


Global wave propagation in three-dimensional aspherical Earth models



Kuangdai Leng
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2017

Dedicated to my wife, Yanni

Acknowledgements

I owe my deepest thanks to my supervisor Prof Tarje Nissen-Meyer, who has been an excellent mentor and friend of mine for my three years at Oxford. Originated from Tarje's DPhil work at Princeton University, the basic concept of this thesis was hatched during a lunch meeting between Tarje and me at Wolfson college in Michaelmas 2014. It was then followed by almost three years of theoretical and technical developments, during which Tarje had been extensively involved in each and every stage, offering constant guidance and critical comments that proved fundamental to the accomplishment of this thesis. Tarje always listened to my opinions carefully, and provided pertinent and inspiring suggestions based on utmost respect to my original ideas. Working with Tarje brings me heart-felt pleasure. I am also grateful to my co-supervisor Prof Karin Sigloch, who had spent plenty of time on establishing into my mind a broader and deeper background in seismology and geophysics, and brought forward insightful remarks to my thesis work from a different perspective. Besides academic supervision, Tarje and Karin have always tried to help me in all aspects, such as finding financial supports, getting into the community, developing teaching skills, and planning early research career. It is a gift for a research student to have a supervisor of both strong academic background and respectable personality, and I am lucky to have two of such.

I also owe a debt of gratitude to my collaborators for their indispensable contributions. Dr Martin van Driel from ETH Zurich developed the mesher for my numerical method; Dr Kasra Hosseini from Oxford offered me great help in data processing and analysis; and Prof David Al-Attar from Cambridge introduced to me a novel theory from his group that

notably enhance the applicability of my method. My code development has been facilitated by communications with the developers of the Salvus project, including Michael Afanasiev, Christian Boehm, Lion Krischer, Dave May, Max Rietmann and Andreas Fichtner, to whom I express my sincere thanks. Computations have been undertaken on the UK National Supercomputing Service (ARCHER), whose fast and accurate technical support is highly appreciated.

My very special thanks are given to my dearest colleagues and friends from the Oxford seismology group, a bunch of enthusiastic and brilliant young people from whom I have received countless help and with whom I have shared countless moments of joy and comfort. I have cheerfully witnessed the robust growth of the group, from four people (besides Tarje, Karin, and myself, we have Maria Tsekhmistrenko as a DPhil student) to more than fifteen, in a short three-year period. I feel blessed to be part of this amazing group, and more, I feel so honoured to become the first DPhil coming out from it. Paula, Kara, Ben×2, Kasra, Harriet, Afsaneh, Marta, Maria, Alex, Eoghan, Kesia, Will, Judit, Natalia, and Chris, your names and faces will always stay on my mind, and I wish you all the best for each and every one of your future endeavours and shining days. Furthermore, I appreciate all help and encouragements from the faculty and the admin and technical staff from my department and my college, especially the head of graduate school Prof Chris Ballentine and my college advisor Prof Ros Rickaby.

I would love to say a big thank you to my family. My parents have been struggling to support me to finish this degree, both financially and mentally. Both of them kept working after reaching retirement age. On March 16th, 2015, my wife gave birth to our first child, Wenchu, and since then she has been filling and accompanying the little boy on her own, with enormous help from her parents. She has endured much more than I do for all these years, and deserves much more than a thesis dedication.

Finally, I thank Prof Heiner Igel and Prof Richard Katz, the two examiners of this thesis. Their thorough critical reading and helpful comments

greatly improve the completeness and accuracy of this thesis.

Every DPhil thesis means not only a phased accomplishment of an academic pursuit but also a new understanding of life. This feels especially true to me who is completing the second one. My literary talent is too mediocre to summarise within a few words my gains at Oxford, but a simple way I found is by answering *yes* to the following questions:

Did the forms of active men rise up living in thy soul; were thy breast warmed by a sympathetic fire; did the vocation which proceeds from within diffuse itself over all thy frame; were the tones of thy voice, the words of thy mouth, delightful to hear; didst thou feel thy own being sufficient for thyself? – Wilhelm Meister’s Apprenticeship, J. W. von Goethe

Abstract

Earth structure and dynamics are largely inferred by seismology as the main tool for data-informed probing of Earth’s interior. Seismic wave propagation in physically realistic Earth models is one of the most fundamental topics in seismology, not only essential to the analysis and interpretation of observed ground motions in terms of their source characteristics and structure-induced propagation effects, but also indispensable for seismic inverse techniques that deliver Earth models by minimising misfits between observations and theoretical predictions. Nevertheless, seismic wave propagation at a global scale still remains as one of the most challenging problems in scientific computing, because of Earth’s multiscale constituents and the observable frequency band of global seismic data for their resolution.

In this thesis, I present a new, computationally efficient numerical approach to simulate global seismic wave propagation in realistic three-dimensional (3-D) Earth models. It is a hybrid method of spectral element method (SEM) and pseudo-spectral method, that is, the azimuth dimension of a 3-D wavefield is characterised in terms of a global Fourier series parameterisation, such that the 3-D wave equation reduces to an algebraic system of coupled 2-D meridian equations, which is then solved by a 2-D SEM named AxiSEM3D. Computational efficiency of such a hybrid method stems from the inherent azimuthal smoothness of 3-D global wavefields, resulted from the predominance of long-wavelength heterogeneity in Earth’s mantle and the axial singularity of a point seismic source.

AxiSEM3D allows for material heterogeneities such as velocity, density, anisotropy and attenuation, as well as for finite undulations on radial discontinuities, both solid-solid and solid-fluid, and thereby a variety of

aspherical Earth features such as ellipticity, topography and bathymetry, variable crustal thickness, and core-mantle boundary topography. Such interface undulations are equivalently interpreted as material perturbations of the contiguous media, based on the “particle relabelling transformation”. Ocean is currently modelled as a hydrodynamic load.

Benchmarked in reference to a discretised 3-D SEM for a variety of Earth models, including 3-D mantle and crustal structures, topography and bathymetry, and ellipticity, AxiSEM3D proves to be a convergent and accurate numerical method. Efficiency comparisons suggest that AxiSEM3D can be 1 to 3 orders of magnitude faster than a conventional 3-D SEM across a period range from 10 s down to 1 s, with its speedup increasing with simulated frequency and decreasing with model complexity. It is sufficiently comprehensive to cover all considered applications at global scale, but is maximally efficient for deep Earth studies with body wave data. The observable frequency range of global seismic data (up to 1 Hz) has been achieved for wavefield modelling upon a 3-D mantle model with moderate computing resources. The MPI-based high-performance C++ code scales up to more than 10,000 cores, available open-source at <https://github.com/kuangdai/AxiSEM3D>.

Three applications are carried out, with different foci from Earth’s surface to the core-mantle boundary. First, surface waves are simulated within a state-of-the-art crustal model; the synthetics are compared to real seismic data to assess two different implementations of a 3-D crust. Second, wave scattering effects and cost impact of a localised small-scale structure with sharp and strong material contrast are investigated, considering a wide range of structure-wavelength combinations. And last, diffracted waves through an ultra-low velocity zone atop the core-mantle boundary are modelled at a 1 Hz frequency.

Contents

1	Introduction	1
1.1	Forward waveform modelling	3
1.2	Aspherical Earth	7
1.3	Seismic tomography	8
1.4	Summary	10
2	Observations and methodology	13
2.1	Global wavefields	13
2.2	Tomographic mantle models	17
2.3	Method architecture	20
2.4	Summary	23
3	Spherical Earth	25
3.1	Theory	25
3.1.1	Equations of motion in 3-D	25
3.1.2	Axisymmetric geometry	26
3.1.3	Dimension-reduced weak forms	29
3.1.4	Fluid core	32
3.1.5	Axial boundary conditions	33
3.1.6	Seismic moment tensor	34
3.1.7	Special case: 1-D models	37
3.2	Mode coupling	39
3.2.1	Direct summation	40
3.2.2	FFT-based mode coupling	41
3.3	Implementation	42
3.3.1	Spectral element discretisation	42

3.3.2	Fourier expansion order of solution	46
3.3.3	Stability	48
3.4	Validation	51
3.4.1	Solution convergence	54
3.4.2	Frequency-dependency of n_u	55
3.4.3	Source influence	59
3.5	Summary	59
4	Aspherical Earth	67
4.1	Theory	67
4.1.1	Particle relabelling transformation	68
4.1.2	Aspherical fluid core	71
4.1.3	Moment tensor within undulating range	72
4.1.4	Ocean load	74
4.2	Validation	75
4.2.1	Ellipticity	75
4.2.2	Topography and bathymetry	77
4.3	Summary	79
5	Computational Efficiency	81
5.1	Tomographic mantle models	82
5.2	General cost and speedup	83
5.3	Parallelisation	90
5.4	Summary	93
6	Applications	95
6.1	Surface waves in 3-D crust	95
6.1.1	Crust implementation	96
6.1.2	Comparison to data	98
6.2	Localised heterogeneity	101
6.3	Simulating a ULVZ at 1 Hz	110
6.4	Summary	114

7	Discussions	119
7.1	Conclusions	119
7.2	Outlook	121
7.2.1	Fluid ocean	121
7.2.2	Seismic inversion	121
7.2.3	Regional to local scale	122
7.2.4	Performance enhancement	122
7.2.5	Usability	123
7.2.6	Rotation and self-gravitation	124
	Bibliography	126

List of Figures

1.1	Three-dimensional global wave propagation in 1-D and 3-D Earth models	4
2.1	Wavefield snapshots in 1-D model PREM and 3-D model S40RTS due to a monopole earthquake	15
2.2	Fourier series approximations of displacement along equatorial great circle of Fig. 2.1d	16
2.3	The Fourier expansion order fields and power spectra of three tomographic mantle models	21
2.4	Architecture of the pseudo-spectral/spectral-element method	24
3.1	Sketch of an Earth model with axisymmetric geometry	27
3.2	Source patterns of an axial seismic moment tensor	38
3.3	Workflow of one time step in AxiSEM3D	43
3.4	Mesh architecture and geometric mapping	45
3.5	Examples of empirical Fourier expansion order fields	49
3.6	Examples of element distributions with respect to elemental Fourier expansion order	50
3.7	Source-centred receiver network	52
3.8	Convergence of global average and maximum TF-misfits between SPECFEM and AxiSEM3D solutions	56
3.9	Convergence behaviour of the Fourier spectral solution	57
3.10	Comparisons of synthetic seismograms (vertical component) for benchmark between AxiSEM3D and SPECFEM	62
3.11	Comparisons of synthetic seismograms (transverse component) for benchmark between AxiSEM3D and SPECFEM	64
3.12	Map of earthquake sources	65

4.1	An illustration of geometric mapping from a selected reference spherical configuration to the physically undulating configuration	69
4.2	Source and stations for ellipticity and topography/bathymetry benchmark	76
4.3	Comparisons of synthetic seismograms (vertical component) for ellipticity benchmark.	78
4.4	Comparisons of synthetic seismograms (vertical component) for topography and bathymetry benchmark	80
5.1	Computational cost of AxiSEM3D using direct summation for mode coupling	84
5.2	Fields of $n_u^{\text{learn}}(s, z)$ obtained by wavefield learning for global mantle and crustal models	88
5.3	Computational cost of AxiSEM3D and its speedup with respect to SPECFEM	89
5.4	An example of domain decomposition	91
5.5	Scaling behaviours of AxiSEM3D on a Cray XC30 supercomputer (ARCHER, UK)	92
6.1	Moho discontinuity of Crust 1.0 and its implementation in SPECFEM	97
6.2	Source and stations for assessment of 3-D crust implementations . . .	99
6.3	Comparisons of synthetics to data (vertical component) on four selected record sections for assessment of 3-D crust implementations . .	103
6.4	Comparisons of synthetics to data (vertical component) on four selected stations for assessment of 3-D crust implementations	105
6.5	Problem setup for a spherical slow blob placed in the mid-mantle of S40RTS	106
6.6	The blob-induced relative cost of AxiSEM3D with respect to the structure-wavelength ratio ζ	108
6.7	Fields of $n_u^{\text{learn}}(s, z)$ obtained by wavefield learning for a spherical slow blob placed in the mid-mantle of S40RTS	109
6.8	Comparisons of synthetics (vertical component) for a spherical slow blob placed in the mid-mantle of S40RTS	111

6.9	Problem setup for a cylinder-shaped ULVZ on CMB of PREM	113
6.10	Synthetic seismograms (radial component) on the in-plane record section for a cylinder-shaped ULVZ atop CMB	116
6.11	Synthetic seismograms (radial component) on the off-plane record section for a cylinder-shaped ULVZ atop CMB	117
7.1	Original and differential wavefields on two PREM models with and without a slow heterogeneity in the lowermost mantle	125

List of Tables

3.1	Average TF-misfits (%) at different seismic periods.	61
3.2	Average TF-misfits (%) for different earthquake sources.	61
4.1	Simulation parameters for ellipticity and topography/bathymetry bench- mark.	76
5.1	Computational cost and speedup of AxiSEM3D using direct summa- tion for mode coupling.	84
5.2	Simulation parameters for efficiency comparisons between AxiSEM3D and SPECFEM.	86
5.3	Computational cost of AxiSEM3D and its speedup with respect to SPECFEM.	89
6.1	Statistics of cross-correlations between synthetics and data for assess- ment of 3-D crust implementations.	102
6.2	Simulation parameters for a spherical slow blob placed in the mid- mantle of S40RTS.	106

Chapter 1

Introduction

For more than a century, global seismology has served as the primary tool for data-informed inference of Earth structure and dynamics. Since the discovery of Earth’s core in early 1900s [1], seismologists laid the foundation to understand Earth’s interior by establishing spherically layered, one-dimensional (1-D) Earth profiles based on seismic traveltimes, a measurement from ray theory [2] that has so far played a central role in seismological investigations of the interior. Until the late 1900s, several 1-D Earth models have been widely accepted, such as PREM [3], IASP91 [4] and AK135 [5]. Most of these 1-D Earth models agree on the existence and average depths of some major radial discontinuities [6], such as the Moho discontinuity (~ 35 km) that separates the crust from the mantle, the transition zone (between ~ 410 km and ~ 660 km) that divides the mantle into the upper and lower parts, and ~ 2740 km that bounds the D” layer atop the core-mantle boundary (CMB). Such first-order layered structures are consistent with the majority of traveltime data within a small margin of error [7].

Starting from 1-D reference Earth models, subsequent efforts have been focused on three-dimensional (3-D) lateral heterogeneities at various scales, with abundant implications for our understanding of plate tectonics and mantle convection. The propagation speeds of the primary (P-) and the shear (S-) waves are the most accessible and informative 3-D seismic properties, for which plentiful models have been proposed. The EMC (Earth Model Collaboration) database on IRIS (ds.iris.edu/ds/products/emc-earthmodels [8]) collects quite a few models of such kind, from global to finer regional scales. Most global mantle models reveal a consensus 3-D picture dominated by long-wavelength thermal and compositional heterogeneities

[9, 10, 11, 12]. The lowermost mantle, for instance, is dominated by the large, low-shear-velocity provinces (LLSVPs) [13, 14], rooted from the CMB and extending several thousands of kilometres laterally beneath Africa and the Pacific. On the edges of the LLSVPs at CMB, smaller-scale flake-like patches with extremely low seismic velocities are detected, known as the ultra low velocity zones (ULVZs) [15, 16], probably from a hundred to a thousand kilometres across and up to a hundred kilometres thick. It is hypothesised that the accumulated subducted slabs that penetrate the deep mantle form the high-velocity blobs around the LLSVPs [6, 17], and, with the aid of a phase transformation from perovskite to post-perovskite [18], the partially melted slabs pool and form the ULVZs [19, 20]. Such a multiscale view of the deep mantle has key implications for the first-order patterns and driving forces of global mantle convection [6, 17, 20, 21, 22]. In addition to seismic velocities, other 3-D Earth features of interest include density [23, 24], seismic anisotropy [25, 26], attenuation [27, 28], topography and bathymetry [29, 30], and undulations on the internal interfaces such as Moho [31, 32], the transition zone boundaries [33, 34] and the CMB [35, 36]. These 3-D structures provide important complementary constraints on the thermal evolution and dynamics of the interior.

Three-dimensional Earth models are primarily delivered by an inverse procedure called seismic tomography [2, 12, 37], which determines one or a series of best-fit 3-D models by minimising the misfit between observations and theoretical predictions. By far, most achievements in tomography have started from a spherically symmetric view and have utilised a certain approximation of the wave physics (or a measurement simplified from full waveforms), such as ray theory [2] (e.g., [38, 39]), normal-mode summation [40] (e.g., [9, 39]), and finite-frequency theory [41] (e.g., [38, 42, 43]). Nevertheless, global wave propagation always remains as one of the most fundamental pursuits in seismology, which seeks to compute the full elastodynamic response of a physically realistic 3-D Earth model to an earthquake source.

Wave propagation is the only comprehensive approach to study wave physics, especially important to non-trivial propagation phenomena such as diffraction [43, 44] and triplication [45], as well explained in Figure 1 of [44]. The fully-captured wave scattering effects are not only helpful in data interpretation, for example, for identification of structure-induced anomalies, but also essential to Earth model assessment,


especially in regard to model uncertainties and methodological limitations. It also bears the next-generation inverse techniques based on full waveforms (e.g., [46, 47]). To probe deeper into the problem, let us first look at forward waveform modelling and then seismic tomography.

1.1 Forward waveform modelling


To analyse and understand observed seismograms in terms of their source characteristics or structure-induced propagation effects, one needs to solve the forward problem of seismic wave propagation. As a first example, Figure 1.1 visualises earthquake waves propagating in two kinds of Earth models: a simple 1-D model whose geometry and material are spherically symmetric with respect to Earth’s centre and a complex 3-D model with lateral heterogeneities in its crust and mantle.

Many numerical methods have been developed for seismic wave propagation in 3-D Earth media [49], such as finite difference [50], spectral element [51, 52, 53, 54], and discontinuous Galerkin [55], with different emphases and scopes of application. Global seismology has benefited from a couple of community software packages. SPEC-FEM3D_GLOBE (SPECFEM) [51, 52] is the most mature and comprehensive spectral element method (SEM) hardcoded for Earth, encompassing 3-D mantle and crust, topography and ocean, ellipticity, rotation and gravity. Salvus [56] is a recently released 2-D/3-D SEM-based toolkit for waveform modelling and inversion from laboratory to global scales, with particular emphasis on usability and extensibility. CSEM [57] and SES3D [58] tackle 3-D media in spherical geometry, aimed mainly at mantle studies. These discretised 3-D methods such as SPECFEM and Salvus are capable of handling complex 3-D structures in arbitrary geometry, and thus may be considered “full” and “exact”, as long as the Earth model can be properly discretised.

Unfortunately, despite the surge in both hardware and software developments, global wave propagation still remains as one of the most challenging problems in scientific computing. Discretised 3-D methods come at a persistently high computational cost [59, 60, 61] that still obliterates any possibility to cover the observable frequency band of global seismic data (up to 1 Hz) with affordable computing resources. The multiscale nature of Earth’s 3-D interior motivates high-frequency simulations. For example, to resolve a mantle plume several hundred kilometres in width, a seismic



(a) 3-D wave in 1-D Earth model



(b) 3-D wave in 3-D Earth model

Figure 1.1. Three-dimensional global wave propagation in 1-D and 3-D Earth models. The source is a shallow earthquake beneath Virginia, US. The snapshots are taken at 1200 s after the source origin. The wavefields are coloured by the norm of the ground motion. Only surface waves are visible. In (a), the model is PREM, spherically symmetric with respect to Earth’s centre, and the resultant wavefield has a simple radiation pattern in the azimuthal direction because of no lateral energy scattering. In (b), the model consists of a 3-D mantle and a 3-D crust, so the resultant wavefield has a very complex radiation pattern; for example, one can observe significant energy trapped near the ocean-continent boundary and scattered within the continental crust. These snapshots are computed by AxiSEM3D [48], a numerical method developed in this thesis.

period down to 5 s is desirable, leading to a few hundred propagated wavelengths, for which one 3-D simulation may consume $10^4 \sim 10^5$ CPU-hours for a 1 hr record length. At a 1 Hz frequency, they would require a formidable $10^7 \sim 10^8$ CPU-hours per 1 hr record length, for instance, $\sim 4.2 \times 10^7$ for the global mantle model S40RTS [39], estimated on the latest UK national supercomputer ARCHER (archer.ac.uk). Concerning seismic inverse problems that involve thousands to millions of forward simulations, such 3-D methods can hardly serve as the underlying wave propagator for realistic problems with decent resolution, a good amount of data and pragmatic computing resources. The bottleneck of computing power is aggravated by an exponentially increasing number of observed broadband data (IRIS annual report 2015, iris.edu), from which abundant information is still excluded, such as a number of source-receiver pairs, amplitude information or high-frequency contents.

Tremendous efforts have been made to decrease the computational cost of discretised 3-D methods, such as GPU [59, 62], algorithmic advances such as improved time schemes [63, 64, 65], cache-optimised tensor products [66, 67], coarse-grained attenuation [68], and a truncated computational domain for regional applications, such as RegSEM [69]. Most of these efforts lead to a computational speedup of 2 to 5. This is most desirable, especially in its multiplicative effect, but does not bear “game-changing” qualities.

Computability can be eased by well-chosen approximations on the physical system. Global seismic data analysis, for instance, largely clusters around seismic traveltimes or waveforms predicted by spherically symmetric Earth models [2, 70]. Under this simplification, methodological and computational advances have remarkably progressed from ray-theoretical solutions [71] to a variety of fast dimension-reduced methods. Depending on their outcome, these dimension-reduced methods can be divided into two categories, waveform methods that produce waveforms at given receiver locations, such as DSM [72], GEMINI [73], Yspec [74], Instaseis [70] and Syngine [75], and wavefield methods that compute the whole wavefield throughout the interior, such as AxiSEM [76], an axisymmetric SEM that bears Instaseis and Syngine, and axisymmetric finite difference methods [77, 78].

If small deviations in the recorded waveforms from such a structural setting are tolerable, one may revert to computing 3-D global wavefields based on PREM-like av-

eraged 1-D models, which can be accomplished at the cost of a 2-D numerical method by exploiting axisymmetry of the wave solution [44]. At higher seismic frequencies, the dimension-reduced methods may become several orders of magnitude faster than full 3-D methods [76]. This has been viable as a basis for global tomography [2], allowing for full-wave based calculation of sensitivity kernels [35] and data-matched high frequency resolutions (e.g., 1 Hz [43]). Spherically symmetric Earth models, as summarised by [76], can thus be viewed as a pragmatic compromise for modelling between broad frequency ranges and reasonably realistic waveforms, so as to exploit a maximal amount of informative broadband data. In addition, these dimension-reduced methods can be coupled with full 3-D methods to efficiently solve multiscale problems with a complex 3-D subdomain inside a 1-D background model, sometimes referred to as the “box scheme” or “wavefield injection” [57, 79, 80, 81]. Other physical approximations may operate on material rheology, such as acoustic approximation [82], or structure-wavefield interaction by means of homogenisation [83].

The computational cost of a discretised 3-D method is largely model-agnostic. It is then implied by the tremendous speedup of dimension-reduced methods (such as AxiSEM [76]) that these 3-D methods may have lavished a large amount of computing power due to wavefield oversampling, not only for radial models but also for 3-D models with smooth, long-wavelength lateral structures. It is then natural to ask if it is possible to bridge the gap between fast methods for spherically symmetric media and slow methods for complex 3-D media by means of developing a method fully adapted to model complexity, without either simplifying the Earth models or compromising on wave physics or solution accuracy. At the same time, all hardware and algorithmic optimisations can still be applied to further enhance its computational efficiency.

With such a motivation, we brought forward a novel hybrid method combining SEM and pseudo-spectral method [48]. It parameterises the azimuthal dimension of global wavefields in terms of Fourier series, so as to exploit their inherent azimuthal smoothness resulting from the point-source singularity and low-order power spectra of global tomographic models. This method features a model-adapted computational cost: for spherically symmetric Earth models, the method naturally degenerates to AxiSEM [76]; while, for typical tomographic mantle models, it proves significantly more efficient than conventional 3-D methods, for example, with a speedup up to two

orders of magnitude at a 10 s period. More attractively, such a speedup is observed to scale with frequency, encouraging high-frequency applications (e.g., below 10 s) that accommodate numerous open geophysical questions. This fully benchmarked and operational numerical method, name AxiSEM3D, is the main contribution of this thesis, along with deep insights into the spatiotemporal complexity of global seismic wavefields as well as several numerical investigations in unprecedented parameter regimes.

1.2 Aspherical Earth

Earth is not of a precisely spherical nature, but undulates, in some cases significantly, along the surface, some of the solid-solid discontinuities such as Moho [31, 84] and the transition zone boundaries [33, 34], and some of the solid-fluid discontinuities such as the ocean floor [30] and the core-mantle boundary (CMB) [35, 36]. These undulating discontinuities must be physically honoured in certain applications because of their striking impact on synthetic waveforms. For instance, surface topography and crustal thickness variation entirely reshape surface waves, as shown in Figure 1.1, not only relevant for crust tomography [84, 85] but also significant to earthquake hazard assessment [86] and nuclear monitoring [87]. Imaging such interface perturbations can also deliver important complementary constraints on Earth’s structure and dynamics [33, 88]. From a methodological aspect, global wave propagation with undulating interfaces has never been benchmarked so far, especially for a realistic 3-D crust that poses significant challenges to numerical methods.

Even for discretised 3-D methods, handling geometric complexity is not a straightforward task. To overcome the outstanding difficulties around discretisation, two main strategies have been relied on: advanced meshing techniques such as unstructured hexahedra [54] and tetrahedra [89], with an attempt to mesh the geometry directly, and effective/equivalent media approaches such as homogenisation [83] and the particle relabelling transformation [90], which convert geometric complexities into more complicated constitutive relations while preserving the resultant wavefield. Our method takes the second route because it does not involve a 3-D mesh. At a continental to global scale, it is reasonable to assume that all undulating radial discontinuities

are topologically equivalent to a sphere (except the ocean floor). Therefore, the particle relabelling transformation is best suited for a dimension-reduced method such as AxiSEM3D, because the transformation is frequency-independent and much easier to be implemented than homogenisation. Applying the particle relabelling transformation, we generalise AxiSEM3D to aspherical geometry [91], rendering it as comprehensive as discretised 3-D methods to include all known types of 3-D complexities of Earth, while being considerably more efficient than the discretised 3-D methods throughout the entire model parameter space and frequency range of geophysical interest.

1.3 Seismic tomography

Most global tomographic models exhibit good agreement on the predominant long-wavelength features in the mantle [10, 11]. For smaller-scale structures, however, these models tend to disagree to a greater extent. Such divergence has persisted to the present day despite advances in data coverage and inverse techniques [92]. Though we seem to face a significant challenge in finding consensus, which may or may not hint at a deeper issue than refined coverage or techniques, acquisition (e.g. transportable arrays) and measurement developments in conjunction with more sophisticated tomographic methods inevitably lead to more profound understanding of observations and more reliable inference on Earth’s interior.

From a theoretical perspective, seismic tomography has been experiencing transformative progress from ray theory [2] to waveform-based inversions [93, 94]. This is partly because, first, most modern measurements of “traveltimes”, such as cross-correlation [37], time-frequency phase delays [94, 95] and instantaneous phase [96], are based upon waveforms, and second, triplicated [45] or core-diffracted phases [43, 97] and noise correlations [84] inherently necessitate waveforms as well. Therefore, fast waveform modelling across a broad frequency band becomes increasingly important as seismologists attempt to extract more information from complex seismic data.

Computational challenges stem not only from data availability but also from the requirement of resolving multi-scale structures, motivating the application of the next-generation tomographic methods. An essential step in seismic tomography is to calculate Fréchet sensitivity kernels for particular misfit functionals, for example, those for

finite-frequency traveltimes [41] and for full waveforms [37, 44]. A fundamental principle behind the “kernels” is that the structural perturbations from the background models should be sufficiently small such that the first-order Born approximation [37], or a linearised relation between the perturbations of waveforms and structures, can decently hold. This is an outstanding question within crustal tomography based on either noise correlation [84] or earthquake surface waves [85], because of the large structural perturbations within the crust.

The last decade has witnessed the trend toward tomography with kernels based on 3-D background models and full-waveform-based misfit functionals, from exploration [98, 99] to regional [46, 93, 94, 100] and global scale [47], mainly through adjoint approaches [101, 102] built upon high performance wave propagators [51, 54]. Thanks to the self-adjointness of elastodynamics, the underlying wave solvers are identical to those used for forward simulations, that is, the 3-D kernels are constructed by convolving a wavefield emanating from a seismic source with a receiver-originating adjoint wavefield that bears the misfit functional. Unfortunately, the formidable cost of wavefield modelling upon 3-D Earth models (that need to be repeated numerous times for each waveform inversion) keeps adjoint tomography far from being available as common practice, especially in view of data availability.

In addition, the ongoing surge of computing power encourages forward modelling of data, that is, searching over vast possibilities for optimised model parameters, which requires efficient wave propagation in conjunction with searching algorithms such as brute-force or Bayesian [103]. Like iterative gradient methods, these non-gradient forward methods drop the Born approximation [37] around a spherically symmetric model, and are thus preferable for strong and sharp heterogeneities. Though avoiding local minima in gradient-based methods, the non-gradient methods come at an even higher algorithmic complexity, and are thus more demanding for an efficient wave solver. By far, these methods are incapable of delivering global models due to an unaffordably large parameter space, but have found their applications to localised structures with small spatial extensions, such as ULVZs [104, 105]. A more efficient wave propagator such as AxiSEM3D will undoubtedly promote the application of both gradient and non-gradient full-waveform inversions.

Moreover, the physics of wave scattering is fundamental to tomography. Imaging

Earth’s interior relies on understanding how waveforms encode information from heterogeneous multiscale structures. To answer the central question over the inherent physical resolution limit of structures encoded in waveforms observed along Earth’s surface, one needs to identify parameter and scattering regimes where structures remain invisible as a function of seismic wavelength, structural scale and strength, and propagation path. This will aid in interpreting tomographic images by acknowledging the scope of “seismic invisibility” [83, 90], and will offer guidance for the selection of seismic data for tomography.

Forward modelling in a tomographic context, however, primarily serves to deliver synthetic waveforms and their gradients for inversion. This is of course appropriate, but depreciates a wealth of complementary information that can be obtained from the evolution of wavefield complexity. Though it has long been evident [49, 106], structure-induced wavefield complexity has never been quantitatively evaluated for deep Earth heterogeneities in a realistic 3-D setting, since such 3-D wavefields were computationally inaccessible. In addition to waveform modelling, AxiSEM3D also serves as a powerful tool to explore the relation between wavefields and structures, not only because of its computational efficiency for vast simulations, but also because the method itself is naturally adapted to model and wavefield complexity in terms of both parameterisation and computational cost. Wave scattering effects will be addressed throughout this thesis.

1.4 Summary

Driven by the rapid growth of broadband seismic data, the advances of inverse methodologies and a more profound understanding of the wave physics that carries every seismological probing of Earth’s interior, modern development of global seismology strongly motivates a new numerical method to simulate global wave propagation in full 3-D Earth models. Such a method should be comprehensive enough to consider all known types of 3-D Earth features, including both laterally heterogeneous material properties (velocity, density, full anisotropy and attenuation) and finite perturbations on any of the radial discontinuities such as the surface, Moho and CMB, while being computationally efficient to reach desirable frequency ranges that

match both the structural scales of geophysical interest and the power of modern, high-performance computing.

The remainder of this thesis strives to deliver a numerical method of such quality, featuring a remarkably improved computational efficiency compared to existing methodologies for wave propagation. Chapter 2 provides a brief but essentially integrated overview of the method, including wavefield and model observations, physical motivations and method architecture. Chapter 3 presents the theory for 3-D Earth models with spherical geometry as well as its numerical implementation and validation. These two chapters are mainly extended from [48], which lays the theoretical and computational foundations of the method. Chapter 4 extends Chapter 3 to aspherical Earth models to consider interface undulations. Chapter 5 summarises the computational cost of the method and its speedup with respect to conventional 3-D methods. Chapter 6 contains three applications of wavefield modelling. These three chapters are mainly extended from [91], which makes the prototype of [48] more comprehensive and efficient, showing a great potential of high-frequency forward studies of Earth's interior. The last chapter concludes this thesis with the main results and method features, ending up with an outlook of future work.

Chapter 2

Observations and methodology

We seek to compute the 3-D global seismic wavefield produced by a realistic earthquake source in an anisotropic, visco-elastic, 3-D Earth model. One may treat the general problem as seismic wave propagation in 3-D heterogeneous media of arbitrary geometry, and solve it with a “universal” 3-D numerical method such as SPEC-FEM3D_GLOBE (SPECFEM) [51, 52]. But one may also take a shortcut by understanding the solution characteristics first and using them to simplify the general wave equation at the outset, for instance, the fixed radiation patterns for spherically symmetric Earth models [44]. Following this idea, we shall start from direct observations of wavefield snapshots in tomographic models.

2.1 Global wavefields

For visualisation purposes, we employ an earthquake source of monopole type, that is, $M_{rr} = 1$ and the other components of the moment tensor vanish, located at a 12 km’s depth beneath Virginia. Such an earthquake source is fully axisymmetric in the source-centred coordinate system, as shown in Fig. 2.1a. See [44] or Section 3.1.6 for a strict definition of monopole sources. We use SPECFEM to compute the wavefields.

Let us first consider spherically symmetric 1-D Earth models. Evidently, the resultant wavefield should be fully axisymmetric due to the axisymmetry of both the source and the model, as illustrated by Fig. 2.1b, computed with PREM [3]. By contrast, the wavefield exhibits high complexity in the meridional slice orthogonal to the azimuthal direction, as shown in Fig. 2.1c. The theory behind this contrast will be explained in Section 3.1.6. Since the solution is rotation invariant, that is,

$u(r, \theta, \phi; t) = u(r, \theta; t)$, the 3-D problem can be rebuilt on a 2-D meridian domain (one slice) spanned by (r, θ) , and solved efficiently with a dimension-reduced method such as AxiSEM [76].

For 3-D models, the solution axisymmetry is violated by lateral heterogeneity, as illustrated by Fig. 2.1d, obtained with the tomographic model S40RTS [39]. However, even by visual inspection, one may observe that the non-axisymmetric deviations are quite “small”, that is, the equatorial pattern (Fig. 2.1d) looks much smoother than the slice pattern (Fig. 2.1e). Therefore, from a numerical perspective, it is unnecessary to describe the azimuthal variations using grid points distributed as densely as required to describe the in-slice variations. Let us instead parameterise the azimuthal dependence by Fourier series, in view of the azimuthal periodicity of solution over $\phi \in [0, 2\pi)$,¹

$$u(r, \theta, \phi; t) \simeq \sum_{|\alpha| \leq N_u} u^\alpha(r, \theta; t) \exp(i\alpha\phi), \quad (2.1)$$

where N_u denotes the global maximum number of terms needed for an accurate description, u^α the α -th order complex Fourier coefficient field, t the time and $i = \sqrt{-1}$.

The separation of variables in (2.1) is adapted to the direction-dependent variability of global wavefields, that is, a small N_u may realise an accurate description of the smooth azimuthal variations, while the 2-D coefficient fields, $u^\alpha = u^\alpha(r, \theta; t)$, are capable of describing any sharp in-slice variations. A similar idea has been proposed in the 1980s [107, 108, 109] for layered media at the local scale, but its numerical implementation could not be found. Figure 2.2 shows the Fourier approximations at the equatorial great circle, where the wavefield in Fig. 2.1d experiences the largest azimuthal variability within the snapshot. It is shown that a Fourier series with $N_u = 40$ visually yields an exact description, while full 3-D discretisation along the great circle in SPECFEM involves 4096 ($256 \times 4 \times 4$) Gauss-Lobatto-Legendre (GLL) points at the considered dominant period (17 s).

¹ In this thesis, we use $ijkl$ as spatial indices and $\alpha\beta\gamma$ as Fourier series indices. Einstein summation convention only applies to spatial indices. Superscripts acting *directly* on field variables denote Fourier series indices; for example, ρ^2 represents the second order Fourier coefficient of ρ , whereas the square of ρ , if needed, will be written as $(\rho)^2$.




Figure 2.1. Wavefield snapshots in 1-D model PREM and 3-D model S40RTS due to a monopole earthquake. These snapshots are computed by SPECFEM at a dominant period of 17 s, taken at ~ 1430 s after source origin. The colour function shows displacement norm (mm). The core regions are not plotted because displacement calculation in the fluid core is not implemented in SPECFEM. (a) shows a sketch of the source-centred coordinate system and the positions of the plotted equatorial and slice cross-sections. (b)-(e) show the snapshots on the two cross-sections for the two Earth models, respectively.




Figure 2.2. Fourier series approximations of displacement along equatorial great circle of Fig. 2.1d. N_u represents the number of terms in the Fourier series, and E_2 the L^2 the norm of fitting error. We use Discrete Fourier Transform (DFT) to compute Fourier coefficients, and then inverse DFT to compute approximated distributions after truncating the Fourier series at different N_u 's.

The Fourier characterisation, eq. (2.1), involves global expansion of functions over the ϕ -dimension, and thus falls into the category of pseudo-spectral methods, which can be equivalently interpreted by means of cardinal functions [110]. For instance, the cardinal version of eq. (2.1) reads

$$u(r, \theta, \phi; t) = \sum_{\beta=0}^{2N_u} u(r, \theta, \phi_\beta; t) C_\beta(\phi), \quad \phi_\beta = \frac{2\pi\beta}{2N_u + 1}, \quad (2.2)$$

where $C_\beta(\phi)$ are shifted sinc functions (cardinal functions associated with Fourier series) such that $C_\beta(\phi_\gamma) = \delta_{\beta\gamma}$. Points at the same ϕ_β assemble what we call slices. We do not use cardinal functions directly in this thesis. But, since they bridge the two field representation methods respectively by means of basis functions and grid points, they may facilitate the understanding of fundamental concepts (such as the mode coupling effect discussed in Section 3.2) and comparisons to discretised 3-D methods. For instance, a good approximation of Fig. 2.1d requires 41 complex terms in the Fourier series, or, interpreted as cardinal functions, 81 slices evenly spaced by $\sim 4.4^\circ$ around the axis, which are significantly less than the grid points required by the 3-D discretisation.

2.2 Tomographic mantle models

Two factors contribute to the azimuthal smoothness (small N_u) of a global seismic wavefield: axial singularity of a seismic point source and lateral smoothness of global tomographic models. In a source-centred coordinate system, the Fourier expansion of a seismic moment tensor only contains the first three terms ($\alpha = 0, \pm 1, \pm 2$), owing to the coordinate singularity of the axis, as explained in Section 3.1.6. For spherically symmetric Earth models, these three terms respectively excite monopole, dipole and quadrupole radiation patterns [44, 111], which means $N_u = 2$ throughout the entire domain. In the above numerical example, the Fourier expansion of the moment tensor only contains the zeroth order term M_{rr} , so only the monopole (rotation invariant) radiation pattern is excited and thus $N_u = 0$.

Theoretically, N_u may become infinitely large for an arbitrary 3-D model due to lateral wave scattering. However, in the case of smooth 3-D structures, the lateral scattering effects are so weak that the wavefield tend to retain the source-dominant

radiation patterns and remain smooth in the azimuthal direction, as illustrated by Fig. 2.1. A good analogy is provided here. Imagine a pebble dropped into a calm and broad lake. If the lake bottom is flat, the ripples will be fully axisymmetric with respect to the dropping point. Instead, if lake bottom is rough, the axisymmetry of the ripples will break down. However, our eyes are still cheated by quasi-concentric “circles”, because the bottom fluctuation is not strong enough to totally reshape the ripples.

Hence, the question arises as to whether Earth’s mantle is smooth enough for such an idea to stand. The answer seems to be affirmative because of the predominance of long-wavelength heterogeneity in the mantle [10], which has been well agreed among most of the recent tomographic models based on different data sets and inverse techniques [11]. To confirm the answer, we analyse the Fourier spectra of three mantle models, S362ANI [34], S20RTS [112] and S40RTS [39].

Consider the shear velocity perturbation $\delta v_s(r, \theta, \phi)$ (δv_{sh} for S362ANI) with respect to the background model. Similar to eq. (2.1), we expand it in Fourier series, that is,

$$\delta v_s(r, \theta, \phi) = \sum_{|\gamma| \leq n_v} \delta v_s^\gamma(r, \theta) \exp(i\gamma\phi), \quad (2.3)$$

where the coefficients δv_s^γ are computed by

$$\delta v_s^\gamma(r, \theta) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \delta v_s(r, \theta, \phi) \exp(-i\gamma\phi) d\phi, \quad (2.4)$$

and n_v denotes the Fourier spectral order required for a good approximation. Note that n_v should be a field in the 2-D meridian domain, $n_v = n_v(r, \theta)$, quantifying the local lateral smoothness of the model. Meanwhile, the maximum value of n_v over the 2-D domain, denoted N_v ,

$$N_v = \max_{r \in [0, R_{\text{earth}}], \theta \in [0, \pi]} n_v(r, \theta), \quad (2.5)$$

may serve as a global quantification. The spectral power associated with (2.4) is defined as

$$\sigma_s^\gamma(r, \theta) = \frac{1}{\gamma + 1} |\delta v_s^\gamma(r, \theta)|^2. \quad (2.6)$$

Following [11], we normalise the spectral power by $\gamma + 1$ to emphasise the wavelength dependence of heterogeneity.

The conventional definition of a power spectrum is based on spherical harmonics expansion [10, 11] or, in a few cases, on Fourier expansion over a sphere-to-square map [113]. In both cases, the spectral power spatially depends on depth only. Here we need to define the power spectra upon Fourier expansion over ϕ , so the spectral power is a function of both r and θ , as indicated by (2.6). It is also stressed that, because the Fourier expansion must be performed under the source-centred coordinate system, both the power spectrum and the field of n_v depend on the source location. Nevertheless, the global maximum order N_v is source invariant, which equals to the maximum degree $l_{\max}$ of the corresponding spherical harmonics expansion. For instance, $N_v = 20$ and 40 respectively for S20RTS and S40RTS, parameterised horizontally by spherical harmonics, and $N_v = 22$ for S362ANI, parameterised by cubic splines both horizontally and vertically but with an equivalent lateral resolution of $l_{\max} \approx 22$.

Given the model and a source, we first rotate the system such that the source lies on or beneath the north pole [44]. Then we compute δv_s^γ by (2.4) and σ_s^γ by (2.6), for $\gamma = 0, 1, \dots$. Eventually, n_v can be determined pointwise such that σ_s^γ becomes negligibly small ($< 10^{-4} \sigma_s^0$) when $\gamma > n_v$. The results are shown in Fig. 2.3. As indicated by Fig. 2.3a, 2.3c and 2.3e, the values of N_v are as expected, that is, $N_v = l_{\max}$. The fields of n_v are source dependent, but the governing pattern holds that $n_v \approx N_v \sin \theta$, the divergence from which reflects the non-uniformity of vertical resolution: for example, a comparison between Fig. 2.3a and 2.3c suggests that S20RTS has a more uniform vertical resolution while S362ANI places more emphasis on the uppermost mantle. Fig. 2.3b, 2.3d and 2.3f manifest concentrated spectral power on the long-wavelength Fourier modes with maximum orders lower than 8, independent of source location.

In summary, small N_v as well as long-wavelength-focused power spectra verify the lateral smoothness of the tomographic models and thus imply the viability and efficiency of the Fourier spectral characterisation over the azimuthal dimension. Certainly, the real Earth may not be as smooth as a tomographic model, because seismic tomography is naturally smoothed by a linearised relation between perturbations of model and data. However, no evidence has been proposed to support a drastically varying interior that can totally destroy the azimuthal radiation pattern. In con-

trast, the fact that global traveltime data can be well predicted by 1-D Earth models supports a smooth interior as tomography delivers.

2.3 Method architecture

Based on the above observations of global wavefields and tomographic models, we propose a hybrid numerical method combining spectral element method and pseudo-spectral method with Fourier basis functions. This section display the method architecture in a nutshell.

Let the equations of motion be loosely written as

$$\mathcal{L}u = f, \quad (2.7)$$

where the operator $\mathcal{L}$, $\mathcal{L} = \mathcal{L}(r, \theta, \phi)$, is determined by the model’s density and velocity fields, and u and f respectively denote the time-dependent displacement and source term fields, $u = u(r, \theta, \phi; t)$ and $f = f(r, \theta, \phi; t)$. We characterise the azimuthal dependency of functions by Fourier series, that is, in addition to (2.1),

$$f(r, \theta, \phi; t) = \sum_{|\beta| \leq N_f} f^\beta(r, \theta; t) \exp(\mathrm{i}\beta\phi), \quad (2.8)$$

$$\mathcal{L}(r, \theta, \phi) = \sum_{|\gamma| \leq N_v} \mathcal{L}^\gamma(r, \theta) \exp(\mathrm{i}\gamma\phi). \quad (2.9)$$

Note that we use “ $\simeq$ ” for (2.1) but “ $=$ ” for (2.8) and (2.9) because the former is an assumption that holds in an asymptotic sense as $N_u \rightarrow \infty$, whereas the latter two are exactly known a priori. The value of N_f in (2.8) has been determined for all commonly used source types in seismology, for example, 2 for point source moment tensors, 1 for point forces (used to compute the Fréchet kernels, see [44]) and 0 for explosions. Also note that N_ρ does not appear in (2.9) because we assume the density field to be smoother than the velocity field, that is, $N_\rho \leq N_v$. This is always the case in typical tomographic models, because 3-D density perturbations do not contribute to traveltime measurements.

Instead of directly solving u from $\mathcal{L}$ and f in the 3-D domain (r, θ, ϕ) , we attempt to solve the coefficients u^α from $\mathcal{L}^\gamma$ and f^β in the 2-D meridian domain (r, θ) . A critical step is to establish the relationship between u^α and $(\mathcal{L}^\gamma, f^\beta)$, that is, the






Figure 2.3. The Fourier expansion order fields and power spectra of three tomographic mantle models. In (a), (c) and (e), n_v is determined such that σ_s^γ becomes negligibly small ($< 10^{-4}\sigma_s^0$) when $\gamma > n_v$ and N_v the maximum n_v over the 2-D domain; the source is located at VIR (37.91° N, 77.93° W, Virginia). (b), (d) and (f) shows the power spectra at the point $(r, \theta) = (6291 \text{ km}, 90^\circ)$, based on different source locations. Only the first 11 orders are plotted because of the vanishing higher order values. Similar patterns can be found at other points in the 2-D domain. The source locations in (b), (d) and (f) include: VIR (37.91° N, 77.93° W, Virginia), ATL (35.38° S, 16.40° W, southern mid-Atlantic ridge) and JPN (27.94° N, 140.56° E, Bonin Islands, Japan). This figure is to be continued on the next page.



(e) n_v of S40RTS



(f) spectral power of S40RTS

Figure 2.3 (continued). The Fourier expansion order fields and power spectra of three tomographic mantle models.

dimension-reduced equations of motion, as will be accomplished in the next chapter. The basic architecture of the method is shown in Fig. 2.4.

2.4 Summary

Numerical experiments using a discretised 3-D method manifest azimuthal smoothness of global seismic wavefields in typical 3-D tomographic mantle models. This is because the lateral wave scattering effects of the smooth models are too weak to substantially change the source-dominant radiation patterns. Model spectra analysis verifies the lateral smoothness of the tomographic models. These observations naturally motivate a sparser characterisation of solution over the azimuthal dimension. Our method parameterises the azimuthal dependency of wavefields in terms of Fourier series, and solves the 3-D wavefields by means of solving their Fourier coefficient fields in a 2-D computational domain.




Figure 2.4. Architecture of the pseudo-spectral/spectral-element method. The double-lined arrow represents the workflow taken by discretised 3-D methods, while our method follows the single-lined arrows. The Fast Fourier Transform (FFT) on the left happens before time loop while the inverse FFT (IFFT) on the right is required only at the receivers, both of which take negligible computing time. This is the most fundamental and understandable workflow, which can be improved by more advanced algorithms for better computational efficiency.

Chapter 3

Spherical Earth

This chapter deals with Earth models of spherical geometry but with 3-D material heterogeneities. It is the basis of a more general theory for aspherical 3-D Earth models, as elaborated in the next chapter, but it is of practical interest in its own right, such as for body wave studies of the deep Earth. We first present the Fourier spectral theory for the reduction of the azimuthal dimension of the 3-D wave equation, and then the implementation of a 2-D SEM to solved the dimension-reduced system. We also validate the implementation by benchmarking solutions with reference to a discretise 3-D method. We considering solution convergence and the impact of seismic frequency and source location.

3.1 Theory

In this section, we establish the dimension-reduced equations of motion on a 2-D meridian domain based on the Fourier series parameterisation of the azimuthal dimension, which forms the central part of our method.

3.1.1 Equations of motion in 3-D

Consider an anisotropic elastic Earth model with spherical volume Ω and surface $\partial\Omega$. A point within Ω is referred to by its position $\mathbf{r}$, and the unit outward normal to $\partial\Omega$ is denoted $\hat{\mathbf{n}}$. Let $\mathbf{u}(\mathbf{r}; t)$ and $\mathbf{f}(\mathbf{r}; t)$ respectively denote the time-dependent fields of displacement and body force, and $\rho(\mathbf{r})$ and $\mathbf{C}(\mathbf{r})$ the fields of density and the fourth-order stiffness tensor, respectively. The time-dependent equations of motion

may be written as

$$\mathcal{L}\mathbf{u} = \mathbf{f}, \text{ in } \Omega, \quad (3.1)$$

$$\mathcal{B}\mathbf{u} = \mathbf{0}, \text{ on } \partial\Omega, \quad (3.2)$$

where the operators are defined as

$$\mathcal{L}\mathbf{u} := \rho\ddot{\mathbf{u}} - \nabla \cdot (\mathbf{C} : \nabla\mathbf{u}), \quad (3.3)$$

$$\mathcal{B}\mathbf{u} := \hat{\mathbf{n}} \cdot (\mathbf{C} : \nabla\mathbf{u}), \quad (3.4)$$

with $\ddot{\mathbf{u}}$ standing for the second-order time derivative of $\mathbf{u}$. Viscoelasticity based on the standard linear solid model can be introduced straightforwardly following [68], which is thus not included in eq. (3.1).

The homogenous Neumann boundary condition, eq. (3.2), can be efficiently handled by the following weak form

$$\langle \rho\ddot{\mathbf{u}}, \mathbf{w} \rangle_{\Omega} + a_{\Omega}(\mathbf{u}, \mathbf{w}) = \langle \mathbf{f}, \mathbf{w} \rangle_{\Omega}, \quad \forall \mathbf{w} \text{ in } \Omega, \quad (3.5)$$

where $\mathbf{w} = \mathbf{w}(\mathbf{r})$ denotes an arbitrary test function (it can be arbitrary only because no Dirichlet boundary condition is present on the surface), and the inner product and the bilinear form in Ω are defined respectively as

$$\langle \mathbf{u}, \mathbf{w} \rangle_{\Omega} := \int_{\Omega} \mathbf{u} \cdot \mathbf{w} \, d^3\mathbf{r}, \quad (3.6)$$

$$a_{\Omega}(\mathbf{u}, \mathbf{w}) := \int_{\Omega} (\nabla\mathbf{u}) : \mathbf{C} : (\nabla\mathbf{w}) \, d^3\mathbf{r}. \quad (3.7)$$

3.1.2 Axisymmetric geometry

The dimension reduction demands axisymmetric geometry of the Earth model, which means the 3-D domain Ω , along with all its internal discontinuities, is invariant by rotation around an axis. We denote the 2-D meridian domain of Ω by D , as it is D-shaped. The boundary of D contains two parts, the meridian of $\partial\Omega$, denoted ∂D , and the axis of rotation that does not have a 3-D counterpart, denoted A , as shown in Fig. 3.1. For clarity, we first exclude the fluid core and present the complete formulation for solid-fluid Earth models in Section 3.1.4.




Figure 3.1. Sketch of an Earth model with axisymmetric geometry. The right-hand side shows the 2-D meridian domain D with physical boundaries $\partial D \cup S_i$ (thick curves) and axial boundary A (dashed thick line). D is rotated about the z -axis to generate the 3-D domain Ω with boundaries $\partial\Omega \cup \Sigma_i$, as shown on the left-hand side. Here we denote the 3-D solid-fluid boundaries and their 2-D meridians by Σ_i and S_i , respectively, referred to later in Section 3.1.4.

We employ a system of cylindrical coordinates, (s, ϕ, z) or (x_1, x_2, x_3) , which is related to the geocentric spherical coordinates (r, θ, ϕ) by

$$s = r \sin \theta, \quad z = r \cos \theta. \quad (3.8)$$

We choose to work under a unit cylindrical basis, denoted $(\hat{\mathbf{s}}, \hat{\boldsymbol{\phi}}, \hat{\mathbf{z}})$ or $(\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \hat{\mathbf{g}}_3)$, which can be linked to the constant Cartesian basis $(\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3)$ by

$$\begin{cases} \hat{\mathbf{s}} = \hat{\mathbf{g}}_1 = \hat{\mathbf{e}}_1 \cos \phi + \hat{\mathbf{e}}_2 \sin \phi, \\ \hat{\boldsymbol{\phi}} = \hat{\mathbf{g}}_2 = -\hat{\mathbf{e}}_1 \sin \phi + \hat{\mathbf{e}}_2 \cos \phi, \\ \hat{\mathbf{z}} = \hat{\mathbf{g}}_3 = \hat{\mathbf{e}}_3. \end{cases} \quad (3.9)$$

We use a cylindrical coordinate system instead of a spherical one because we will need some non-spherical spectral elements to discretise the 2-D domain D by an unstructured mesh.

The gradient operator under such a frame can be written as $\nabla = \hat{\mathbf{s}}\partial_s + \frac{1}{s}\hat{\boldsymbol{\phi}}\partial_\phi + \hat{\mathbf{z}}\partial_z$. Because $\partial_\phi \hat{\mathbf{s}} = \hat{\boldsymbol{\phi}}$ and $\partial_\phi \hat{\boldsymbol{\phi}} = -\hat{\mathbf{s}}$, the gradient of a generic vector $\mathbf{v}$, $\mathbf{v} = v_j \hat{\mathbf{g}}_j$, can be shown as

$$\nabla \mathbf{v} = \left[v_{j,i} + \left(\frac{1}{s} - 1 \right) \delta_{i2} v_{j,2} - \frac{1}{s} \delta_{i2} \epsilon_{3jk} v_k \right] \hat{\mathbf{g}}_i \hat{\mathbf{g}}_j, \quad (3.10)$$

where δ_{ij} denotes the Kronecker delta, ϵ_{ijk} the Levi-Civita symbol and $v_{j,k}$ an abbreviation of $\partial_{x_k} v_j$. The above index notation is non-standard (compared to covariant and contravariant formulations) but effectively shortens tensorial equations in the weighted Sobolev space measured by $s \, ds \, dz$ [114]. It is also obvious that the unit outward normal to ∂D is always perpendicular to $\hat{\boldsymbol{\phi}}$, that is, $\hat{n}_2 \equiv 0$.

The axisymmetric geometry of Ω directly leads to the following integral equality, serving as a key to reduce the azimuthal dimension, that is, for any functions $g(\phi)$ and $h(s, z)$,

$$\int_{\Omega} g(\phi) h(s, z) s \, ds \, d\phi \, dz = \int_{-\pi}^{\pi} g(\phi) \, d\phi \int_D h(s, z) s \, ds \, dz. \quad (3.11)$$

3.1.3 Dimension-reduced weak forms

Following the basic ideas in the previous chapter, we characterise the azimuthal dependency of functions in terms of Fourier series, that is,

$$\mathbf{u} = \hat{\mathbf{g}}_i(\phi) u_i(s, \phi, z; t) \simeq \hat{\mathbf{g}}_i(\phi) \sum_{|\alpha| \leq n_u} u_i^\alpha(s, z; t) \Psi^\alpha(\phi), \quad (3.12)$$

$$\mathbf{f} = \hat{\mathbf{g}}_i(\phi) f_i(s, \phi, z; t) = \hat{\mathbf{g}}_i(\phi) \sum_{|\alpha| \leq n_f} f_i^\alpha(s, z; t) \Psi^\alpha(\phi), \quad (3.13)$$

$$\mathbf{w} = \hat{\mathbf{g}}_i(\phi) w_i(s, \phi, z) = \hat{\mathbf{g}}_i(\phi) \sum_{|\beta| \leq n_u} w_i^\beta(s, z) \Psi^\beta(\phi), \quad (3.14)$$

where $\Psi^\alpha(\phi)$, $\alpha = -n_u, \dots, n_u$, denote the Fourier modes,

$$\Psi^\alpha(\phi) := \exp(i\alpha\phi), \quad i = \sqrt{-1}, \quad (3.15)$$

which form an orthogonal set of basis functions,

$$\int_{-\pi}^{\pi} \Psi^\alpha(\phi) \Psi^{-\beta}(\phi) d\phi = 2\pi \delta_{\alpha\beta}. \quad (3.16)$$

Note that in eq. (3.14), we choose a test function space coinciding with the solution space by limiting $|\beta| \leq n_u$, as necessitated by the uniqueness of solution. Derivation of the source term coefficients f_i^α will be elaborated in Section 3.1.6.

The Fourier expansion order of solution, n_u , is a vital parameter, because it defines the wavefield resolution in the ϕ -direction and thus governs both solution accuracy and computational cost. The azimuthal complexity of a wavefield is position dependent, for example, near-axis regions demand much fewer terms than regions far away from the axis. Hence, it is reasonable to assign n_u locally, that is, $n_u = n_u(s, z)$. The scheme to determine $n_u(s, z)$ will be elaborated in Section 3.2. In the rest of this section, n_u is treated as a constant. One may understand it as the global maximum order N_u used in Section 2.1. Such a notational simplification does not affect the generality of the theory.

Substituted with eq. (3.12), the displacement gradient eq. (3.10) yields the following Fourier series form

$$\nabla \mathbf{u} \simeq \hat{\mathbf{g}}_i \hat{\mathbf{g}}_j \sum_{|\alpha| \leq n_u} u_{j;i}^\alpha(s, z; t) \Psi^\alpha(\phi), \quad (3.17)$$

where we define

$$u_{j;i}^\alpha := u_{j,i}^\alpha + \frac{1}{s} \Gamma_{ijk}^\alpha u_k^\alpha, \quad \Gamma_{ijk}^\alpha := \delta_{i2} (\mathbf{i} \alpha \delta_{jk} - \epsilon_{3jk}), \quad (3.18)$$

in view of that $u_{j,2}^\alpha \equiv 0$. Here the role of Γ_{ijk}^α is quite similar to the Christoffel symbols in covariant and contravariant derivatives. Likewise, for the test function $\mathbf{w}$, we have

$$\nabla \mathbf{w} = \hat{\mathbf{g}}_i \hat{\mathbf{g}}_j \sum_{|\beta| \leq n_u} w_{j;i}^\beta (s, z) \Psi^\beta(\phi). \quad (3.19)$$

Now we derive the dimension-reduced weak forms. Substituted with the Fourier characterisations, eqs. (3.13) and (3.14), the source term in eq. (3.5) can be reduced in the following steps:

$$\begin{aligned} \langle \mathbf{f}, \mathbf{w} \rangle_\Omega &= \int_\Omega \mathbf{f} \cdot \mathbf{w} d^3 \mathbf{r} \\ &= \int_\Omega \sum_{|\alpha| \leq n_f} f_i^\alpha \Psi^\alpha \sum_{|\beta| \leq n_u} w_i^\beta \Psi^\beta s ds d\phi dz \\ &= \sum_{|\alpha| \leq n_f} \sum_{|\beta| \leq n_u} \int_D f_i^\alpha w_i^\beta s ds dz \int_{-\pi}^{\pi} \Psi^\alpha \Psi^\beta d\phi \\ &= 2\pi \sum_{|\beta| \leq n_u} \int_D f_i^{-\beta} w_i^\beta s ds dz, \end{aligned} \quad (3.20)$$

where the third line is based on the integral equality eq. (3.11) and the fourth line the orthogonality condition eq. (3.16). Likewise, the mass and stiffness terms are respectively reduced by

$$\begin{aligned} \langle \rho \ddot{\mathbf{u}}, \mathbf{w} \rangle_\Omega &= \int_\Omega \rho \ddot{\mathbf{u}} \cdot \mathbf{w} d^3 \mathbf{r} \\ &= \int_\Omega \sum_{|\gamma| \leq n_\rho} \rho^\gamma \Psi^\gamma \sum_{|\alpha| \leq n_u} \ddot{u}_i^\alpha \Psi^\alpha \sum_{|\beta| \leq n_u} w_i^\beta \Psi^\beta s ds d\phi dz \\ &= \sum_{|\gamma| \leq n_\rho} \sum_{|\alpha| \leq n_u} \sum_{|\beta| \leq n_u} \int_D \rho^\gamma u_i^\alpha w_i^\beta s ds dz \int_{-\pi}^{\pi} \Psi^\gamma \Psi^\alpha \Psi^\beta d\phi \\ &= 2\pi \sum_{|\alpha| \leq n_u} \sum_{|\beta| \leq n_u} \int_D \rho^{-(\alpha+\beta)} u_i^\alpha w_i^\beta s ds dz, \end{aligned} \quad (3.21)$$

and

$$\begin{aligned}
a_{\Omega}(\mathbf{u}, \mathbf{w}) &= \int_{\Omega} (\nabla \mathbf{u}) : \mathbf{C} : (\nabla \mathbf{w}) d^3 \mathbf{r} \\
&= \int_{\Omega} \sum_{|\gamma| \leq n_c} C_{ijkl}^{\gamma} \Psi^{\gamma} \sum_{|\alpha| \leq n_u} u_{i;j}^{\alpha} \Psi^{\alpha} \sum_{|\beta| \leq n_u} w_{k;l}^{\beta} \Psi^{\beta} s ds dz \\
&= \sum_{|\gamma| \leq n_c} \sum_{|\alpha| \leq n_u} \sum_{|\beta| \leq n_u} \int_D C_{ijkl}^{\gamma} u_{i;j}^{\alpha} w_{k;l}^{\beta} s ds dz \int_{-\pi}^{\pi} \Psi^{\gamma} \Psi^{\alpha} \Psi^{\beta} d\phi \\
&= 2\pi \sum_{|\alpha| \leq n_u} \sum_{|\beta| \leq n_u} \int_D C_{ijkl}^{-(\alpha+\beta)} u_{i;j}^{\alpha} w_{k;l}^{\beta} s ds dz.
\end{aligned} \tag{3.22}$$

Due to the arbitrariness of $\mathbf{w}$, we may drop the mode summations over β in eqs. (3.20)-(3.22), and finally obtain the dimension-reduced weak forms in the 2-D meridian domain D , with the ϕ -dimension reduced,

$$\sum_{|\alpha| \leq n_u} \left[\langle \rho^{-(\alpha+\beta)} \ddot{\mathbf{u}}^{\alpha}, \mathbf{w}^{\beta} \rangle_D + a_D^{-(\alpha+\beta)}(\mathbf{u}^{\alpha}, \mathbf{w}^{\beta}) \right] = \langle \mathbf{f}^{-\beta}, \mathbf{w}^{\beta} \rangle_D, \tag{3.23}$$

$\forall \mathbf{w}^{\beta} \text{ in } D, \forall \beta \in \{-n_u, \dots, n_u\},$

where the inner product and the bilinear form in D , in relation to eq. (3.6) and eq. (3.7), are defined respectively as

$$\langle \mathbf{u}^{\alpha}, \mathbf{w}^{\beta} \rangle_D := \int_D u_i^{\alpha} w_i^{\beta} s ds dz, \tag{3.24}$$

$$a_D^{\gamma}(\mathbf{u}^{\alpha}, \mathbf{w}^{\beta}) := \int_D u_{i;j}^{\alpha} C_{ijkl}^{\gamma} w_{k;l}^{\beta} s ds dz. \tag{3.25}$$

The parameters involved in eq. (3.23), ρ^{γ} and C_{ijkl}^{γ} , are respectively the γ -th order Fourier coefficients of ρ and C_{ijkl} , that is,

$$\rho^{\gamma}(s, z) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \rho(s, \phi, z) \Psi^{-\gamma}(\phi) d\phi, \tag{3.26}$$

$$C_{ijkl}^{\gamma}(s, z) = \frac{1}{2\pi} \int_{-\pi}^{\pi} C_{ijkl}(s, \phi, z) \Psi^{-\gamma}(\phi) d\phi, \tag{3.27}$$

or, inversely,

$$\rho(s, \phi, z) = \sum_{|\gamma| \leq n_{\rho}} \rho^{\gamma}(s, z) \Psi^{\gamma}(\phi), \tag{3.28}$$

$$C_{ijkl}(s, \phi, z) = \sum_{|\gamma| \leq n_c} C_{ijkl}^{\gamma}(s, z) \Psi^{\gamma}(\phi). \tag{3.29}$$

where n_{ρ} and n_c are the maximum Fourier expansion orders required for an accurate description of the model, defined pointwise like n_u . Apparently, for each term in

eq. (3.23), the summation of the superscripts should always be 0, as required by the orthogonality condition eq. (3.16).

With β evaluated between $[-n_u, n_u]$, eq. (3.23) yields $2n_u + 1$ coupled weak forms to constrain the $2n_u + 1$ Fourier coefficients $\mathbf{u}^\alpha$. We shall discuss the implementation of a 2-D SEM to solve eq. (3.23) in Section 3.3. The structure of the dimension-reduced system is discussed in Section 3.2, based on a more intuitive strong form equivalent to eq. (3.23).

3.1.4 Fluid core

The Earth model should contain a couple of solid-fluid boundaries, collectively denoted Σ , as shown in Fig. 3.1. The method requires axisymmetry of Σ such that it can be generated by rotating the meridian arc S about the axis A .

We follow a potential formulation in the fluid domain Ω_f [44, 115], assuming incompressibility of the outer core. Let $\mathbf{u} = \nabla\chi/\rho$, such that in Ω_f , the 3-D equations of motion, eq.(3.1), can be recast as

$$\frac{\ddot{\chi}}{\kappa} = \nabla \cdot \left(\frac{\nabla\chi}{\rho} \right), \text{ in } \Omega_f, \quad (3.30)$$

subject to the following solid-fluid boundary conditions,

$$\hat{\mathbf{n}} \cdot \mathbf{u} = \frac{\chi_{,r}}{\rho} \quad \text{and} \quad \hat{\mathbf{n}} \cdot \mathbf{C} : (\nabla\mathbf{u}) = \ddot{\chi}\hat{\mathbf{n}}, \quad \text{on } \Sigma, \quad (3.31)$$

where κ denotes the bulk modulus of the fluid media. Note that κ and ρ are spherically symmetric. The completeness of the above potential representation is guaranteed by the absence of toroidal motion within an inviscid fluid core. Accordingly, the 3-D weak form, eq.(3.5), can be generalised as

$$\langle \rho \ddot{\mathbf{u}}, \mathbf{w} \rangle_{\Omega_s} + a_{\Omega_s}(\mathbf{u}, \mathbf{w}) + \langle \ddot{\chi} \hat{\mathbf{n}}, \mathbf{w} \rangle_{\Sigma} = \langle \mathbf{f}, \mathbf{w} \rangle_{\Omega_s}, \quad \forall \mathbf{w} \text{ in } \Omega_s, \quad (3.32)$$

$$\langle \ddot{\chi}/\kappa, w \rangle_{\Omega_f} + b_{\Omega_f}(\chi, w) - \langle \mathbf{u}, w \hat{\mathbf{n}} \rangle_{\Sigma} = 0, \quad \forall w \text{ in } \Omega_f, \quad (3.33)$$

where the bilinear form in Ω_f is defined as

$$b_{\Omega_f}(\chi, w) := \int_{\Omega_f} \frac{1}{\rho} (\nabla\chi) \cdot (\nabla w) d^3\mathbf{r}. \quad (3.34)$$

Solid-fluid coupling effects are embodied by the third terms on the left-hand side of eq.(3.32) and eq.(3.33).

Similar to eq.(3.12), we characterise the fluid potential χ in terms of Fourier series,

$$\chi(s, \phi, z; t) \simeq \sum_{|\alpha| \leq n_u} \chi^\alpha(s, z; t) \Psi^\alpha(\phi), \quad (3.35)$$

$$w(s, \phi, z) = \sum_{|\beta| \leq n_u} w^\beta(s, z) \Psi^\beta(\phi). \quad (3.36)$$

Following the same steps in eqs.(3.20)-(3.22), we may extend the reduced weak forms, eq.(3.23), to solid-fluid media:

$$\begin{aligned} & \sum_{|\alpha| \leq n_u} \left[\langle \rho^{-(\alpha+\beta)} \ddot{\mathbf{u}}^\alpha, \mathbf{w}^\beta \rangle_{D_s} + a_{D_s}^{-(\alpha+\beta)} (\mathbf{u}^\alpha, \mathbf{w}^\beta) \right] \\ & + \langle \ddot{\chi}^{-\beta} \hat{\mathbf{n}}, \mathbf{w}^\beta \rangle_S = \langle \mathbf{f}^{-\beta}, \mathbf{w}^\beta \rangle_{D_s}, \quad \forall \mathbf{w}^\beta \text{ in } D_s, \quad \forall \beta \in \{-n_u, \dots, n_u\}, \end{aligned} \quad (3.37)$$

$$\begin{aligned} & \langle \ddot{\chi}^{-\beta} / \kappa^0, w^\beta \rangle_{D_f} + b_{D_f}^0(\chi^{-\beta}, w^\beta) - \langle \mathbf{u}^{-\beta}, w^\beta \hat{\mathbf{n}} \rangle_S \\ & = 0, \quad \forall w^\beta \text{ in } D_f, \quad \forall \beta \in \{-n_u, \dots, n_u\}, \end{aligned} \quad (3.38)$$

where the bilinear form in D_f is defined as

$$b_{D_f}^0(\chi^\alpha, w^\beta) := \int_{D_f} \frac{1}{\rho^0} \left(\chi_{,i}^\alpha w_{,i}^\beta - \alpha\beta \chi^\alpha w^\beta \right) s \, ds dz. \quad (3.39)$$

Equation (3.38) implies no mode coupling in the fluid domain, thanks to the spherical symmetry of the outer core, that is, $\kappa(s, \phi, z) = \kappa^0(s, z)$ and $\rho(s, \phi, z) = \rho^0(s, z)$.

3.1.5 Axial boundary conditions

Special attention must be paid to the axis A , a boundary of the 2-D computational domain D that does not have a physical counterpart in the 3-D Earth model. The *essential* boundary conditions on A are of Dirichlet type, which must be imposed on eq. (3.23) explicitly to ensure the well-posedness of the Galerkin problem.

If the functional form of the solution is analytically known, one may derive the axial conditions directly by evaluating the asymptotic behaviour of the solution as $s \rightarrow 0$ or $\theta \rightarrow 0$, for example, from the normal mode solution of 1-D models [44]. We cannot follow this path because no analytical solution has been reported for 3-D models. Alternatively, one may impose the axial conditions such that the potential energy of the system is bounded [116, 117]. In fact, because the wavefield is globally

continuous (except on solid-fluid boundaries), the boundedness of the potential energy is equivalent to the local boundedness of the displacement gradient, that is, the numerator of the second item in eq. (3.18) should vanish,

$$(\mathbb{i}\alpha\delta_{jk} - \epsilon_{3jk}) u_k^\alpha = 0, \quad j = 1, 2, 3 \quad , \quad (3.40)$$

or

$$\begin{cases} u_s^0 = u_\phi^0 = 0, & \alpha = 0; \\ u_s^1 + \mathbb{i}u_\phi^1 = 0, \quad u_z^1 = 0, & \alpha = 1; \\ u_s^\alpha = u_\phi^\alpha = u_z^\alpha = 0, & \alpha \geq 2. \end{cases} \quad (3.41)$$

Interestingly, the result turns out to be identical to 1-D models [44]. Applying L'Hôpital's rule to eq. (3.18), the displacement gradient on the axis turns into

$$u_{j;i}^\alpha = u_{j,i}^\alpha + \Gamma_{ijk}^\alpha u_{k,1}^\alpha. \quad (3.42)$$

3.1.6 Seismic moment tensor

Now we evaluate the source term with respect to a time-dependent point source moment tensor $\mathbf{M}(t)$, positioned at $\mathbf{r} = \mathbf{r}_s$. The corresponding body force field can be written as

$$\mathbf{f}(\mathbf{r}; t) = -\mathbf{M}(t) \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_s), \quad (3.43)$$

where δ represents the Dirac-delta distribution. The right-hand side of the 3-D weak form eq. (3.5) can then be written as

$$\langle \mathbf{f}, \mathbf{w} \rangle_\Omega = \mathbf{M} : \nabla \mathbf{w}(\mathbf{r}_s). \quad (3.44)$$

Note that for a point source, $\mathbf{M}$ in eq. (3.43) or eq. (3.44) is a position-independent constant. However, its cylindrical components, M_{ij} , should depend on the ϕ -coordinate, as the cylindrical basis vectors $\hat{\mathbf{g}}_i$ are ϕ -dependent. To start from constant components, we first represent $\mathbf{M}$ in the Cartesian coordinate system,

$$\mathbf{M} = \hat{M}_{ij} \hat{\mathbf{e}}_i \hat{\mathbf{e}}_j. \quad (3.45)$$

The cylindrical components $M_{ij}(\phi)$ can be linked to the constant Cartesian components $\hat{M}_{ij}$ by a rotation matrix $\mathbf{Q}(\phi)$, that is,

$$\mathbf{M} = M_{ij} \hat{\mathbf{g}}_i \hat{\mathbf{g}}_j, \quad M_{ij}(\phi) = Q_{ik}^T(\phi) \hat{M}_{kl} Q_{lj}(\phi), \quad (3.46)$$

where

$$Q_{ij}(\phi) := \hat{\mathbf{e}}_i \cdot \hat{\mathbf{g}}_j = \begin{bmatrix} \cos \phi & \sin \phi & 0 \\ -\sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (3.47)$$

Because the source azimuth ϕ_s becomes indeterminate on the axis, we consider axial and non-axial sources separately.

3.1.6.1 General sources

Substituted with eq. (3.46) and eq. (3.19), eq. (3.44) yields

$$\langle \mathbf{f}, \mathbf{w} \rangle_\Omega = M_{ij}(\phi_s) \sum_{|\beta| \leq n_u} w_{i;j}^\beta(s_s, z_s) \Psi^\beta(\phi_s). \quad (3.48)$$

Comparing eq. (3.20) to (3.23), the right-hand side of the 2-D weak forms, eq. (3.23), can be derived from eq. (3.48) simply by scaling it with $(2\pi)^{-1}$ and dropping the summation over β , that is,

$$\langle \mathbf{f}^{-\beta}, \mathbf{w}^\beta \rangle_D = \frac{1}{2\pi} M_{ij}(\phi_s) w_{i;j}^\beta(s_s, z_s) \Psi^\beta(\phi_s). \quad (3.49)$$

Without loss of generality, we assume $\phi_s = 0$, and the above equation is simplified to

$$\langle \mathbf{f}^{-\beta}, \mathbf{w}^\beta \rangle_D = \frac{1}{2\pi} \hat{M}_{ij} w_{i;j}^\beta(s_s, z_s). \quad (3.50)$$

Equation (3.50) exposes an unfavourable fact that the reduced 2-D source terms do not decay as β increases. This is quite understandable because the Dirac-delta distribution $\delta(\phi - \phi_s)$ involved in the 3-D source term has a white Fourier spectrum. Therefore, no matter how large we choose the solution order n_u , we cannot model the source exactly, which causes uncontrollable errors that may kill the numerical advantage of the pseudo-spectral methodology. Therefore, our method is currently inadequate for off-axis sources.

3.1.6.2 Axial sources

If the source is located on the axis, eq. (3.49) does not make sense because ϕ_s becomes indeterminate when $s_s = 0$. The axial source terms may be derived in various ways, among which the simplest is from an asymptotic perspective. Because the source can

be imagined approaching to the axis from all azimuths, the axial source term should equal to the integration of eq. (3.49) with respect to ϕ_s scaled by $(2\pi)^{-1}$, that is,

$$\langle \mathbf{f}^{-\beta}, \mathbf{w}^\beta \rangle_D = \frac{1}{2\pi} \hat{M}_{ij} \Lambda_{ijkl}^\beta w_{k;l}^\beta(0, z_s), \quad (3.51)$$

where Λ_{ijkl}^β is a constant fourth order tensor given by

$$\Lambda_{ijkl}^\beta = \frac{1}{2\pi} \int_{-\pi}^{\pi} Q_{ik}(\phi_s) Q_{jl}(\phi_s) \Psi^\beta(\phi_s) d\phi_s. \quad (3.52)$$

Clearly, since $\mathbf{Q}$ only involves Ψ^0 and $\Psi^{\pm 1}$, in view of the orthogonality condition eq. (3.16), Λ_{ijkl}^β should vanish if $|\beta| > 2$. Therefore, a seismic moment tensor can be *exactly* described in the 2-D domain as long as $n_u \geq 2$.

With Λ_{ijkl}^β evaluated respectively for $\beta = 0, 1, 2$ and $w_{k;l}^\beta$ substituted by the axial formulation of displacement gradient, eq. (3.42), the axial source term eq. (3.51) can be finally simplified to

- Monopole term ($\beta = 0$):

$$2\pi \langle \mathbf{f}^0, \mathbf{w}^0 \rangle_D = \left(\hat{M}_{xx} + \hat{M}_{yy} \right) w_{s,s}^0 + \hat{M}_{zz} w_{z,z}^0 \quad (3.53)$$

- Dipole term ($\beta = 1$):

$$2\pi \langle \mathbf{f}^{-1}, \mathbf{w}^1 \rangle_D = \frac{\hat{M}_{xz} - \mathbf{i} \hat{M}_{yz}}{2} (w_{s,z}^1 + 2w_{z,s}^1 + \mathbf{i} w_{\phi,z}^1) \quad (3.54)$$

- Quadrupole term ($\beta = 2$):

$$2\pi \langle \mathbf{f}^{-2}, \mathbf{w}^2 \rangle_D = \left(\frac{\hat{M}_{yy} - \hat{M}_{xx}}{2} + \mathbf{i} \hat{M}_{xy} \right) (w_{s,s}^2 + \mathbf{i} w_{\phi,s}^2) \quad (3.55)$$

In eqs. (3.53)-(3.55), all $w_{k;l}^\beta$'s are evaluated at the source location $(0, z_s)$. The above results are equivalent to the trigonometric formulations in [44]. Equations (3.53)-(3.55) may also be expressed in terms of $\mathbf{f}$, that is,

- Monopole term ($\beta = 0$):

$$-2\pi \begin{bmatrix} f_s^0 \\ f_\phi^0 \\ f_z^0 \end{bmatrix} = \begin{bmatrix} \left(\hat{M}_{xx} + \hat{M}_{yy} \right) \delta'(s) \delta(z - z_s) \\ 0 \\ \hat{M}_{zz} \delta(s) \delta'(z - z_s) \end{bmatrix} \quad (3.56)$$

- Dipole term ($\beta = 1$):

$$-2\pi \begin{bmatrix} f_s^{-1} \\ f_\phi^{-1} \\ f_z^{-1} \end{bmatrix} = \frac{\hat{M}_{xz} - \mathbf{i}\hat{M}_{yz}}{2} \begin{bmatrix} \delta(s) \delta'(z - z_s) \\ \mathbf{i}\delta(s) \delta'(z - z_s) \\ 2\delta'(s) \delta(z - z_s) \end{bmatrix} \quad (3.57)$$

- Quadrupole term ($\beta = 2$):

$$-2\pi \begin{bmatrix} f_s^{-2} \\ f_\phi^{-2} \\ f_z^{-2} \end{bmatrix} = \left(\frac{\hat{M}_{yy} - \hat{M}_{xx}}{2} + \mathbf{i}\hat{M}_{xy} \right) \delta'(s) \delta(z - z_s) \begin{bmatrix} 1 \\ \mathbf{i} \\ 0 \end{bmatrix} \quad (3.58)$$

where δ' denotes the “derivative” of Dirac-delta distribution. It is noteworthy that δ and δ' appear in $\mathbf{f}$ *only* along the in-plane directions, s and z , which essentially explains from a source perspective why the wavefields can remain smooth in the azimuthal direction (for example Fig. 2.1b and Fig. 2.1d), while experiencing rather high in-plane complexity (for example Fig. 2.1c and Fig. 2.1e).

The source terms for $\beta = 0, 1, 2$ are classified as monopole, dipole and quadrupole, respectively, based on their induced radiation patterns within an infinitesimal neighborhood of the source point. Such radiation patterns may be termed as *source patterns*, as they are intrinsic (model-independent) source properties determined by the moment tensor, similar to focal mechanisms, as summarised in Fig. 3.2. For 1-D Earth models, the global radiation patterns of wavefields strictly follow the source patterns of the same angular order [44], but deviate from them for 3-D models to some degree.

3.1.7 Special case: 1-D models

For spherically symmetric models, $n_c = n_\rho = 0$, and the 2-D weak forms, eq. (3.23), degenerate to

$$\langle \rho^0 \ddot{\mathbf{u}}^{-\beta}, \mathbf{w}^\beta \rangle_D + a_D^0(\mathbf{u}^{-\beta}, \mathbf{w}^\beta) = \langle \mathbf{f}^{-\beta}, \mathbf{w}^\beta \rangle_D, \quad \forall \mathbf{w}^\beta \text{ in } D, \quad \beta = 0, 1, 2, \quad (3.59)$$

Because the source term $\mathbf{f}^\beta$ vanishes for $\beta > 2$, as proved in Section 3.1.6, we only need to solve independently the zeroth, first and second order Fourier modes, or the monopole, dipole and quadrupole radiation patterns, as shown in Fig. 3.2. Equation (3.59) epitomises the results of [44] that lay the theoretical foundations for the axisymmetric code AxiSEM [76].




Figure 3.2. Source patterns of an axial seismic moment tensor. There are four distinct source patterns, two monopole (top), one dipole (middle) and one quadrupole. For 1-D models, the source patterns also represent the global radiation patterns. This figure is adapted from [76].

3.2 Mode coupling

The coupled system of 2-D weak forms, eq. (3.23), is efficient in handling the stress-free boundary condition on the surface and is convenient for SEM implementation, but the inherent structure of the dimension-reduced system is better revealed by an equivalent strong form, as discussed below.

Applying the divergence theorem to eq. (3.23), recalling that $\hat{n}_2 \equiv 0$ and $u_{j,2}^\alpha \equiv 0$, we obtain the reduced strong form in correspondence to eq. (3.1) and eq. (3.2), that is, $\forall \beta \in \{-n_u, \dots, n_u\}$,

$$\sum_{|\alpha| \leq n_u} \mathcal{L}^{-(\alpha+\beta)} \mathbf{u}^\alpha = \mathbf{f}^{-\beta}, \text{ in } D, \quad (3.60)$$

$$\sum_{|\alpha| \leq n_u} \mathcal{B}^{-(\alpha+\beta)} \mathbf{u}^\alpha = \mathbf{0}, \text{ on } \partial D, \quad (3.61)$$

where the source terms $\mathbf{f}^{-\beta}$ are given by eqs. (3.56)-(3.58) for seismic moment tensors, and the reduced 2-D operators are defined as

$$\mathcal{L}^\gamma \mathbf{u}^\alpha := \hat{\mathbf{g}}_i \left[\rho^\gamma \ddot{u}_i^\alpha - (C_{ijkl}^\gamma u_{k;l}^\alpha)_{;j} \right], \quad (3.62)$$

$$\mathcal{B}^\gamma \mathbf{u}^\alpha := \hat{\mathbf{g}}_i \left[\hat{n}_j (C_{ijkl}^\gamma u_{k;l}^\alpha) \right]. \quad (3.63)$$

Note that the surface tractions do not vanish on the axis A , but they do not have to be prescribed explicitly in eq. (3.61) because A is a Dirichlet boundary, implicitly conditioned by eq. (3.40).

Indeed, eq. (3.60) defines a linear system of 2-D hyperbolic PDEs with a Hermitian *Toeplitz* [118] operator matrix of size $(2n_u + 1)^2$, for instance, when $n_u = 3$,

$$\begin{bmatrix} \mathcal{L}^0 & \mathcal{L}^{-1} & \mathcal{L}^{-2} & \mathcal{L}^{-3} & \mathcal{L}^{-4} & \mathcal{L}^{-5} & \mathcal{L}^{-6} \\ \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & \mathcal{L}^{-2} & \mathcal{L}^{-3} & \mathcal{L}^{-4} & \mathcal{L}^{-5} \\ \mathcal{L}^2 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & \mathcal{L}^{-2} & \mathcal{L}^{-3} & \mathcal{L}^{-4} \\ \mathcal{L}^3 & \mathcal{L}^2 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & \mathcal{L}^{-2} & \mathcal{L}^{-3} \\ \mathcal{L}^4 & \mathcal{L}^3 & \mathcal{L}^2 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & \mathcal{L}^{-2} \\ \mathcal{L}^5 & \mathcal{L}^4 & \mathcal{L}^3 & \mathcal{L}^2 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} \\ \mathcal{L}^6 & \mathcal{L}^5 & \mathcal{L}^4 & \mathcal{L}^3 & \mathcal{L}^2 & \mathcal{L}^1 & \mathcal{L}^0 \end{bmatrix} \begin{bmatrix} \mathbf{u}^{-3} \\ \mathbf{u}^{-2} \\ \mathbf{u}^{-1} \\ \mathbf{u}^0 \\ \mathbf{u}^1 \\ \mathbf{u}^2 \\ \mathbf{u}^3 \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{f}^{-2} \\ \mathbf{f}^{-1} \\ \mathbf{f}^0 \\ \mathbf{f}^1 \\ \mathbf{f}^2 \\ \mathbf{0} \end{bmatrix}. \quad (3.64)$$

On the right-hand side, $\mathbf{f}^{-3} = \mathbf{f}^3 = \mathbf{0}$, because $n_f = 2$ for a seismic moment tensor. The non-diagonality of the operator matrix indicates the coupling of the Fourier

modes. From the viewpoint of cardinal functions, the non-diagonal operators “scatter” the wave energy from one azimuthal slice to another. Such a mode-coupling effect deserves utmost care because it could deny the viability of the method if the computational cost surged with n_u .

For a given azimuthal resolution, let N denote the number of grid points in a discretised 3-D mesh corresponding to n_u . For example, we may count both N and n_u at the equatorial great circle. We have illustrated in Section 2.1 that $n_u \ll N$ for typical tomographic models, our initial motivation for the pseudo-spectral characterisation. Nonetheless, eq. (3.60) or eq. (3.64) seems to suggest a computational complexity of $O(n_u^2)$ for the azimuthal dimension, in contrast to $O(N)$ for discretised 3-D methods. Fortunately, this is untrue thanks to the special structure of the operator matrix. The effect of mode coupling can be realised in two different ways, direct summation and fast Fourier transform (FFT), as compared below.

3.2.1 Direct summation

What is not included in eq. (3.60) is model complexity. Recalling eq. (3.28) and eq. (3.29), we note that $\rho^\gamma = 0$ for $\gamma > n_\rho$ and $C_{ijkl}^\gamma = 0$ for $\gamma > n_c$. Therefore, the operator matrix is not only Hermitian Toeplitz but also banded, with a bandwidth of $\max(2n_c + 1, 2n_\rho + 1)$. Based on our knowledge of Earth’s nature, we assume $n_c \geq n_\rho$, and eq. (3.60) turns into

$$\sum_{\substack{|\alpha| \leq n_u \\ |\alpha + \beta| \leq n_c}} \mathcal{L}^{-(\alpha + \beta)} \mathbf{u}^\alpha = \mathbf{f}^{-\beta}, \text{ in } D. \quad (3.65)$$

For instance, when $n_c = 1$, the above example eq. (3.64) reads

$$\begin{bmatrix} \mathcal{L}^0 & \mathcal{L}^{-1} & 0 & 0 & 0 & 0 & 0 \\ \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & 0 & 0 & 0 & 0 \\ 0 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & 0 & 0 & 0 \\ 0 & 0 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & 0 & 0 \\ 0 & 0 & 0 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} & 0 \\ 0 & 0 & 0 & 0 & \mathcal{L}^1 & \mathcal{L}^0 & \mathcal{L}^{-1} \\ 0 & 0 & 0 & 0 & 0 & \mathcal{L}^1 & \mathcal{L}^0 \end{bmatrix} \begin{bmatrix} \mathbf{u}^{-3} \\ \mathbf{u}^{-2} \\ \mathbf{u}^{-1} \\ \mathbf{u}^0 \\ \mathbf{u}^1 \\ \mathbf{u}^2 \\ \mathbf{u}^3 \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{f}^{-2} \\ \mathbf{f}^{-1} \\ \mathbf{f}^0 \\ \mathbf{f}^1 \\ \mathbf{f}^2 \\ \mathbf{0} \end{bmatrix}. \quad (3.66)$$

Apparently, a direction summation of the Fourier modes will involve an azimuthal cost of $O(n_c n_u)$ as long as $n_u \geq n_c$. In other words, the cost scales with n_u for a

specific model ($n_c = \text{const}$). The condition $n_u \geq n_c$ can be easily achieved for smooth tomographic mantle models with small n_c , as shown in Fig. 2.3.

The banded structure also suggests a model-dependent computational cost. For models with smooth lateral variations (small n_c), the cost can be greatly lowered by a narrow bandwidth (weak mode coupling), as well as its induced lower complexity of the resultant wavefield (small n_u).

3.2.2 FFT-based mode coupling

The direct summation approach will however cease to work efficiently if more complicated structures are involved, such as the global crustal model Crust 1.0 [31], with n_c slightly smaller than 180.

Because the operator matrix is Hermitian Toeplitz, we may alternatively implement the summation by means of fast Fourier transform (FFT), commonplace in pseudo-spectral methods for partial differential equations with spatially-varying parameters [110]. The basic idea is to compute the spatial derivatives (such as $\nabla \mathbf{u}$) in the Fourier space, but the multiplications with material properties (such as $\mathbf{C} : \nabla \mathbf{u}$) in the physical space or, to be exact, on the $2n_u + 1$ cardinal grid points evenly distributed along the ϕ -direction. FFT is responsible for efficient transitions between the Fourier and the physical spaces. Compared to the direct summation approach that has an algorithmic complexity of $O(n_u n_c)$ for $n_c \leq n_u$, and $O(n_u^2)$ for $n_c > n_u$, the FFT-based approach has a complexity of $O(n_u \log n_u)$ ¹.

In our SEM implementation, we finally choose the FFT-based mode coupling method for three reasons. First, the computational efficiency will be greatly promoted for non-smooth models with a large n_c . Second, because the cost no longer depends *explicitly* on the model, we may maximise the computational efficiency by learning the azimuthal complexity of wavefields, as explained later in Section 3.3.2. Last, FFT is highly optimised in modern numerical libraries such as FFTW3 [119], and the advantage of the direct summation approach is observed to be trivial even for

¹ In FFTW3 [119], $O(N \log N)$ is the slowest case when the logical size N is a large prime number. In AxiSEM3D, we boost the performance by rounding up the values of n_u such that N always has small prime factors.

a smooth mantle model such as S20RTS [112]. A typical time step is illustrated in Fig. 3.3.

3.3 Implementation

In this section, we introduce the implementation of a 2-D SEM to solve the dimension-reduced weak forms, eq. (3.23). We follow the axisymmetric 2-D spectral element framework AxiSEM, which has been established, tested and improved via successive efforts by [65, 67, 68, 111]. Our 3-D extension is thus named AxiSEM3D. Most of the basic constituents remain unchanged for the 3-D extension, such as spatial discretisation, derivative and quadrature, and time-marching schemes.

3.3.1 Spectral element discretisation

As shown in Fig. 3.4, the 2-D meridian domain is discretised by a finite element mesh, each element mapped onto a reference square on which all spatial operations take place. We use different types of geometric mapping ($s = s(\xi, \eta)$ and $z = z(\xi, \eta)$) for different regions of the mesh, such that the surface curvature can be honoured analytically [67]. Given the element shape, we can compute the elemental Jacobian by the following relation,

$$\mathcal{J}(\xi, \eta) = \frac{\partial(s, z)}{\partial(\xi, \eta)} = \det \begin{bmatrix} \partial_\xi s & \partial_\eta s \\ \partial_\xi z & \partial_\eta z \end{bmatrix}. \quad (3.67)$$

A generic function $f(s, z)$ on an element D_e maps onto the reference square by $f(\xi, \eta) = f(s(\xi, \eta), z(\xi, \eta))$, which is approximated by the following piecewise interpolation,

$$f(\xi, \eta) = f(s(\xi, \eta), z(\xi, \eta)) \approx \sum_{\mathbf{p}, \mathbf{q}=0}^N f_{\mathbf{pq}} l_{\mathbf{p}}^N(\xi) l_{\mathbf{q}}^N(\eta), \quad (3.68)$$

where $f_{\mathbf{pq}}$ is the value on the Gauss-Lobatto-Legendre (GLL) point $(\xi_{\mathbf{p}}, \eta_{\mathbf{q}})$, that is, $f_{\mathbf{pq}} = f(\xi_{\mathbf{p}}, \eta_{\mathbf{q}})$, and $l_{\mathbf{p}}^N$ the $\mathbf{p}$ -th Lagrange interpolating function of polynomial order N , such that $l_{\mathbf{p}}^N(\xi_{\mathbf{q}}) = l_{\mathbf{p}}^N(\eta_{\mathbf{q}}) = \delta_{\mathbf{pq}}$ (diagonal mass matrix). Spatial derivatives of




Figure 3.3. Workflow of one time step in AxiSEM3D. Two additional field variables are introduced: the stress tensor $\boldsymbol{\sigma}$ and the stiffness force $\mathbf{k}$ (the body force reacted to $\boldsymbol{\sigma}$). The FFT prefix “F2P” means “from Fourier to physical space”, and similarly for “P2F”. All 3-D operations are conducted in the physical space (double-lined), on the cardinal grid points anchored at ϕ_J , $\phi_J = 2\pi J/(2n_u + 1)$, $J = 0, 1, \dots, 2n_u$.

$f(s(\xi, \eta), z(\xi, \eta))$ can then be computed as

$$\frac{\partial f}{\partial s} \approx \sum_{p,q=0}^N \left[\frac{\partial z}{\partial \eta} \frac{\partial l_p^N(\xi)}{\partial \xi} l_q^N(\eta) - \frac{\partial z}{\partial \xi} l_p^N(\eta) \frac{\partial l_q^N(\eta)}{\partial \eta} \right] f_{pq} \mathcal{J}^{-1}, \quad (3.69)$$

$$\frac{\partial f}{\partial z} \approx \sum_{p,q=0}^N \left[\frac{\partial s}{\partial \xi} l_p^N(\eta) \frac{\partial l_q^N(\eta)}{\partial \eta} - \frac{\partial s}{\partial \eta} \frac{\partial l_p^N(\xi)}{\partial \xi} l_q^N(\eta) \right] f_{pq} \mathcal{J}^{-1}. \quad (3.70)$$

The integral of $f(s(\xi, \eta), z(\xi, \eta))$ over the element D_e can be realised by the GLL quadrature

$$\int_{D_e} f(s, z) s \, ds \, dz \approx \sum_{p,q=0}^N \sigma_p^N \sigma_q^N s(\xi_p, \eta_q) f_{pq} \mathcal{J}(\xi_p, \eta_q), \quad (3.71)$$

where the weights are determined by $\sigma_{p,q}^N = \int_{-1}^1 l_{p,q}^N(\xi) \, d\xi$. Note that eqs.(3.69)-(3.71) have to be modified to their Gauss-Lobatto-Jacobi counterparts for the axial elements [67] to handle the singularity $s = 0$.

Equations (3.67)-(3.71) fulfil the requirements for the spatial discretisation of the dimension-reduced weak forms, eq.(3.23). For instance, using the direct summation approach for mode coupling (Section 3.2.1), the discretised stiffness term in eq.(3.25) is computed by the the following steps:

1. given the Fourier coefficients of displacement u_i^α , $-n_u \leq \alpha \leq n_u$, on the GLL points of element D_e , compute the deformation gradient $u_{i,j}^\alpha$ by eq.(3.18), where the derivatives $u_{i,j}^\alpha$ are evaluated by eq.(3.69) and eq.(3.70);
2. compute the Fourier coefficients of stress pointwise, that is, $\sigma_{kl}^\beta = \sum_{|\alpha| \leq n_u} C_{ijkl}^{-(\alpha+\beta)} u_{i,j}^\alpha$, for all $\beta = -n_u, \dots, n_u$; mode coupling takes place in this step;
3. similar to the first step, compute the deformation gradient (denoted ς_{kl}^β) with respect to a discretised delta function w_k^β that equals to 1 in the x -direction ($x = s, z$) at a GLL point P and vanishes otherwise;
4. compute the contraction $\sigma_{kl}^\beta \varsigma_{kl}^\beta$ pointwise and then its elemental integral by means of the quadrature eq.(3.71); the result yields the β -th order Fourier coefficient of the stiffness term in the x -direction at the GLL point P , contributed by element D_e ;
5. assemble the elemental stiffness terms at shared GLL points to compute the global stiffness terms.




Figure 3.4. Mesh architecture and geometric mapping. In this example, the 2-D meridian domain D is discretised by a mesh with 128 elements on the outer surface, the equal of $\text{NEX} = 64$ in SPECFEM (see [60] for details of the SPECFEM parameter NEX). Each element is mapped onto a reference square spanning $-1 \leq \xi, \eta \leq 1$, on which all spatial operations take place.

In practice, the above process can be facilitated by a variety of properties and techniques, such as the sparsity of C_{ijkl} , Hermitian of the reduced system and pre-computable terms. For FFT-based mode coupling (Section 3.2.2), the second step above is replaced by the lowermost three arrows of the workflow shown in Fig. 3.3, all performed pointwise in the 2-D domain.

For the complete implementation, we refer the reader to [67] for isotropic media or [111] for anisotropic media, where the implementation of eq.(3.59), or the degenerated form of eq.(3.23) for 1-D Earth models, is elaborated. The 3-D generalisations are purely algebraic, basically nothing but adding an outer loop of Fourier expansion orders when evaluating the field variables. Such simplicity largely relies on an explicit time scheme (and thus a diagonal mass matrix) that preserves the fundamental framework.

3.3.2 Fourier expansion order of solution

The most important issue around the application of AxiSEM3D is how to construct the field $n_u = n_u(s, z)$, namely, the expansion order prescribed pointwise on the Fourier series parameterisation of solution, see eq. (3.12). For a given 3-D Earth model, a reasonable choice of $n_u(s, z)$ is critical to both solution quality and computational efficiency. This is the only input parameter required by AxiSEM3D in addition to those of a conventional 3-D SEM such as SPECSEM. Apparently, $n_u(s, z)$ should increase with the local azimuthal complexity of the wavefield, which is however unknown before simulation, influenced by a variety of factors such as model, source location and depth, and record length.

In this subsection, we provide two approaches to determine $n_u(s, z)$, an empirical equation designed for tomographic mantle models and the technique of “wavefield learning” that determines a near-optimal $n_u(s, z)$ sufficiently large for the convergence of solutions.

3.3.2.1 Empirical equation

For typical tomographic mantle models, we obtain an efficient empirical equation based on massive trial simulations, taking into account some first-order propagation

effects:

$$n_u = \max \left(\lceil n_u^{\text{ref}} F_s F_\theta F_d \rceil, n_u^{\text{min}} \right), \quad (3.72)$$

where n_u^{ref} and n_u^{min} denote respectively a basic reference value and the global minimum value, the only two parameters from users, whereas F_s , F_θ and F_d are three scale factors acting on n_u^{ref} for different purposes, and the brackets $\lceil \rceil$ mean rounding up to the next integer. Factor F_s simply reflects the fact that points away from the axis should demand higher orders (more slices) than those close to the axis (note that $n_u \equiv 1$ on the axis, see eq. (3.41)),

$$F_s = F_s(s) = \frac{s}{R_{\text{earth}}}. \quad (3.73)$$

Factor F_θ increases n_u by the epicentral distance θ , since the further the waves travel, the more 3-D structures they “see”,

$$F_\theta = F_\theta(\theta) = \begin{cases} 1, & \theta \leq \frac{\pi}{4}; \\ 1 + 4 \left(\frac{4\theta - \pi}{3\pi} \right)^3, & \frac{\pi}{4} < \theta \leq \pi. \end{cases} \quad (3.74)$$

Lastly, factor F_d enhances the resolution of surface waves,

$$F_d = F_d(d) = \begin{cases} 2, & d \leq d_0; \\ 4 - 2\frac{d}{d_0}, & d_0 < d \leq 1.5d_0; \\ 1, & d > 1.5d_0, \end{cases} \quad (3.75)$$

where d denotes the depth, $d = R_{\text{earth}} - r$, and d_0 a characteristic depth specifying the enhanced range, which must be considerably greater than the penetration depth of the longest wavelength of interest. For a safe play, we choose $d_0 = 200\text{km}$. Note that such enhancement will not increase the computational cost much because it only affects a small fraction of elements near the surface. For deep earthquake sources and applications where surface waves are not of interest, one may set $F_d = 1$. The constants in eqs.(3.73)-(3.75) may be adjusted for specific applications, for example, doubling F_d for a higher surface wave resolution. But, in general, one can simply increase/decrease n_u^{ref} for better accuracy/performance.

Fig. 3.5 shows examples of eq. (3.72) where $n_u^{\text{ref}} = 24$ and $n_u^{\text{min}} = 8$, with and without surface wave enhancement. Mapping Fig. 3.5 onto Fig. 3.4, we obtain the corresponding n_u fields defined on a discretised mesh. Note that n_u is an indicator

of computational cost, which means the elemental cost should vary from element to element. Before propagating waves, we may pre-examine the distributions of number of elements with respect to their Fourier expansion orders, such as Fig. 3.6, from which the relative cost for different choices of $n_u(s, z)$ can be roughly estimated. For instance, Fig. 3.6 indicates that the surface wave enhancement by eq. (3.75) primarily doubles the order of $\sim 20\%$ elements, and thus is less likely to produce a cost surge.

3.3.2.2 Wavefield learning

The technique of “wavefield learning” can determine a computationally efficient $n_u(s, z)$ that guarantees the convergence of solutions. The idea is straightforward. We run a trial simulation with a starting n_u field, denoted $n_u^{\text{start}}(s, z)$, which is sufficiently large to capture the wavefield accurately throughout the record length. The solver then monitors the energy in each Fourier mode over the whole time loop and identifies those high-order modes that are never excited until the end of the simulation. At the end of the simulation, the “learnt” n_u field, denoted $n_u^{\text{learn}}(s, z)$, is written to a file, which can be used repeatedly afterwards for simulations with similar parameters, such as those with similar or smoother models or nearby events. Apparently, the solution should have converged if $n_u^{\text{learn}}(s, z)$ turns out smaller than $n_u^{\text{start}}(s, z)$ pointwise, making AxiSEM3D stand alone without a reference solution from a full 3-D method. For a specific research purpose, a few times of learnings are usually enough to bound a sub-optimal $n_u(s, z)$.

The computational cost of the learning simulations may vary a lot, mainly depending on the choice of $n_u^{\text{start}}(s, z)$ and the learning frequency. Fortunately, we will show it later in Section 6.2 that learnings can be performed at a lower frequency, with the learnt results applicable to higher frequencies. Moreover, a learning simulation itself is usually less expensive than the corresponding 3-D simulation unless $n_u^{\text{start}}(s, z)$ is unwisely large. We will show plenty of applications of wavefield learning in Chapter 5 and Chapter 6.

3.3.3 Stability

Numerical stability could be another problem around the 3-D extension. The stability of the axisymmetric method has been comprehensively considered for 1-D Earth




Figure 3.5. Examples of empirical Fourier expansion order fields. It is computed by eq. (3.72) with $n_u^{\text{ref}} = 24$ and $n_u^{\text{min}} = 8$. The source is located near the top of the axis. The maximum value appears near $\theta = 135^\circ$ on the surface, a combined effect of F_s and F_θ . For deep earthquakes or the cases where surface waves are not of interest, one may use (a); otherwise, (b) should be employed to enhance surface wave resolution. Integer rounding ($\lceil \cdot \rceil$) is ignored in the plots for a clearer view.




Figure 3.6. Examples of element distributions with respect to elemental Fourier expansion order. The statistics are based on the two n_u fields shown in Fig. 3.5, with and without surface wave enhancement, which are mapped onto the mesh shown in Fig. 3.4. The surface wave enhancement primarily doubles the order of $\sim 20\%$ elements.

models [65], with special attention paid to the thin crustal layers, the internal discontinuities and the fluid outer core. When extended to 3-D, the stability condition generally becomes more strict. Though the complete theory has remained an open question for general 3-D media [120], we benefit from the smoothness of lateral heterogeneity, especially the absence of horizontal discontinuities, allowing us to follow the 1-D condition exactly [120, 121]. Numerical experiments show that time intervals determined by 1-D reference models still work well for 3-D tomographic models. Note that time-stepping also serves as a small source of computational advantage, because the Courant number can be slightly larger in 2-D than in 3-D [122]. However, the stability condition becomes much harsher in the presence of a 3-D crust, as elaborated later in Chapter 5.

3.4 Validation

In this section, we benchmark AxiSEM3D in reference to a discretised 3-D SEM, SPECSEM3D_GLOBE (SPECSEM) [51, 52]. We use three global tomographic models for validation, S362ANI [34], S20RTS [112] and S40RTS [39], considering anisotropy [111] and attenuation [68] but, for the time being, ignoring 3-D crustal structures and the internal topography at depths ~ 410 km and ~ 660 km in S362ANI. We shall discuss 3-D crust and internal topography in the next chapter. The conclusions are almost the same for the three models, so we only display the results for S362ANI. Remarkably, this is the first time two different numerical methods are benchmarked for global wave propagation in 3-D Earth models.

To ensure global coverage, we distribute a fictitious network of receivers evenly on the surface in the source-centred coordinate system, spaced by 15° in the θ -direction (epicentral distance) and 45° in the ϕ -direction (azimuth), containing 88 receivers in total. An example is shown in Fig. 3.7.

The accuracy is quantified in terms of time-frequency misfits (TF-misfits) [95]. For each pair of synthetics (SPECSEM vs. AxiSEM3D), we compute the globally normalised envelope and phase misfits (m_e and m_p) on three windows, the body wave governed window $[0, t_{\text{surf}}]$, the surface wave governed window $[t_{\text{surf}}, t_{\text{max}}]$ and the full-length window $[0, t_{\text{max}}]$, where t_{surf} denotes the approximate traveltimes of the first




Figure 3.7. Source-centred receiver network. The source location is VIR (37.91° N, 77.93° W, Virginia). The coloured dots and curves indicate the receivers whose seismograms will be displayed later in this section to visualise benchmark results. We focus on the meridian $\phi = 180^\circ$ because most of its receivers are located on continents and the misfits along this azimuth turn out to be comparatively large.

surface wave arrival and $t_{\max}$ the total record length. We do not treat P-wave and S-wave separately because the accuracy difference turns out to be insignificant.

Before accuracy judgement, an additional modification is applied to the TF-misfits. We compute synthetics for the corresponding 1-D reference model using both SPECFEM and AxiSEM [76] (the validated 1-D version of AxiSEM3D), and then compute the 1-D TF-misfits. We subtract the 1-D misfits from the 3-D ones, based on the fact that the Fourier spectral theory holds exact for 1-D models, that is, eq.(3.59) is fully equivalent to eq.(3.5), and thus these 1-D misfits should result from pure numerical issues, such as distinct meshing schemes, source and receiver mislocations, and accumulative floating-point errors. Indeed, the 1-D phase misfits turn out invariably small ($<2\%$), but the 1-D envelope misfits might be comparatively large (up to 10%) on rare occasions for unidentified reasons.

The results are deemed acceptable if the following three criteria are satisfied simultaneously:

- the global averages of the modified TF-misfits over the receivers, both envelope and phase, should not exceed 5%;
- the local modified TF-misfits should not exceed 10% (note that 10% is a critical value below which the goodness-of-fit may be verbally classified as “excellent” for a comparison of synthetics to data [95, 123]);
- and the local raw 3-D TF-misfits should not exceed 20%.

The above criteria apply to all the three displacement components (radial, transverse and vertical) on all the three predefined windows. The choice of the threshold values largely relies on a reasonable link between quantitative TF-misfits and visible waveform differences. The values themselves are not significant as long as the method proves convergent.

We stress that this section should be recognised as cross-validation rather than answer checking, as SPECFEM has remained unbenchmarked for 3-D Earth models ever before. On the whole, SPECFEM and AxiSEM3D well validate each other, notwithstanding some larger local misfits due to diverse numerical reasons. Beware that such misfits may stem from *either* side. For instance, we have observed large envelope and phase misfits for surface waves at low frequencies (seismic period > 34 s

or $NEX < 128$), which possibly result from the poorly honoured surface curvature in SPECFEM.

3.4.1 Solution convergence

Solution convergence is one of the most basic attributes of a numerical method. Since the convergence of SEM has been well established, for AxiSEM3D specifically, it is necessary to show that the solution converges to the real one as the Fourier expansion order increases. Theoretically, this is guaranteed by the convergence of the Fourier series [110, 114]. In this subsection, we investigate such convergence behaviour numerically.

Let us consider a typical forward problem, with the source located beneath Virginia (37.91° N, 77.93° W) at a depth of 12 km, and the seismic period, denoted T_{simu} , targeted at 17 s ($NEX = 256$). We first need to determine $n_u(s, z)$ properly, basically a trade-off between accuracy and cost. Here we employ the empirical equation (3.72), with n_u^{min} fixed at 8, and n_u^{ref} increased from 8 to 40 to exhibit the convergence procedure.

Figure 3.8 shows how the global average and maximum TF-misfits decrease with n_u^{ref} , manifesting the convergence of solutions at a global level. Quantitatively, the accuracy criteria are satisfied when $n_u^{\text{ref}} = 24$, after which the convergence rate becomes comparatively slow. Figure 3.9 shows waveform comparisons at three receivers located respectively at the distances of 45° , 90° and 135° (the three red dots in Fig. 3.7). The convergence is obvious, as the local TF-misfits monotonically decrease with n_u^{ref} and substantially keep constant after n_u^{ref} reaches 24.

In addition, some interesting propagation effects are observed from Fig. 3.9. Apparently, the required n_u^{ref} varies with both time windows (seismic phases) and receiver locations, which means different spatio-temporal sections of the wavefield have different sensitivity to the 3-D structures. In view of the spatially varying azimuthal resolution caused by $n_u(s, z)$, here we only discuss the phase-dependent sensitivity within each waveform. First, the surface waves are more sensitive than the body waves, which makes sense because surface waves sample the structures more extensively than body waves do; this is also the reason for the surface enhancement in the empirical equation eq. (3.72), as realised by eq. (3.75) and illustrated in Fig. 3.5.

Second, the direct body wave arrivals are the most insensitive seismic phases, as some of them can even be accurately modelled with PREM. These observations reflect the wave scattering effects accumulated along different paths. Exploiting such structure-induced wavefield complexity or sparsity can largely promote the efficiency for both forward modelling and seismic tomography. For instance, starting from 1-D background models should be plausible for tomography based on direct P- and S-waves, but probably not reasonable for surface wave tomography.

3.4.2 Frequency-dependency of n_u

The cost scalability with respect to frequency, that is, how the computation cost increases with frequency, is critical to high-frequency applications. Let ω denote the seismic frequency. A discretised 3-D method has a computational cost of $O(\omega^4)$, that is, three spatial dimensions plus one time dimension. Section 3.2 proves that, using the direct summation approach for mode coupling, AxiSEM3D has a computational cost of $O(n_u \omega^3)$ for a fixed Earth model, given the condition that the azimuthal wavefield sampling can accurately describe the model. Hence, the relation between n_u and ω dominates the relative frequency scalability of the two methods. Unfortunately, such a relation has not been theoretically established. Here we investigate it numerically. For this purpose, we fix the order parameter n_u^{ref} at 24, but change the seismic period T_{simu} from 34 s down to 11 s, based on the same forward problem used in the previous subsection.

Table 3.1 summarises the average TF-misfits at different seismic periods. It is shown that the misfits remain at the same level (sometimes even decrease) as T_{simu} decreases, which means n_u might be frequency independent, or, to say the least, should have a weak dependency on frequency. Such a numerical observation greatly benefits high frequency applications owing to an implied cost scaling factor of ω^3 . We shall verify this observation with profound wavefield learning results in Chapter 5 and Chapter 6.

To visualise the goodness-of-fit, we display waveform comparisons at $T_{\text{simu}} \approx 11\text{s}$ in Fig. 3.10 and Fig. 3.11. The former shows the vertical component on the meridional section $\phi = 180^\circ$ (green arc in Fig. 3.7), where a few significant seismic phases are magnified. The latter shows the transverse component on the two circles of latitude




Figure 3.8. Convergence of global average and maximum TF-misfits between SPECfEM and AxiSEM3D solutions. Simulation parameters: model = S362ANI, $T_{\text{simu}} \approx 17\text{s}$, source = VIR (12 km, 37.91° N, 77.93° W, Virginia). The plotted misfits are computed on the full-length window with $t_{\text{max}} = 1\text{h}$ and modified by 1-D results. The parentheses near a data point show the distance and component (R = radial, T = transverse, Z = vertical) where the maximum misfit is observed.




Figure 3.9. Convergence behaviour of the Fourier spectral solution. Simulation parameters: model = S362ANI, $T_{\text{simu}} \approx 17\text{s}$, source = VIR (12 km, 37.91° N, 77.93° W, Virginia). The three receivers (red dots in Fig. 3.7) are all located on the meridian $\phi = 180^\circ$. The first waveform pair in each figure is special, where we use 3-D model (S362ANI) in SPECFEM but 1-D reference model (PREM) in AxiSEM3D, which gives the upper bound of the misfits. The subsequent pairs are all based on the 3-D model, with the leading integer (8, 16, 24 or 32) specifying n_u^{ref} in (3.72). The symbols “Ex” and “Py” should be interpreted as $m_e = x\%$ and $m_p = y\%$, respectively. For each pair, such symbols from left to right indicate the TF-misfits over the three windows $[0, t_{\text{surf}}]$, $[t_{\text{surf}}, t_{\text{max}}]$ and $[0, t_{\text{max}}]$, respectively, where $t_{\text{max}} = 1\text{h}$ and t_{surf} an approximate arrival time for surface waves (green vertical line). Note that we only show the full-length window misfits for the radial and vertical components at $\theta = 135^\circ$ because the Reighley wave has not arrived ($t_{\text{surf}} > t_{\text{max}}$). Low frequency components below 0.005 Hz are filtered out. This figure is to be continued on the next page.




Figure 3.9 (continued). Convergence behaviour of the Fourier spectral solution.

at $\theta = 60^\circ$ and $\theta = 120^\circ$ (blue rings in Fig. 3.7), respectively. Note that though we only compare waveforms on the surface, the well-fitted body wave phases, both direct and multiple, should imply a well-fitted wavefield throughout the entire interior, and, based on the self-adjointness of the elastic wave operator, the capability of our method to deliver accurate 3-D sensitivity kernels.

3.4.3 Source influence

Because our method necessitates a source-centred coordinate system, we also need to verify the source dependency of its accuracy. We select eight earthquakes located worldwide, including four shallow, two intermediate and two deep earthquakes, as shown in Fig. 3.12. The period is fixed at 17 s. The misfits summarised in Table 3.2 indicate that source location and depth scarcely affect the goodness-of-fit, which, in turn, suggests a uniform orientational complexity of the tomographic models at such as period range.

3.5 Summary

The nature of seismic point sources and realistic Earth models bestows azimuthal smoothness on global seismic wavefields (e.g. Fig. 2.1). To honour and exploit this property, we parameterise 3-D wavefields in terms of Fourier series over the azimuthal dimension while retaining a spectral-element discretisation in the 2-D meridian domain, such that the number of unknowns and thus the computational cost can be remarkably reduced compared to a full 3-D discretisation.

We reduce the 3-D equations of motion to an algebraic system of coupled 2-D meridian equations, and implement a 2-D SEM, AxiSEM3D, to solve the dimension-reduced system. Numerical benchmarks show excellent agreement between AxiSEM3D and SPECFEM3D_GLOBE, an independent full 3-D SEM, for three state-of-the-art global tomographic models across a frequency range typically employed in global waveform modelling and tomography, considering material anisotropy, attenuation and the fluid outer core. To the best of our knowledge, this is the first benchmark between two independent methods for global wave propagation in 3-D Earth models.

It is also shown that different seismic phases observed at different locations have different sensitivities to the 3-D structures, providing useful guidance to both forward modelling and seismic tomography.

The cost characteristics of the method are estimated from the perspective of the spatially varying Fourier expansion order. The precise cost will be summarised in Chapter 5, together with the extension for aspherical Earth models.

Table 3.1. Average TF-misfits (%) at different seismic periods.

$T_{\text{simu}} / \text{s}$	NEX	m_e^{body}	m_p^{body}	m_e^{surf}	m_p^{surf}	m_e	m_p
34	128	1 (T)	0 (T)	1 (T)	1 (T)	1 (T)	1 (T)
23	192	1 (T)	1 (T)	1 (T)	2 (T)	2 (T)	2 (T)
17	256	2 (T)	1 (T)	2 (R)	3 (T)	2 (T)	3 (T)
14	320	2 (T)	1 (T)	2 (R)	3 (T)	2 (T)	3 (T)
11	384	2 (T)	1 (T)	1 (R)	2 (T)	2 (T)	2 (T)

Simulation parameters: model = S362ANI, $n_u^{\text{ref}} = 24$, source = VIR (12 km, 37.91° N, 77.93° W, Virginia). The parentheses indicate the displacement component (R = radial, T = transverse, Z = vertical) of the leading misfit, the maximum among the three directions.

Table 3.2. Average TF-misfits (%) for different earthquake sources.

Label	d / km	m_e^{body}	m_p^{body}	m_e^{surf}	m_p^{surf}	m_e	m_p
VIR	12.00	3 (T)	1 (T)	2 (R)	3 (T)	2 (T)	3 (T)
ITA	15.00	1 (T)	0 (T)	2 (R)	1 (T)	2 (T)	1 (T)
PAK	15.38	1 (T)	0 (T)	3 (T)	3 (T)	2 (T)	3 (T)
ATL	21.81	1 (T)	0 (T)	2 (R)	2 (T)	1 (T)	2 (T)
ALS	129.76	1 (T)	0 (Z)	1 (R)	1 (T)	1 (T)	1 (Z)
NZL	165.12	1 (T)	0 (T)	2 (T)	0 (T)	2 (T)	0 (T)
BOL	647.10	2 (T)	1 (Z)	3 (T)	1 (T)	2 (T)	1 (Z)
JPN	680.70	1 (T)	0 (T)	2 (R)	1 (R)	1 (T)	1 (T)

Simulation parameters: model = S362ANI, $T_{\text{simu}} \approx 17\text{s}$, $n_u^{\text{ref}} = 24$. The earthquakes are shown in Fig. 3.12. Other table settings are identical to Table 3.1.




Figure 3.10. Comparisons of synthetic seismograms (vertical component) for benchmark between AxiSEM3D and SPECFEM. Simulation parameters: model = S362ANI, $T_{\text{simu}} \approx 11\text{s}$, $n_u^{\text{ref}} = 24$, source = VIR (12 km, 37.91° N, 77.93° W, Virginia). The receivers are all located on the meridian $\phi = 180^\circ$ (green arc in Fig. 3.7). The amplitudes have been scaled with a factor of $\sqrt{\theta}$ to identify most phases at all distances. Low frequency components below 0.005 Hz are filtered out. The symbols after each waveform pair indicate the TF-misfits over the full-length window, for example, “E5, P2” means $m_e = 5\%$ and $m_p = 2\%$. We additionally plot the synthetics for the 1-D reference model PREM in the magnified windows to emphasise the lateral scattering effects and, incidentally, to exhibit “perfect” benchmark for 1-D models. The traveltime curves are computed based on the 1-D reference model AK135 [5] using OBSPY.TAUP [124]. This figure is to be continued on the next page.




Figure 3.10 (continued). Comparisons of synthetic seismograms (vertical component) for benchmark between AxiSEM3D and SPECSEM.




Figure 3.11. Comparisons of synthetic seismograms (transverse component) for benchmark between AxiSEM3D and SPECFEM. The receivers are located on the two circles of latitude at $\theta = 60^\circ$ and $\theta = 120^\circ$ (blue rings in Fig. 3.7), respectively. Other figure settings are identical to Fig. 3.10.




Figure 3.12. Map of earthquake sources. The shallow, intermediate and deep sources are coloured respectively blue, green and red. The beachball radii are scaled by magnitude.

Chapter 4

Aspherical Earth

The capability of considering aspherical Earth features is essential to a global wave propagator, as reviewed in Section 1.2. This chapter generalises the theory of the previous chapter to Earth models with undulating interfaces by applying the particle relabelling transformation, which preserves the functional form of the dimension-reduced weak forms, such that the extension of the SEM implementation becomes straightforward.

4.1 Theory

In this section, we focus on the theory of “particle relabelling transformation” [90] that enables our method to consider undulating discontinuities. The basic idea is easily understood. Continuum mechanics typically describes deformation processes starting from a static equilibrium state. However, one can always associate a deformation process to any compatible reference state to reach the same objective description [90]. The concept of particle relabelling has been implied within continuum mechanics literature for a long time, but the actual description of such transformations has only been given recently [90, 125]. In our case, we transform the undulating physical model (static equilibrium state) into a spherically symmetric reference model (a selected compatible reference state), whose density and elasticity properties are determined such that the two models are identical for wave propagation. In other words, we handle interface undulations in terms of volumetric perturbations, and thus retain the form of eq. (3.23) as well as its SEM implementation. In this sense, the transformation can be regarded as an equivalent media approach.

The “particle relabelling transformation” is frequency-independent: it does not interact with wave propagation effects but only depends on the given physical model and the selected geometry of the reference model. Such a notable advantage allows for an easy implementation at the preprocessing stage. However, because the transformation requires compatibility or homeomorphism between the physical and the reference models, it cannot handle discontinuities non-homeomorphic to a sphere, such as the ocean floor. In our method, we model the ocean as a hydrodynamic load, the theory of which is also presented in this section.

4.1.1 Particle relabelling transformation

The Fourier series parameterisation, eq. (3.12), relies on the spherical symmetry of the 3-D domain Ω . To honour aspherical Earth features such as ellipticity and an undulating Moho, we apply the particle relabelling transformation [90].

Consider an Earth model subject to finite undulations on a few interfaces, which can be the surface or any internal solid-solid or solid-fluid discontinuities, such as Moho or CMB. We denote the volume of this physically undulating Earth model by Ω . The 3-D weak form for this physical model is given by eq. (3.5). We map all the particles in this physical model, in one-to-one correspondence, onto an imaginary reference model with spherically symmetric geometry, denoted $\tilde{\Omega}$. Apparently, the choice of the reference spherical geometry is not unique. To minimise the computational cost of the transformation, we constrain the geometric mapping in the radial direction $\hat{\mathbf{r}}$, quantified by the following relation,

$$\xi(\mathbf{r}) = \mathbf{r} + \tau(\mathbf{r}) \hat{\mathbf{r}}, \quad (4.1)$$

where $\tau(\mathbf{r})$ gives the radial shift at position $\mathbf{r}$ in $\tilde{\Omega}$. Literally, eq. (4.1) means, through the geometric mapping $\xi : \tilde{\Omega} \rightarrow \Omega$, the undulating configuration Ω can be constructed by radially stretching all the particles in the reference configuration $\tilde{\Omega}$ from position $\mathbf{r}$ to position $\mathbf{r} + \tau(\mathbf{r}) \hat{\mathbf{r}}$, as illustrated in Fig. 4.1.

The undulating configuration can be considered as a kinematically permissible deformation state from the reference configuration, of which the corresponding deformation gradient, denoted $\mathbf{F}$ transpose, can be written in the spherical coordinate




Figure 4.1. An illustration of geometric mapping from a selected reference spherical configuration to the physically undulating configuration. With the geometric mapping constrained in the radial direction, a particle P at position $\mathbf{r}$ in the reference configuration is stretched to position $\mathbf{r} + \tau(\mathbf{r}) \hat{\mathbf{r}}$ to construct the physical configuration. For an undulating internal interface, one needs to decide an upper and a lower boundary (not necessarily model discontinuities) at which the undulation vanishes. Linear interpolation may be applied in between. The choice of the undulated range is not unique: a narrower one tends to involve fewer stretched elements and thus less additional computational cost for the particle relabelling transformation, but may produce more squeezed elements that diminish the global time step.

system (r, θ, ϕ) in Fig. 3.1 as

$$\mathbf{F}(\tau, \mathbf{r}) = [\nabla(\tau \hat{\mathbf{r}})]^T = \begin{pmatrix} \hat{\mathbf{r}} & \hat{\boldsymbol{\theta}} & \hat{\boldsymbol{\phi}} \end{pmatrix} \begin{pmatrix} \partial_r \tau & \frac{\partial_\theta \tau}{r} & \frac{\partial_\phi \tau}{r \sin \theta} \\ 0 & \frac{\tau}{r} & 0 \\ 0 & 0 & \frac{\tau}{r} \end{pmatrix} \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\theta}} \\ \hat{\boldsymbol{\phi}} \end{pmatrix}. \quad (4.2)$$

According to eq. (152) of [90], the 3-D weak form for the physical model, eq. (3.5), proves equivalent to the following weak form established on the reference configuration $\tilde{\Omega}$,

$$\int_{\tilde{\Omega}} \tilde{\rho} \partial_t^2 \tilde{\mathbf{u}} \cdot \tilde{\mathbf{w}} \, d^3 \mathbf{r} + \int_{\tilde{\Omega}} \nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}} \, d^3 \mathbf{r} = \int_{\tilde{\Omega}} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}} \, d^3 \mathbf{r}, \quad (4.3)$$

in the sense that the wavefield solution of eq. (4.3), denoted $\tilde{\mathbf{u}}$, is identical to that of eq. (3.5):

$$\tilde{\mathbf{u}}(\mathbf{r}; t) = \mathbf{u}(\xi(\mathbf{r}); t). \quad (4.4)$$

The parameters in eq. (4.3) are respectively determined by

$$\tilde{\rho}(\mathbf{r}) = \rho(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}) \quad (4.5)$$

for density, and

$$\tilde{\mathbf{C}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{C}(\xi(\mathbf{r})) \cdot \mathbf{J}^{-T}(\tau, \mathbf{r}) \det \mathbf{J}(\tau, \mathbf{r}) \quad (4.6)$$

for elastic stiffness tensor, and

$$\tilde{\mathbf{f}}(\mathbf{r}) = \mathbf{f}(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}) \quad (4.7)$$

for the source term, with $\mathbf{J}$ denoting the Jacobian of the mapping,

$$\mathbf{J}(\tau, \mathbf{r}) = \mathbf{I} + \mathbf{F}(\tau, \mathbf{r}), \quad (4.8)$$

and $\mathbf{I}$ the second order identity tensor. Note that $\rho(\xi(\mathbf{r}))$ is the density of the physical model at position $\mathbf{r} + \tau \hat{\mathbf{r}}$, similar for $\mathbf{C}(\xi(\mathbf{r}))$ and $\mathbf{f}(\xi(\mathbf{r}))$.

The inverse of $\mathbf{J}$ can be explicitly written as

$$\mathbf{J}^{-1}(\tau, \mathbf{r}) = \frac{1}{h} \begin{pmatrix} \hat{\mathbf{r}} & \hat{\boldsymbol{\theta}} & \hat{\boldsymbol{\phi}} \end{pmatrix} \begin{pmatrix} 1 + \frac{\tau}{r} & -\frac{\partial_\theta \tau}{r} & -\frac{\partial_\phi \tau}{r \sin \theta} \\ 0 & 1 + \partial_r \tau & 0 \\ 0 & 0 & 1 + \partial_r \tau \end{pmatrix} \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\theta}} \\ \hat{\boldsymbol{\phi}} \end{pmatrix}, \quad (4.9)$$

where $h = (r + \tau)(1 + \partial_r \tau)/r$. Note that $\mathbf{J}^{-1}$ shares the same sparsity with $\mathbf{F}$, and this is the reason why we constrain the geometric mapping, eq. (4.1), in the radial direction.

The generalised weak form, eq. (4.3), is formally identical to the original one, eq. (3.5), except for the parameters. Thus, we may reduce it to an algebraic system of 2-D weak forms following the same approach in Section 3.1.3, and the result should be formally identical to eq. (3.23), but with the Fourier coefficients replaced by those of the transformed parameters ($\tilde{\rho}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{f}}$). However, additional attentions must be paid to solid-fluid discontinuities, as addressed in Section 4.1.2. Also, the closed form of the source terms must be modified, as given in Section 4.1.3.

The price for the particle relabelling transformation is that the elastic stiffness tensor of the reference model will lose its minor symmetry, that is, $\tilde{C}_{ijkl} \neq \tilde{C}_{jikl} \neq \tilde{C}_{ijlk}$, as can be inferred from eq. (4.6). Because the major symmetry ($\tilde{C}_{ijkl} = \tilde{C}_{klij}$) still holds, $\tilde{C}_{ijkl}$ contains 45 ostensibly independent elements, in contrast to 5 for transversely isotropy and 2 for isotropy, the two most frequently used types in seismology. Therefore, the double contractions in eq. (3.23), $\nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}}$, will become very expensive if implemented as is. In practice, we do not compute $\tilde{\mathbf{C}}$ by eq. (4.6). Instead, we first decompose the deformation gradient $\nabla \tilde{\mathbf{u}}$ in the spherical coordinate system under which both $\mathbf{J}^{-1}$ and $\mathbf{C}$ are sparse, and perform the contractions $\nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}}$ by means of $(\mathbf{J}^{-T} \cdot \nabla \tilde{\mathbf{u}}) : \mathbf{C} : (\mathbf{J}^{-T} \cdot \nabla \tilde{\mathbf{w}})$.

4.1.2 Aspherical fluid core

An Earth model should contain various solid-fluid discontinuities, collectively denoted Σ , as shown in Fig. 3.1. Following [48], we use the scalar potential formulation in the fluid domains. The 3-D weak form for a solid-fluid Earth model may be written as

$$\begin{aligned} & \int_{\Omega_s} \rho \partial_t^2 \mathbf{u} \cdot \mathbf{w} \, d^3 \mathbf{r} + \int_{\Omega_s} \nabla \mathbf{u} : \mathbf{C} : \nabla \mathbf{w} \, d^3 \mathbf{r} \\ &= \int_{\Omega_s} \mathbf{f} \cdot \mathbf{w} \, d^3 \mathbf{r} - \int_{\Sigma} \partial_t^2 \chi \hat{\mathbf{n}} \cdot \mathbf{w} \, d^2 \mathbf{r}, \end{aligned} \quad (4.10)$$

in the solid domain Ω_s , and

$$\int_{\Omega_f} \frac{1}{\kappa} \partial_t^2 \chi w \, d^3 \mathbf{r} + \int_{\Omega_f} \frac{1}{\rho} \nabla \chi \cdot \mathbf{I} \cdot \nabla w \, d^3 \mathbf{r} = \int_{\Sigma} \mathbf{u} \cdot w \hat{\mathbf{n}} \, d^2 \mathbf{r}, \quad (4.11)$$

in the fluid domain Ω_f , where χ denotes the scalar potential such that $\mathbf{u} = \nabla\chi/\rho$ in Ω_f , w the test function in Ω_f , and κ the bulk modulus of the fluid media. Here we insert a redundant identity tensor $\mathbf{I}$ in the second term of eq. (4.11) to make its functional form more general, so as to facilitate the comparison with its aspherical counterpart given below.

When undulations occur on the solid-fluid discontinuities, the above weak forms need to be modified in analogy to eq. (4.3). For the solid domain, eq. (4.10) becomes

$$\begin{aligned} & \int_{\tilde{\Omega}_s} \tilde{\rho} \partial_t^2 \tilde{\mathbf{u}} \cdot \tilde{\mathbf{w}} \, d^3\mathbf{r} + \int_{\tilde{\Omega}_s} \nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}} \, d^3\mathbf{r} \\ &= \int_{\tilde{\Omega}_s} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}} \, d^3\mathbf{r} - \int_{\tilde{\Sigma}} \partial_t^2 \tilde{\chi} \tilde{\mathbf{n}} \cdot \tilde{\mathbf{w}} \, d^2\mathbf{r}, \end{aligned} \quad (4.12)$$

where $\tilde{\rho}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{f}}$ are computed by eqs. (4.5)-(4.7), and the vector $\tilde{\mathbf{n}}$, orthogonal to the undulating interface Σ in the physical model (but not in unit length), can be shown to be

$$\tilde{\mathbf{n}}(\mathbf{r}) = \mathbf{J}^{-T}(\tau, \mathbf{r}) \cdot \hat{\mathbf{n}}(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}). \quad (4.13)$$

Similarly, for the fluid domain, eq. (4.11) will become

$$\int_{\tilde{\Omega}_f} \frac{1}{\tilde{\kappa}} \partial_t^2 \tilde{\chi} \tilde{w} \, d^3\mathbf{r} + \int_{\tilde{\Omega}_f} \frac{1}{\tilde{\rho}} \nabla \tilde{\chi} \cdot \tilde{\mathbf{I}} \cdot \nabla \tilde{w} \, d^3\mathbf{r} = \int_{\tilde{\Sigma}} \tilde{\mathbf{u}} \cdot \tilde{w} \tilde{\mathbf{n}} \, d^2\mathbf{r}, \quad (4.14)$$

where

$$\tilde{\kappa}(\mathbf{r}) = \kappa(\xi(\mathbf{r})) \det^{-1} \mathbf{J}(\tau, \mathbf{r}), \quad (4.15)$$

$$\tilde{\rho}(\mathbf{r}) = \rho(\xi(\mathbf{r})) \det^{-1} \mathbf{J}(\tau, \mathbf{r}), \quad (4.16)$$

$$\tilde{\mathbf{I}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{J}^{-T}(\tau, \mathbf{r}). \quad (4.17)$$

Again, the 3-D weak forms for an aspherical solid-fluid Earth, eqs. (4.12) and (4.14), are formally identical to those for a spherical solid-fluid Earth, eqs. (4.10) and (4.11), respectively. Therefore, we may consider an undulating solid-fluid discontinuity by modifying both the normal vector of its spherical counterpart and the material properties of the contiguous media in the reference model.

4.1.3 Moment tensor within undulating range

For a point source moment tensor, denoted $\mathbf{M}$, the right hand side of the 3-D weak form, eq. (3.5), yields $\mathbf{M} : \nabla \mathbf{w}(\mathbf{r}_s)$, where $\mathbf{r}_s$ denotes the source location. Accordingly,

the right hand side of the dimension-reduced weak forms, eq. (3.23), can be written as [48]

- Monopole term ($\beta = 0$):

$$2\pi \int_D f_i^0 w_i^0 s \, ds dz = (M_{xx} + M_{yy}) w_{s,s}^0 + M_{zz} w_{z,z}^0, \quad (4.18)$$

- Dipole term ($\beta = 1$):

$$2\pi \int_D f_i^{-1} w_i^1 s \, ds dz = \frac{M_{xz} - \mathbf{i} M_{yz}}{2} (w_{s,z}^1 + 2w_{z,s}^1 + \mathbf{i} w_{\phi,z}^1), \quad (4.19)$$

- Quadrupole term ($\beta = 2$):

$$2\pi \int_D f_i^{-2} w_i^2 s \, ds dz = \left(\frac{M_{yy} - M_{xx}}{2} + \mathbf{i} M_{xy} \right) (w_{s,s}^2 + \mathbf{i} w_{\phi,s}^2), \quad (4.20)$$

where M_{xx} , M_{yy} , ... are the Cartesian components of $\mathbf{M}$, and the derivatives of the Fourier coefficients of $\mathbf{w}$ are evaluated at $\mathbf{r}_s$. Note that the source term vanishes when $|\beta| > 2$.

If the source is located within an undulated range, for instance, within a 3-D crust, the right hand side of eq. (4.3) will become

$$\int_{\tilde{\Omega}} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}} \, d^3 \mathbf{r} = \mathbf{M} : (\mathbf{J}^{-\text{T}}(\tau_s, \mathbf{r}_s) \cdot \nabla \tilde{\mathbf{w}}(\mathbf{r}_s)). \quad (4.21)$$

Accordingly, the above closed forms, eqs. (4.18)-(4.20), should be respectively generalised as

- Monopole term ($\beta = 0$):

$$2\pi \int_{\tilde{D}} \tilde{f}_i^0 \tilde{w}_i^0 s \, ds dz = \left(1 + \frac{\tau_s}{z_s} \right) (M_{xx} + M_{yy}) \tilde{w}_{s,s}^0 + M_{zz} \tilde{w}_{z,z}^0 \\ + \frac{\partial \tau}{\partial x} M_{xz} \tilde{w}_{z,z}^0 + \frac{\partial \tau}{\partial y} M_{yz} \tilde{w}_{z,z}^0 + \frac{\partial \tau}{\partial z} M_{zz} \tilde{w}_{z,z}^0, \quad (4.22)$$

- Dipole term ($\beta = 1$):

$$2\pi \int_{\tilde{D}} \tilde{f}_i^{-1} \tilde{w}_i^1 s \, ds dz = \left(1 + \frac{\tau_s}{z_s} \right) (M_{xz} - \mathbf{i} M_{yz}) \tilde{w}_{z,s}^1 \\ + \frac{M_{xz} - \mathbf{i} M_{yz}}{2} (\tilde{w}_{s,z}^1 + \mathbf{i} \tilde{w}_{\phi,z}^1) \\ + \frac{\partial \tau}{\partial x} \frac{M_{xx} - \mathbf{i} M_{xy}}{2} (\tilde{w}_{s,z}^1 + \mathbf{i} \tilde{w}_{\phi,z}^1) \\ + \frac{\partial \tau}{\partial y} \frac{M_{xy} - \mathbf{i} M_{yy}}{2} (\tilde{w}_{s,z}^1 + \mathbf{i} \tilde{w}_{\phi,z}^1) \\ + \frac{\partial \tau}{\partial z} \frac{M_{xz} - \mathbf{i} M_{yz}}{2} (\tilde{w}_{s,z}^1 + \mathbf{i} \tilde{w}_{\phi,z}^1) \quad (4.23)$$

- Quadrupole term ($\beta = 2$):

$$2\pi \int_{\tilde{D}} \tilde{f}_i^{-2} \tilde{w}_i^2 s \, ds dz = \left(1 + \frac{\tau_s}{z_s}\right) \left(\frac{M_{yy} - M_{xx}}{2} + iM_{xy}\right) (\tilde{w}_{s,s}^2 + i\tilde{w}_{\phi,s}^2), \quad (4.24)$$

where the derivatives of τ are also evaluated at $\mathbf{r}_s$. The above closed forms degenerate to eqs. (4.18)-(4.20) when $\tau_s = 0$ and $\partial_{x_i} \tau|_{\mathbf{r}_s} = 0$, that is, the geometric mapping vanishes in a finite neighbourhood of the source point.

4.1.4 Ocean load

Ocean implementation is another fundamental issue. The most realistic approach is to add patched fluid domains atop the surface. However, this will introduce a multiply-connected solid-fluid discontinuity that falls out of the scope of the particle relabelling transformation. This is currently a limitation of AxiSEM3D. For the time being, we approximate the ocean as a hydrodynamic load on the ocean floor, following SPECfEM. The load approximation is accurate when the wavelengths of interest are larger than the ocean depth. It is suggested by [52] that such an approximation should be acceptable when the period is above ~ 5 s.

When approximated as load, the ocean introduces an additional term [52] to the 3-D weak form, eq. (3.5),

$$\begin{aligned} & \int_{\Omega} \rho \partial_t^2 \mathbf{u} \cdot \mathbf{w} \, d^3 \mathbf{r} + \int_{\Omega} \nabla \mathbf{u} : \mathbf{C} : \nabla \mathbf{w} \, d^3 \mathbf{r} \\ &= \int_{\Omega} \mathbf{f} \cdot \mathbf{w} \, d^3 \mathbf{r} - \int_{\partial\Omega} \rho_w h_w (\partial_t^2 \mathbf{u} \cdot \hat{\mathbf{n}}) (\mathbf{w} \cdot \hat{\mathbf{n}}) \, d^2 \mathbf{r}, \end{aligned} \quad (4.25)$$

where ρ_w denotes the density of sea water and h_w the ocean depth. When surface topography is considered, the normal vector $\hat{\mathbf{n}}$ has to be replaced by $\tilde{\mathbf{n}}$ in eq. (4.13), in addition to ρ , $\mathbf{C}$ and $\mathbf{f}$. In SEM, the ocean load can be easily implemented as an anisotropic mass matrix on the ocean floor.

Taking ρ_w as constant, h_w is the only 3-D parameter involved in the additional term. Applying the global Fourier series parameterisation to h_w , the above weak form reduces to the following algebraic system of 2-D weak forms in D , as a generalisation of eq. (3.23),

$$\begin{aligned} & \sum_{|\alpha| \leq n_u} \left(\int_D \rho^{-(\alpha+\beta)} \partial_t^2 u_i^\alpha w_i^\beta s \, ds dz + \int_D u_{i;j}^\alpha C_{ijkl}^{-(\alpha+\beta)} w_{k;l}^\beta s \, ds dz \right) \\ &= \int_D f_i^{-\beta} w_i^\beta s \, ds dz - \rho_w \sum_{|\alpha| \leq n_u} \int_{\partial D} h_w^{-(\alpha+\beta)} \partial_t^2 u_i^\alpha w_j^\beta \hat{n}_i \hat{n}_j s \, ds dz. \end{aligned} \quad (4.26)$$

4.2 Validation

In this section, we validate the implementation of the particle relabelling transformation by benchmarking AxiSEM3D in reference to SPECFEM. Here we consider two representative aspherical Earth features: (a) Earth’s ellipticity, which undulates the surface and all the internal interfaces (including solid-fluid interfaces), and (b) surface topography and bathymetry, which involve much larger undulations on the surface as well as a 3-D ocean load. Full crustal structure will be elaborated in Chapter 6.

The simulation parameters used in this section are summarised in Table 4.1. To ensure global coverage, we adopt an imaginary network of stations, as shown in Fig. 4.2. This is denser than Fig. 3.7 used in the previous chapter for a better sampling for surface waves. Such coverage on the surface should be able to reflect the goodness-of-fit of the entire interior wavefields. Again, we quantify the goodness-of-fit of waveforms in terms of time-frequency misfits. For each pair of synthetic seismograms produced by SPECFEM and AxiSEM3D, we compute the globally normalised envelope and phase misfits, denoted m_e and m_p , respectively. The accuracy is deemed acceptable when the global averages of m_e and m_p over all the stations are smaller than 5%, and the maximum m_e and m_p smaller than 10%. Such criteria apply to all the three displacement components (radial, transverse and vertical). These threshold values are smaller than those in the previous chapter for material heterogeneities because, when computing the TF-misfits, we adopt a bandpass filter to the waveforms, with its higher cutoff period fixed at 1.5 times of the nominal simulation period. This is because the SPECFEM convention of seismic period adopts a one-element-per-wavelength resolution that appears insufficient to accurately model surface waves. Doing so, we no longer subtract the 3-D TF-misfits by the corresponding 1-D ones.

In this section, we are not concerned with computational cost. We use a constant Fourier expansion order (n_u in eq. (3.12)) that is sufficiently large for the convergence of the AxiSEM3D solution. Computational efficiency will be discussed in Chapter 5.

4.2.1 Ellipticity

Since the entire Earth model has to be flattened from core to surface, Earth’s ellipticity represents an integrated test of the particle relabelling transformation.

Table 4.1. Simulation parameters for ellipticity and topography/bathymetry benchmark.

Parameter	Value	Notes
1-D reference	PREM	transversely isotropic
ellipticity	WGS 84	geographic-geocentric or full correction
topography	ETOPO2v2	surface undulation from -7.7 km to 5.5 km
ocean	ETOPO2v2	approximated as load
period	5 s	$NEX = 896$ in SPECFEM
record length	1 hr	covering most teleseismic phases of interest
earthquake	Virginia	12 km, 37.91° N, 77.93° W
stations	Fig. 4.2	17×18 imaginary stations




Figure 4.2. Source and stations for ellipticity and topography/bathymetry benchmark. We compare synthetic seismograms recorded at 17×18 imaginary stations evenly distributed on the surface (grey), spaced by 10° in the direction of epicentral distance (θ) and by 20° in the azimuthal direction (ϕ). The seismograms recorded at the coloured stations will be shown later in this section.


In practice, there are two levels of ellipticity corrections, the geographic-geocentric (colatitude- θ) correction and the full ellipticity correction. AxiSEM3D contains both the two levels. In the former, the Earth geometry is kept spherical but the locations of source and stations are corrected. In the latter, the entire model is actually flattened to satisfy the hydrostatic equilibrium condition governed by Clairaut’s equation [37], and thus the particle relabelling transformation is activated throughout the entire computational domain.

Excellent agreement has been achieved between SPECFEM and AxiSEM3D. The global averages of m_e and m_p turn out 3% and 1%, respectively, and the maximum local values 6% and 3%, respectively, computed on the 1-hr full-length window. To visualise how small the misfits are, we compare the waveforms in Fig. 4.3, observed on a record section where the maximum phase misfit occurs, as indicated by the blue triangles in Fig. 4.2. From the magnified windows in Fig. 4.3, one can observe the effects of the two levels of ellipticity corrections. It is shown that ellipticity mainly causes phase shifts for the body waves but alters the surface wave waveforms to some degree.

4.2.2 Topography and bathymetry

The 3-D geometric perturbation due to ellipticity varies smoothly in space, as reflected by the small ellipticity corrections in Fig. 4.3. Therefore, we further consider a benchmark solution for topography and bathymetry using the 2-minute Gridded Global Relief Data, ETOPO2v2 [126]. Sampled at a 5 s period (about 15 km between two sampling points on the surface), the surface topography varies from about -7.7 km (Mariana Trench) to 5.5 km (Tibet), capable of generating obvious anomalies on surface waves. The bottom of the asthenosphere (depth = 220 km in PREM) is selected to be the “lower zero boundary” in Fig. 4.1, at which the undulation vanishes. Between the surface and 220 km, the undulation decreases linearly with depth. Ocean is included in the simulations.

Similar to ellipticity, the agreement between SPECFEM and AxiSEM3D turns out to be excellent for topography and bathymetry. The global averages for m_e and m_p are 5% and 3%, respectively, and the maximum local values 9% and 6%, respectively. These misfits turn out larger than those for ellipticity because of two reasons. First,



the sampling points on the surface are different in SPECFEM and AxiSEM3D and second, the SPECFEM directly interprets elevation as ocean depth, while we use the thickness of the water layer in Crust 1.0.

The seismograms are compared in Fig. 4.4 for an azimuthal record section at the distance of 70° , as indicated by the green triangles in Fig. 4.2, with emphasis on the impact of ocean. Apparently, the ocean greatly strengthens surface wave dispersion and causes remarkable phase delays over the oceanic paths, for example, nearly 10 min to the high-frequency surface waves at the azimuth of 80° .

4.3 Summary

Based on the particle relabelling transformation, we extend the theory for spherical geometry to aspherical geometry by mapping the undulated physical model onto a spherical reference model. The material properties of the spherical reference model are determined by those of the undulated physical model and the selected geometric mapping, such that the resultant wavefield remains unchanged. The particle relabelling transformation is frequency-independent and easy to be implemented at the stage of preprocessing. However, it introduces complicated and ostensibly unphysical anisotropy to the elasticity tensor that demands additional computational cost.

Because the geometric complexity is handled in terms of material complexity, the SEM implementation (Section 3.3) for spherically-shaped Earth models can be retained. Benchmark solutions for ellipticity and topography/bathymetry validate the implementation of the particle relabelling transformation in AxiSEM3D. To our knowledge, it is the first time this theory is applied in computational seismology. Apparently, it can also be useful to other numerical methods with difficulties in handling geometric complexity, such as the finite difference method.

Because the particle relabelling transformation is incapable of considering multiply-connected interfaces such as Earth's ocean floor, the ocean is modelled as a hydrodynamic load in AxiSEM3D. It is numerically observed that the ocean load significantly intensifies surface wave dispersion. Full crustal structures will be elaborated in Chapter 6 as a comprehensive application.




Figure 4.4. Comparisons of synthetic seismograms (vertical component) for topography and bathymetry benchmark. The simulation parameters are summarised in Table 4.1. The plotted stations are located at the distance $\theta = 70^\circ$, coloured green in Fig. 4.2. The symbols after each trace indicate the TF-misfits over the full-length window, as explained in Fig. 4.3. Here we focus on the surface wave window because the body waves are barely affected by topography and bathymetry. The blue and red traces consider both topography and ocean, respectively computed by SPECfEM and AxiSEM3D, while the green ones ignore the ocean, computed by SPECfEM (same by AxiSEM3D).

Chapter 5

Computational Efficiency

This chapter investigates the computational efficiency of AxiSEM3D and its speedup with respect to a discretised 3-D method, covering all known types of 3-D Earth features and the entire observable frequency range of global seismic data. Before showing the actually measured costs, we first review our existing knowledge about the cost characteristics of AxiSEM3D.

A model-dependent computational cost distinguishes AxiSEM3D from discretised 3-D methods. Generally, the stronger the 3-D structures are, the more expensive the simulations become. Such a structural-adapted cost characteristic is embodied by the Fourier expansion order of solution, n_u in eq. (3.12), which can be adapted *locally* to the azimuthal complexity of wavefields, that is, $n_u = n_u(s, z)$. For the simplest case of PREM-like 1-D Earth models, the wavefield due to an earthquake takes on order two, that is, $n_u(s, z) = 2$ throughout space and time, for which AxiSEM3D performs as efficiently as the original AxiSEM [76], as proved in Section 3.1.7. For 3-D Earth models, two methods are developed to determine $n_u(s, z)$ in Section 3.3.2, an empirical equation designed for tomographic mantle models and the technique of “wavefield learning” that determines a near-optimal $n_u(s, z)$ sufficiently large for the convergence of solutions via a trial simulation.

Furthermore, we have discussed the frequency scalability of computation cost in Section 3.4.2. A discretised 3-D method has a cost of $O(\omega^4)$. For AxiSEM3D, if the direct summation approach (Section 3.2.1) is used for Fourier mode coupling, the cost will be $O(n_c n_u \omega^3)$, provided that the sampled wavefield complexity is higher than the model complexity, that is, $n_u \geq n_c$ pointwise. Here n_c denotes the Fourier expansion order of the model. on the other hand, if the FFT-based approach (Section 3.2.2)

is used, the cost will become $O(n_u \log(n_u) \omega^3)$ in the worst case. In Section 3.4.2, it is numerically observed that, using the empirical equation eq. (3.72), n_u appears insensitive to frequency across the period range from 68 s to 11 s [48], suggesting a frequency scalability of ω^3 for AxiSEM3D. Such a scaling behaviour is attractive to high-frequency applications.

This chapter will verify the above characteristics and observations with actually measured computational cost. We will also introduce the MPI-based parallelisation scaled up to more than 10,000 cores.

5.1 Tomographic mantle models

This section summarises the computational cost for the three tomographic mantle models used for method validation in Section 3.4, S362ANI [34], S20RTS [112] and S40RTS [39], using the empirical equation eq. (3.72) to determine $n_u(s, z)$ and the direct summation approach for mode coupling. These results were obtained with a prototype Fortran code directly extended from AxiSEM [76], which has been replaced by a recently released C++ code with better performance and usability. The direct summation approach has been deprecated in the new code because it turns out slower than the FFT-based approach in most cases. However, these early results are still worthy of discussion because they clearly reveal the model-adaptivity of the computational cost.

Repeating the convergence test in Section 3.4.1, we find that the accuracy criteria can be satisfied when n_u^{ref} reaches 24 for S20RTS and 32 for S40RTS. The n_u^{ref} obtained at $T_{\text{simu}} \approx 17$ s turns out to be sufficient at $T_{\text{simu}} \approx 11$ s for S20RTS and S40RTS, same as for S362ANI, reconfirming the negligible frequency dependency of n_u . Recalling Fig. 2.3, the maximum values of n_c for the three models are respectively 22 (S362ANI), 20 (S20RTS) and 40 (S40RTS), while the required n_u^{ref} s turn out to be 24, 24 and 32 respectively, so we can assert that n_u should increase with n_c . This is actually as expected because the complexity of the wavefield should grow with that of the model. Nevertheless, only three models are far from enough to establish a solid relation between n_u and n_c . Eventually, we may rewrite the cost of the direct summation approach as $O(n_c^q \omega^3)$, where q is a positive real constant greater than 1. In other words, the cost strongly depends on the lateral complexity of the model, inasmuch as

n_c controls both the size (n_u) and the bandwidth of the dimension-reduced system in eq. (3.65).

The measured CPU-hours for AxiSEM3D and SPECFEM are summarised in Table 5.1, which well support our premises at the beginning of this section:

- AxiSEM3D prevails over SPECFEM for the three tomographic models across the considered frequency band;
- the maximum speedup reaches 85 for S20RTS at 11s;
- the speedup steadily increases with frequency;
- the cost for S20RTS and S362ANI are similar while that for S40RTS nearly doubles.

It is noted that the cost of AxiSEM3D can be further reduced by approximately 30% for deep earthquakes, such as BOL and JPN in Fig. 3.12, as one may switch off the resolution enhancement for surface waves by fixing $F_d = 1$ in eq. (3.72).

Figure 5.1 plots the cost measurements with respect to frequency. As emphasised by the dashed vertical lines, the advantage of AxiSEM3D grows with frequency. In light of this trend, we might achieve real-time simulations (seismogram length equals wallclock simulation time) at 5 s only with a few hundred cores. We will verify such attractive extrapolation in the next section. The low-frequency trend, however, predicts a critical period above which SPECFEM should excel (a result of the constant prefactor n_c^q), which, however, remain undiscovered due to the mesh restriction of SPECFEM ($T_{\text{simu}} \leq 68$ s or $\text{NEX} \geq 64$). Most importantly, the two AxiSEM3D curves for the tomographic models lie between the AxiSEM curve for 1-D models and the SPECFEM curve for 3-D models, supporting our assertion in the introduction that AxiSEM3D bridges the gap between fast, dimension-reduced methods and slow, discretised 3-D methods.

5.2 General cost and speedup

In this section, we show the computational efficiency of our production-level high-performance code released at <https://github.com/kuangdai/AxiSEM3D>. The FFT-based approach introduced in Section 3.2.2 is used for Fourier mode coupling. Parallelisation overheads are also taken into account.

Table 5.1. Computational cost and speedup of AxiSEM3D using direct summation for mode coupling.

T_{simu}/s	SPECFEM	S20RTS	S362ANI	S40RTS
34	73.6	2.1 (34)	2.3 (32)	4.2 (17)
23	188.8	4.1 (47)	4.3 (44)	7.9 (24)
17	324.8	6.8 (48)	7.2 (45)	11.9 (27)
14	768.0	12.6 (61)	13.3 (58)	22.3 (34)
11	2046.7	24.2 (85)	25.6 (80)	42.9 (48)

The second column shows the SPECFEM CPU-hours measured with S362ANI, which remain the same for S20RTS and S40RTS. The following three columns show the AxiSEM3D CPU-hours and the corresponding speedups (parenthesised) relative to SPECFEM. The empirical equation eq. (3.72) is used to determine $n_u(s, z)$, with $n_u^{\text{ref}} = 24, 24$, and 32 respectively for S362ANI, S20RTS and S40RTS. The direct summation approach is used for mode coupling. The CPU-hours are obtained on ARCHER, the latest UK national supercomputing service (archer.ac.uk). The system default (suggested) compiler options are applied for both codes, with enforced cache-optimised loop vectorisation [66]. For SPECFEM, all possible numbers of processors are tested for its best performance, while the AxiSEM3D results are obtained from serial simulations with one processor.




Figure 5.1. Computational cost of AxiSEM3D using direct summation for mode coupling. The vertical axis represents the number of processors required to achieve real-time simulations, that is, seismogram length equals wallclock simulation time, assuming perfect parallelisation scalability. Alternatively, it can be interpreted as the total CPU-hours required to produce 1-hr length synthetics. The data used to produce this figure are given in Table 5.1. The curve for S20RTS is not shown because it is close to that for S362ANI.

Three-dimensional crustal structures are included, for which we employ a state-of-the-art global model, Crust 1.0 [31]. This model is defined on a 1×1 degree grid horizontally, each grid point containing 8 vertical layers: water, ice, upper/middle/lower sediments and upper/middle/lower crystalline crust. Sampled at a 5 s period, the crustal thickness varies from about 8 km to 75 km, introducing very large undulations to the Moho discontinuity. The implementation of the 3-D crust involves quite a few technical details, as elaborated later in Section 6.1. Here we are concerned only with its cost impact. A 3-D crust can considerably increase the computational cost of AxiSEM3D for three reasons. First, the thin oceanic crust greatly diminishes the global time step. Second, the strongly scattered surface waves require much higher Fourier expansion orders near the surface. And third but minor, the particle relabelling transformation demands additional computational cost. In discretised 3-D methods, the first issue can be avoided by an unstructured meshing scheme within the crust, while the second and the third do not apply. Therefore, the 3-D crust also decreases the speedup of AxiSEM3D with respect to 3-D methods.

The actual cost measurements are performed on three representative Earth models, (A) PREM, (B) PREM + S40RTS and, (C) full 3-D: PREM + S40RTS + Crust 1.0 + ellipticity + topography + ocean. The model parameters are summarised in Table 5.2. Each of these models introduces a complexity that notably increases the computational cost of AxiSEM3D: (B) adds 3-D volumetric heterogeneities to (A), changing the problem from 2-D (axisymmetric) to 3-D; and (C) adds geometric complexities to (B), decreasing the global time step, increasing $n_u(s, z)$ near the surface, and activating particle relabelling transformation throughout the domain for ellipticity correction. These three models should be sufficient to constrain other combinations with different model constituents. The simulation period ranges from 50 s down to 1 s, which not only reaches the highest observable frequency of global seismic data, but also utilises hardware at all scales, from laptop (50 s) to large-scale supercomputer (1 s). The shallow earthquake in Section 4.2 is adopted because the solution accuracy is governed by surface waves.

The field of $n_u(s, z)$ is determined by wavefield learning, except for PREM where $n_u(s, z) = 2$. Some of the learning results are shown in Fig. 5.2. As suggested by Fig. 5.2a, for the mantle model (B), $n_u(s, z)$ increases as the period drops from 20 s

Table 5.2. Simulation parameters for efficiency comparisons between AxiSEM3D and SPECFEM.


Parameter	Value	Notes
1-D reference	PREM	transverse isotropic
mantle	S40RTS	degree 40, isotropic
ellipticity	WGS 84	full correction
topography	ETOPO2v2	surface undulation from -7.7 km to 5.5 km
ocean	ETOPO2v2	approximated as load
crust	Crust 1.0	including variable crustal thickness and material perturbations
attenuation	PREM data	modelled as standard linear solids
period	10 s	NEX = 448 in SPECFEM
record length	1.5 hr	covering surface waves at all distances
earthquake	Japan	12 km, 36.97° N, 142.39° E

to 5 s, but only for a small fraction of elements near the surface. Such an observation on the azimuthal wavefield complexity verifies the weak frequency dependency of n_u observed in Section 3.4.2. Figure 5.2b shows that $n_u(s, z)$ becomes much larger near the surface in the presence of a 3-D crust, manifesting its strong scattering effects on the surface waves. The impact of the crust is concentrated but not confined to near the surface. For the deep interior, $n_u(s, z)$ does not increase below 10 s. These observations indicate that for global mantle and crustal models with finite power spectra, there exists a frequency above which the azimuthal complexity of wavefields may cease to develop for body waves, and the computational cost of AxiSEM3D will start to scale with $\omega^{3+\epsilon}$, with ϵ accounting for surface waves, significantly smaller than 1.


The CPU-hours required to compute a 1-hr record length are summarised in Table 5.3 and plotted in Fig. 5.3. It turns out that AxiSEM3D can be 1 to 3 orders of magnitude faster than SPECFEM for the 3-D mantle model (B), and 0 to 2 orders of magnitude faster for the full 3-D model (C). In other words, the geometric complexities (mainly the 3-D crust) increase the cost of AxiSEM3D by an order of magnitude. For both (B) and (C), the speedup roughly scales with frequency. Apparently, AxiSEM3D is most efficient for high-frequency body-wave studies. Similar to Fig. 5.1, Fig. 5.3 stresses again the model-adaptivity of AxiSEM3D for a larger model parameter space.

In addition, the computational cost of AxiSEM3D may be further reduced by the following factors:

1. smoother models: for example, substituting S40RTS by S20RTS will reduce the cost for (B) by $\sim 20\%$, and substituting Crust 1.0 by Crust 2.0 [127] will reduce the cost for (C) by 10% to 20%, depending on the period;
2. ignoring surface waves: for deep events or applications where surface waves are not of interest, the cost for (B) can be reduced by 20% to 40%, depending on the period; it might be unprofitable to include a 3-D crust for such cases;
3. ellipticity: the cost for (C) can be reduced by $\sim 15\%$ by turning off full ellipticity;
4. record length: the computational cost scales with record length in a super-linear style, as the azimuthal complexity of wavefields also evolves with time, that is,



(a) PREM + S40RTS



(b) PREM + S40RTS + Crust 1.0 + ellipticity + topography + ocean

Figure 5.2. Fields of $n_u^{\text{learn}}(s, z)$ obtained by wavefield learning for global mantle and crustal models. In (a), we only involve S40RTS for the mantle, and in (b), we make a full 3-D model using S40RTS for the mantle and Crust 1.0 for the crust, including ellipticity, topography and ocean. The model parameters are summarised in Table 5.2. All the learning simulations start from a constant n_u field, $n_u^{\text{start}}(s, z) = 1,200$. As an indicator of computational cost, Σn_u gives the sum of $n_u^{\text{learn}}(s, z)$ over all the grid points in the mesh. The same colour scale is used in (a) and (b).

Table 5.3. Computational cost of AxiSEM3D and its speedup with respect to SPECfEM.

period / s	(A) PREM	(B) S40RTS	(C) Full 3-D
50	0.22 (250)	10 (6)	55 (1.5)
20	0.88 (393)	24 (14)	188 (2.4)
10	3 (1254)	62 (62)	733 (6)
5	20 (3410)	642 (105)	2926 (28)
2	301 (8749*)	6768 (389*)	55370 (57*)
1	2427 (17356*)	55973 (753*)	442958 [†] (114* [†])

This table gives the CPU-hours required by AxiSEM3D to compute 1-hr length seismograms, followed by the speedups (parenthesised) relative to SPECfEM. The values marked by * are estimated by extrapolation because of unaffordable SPECfEM solutions, so are those marked by [†] because of unaffordable AxiSEM3D solutions. The CPU-hours are measured on ARCHER, the latest UK national supercomputing service (archer.ac.uk), taking into account parallelisation overheads.




Figure 5.3. Computational cost of AxiSEM3D and its speedup with respect to SPECfEM. The vertical axis represents the CPU-hours required to compute 1-hr length seismograms. The data used to produce this figure are given in Table 5.3. The dashed parts are estimated by extrapolation because of unaffordable solutions. For SPECfEM, we only plot the CPU-hours for models with a 1-D PREM crust, which will be increased by 20% to 30% for Crust 1.0, depending on the period.

longer recorder length leads to higher n_u ; for example, halving the record length will nearly quarter the cost for both (B) and (C).

5.3 Parallelisation

AxiSEM3D is parallelised via the message passing interface (MPI), using non-blocking point-to-point communication modes. Load balancing across processors is the most important issue around parallelisation. In ordinary SEMs such as AxiSEM and SPECFEM, since the computational cost of an individual element is constant (not considering the fluid core), load balancing can be easily achieved by equally distributing the elements across the processors. In AxiSEM3D, however, the computational cost varies from element to element, due to the variability of the Fourier expansion order $n_u(s, z)$. Therefore, a heterogeneous domain decomposition is required. In brief, we use the package METIS [128] to partition the 2-D computational domain, with a weight assigned to each element. The weights are determined by measuring the elemental costs during runtime¹. A complete procedure of domain decomposition is described with an example in Fig. 5.4.

The scaling behaviour of such a measure-and-decompose scheme is tested on the supercomputer ARCHER (archer.ac.uk). Both strong (fixed total workload) and weak (fixed workload per core) scalabilities are examined. We first consider a 1-D PREM model, where $n_u(s, z)$ is constant over the domain, and the decomposition is thus unweighted. Perfect strong scaling is observed up to 12 288 cores for a 1 Hz simulation, as shown in Fig. 5.5a. Weak scaling turns out slightly sub-optimal, with a 4% parallelisation overhead at 12 288 cores, as shown in Fig. 5.5b.

For a 3-D Earth model, both strong and weak scaling properties are expected to decrease for two reasons. First, the runtime measurement of elemental costs (Fig. 5.4c) fluctuates. Second, extra workload emerges on the edge-cuts (Fig. 5.4d) after domain decomposition, which cannot be balanced a priori (in AxiSEM or SPECFEM, the

¹ Most of the computing time in AxiSEM3D is consumed by FFT and matrix arithmetics, for which we rely on the high performance packages FFTW3 [119] and EIGEN3 [129], respectively. The cost of an element becomes unpredictable due to the low-level optimisations in these packages, and thus can be obtained only through runtime measurements.




Figure 5.4. An example of domain decomposition. The global mantle model S40RTS is considered, with a 50 s period. (a) shows the 2-D spectral element mesh, which contains 4 element classes (in the legend, ISO means isotropic and TISO transversely isotropic). (b) shows a $n_u(s, z)$ field computed by the empirical equation, eq. (3.72). The computational cost of an element depends on both its class and the value of n_u . (c) shows the elemental cost measured by the solver. Note that an element with a larger n_u is not necessarily more expensive than those of the same class but with a smaller n_u , mainly because the cost of FFT does not monotonically increase with the logical size, and this is why we need to measure elemental costs at runtime. Finally, (d) shows the weighted domain decomposition produced by METIS, weighted by the measured costs in (c). The varying numbers of elements on different processors clearly reflect the effect of cost weighting.




Figure 5.5. Scaling behaviours of AxiSEM3D on a Cray XC30 supercomputer (ARCHER, UK). Two Earth models are tested, the 1-D model PREM with $n_u(s, z) = 2$, and the 3-D model S40RTS with $n_u(s, z)$ determined by the empirical equation, eq. (3.72). For strong scaling, the problem (total workload) is fixed, that is, a fixed period of 1 s for PREM and 2 s for S40RTS. For weak scaling, we change the period such that the workload per core remains nearly constant, where the workload is estimated by the summation of n_u . The times to solution (y-axis) are wall-clock times measured for 2000 time steps.

processor boundaries are also pre-balanced). The 3-D scaling behaviours shown in Fig. 5.5 are observed using 3-D mantle model S40RTS, with an empirical $n_u(s, z)$ sufficiently large for a converged solution. For strong and weak scaling, the overheads are respectively 6% and 10% at 12 288 cores. Compared to the speedups shown in Table 5.3, these overheads are negligible.

It is noted that the above scalabilities are evaluated in reference to simulations with 384 cores. Compared to serial simulations, AxiSEM3D still suffers a parallelisation overhead of 20% to 30% for 3-D Earth models due to unstable cost measurement. This is an issue awaiting future optimisation.

5.4 Summary

This chapter summarises the computational efficiency of AxiSEM3D and its speedup with respect to a discretised 3-D method, based on cost measurements on a supercomputer. We adopt a large model parameter space, considering 3-D mantle and crustal structures, topography and ocean, and ellipticity. We also adopt a wide period range from 50 s down to 1 s, reaching the highest observable frequency of global seismic data and utilising hardware at all scales.

The measured computational cost suggests that AxiSEM3D is substantially more efficient than discretised 3-D methods for 3-D mantle structures without a 3-D crust, with an enormous speedup of 1 to 3 orders of magnitude for the considered period range. Such a speedup basically scales with frequency. In other words, AxiSEM3D drastically improves the computability of global wave propagation in 3-D mantle structures, which can greatly promote forward and inverse studies of the deep Earth based on high-frequency body waves. When 3-D crustal structures are involved, the speedup is decreased by an order of magnitude mainly for two reasons. First, the thin oceanic crust greatly diminishes the global time step. And second, crustal structures introduce a much higher azimuthal complexity to surface waves, leading to much higher Fourier expansion orders near the surface. Nevertheless, the remaining speedup is still beyond or comparable to other hardware and algorithmic optimisations. In summary, AxiSEM3D bridges the gap between fast dimension-reduced methods (such as AxiSEM [76]) and slow discretised 3-D methods (such as SPECFEM [51, 52]) by adapting itself to the complexity of the Earth model.

The uneven elemental cost makes parallelisation more difficult for AxiSEM3D than for other conventional SEMs such as AxiSEM and SPECFEM. By means of a weight-balanced domain decomposition scheme, we achieve satisfactory parallel scalabilities up to 12,288 cores.

Chapter 6

Applications

This chapter shows three applications of AxiSEM3D. In the first application, we model surface waves with a global crustal model, Crust 1.0, and compare the synthetics to data. It is shown that the 3-D crust significantly enhances the accuracy for surface wave modelling. In the second application, we investigate a small-scale localised heterogeneity with sharp material contrast against the 3-D background model. The wave scattering effects for a wide range of structure-wavelength combinations are examined by wavefield learning, verifying the robustness of AxiSEM3D in handling localised 3-D structures. The third application simulates an ultra-low velocity zone (ULVZ) at a 1 Hz frequency, showing the great potential of AxiSEM3D for deep Earth studies at high frequencies.

6.1 Surface waves in 3-D crust

Three-dimensional crustal structures have a major influence on surface wave modelling. Compared to smooth long-wavelength mantle structures, Earth’s crust varies much more drastically in space, in terms of both amplitude and gradient. The most outstanding challenge comes from a severely varying crustal thickness, characterising the Moho discontinuity with large undulations and sharp slopes beneath the ocean-continent boundaries. For instance, Fig. 6.1a shows the undulating Moho in a state-of-the-art 3-D crustal model, Crust 1.0 [31], defined on a 1×1 degree grid. Sampled at a 5 s period (about 15 km between two sampling points on the surface), the crustal thickness varies from about 8 km to 75 km. In this section, we describe the different implementations of a global crustal model in AxiSEM3D and the built-in


SPECFEM mesher, and assess the results by comparing them to real seismic data.

6.1.1 Crust implementation


Global crustal models involve topography on both Earth’s surface and Moho, and material heterogeneity between them. The undulating surface and Moho can be honoured by stretching the spherical layers of the 1-D reference model. Those elements across the sharp slopes beneath the ocean-continent boundaries may easily get distorted after stretching, which demand special treatment. Stability condition also becomes more stringent, as governed by the thin oceanic crust. Driven by different motivations, SPECFEM and AxiSEM3D select different radial layers to match the undulating Moho.

Because the particle relabelling transformation requires the homeomorphism (topological equivalence) of any radial discontinuities to a sphere, AxiSEM3D stretches the spherical Moho of the 1-D reference model, for example, 24.4 km in PREM. Such implementation is physically sound because it keeps the integrity of the Moho discontinuity. However, the squeezed upper and lower crustal layers beneath the ocean will considerably diminish the global time step. If an element is distorted after stretching, the Jacobian of the geometric mapping in eq. (4.8) will lose its positive definiteness. If such a situation happens, AxiSEM3D will automatically adjust the position of the Moho, but only locally at the slope that actually causes the distortion, such as part of the Andean Mountains.

The selection of the stretched layers can be more flexible in a discretised 3-D method. For instance, the built-in SPECFEM mesher stretches different radial layers in different regions, as explained in Fig. 6.1b. Such implementation effectively avoids distorted elements near the sharp slopes and squeezes only the upper crustal layer beneath the ocean to have a less decreased global time step. The time step of AxiSEM3D is about one third of that of SPECFEM for Crust 1.0. However, the disadvantages are complementary. First, almost one-fourth of the Moho discontinuity is smoothed away where it traverses element interiors. These dishonoured regions (coloured white in Fig. 6.1b) largely overlap with seismically active regions along the plate boundaries, which may introduce large errors to near-field wavefields. Second, the radial sampling of the oceanic crust might be too sparse at some regions. To overcome the



(a) Moho of Crust 1.0



(b) SPECFEM implementation

Figure 6.1. Moho discontinuity of Crust 1.0 and its implementation in SPECFEM. (a) shows topography on Moho, coloured by the undulation with respect to the PREM Moho (24.4 km). (b) shows the SPECFEM implementation. In (b), the mid-crustal discontinuity (15 km in PREM) is stretched upward to match the 3-D Moho for regions where the crustal thickness is smaller than 15 km (coloured blue), mainly beneath the ocean; an artificial discontinuity (35 km) is stretched downward for regions where the crustal thickness is greater than 35 km (coloured red), mainly beneath the continents; for the rest (coloured white), mainly beneath the ocean-continent boundaries, the Moho cuts through spectral elements and is thus smoothed by means of material interpolations.

difficulties around 3-D crust discretisation, an unstructured meshing scheme is under development within the Salvus project [56], which honours the Moho exactly while avoiding element distortion and multiply squeezed elements in the oceanic crust.

6.1.2 Comparison to data

In this subsection, we produce synthetic seismograms for a full 3-D Earth model, using S40RTS for the mantle and Crust 1.0 for the crust, including ellipticity, topography and ocean, same as the full 3-D model used for efficiency comparison in Section 5.2. The simulation period is 10 s, as SPECFEM does not support any higher frequencies due to a known technical issue in its built-in crustal mesher. The simulation parameters are summarised in Table 5.2. We compare both SPECFEM and AxiSEM3D synthetics to real seismic data to assess the two different crust implementations. It should be noted that such evaluation can only be rough because the misfits between synthetics and data are subject to uncertainties such as event information, tomographic models, and noise in data.

The source and stations are plotted in Fig. 6.2. The discrepancy between SPECFEM and AxiSEM3D should be path-dependent: the AxiSEM3D implementation is expected to be more realistic where the path largely traverses the ocean and the ocean-continent boundaries. Therefore, to analyse the data in more detail, we divide the stations into four groups by their azimuths: (a) the Pacific ocean group, labeled PA, with pure oceanic paths; (b) the Australian group, labeled AU, with mixed paths across the ocean-continent boundaries almost perpendicularly; (c) the Indian ocean group, labeled IN, with mixed paths going along the ocean-continent boundaries; and (d) the European group, labeled EU, mainly with continental paths.

For each station, we have four traces in total: two synthetics for the full 3-D model computed by SPECFEM and AxiSEM3D, one synthetic for PREM computed by SPECFEM (same by AxiSEM3D), and the real data. To reduce the uncertainties in data, we only consider the stations where the cross-correlation coefficient on the direct P-wave window between PREM synthetics and data turns out greater than 0.7. This is because the direct P-arrival is less affected by 3-D structures and, particularly its shape, should be better predicted by PREM, as numerically verified in Section 3.4.1. We quantify the goodness-of-fit by the cross-correlation on the full-length window,




Figure 6.2. Source and stations for assessment of 3-D crust implementations. We compare synthetic seismograms to data recorded at 616 stations (grey). Based on their azimuths, we divide the stations into four groups, labeled respectively PA (Pacific ocean, blue), AU (Australia, green), IN (Indian ocean, orange) and EU (Europe, cyan), separated by the dashed great circles. We will show seismograms at the coloured record sections and stations later in this subsection.

so the uncertainties from the source, both its origin time and magnitude, are automatically removed. Here we do not use TF-misfits because the waveform differences at some stations might be too large for TF-misfits to work robustly. The statistics of cross-correlations are summarised in Table 6.1. We draw the following conclusions from the table:

1. the full 3-D model fits surface wave data better than PREM does;
2. compared to pure 3-D mantle models (Chapter 3), the 3-D crustal model leads to a larger disagreement between SPECfEM and AxiSEM3D;
3. the synthetics computed by AxiSEM3D have a statistically higher cross-correlation to data than those by SPECfEM for the considered event;
4. as expected, the superiority of AxiSEM3D is visible for the oceanic group PA as well as the mixed-path groups AU and IN.
5. SPECfEM and AxiSEM3D have very similar accuracy for the continental group EU; when filtered down to 20 s, the average cross-correlation of SPECfEM becomes larger, implying that the third element layer (24.4 km to 35 km in PREM) contained in the SPECfEM continental crust may have begun to take effect at this period.

For each of the four groups, we show a well-fitted record section in Fig. 6.3, filtered down to 40 s. At this period, SPECfEM and AxiSEM3D agree well with each other. The seismograms at four selected stations are displayed in Fig. 6.4, filtered down to both 40 s and 20 s. Fig. 6.4a-6.4c are selected from the EU group. Fig. 6.4a shows an example of good agreement between SPECfEM and AxiSEM3D, and both of them well predict the data. Fig. 6.4b and 6.4c respectively represent the case where SPECfEM and AxiSEM3D predominate over each other, with the major differences concentrated on the amplitudes of high-frequency surface waves. Fig. 6.4d shows an example where the full 3-D model successfully fits the data at 40 s, but fails to capture some of the high-frequency surface waves at 20 s, which are however clearly observed in the 1-D synthetics without ocean. Such a situation only occurs at some of the island stations in our dataset, which can be remedied by turning off the ocean load. Because the load approximation of ocean should be accurate at this frequency range, the errors may result from the fact that island stations are actually placed on the ocean floor in the simulations, as those small islands may not be sampled either

by Crust 1.0 or the computing mesh. The wave physics behind these high-frequency reverberations needs more thorough analysis. Finally, as suggested by Fig. 6.4, the 3-D simulations basically reproduce the body wave data, indicating that at such a period S40RTS should be highly credible for the sampled regions.

6.2 Localised heterogeneity

Apart from smooth, long-wavelength material heterogeneities and undulating interfaces, localised small-scale structures represent another type of 3-D feature of interest, such as the ULVZs [104, 105] and mid-mantle scatterers [130, 131]. Depending on the structure-wavelength ratio, such small-scale structures may enter the regime of backscattering, reflecting waves to all directions like a secondary source. Such backscattering effects can be well simulated by discretised 3-D methods, almost without additional computational cost, but are generally believed to be a bottleneck of pseudo-spectral approaches, due to their infinite power spectra [110]. Being convergent in theory, our method is undoubtedly capable of incorporating small-scale 3-D heterogeneities, but their cost impact has to be carefully examined.

For this purpose, we place a spherical slow blob in the mid-mantle of S40RTS. The centre of the blob is located 30° in epicentral distance and 1200 km in depth. Within the blob, the S- and P-wave velocities are respectively reduced by 50% and 30%, and the density by 20%. Such strength of heterogeneity should be strong enough to exceed typical conditions in deep Earth. We consider 4 simulation periods and 7 blob radii, that is, 28 simulations in total. Again, wavefield learning is applied to ensure the convergence of solutions. Because surface waves are not affected by the blob, we use a 1-D PREM crust and a deep earthquake. The simulation parameters are summarised in Table 6.2. The geometry of the blob, the source and the stations are shown in Fig. 6.5.


The influence of the blob depends not only on its size, but also on the wavelength. To quantify the strength of wave scattering, we define the structure-wavelength ratio, denoted ζ ,

$$\zeta = \frac{2R_{\text{blob}}}{3 \text{ km/s} \times T_{\text{simu}}}, \quad (6.1)$$


Table 6.1. Statistics of cross-correlations between synthetics and data for assessment of 3-D crust implementations.

period 40 s	all	PA	AU	IN	EU
number of stations	462	49	128	124	161
higher CC for Axi3D	62%	81%	58%	78%	48%
CC SPEC / Axi3D	0.98 ± 0.01	0.99 ± 0.01	0.96 ± 0.02	0.98 ± 0.01	0.98 ± 0.01
CC SPEC / data	0.65 ± 0.14	0.78 ± 0.09	0.52 ± 0.19	0.62 ± 0.17	0.74 ± 0.10
CC Axi3D / data	0.69 ± 0.13	0.87 ± 0.06	0.57 ± 0.14	0.66 ± 0.15	0.75 ± 0.11
CC PREM / data	0.17 ± 0.20	0.13 ± 0.15	0.34 ± 0.22	0.37 ± 0.18	-0.11 ± 0.27
period 20 s	all	PA	AU	IN	EU
number of stations	450	44	112	127	167
higher CC for Axi3D	54%	86%	66%	55%	38%
CC SPEC / Axi3D	0.72 ± 0.11	0.96 ± 0.02	0.85 ± 0.07	0.70 ± 0.11	0.59 ± 0.16
CC SPEC / data	0.29 ± 0.13	0.61 ± 0.11	0.21 ± 0.18	0.21 ± 0.15	0.33 ± 0.11
CC Axi3D / data	0.31 ± 0.10	0.68 ± 0.08	0.25 ± 0.12	0.22 ± 0.09	0.31 ± 0.09
CC PREM / data	0.18 ± 0.15	0.09 ± 0.24	0.30 ± 0.18	0.15 ± 0.14	0.01 ± 0.22

The simulation parameters are summarised in Table 5.2. The stations as well as the four groups (PA, AU, IN, and EU) are shown in Fig. 6.2. The simulation period is 10 s, and we filter the seismograms down to both 40 s and 20 s for data comparisons. The cross-correlations are computed using the vertical component, but similar conclusions can be drawn from the radial and transverse components. The numbers before and after $\pm$ are respectively mean values and standard deviations of the cross-correlation coefficients. The second row “higher CC for AxiSEM3D” shows the percentage of stations where the cross-correlation of AxiSEM3D to data is higher than that of SPECFEM to data. The statistics in favour of SPECFEM are emphasised in bold.




(a) record section PA




(b) record section AU

Figure 6.3. Comparisons of synthetics to data (vertical component) on four selected record sections for assessment of 3-D crust implementations. The simulation parameters are summarised in Table 5.2. The seismograms are filtered from 40 to 100 s. The plotted record sections are shown in Fig. 6.2. The synthetics are shifted by cross-correlating the direct P-wave window, and plotted in their original amplitudes. The symbols after each trace indicate the TF-misfits for the full 3-D model over the plotted window, for example, “S: E22, P19” and “A: E21, P6” after the top trace in (a) means $m_e \approx 22\%$ and $m_p \approx 19\%$ for SPECFEM to data, and $m_e \approx 21\%$ and $m_p \approx 6\%$ for AxiSEM3D to data. To emphasise surface waves, the time origin of each trace is fixed at the traveltimes of the surface wave of speed 5 km/s in PREM. This figure is to be continued on the next page.




(c) (continued) record section IN




(d) (continued) record section EU


Figure 6.3 (continued). Comparisons of synthetics to data (vertical component) on four selected record sections for assessment of 3-D crust implementations.




(a) station FNA, Florina, Greece



(b) station LRVF, Nouvelle-Aquitaine, France



(c) station BFO, Black Forest, Germany




(d) station KNTN, Kanton Island, Kiribati

Figure 6.4. Comparisons of synthetics to data (vertical component) on four selected stations for assessment of 3-D crust implementations. The simulation parameters are summarised in Table 5.2. For each station, the top trace is filtered from 40 to 100 s and the bottom from 20 to 100 s. The plotted stations are shown in Fig. 6.2, coloured red. The synthetics are shifted by cross-correlating the direct P-wave window, and plotted in their original amplitudes. The symbols at the beginning of each trace indicate the TF-misfits for the full 3-D model over the plotted window, as explained in Fig. 6.3.

Table 6.2. Simulation parameters for a spherical slow blob placed in the mid-mantle of S40RTS.

Parameter	Value	Notes
1-D reference	PREM	fully benchmarked
mantle	S40RTS	fully benchmarked
spherical blob	position	1200 km, 2.3389° S, 39.2° W; epicentral distance = 30°, azimuth = 110°
	radii / km	0, 5, 10, 20, 40, 80, and 160
	material	S- and P-wave velocities reduced by 50% and 30% respectively, density by 20%
period	4 periods	50 s, 20 s, 10 s, and 5 s
record length	1.0 hr	covering all affected phases
earthquake	Bolivia	647.1 km, 13.82° S, 67.25° W.
stations	Fig. 6.5	only for visualisation purpose



(a) slice view



(b) geography

Figure 6.5. Problem setup for a spherical slow blob placed in the mid-mantle of S40RTS. (a) shows the slice across the blob centre, located 30° in distance and 1200 km in depth. Within the blob, the S- and P-wave velocities are respectively reduced by 50% and 30%, and the density by 20%. The reductions fade in a Gaussian style outside the blob, with a halfwidth at half maximum of 1/5 of its radius. The path labelled “P” corresponds to the direct P-arrival and “P_{BS}” the backscattered P-arrival. (b) shows the geography of the blob, the source and the stations. The record section coloured blue has the same azimuth as the blob centre, with distances ranging from 30° to 50°. The one coloured green is located 40° in distance, with azimuths ranging from -15° to +15° with respect to the blob centre.

where R_{blob} is the radius of the blob, T_{simu} the simulation period, and 3 km/s the S-wave velocity within the blob. In general, the propagation effects of the blob should increase with ζ .

Figure 6.6 summarises the blob-induced relative cost with respect to different structure-wavelength ratios. It first confirms the competency of AxiSEM3D in handling such localised small-scale structures, as the maximum slowdown due to the blob is only 1.87. This maximum slowdown occurs when $R_{\text{blob}} = 160$ km, which is expected, but at a period of 10 s instead of 5 s (the shortest). More interestingly, it is shown that the relative cost becomes more and more insensitive to the blob structure as the period decreases, which again benefits high frequency applications. To look into the reason behind such observations, we resort to the actual distributions of $n_u^{\text{learn}}(s, z)$ obtained from wavefield learning, as shown in Fig. 6.7, which visualise the impact of the blob on the azimuthal complexity of wavefields. Above all, as the structure-wavelength ratio increases from zero to a very large value, the wave scattering regime shifts from forward-dominant (refraction and diffraction) to backward-dominant (reflection).

Let us first read Fig. 6.7 row-wise. For a fixed period, the influence of the blob increases with its size. For the sub-wavelength regime where $\zeta < 1.0$, the blob has very little influence on wave propagation. As R_{blob} increases, the azimuthal scattering effects become stronger, but are highly concentrated in a smeared neighbourhood around the blob, as suggested by the comet-like patterns such as (d4). AxiSEM3D takes advantage of such concentrated scattering effects by locally adapting $n_u(s, z)$, and thus enjoys a very limited increase to the total computational cost. For example, the maximum slowdown due to the blob is only 1.87, observed in (c6). It can also be inferred that the relative cost increases with Σn_u in a weak super-linear style because of the super-linearity of FFT.

Now we read Fig. 6.7 column-wise to see how $n_u(s, z)$ evolves with frequency for a fixed blob radius. As suggested by the first column, for bare S40RTS, $n_u(s, z)$ increases when the period drops from 50 s to 20 s, but then remains constant for the shorter periods. This is consistent with the conclusion we drew in Sections 3.4.2 and 5.2. As suggested by the remaining columns, as the frequency increases, the blob leads to an ever-growing n_u in a small neighbourhood around it, but within an increasingly




Figure 6.6. The blob-induced relative cost of AxiSEM3D with respect to the structure-wavelength ratio ζ . The simulation parameters are summarised in Table 6.2. The relative cost is obtained by normalising the actual computational cost to that of the pure S40RTS, so the part above 1.0 is caused by the blob. Note that the ζ -axis is logarithmic.




Figure 6.7. Fields of $n_u^{\text{learn}}(s, z)$ obtained by wavefield learning for a spherical slow blob placed in the mid-mantle of S40RTS. The problem is described in Fig. 6.5 and the simulation parameters are summarised in Table 6.2. The simulation period T_{simu} varies column-wise and the blob radius R_{blob} varies row-wise. As an indicator of computational cost, Σn_u gives the sum of $n_u^{\text{learn}}(s, z)$ over all the grid points in the mesh. For each row, the actually measured computational costs are normalised to the first column to emphasise the cost impact of the blob.


concentrated range of influence. These two counterbalancing tendencies will give rise to a frequency at which the blob-induced cost reaches its maximum, for example, 20 s for $R_{\text{blob}} = 80$ km and 10 s for $R_{\text{blob}} = 160$ km. Such concentrated lateral scattering effects make the computational cost insensitive to small-scale structures, especially for high-frequency simulations.

The concentrated scattering also suggests a fast healing effect of the azimuthal wave complexity (different from wavefront healing), which further hints a type of seismic invisibility on the azimuthal direction. At a small ζ where forward scattering dominates, such as (c3) and (d2), the azimuthal complexity of the wavefield has readily recovered on the surface, which means the wave has “forgotten” this part of information it encoded from the structure. In this sense, constraining the transverse dimension of the structure (orthogonal to source-receiver plane) can be more difficult, especially at a higher frequency. This is partially consistent with finite-frequency theory [41], that is, the higher the frequency, the narrower the sensitivity kernels. But more, we claim that within a cross-section of a banana-doughnut kernel, the resolution should be higher in the vertical (in-plane) direction than in the transverse (off-plane) direction. In this sense, the resolution or uncertainty of a tomographic model should depend on not only the path coverage but also the path orientations. For instance, if the seismic paths are orientated dominantly from north to south, the resultant model should be subject to a higher uncertainty from east to west.


Finally, we show the seismograms on two record sections in Fig. 6.8. Full 3-D backscattering effects can be observed in Fig. 6.8b: as the receiver moves away from the central azimuth, the backscattered energy (amplitude of P_{BS}) decays while the differential traveltimes between P and P_{BS} enlarges.

6.3 Simulating a ULVZ at 1 Hz

ULVZs are small-scale slow heterogeneities atop the core-mantle boundary (CMB), significant to the understanding of Earth’s thermal evolution and mantle convection, as reviewed at the beginning of the introduction. These small-scale structures typically extend a few hundred kilometres horizontally and a few dozen kilometres vertically, with sharp velocity contrast (up to -50% for v_s) against the long-wavelength background structures such as the large low-shear-velocity provinces (LLSVPs). Such



(a) record section with $\phi = 110^\circ$



(b) record section with $\theta = 40^\circ$

Figure 6.8. Comparisons of synthetics (vertical component) for a spherical slow blob placed in the mid-mantle of S40RTS. The simulation parameters are summarised in Table 6.2, with $R_{\text{blob}} = 160$ km and $T_{\text{simu}} = 5$ s, or (d6) in Fig. 6.7. The two record sections are shown in Fig. 6.5. The blue traces are computed without the blob (pure S40RTS), and the red and green ones with the blob. SPECFEM is used for solution benchmark, but the SPECFEM traces (green) are almost invisible because the misfits are too small. The P-reflections from the blob are highlighted, coming right after the direct P-arrivals. As observed in (b), as the receiver moves away from the central azimuth, the backscattered energy (amplitude of P_{BS}) decays while the differential traveltimes between P and P_{BS} enlarges, visualising full 3-D backscattering effects.

a structural scale and velocity contrast cannot be resolved by linearised tomography. High-frequency forward modelling makes an ideal tool to study ULVZs. However, insufficient computing power has limited previous studies to either an axisymmetric in-plane 2-D scenario [105] or a lower frequency and a smaller model parameter space [104]. In this section, we present a preliminary forward study of a full 3-D ULVZ, targeted at a 1 Hz frequency.

The problem is described with Fig. 6.9. Basically, we follow [105] to investigate the core-diffracted phase SPdKS scattered by the ULVZ, except that we consider a 3-D ULVZ instead of an in-plane 2-D one. We use a monopole (axisymmetric) earthquake source and employ PREM as the background model, so that all azimuthal waveform variations result from the ULVZ. We adopt a depth-dependent n_u field, as shown in Fig. 6.9b. For such a large n_u field, a 1 Hz simulation with a 1-hr record length will consume 150,000 CPU-hours, which can be reduced by 60% if wavefield learning is used instead.

With such a depth-dependent n_u field, we intend to address the point that wavefield learning can deliver a near-optimal n_u field, but is not the only way to achieve solution convergence. In practice, one may design the n_u field based on a good understanding of the model as well as its wave scattering effects. In this example, we use $n_u = 1,000$ near the CMB, which gradually decreases to 100 across a depth range of 800 km and then keeps constant in the middle and upper mantle. It is designed based on our knowledge obtained in the last section, that is, the azimuthal wave scattering effects of a localised structure such as the ULVZ are concentrated in the near-field for a large structure-wavelength ratio. We can further optimise this n_u field by varying it with distance. If the 3-D ULVZ is placed in S40RTS rather than PREM, an efficient n_u field can be determined by superposing large near-field values (such as 1,000) onto Fig. 5.2a, the wavefield learning results for S40RTS.

The synthetics are shown in Fig. 6.10 and Fig. 6.11. We first look at the broadband waveforms filtered between 1 s and 12 s, that is, the left-hand side of both Fig. 6.10 and 6.11. The anomalies around the SPdKS phase are clearly manifested. As indicated by Fig. 6.10, even for the source-receiver plane across the ULVZ centre, the in-plane 2-D approximation already causes notable differences, especially for amplitudes, compared to the 3-D solution. Worse, such differences increase as the station moves away from




Figure 6.9. Problem setup for a cylinder-shaped ULVZ on CMB of PREM. The left-hand side of (a) shows the slice across the ULVZ centre, which looks exactly the same as [105] except that we implement the ULVZ as a 3-D cylinder, with a diameter (l) of 8° on CMB and a height (h) of 20 km, and with its edge located at 10° in distance. Several record sections are placed between 100° and 120° to observe the SPdKS anomalies, with a maximum azimuthal shift of 20° , as shown on the right-hand side of (a). An axisymmetric earthquake is used to get rid of source-induced radiation patterns. We adopt a depth-dependent n_u field, which equals to 1,000 near CMB and gradually decreases to 100 when moving away from CMB, as shown in (b). Such a n_u field leads to solution convergence but is obviously not optimal.

the central azimuth, as shown in Fig. 6.11. For example, the top trace in (b) shows that the impact of the ULVZ has remarkably weakened at the azimuth of 20° , where the corresponding ray path no longer penetrates the ULVZ, see Fig. 6.9a. Such information can be useful in constraining the azimuthal extent of the ULVZ [104], but cannot be utilised by axisymmetric approaches.

The high-frequency waveforms basically yield the same general conclusions, but display more details. With a dominant period around 1.5 s, they may resolve structures less than 10 km. We associate the high-frequency anomalies to waves bouncing vertically within the ULVZ. If this is the case, these data may help constraining both its thickness and length. To make solid interpretations of these high-frequency waveforms, more simulations with different model parameters and corresponding data are needed. In summary, AxiSEM3D opens doors to such kind of high-frequency forward studies for realistic 3-D Earth structures with affordable computing resources.

6.4 Summary

In this chapter, we show three applications of AxiSEM3D, with different foci from the surface to CMB.

In the first application, we model surface waves within a state-of-the-art global crustal model, Crust 1.0, and compare the synthetics to data. The large undulations and sharp slopes on the Moho discontinuity pose a great challenge to the implementation of the 3-D crust, for which different strategies are employed in SPECFEM and AxiSEM3D. In brief, AxiSEM3D honours the Moho in a more realistic way. Comparisons of synthetics to data show that the 3-D crust significantly enhances the accuracy for surface wave modelling. AxiSEM3D and SPECFEM agree well with each other at a lower frequency, but deviate to some degree at a higher frequency because of the different implementations. Comparisons to data show that AxiSEM3D achieves statistically higher cross-correlations than SPECFEM does for oceanic and mixed ocean-continent paths.

In the second application, we investigate the scattering effects of a localised high-amplitude heterogeneity in the mid-mantle. It is verified that AxiSEM3D can efficiently capture both forward and backward wave scattering across a wide range of structure-wavelength combinations, owing to the concentrated azimuthal scattering

effects in a near-field of the structure. Because such concentration strengthens with frequency, there exists a critical frequency at which the additional cost induced by the local structure reaches its maximum, an observation in favour of high-frequency applications. In addition, the revealed wavefield complexity suggests a type of seismic invisibility on the azimuthal direction.

The last application shows a preliminary study on a ULVZ at a 1 Hz frequency. We put a cylinder-shaped ULVZ atop CMB and observe the anomalies around the SPdKS phase. We verify the advantages of full 3-D simulations over 2-D in-plane ones: first, even for the source-receiver plane, the 2-D in-plane approximation already produces notable waveform differences, and second, with full 3-D simulations, off-plane data can be utilised to constrain the azimuthal dimension of structures, in spite of the healing effect of the azimuthal wave complexity. With a speedup of 3 orders of magnitude at 1 Hz compared to 3-D methods, AxiSEM3D opens doors to high-frequency forward studies for realistic 3-D Earth structures with affordable computing resources.




Figure 6.10. Synthetic seismograms (radial component) on the in-plane record section for a cylinder-shaped ULVZ atop CMB. The record section is aligned within the source-ULVZ plane. The left-hand side shows the broadband waveforms filtered between 1 s and 12 s, and the right-hand side the high-frequency waveforms filtered between 1 s and 3 s. Three models are compared, PREM-only (blue), PREM + 3-D ULVZ (red), and PREM + in-plane 2-D ULVZ (dotted green). Because both the source and the background model are axisymmetric, the waveforms for PREM-only and PREM + in-plane 2-D ULVZ do not vary with azimuth in (b). The traveltime curves for SPdKS are approximately drawn based on the broadband waveforms.




Figure 6.11. Synthetic seismograms (radial component) on the off-plane record section for a cylinder-shaped ULVZ atop CMB. The record section is orthogonal to the source-ULVZ plane, located the distance of 108° . Other figure settings are identical to Fig. 6.10.

Chapter 7

Discussions

7.1 Conclusions

Global wave propagation in 3-D Earth models is one of the most fundamental problems in seismology, carrying both the physics of wave scattering and the next-generation full-waveform inverse techniques. Though a number of 3-D numerical methods have been developed for this purpose, capable of considering all kinds of 3-D Earth features, the problem of insufficient computing power still remains outstanding, in view of Earth's multiscale nature of interest and the availability of broadband seismic data for its resolution.

This thesis develops a novel numerical method for global wave propagation in 3-D Earth models. Distinct from hardware and algorithmic optimisations, our method is physics-based, which substantially reduces the scale of the problem by exploiting the inherent spatial sparsity of global wavefields. We observe that global wavefields experience different complexity along different directions, that is, in the source-centred coordinate system, they vary much more smoothly along the azimuthal direction than within the meridian slices. This results from the fact that the lateral wave scattering effects of smooth Earth structures are too weak to substantially change the source-dominant radiation patterns.

To take advantage of such smoothness, we parameterise the azimuthal dimension of a global wavefield in Fourier series. In doing so, we reduce the 3-D wave equation to a coupled algebraic system of 2-D wave equations of the Fourier coefficients, defined on a 2-D meridian domain. We implement a 2-D spectral element method (SEM), AxiSEM3D, to solve such a system. Controlling both solution quality and

computational cost, the Fourier expansion order of wavefields is the most important parameter for AxiSEM3D, which can be locally adapted to the azimuthal complexity of wavefields. Two approaches are developed for its determination, an empirical equation for tomographic mantle models, and the technique of “wavefield learning” that determines a near-optimal expansion order field sufficiently large for the convergence of solutions. We employ the theory of particle relabelling transformation to consider finite undulations on any of the radial discontinuities.

Irrespective of the pseudo-spectral characterisation, as a stand-alone and convergent 3-D numerical method for global wave propagation, AxiSEM3D is characterised by the following properties:

1. *Comprehensive*: AxiSEM3D allows not only for material heterogeneities such as velocity, density, anisotropy and attenuation, but also for finite undulations on vertical discontinuities, and thereby a variety of aspherical Earth features such as ellipticity, topography and bathymetry, variation of crustal thickness, and core-mantle boundary topography. AxiSEM3D also proves robust to localised small-scale heterogeneities with sharp material contrast.
2. *Accurate*: AxiSEM3D is fully cross-validated with an independent discretised 3-D SEM, SPECFEM3D_GLOBE, for both volumetric and geometric complexities. The two methods slightly differ around the implementation of a 3-D crust, and comparisons of synthetics to data show that AxiSEM3D statistically resembles observed data better than SPECFEM does. To our knowledge, this is the first time two numerical methods are cross-validated for global wavefield modelling in realistic 3-D Earth structures.
3. *Fast*: For a period ranging from 50 s to 1 s, the speedup of AxiSEM3D with respect to a 3-D SEM can reach 1 to 3 orders of magnitude for the state-of-the-art tomographic mantle models, and 0 to 2 orders of magnitude for full 3-D Earth models with crustal structures. Such a speedup approximately scales with frequency. The observable frequency band (up to 1Hz) of global seismic data has been covered with affordable computing resources, opening doors to high-frequency body wave studies of the deep Earth.
4. *Adaptive*: Adaptivity is a central focus of AxiSEM3D. Through the spatially varying Fourier expansion order field, AxiSEM3D takes full advantage of the

spatio-temporal wavefield complexity. Model adaptivity is its most outstanding feature, bridging the gap between fast dimension-reduced methods for 1-D Earth models such as AxiSEM and slow discretised 3-D methods such as SPECFEM. In addition to model complexity, computational cost can be further saved by exploiting various sparsities in wavefields caused by record length, seismic phases of interest, source-structure-receiver geometry and source depth.

The MPI-based high-performance C++ code scales up to more than 10,000 cores, available open-source at <https://github.com/kuangdai/AxiSEM3D>.

7.2 Outlook

7.2.1 Fluid ocean

The ocean has a great impact on surface waves, as shown in Fig. 4.4. Physical realisation of a fluid ocean is essential to surface wave modelling, especially for high-frequency applications where the load approximation becomes inaccurate [52]. It also provides a means to study waves propagated in the ocean, such as the T-phases [132], tsunami [133], and seismic noise excitation [134].

The ocean floor is a multiply-connected solid-fluid discontinuity, which cannot be handled by the particle relabelling transformation. Therefore, a 3-D fluid ocean is currently not implemented in AxiSEM3D. We are examining possible solutions such as generalising the particle relabelling transformation for incompatible geometric mapping, or replacing the Fourier basis by discontinuous wavelets on the surface. An integrated 3-D ocean can be implemented in AxiSEM3D, which may help quantifying the accuracy of the load approximation, and may be applied to study oceanic events at a regional scale.

7.2.2 Seismic inversion

The vast speedup AxiSEM3D brings to forward modelling will indubitably facilitate seismic inverse problems.

First, it accelerates the computation of the spatio-temporal sensitivity kernels of broadband waveforms to both volumetric and geometric perturbations [44], with reference to a 3-D background model. The 3-D kernels are significant to surface wave

tomography, because the crustal perturbations are too large to be described by the linearised Born approximation around a 1-D background model [85].

AxiSEM3D may also serve as the underlining forward solver for adjoint tomography. Compared to 3-D solvers such as SPECFEM [54, 135] and Salvus [56], AxiSEM3D cannot simulate multiple sources simultaneously, as the source has to be placed on the axis. This limitation excludes finite-fault applications but does not affect adjoint methods, because we do not need the complete adjoint wavefields but only those components that actually contribute to the sensitivity kernels, which, fortunately, only involve certain lower Fourier expansion orders. This allows us to propagate back the adjoint wavefields from all the stations simultaneously, at a computational cost similar to (or even cheaper than) that of the forward simulation. In addition, because the Fourier series parameterisation naturally compresses wavefields based on their spatially-varying smoothness, AxiSEM3D requires orders-of-magnitude less memory than 3-D methods to store the entire forward and adjoint wavefields. This may enable the Gauss-Newton iterations based on the approximate Hessian [136] in addition to the gradient of the misfit functional, leading to accelerated convergence and enabling resolution study of the resultant model.

7.2.3 Regional to local scale

Both the mesher and the solver of AxiSEM3D are designed with the flexibility to include only a part of the D-shaped computational domain. Currently, AxiSEM3D is ready for both regional- and local-scale simulations. Absorbing boundary conditions have not been implemented, but it can be remedied by enlarging the actually needed computational domain. We will implement absorbing boundary conditions in the near future, such as the Engquist's [137] and the perfectly matched layer [138]. An interesting local-scale application has been carried out to study the remote warning system among African elephants by means of foot stomps.

7.2.4 Performance enhancement

The performance of AxiSEM3D can be further enhanced by a few technical refinements in the near future. The most promising two are described below.

First, AxiSEM3D can be accelerated with better load balancing for parallelisation. The current measure-and-decompose scheme described in Section 5.3 does not achieve perfect load balancing because of fluctuating runtime measurement. The parallelisation overhead is about 20% to 30%, compared to serial simulations. Such overhead may be minimised by means of profound system profiling for the packages EIGEN3 and FFTW3 at installation time.

Local time stepping [63] is another potential source for performance boost, especially for 3-D crustal models. For example, when Crust 1.0 is considered, the global time step of AxiSEM3D is nearly one third of that of SPECFEM, because one more layer is squeezed underneath the ocean. However, the bottleneck elements within the oceanic crust only take up less than 1% of the total. This makes a perfect scenario for the application of local time stepping. We strive to double the efficiency of AxiSEM3D in the presence of a 3-D crust (the third column in Table 5.3).

7.2.5 Usability

Easy implementation of user-defined models is the key to usability. The 2-D meshes used in AxiSEM3D are generated by the Salvus Mesher [56], a Python toolkit that can easily create local- to global-scale 2-D meshes for spherically symmetric (layered) models of Earth and other planets, such as Mars [139]. Three-dimensional structures are then imposed onto the 2-D meshes by AxiSEM3D. We provide two approaches to implement a user-defined 3-D model, either volumetric ones such as S40RTS or geometric ones such as an undulating CMB. First, AxiSEM3D accepts discrete models defined on structured grid points, supporting several formats such as the IRIS Earth Model Collaboration database (<http://ds.iris.edu/ds/products/emc/>) and the SubMachine database (<https://www.earth.ox.ac.uk/~smachine/>). Second, because AxiSEM3D is written in modern C++ with an object-oriented framework, a user-defined model can be implemented by class inheritance, with its encapsulated functionalities and data. This approach requires moderate programming experience with C++, but maximises the flexibility for model description.

In addition, we attach great importance to wavefield visualisation, as wavefields contain much more information than seismograms do. Wavefield snapshots and animations can show details of wave scattering effects at all scales, which can be useful

in both data interpretation and model assessment. The visualisation module will be released in the near future. For example, Fig. 7.1 shows wavefield snapshots in two PREM models with and without a slow heterogeneity in the lowermost mantle. The differential wavefield between the two clearly visualises the scattering effects of the heterogeneity. It can help not only detecting or understanding data anomalies in terms of their structure-induced propagation effects, but also identifying measurements (absolute or differential) that eliminate or alleviate structural trade-offs. If the seismic array is dense enough, such as the USArray, we may even back-propagate the differential wavefield observed on the surface to directly identify the causal structure in the deep interior.

7.2.6 Rotation and self-gravitation

Earth’s rotation modifies global wavefields by two means, ellipticity and Coriolis effect. Ellipticity has been incorporated in AxiSEM3D by applying the particle relabelling transformation, but the Coriolis effect has not been considered so far. Though Coriolis forces can be implemented in AxiSEM3D without much difficulty, this task is given a low priority because of its small impact [52], especially for high-frequency applications at which AxiSEM3D is aimed.

Self-gravitation may have a non-trivial effect on the low-frequency surface waves and normal modes in response to large earthquakes [40]. To consider self-gravitation, SPECFEM and Salvus adopt the Cowling’s approximation [52, 115] that ignores the effects of mass redistribution. Though the Cowling’s approximation is fully compatible with the AxiSEM3D framework, we decide to incorporate full gravity effects by coupling the wave equation and the Poisson’s equation of gravity. The object-orientated framework is friendly to the introduction of multi-physics.




Figure 7.1. Original and differential wavefields on two PREM models with and without a slow heterogeneity in the lowermost mantle. At the moment of 375 s, the differential wavefield is zero everywhere because the P-wavefront has not arrived the heterogeneity, so we use the upper-right slot to display the setup. Each column shares the same model while each row the same time-stamp. The heterogeneity behaves like a source of the differential wavefield.

Bibliography

- [1] Stephen G Brush. Discovery of the Earth’s core. *American Journal of Physics*, 48(9):705–724, 1980.
- [2] N Rawlinson, S Pozgay, and S Fishwick. Seismic tomography: a window into deep Earth. *Physics of the Earth and Planetary Interiors*, 178(3):101–135, 2010.
- [3] Adam M Dziewonski and Don L Anderson. Preliminary reference earth model. *Physics of the earth and planetary interiors*, 25(4):297–356, 1981.
- [4] BLN Kennett and ER Engdahl. Traveltimes for global earthquake location and phase identification. *Geophysical Journal International*, 105(2):429–465, 1991.
- [5] BLN Kennett, ER Engdahl, and R Buland. Constraints on seismic velocities in the Earth from traveltimes. *Geophysical Journal International*, 122(1):108–124, 1995.
- [6] George R Helffrich and Bernard J Wood. The Earth’s mantle. *Nature*, 412(6846):501, 2001.
- [7] Luciana Astiz, Paul Earle, and Peter Shearer. Global stacking of broadband seismograms. *Seismological Research Letters*, 67(4):8–18, 1996.
- [8] Chad Trabant, Alexander R Hutko, Manochehr Bahavar, Richard Karstens, Timothy Ahern, and Richard Aster. Data products at the IRIS DMC: Stepping stones for research and other applications. *Seismological Research Letters*, 83(5):846–854, 2012.
- [9] John H Woodhouse and Adam M Dziewonski. Mapping the upper mantle: Three-dimensional modeling of Earth structure by inversion of seismic waveforms. *Journal of Geophysical Research: Solid Earth*, 89(B7):5953–5986, 1984.

- [10] Wei-jia Su and Adam M Dziewonski. Predominance of long-wavelength heterogeneity in the mantle. *Nature*, 352(6331):121–126, 1991.
- [11] Thorsten W Becker and Lapo Boschi. A comparison of tomographic and geodynamic mantle models. *Geochemistry, Geophysics, Geosystems*, 3(1), 2002.
- [12] Barbara Romanowicz. Using seismic waves to image earth’s internal structure. *Nature*, 451(7176):266, 2008.
- [13] Vedran Lekic, Sanne Cottaar, Adam Dziewonski, and Barbara Romanowicz. Cluster analysis of global lower mantle tomography: A new class of structure and implications for chemical heterogeneity. *Earth and Planetary Science Letters*, 357:68–77, 2012.
- [14] Akiko To, Barbara Romanowicz, Yann Capdeville, and Nozomu Takeuchi. 3D effects of sharp boundaries at the borders of the African and Pacific Superplumes: Observation and modeling. *Earth and Planetary Science Letters*, 233(1):137–153, 2005.
- [15] Lianxing Wen and Donald V Helmberger. Ultra-low velocity zones near the core-mantle boundary from broadband PKP precursors. *Science*, 279(5357):1701–1703, 1998.
- [16] Edward J Garnero, Justin Revenaugh, Quentin Williams, Thorne Lay, and Louise H Kellogg. Ultralow velocity zone at the core-mantle boundary. *The core-mantle boundary region*, pages 319–334, 1998.
- [17] Sh Maruyama, M Santosh, and D Zhao. Superplume, supercontinent, and post-perovskite: mantle dynamics and anti-plate tectonics on the core–mantle boundary. *Gondwana Research*, 11(1):7–37, 2007.
- [18] Motohiko Murakami, Kei Hirose, Katsuyuki Kawamura, Nagayoshi Sata, and Yasuo Ohishi. Post-perovskite phase transition in MgSiO₃. *Science*, 304(5672):855–858, 2004.
- [19] Wendy L Mao, Ho-kwang Mao, Wolfgang Sturhahn, Jiyong Zhao, Vitali B Prakapenka, Yue Meng, Jinfu Shu, Yingwei Fei, and Russell J Hemley.

- Iron-rich post-perovskite and the origin of ultralow-velocity zones. *Science*, 312(5773):564–565, 2006.
- [20] Allen K McNamara, Edward J Garnero, and Sebastian Rost. Tracking deep mantle reservoirs with ultra-low velocity zones. *Earth and Planetary Science Letters*, 299(1):1–9, 2010.
- [21] Bernhard Steinberger and Trond H Torsvik. A geodynamic model of plumes from the margins of Large Low Shear Velocity Provinces. *Geochemistry, Geophysics, Geosystems*, 13(1), 2012.
- [22] Adam M Dziewonski, Vedran Lekic, and Barbara A Romanowicz. Mantle anchor structure: an argument for bottom up tectonics. *Earth and Planetary Science Letters*, 299(1):69–79, 2010.
- [23] Miaki Ishii and Jeroen Tromp. Normal-mode and free-air gravity constraints on lateral variations in velocity and density of Earth’s mantle. *Science*, 285(5431):1231–1236, 1999.
- [24] Paula Koelemeijer, Arwen Deuss, and Jeroen Ritsema. Density structure of Earth’s lowermost mantle from stoneley mode splitting observations. *Nature Communications*, 8, 2017.
- [25] MK Savage. Seismic anisotropy and mantle deformation: what have we learned from shear wave splitting? *Reviews of Geophysics*, 37(1):65–106, 1999.
- [26] Maureen D Long and Thorsten W Becker. Mantle dynamics and seismic anisotropy. *Earth and Planetary Science Letters*, 297(3):341–354, 2010.
- [27] Y Gung and B Romanowicz. Q tomography of the upper mantle using three-component long-period waveforms. *Geophysical Journal International*, 157(2):813–830, 2004.
- [28] Colleen A Dalton, Göran Ekström, and Adam M Dziewoński. The global attenuation structure of the upper mantle. *Journal of Geophysical Research: Solid Earth*, 113(B9), 2008.

- [29] Yang Luo, Jeroen Tromp, Bertrand Denel, and Henri Calandra. 3D coupled acoustic-elastic migration with topography and bathymetry based on spectral-element and adjoint methods. *Geophysics*, 78(4):S193–S202, 2013.
- [30] Lucia Gualtieri, Éléonore Stutzmann, Véronique Farra, Yann Capdeville, M Schimmel, Fabrice Ardhuin, and Andrea Morelli. Modelling the ocean site effect on seismic noise body waves. *Geophysical Journal International*, 197(2):1096–1106, 2014.
- [31] Gabi Laske, Guy Masters, Zhitu Ma, and Mike Pasyanos. Update on CRUST 1.0 - a 1-degree global model of Earth’s crust. In *Geophysical Research Abstracts*, volume 15, page 2658, 2013.
- [32] Robert S White, Dan McKenzie, and R Keith O’Nions. Oceanic crustal thickness from seismic measurements and rare earth element inversions. *Journal of Geophysical Research: Solid Earth*, 97(B13):19683–19715, 1992.
- [33] Arwen Deuss, Simon AT Redfern, Kit Chambers, and John H Woodhouse. The nature of the 660-kilometer discontinuity in Earth’s mantle from global seismic observations of PP precursors. *Science*, 311(5758):198–201, 2006.
- [34] B Kustowski, G Ekström, and AM Dziewoński. Anisotropic shear-wave velocity structure of the Earth’s mantle: A global model. *Journal of Geophysical Research: Solid Earth (1978–2012)*, 113(B6), 2008.
- [35] Andrea Colombi, Tarje Nissen-Meyer, Lapo Boschi, and Domenico Giardini. Seismic waveform sensitivity to global boundary topography. *Geophysical Journal International*, 191(2):832–848, 2012.
- [36] Gaia Soldati, Paula Koelemeijer, Lapo Boschi, and Arwen Deuss. Constraints on core-mantle boundary topography from normal mode splitting. *Geochemistry, Geophysics, Geosystems*, 14(5):1333–1342, 2013.
- [37] Guust Nolet. A breviary of seismic tomography. *Imaging the Interior*, 2008.
- [38] R Montelli, G Nolet, G Masters, FA Dahlen, and S-H Hung. Global P and PP traveltimes tomography: rays versus waves. *Geophysical Journal International*, 158(2):637–654, 2004.

- [39] J Ritsema, A Deuss, HJ Van Heijst, and JH Woodhouse. S40RTS: a degree-40 shear-velocity model for the mantle from new rayleigh wave dispersion, teleseismic traveltime and normal-mode splitting function measurements. *Geophysical Journal International*, 184(3):1223–1236, 2011.
- [40] FAand Dahlen and Jeroen Tromp. *Theoretical global seismology*. Princeton university press, 1998.
- [41] FA Dahlen, S-H Hung, and Guust Nolet. Fréchet kernels for finite-frequency traveltimes – I. theory. *Geophysical Journal International*, 141(1):157–174, 2000.
- [42] Karin Sigloch. Mantle provinces under North America from multifrequency P wave tomography. *Geochemistry, Geophysics, Geosystems*, 12(2), 2011.
- [43] Kasra Hosseini and Karin Sigloch. Multifrequency measurements of core-diffracted P waves (Pdiff) for global waveform tomography. *Geophysical Journal International*, 203(1):506–521, 2015.
- [44] Tarje Nissen-Meyer, FA Dahlen, and Alexandre Fournier. Spherical-earth fréchet sensitivity kernels. *Geophysical Journal International*, 168(3):1051–1066, 2007.
- [45] Simon C Stähler, Karin Sigloch, and Tarje Nissen-Meyer. Triplicated P-wave measurements for waveform tomography of the mantle transition zone. *Solid Earth*, 3(2):339, 2012.
- [46] Hejun Zhu, Ebru Bozdağ, and Jeroen Tromp. Seismic structure of the European upper mantle based on adjoint tomography. *Geophysical Journal International*, 201(1):18–52, 2015.
- [47] Ebru Bozdağ, Daniel Peter, Matthieu Lefebvre, Dimitri Komatitsch, Jeroen Tromp, Judith Hill, Norbert Podhorszki, and David Pugmire. Global adjoint tomography: first-generation model. *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, 207(3):1739–1766, 2016.

- [48] Kuangdai Leng, Tarje Nissen-Meyer, and Martin van Driel. Efficient global wave propagation adapted to 3-D structural complexity: a pseudospectral/spectral-element approach. *Geophysical Journal International*, 207(3):1700–1721, 2016.
- [49] Heiner Igel. *Computational seismology: a practical introduction*. Oxford University Press, 2017.
- [50] Heiner Igel, Tarje Nissen-Meyer, and Gunnar Jahnke. Wave propagation in 3D spherical sections: effects of subduction zones. *Physics of the Earth and Planetary Interiors*, 132(1):219–234, 2002.
- [51] Dimitri Komatitsch and Jeroen Tromp. Spectral-element simulations of global seismic wave propagation – I. validation. *Geophysical Journal International*, 149(2):390–412, 2002.
- [52] Dimitri Komatitsch and Jeroen Tromp. Spectral-element simulations of global seismic wave propagation – II. three-dimensional models, oceans, rotation and self-gravitation. *Geophysical Journal International*, 150(1):303–318, 2002.
- [53] Emmanuel Chaljub, Dimitri Komatitsch, Jean-Pierre Vilotte, Yann Capdeville, Bernard Valette, and Gaetano Festa. Spectral-element analysis in seismology. *Advances in Geophysics*, 48:365–419, 2007.
- [54] Daniel Peter, Dimitri Komatitsch, Yang Luo, Roland Martin, Nicolas Le Goff, Emanuele Casarotti, Pieyre Le Loher, Federica Magnoni, Qinya Liu, Céline Blitz, et al. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes. *Geophysical Journal International*, 186(2):721–739, 2011.
- [55] Michael Dumbser and Martin Käser. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – II. the three-dimensional isotropic case. *Geophysical Journal International*, 167(1):319–336, 2006.
- [56] Michael Afanasiev, Christian Boehm, Martin van Driel, Lion Krischer, Dave May, Max Rietmann, and Andreas Fichtner. Salvus: A flexible high-performance and open-source package for waveform modelling and inversion

- from laboratory to global scales. *Geophysical Research Abstracts*, 19:EGU2017–9456, 2017.
- [57] Y Capdeville, A To, and B Romanowicz. Coupling spectral elements and modes in a spherical Earth: an extension to the ‘sandwich’ case. *Geophysical Journal International*, 154(1):44–57, 2003.
 - [58] Alexey Gokhberg and Andreas Fichtner. Full-waveform inversion on heterogeneous HPC systems. *Computers & Geosciences*, 89:260–268, 2016.
 - [59] Max Rietmann, Peter Messmer, Tarje Nissen-Meyer, Daniel Peter, Piero Basini, Dimitri Komatitsch, Olaf Schenk, Jeroen Tromp, Lapo Boschi, and Domenico Giardini. Forward and adjoint simulations of seismic wave propagation on emerging large-scale GPU architectures. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, page 38. IEEE Computer Society Press, 2012.
 - [60] Laura Carrington, Dimitri Komatitsch, Michael Laurenzano, Mustafa M Tikir, David Michéa, Nicolas Le Goff, Allan Snively, and Jeroen Tromp. High-frequency simulations of global seismic wave propagation using SPECFEM3D_GLOBE on 62k processors. In *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, page 60. IEEE Press, 2008.
 - [61] Seiji Tsuboi, Kazuto Ando, Takayuki Miyoshi, Daniel Peter, Dimitri Komatitsch, and Jeroen Tromp. A 1.8 trillion degrees-of-freedom, 1.24 petaflops global seismic wave simulation on the K computer. *The International Journal of High Performance Computing Applications*, 30(4):411–422, 2016.
 - [62] Dimitri Komatitsch. Fluid–solid coupling on a cluster of GPU graphics cards for seismic wave propagation. *Comptes Rendus Mécanique*, 339(2):125–135, 2011.
 - [63] Max Rietmann, Daniel Peter, Olaf Schenk, Bora Uçar, and Marcus Grote. Load-balanced local time stepping for large-scale wave propagation. In *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International*, pages 925–935. IEEE, 2015.

- [64] Xiao Ma, Dinghui Yang, Guojie Song, and Meixia Wang. A low-dispersive symplectic partitioned Runge–Kutta method for solving seismic-wave equations: I. scheme and theoretical analysis. *Bulletin of the Seismological Society of America*, 2014.
- [65] Tarje Nissen-Meyer, Alexandre Fournier, and FA Dahlen. A 2-d spectral-element method for computing spherical-earth seismograms – II. waves in solid–fluid media. *Geophysical Journal International*, 174(3):873–888, 2008.
- [66] Michel O Deville, Paul F Fischer, and Ernest H Mund. *High-order methods for incompressible fluid flow*, volume 9. Cambridge University Press, 2002.
- [67] Tarje Nissen-Meyer, Alexandre Fournier, and Francis Anthony Dahlen. A two-dimensional spectral-element method for computing spherical-earth seismograms – I. moment-tensor source. *Geophysical Journal International*, 168(3):1067–1092, 2007.
- [68] Martin van Driel and Tarje Nissen-Meyer. Optimized viscoelastic wave propagation for weakly dissipative media. *Geophysical Journal International*, 199(2):1078–1093, 2014.
- [69] Paul Cupillard, Elise Delavaud, Gaël Burgos, Geatano Festa, Jean-Pierre Vilotte, Yann Capdeville, and Jean-Paul Montagner. RegSEM: a versatile code based on the spectral element method to compute seismic wave propagation at the regional scale. *Geophysical Journal International*, 188(3):1203–1220, 2012.
- [70] M van Driel, L Krischer, SC Stähler, K Hosseini, and T Nissen-Meyer. Instaseis: instant global seismograms based on a broadband waveform database. *Solid Earth Discussions*, 7(1):957–1005, 2015.
- [71] Harold Jeffreys and Keith Edward Bullen. *Seismological tables*. Office of the British Association, 1958.
- [72] Kenji Kawai, Nozomu Takeuchi, and Robert J Geller. Complete synthetic seismograms up to 2 Hz for transversely isotropic spherically symmetric media. *Geophysical Journal International*, 164(2):411–424, 2006.

- [73] Wolfgang Friederich and Jörg Dalkolmo. Complete synthetic seismograms for a spherically symmetric earth by a numerical computation of the Green’s function in the frequency domain. *Geophysical Journal International*, 122(2):537–550, 1995.
- [74] David Al-Attar and John H Woodhouse. Calculation of seismic displacement fields in self-gravitating earth models—applications of minors vectors and symplectic structure. *Geophysical Journal International*, 175(3):1176–1208, 2008.
- [75] Lion Krischer, Alexander R Hutko, Martin van Driel, Simon Stähler, Manochehr Bahavar, Chad Trabant, and Tarje Nissen-Meyer. On-demand custom broadband synthetic seismograms. *Seismological Research Letters*, 2017.
- [76] Tarje Nissen-Meyer, Martin van Driel, SC Stähler, Kasra Hosseini, Stefanie Hempel, Ludwig Auer, A Colombi, and Anne Fournier. AxiSEM: broadband 3-d seismic wavefields in axisymmetric media. *Solid Earth*, 5(1):425–445, 2014.
- [77] Heiner Igel and Michael Weber. SH-wave propagation in the whole mantle using high-order finite differences. *Geophysical Research Letters*, 22(6):731–734, 1995.
- [78] Genti Toyokuni and Hiroshi Takenaka. FDM computation of seismic wavefield for an axisymmetric earth with a moment tensor point source. *Earth, planets and space*, 58(8):e29–e32, 2006.
- [79] Yder Masson, Paul Cupillard, Yann Capdeville, and Barbara Romanowicz. On the numerical implementation of time-reversal mirrors for tomographic imaging. *Geophysical Journal International*, page ggt459, 2013.
- [80] Vadim Monteiller, Sébastien Chevrot, Dimitri Komatitsch, and Yi Wang. Three-dimensional full waveform inversion of short-period teleseismic wavefields based upon the SEM–DSM hybrid method. *Geophysical Journal International*, 202(2):811–827, 2015.
- [81] Vadim Monteiller, Sébastien Chevrot, Dimitri Komatitsch, and Nobuaki Fuji. A hybrid method to compute short-period synthetic seismograms of teleseismic body waves in a 3-D regional model. *Geophysical Journal International*, 192(1):230–247, 2012.

- [82] Hejun Zhu, Yang Luo, Tarje Nissen-Meyer, Christina Morency, and Jeroen Tromp. Elastic imaging and time-lapse migration based on adjoint methods. *Geophysics*, 74(6):WCA167–WCA177, 2009.
- [83] Yann Capdeville, Éléonore Stutzmann, Nian Wang, and Jean-Paul Montagner. Residual homogenization for seismic forward and inverse problems in layered media. *Geophysical Journal International*, 194(1):470–487, 2013.
- [84] Piero Basini, Tarje Nissen-Meyer, Lapo Boschi, Emanuele Casarotti, Julie Verbeke, Olaf Schenk, and Domenico Giardini. The influence of nonuniform ambient noise on crustal tomography in Europe. *Geochemistry, Geophysics, Geosystems*, 14(5):1471–1492, 2013.
- [85] Ying Zhou, FA Dahlen, Guust Nolet, and Gabi Laske. Finite-frequency effects in global surface-wave tomography. *Geophysical Journal International*, 163(3):1087–1111, 2005.
- [86] Julie J Nazareth and Egill Hauksson. The seismogenic thickness of the southern California crust. *Bulletin of the Seismological Society of America*, 94(3):940–960, 2004.
- [87] Tae-Kyung Hong, Chang-Eob Baag, Hoseon Choi, and Dong-Hoon Sheen. Regional seismic observations of the 9 October 2006 underground nuclear explosion in North Korea and the influence of crustal structure on regional phases. *Journal of Geophysical Research: Solid Earth*, 113(B3), 2008.
- [88] Andrea Colombi, Tarje Nissen-Meyer, Lapo Boschi, and Domenico Giardini. Seismic waveform inversion for core–mantle boundary topography. *Geophysical Journal International*, page ggu112, 2014.
- [89] Christian Pelties, Josep Puente, Jean-Paul Ampuero, Gilbert B Brietzke, and Martin Käser. Three-dimensional dynamic rupture simulation with a high-order discontinuous Galerkin method on unstructured tetrahedral meshes. *Journal of Geophysical Research: Solid Earth*, 117(B2), 2012.
- [90] David Al-Attar and Ophelia Crawford. Particle relabelling transformations in elastodynamics. *Geophysical Journal International*, 205(1):575–593, 2016.

- [91] Kuangdai Leng, Tarje Nissen-Meyer, Martin van Driel, Kasra Hosseini, and David Al-Attar. AxiSEM3D: broadband seismic wavefields in 3-D global Earth models with undulating discontinuities. *Geophysical Journal International*, forthcoming.
- [92] L Auer, L Boschi, TW Becker, T Nissen-Meyer, and D Giardini. Savani: A variable resolution whole-mantle model of anisotropic shear velocity variations based on multiple data sets. *Journal of Geophysical Research: Solid Earth*, 119(4):3006–3034, 2014.
- [93] Carl Tape, Qinya Liu, Alessia Maggi, and Jeroen Tromp. Adjoint tomography of the southern California crust. *Science*, 325(5943):988–992, 2009.
- [94] Andreas Fichtner, Brian LN Kennett, Heiner Igel, and Hans-Peter Bunge. Theoretical background for continental-and global-scale full-waveform inversion in the time–frequency domain. *Geophysical Journal International*, 175(2):665–685, 2008.
- [95] Miriam Kristeková, Jozef Kristek, and Peter Moczo. Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals. *Geophysical Journal International*, 178(2):813–825, 2009.
- [96] Ebru Bozdağ, Jeannot Trampert, and Jeroen Tromp. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements. *Geophysical Journal International*, 185(2):845–870, 2011.
- [97] Pawan Bharadwaj, Tarje Nissen-Meyer, Gerard Schuster, and P Martin Mai. Enhancing core-diffracted arrivals by supervirtual interferometry. *Geophysical Journal International*, page ggt452, 2013.
- [98] R-E Plessix. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2):495–503, 2006.
- [99] Yang Luo, Hejun Zhu, Tarje Nissen-Meyer, Christina Morency, and Jeroen Tromp. Seismic modeling and imaging based upon spectral-element and adjoint methods. *The Leading Edge*, 28(5):568–574, 2009.

- [100] Qinya Liu and Jeroen Tromp. Finite-frequency sensitivity kernels for global seismic wave propagation based upon adjoint methods. *Geophysical Journal International*, 174(1):265–286, 2008.
- [101] Albert Tarantola. Theoretical background for the inversion of seismic waveforms including elasticity and attenuation. *Pure and Applied Geophysics*, 128(1-2):365–399, 1988.
- [102] Jeroen Tromp, Carl Tape, and Qinya Liu. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophysical Journal International*, 160(1):195–216, 2005.
- [103] Jonas Mockus. *Bayesian approach to global optimization: theory and applications*, volume 37. Springer Science & Business Media, 2012.
- [104] Sanne Cottaar and Barbara Romanowicz. An unusually large ULVZ at the base of the mantle near Hawaii. *Earth and Planetary Science Letters*, 355:213–222, 2012.
- [105] Michael S Thorne, Edward J Garnero, Gunnar Jahnke, Heiner Igel, and Allen K McNamara. Mega ultra low velocity zone and mantle flow. *Earth and Planetary Science Letters*, 364:59–67, 2013.
- [106] Keiiti Aki and Paul G Richards. *Quantitative seismology, 1st Edition*. University Science Books, 1980.
- [107] AS Alekseev and BG Mikhailenko. The solution of dynamic problems of elastic wave-propagation in inhomogeneous-media by a combination of partial separation of variables and finite-difference methods. *Journal of Geophysics - Zeitschrift fur Geophysik*, 48(3):161–172, 1980.
- [108] BG Mikhailenko. Synthetic seismograms for complex three-dimensional geometries using an analytical-numerical algorithm. *Geophysical Journal International*, 79(3):963–986, 1984.
- [109] BG Mikhailenko and VI Korneev. Calculation of synthetic seismograms for complex subsurface geometries by a combination of finite integral fourier-transforms

- and finite-difference techniques. *Journal of Geophysics - Zeitschrift für Geophysik*, 54(3):195–206, 1984.
- [110] John P Boyd. *Chebyshev and Fourier spectral methods*. Courier Corporation, 2001.
 - [111] Martin van Driel and Tarje Nissen-Meyer. Seismic wave propagation in fully anisotropic axisymmetric media. *Geophysical Journal International*, 199(2):880–893, 2014.
 - [112] Jeroen Ritsema and HJ Van Heijst. Seismic imaging of structural heterogeneity in Earth’s mantle: evidence for large-scale mantle flow. *Science Progress*, 83:243–259, 1999.
 - [113] S Chevrot, JP Montagner, and R Snieder. The spectrum of tomographic earth models. *Geophysical Journal International*, 133(3):783–788, 1998.
 - [114] Christine Bernardi, Monique Dauge, Yvon Maday, and Meiji Azaïez. *Spectral methods for axisymmetric domains*. Gauthier-Villars Paris, 1999.
 - [115] Emmanuel Chaljub and Bernard Valette. Spectral element modelling of three-dimensional wave propagation in a self-gravitating Earth with an arbitrarily stratified outer core. *Geophysical Journal International*, 158(1):131–141, 2004.
 - [116] JM Lopez and Jie Shen. An efficient spectral-projection method for the Navier–Stokes equations in cylindrical geometries: I. axisymmetric cases. *Journal of Computational Physics*, 139(2):308–326, 1998.
 - [117] Alexandre Fournier, Hans-Peter Bunge, Rainer Hollerbach, and Jean-Pierre Vilotte. A Fourier-spectral element algorithm for thermal convection in rotating axisymmetric containers. *Journal of Computational Physics*, 204(2):462–489, 2005.
 - [118] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.

- [119] Matteo Frigo and Steven G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [120] David Gottlieb and Jan S Hesthaven. Spectral methods for hyperbolic problems. *Journal of Computational and Applied Mathematics*, 128(1):83–131, 2001.
- [121] David Gottlieb, Steven A Orszag, and Eli Turkel. Stability of pseudospectral and finite-difference methods for variable coefficient problems. *mathematics of computation*, 37(156):293–305, 1981.
- [122] Dimitri Komatitsch and Jean-Pierre Vilotte. The spectral element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures. *Bulletin of the seismological society of America*, 88(2):368–392, 1998.
- [123] John G Anderson. Quantitative measure of the goodness-of-fit of synthetic seismograms. In *13th World Conference on Earthquake Engineering Conference Proceedings, Vancouver, Canada, Paper*, volume 243, 2004.
- [124] Lion Krischer, Tobias Megies, Robert Barsch, Moritz Beyreuther, Thomas Lecocq, Corentin Caudron, and Joachim Wassermann. ObsPy: a bridge for seismology into the scientific python ecosystem. *Computational Science & Discovery*, 8(1):014003, 2015.
- [125] Nozomu Takeuchi. Finite boundary perturbation theory for the elastic equation of motion. *Geophysical Journal International*, 160(3):1044–1058, 2005.
- [126] National Geophysical Data Center. 2-minute gridded global relief data (etopo2v2). 2006.
- [127] C Bassin, G Laske, and G Masters. The current limits of resolution for surface wave tomography in North America. *Eos*, 81, 2000.
- [128] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.

- [129] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [130] Jennifer Jenkins, Arwen Deuss, and Sanne Cottaar. Converted phases from sharp 1000 km depth mid-mantle heterogeneity beneath Western Europe. *Earth and Planetary Science Letters*, 459:196–207, 2017.
- [131] Jun Korenaga. Constraining the geometries of small-scale heterogeneities: A case study from the Mariana region. *Journal of Geophysical Research: Solid Earth*, 120(11):7830–7851, 2015.
- [132] Jeffry L Stevens, G Eli Baker, Ron W Cook, LD Gerald, Lewis P Berger, and Steven M Day. Empirical and numerical modeling of T-phase propagation from ocean to land. In *Monitoring the Comprehensive Nuclear-Test-Ban Treaty: Hydroacoustics*, pages 531–565. Springer, 2001.
- [133] Zygmunt Kowalik, William Knight, Tom Logan, and Paul Whitmore. Numerical modeling of the global tsunami: Indonesian tsunami of 26 December 2004. *Science of Tsunami Hazards*, 23(1):40–56, 2005.
- [134] Toshiro Tanimoto. Equivalent forces for colliding ocean waves. *Geophysical Journal International*, 181(1):468–478, 2010.
- [135] Jeroen Tromp, Dimitri Komatitsch, and Qinya Liu. Spectral-element and adjoint methods in seismology. *Communications in Computational Physics*, 3(1):1–32, 2008.
- [136] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [137] Björn Engquist and Andrew Majda. Absorbing boundary conditions for numerical simulation of waves. *Proceedings of the National Academy of Sciences*, 74(5):1765–1766, 1977.
- [138] Dimitri Komatitsch and Jeroen Tromp. A perfectly matched layer absorbing boundary condition for the second-order seismic wave equation. *Geophysical Journal International*, 154(1):146–153, 2003.

- [139] Ebru Bozdağ, Youyi Ruan, Nathan Mettetz, Amir Khan, Kuangdai Leng, Martin van Driel, Mark Wieczorek, Attilio Rivoldini, Carène S Larmat, Domenico Giardini, et al. Simulations of seismic wave propagation on Mars. *Space Science Reviews*, pages 1–24, 2016.