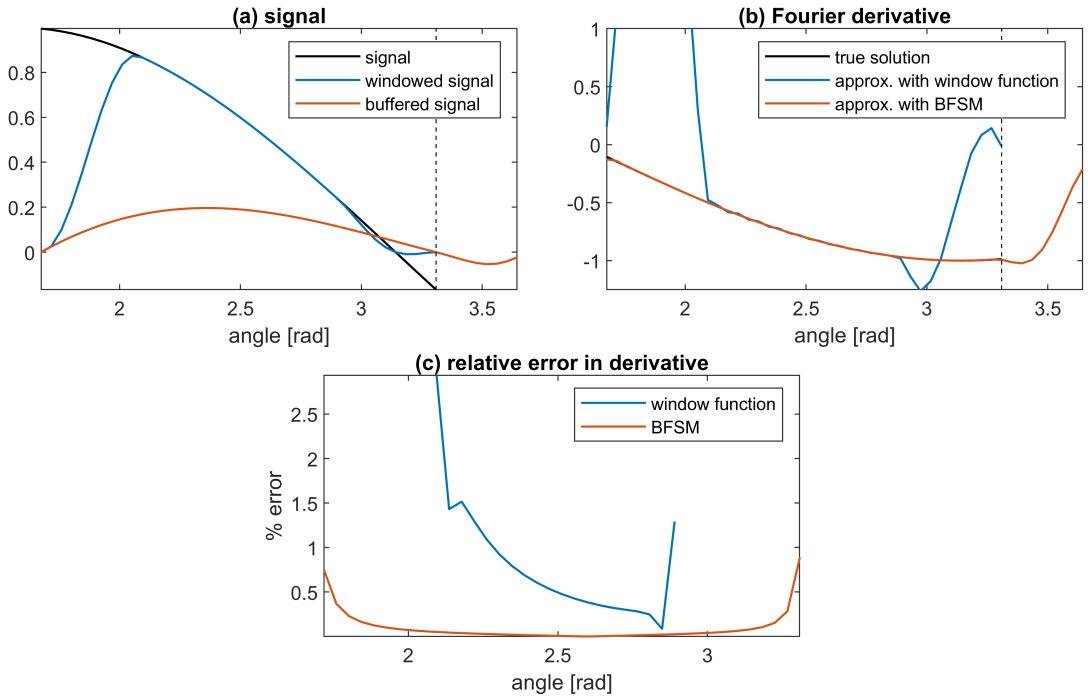



**Figure 7.4:** Comparison of the windowed Fourier derivative and the BFSM using the same segment of a periodic function as shown in Figure 7.3. (a) Black line: Original function segment. Blue line: function segment multiplied with Tukey window. Red line: Function segment after adding linear function and attaching a spline (extension indicated by dashed line). (b) Respective Fourier derivatives (extension indicated by dashed line). (c) Absolute error of the windowed Fourier derivative and the BFSM derivative. Note that the side-lobes of the windowed Fourier derivative were omitted as they would be cancelled by adjacent windows.

in 2 steps: First, a linear function is added to the signal such that the resulting function both starts and ends at 0. This must later be corrected for in the derivative by simply subtracting a constant. Second, an $8^{\text{th}}$ order spline interpolation is used to calculate a smooth segment joining the beginning and the end of the function. After these adjustments the Fourier derivative can be calculated at $5^{\text{th}}$ order accuracy, according to Fu and Liu (2011). The resulting derivative is compared to the previous method in Figure 7.4.

## 7.4 Interpolation

In standard ACA, the azimuthal sampling rate must be constant because the Fourier transform strictly requires regular sampling. In the STFT, the Fourier

transform is applied to individual windows thus technically allowing for different sampling rates. Figure 7.5 shows a hypothetical set-up for the windows in one GLL point. However, in such a set-up the sampling points in the overlap region between windows are not co-located. So, an interpolation method is required both to separate and recombine the windows in physical space.

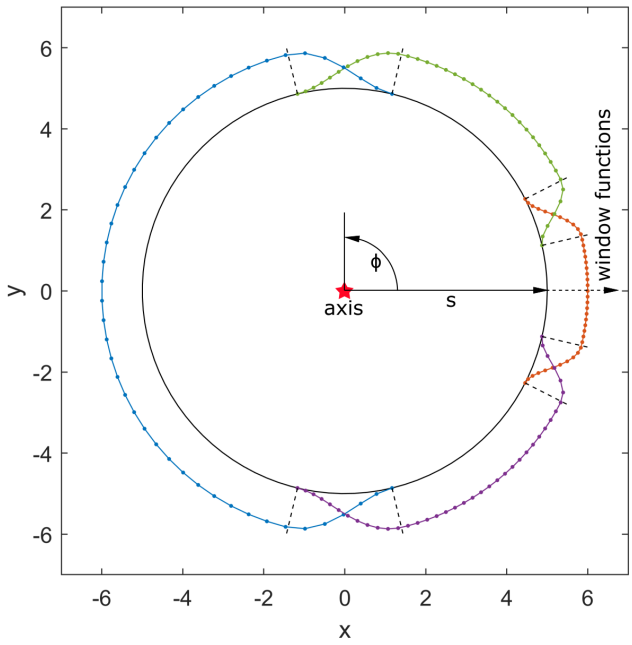



**Figure 7.5:** Hypothetical set-up of azimuthal windows in one GLL point. This example could be used to handle a localised complex region around $\phi = 0$. The spacing of the data points is representative for the azimuthal sampling interval. Asymmetric Tukey windows are used to process the overlapping regions.

As shown in the previous section, $8^{\text{th}}$ order spline interpolation is required to calculate accurate Fourier derivatives for the windows using the BFSM. Since this method is already implemented, it would be desirable to reuse it for the interpolation between windows. There is a potential issue with this however: spline interpolation tends to introduce ringing which may cause problems in the form of propagating errors. An alternative method is Whittaker-Shannon interpolation, which does not have this problem and which further maintains the integrity of the ACA method by preserving the spectrum of the interpolated function. Unfortunately, since the code is numerically unstable even in test cases which avoid interpolation, it is currently not possible to test the stability of either approach.

## 7.5 Changes to the Method Architecture of AxiSEM3D

The work flow for one time step in the original AxiSEM3D code is shown in Figure 3.4. By adding windows and the BFSM, I implemented the altered work flow shown in Figure 7.6. There are six additional steps: four interpolations are performed (when moving from element window to element and back, and when moving from point window to point and back) and buffering is performed twice (before the gradient and before the divergence calculation).

The functionality of the original points and of the original elements is mostly contained in the new point windows and element windows, respectively. The new points and the new elements serve to communicate between adjacent windows. A notable change in the point and element set up is that, for the BFSM to work, all point windows in one element window must have equal azimuthal sampling rates. This may create a mismatch of azimuthal sampling rates at element boundaries which needs to be processed within the points - ideally by using as few interpolations as possible. A solution to this is to have the points store fields in irregularly (following the finest local azimuthal sampling rate) sampled physical space. Here, the stiffness from all point windows is added up, after being multiplied by the respective Tukey window functions. In the elements, overlaps can be processed more cheaply: the weighted averages between the current and the adjacent window are calculated on a window-by-window basis. The ratios for the weighted averages are given by the tapered slopes of the Tukey window.

Evidently, a sizeable overhead is created by all the extra calculations required to handle azimuthal windows. Hence, great care was taken to ensure the algorithm skips superfluous steps. Both points and elements check for co-located sampling points and skip the interpolation if they exist. The BFSM is also skipped if a window is recognised as “global”, i.e. covering the whole azimuth. Thus the method defaults to standard AxiSEM3D in regions without windows.

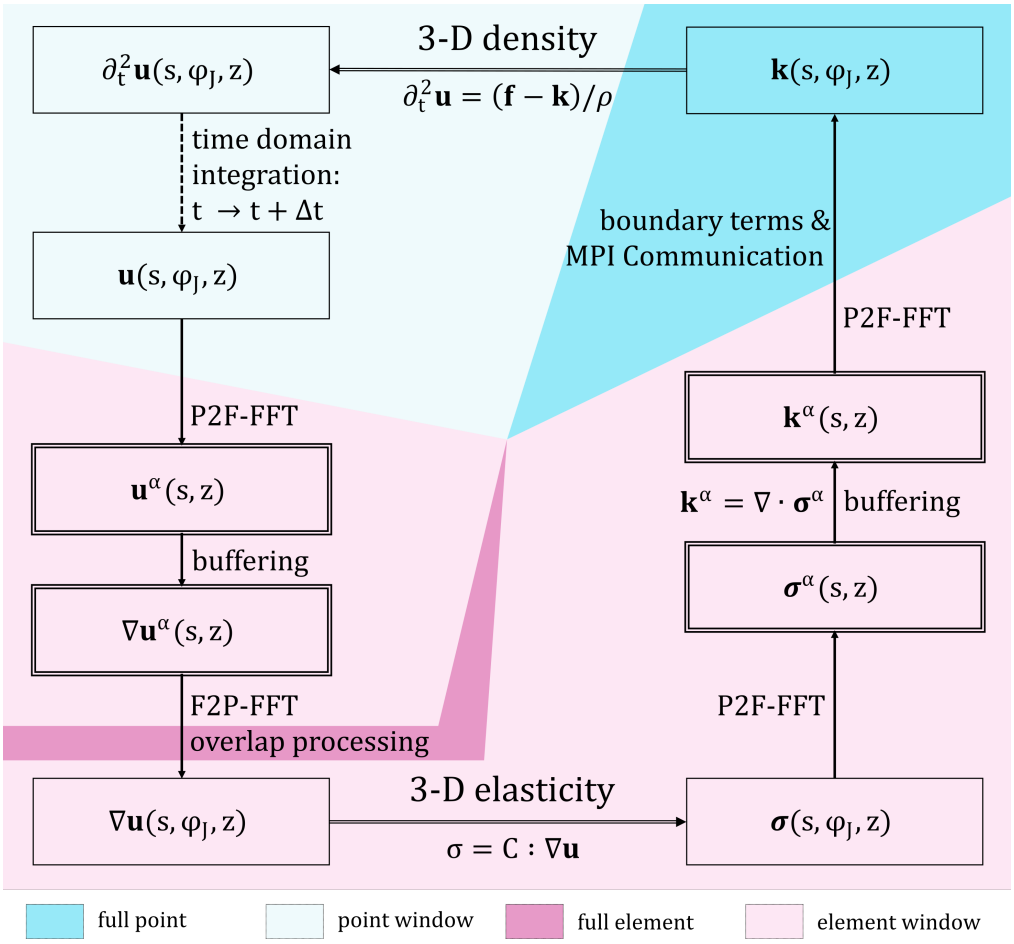



**Figure 7.6:** Method architecture of AxiSEM3D with windows. When comparing this to Figure 3.4, the meaning of the single and double bounded boxes is flipped. Here double lines indicate fields in Fourier domain because the physical domain is now the dominant domain.

## 7.6 Numerical Instability

As of now, a simple test case shows that the suggested STFT method is numerically unstable. The test simulation is set up as follows:

- homogeneous elastic model
- explosive source at 1 km depth with a 1 Hz Gaussian STF
- no windowing near the axis up to a distance where the in-plane limited $n_r$ reaches 108
- 6 equally sized windows each with $n_r$=25

- sampling points in overlapping regions coincide, so no interpolation is performed to combine or separate windows
- overlaps comprised of 7 sampling points
- the BFSM is skipped because no edge artefacts should arise from Fourier transforming a constant (scenario is fully 1D)
- the Courant number has been set to 0.1

The instability starts in the overlapping regions and grows slowly over the course of a few hundred time-steps. Time-step by time-step testing was performed to locate the the first deviation form the manually calculated solution. These initial deviations occur during the recombination of windows and they are limited to numerical round-off errors. However, these errors are amplified in each time step.

The cause of this instability is yet unknown. Several attempts were made to introduce rounding to the calculation which appears to cause the error, but without success. The greatest concern during the method development was the potential impact of the interpolation method on numerical stability. But interpolation is not used at all in this test case, making the observed behaviour quite puzzling.

## 7.7 A Way Forward

Despite extensive tests, it is still possible that the instability is caused by a coding error rather than some inherent property of the method. In either case, the apparent sensitivity of the method to minor errors does not bode well for the prospect of combining the overlaps with an interpolation method. But there are other potential applications of the azimuthal window set up with the BFSM: azimuthal solid-fluid interfaces and azimuthal domain boundaries implemented by masking. Both applications do not require an overlap region, so they may not trigger the observed instability. Failing this, another option may be to further hybridise AxiSEM3D, allowing a mixture of ACA and a full 3D method across the domain. In this way local solids and fluids can be processed and ACA can be used where the overall wavefield smoothness is sufficient for it to be beneficial.

*It takes an earthquake to remind us that we walk on the crust of an unfinished planet.*

— Charles Kuralt

# 8

# Using ACA for Hazard Prediction: A Case Study of the San Francisco Bay Area

## Contents


**8.1 Introduction . . . . . 98**
8.1.1 Local Geology . . . . . 100
8.1.2 Previous Simulations of Hayward Fault Earthquakes . . 101
**8.2 Simulation Set-Up . . . . . 102**
**8.3 Comparison of a Magnitude 6.5 Point Source and a Finite Rupture Earthquake Scenario . . . . . 104**
8.3.1 Ground Motion . . . . . 104
8.3.2 Cost . . . . . 105
**8.4 Testing the Effect of Extremely Low S-Wave Velocities 108**
**8.5 Analysis of the Berkeley Earthquake on 4 Jan 2018 . . 110**
**8.6 Conclusions . . . . . 118**


## 8.1 Introduction

So far, I have only been looking at local-scale applications for ACA, studying test models for petroleum exploration. Seismic hazard prediction is another area which benefits from efficient modelling methods and which may be another use-case for AxiSEM3D with its latest adaptations. For this final study, I focus on a high risk area which is notoriously hard to model: the San Francisco Bay Area (SFBA). A 3D

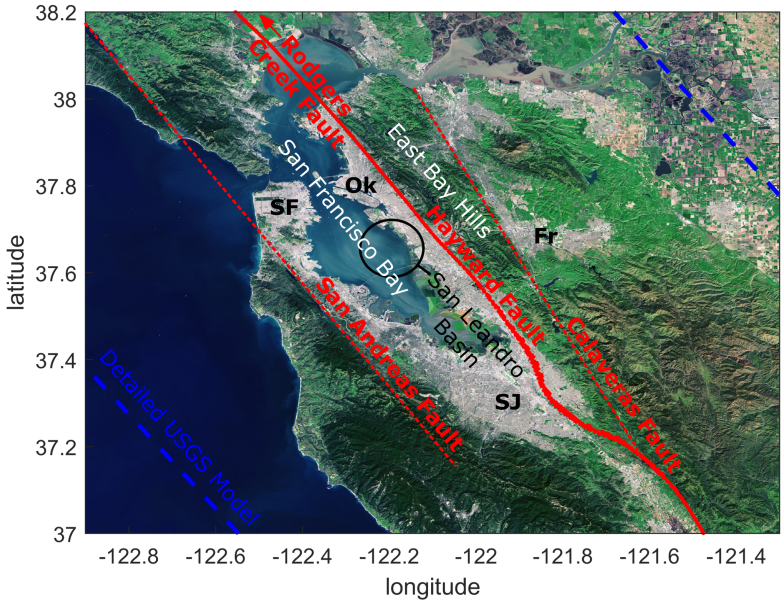



**Figure 8.1:** Annotated satellite photograph of the San Francisco Bay Area (modified from European Space Agency (ESA), 2015). Solid red line shows the position of the Hayward Fault near the surface. Dashed red lines show approximate positions of other prominent faults in the area. Dashed blue lines shows the borders of the detailed USGS velocity model. Black annotations are the abbreviated names of the largest cities in the area: SF - San Francisco, Ok - Oakland, Fr - Fremont, SJ - San Jose. White line marks the position of the cross section shown in Figure 8.2.

velocity model of this region, along with numerous fault models, is available from the USGS Earthquake Hazard Program (U.S. Geological Survey (USGS), 2018a). The SFBA is at significant risk to experience a large (magnitude 6-7) earthquake from the nearby Hayward fault (see Figure 8.1) within the next decade (Field et al., 2017). Many studies have attempted to model the effect such an earthquake would have on the densely populated region around the San Francisco Bay. However, the very low shear wave velocities in the bay muds of the San Leandro Basin (as low as 80 m/s, according to the USGS velocity model) make an accurate simulation of such an event prohibitively expensive. As a workaround, all currently published studies impose a lower limit on shear wave velocity, despite indications that this leads to an underestimate of peak ground velocities in the densely populated areas (Rodgers et al., 2019). In this chapter, I present the results of long-period (1 s) simulations without an imposed limit on shear wave velocity.

### 8.1.1 Local Geology

Figure 8.2 shows a cross section through the USGS SFBA velocity model (descriptions taken from Graymer et al., 2005). The near-surface velocities and topography, which are particularly relevant for predicting and interpreting ground motion, are shown in Figure 8.3. The geology of the SFBA is split into two along the SE-NW trending Hayward Fault. To the north-east lie the East Bay Hills where the bedrock consists in the Great Valley Sequence. This sedimentary sequence is characterised by low wavespeeds that increase rather slowly with depth. The bay itself if formed by

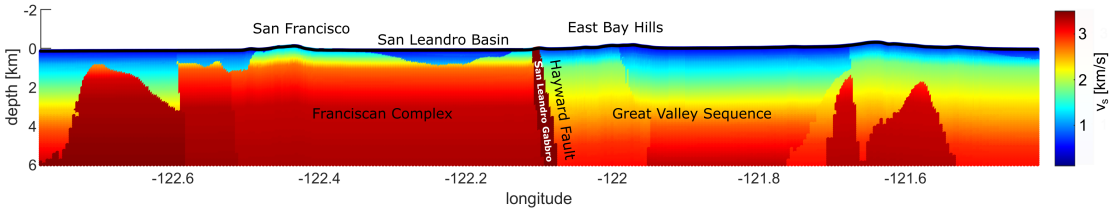



**Figure 8.2:** Vertical cross section of the detailed USGS velocity model (U.S. Geological Survey (USGS), 2018a) at latitude 37.7°$N$ (position marked in Figure 8.3, left). Colour scale shows S-wavespeeds in km/s.

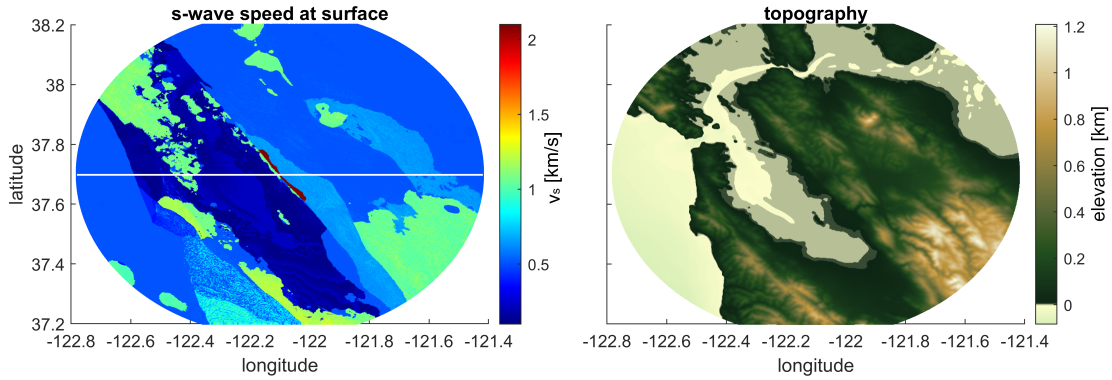



**Figure 8.3: Left:** Horizontal cross section of the detailed USGS velocity model at the surface. White line indicates position of cross section in Fig. 8.2. Colour scale shows S-wavespeeds in km/s. **Right:** Surface topography in the SFBA. Colour scale shows elevation in km.

the Franciscan Complex: low-metamorphosed sedimentary rocks with a thin layer of very low velocity surface deposits, followed by higher velocity rock underneath. It is the San Franciscan Complex surface deposits that make the modelling of the SFBA so difficult. They reach their maximum thickness of >500 m in the San Leandro

Basin. Directly at the Hayward fault, there is a third geological unit called the San Leandro Gabbro, an ophiolite body with high wavespeeds (Ponce et al., 2003).

The Hayward Fault itself is an active transverse fault with a strike around 318° and a dip varying from 90° to 75°, following pre-existing structures (Ponce et al., 2004). It runs parallel to the San Andreas Fault as part of the same large fault complex. Numerous small (magnitude $< 4$) earthquakes originate from the Hayward fault every year. Large earthquakes have occurred roughly every 150 years with the last one being the magnitude 6.5 Hayward Earthquake in 1868 (Hough and Martin, 2015). This gives reason to believe that the next such event will happen soon. Moreover, since the link between the Hayward fault and the Rodgers-Creek fault has been discovered, magnitude estimates for such an event have gone up to 7 (Watt et al., 2016).

### 8.1.2 Previous Simulations of Hayward Fault Earthquakes

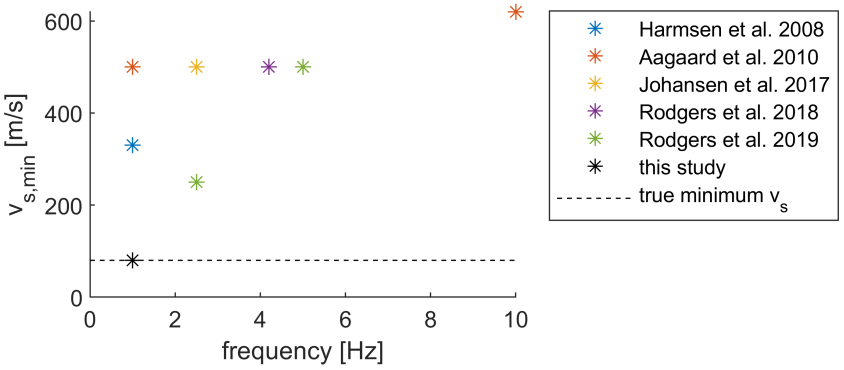



**Figure 8.4:** Input parameters of studies conducted of Hayward fault ruptures. $v_{s,min}$ denotes the lower limit the studies imposed on S-wavespeed to save computational cost.

Several studies of numerically simulated ground motion in the SFBA have been conducted. Due to limits in computational resources, all of these studies were forced to impose a lower limit on S-wave velocity, $v_{s,\min}$. A comparison of the parameter space of these studies is shown in Figure 8.4. Harmsen et al. (2008) simulated a series of finite ruptures with different rupture parameters, such as rupture length, hypocentre and rupture velocity. They found that the resulting peak ground velocities are highly dependent on these parameters, pointing out

the need for probabilistic source modelling. This was confirmed by Aagaard et al. (2010) who combined the results from 5 different research groups with different modelling methods. They accounted for creep in their simulations which slightly lowered the observed PGVs. In another study, Rodgers et al. (2019) investigated the effects of neglecting fault geometry, putting a lower limit on shear wave velocity and using low source frequencies. Their results showed that honouring the fault geometry is vital for accurate PGV predictions. Also, both lowering the frequencies and limiting the minimum shear wave velocity lead to a significant underestimate of the PGV values. It appeared that increasing $v_{s,\min}$ in an effort to lower the computational cost causes a larger error than decreasing the frequency. With this in mind and given limited computational resources, this chapter adds another data point along the $v_{s,\min}$-$f_{\max}$ trade-off, as indicated in Figure 8.4.

## 8.2 Simulation Set-Up

I simulate a hypothetical magnitude 6.5 earthquake originating from the Hayward fault at latitude 37.7°, longitude -122.1°, and 10 km depth. For all simulations, the USGS model is treated volumetrically, the topography and the Moho are implemented using particle relabelling, and the bay and the ocean are air-filled. Attenuation is fully implemented, also using the values from the 3D model. Figure

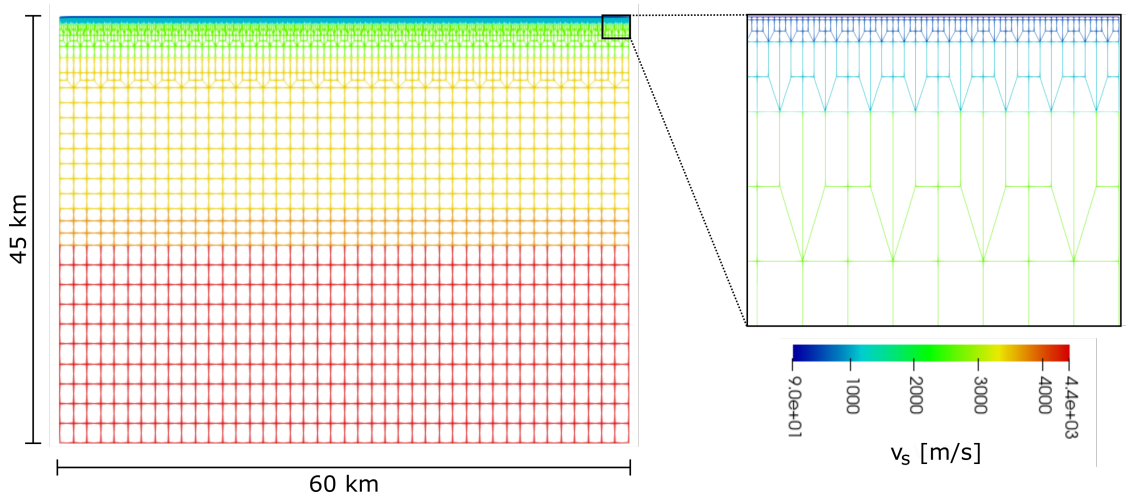



**Figure 8.5:** Mesh used for the simulation of the Hayward fault scenario of the SFBA. Colour scale shows S-wave velocity.

8.5 shows the mesh and the background model used for the simulation. Note the large variations in element sizes which are necessitated by the low S-wave velocities at the top of the model. Because of the mesh refinements, the whole model only contains 6300 elements. The immense cost of the simulations mostly stems from the very small time step (0.0002 s). The time step is chosen based on the Courant-Friedrichs–Lewy stability condition with a Courant number of 0.6 (i.e. $\Delta t = 0.6\frac{\Delta x_{\min}}{\Delta v_{\max}}$) and, here, it is limited by the region where the high-wavespeed San Leandro Gabbro intersects the small top-layer elements.

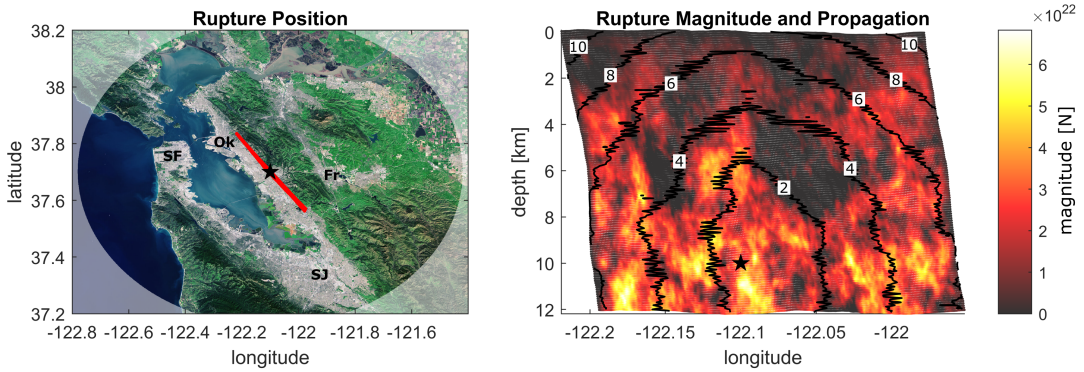



**Figure 8.6: Left:** Satellite photograph showing the modelled area, the section of the Hayward fault involved in the simulated rupture (red line) and the epicentre (black star). **Right:** Scatter of point sources used to simulate the rupture. Colour scale indicates magnitude in Newtons. Contours show rupture propagation with numbers indicating the time since the start of the rupture in seconds.

In the first set of simulations, I compare a moment tensor point source and a finite rupture. The rupture is implemented as a cloud of tightly spaced point sources, as shown in Figure 8.6. This rupture was generated using the rupture generation tool from the Southern California Earthquake Center (SCEC) Broadband Platform (Maechling et al., 2015). The rupture parameters are chosen arbitrarily within a reasonable range for the Hayward fault (e.g. Rodgers et al., 2019; Waldhauser and Ellsworth, 2002), as follows: strike 318°, dip 75°, rake 180°, rupture length 32 km, depth range 0-12 km relative to the surface. Based on these inputs, the moment tensor for the single source are calculated and rupture generation is performed. The point cloud from the rupture generation is further projected onto the Hayward fault,

using the coordinates given in the USGS model. Each point source is a 1 Hz error function with a time shift as indicated by the contours in Figure 8.6 (right).

## 8.3 Comparison of a Magnitude 6.5 Point Source and a Finite Rupture Earthquake Scenario

### 8.3.1 Ground Motion

Figure 8.7 shows the peak ground velocities (PGVs) calculated for the point source and the finite rupture, respectively. The velocities from the point source are significantly higher than the ones from the rupture. This is likely an effect of the rupture releasing the energy slowly over the course of about 10 s whereas the point source releases the entire moment at once. Also, with the point source the near field effect of the dipole source can be observed as four lobes with large PGV values centred on the source. This somewhat obscures the effect of the local geology and topography.

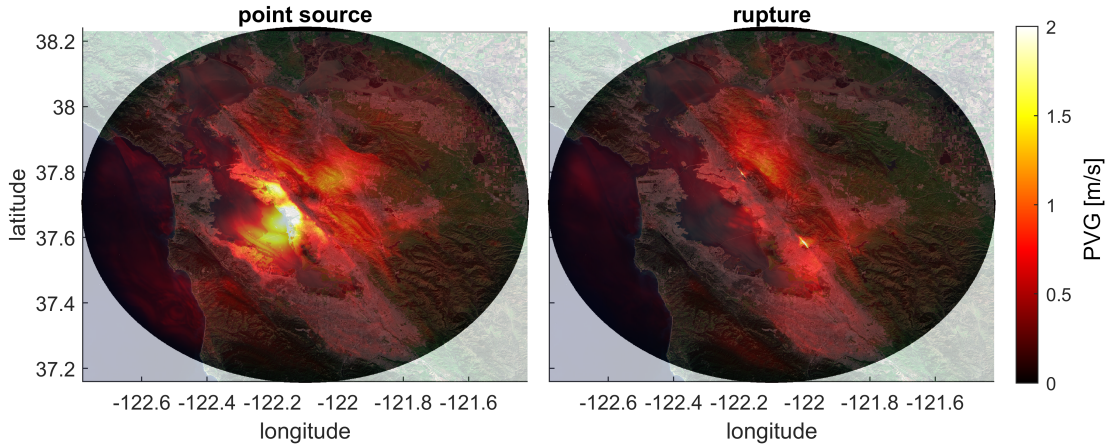



**Figure 8.7:** Peak ground velocities recoded on the surface for the point source and the finite rupture scenarios of the Hayward fault. Colour scale shows maximum recorded magnitude of the velocity vector in m/s.

One effect from the local geology that remains obvious, even in the near field is the amplification of PGVs by the bay muds: to the east of the Haywards fault we get PGVs twice as high as to the west of it. In the far field, the PGV seems largely determined by the near-surface wavespeeds. Comparing both PGV maps

with Figure, 8.3 many of the high wavespeed patches showing up in green in Figure 8.3, can also be recognised in Figure 8.7 as zones of particularly low PGV.

The PGV map resulting from the rupture lacks obvious near-field effects. Instead, we observe that the difference between the east and west side of the Hayward fault has diminished. This is due to the west side being the hanging wall of the Hayward fault. Dipping faults are known to transfer more seismic energy to the hanging wall than to the footwall for simple geometric reasons (Donahue and Abrahamson, 2014). This counteracts the effect of the bay muds. Another noticeable feature of the rupture are the particularly high PGV values at either end of the rupture. This is the region where the rupture extends beyond the San Leandro Gabbro, which apparently has a mitigating effect of ground motion in close proximity of the fault.

Similar observations can be made from the individual wavefield snapshots shown in Figure 8.8. Further, looking at the propagation of the wavefronts, we see waves hugging the Hayward Fault as they follow the fault to the north-west (particularly visible on both 12.7 s snapshots). The acute angle between these wavefronts and the fault is on the east side of the fault, suggesting that the propagation velocity in the west is faster than in the east. This is because the Franciscan Complex in the west, while having very low wavespeed sediments on top, has higher wavespeeds at depth than the Great Valley Sequence to the east (see Fig. 8.2). A similar effect can be seen at 17.7 s at the San Andreas Fault which, at depth, has lower wavespeeds to the west than to the east. Returning to the San Leandro Basin, we see that waves appear to get trapped there, so that the area reverberates for a long time.

### 8.3.2 Cost

This is the first time a finite rupture was implemented in AxiSEM3D, made possible by the off-axis sources developed by Szenicer et al. (2020). Off-axis sources in AxiSEM3D have the disadvantage of requiring very high azimuthal resolution in their immediate vicinity. Also, if wavefronts do not originate from the axis, this compromises the whole idea behind ACA. But in the case of a rupture centred on the axis, azimuthal smoothness persists to a surprisingly high degree.

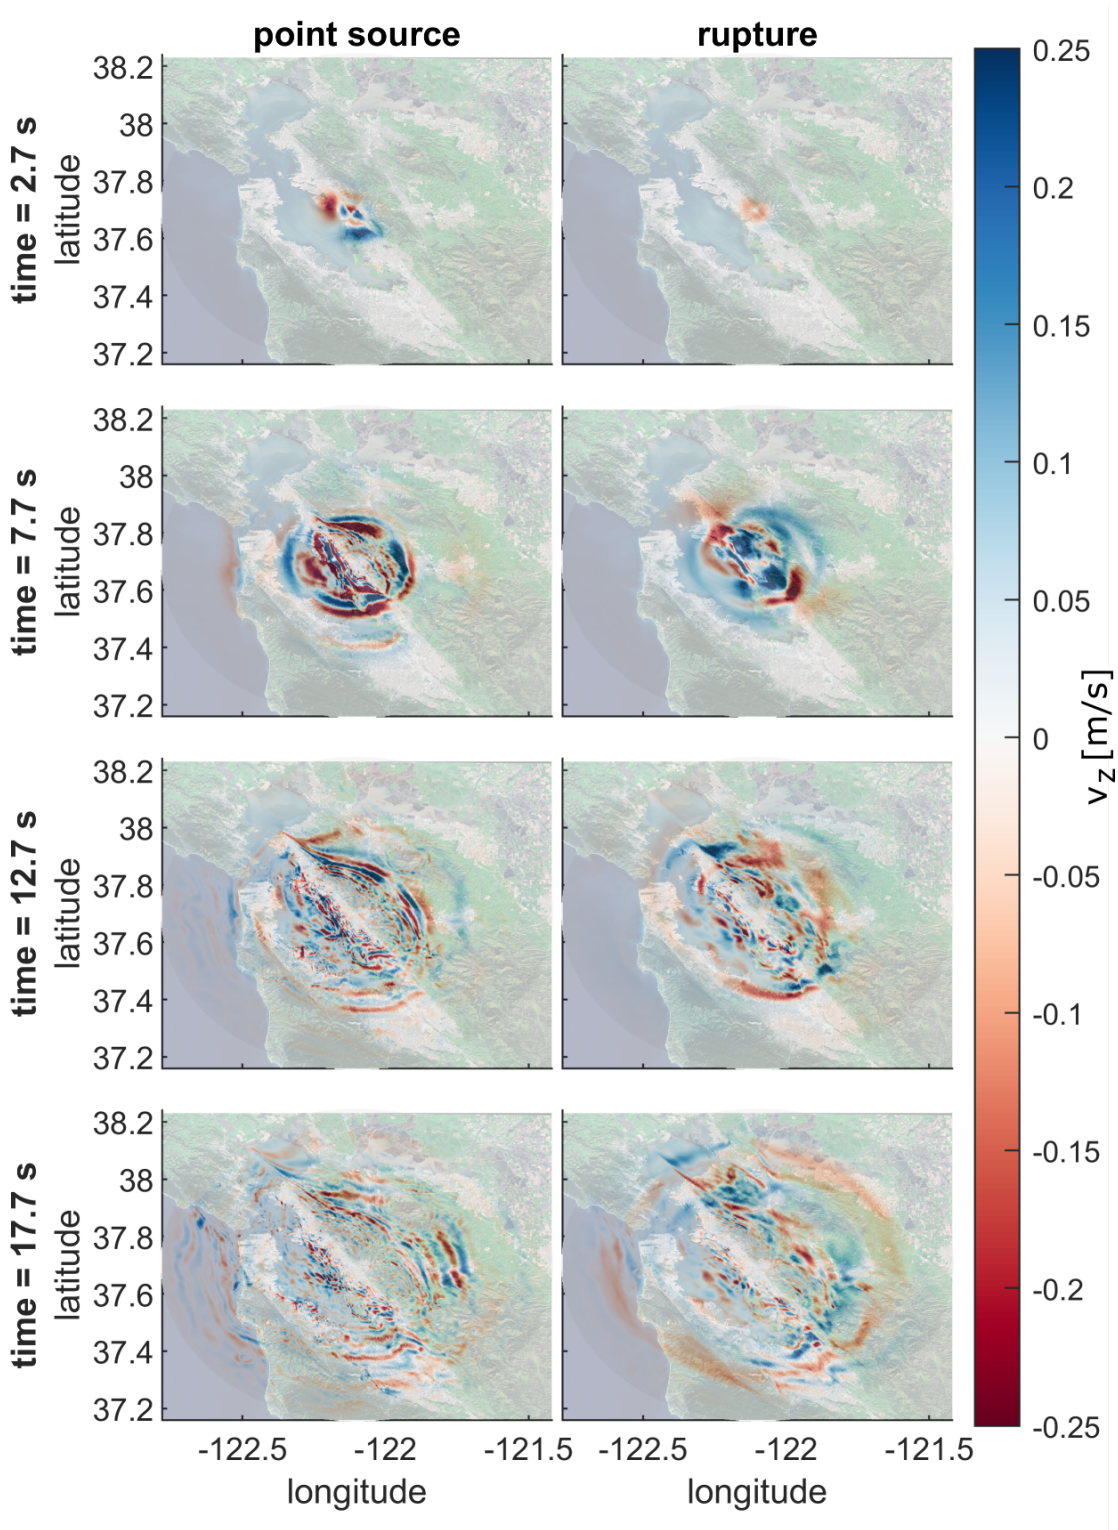



**Figure 8.8:** Snapshots of the surface wavefield of the point source and the finite rupture scenario for the Hayward fault. Colour scale shows vertical velocity in m/s.

After some trial and error, it was possible to manually create a complexity map which functions as an initial guess for the wavefield scanning results. This map reduces the cost of the high resolution test run dramatically. A simulation at full 3D resolution would have $1.8 \times 10^9$ DOF, but the adapted test run is 7 times cheaper with $2.5 \times 10^8$ DOF. The wavefield scanning algorithm can reduce this even further, to $1.0 \times 10^8$ DOF for the rupture and $8.5 \times 10^7$ DOF for the point source. The corresponding complexity maps are shown in Figure 8.9. Note that the maximum azimuthal wavenumber is given as a logarithm to keep structures other than the rupture visible.

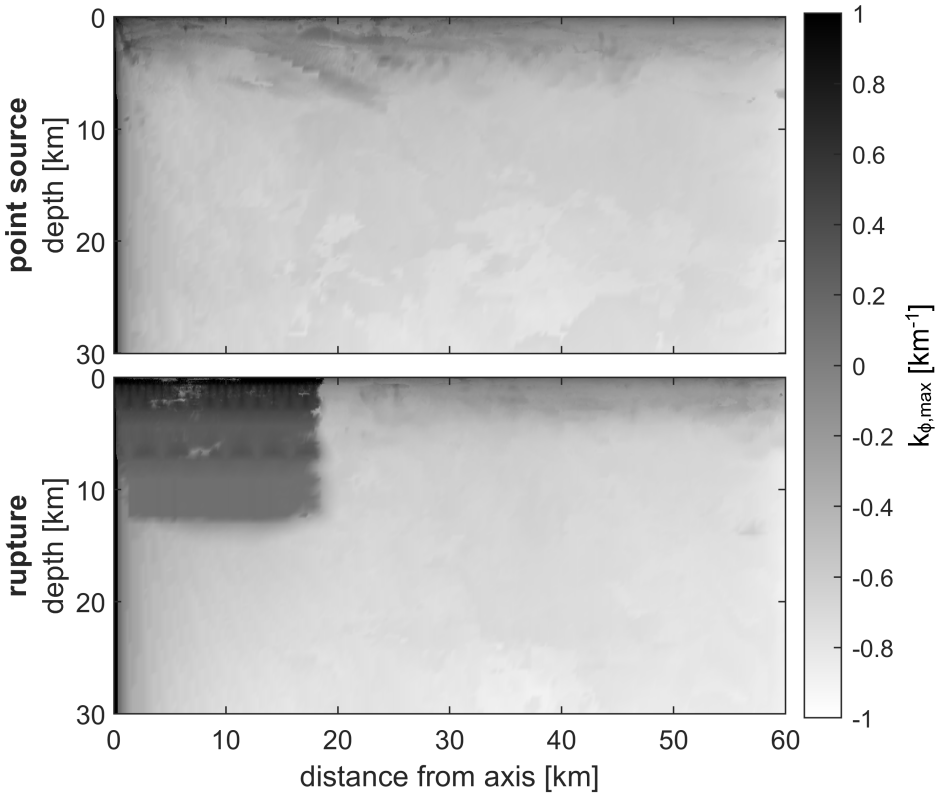



**Figure 8.9:** Complexity maps of the point source scenario and the finite rupture scenario for the Hayward fault. Colour scale shows logarithm of the maximum azimuthal wavenumber in km$^{-1}$.

The rupture is clearly recognisable as the dark rectangle in the top left corner of the bottom panel. In the direct vicinity of the rupture, the azimuthal complexity is saturated, i.e. the azimuthal resolution must be the same as the in-plane resolution. There is a repetitive pattern visible in the dark rectangle which is caused by mesh

refinements. These refinements cause the in-plane resolution to vary (see Fig. 8.5). Since the azimuthal resolution is capped at local in-plane resolution and ACA cannot reduce the resolution any further, the pattern remains visible on the complexity map.

Interestingly, outside the rupture vicinity, the wavefield complexity of the two simulations seems to be roughly equal. The highest $n_r$ recorded is 3461 in the point source scenario and 13179 in the rupture scenario, respectively. But in terms of the sum total DOF, the rupture simulation only requires 15% more than the point source simulation. Nevertheless, the cost of AxiSEM3D simulations with the rupture was about 3 times as high as the simulation with point source. This is in part due to the handling of moment tensor sources which has not been optimised for scalability in the presence of many sources.

It is not possible to do 3D scanning for this run because the memory cost of a wavelet transform for an $n_r \geq 10000$ is prohibitive.

## 8.4 Testing the Effect of Extremely Low S-Wave Velocities

In previous studies of the SFBA, S-wave velocity had to be capped at a minimum value in order to keep simulation cost feasible. While keeping relatively long periods (minimum of 1 s), it is possible to run simulations without such a cap in AxiSEM3D. This allows me to make the first proper assessment of the effect of this commonly used approximation. Figure 8.10 shows PGV maps from the previously introduced Hayward rupture scenario, comparing different S-wave velocity caps.

In both capped scenarios, the PGV is underestimated west of the Hayward fault. The largest absolute errors are measured in the San Leandro Basin and at the coast north of San Jose: here the PGV is underestimated by 50% and 60% when using the 250 m/s cap and the 500 m/s cap, respectively. In relative terms, large errors also occur off the Pacific coast, to the east of the San Andreas Fault. The lower velocity cap of 250 m/s results in fairly wide-spread errors of 10-20% in the urban areas around the SFBA: Berkeley, Oakland, Fremont, San Jose and up the west side reaching Palo Alto and San Mateo. San Francisco is barely

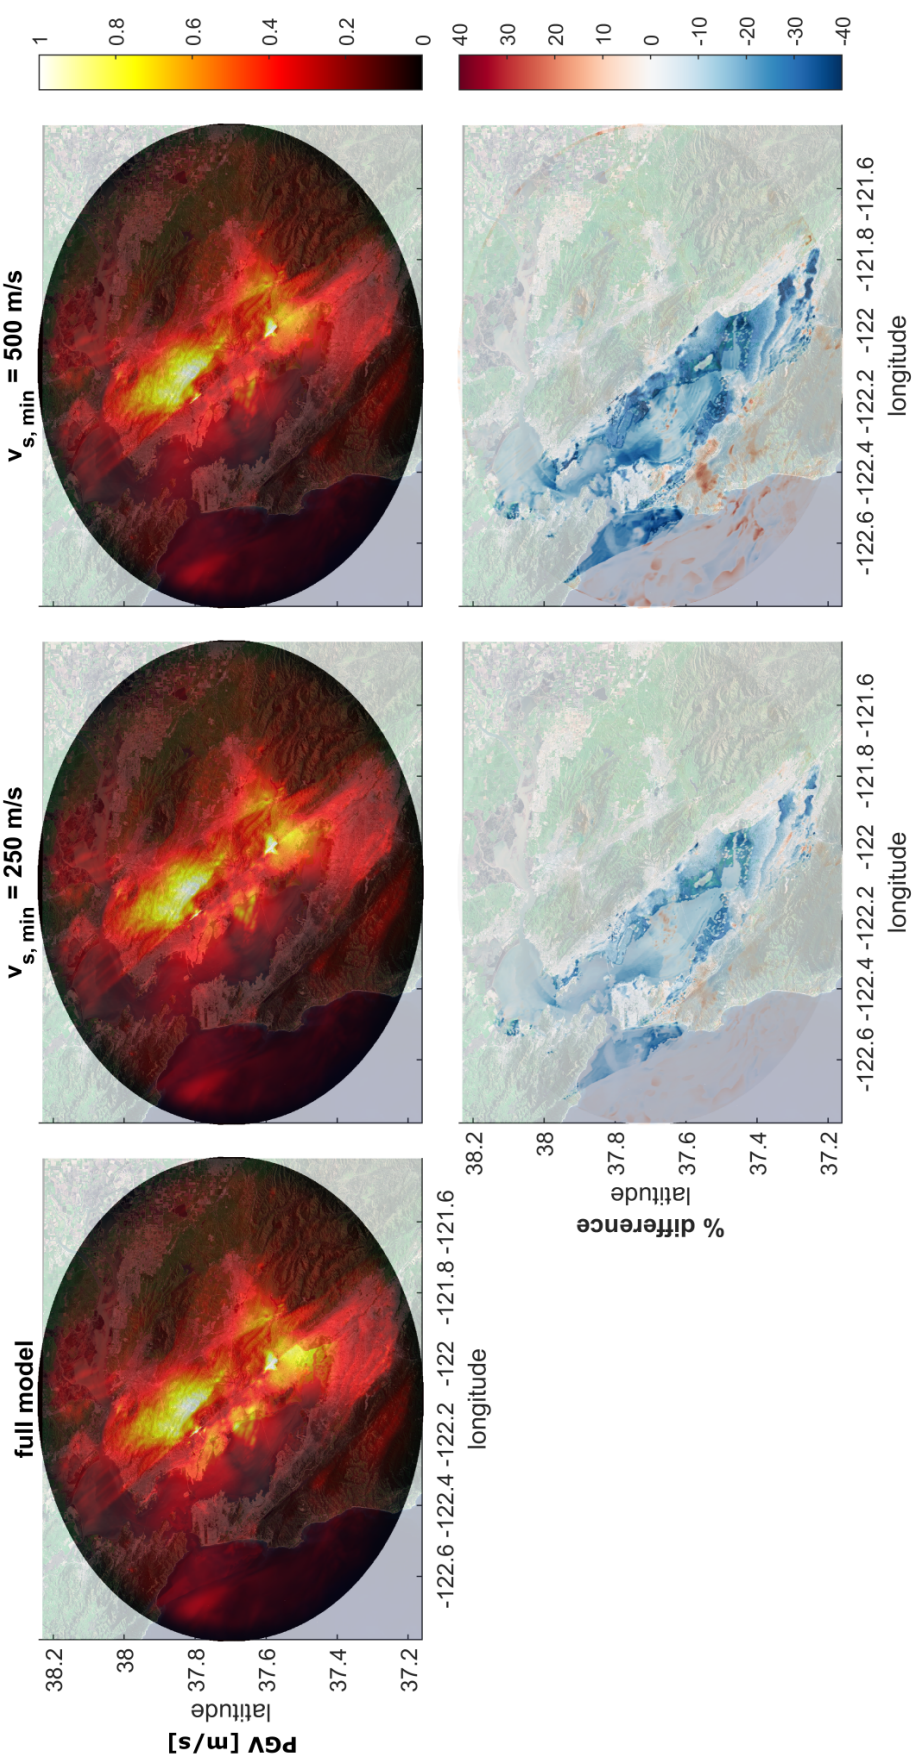



**Figure 8.10:** Comparison of results from simulating a Hayward fault rupture after imposing different lower limits to the S-wave velocity in the USGS SFBA model. Top row shows PGV maps generated using no $v_s$ cap and a minimum $v_s$ of 250 m/s and 500 m/s, respectively. Bottom row shows the difference between the two capped PGV maps and the uncapped PGV map relative to the uncapped PGV.

affected by the velocity cap. the higher cap of 500 m/s results in a very similar pattern but with errors between 30 and 40%. The affected area also increases slightly, mostly to include the bay itself.

A similar study has been conducted by Rodgers et al. (2019), who compare the results of a 250 m/s cap and a 500 m/s cap for a magnitude 6.5 Hayward fault rupture. They simulate up to 2.5 Hz but they are not able to remove the S-wave velocity cap completely. Between the two velocity caps, Rodgers et al. find an underestimation of about 25% west of the Hayward fault and no notable difference to the east. This can be confirmed in this study, where PGVs of the two capped models differ by about 20% around the SFBA. The larger differences found by Rodgers et al. are likely due to the larger frequency spectrum they consider.

## 8.5 Analysis of the Berkeley Earthquake on 4 January 2018

In this section, I compare synthetics created with AxiSEM3D with measured data from a real event in the SFBA. The last major event on the Hayward fault was the 1906 Earthquake, for which there are no publicly available seismogram data. But since then there have been many smaller events, one of which is the magnitude 4.4 earthquake which occurred on the $4^{\text{th}}$ of January 2018 near Berkeley. This event was recorded by a number of stations in the area.

I include all five stations from the Berkeley Digital Seismic Network (BK) which are available through ObsPy (triangles in Fig. 8.11) and a selection of five stations from other networks available from Strong Motion Center (2018) (circles in Fig. 8.11). A moment tensor solution is available on the SCEC website (Northern California Earthquake Data Center (NCEDC), 2018). Lacking any information about the rupture associated with this earthquake, I revert to using a point-source approximation. Figure 8.11 shows the positions of the stations and the epicentre along with the S-wave velocity at the surface.

Four of the five stations from the BK Network are very close to each other and to the epicentre, but their positioning is interesting. Station BRK is to the

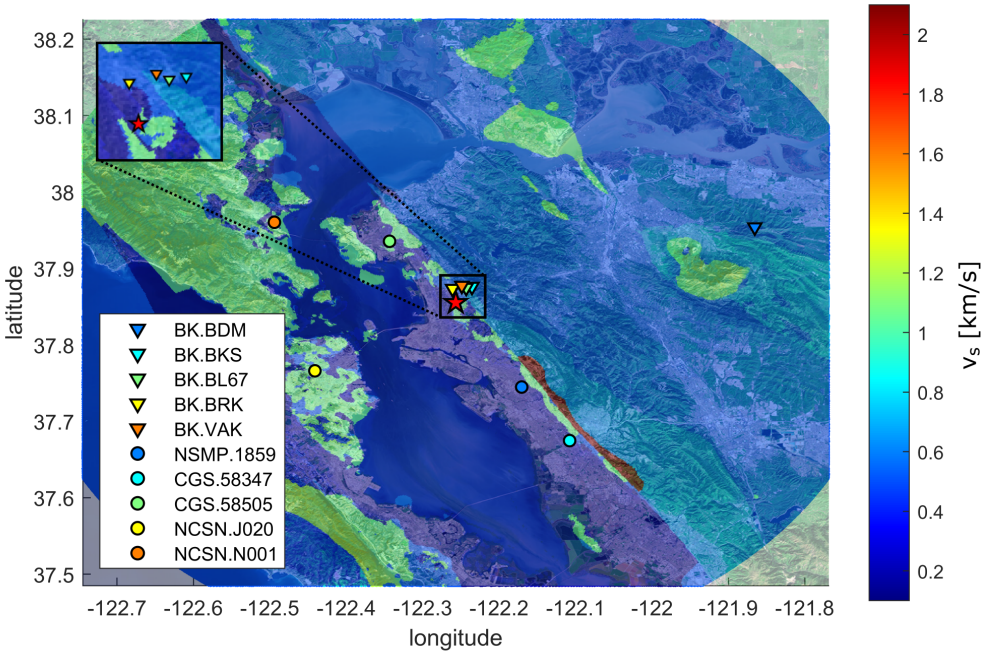



**Figure 8.11:** Satellite photograph of the SFBA with positions of the stations used for the synthetics to data comparison. Triangles show stations from the BK Network and circles show stations from other networks. Star shows position of epicentre of the 2018 Berkeley earthquake. Colour overlay shows surface $v_s$ in km/s.

west of the Hayward fault in a region with shallow bay muds, as is the epicentre of the earthquake. Station VAK is the closest to the Hayward fault, and BL67 and BKS are placed to the east of the fault at increasing distances. To be able to judge the accuracy of the simulation around the bay, I included data from other networks in this comparison. However, these are likely of poorer quality than the BK Network data. The BK Network stations are situated in quiet areas in boreholes, basements or abandoned mine shafts. Meanwhile, the additional stations are placed on surface level in urban areas and, thus, their data may be more noisy. The instrument responses were removed from the data. Since the data available from the non-BK stations has been pre-processed with a bandpass filter between 0.3 and 23 Hz, the same filter has been applied to synthetics of these stations. Further, all traces, including the synthetics, were filtered using a low-pass filter with a pass band frequency of 1 Hz and a stop band at 1.5 Hz.

In Figure 8.12 the synthetics are compared to the recorded data from the

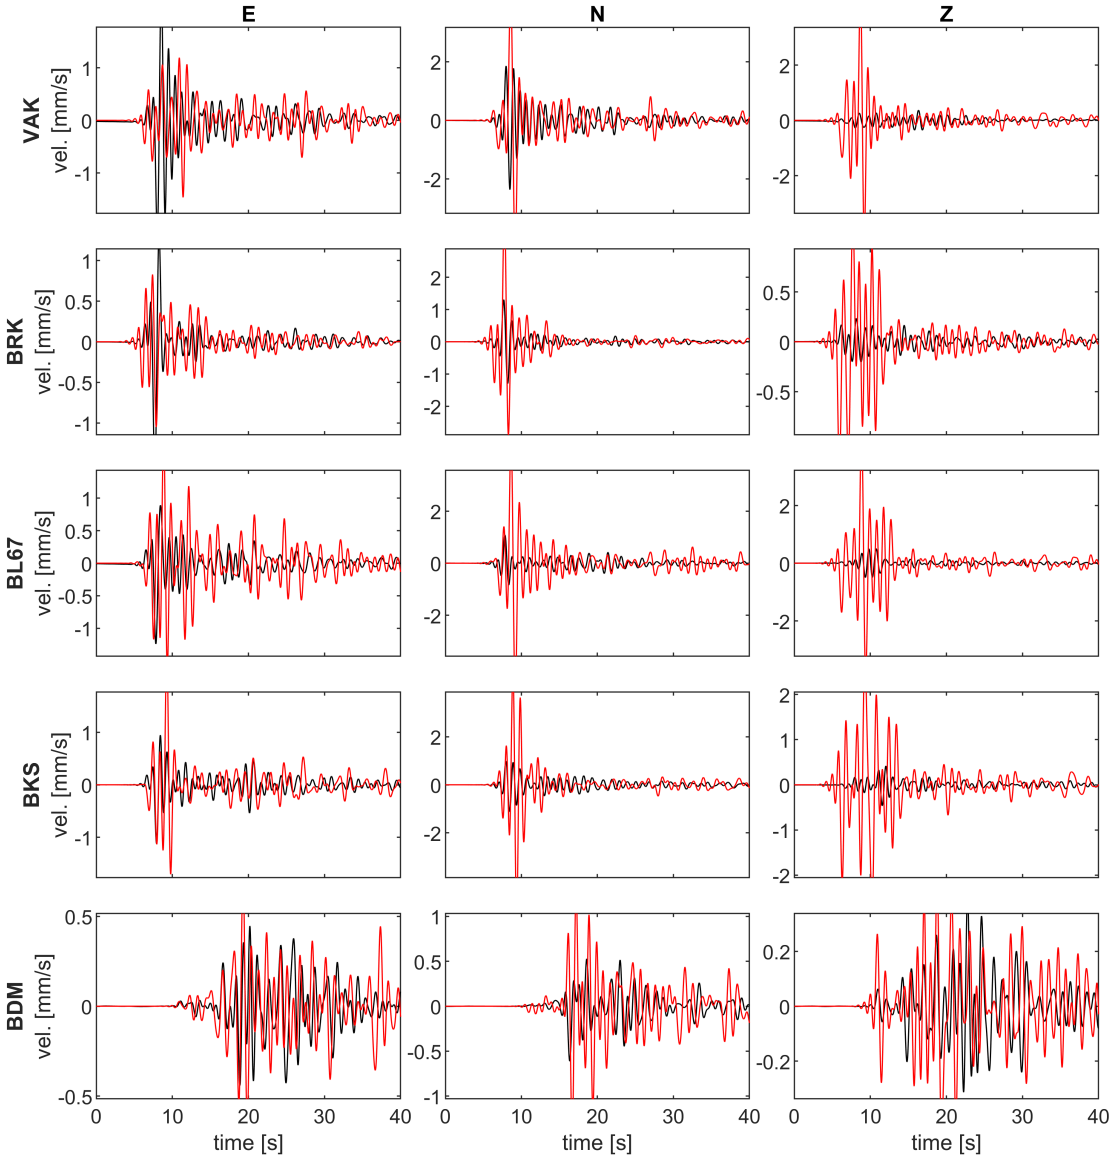



**Figure 8.12:** Comparison of synthetic (red) and measured (black) traces of the 2018 Berkeley earthquake. Traces show the absolute amplitudes of the east (E), north (N) and vertical (Z) components of the ground velocity for the 5 stations in the BK Network shown as triangles in Figure 8.11.

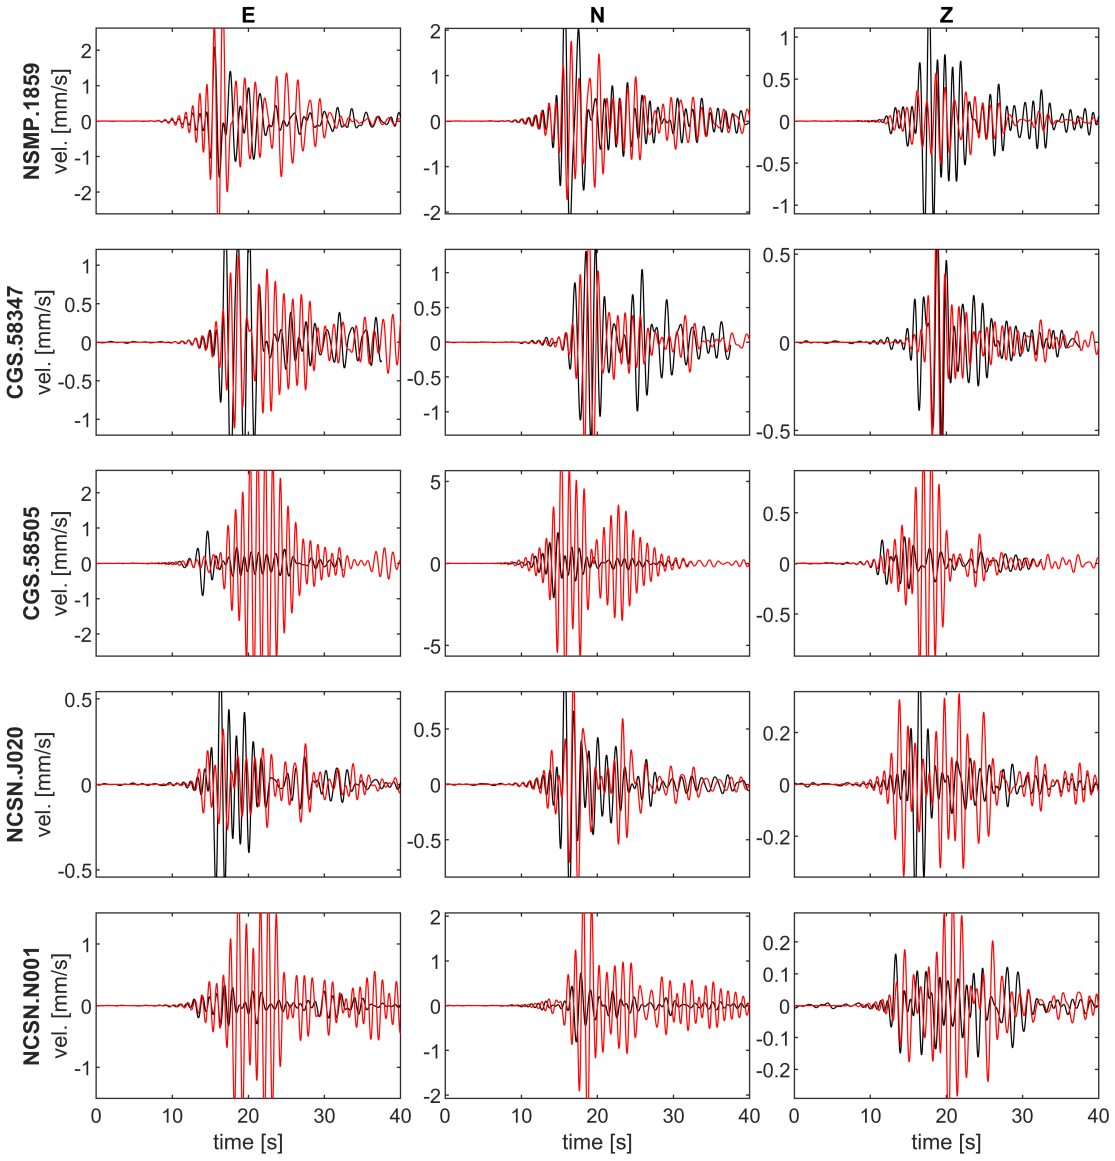



**Figure 8.13:** Comparison of synthetic (red) and measured (black) traces of the 2018 Berkeley earthquake. Traces show the absolute amplitudes of the east (E), north (N) and vertical (Z) components of the ground velocity for the 5 stations in various networks shown as circles in in Figure 8.11.

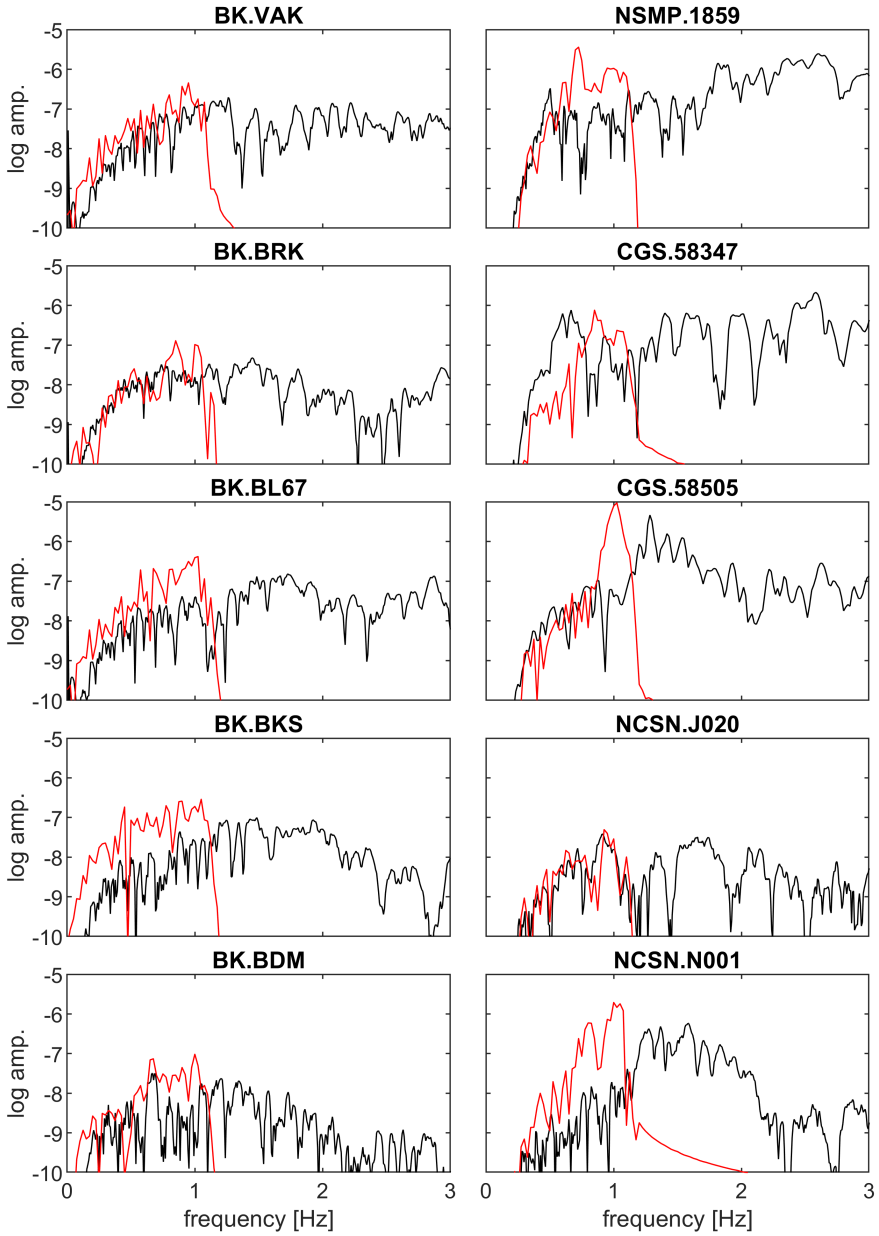



**Figure 8.14:** Comparison of synthetic (red) and measured (black) spectra of the 2018 Berkeley earthquake. Y-axis shows the logarithm of the power spectrum (in [m/s/Hz]) of the east components shown in Figures 8.12 and 8.13. The measured data has not been low-pass filtered.

BK Network. The arrival times fit well at all five stations. This is particularly visible in the east component (left column of panels). In the vertical component, the synthetics seem to arrive early but this only appears this way because the real data has much lower amplitudes. Regarding amplitudes, we get similar values for the east component, but the synthetics overestimate both the north component and the vertical component. Interestingly, the amplitude match is better for the distant BDM station than for the four stations which are close to the source. The north-components of the BRK and the BKS stations stand out with a particularly close match in phase, whereas in the east-component of BL67 the phase appears shifted by a full 180°.

The data from the additional stations is compared to synthetics in Figure 8.13. Again, we see a good fit in terms of arrival times but, with this dataset, there seems to be no systematic deviation in the north and vertical component amplitudes. Instead, the amplitudes at the two stations to the north-west, CGS.58505 and NCSN.N001, is overestimated by the synthetics.

Figure 8.14 shows the power spectra of the east components of all ten stations. The measured data was not low-pass filtered in this comparison to show the relative frequency content of the event that can not yet be simulated due to the high cost. For the aforementioned CGS.58505 and NCSN.N001 stations, we see that the overestimated amplitudes are limited to a narrow frequency band around 1 Hz. Similar amplitudes are found in the measured data around 1.2 Hz, which is just above the range of frequencies I managed to simulate. The closest fits in terms of frequency contents are seen for stations BK.BRK and NCSN.J020.

The consistent fit of arrival times shows that the simulation set-up is correct and that the USGS model is largely accurate. One exception to this may be indicated by the overestimation of amplitudes to the north-west. Based on my previous interpretations of the PGV maps, such overestimation may be caused by using wavespeeds that are too low near the surface. Alternatively, the discrepancy may be caused by neglecting non-linear soil mechanics which can have a damping effect (e.g. Hartzell et al., 2004). Another systematic discrepancy is seen in the north and

vertical components of the four BK Network stations clustered near the Hayward fault. An initial guess for the cause of this was that the source moment tensor may be rotated with respect to the true source. However, this seems unlikely given the fact that the same systematic discrepancy cannot be seen in the stations from the additional networks. Instead, the differences may arise from small-scale local geological features, including local anisotropy. Also, an issue specifically affecting the accuracy of the synthetics of the BK Network is that these stations are buried beneath 20-100 m of overburden and that the burial depth is not reported (except for the BK.BDM station). Within the top km of the subsurface, seismic amplitudes decrease rapidly with depth. So, an error in station depth results in a significant amplitude error in the synthetics.

Apart from the two north-west stations and the four clustered BK Network stations, the amplitudes of the main arrivals and of the coda fit reasonably well. The phase and envelope fit are very hit-and-miss, however. This could, in part, be due to the USGS SFBA model neglecting small-scale local geology. Another likely significant contribution to the observed differences arises from treating most of the model as a volumetric grid. It has been shown repeatedly that not treating sharp velocity contrast rigorously results in errors in the synthetics (Chaljub et al., 2015; Maufroy et al., 2015). The 2D mesh used in ACAs greatly limits our capability of adhering to discontinuities: only continuous, near-horizontal discontinuities can be implemented properly. The USGS SFBA model, with its numerous transverse faults, contains many velocity contrasts which cannot be meshed under the given limitations. Although, it should be noted that even without these limitations meshing such a complex model is very difficult. So far, there has not been any published attempt at meshing the SFBA rigorously, which comes to show how challenging such an endeavour would be.

Finally, in Figure 8.15, I compare the synthetic PGV maps with the PGVs recorded by an array of stations in the area. The data are taken from the ShakeMap data service of the USGS (U.S. Geological Survey (USGS), 2018b). The stations have rather variable data quality, with some of them being installed in buildings

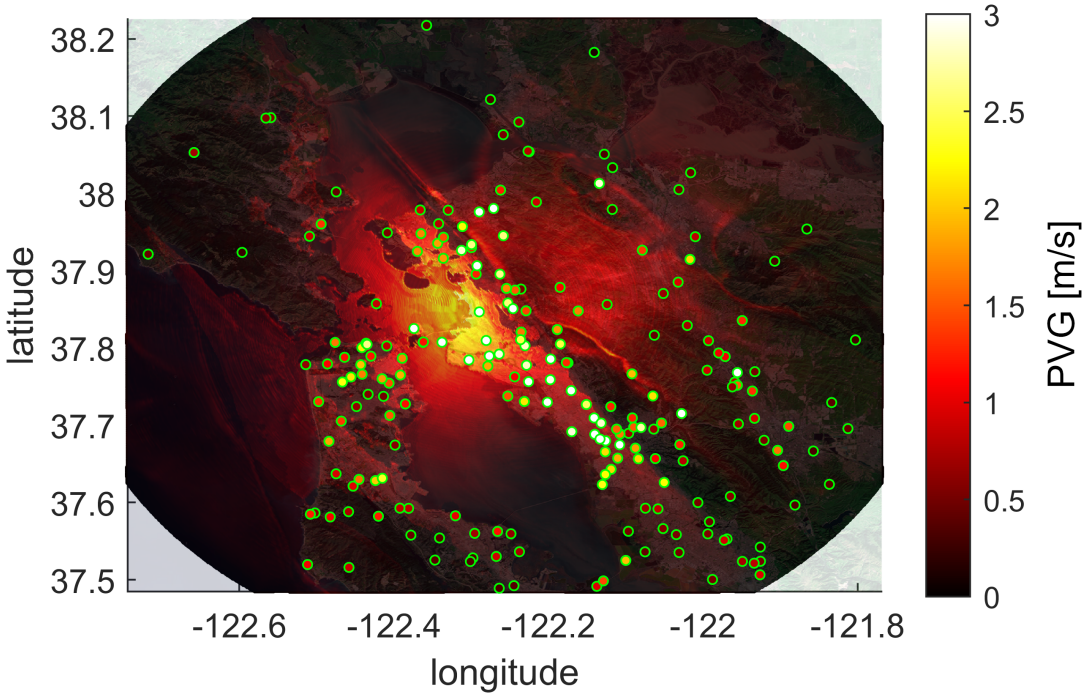



**Figure 8.15:** Comparison of synthetic PGV (background) and measured PGV (green markers) of the 2018 Berkeley earthquake. Colour scale indicated PGV in cm/s.

in urban centres. Consequently, I will stick to observing general trends without putting weight on individual data points.

Many stations agree roughly with the simulated PGV. Notably, there seems to be no systematic under-prediction of PGV values on the hanging wall side (east) of the Hayward fault. This indicates that the point source approximation is suitable for the 2018 event. However, the recorded maximum values are generally higher than the simulated maximum values, contradicting the trend seen in the traces in Figures 8.12 and 8.13. This is likely because the data shown here is broadband whereas the traces seen earlier only contain frequencies up to 1.2 Hz. The high amplitudes recorded by the ShakeMap array probably occur at higher frequencies. Further, there is one area where the measured PGV are systematically higher than the predicted PGV: the vicinity of San Leandro, which is the area with particularly deep bay muds. This is despite the low S-wave velocities of the bay muds being fully included in this simulation. Again, a likely explanation is the lack of higher

frequencies in the numerical simulation. Some evidence for this can be seen in Figure 8.14 where stations NSMP.1859 and CGS.58347, both of which are situated near San Leandro, exhibit maximum amplitudes around 2.5 Hz. The amplifying effect of the bay muds has been shown to more strongly affect higher frequencies (Rodgers et al., 2019). Moreover, thin sediment layers can trap high frequencies, causing amplification of narrow frequency bands. This demonstrates the necessity of pushing efficient modelling methods further, until the SFBA can be modelled accurately in the entire broadband frequency range.

## 8.6 Conclusions

A thin layer of very low wavespeed sediments makes earthquakes in the SFBA hard to model. All previous studies of this region have therefore imposed a lower limit on S-wave velocities to reduce computational cost. In this chapter, a hypothetical magnitude 6.5 earthquake and a real earthquake which occurred in 2018 in the SFBA were modelled without such a limit using frequencies up to 1 Hz. I found that, even at low frequencies, the PGV in the bay is underestimated by 20-40% if velocity limits of 250-500 m/s are imposed.

Further, the results for the 2018 event show that a point source approximation is suitable for this event. Arrival times are in good agreement between the synthetic and the measured traces. Envelope and phase fits are less convincing, which may be due to a number of reasons including small-scale geology which may not be included in the model, and the fact that the many steep discontinuities in the SFBA have not been meshed. Discrepancies between the modelled and the recorded amplitudes at stations in the north-west of the model suggest that the USGS velocity model may underestimate the surface velocities in this region. Recorded PGVs and simulated PGV maps are in rough agreement, but I conclude that higher frequencies need to be simulated to accurately predict ground motion for a given event.

On a more technical note, the results also demonstrate the feasibility of modelling finite ruptures in AxiSEM3D using the off-axis source implementation from Szenicer et al. (2020). Azimuthal wavefield complexity is extremely high in the direct

vicinity of the rupture. At greater distance, however, the wavefield complexity of the rupture simulation and a point source simulation are roughly equal. For this reason, wavefield scanning performs very well, reducing the DOF of the rupture simulation by a factor 18 compared to an unoptimised 3D run. However, several trial runs were required to create a working initial guess for the complexity map, since running the test simulation at full 3D resolution would have been too costly. I found that the biggest factor in the computational cost of the simulations is a very small timestep, suggesting that local time stepping would be particularly beneficial for simulations of the SFBA.

*A thing is not proved just because no one has ever questioned it. What has never been gone into impartially has never been properly gone into. Hence skepticism is the first step toward truth. It must be applied generally, because it is the touchstone.*

— Denis Diderot

# 9 Discussion

## Contents


**9.1 The Potential of ACA for Exploration-Type Settings . 120**
**9.2 Cost Scaling of ACA for Global-Scale and Local-Scale Models . . . . . 121**
**9.3 Other Lessons from 2D and 3D Wavefield Scanning . . 122**
**9.4 Overcoming the Limitations of ACA . . . . . 123**
**9.5 Making Expensive Simulations Possible Through ACA 124**


## 9.1 The Potential of ACA for Exploration-Type Settings

After adding absorbing boundaries and optimising Cartesian models in AxiSEM3D, I tested whether ACA is at all useful in very complex exploration-type settings. Contrary to expectations, I find that ACA leads to a significant speed up for a source-centred salt body, and I demonstrate that the performance is not decreased much if the salt-body is not source centred. This bodes well for the use of ACA on general exploration-type settings. However, there is currently still one road block on the way to making this a reality: accessibility.

The set-up for an AxiSEM3D simulation is quite different from the set up of classic seismic 3D simulations. In Chapter 5, I show the optimal but uncommon

(at least in resource exploration) use-case of having a single seismic source in the centre of a model. For a source near the edge of a model, the model would have to be extended manually to allow for proper placement of the source. For a grid of sources this would have to be done repeatedly. This is quite tedious - especially from the point of view of someone with a highly optimised FDM that they are used to work with. And, as pointed out in the literature review, even standard SEMs are considered "too complicated" by the exploration community. So, we cannot expect AxiSEM3D in its current state to gain a following there. If this is still a goal we wish to pursue, we need to develop pre-processing software that automates the AxiSEM3D-specific model generation, ideally including estimates for cost savings and an approximation for the complexity map.

Another interesting line of approach would be to develop an ACA-FDM. Finite difference methods are faster than spectral element methods for a given level of accuracy. Hence, if ACA is to be used to achieve maximum computational efficiency, moving towards an ACA-FDM is the natural progression.

Finally, an important addition to increase the scope of ACA is the implementation of azimuthally truncated domains. This would allow efficient modelling of a larger range of geometries, without the need for model extensions and other trickery which may scare away inexperienced users. The BFSM, introduced in Chapter 7, combined with the absorbing boundaries, introduced in Chapter 4, may offer a way to do this.

## 9.2 Cost Scaling of ACA for Global-Scale and Local-Scale Models

Cost scaling of ACA with frequency is not as favourable on local scales as it is on global scales. My tests showed scaling with the frequency to the power 4, which is the same for standard 3D modelling, rather to the power 3, which was reported by Leng et al. (2016) for global-scale models. We attribute this to the relative levels of model smoothness of local scale and global scale models. Even in a version of the SEG salt model which was subjected to intense Gaussian smoothing, the power of cost scaling did not decrease. It is likely that a reasonable degree

of smoothing on local scales can simply not match the levels of smoothness of global tomographic models like s20rts which was used by Leng et al. (2016). 3D scanning revealed that the maximum azimuthal complexity occurs near abrupt velocity contrasts, in the medium which has the minimum wavespeed. By the very nature of tomography, velocity contrasts sharper than a single wavelength do not occur in global tomographic models. This contributes to an explanation of the observed difference in ACA cost scaling. The parameters which control the scaling of azimuthal complexity with frequency should be subject to further study.

## 9.3 Other Lessons from 2D and 3D Wavefield Scanning

When starting to look into complexity maps, we were trying to find a method to quantify the error in full waveform inversions. Since Kernels are (usually) based on the Born approximation, strong off-plane scattering would cause errors. The complexity maps produced by AxiSEM3D quantify off-plane scattering and thus they may contain information about the feasibility of FWIs. I did not find conclusive evidence to support this. Especially, having learned that wavefield complexity is, in most cases, caused by strong velocity contrasts, I see no advantage in looking at complexity maps over looking at the model itself. Also, since the Born approximation becomes invalid in the presence of strong scattering, no matter whether it is in-plane or off-plane, low azimuthal wavefield complexity could only ever be a necessary but insufficient condition for a successful inversion.

Rather, the information in complexity maps should be acknowledged for what it is: a map of the maximum wavenumbers recorded in $\phi$-direction (or, equivalently, along the azimuth or off-plane) over the course of the simulation. This could be used to quantify the angle of propagation of the most off-plane phase (i.e. the phase that is closest to propagating at 90° to the initial propagation direction) in a manner akin to the split between horizontal and vertical wavenumbers. But even without any further manipulation, the data is useful. It informs the necessary

azimuthal sampling rate, and allows us to get a measure of efficiency for a classic 3D mesh, for complexity-adapted cylindrical meshes and for an ACA set-up.

We are not yet at the point where we can predict 2D complexity maps and thus avoid the ACA test runs. Yet, the importance of addressing this in the future development of ACA cannot be overstated. To give a real-life example: The SFBA simulations of Chapter 8 are prohibitively expensive at full azimuthal resolution. Thus, an arbitrary upper limit had to be imposed on the resolution of the test run. The first test run required 14 days to be completed on a supercomputer. By the end, the data showed that the starting value for azimuthal resolution had not been high enough. Consequently, the run had to be repeated at higher resolution taking another 14 days. But there is reason for optimism: the results from Chapters 5 and 6 show that the relationship between source, model and complexity map is systematic and probably predictable with the right method. Machine learning seems like the ideal tool for this. If this can be implemented, complexity map prediction would be a central part of the aforementioned pre-processing suite, greatly increasing the accessibility of ACA.

## 9.4 Overcoming the Limitations of ACA

Currently, the azimuthal resolution of an ACA simulation must be constant along the azimuth. Equally, solid and fluid domains must be azimuthally continuous and so do discontinuities if they are to be modelled rigorously. In Chapter 7, I set out to write an ACA solver which does not have these limitations. Unfortunately, the resulting solver turns out to be unstable. Having seen the enormous level of complexity this algorithm reached (e.g. requiring four $8^{\text{th}}$ order spline interpolations per point window per time-step) by the end of its development, I conclude that ACA is not the right tool for settings that require variable azimuthal sampling. The efficiency gain from optimising the sampling rate of the ACA is disproportionately small if compared to the redundancy from overlapping windows and the numerous interpolations. A better alternative may be a hybrid approach which combines ACA in model regions with azimuthally continuous media with a fully spatial modelling

method with a cylindrical mesh in regions with azimuthally discontinuous media. This would require interpolation on fewer points (only along the domain boundaries) only twice per time step and no redundancy from overlaps.

It was not possible to pursue this within the time frame of this PhD project, but there may be a chance to salvage some of the developments made for STFT-AxiSEM3D: as mentioned earlier, the BFSM may be a way to azimuthally truncate the modelling domain if combined with absorbing boundaries. Also, although much more complicated, azimuthal solid-fluid interfaces are still within the realm of possibilities using the principles explained in Chapter 8. It is specifically the overlaps which cause the observed instability, and overlaps are only required to vary azimuthal sampling rates - not for solid-fluid boundaries.

## 9.5 Making Expensive Simulations Possible Through ACA

After the favourable scaling of the cost of ACA with frequency on global scales has been demonstrated (Leng et al., 2016), hope is placed in ACA as a potential method to enable us to perform previously prohibitively expensive simulations. As shown in Chapter 8, such hope is warranted. But some lessons on the limitations of ACA cost scaling have also been learned in Chapter 5 and when collecting the data for Chapter 8. Here, I present a perspective from a long-term user and co-developer on the efficiency and applicability of ACA.

Arguably, one of the biggest advantages of ACA is the easy simulation set-up. Setting the azimuthal resolution to 1 allows a lot of the testing that comes with the implementation of a new model to be performed at a fraction of the simulation cost, without having to use a coarser test mesh or other workarounds. This greatly simplifies the implementation of models, particularly if they are computationally expensive. A notable exception from the easy early debugging is particle-relabelling. It is quite common to have a model become unstable after changing the azimuthal resolution. This also changes the sampling of the geometric model and, thus, a problematic section of a model may be sampled at some azimuthal resolutions, but

not at others. From a user-perspective, a debugging tool for these relabelling errors should be at the top of the list for the future development of AxiSEM3D.

As explained previously, cost scaling on local scales does not work as favourably as on global scales. Apparently, on global scales, a scaled-up complexity map from a low resolution test run can be used for higher resolution runs, and the scaling factor can even be lower than the frequency ratios of the test run and the up-scaled run. However, as seen in Chapter 5, doing this when dealing with local-scale complexity leads to noticeable errors even without further lowering the scaling factor. This makes sense, considering that higher frequency seismic waves can resolve finer structures. Without a complexity map prediction method, users are currently limited to up-scaling a complexity map, adding an error margin to the scaled map and then proceeding by trial and error. This results in a cost scaling that is actually somewhat worse than the frequency to the power 4 (the scaling of standard 3D modelling) for the initial test run at the increased frequency. But this is counteracted by the fact that the low-frequency run is already optimised and thus likely an order of magnitude cheaper than an equal low-frequency test run in standard 3D modelling. On the other hand, there is always the chance of making the error margin too small and having to repeat the entire simulation. This is especially problematic for large simulations, where increasing the error margin would add days to the runtime, but an under-sampled simulation would result in a set-back of several weeks.

In the case of the infamously hard-to-model 3D model of the San Francisco Bay Area, it was indeed possible to run several simulations with AxiSEM3D - something that to my knowledge has not been achieved before. A manually optimised run (single source) simulating an area within a 60 km radius of the source for 40 s cost about 28,000 CPU hours. For the trial-and-error that comes with finding suitable complexity maps, two runs of similar cost had to be performed before the main run. The main reasons for the success of this project are the very efficient way AxiSEM3D is written and the relative simplicity with which ACA runs can be optimised. A standard 3D method with local mesh refinements around the bay muds, especially if combined with local time stepping and a cylindrical set-up to make use of azimuthal

smoothness, could likely run the SFBA simulations more efficiently than an ACA method. But the work required for the set-up would be immense. As such, ACA indeed enables us to perform previously prohibitively expensive simulations - at the very least by lowering the point of entry for model-specific optimisation.

*Data are just summaries of thousands of stories – tell a few of those stories to help make the data meaningful.*

— Chip and Dan Heath

# 10 Conclusions

In preparation for this thesis, the functionality of the ACA-based wave propagation code, AxiSEM3D, was expanded adding absorbing boundaries and properly implementing Cartesian input models. This paved the way for testing ACA in the context of local-scale an regional scale models. I found that, provided the study has a suitable geometry (i.e. a source roughly in the centre of the modelling domain), ACA can significantly lower the cost of seismic simulations in complex settings. Depending on the smoothness of the model, reductions in CPU hours of a factor 7-15 have been observed. Cost scaling with frequency is, however, the same as for standard 3D modelling (i.e. frequency to the power 4) for all tested local-scale models. The discrepancy between this behaviour and the more favourable cost scaling reported for tomographic models by Leng et al. (2016) is likely due to the presence of sharp velocity contrasts (in the order of a wavelength or smaller) in local-scale models.

To improve our understanding of the measure of azimuthal complexity used by ACA, I developed a 3D scanning method which is capable of resolving the complexity along the azimuthal dimension in addition to the radial and vertical dimensions. Using this method, I found that shifting complex structures from the source centred axis to an off-axis position does not significantly increase azimuthal complexity. It does, however, increase absolute simulation cost, since the complexity affects ACA cost more strongly the further it is from the axis. Furthermore, I found

that the largest contribution to wavefield complexity comes from the low-velocity side of abrupt velocity contrasts in the 3D model.

The 3D scanning method also provides a measure for the inefficiency due to having fixed azimuthal resolution in ACA. This, along with other current limitations, inspired the creation of a new numerical method capable of performing ACA locally. Such a method was developed using the short-time Fourier transform in combination with the buffered Fourier spectral method to accurately handle non-periodic boundary conditions. The resulting solver is unstable, but it may serve as a starting point for the development of azimuthal masking (i.e. modelling a wedge instead of a whole cylinder) or alternative methods for local ACA.

Finally, ACA was used to model earthquake scenarios in the San Francisco Bay Area. A known problem in modelling this area is the presence of a thin layer of very low wavespeed bay muds which make full simulations prohibitively expensive. I was able to simulate this region without imposing a lower limit on the wavespeeds in the bay muds - something which, to my knowledge, has not been done before. A comparison with simulations with such a velocity cap shows that limiting wavespeeds leads to an average underestimation of peak ground velocities by 20% and 40% for caps at 250 m/s and 500 m/s, respectively. Further, I compare synthetics and real data for the magnitude 4.4 Berkeley earthquake in January 2018. Even with bay muds included, significant differences persist. I attribute the observed differences to soil effects, small-scale local geology not featured in the model, and to the mesh which does not adhere to the many steep discontinuities present in the area.

# References


Aagaard, Brad T. et al. (2010). "Ground-Motion Modeling of Hayward Fault Scenario Earthquakes, Part I: Construction of the Suite of Scenarios". In: *Bulletin Of The Seismological Society Of America* 100.6, pp. 2927–2944.

Alford, R. M., K. R. Kelly, and D. M. Boore (1974). "Accuracy of Finite-Difference Modeling of the Acoustic Wave Equation". In: *Geophysics* 39.6, pp. 834–842.

Alkhalifah, Tariq and Zedong Wu (2017). "Migration Velocity Analysis Using Pre-Stack Wave Fields". In: *Geophysical Prospecting* 65.3, pp. 639–649.

Alterman, Z. S. and J. Aboudi (1968). "Pulse Propagation in a Laterally Heterogeneous Fluid Sphere by Finite Difference Methods". In: *Journal Of Physics Of The Earth* 16.Special, pp. 173–193.

Alterman, Z. S., J. Aboudi, and F. C. Karal (1970). "Pulse Propagation in a Laterally Heterogeneous Solid Elastic Sphere". In: *Geophysical Journal International* 21.3, pp. 243–260.

Alterman, Z. S. and P. Kornfeld (1968). "Finite Difference Solution for Pulse Propagation in a Sphere". In: *Israel Journal Of Technology* 6, pp. 138–149.

Alterman, Z. and F. C. Karal (1968). "Propagation of Elastic Waves in Layered Media by Finite Difference Methods". In: *Bulletin Of The Seismological Society Of America* 58.1, pp. 367–398.

Al-Attar, David and Ophelia Crawford (2016). "Particle Relabelling Transformations in Elastodynamics". In: *Geophysical Journal International* 205.1, pp. 575–593.

Al-Attar, David et al. (2018). "Hamilton's Principle and Normal Mode Coupling in an Aspherical Planet with a Fluid Core". In: *Geophysical Journal International* 214.1, pp. 485–507.

Bao, Hesheng et al. (1998). "Large-scale Simulation of Elastic Wave Propagation in Heterogeneous Media on Parallel Computers". In: *Computer Methods In Applied Mechanics And Engineering* 152.1-2, pp. 85–102.

Berenger, Jean-Pierre (1994). "A Perfectly Matched Layer for the Absorption of Electromagnetic Waves". In: *Journal Of Computational Physics* 114.2, pp. 185–200.

Boore, David M. (1970). "Love Waves in Nonuniform Wave Guides: Finite Difference Calculations". In: *Journal Of Geophysical Research* 75.8, pp. 1512–1527.

— (1972a). "A Note on the Effect of Simple Topography on Seismic SH Waves". In: *Bulletin Of The Seismological Society Of America* 62.1, pp. 275–284.

— (1972b). "Finite Difference Methods for Seismic Wave Propagation in Heterogeneous Materials". In: *Methods in Computational Physics: Advances in Research and Applications*. Ed. by Bruce A. Bolt. Vol. 11. Seismology: Surface Waves and Earth Oscillations. Elsevier, pp. 1–37.

Capdeville, Y., J. P. Vilotte, and J. P. Montagner (2003). "Coupling the Spectral Element Method with a Modal Solution for Elastic Wave Propagation in Global Earth Models". In: *Geophysical Journal International* 152.1, pp. 34–67.

Capdeville, Yann, Laurent Guillot, and Jean-Jacques Marigo (2010). "1-D Non-Periodic Homogenization for the Seismic Wave Equation". In: *Geophysical Journal International* 181.2, pp. 897–910.

Capdeville, Yann et al. (2013). "Residual Homogenization for Seismic Forward and Inverse Problems in Layered Media". In: *Geophysical Journal International* 194.1, pp. 470–487.

Carcione, José M. (1996). "A 2D Chebyshev Differential Operator for the Elastic Wave Equation". In: *Computer Methods In Applied Mechanics And Engineering* 130.1, pp. 33–45.

Cerjan, Charles et al. (1985). "A Nonreflecting Boundary Condition for Discrete Acoustic and Elastic Wave Equations". In: *Geophysics* 50.4, pp. 705–708.

Chaljub, E. et al. (2010). "Quantitative Comparison of Four Numerical Predictions of 3D Ground Motion in the Grenoble Valley, France". In: *Bulletin Of The Seismological Society Of America* 100.4, pp. 1427–1455.

Chaljub, Emmanuel and Albert Tarantola (1997). "Sensitivity of SS Precursors to Topography on the Upper-mantle 660-km Discontinuity". In: *Geophysical Research Letters* 24.21, pp. 2613–2616.

Chaljub, Emmanuel and Bernard Valette (2004). "Spectral Element Modelling of Three-Dimensional Wave Propagation in a Self-gravitating Earth with an Arbitrarily Stratified Outer Core". In: *Geophysical Journal International* 158.1, pp. 131–141.

Chaljub, Emmanuel et al. (2015). "3-D Numerical Simulations of Earthquake Ground Motion in Sedimentary Basins: Testing Accuracy through Stringent Models". In: *Geophysical Journal International* 201.1, pp. 90–111.

Chauris, Hervé and Emmanuel Cocher (2017). "From Migration to Inversion Velocity Analysis". In: *Geophysics* 82.3, S207–S223.

Clayton, Robert and Björn Engquist (1977). "Absorbing Boundary Conditions for Acoustic and Elastic Wave Equations". In: *Bulletin Of The Seismological Society Of America* 67.6, pp. 1529–1540.

Clough, R. W. (1960). "The Finite Element Method in Plane Stress Analysis". In: *Proceedings Of 2nd Asce Conference On Electronic Computation, Pittsburgh Pa.*

Costantino, Carl J. (1967). "Finite Element Approach to Stress Wave Problems". In: *Journal Of The Engineering Mechanics Division* 93.2, pp. 153–176.

Crase, E. et al. (1990). "Robust Elastic Nonlinear Waveform Inversion: Application to Real Data". In: *Geophysics* 55.5, pp. 527–538.

Crase, Edward (1990). "High-Order (Space and Time) Finite-Difference Modeling of the Elastic Wave Equation". In: *SEG Technical Program Expanded Abstracts 1990*. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, pp. 987–991.

Cummins, Phil R., Nozomu Takeuchi, and Robert J. Geller (1997). "Computation of Complete Synthetic Seismograms for Laterally Heterogeneous Models Using the Direct Solution Method". In: *Geophysical Journal International* 130.1, pp. 1–16.

Cupillard, Paul and Yann Capdeville (2018). "Non-Periodic Homogenization of 3-D Elastic Media for the Seismic Wave Equation". In: *Geophysical Journal International* 213.2, pp. 983–1001.

Dahlen, F. A. (1968). "The Normal Modes of a Rotating, Elliptical Earth". In: *Geophysical Journal International* 16.4, pp. 329–367.

De la Puente, Josep et al. (2007). "An Arbitrary High-order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes - IV. Anisotropy". In: *Geophysical Journal International* 169.3, pp. 1210–1228.

Donahue, Jennifer L. and Norman A. Abrahamson (2014). "Simulation-Based Hanging Wall Effects". In: *Earthquake Spectra* 30.3, pp. 1269–1284.

Dumbser, Michael and Martin Käser (2006). "An Arbitrary High-Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes - II. The Three-dimensional Isotropic Case". In: *Geophysical Journal International* 167.1, pp. 319–336.

Dumbser, Michael, Martin Käser, and Eleuterio F. Toro (2007). "An Arbitrary High-Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes - V. Local Time Stepping and P-adaptivity". In: *Geophysical Journal International* 171.2, pp. 695–717.

Duru, Kenneth et al. (2017). "A New Discontinuous Galerkin Spectral Element Method for Elastic Waves with Physically Motivated Numerical Fluxes". In: *Waves17 International Conference On Mathematical And Numerical Aspects Of Wave Propagation*. arXiv: 1802.06380.

European Space Agency (ESA) (2015). *San Francisco Bay Area, USA*. URL: https://www.esa.int/ESA_Multimedia/Images/2015/06/San_Francisco_Bay_Area_USA (visited on 12/01/2021).

Everstine, G. C. (1981). "A Symmetric Potential Formulation for Fluid-Structure Interaction". In: *Journal Of Sound And Vibration* 79.1, pp. 157–160.

Faccioli, E. et al. (1997). "2D and 3D Elastic Wave Propagation by a Pseudo-Spectral Domain Decomposition Method". In: *Journal Of Seismology* 1.3, pp. 237–251.

Faccioli, Ezio et al. (1996). "Spectral-Domain Decomposition Methods for the Solution of Acoustic and Elastic Wave Equations". In: *Geophysics* 61.4, pp. 1160–1174.

Festa, Gaetano and Jean-Pierre Vilotte (2005). "The Newmark Scheme As Velocity-Stress Time-Staggering: An Efficient PML Implementation for Spectral Element Simulations of Elastodynamics". In: *Geophysical Journal International* 161.3, pp. 789–812.

Field, Edward H. et al. (2017). "A Synoptic View of the Third Uniform California Earthquake Rupture Forecast (UCERF3)". In: *Seismological Research Letters* 88.5, pp. 1259–1267.

Fornberg, Bengt (1987). "The Pseudospectral Method: Comparisons with Finite Differences for the Elastic Wave Equation". In: *Geophysics* 52.4, pp. 483–501.

— (1988). "The Pseudospectral Method: Accurate Representation of Interfaces in Elastic Wave Calculations". In: *Geophysics* 53.5, pp. 625–637.

Friederich, Wolfgang and Jörg Dalkolmo (1995). "Complete Synthetic Seismograms for a Spherically Symmetric Earth by a Numerical Computation of the Green's Function in the Frequency Domain". In: *Geophysical Journal International* 122.2, pp. 537–550.

Friederich, Wolfgang, Erhard Wielandt, and Stefan Stange (1993). "Multiple Forward Scattering of Surface Waves: Comparison with an Exact Solution and Born Single-Scattering Methods". In: *Geophysical Journal International* 112.2, pp. 264–275.

Fu, Huankun and Chaoqun Liu (2011). "A Buffered Fourier Spectral Method for Non-periodic PDE". In: *International Journal Of Numerical Analysis And Modeling* 9.

Furumura, T., B. L. N. Kennett, and M. Furumura (1998). "Seismic Wavefield Calculation for Laterally Heterogeneous Whole Earth Models Using the Pseudospectral Method". In: *Geophysical Journal International* 135.3, pp. 845–860.

Gao, Yingjie et al. (2017). "Comparison of Artificial Absorbing Boundaries for Acoustic Wave Equation Modelling". In: *Exploration Geophysics* 48.1, pp. 76–93.

Gauthier, Odile, Jean Virieux, and Albert Tarantola (1985). "Nonlinear Inversion of Seismic Reflection Data". In: *SEG Technical Program Expanded Abstracts 1985*. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, pp. 390–393.

— (1986). "Two-Dimensional Nonlinear Inversion of Seismic Waveforms: Numerical Results". In: *Geophysics* 51.7, pp. 1387–1403.

Gazdag, Jenö (1981). "Modeling of the Acoustic Wave Equation with Transform Methods". In: *Geophysics* 46.6, pp. 854–859.

Geller, Robert J. and Takao Ohminato (1994). "Computation of Synthetic Seismograms and Their Partial Derivatives for Heterogeneous Media with Arbitrary Natural Boundary Conditions Using the Direct Solution Method". In: *Geophysical Journal International* 116.2, pp. 421–446.

Geller, Robert J. and Nozomu Takeuchi (1995). "A New Method for Computing Highly Accurate DSM Synthetic Seismograms". In: *Geophysical Journal International* 123.2, pp. 449–470.

— (1998). "Optimally Accurate Second-order Time-domain Finite Difference Scheme for the Elastic Equation of Motion: One-dimensional Cases". In: *Geophysical Journal International* 135.1, pp. 48–62.

Gilbert, Freeman (1971). "Excitation of the Normal Modes of the Earth by Earthquake Sources". In: *Geophysical Journal Of The Royal Astronomical Society* 22.2, pp. 223–226.

Graves, Robert et al. (2011). "CyberShake: A Physics-Based Seismic Hazard Model for Southern California". In: *Pure And Applied Geophysics* 168.3, pp. 367–381.

Graymer, R.W. et al. (2005). "Three-dimensional Geologic Map of the Hayward Fault, Northern California: Correlation of Rock Units with Variations in Seismicity, Creep Rate, and Fault Dip". In: *Geology* 33.6, pp. 521–524.

Haindl, C., K. Leng, and T. Nissen-Meyer (2021). "A 3D Complexity-Adaptive Approach to Explore Sparsity in Elastic Wave Propagation". In: *Geophysics* 86.5.

Harmsen, Stephen, Stephen Hartzell, and Pengcheng Liu (2008). "Simulated Ground Motion in Santa Clara Valley, California, and Vicinity from M≥6.7 Scenario Earthquakes". In: *Bulletin Of The Seismological Society Of America* 98.3, pp. 1243–1271.

Hartzell, Stephen, L. F. Bonilla, and Robert A. Williams (2004). "Prediction of Nonlinear Soil Effects". In: *Bulletin Of The Seismological Society Of America* 94.5, pp. 1609–1629.

Hedjazian, N, Y Capdeville, and T Bodin (2021). "Multiscale Seismic Imaging with Inverse Homogenization". In: *Geophysical Journal International* 226.1, pp. 676–691.

Hestholm, Stig (2003). "Elastic Wave Modeling with Free Surfaces: Stability of Long Simulations". In: *Geophysics* 68.1, pp. 314–321.

Higdon, Robert L. (1986). "Absorbing Boundary Conditions for Difference Approximations to the Multidimensional Wave Equation". In: *Mathematics Of Computation* 47.176, pp. 437–459.

Hilterman, Fred J. (1970). "Three-Dimensional Seismic Modeling". In: *Geophysics* 35.6, pp. 1020–1037.

Hough, Susan E. and Stacey S. Martin (2015). "The 1868 Hayward Fault, California, Earthquake: Implications for Earthquake Scaling Relations on Partially Creeping Faults". In: *Bulletin Of The Seismological Society Of America* 105.6, pp. 2894–2909.

Huang, Jau-Inn, Yu-Hua Chu, and Chuan-Sheng Yin (1988). "Lattice-Gas Automata for Modeling Acoustic Wave Propagation in Inhomogeneous Media". In: *Geophysical Research Letters* 15.11, pp. 1239–1241.

Hughes, T. J. R. (1987). *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice Hall.

Igel, Heiner (1999). "Wave Propagation in Three-Dimensional Spherical Sections by the Chebyshev Spectral Method". In: *Geophysical Journal International* 136.3, pp. 559–566.

Igel, Heiner, Peter Mora, and Bruno Riollet (1995). "Anisotropic Wave Propagation through Finite-Difference Grids". In: *Geophysics* 60.4, pp. 1203–1216.

Igel, Heiner and Michael Weber (1995). "SH-Wave Propagation in the Whole Mantle Using High-Order Finite Differences". In: *Geophysical Research Letters* 22.6, pp. 731–734.

Igel, Heiner et al. (2000). "The COSY Project: Verification of Global Seismic Modeling Algorithms". In: *Physics Of The Earth And Planetary Interiors* 119.1, pp. 3–23.

Jeon, Hohyub et al. (2020). "Area-Efficient Short-Time Fourier Transform Processor for Time–Frequency Analysis of Non-Stationary Signals". In: *Applied Sciences* 10, p. 7208.

Käser, Martin and Michael Dumbser (2006). "An Arbitrary High-Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes - I. The Two-dimensional Isotropic Case with External Source Terms". In: *Geophysical Journal International* 166.2, pp. 855–877.

Käser, Martin et al. (2007). "An Arbitrary High-Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes - III. Viscoelastic Attenuation". In: *Geophysical Journal International* 168.1, pp. 224–242.

Kawai, Kenji, Nozomu Takeuchi, and Robert J. Geller (2006). "Complete Synthetic Seismograms up to 2 Hz for Transversely Isotropic Spherically Symmetric Media". In: *Geophysical Journal International* 164.2, pp. 411–424.

Kelly, K. R. et al. (1976). "Synthetic Seismograms: A Finite-Difference Approach". In: *Geophysics* 41.1, pp. 2–27.

Komatitsch, Dimitri, Christophe Barnes, and Jeroen Tromp (2000). "Wave Propagation near a Fluid-Solid Interface: A Spectral-Element Approach". In: *Geophysics* 65.2, pp. 623–631.

Komatitsch, Dimitri, Fabien Coutel, and Peter Mora (1996). "Tensorial Formulation of the Wave Equation for Modelling Curved Interfaces". In: *Geophysical Journal International* 127.1, pp. 156–168.

Komatitsch, Dimitri and Jeroen Tromp (2002a). "Spectral-Element Simulations of Global Seismic Wave Propagation - I. Validation". In: *Geophysical Journal International* 149.2, pp. 390–412.

— (2002b). "Spectral-Element Simulations of Global Seismic Wave Propagation - II. Three-dimensional Models, Oceans, Rotation and Self-gravitation". In: *Geophysical Journal International* 150.1, pp. 303–318.

Komatitsch, Dimitri and Jean-Pierre Vilotte (1998). "The Spectral Element Method: An Efficient Tool to Simulate the Seismic Response of 2D and 3D Geological Structures". In: *Bulletin Of The Seismological Society Of America* 88.2, pp. 368–392.

Kosloff, D. et al. (1990). "Solution of the Equations of Dynamic Elasticity by a Chebychev Spectral Method". In: *Geophysics* 55.6, pp. 734–748.

Kosloff, Dan D. and Edip Baysal (1982). "Forward Modeling by a Fourier Method". In: *Geophysics* 47.10, pp. 1402–1412.

Kosloff, Dan, Moshe Reshef, and Dan Loewenthal (1984). "Elastic Wave Calculations by the Fourier Method". In: *Bulletin Of The Seismological Society Of America* 74.3, pp. 875–891.

Kosloff, R. and D. Kosloff (1986). "Absorbing Boundaries for Wave Propagation Problems". In: *Journal Of Computational Physics* 63.2, pp. 363–376.

Kreiss, Heinz-Otto and Joseph Oliger (1972). "Comparison of Accurate Methods for the Integration of Hyperbolic Equations". In: *Tellus* 24.3, pp. 199–215.

Kristeková, Miriam, Jozef Kristek, and Peter Moczo (2009). "Time-Frequency Misfit and Goodness-Of-Fit Criteria for Quantitative Comparison of Time Signals". In: *Geophysical Journal International* 178.2, pp. 813–825.

Leng, K. (2017). "Global Wave Propagation in Three-Dimensional Aspherical Earth Models". PhD thesis. University of Oxford.

Leng, K., J. Korenaga, and T. Nissen-Meyer (2020). "3-D Scattering of Elastic Waves by Small-Scale Heterogeneities in the Earth's Mantle". In: *Geophysical Journal International* 223.1, pp. 502–525.

Leng, Kuangdai, Tarje Nissen-Meyer, and Martin van Driel (2016). "Efficient Global Wave Propagation Adapted to 3-D Structural Complexity: A Pseudospectral/spectral-element Approach". In: *Geophysical Journal International* 207.3, pp. 1700–1721.

Leng, Kuangdai et al. (2019). "AxiSEM3D: Broad-Band Seismic Wavefields in 3-D Global Earth Models with Undulating Discontinuities". In: *Geophysical Journal International* 217.3, pp. 2125–2146.

Leveille, Jacques P. et al. (2011). "Subsalt Imaging for Exploration, Production, and Development: A Review". In: *Geophysics* 76.5, WB3–WB20.

Liu, Tianshi and Haiming Zhang (2017). "Synthetic Seismograms for Finite Sources in Spherically Symmetric Earth Using Normal-mode Summation". In: *Earthquake Science* 30.3, pp. 125–133.

Liu, Yang and Mrinal K. Sen (2010). "A Hybrid Scheme for Absorbing Edge Reflections in Numerical Modeling of Wave Propagation". In: *Geophysics* 75.2, A1–A6.

Lysmer, J. and L. A. Drake (1972). "A Finite Element Method for Seismology". In: *Methods in Computational Physics: Advances in Research and Applications*. Ed. by BRUCE A. Bolt. Vol. 11. Seismology: Surface Waves and Earth Oscillations. Elsevier, pp. 181–216.

Madariaga, Raúl (1976). "Dynamics of an Expanding Circular Fault". In: *Bulletin Of The Seismological Society Of America* 66, pp. 639–666.

Maechling, P. J. et al. (2015). "Scec Broadband Platform: System Architecture and Software Implementation". In: *Seismological Research Letters* 86.1.

Maenchen, G. (1963). *The Tensor Code*. Report. University of California Lawrence Radiation Laboratory.

Marfurt, Kurt J. (1984). "Accuracy of Finite-Difference and Finite-Element Modeling of the Scalar and Elastic Wave Equations". In: *Geophysics* 49.5, pp. 533–549.

Maufroy, E. et al. (2015). "Earthquake Ground Motion in the Mygdonian Basin, Greece: The E2VP Verification and Validation of 3D Numerical Simulation up to 4 Hz". In: *Bulletin Of The Seismological Society Of America* 105.3, pp. 1398–1418.

Mizutani, Hiromitsu, Robert J. Geller, and Nozomu Takeuchi (2000). "Comparison of Accuracy and Efficiency of Time-domain Schemes for Calculating Synthetic Seismograms". In: *Physics Of The Earth And Planetary Interiors* 119.1-2, pp. 75–97.

Moczo, Peter, Johan Robertsson, and L. Eisner (2007). "The Finite-difference Time-domain Method for Modeling of Seismic Wave Propagation". In: *Advances in Geophysics*. Vol. 48, pp. 421–516.

Moczo, Peter et al. (2011). "3-D Finite-Difference, Finite-Element, Discontinuous-Galerkin and Spectral-Element Schemes Analysed for Their Accuracy with Respect to P-wave to S-wave Speed Ratio: Accuracy of the Seismic Numerical Schemes". In: *Geophysical Journal International* 187.3, pp. 1645–1667.

Mora, Peter (1989). "Modeling Anisotropic Seismic Waves in 3-D". In: *SEG Technical Program Expanded Abstracts 1989*. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, pp. 1039–1043.

— (1992). "The Lattice Boltzmann Phononic Lattice Solid". In: *Journal Of Statistical Physics* 68.3, pp. 591–609.

Mora, Peter and Bertrand Maillot (1990). "Seismic Modeling Using the Phononic Lattice Solid Method". In: *SEG Technical Program Expanded Abstracts 1990*. SEG Technical Program Expanded Abstracts. Society of Exploration Geophysicists, pp. 1008–1011.

Mullen, Robert and Ted Belytschko (1982). "Dispersion Analysis of Finite Element Semidiscretizations of the Two-Dimensional Wave Equation". In: *International Journal For Numerical Methods In Engineering* 18.1, pp. 11–29.

Nissen-Meyer, T., F. A. Dahlen, and A. Fournier (2007a). "Spherical-Earth Fréchet Sensitivity Kernels". In: *Geophysical Journal International* 168.3, pp. 1051–1066.

Nissen-Meyer, T., Alexandre Fournier, and F. A. Dahlen (2008). "A 2-D Spectral-element Method for Computing Spherical-Earth Seismograms - II. Waves in Solid–fluid Media". In: *Geophysical Journal International* 174.3, pp. 873–888.

Nissen-Meyer, T. et al. (2014). "AxiSEM: Broadband 3-D Seismic Wavefields in Axisymmetric Media". In: *Solid Earth* 5.1, pp. 425–445.

Nissen-Meyer, T. et al. (2019). "Modelling the Effects of 3D Shallow Scatterers and Atmospheric Sources on Martian Seismic Signals at High Frequencies". In: *AGU Fall Meeting Abstracts*, DI51A–0012.

Nissen-Meyer, Tarje, Alexandre Fournier, and F. A. Dahlen (2007b). "A Two-Dimensional Spectral-Element Method for Computing Spherical-Earth Seismograms - I. Moment-Tensor Source". In: *Geophysical Journal International* 168.3, pp. 1067–1092.

Northern California Earthquake Data Center (NCEDC) (2018). *NCSS Moment Tensor Solution for Event Nc72948801*. URL: `http://www.ncedc.org/mt/nc72948801_MT.html` (visited on 12/05/2021).

O'Brien, Gareth S. (2014). "Elastic Lattice Modelling of Seismic Waves Including a Free Surface". In: *Computers & Geosciences* 67, pp. 117–124.

O'Brien, Gareth S. and Christopher J. Bean (2004). "A 3D Discrete Numerical Elastic Lattice Method for Seismic Wave Propagation in Heterogeneous Media with Topography". In: *Geophysical Research Letters* 31.14.

Orszag, Steven A. (1971). "Numerical Simulation of Incompressible Flows Within Simple Boundaries. I. Galerkin (Spectral) Representations". In: *Studies In Applied Mathematics* 50.4, pp. 293–327.

Padovani, E., E. Priolo, and G. Seriani (1994). "Low and High Order Finite Element Method: Experience in Seismic Modeling". In: *Journal Of Computational Acoustics* 02.04, pp. 371–422.

Patera, Anthony T (1984). “A Spectral Element Method for Fluid Dynamics: Laminar Flow in a Channel Expansion”. In: *Journal Of Computational Physics* 54.3, pp. 468–488.

Peter, Daniel et al. (2011). “Forward and Adjoint Simulations of Seismic Wave Propagation on Fully Unstructured Hexahedral Meshes”. In: *Geophysical Journal International* 186.2, pp. 721–739.

Ponce, D. A., T. G. Hildenbrand, and R. C. Jachens (2003). “Gravity and Magnetic Expression of the San Leandro Gabbro with Implications for the Geometry and Evolution of the Hayward Fault Zone, Northern California”. In: *Bulletin Of The Seismological Society Of America* 93.1, pp. 14–26.

Ponce, D. A. et al. (2004). “Gravity, Magnetic, and High-precision Relocated Seismicity Profiles Suggest a Connection between the Hayward and Calaveras Faults, Northern California”. In: *Geochemistry, Geophysics, Geosystems* 5.7.

Raknes, Espen Birger, BÃ¸rge Arntsen, and Wiktor Weibull (2015). “Three-dimensional Elastic Full Waveform Inversion Using Seismic Data from the Sleipner Area”. In: *Geophysical Journal International* 202.3, pp. 1877–1894.

Reed, W. H. and T. R. Hill (1973). *Triangular Mesh Methods for the Neutron Transport Equation*. Tech. rep. LA-UR–73-479. Los Alamos Scientific Lab.

Reynolds, Albert C. (1978). “Boundary Conditions for the Numerical Solution of Wave Propagation Problems”. In: *Geophysics* 43.6, pp. 1099–1110.

Robertsson, Johan O. A. (1996). “A Numerical Free-surface Condition for Elastic/Viscoelastic Finite-difference Modeling in the Presence of Topography”. In: *Geophysics* 61.6, pp. 1921–1934.

Rodgers, Arthur J., Arben Pitarka, and David B. McCallen (2019). “The Effect of Fault Geometry and Minimum Shear Wavespeed on 3D Ground-Motion Simulations for an Mw 6.5 Hayward Fault Scenario Earthquake, San Francisco Bay Area, Northern California”. In: *Bulletin Of The Seismological Society Of America* 109.4, pp. 1265–1281.

Rothman, Daniel H. (1987). “Modeling Seismic P-Waves with Cellular Automata”. In: *Geophysical Research Letters* 14.1, pp. 17–20.

Schlue, J. W. (1979). “Finite Element Matrices for Seismic Surface Waves in Three-Dimensional Structures”. In: *Bulletin Of The Seismological Society Of America* 69.5, pp. 1425–1438.

Seriani, G. and S. P. Oliveira (2008). “Dispersion Analysis of Spectral Element Methods for Elastic Wave Propagation”. In: *Wave Motion* 45.6, pp. 729–744.

Seriani, Géza and Saulo P. Oliveira (2007). “Optimal Blended Spectral-element Operators for Acoustic Wave Modeling”. In: *Geophysics* 72.5, SM95–SM106.

Seriani, Géza and Enrico Priolo (1994). “Spectral Element Method for Acoustic Wave Simulation in Heterogeneous Media”. In: *Finite Elements In Analysis And Design*. Special Issue Selection of Papers Presented at ICOSAHOM’92 16.3, pp. 337–348.

Serón, F. J., J. Badal, and F. J. Sabadell (1996). “A Numerical Laboratory for Simulation and Visualization of Seismic Wavefields”. In: *Geophysical Prospecting* 44.4, pp. 603–642.

Serón, F. J. et al. (1990). “Finite-Element Method for Elastic Wave Propagation”. In: *Communications In Applied Numerical Methods* 6.5, pp. 359–368.

Smith, Warwick D. (1975). “The Application of Finite Element Analysis to Body Wave Propagation Problems”. In: *Geophysical Journal International* 42.2, pp. 747–768.

Strong Motion Center (2018). *Berkeley Earthquake of 04 Jan 2018*. URL: https://www.strongmotioncenter.org/cgi-bin/CESMD/iqr_dist_DM2.pl?IQRID=Berkeley_04Jan2018_72948801&SFlag=0&Flag=2 (visited on 12/29/2021).

Szenicer, Alexandre, Kuangdai Leng, and Tarje Nissen-Meyer (2020). "A Complexity-driven Framework for Waveform Tomography with Discrete Adjoints". In: *Geophysical Journal International* 223.2, pp. 1247–1264.

Takeuchi, Nozomu and Robert J. Geller (2000). "Optimally Accurate Second Order Time-domain Finite Difference Scheme for Computing Synthetic Seismograms in 2-D and 3-D Media". In: *Physics Of The Earth And Planetary Interiors* 119.1, pp. 99–131.

Takeuchi, Nozomu, Robert J. Geller, and Phil R. Cummins (2000). "Complete Synthetic Seismograms for 3-D Heterogeneous Earth Models Computed Using Modified DSM Operators and Their Applicability to Inversion for Earth Structure". In: *Physics Of The Earth And Planetary Interiors* 119.1, pp. 25–36.

Tesoniero, Andrea et al. (2020). "Full Wave Sensitivity of SK(K)S Phases to Arbitrary Anisotropy in the Upper and Lower Mantle". In: *Geophysical Journal International* 222.1, pp. 412–435.

Toomey, Aoife and Christopher J. Bean (2000). "Numerical Simulation of Seismic Waves Using a Discrete Particle Scheme". In: *Geophysical Journal International* 141.3, pp. 595–604.

Tromp, Jeroen and F. A. Dahlen (1990). "Summation of the Born Series for the Normal Modes of the Earth". In: *Geophysical Journal International* 100.3, pp. 527–533.

Turner, M. J. et al. (1956). "Stiffness and Deflection Analysis of Complex Structures". In: *Journal Of The Aeronautical Sciences* 23.9, pp. 805–823.

U.S. Geological Survey (USGS) (2018a). *3-D Geologic and Seismic Velocity Models of the San Francisco Bay Region*. URL: https://www.usgs.gov/natural-hazards/earthquake-hazards/science/3-d-geologic-and-seismic-velocity-models-san-francisco (visited on 12/01/2021).

— (2018b). *ShakeMap: M 4.4 - 2km SE of Berkeley, CA*. URL: https://earthquake.usgs.gov/earthquakes/eventpage/nc72948801/shakemap/intensity (visited on 12/16/2021).

Van Driel, M. et al. (2020). "Accelerating Numerical Wave Propagation Using Wavefield Adapted Meshes. Part I: Forward and Adjoint Modelling". In: *Geophysical Journal International* 221.3, pp. 1580–1590.

Vigh, Denes et al. (2019). "Is the Salt-related Full-waveform Inversion Sorted Out?" In: *SEG Technical Program Expanded Abstracts 2019*. Society of Exploration Geophysicists, pp. 1265–1269.

Virieux, J. and S. Operto (2009). "An Overview of Full-waveform Inversion in Exploration Geophysics". In: *Geophysics* 74.6, WCC1–WCC26.

Virieux, Jean (1984). "SH-Wave Propagation in Heterogeneous Media: Velocity-Stress Finite-Difference Method". In: *Geophysics* 51.4, pp. 889–901.

— (1986). "P-SV Wave Propagation in Heterogeneous Media: Velocity-Stress Finite-Difference Method". In: *Geophysics* 51.4, pp. 889–901.

Virieux, Jean, Henri Calandra, and René-Édouard Plessix (2011). "A Review of the Spectral, Pseudo-Spectral, Finite-Difference and Finite-Element Modelling Techniques for Geophysical Imaging: Spectral, Pseudo-Spectral, Finite-Difference and Finite-Element Modelling". In: *Geophysical Prospecting* 59.5, pp. 794–813.

Virieux, Jean and Raul Madariaga (1982). "Dynamic Faulting Studied by a Finite Difference Method". In: *Bulletin Of The Seismological Society Of America* 72.2, pp. 345–369.

Virieux, Jean et al. (2009). "Seismic Wave Modeling for Seismic Imaging". In: *The Leading Edge* 28.5, pp. 538–544.

Waldhauser, Felix and William L. Ellsworth (2002). "Fault Structure and Mechanics of the Hayward Fault, California, from Double-difference Earthquake Locations". In: *Journal Of Geophysical Research: Solid Earth* 107.3, pp. 1–15.

Warner, Michael et al. (2013). "Anisotropic 3D Full-Waveform Inversion". In: *Geophysics* 78.2, R59–R80.

Watt, Janet et al. (2016). "Missing Link between the Hayward and Rodgers Creek Faults". In: *Science Advances* 2.10, e1601441.

Woodhouse, J. H. and F. A. Dahlen (1978). "The Effect of a General Aspherical Perturbation on the Free Oscillations of the Earth". In: *Geophysical Journal Of The Royal Astronomical Society* 53.2, pp. 335–354.

Xia, Muming et al. (2017). "A General 3D Lattice Spring Model for Modeling Elastic Waves". In: *Bulletin Of The Seismological Society Of America* 107.5, pp. 2194–2212.

Yang, Jidong et al. (2017). "2D Isotropic Elastic Gaussian-beam Migration for Common-shot Multicomponent Records". In: *Geophysics* 83.2, S127–S140.

Zergani, Sara, Z A Aziz, and K K Viswanathan (2016). "Elastic and Seismic Model for the Generation of Tsunamis Via Lattice Boltzmann Method". In: *Global Journal Of Pure And Applied Mathematics* 12.3, pp. 1979–1999.