AxiSEM3D: broad-band seismic wavefields in 3-D global earth models with undulating discontinuities

Kuangdai Leng,[1,2] Tarje Nissen-Meyer,[1] Martin van Driel,[3] Kasra Hosseini[1] and David Al-Attar[4]

[1]*Department of Earth Sciences, University of Oxford, South Parks Road, Oxford OX*1 3*AN, UK. E-mail: kuangdai.leng@earth.ox.ac.uk*
[2]*Key Laboratory of Transportation Tunnel Engineering, Southwest Jiaotong University, Chengdu 610031, China*
[3]*Institute of Geophysics, ETH Zurich, Sonneggstrasse 5, 8092 Zurich, Switzerland*
[4]*Bullard Laboratories, University of Cambridge, Madingley Road, Cambridge CB*3 *OEZ, UK*

SUMMARY

We present a novel numerical method to simulate global seismic wave propagation in realistic aspherical 3-D earth models across the observable frequency band of global seismic data. Our method, named AxiSEM3D, is a hybrid of spectral element method (SEM) and pseudospectral method. It describes the azimuthal dimension of global wavefields with a substantially reduced number of degrees of freedom via a global Fourier series parametrization, of which the number of terms can be locally adapted to the inherent azimuthal complexity of the wavefields. 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, surface topography, variable crustal thickness, undulating transition zone and core–mantle boundary topography. Undulating discontinuities are honoured by means of the 'particle relabelling transformation', so that the spectral element mesh can be kept spherical. The implementation of the particle relabelling transformation is verified by benchmark solutions against a discretized 3-D SEM, considering ellipticity, topography and bathymetry (with the ocean approximated as a hydrodynamic load) and a tomographic mantle model with an undulating transition zone. For the state-of-the-art global tomographic models with aspherical geometry but without a 3-D crust, efficiency comparisons suggest that AxiSEM3D can be two to three orders of magnitude faster than a discretized 3-D method for a seismic period at 5 s or below, with the speed-up increasing with frequency and decreasing with model complexity. We also verify AxiSEM3D for localized small-scale heterogeneities with strong perturbation strength. With reasonable computing resources, we have achieved a corner frequency of up to 1 Hz for 3-D mantle models.

Key words: Structure of the Earth; Numerical solutions; Computational seismology; Theoretical seismology; Wave propagation.

1 INTRODUCTION

For more than a century, seismology has served as the primary tool for data-informed inference of Earth structure and dynamics. Seismic tomography has revealed a consensual 3-D picture of the Earth's interior, dominated by long-wavelength thermal and compositional heterogeneities (Su & Dziewonski 1991; Becker & Boschi 2002; Romanowicz 2003). Most achievements in tomography have started from a radially symmetric Earth structure (e.g. PREM, Dziewonski & Anderson 1981) and have utilized a certain measurement extracted from full waveforms, such as the

ray theory (e.g. Auer *et al.* 2014), normal-mode summation (e.g. Ishii & Tromp 2004), finite-frequency theory (e.g. Hosseini & Sigloch 2015) as well as their combinations (e.g. Ritsema *et al.* 2011; Zaroli *et al.* 2015). Nevertheless, aimed at computing the full elastodynamic response of a physically realistic 3-D earth model to an earthquake source, global wave propagation has always remained as one of the most fundamental pursuits in seismology, particularly encouraged by the ongoing prosperity of broad-band waveform data (see e.g. IRIS annual report 2015, http://www.iris.edu/hq/publications/annual_reports) and high-performance computing.

Numerical waveform modelling is the most comprehensive approach to study wave physics, inevitable for non-trivial wave phenomena such as diffraction (Hosseini & Sigloch 2015) and triplication (Stähler *et al.* 2012). It fully captures structure-induced wave anomalies in different regimes of wave scattering as a function of wavelength, structural scale and strength, and propagation path (Aki & Richards 1980; Igel 2017). Such wave scattering effects are fundamental to all seismological probings of the Earth's interior that rely on understanding how waveforms encode information from heterogeneous multiscale structures. For instances, waveform modelling has been applied to explain the seismic insensitivity of mantle plumes (Hwang *et al.* 2011; Maguire *et al.* 2016), to investigate the polarity of phase conversions at subducted slabs (Haugland *et al.* 2017), and to verify complex seismic probes of the ultralow velocity zones (ULVZs; Thorne *et al.* 2013).

Seismic tomography has delivered plentiful 3-D earth models using different inverse approaches and data sets [see e.g. the collections of IRIS-EMC (http://ds.iris.edu/ds/products/emc/, Trabant *et al.* 2012) and SubMachine (http://submachine.earth.ox.ac.uk/, Hosseini *et al.* 2018)]. As these tomographic models diverge at shorter wavelengths, a central question arises as to what extent our knowledge about the Earth's 3-D structure has approached to the underneath truth since the early-established long-wavelength pictures such as M84 (Woodhouse & Dziewonski 1984). This question seems to remain poorly answered because comparisons among tomographic models have been limited within the model space, mainly by visual inspection or power spectra analysis (e.g. Becker & Boschi 2002; Auer *et al.* 2014). More objective and quantitative model assessment can be achieved in the data space, that is, by propagating waves in those 3-D tomographic models and comparing the synthetics to seismic data.

Wave propagation also bears the next-generation full-waveform-based inverse techniques, such as adjoint tomography (e.g. Tape *et al.* 2009; Zhu *et al.* 2015; Bozdağ *et al.* 2016a), which iteratively refines a 3-D earth model based on 3-D sensitivity kernels, and forward modelling of data (e.g. Cottaar & Romanowicz 2012; Thorne *et al.* 2013), which searches over vast possibilities for optimized model parameters. Compared to a conventional linearized tomography, both iterative gradient (Tromp *et al.* 2008) and non-gradient (chapter 2, Tarantola 2005) approaches can access the regime of strong wave scattering, desirable or even necessary for strong and sharp heterogeneities, for instance, crustal tomography using noise correlation (Basini *et al.* 2013) or earthquake waves (Tape *et al.* 2009) as well as imaging localized small-scale structures such as the ULVZs (Cottaar & Romanowicz 2012; Thorne *et al.* 2013) and mantle scatterers (Korenaga 2015; Haugland *et al.* 2017). In both these nonlinear approaches, however, one may need to propagate waves in different 3-D earth models hundreds to thousands of times for each event.

Many numerical methods have been developed for seismic wave propagation in 3-D earth media (Moczo *et al.* 2011; Igel 2017), such as finite difference (Igel *et al.* 2002; Kristek & Moczo 2003), spectral element (Komatitsch & Tromp 2002a; Capdeville *et al.* 2003; Chaljub *et al.* 2003) and discontinuous Galerkin (Dumbser & Käser 2006; Etienne *et al.* 2010). Global seismology has benefited from a couple of community software packages, including SPECFEM3D_GLOBE (SPECFEM, Komatitsch & Tromp 2002a,b), the most mature and comprehensive spectral element method (SEM) hardcoded for the Earth, encompassing 3-D mantle and crust, topography and ocean, ellipticity, rotation and gravity and Salvus (Afanasiev *et al.* 2017), 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, and CSEM (Capdeville *et al.* 2003) and SES3D (Gokhberg & Fichtner 2016) that tackle 3-D media in spherical geometry, aimed mainly at mantle studies. Discretized 3-D methods such as SPECFEM and Salvus are capable of handling complex 3-D structures in arbitrary geometry and may thus be considered 'full' and 'exact', as long as the earth model can be properly discretized.

However, despite the surge in both hardware and software developments, global wave propagation has remained as one of the most challenging problems in scientific computing. Discretized 3-D methods come at a persistently high computational cost (Carrington *et al.* 2008; Rietmann *et al.* 2012; Tsuboi *et al.* 2016) that still obliterates any possibility to cover the observable frequency band of global seismic data (up to 1 Hz) with realistic computing resources. The multiscale nature of the Earth's 3-D interior requires high-frequency simulations. For example, to resolve a mantle plume a few hundred kilometres in width, a seismic period down to 5 s should be desirable, at which one 3-D simulation may consume $10^4 \sim 10^5$ CPU-hr for a 1 hr record length. At a 1 Hz frequency, it may consume a formidable amount of $10^7 \sim 10^8$ CPU-hr per 1 hr record length. Concerning seismic inverse problems that involve a large number of forward simulations, these 3-D methods can hardly serve as the underlying wave propagator for realistic problems with a high-frequency resolution and a good amount of data. Computational cost may be reduced by hardware accelerators such as GPU (Komatitsch *et al.* 2010; Rietmann *et al.* 2012), algorithmic advances such as local time stepping (Rietmann *et al.* 2015) and coarse-grained attenuation (van Driel & Nissen-Meyer 2014), or a truncated computational domain such as RegSEM (Cupillard *et al.* 2012). Practically, most of these efforts may lead to a computational speed-up of 2 to 5, desirable but insufficient to substantially ease the computability.

As justified by the fact that a large part of traveltime data can be predicted by an averaged 1-D Earth structure within a small margin of error (Astiz *et al.* 1996), spherically symmetric earth models come as a pragmatic compromise for modelling between broad frequency ranges and reasonably realistic waveforms. Based on this presumption, many dimension-reduction methods have been developed. Capable of decreasing the computational cost by a few orders of magnitude, the dimension-reduction methods have taken a predominant role in global seismology, maximizing data utilization in terms of both data volume and broad-band waveform contents (e.g. Colombi *et al.* 2012; Hosseini & Sigloch 2015). Depending on how their computational cost scales with the number of output sampling points, these dimension-reduction methods can be divided into two categories: waveform methods that produce waveforms at given receiver locations, such as DSM (Kawai *et al.* 2006), GEMINI (Friederich & Dalkolmo 1995), Yspec (Al-Attar & Woodhouse 2008), Instaseis (van Driel *et al.* 2015) and Syngine (Krischer *et al.* 2017), and wavefield methods that deliver the whole wavefield throughout the interior, such as AxiSEM (Nissen-Meyer *et al.* 2007b, 2008, 2014), an axisymmetric SEM that carries Instaseis and Syngine, and axisymmetric finite difference methods (e.g. Igel & Weber 1995; Chaljub & Tarantola 1997; Toyokuni & Takenaka 2006). In addition, a dimension-reduction method can be coupled with a full 3-D method to efficiently solve multiscale problems with a 3-D subdomain inside a 1-D background model (e.g. Capdeville *et al.* 2003; Masson *et al.* 2014; Monteiller *et al.* 2015).

Motivated to bridge the huge gap between expensive methods for 3-D media and fast methods for spherically symmetric media, Leng *et al.* (2016) brought forward a novel hybrid method combining SEM and pseudospectral method. It parametrizes 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 the smoothness of global tomographic models. This method features a model-adapted computational cost: for spherically symmetric earth models, the method naturally degenerates to the axisymmetric SEM, AxiSEM (Nissen-Meyer *et al.* 2007b); while, for typical tomographic mantle models, it proves significantly more efficient than conventional 3-D methods, for example, with a speed-up of two orders of magnitude near a 10 s period. More attractively, such a speed-up is observed to scale with seismic frequency, attractive to high-frequency applications (e.g. below 5 s) that harbour numerous open geophysical questions.

This paper strives to extend the prototype of Leng *et al.* (2016) to become as comprehensive as discretized 3-D methods to include all known types of 3-D Earth features, while being considerably more efficient than the 3-D methods across a large model parameter space and a broad frequency range. From a model perspective, we will generalize the method to aspherical geometry and verify its robustness to localized small-scale structures. From a computational perspective, we will introduce a couple of techniques that enhance its usability and efficiency. We will also carry out comprehensive verification in reference to a discretized 3-D method.

First, we will extend our method to aspherical earth models with undulating interfaces. The 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 (e.g. Basini *et al.* 2013; Laske *et al.* 2013) and the transition zone boundaries (e.g. Deuss *et al.* 2006; Kustowski *et al.* 2008) and some of the solid–fluid discontinuities such as the ocean floor (e.g. Gualtieri *et al.* 2014) and CMB (e.g. Soldati *et al.* 2013; Colombi *et al.* 2014). These undulating discontinuities must be physically honoured in certain applications because of their striking impact on the synthetic waveforms, for instance, surface topography and crustal thickness variation that entirely reshape surface waves (Zhou *et al.* 2005; Tape *et al.* 2009). Imaging interface perturbations can also deliver important complementary constraints on Earth structure and dynamics (Deuss *et al.* 2006; Colombi *et al.* 2014). From a methodological aspect, global wave propagation with undulating interfaces has not been cross-validated so far, especially for the 3-D crust that poses significant challenges to model discretization.

Even for a discretized 3-D method, handling geometric complexity is not a straightforward task. To overcome the outstanding difficulties around discretization, two main strategies have been employed. First, advanced meshing techniques can be used to mesh the geometry directly, such as unstructured hexahedra (Peter *et al.* 2011) and tetrahedra (Pelties *et al.* 2012). These meshing techniques are capable of tackling arbitrary 3-D geometry, but are not available to pseudospectral schemes without a 3-D mesh, such as normal-mode summation (chapter 13, Dahlen & Tromp 1998), finite difference approaches with staggered grids (chapter 4, Igel 2017) and our hybrid method. For these approaches, certain types of geometric complexities can still be considered exactly via several medium transformations, such as the non-periodic homogenization technique (Capdeville *et al.* 2010, 2013), curvilinear coordinate transformations (Hestholm & Ruud 1994; Komatitsch *et al.* 1996; Zhang & Chen 2006; Tarrass *et al.* 2011; Zhang *et al.* 2012) and the particle relabelling transformation (Takeuchi 2005; Al-Attar & Crawford 2016) used in this paper. Starting from different physical or mathematical considerations, such medium transformations convert the geometric complexities into volumetric complexities of the contiguous media while preserving the resultant wavefields. At a continental to global scale, it is reasonable to assume that

the radial discontinuities (except the ocean floor) are topologically equivalent to a sphere. Under this assumption, we base our aspherical extension on the particle relabelling transformation (Al-Attar & Crawford 2016). It establishes a new 3-D wave equation on an imaginary spherically symmetric earth model; this wave equation is formally identical to the original one for the physically undulating earth model and yields an equivalent wavefield solution. Such a transformation is frequency-independent and easy to be implemented.

Apart from long-wavelength mantle heterogeneities and undulating interfaces, localized small-scale heterogeneities represent another type of known 3-D Earth structure, such as the ULVZs (Cottaar & Romanowicz 2012; Thorne *et al.* 2013) and mid-mantle scatterers (Korenaga 2015; Jenkins *et al.* 2017). Depending on the structure–wavelength ratio and the perturbation strength, such small-scale heterogeneities may enter the regime of backscattering, reflecting energy to all directions like a secondary source. Such scattering effects can be simulated by discretized 3-D methods, almost with an indifferent computational cost. However, they are generally believed to be a bottleneck of pseudospectral schemes due to their white power spectra (chapter 9, Boyd 2001). Being convergent in theory, our method is capable of incorporating such small-scale heterogeneities, but their cost impact has to be carefully examined. For this purpose, we impose a strong mantle scatterer onto PREM, considering a wide range of structure–wavelength combinations.

Besides the above two model-driven extensions, three methodological improvements will be introduced. The first is called 'wavefield scanning', which maximizes the computational efficiency of our method by locally optimizing the Fourier expansion order of solution. It also provides a natural way to study azimuthal wave scattering and healing. Second, we will introduce an improved Fourier-mode coupling approach that greatly promotes the performance for non-smooth models. Finally, the parallel scalability of our message passing interface (MPI)-based implementation up to more than 12 000 cores will be presented, based on a heterogeneous domain decomposition scheme for load-balancing. Efficiency comparisons will be carried out between our method and a discretized 3-D SEM, SPECFEM, considering a large parameter space for both realistic Earth structures and seismic frequencies.

2 THEORY

For spherically shaped earth models with 3-D material heterogeneities, the theory and the SEM implementation of our hybrid method have been elaborated in Leng *et al.* (2016). In this section, we focus on the 'particle relabelling transformation' (Al-Attar & Crawford 2016) for aspherical Earth geometry. A new approach for Fourier-mode coupling will also be introduced.

2.1 3-D earth with spherical geometry

We first briefly review the theory for 3-D earth models with spherical geometry, established in Leng *et al.* (2016). 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 $\partial\Omega$ by $\hat{\mathbf{n}}$. Let $\mathbf{u}(\mathbf{r};t)$ and $\mathbf{f}(\mathbf{r};t)$, respectively, denote the time-dependent displacement and body force, $\rho(\mathbf{r})$ the density, and $\mathbf{C}(\mathbf{r})$ the fourth-order elastic stiffness tensor.

We select a coordinate system in which the source lies on or beneath the north pole. Such a source-centred coordinate system is framed by a unit cylindrical basis denoted by $\left(\hat{\mathbf{s}}, \hat{\boldsymbol{\phi}}, \hat{\mathbf{z}}\right)$ or
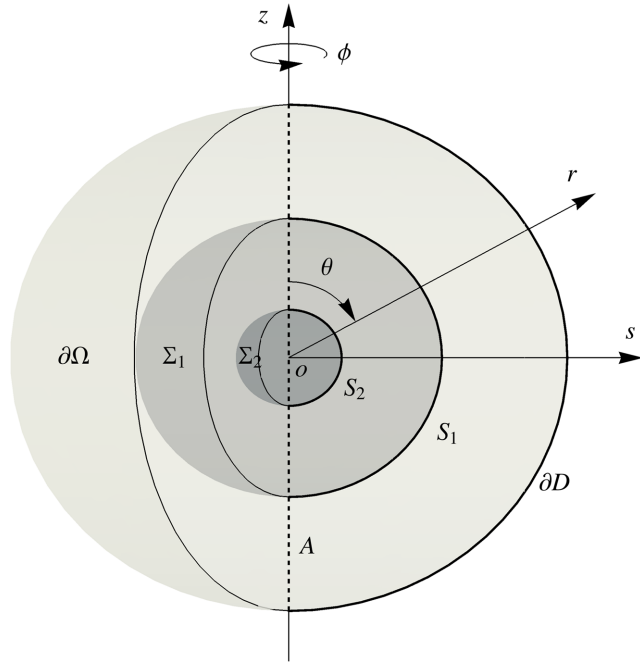
Figure 1. Sketch of an earth model with axisymmetric geometry. The right-hand side shows the 2-D meridian domain D with physical boundaries $\partial D \cup S_i$ (thick curves) and axial boundary A (dashed thick line). Rotating D about A by 2π will generate the 3-D domain Ω, with boundaries $\partial\Omega \cup \Sigma_i$, as shown on the left-hand side. Here we denote the 3-D solid–fluid discontinuities and their 2-D meridians by Σ_i and S_i, respectively, which are not included in this section but later referred to in Appendix A.

$(\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2, \hat{\mathbf{g}}_3)$. Because of its spherically symmetric geometry, including the surface and all the internal discontinuities, the 3-D domain Ω can be generated by rotating a 2-D meridian domain D ('D'-shaped) about the axis A by an angle of 2π, as shown in Fig. 1.

We start from the following weak form of the 3-D equations of motion,

$$\int_\Omega \rho \partial_t^2 \mathbf{u} \cdot \mathbf{w} \, d^3\mathbf{r} + \int_\Omega \nabla\mathbf{u} : \mathbf{C} : \nabla\mathbf{w} \, d^3\mathbf{r} = \int_\Omega \mathbf{f} \cdot \mathbf{w} \, d^3\mathbf{r}, \tag{1}$$

where $\mathbf{w} = \mathbf{w}(\mathbf{r})$ denotes an arbitrary test function in Ω.

As observed by Leng *et al.* (2016), global wavefields in a typical tomographic mantle model take on different variabilities along different directions, that is, in the source-centred coordinate system, they vary much more smoothly along the azimuthal direction than within the meridian plane. This results from the fact that the lateral wave scattering effects of a smooth Earth structure are too weak to substantially change the source-dominant radiation patterns. To take advantage of such inherent azimuthal smoothness, Leng *et al.* (2016) introduced a global Fourier series parametrization of functions over the ϕ-direction. For example, the displacement field $\mathbf{u}$ will be approximated by the following truncated Fourier series,[1]

$$\mathbf{u} = \hat{\mathbf{g}}_i(\phi) u_i(s, \phi, z; t) \simeq \hat{\mathbf{g}}_i(\phi) \sum_{|\alpha| \leq n_u} u_i^\alpha(s, z; t) e^{\hat{\imath}\alpha\phi}, \tag{2}$$

where $\hat{\imath} = \sqrt{-1}$ and $u_i^\alpha(s, z; t)$, a 2-D field defined in D, denotes the α-th order Fourier coefficient of $u_i(s, \phi, z; t)$. Doing the same for

the test function $\mathbf{w}$, the source term $\mathbf{f}$ and the material properties ρ and $\mathbf{C}$, we may reduce the 3-D weak form in Ω, eq. (1), to a coupled algebraic system of 2-D weak forms in D, with the ϕ-dimension reduced (see Leng *et al.* (2016) for derivations),

$$\sum_{|\alpha| \leq n_u} \left(\int_D \rho^{-(\alpha+\beta)} \partial_t^2 u_i^\alpha w_i^\beta s \, ds\,dz + \int_D u_{i;j}^\alpha C_{ijkl}^{-(\alpha+\beta)} w_{k;l}^\beta s \, ds\,dz \right)$$
$$= \int_D f_i^{-\beta} w_i^\beta s \, ds\,dz, \quad \forall \beta \in \{-n_u, \ldots, n_u\}, \tag{3}$$

where $u_{i;j}^\alpha$ denote the Fourier coefficients of the deformation gradient $\nabla\mathbf{u}$ computed by

$$u_{j;i}^\alpha = \partial_{x_i} u_j^\alpha + \frac{1}{s} \Gamma_{ijk}^\alpha u_k^\alpha, \quad \Gamma_{ijk}^\alpha = \delta_{i2} \left(\hat{\imath}\alpha\delta_{jk} - \epsilon_{3jk} \right), \tag{4}$$

with δ_{ij} being the Kronecker delta and ϵ_{ijk} the Levi-Civita symbol. For an earthquake source represented by a point moment tensor, the right-hand side of eq. (3) has a closed form (see Appendix B), which vanishes when $|\beta| > 2$; for a point force vector such as a meteorite impact or an adjoint source for waveform inversion, it vanishes when $|\beta| > 1$.

Since the unknowns (u_i^α) and the parameters (ρ^γ, C_{ijkl}^γ and f_i^β) of eq. (3) are all defined in the 2-D meridian domain D, it can be solved with a 2-D numerical method. Leng *et al.* (2016) implemented a 2-D SEM to solve this problem, named AxiSEM3D, as it is a natural extension of AxiSEM (Nissen-Meyer *et al.* 2007b).

## 2.2 Particle relabelling transformation

The Fourier series parametrization, eq. (2), relies on 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 (Al-Attar & Crawford 2016).

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. 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, so as to retain the form of eq. (3) as well as its SEM implementation. The idea behind such transformations comes from the work of Woodhouse (1976) on boundary perturbation theory for seismic free oscillations. This approach was then extended to exact wavefield calculations by Takeuchi (2005) in a Cartesian geometry and by Al-Attar & Crawford (2016) who established a general theoretical framework for the transformations.

Frequency-independence is a notable advantage of the particle relabelling transformation: it does not interact with wave propagation effects but only depends on the given physical model and the selected geometric configuration of the reference model, allowing for an easy implementation at the pre-processing stage. However, because the transformation requires compatibility or homeomorphism between the physical and the reference configuration, we cannot utilize it to handle discontinuities non-homeomorphic to a sphere, such as the patched (multiply connected) ocean floor. In AxiSEM3D, we model the ocean as a hydrodynamic load (Komatitsch & Tromp 2002b), with the formulations given in Appendix C.

---

[1] In this paper, subscripts of field variables denote spatial indices, such as *ijkl*, and superscripts Fourier series indices (not powers), such as $\alpha\beta\gamma$. Einstein summation convention only applies to spatial indices.
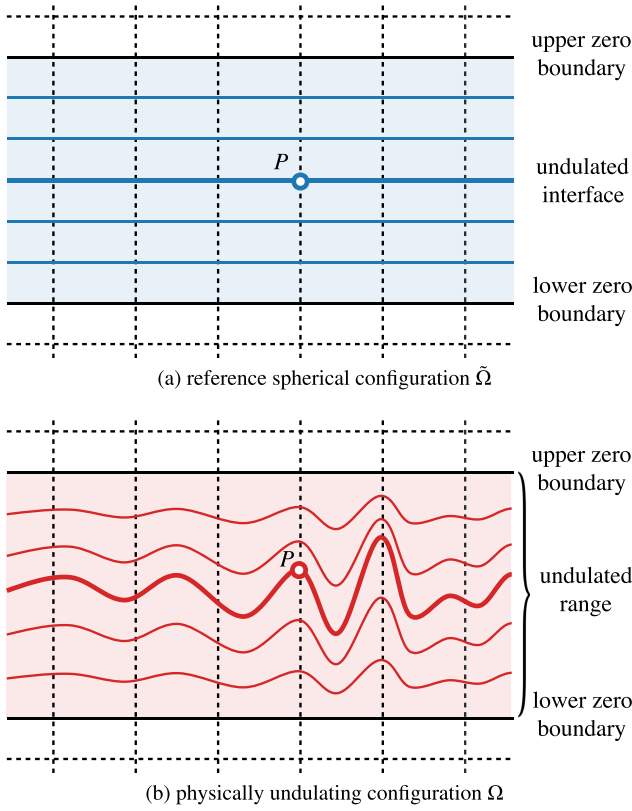
Figure 2. An illustration of geometric mapping from a selected reference spherical configuration to the physically undulating configuration. With the geometric mapping constrained in the radial direction, a particle P at position $\mathbf{r}$ in the reference configuration is stretched to position $\mathbf{r} + \tau(\mathbf{r})\hat{\mathbf{r}}$ 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 at the cost of more squeezed elements that may diminish the global time step.

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. (1). 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 minimize the computational cost related to the transformation, we constrain the geometric mapping in the radial direction $\hat{\mathbf{r}}$, quantified by the following relation,

$$\xi(\mathbf{r}) = \mathbf{r} + \tau(\mathbf{r})\hat{\mathbf{r}}, \tag{5}$$

where $\tau(\mathbf{r})$ gives the radial shift at position $\mathbf{r}$ in $\tilde{\Omega}$. Literally, eq. (5) means, through the geometric mapping $\xi : \tilde{\Omega} \to \Omega$, 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. 2.

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 system $(r, \theta,$

$\phi)$ in Fig. 1 as

$$\mathbf{F}(\tau, \mathbf{r}) = \left[\nabla(\tau\hat{\mathbf{r}})\right]^{\mathrm{T}} = \begin{pmatrix} \hat{\mathbf{r}} & \hat{\boldsymbol{\theta}} & \hat{\boldsymbol{\phi}} \end{pmatrix} \begin{pmatrix} \partial_r \tau & \dfrac{\partial_\theta \tau}{r} & \dfrac{\partial_\phi \tau}{r\sin\theta} \\ 0 & \dfrac{\tau}{r} & 0 \\ 0 & 0 & \dfrac{\tau}{r} \end{pmatrix} \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\theta}} \\ \hat{\boldsymbol{\phi}} \end{pmatrix}. \tag{6}$$

According to eq. (152) of Al-Attar & Crawford (2016), the 3-D weak form for the physical model, eq. (1), proves equivalent to the following weak form established on the reference configuration $\tilde{\Omega}$,

$$\int_{\tilde{\Omega}} \tilde{\rho} \partial_t^2 \tilde{\mathbf{u}} \cdot \tilde{\mathbf{w}} \, \mathrm{d}^3\mathbf{r} + \int_{\tilde{\Omega}} \nabla\tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla\tilde{\mathbf{w}} \, \mathrm{d}^3\mathbf{r} = \int_{\tilde{\Omega}} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}} \, \mathrm{d}^3\mathbf{r}, \tag{7}$$

in the sense that the solution of eq. (7), denoted $\tilde{\mathbf{u}}$, is identical to that of eq. (1):

$$\tilde{\mathbf{u}}(\mathbf{r};t) = \mathbf{u}(\xi(\mathbf{r});t). \tag{8}$$

The parameters in eq. (7) are, respectively, determined by

$$\tilde{\rho}(\mathbf{r}) = \rho(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}) \tag{9}$$

for density, and

$$\tilde{\mathbf{C}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{C}(\xi(\mathbf{r})) \cdot \mathbf{J}^{-\mathrm{T}}(\tau, \mathbf{r}) \det \mathbf{J}(\tau, \mathbf{r}) \tag{10}$$

for elastic stiffness tensor, and

$$\tilde{\mathbf{f}}(\mathbf{r}) = \mathbf{f}(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}) \tag{11}$$

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

$$\mathbf{J}(\tau, \mathbf{r}) = \mathbf{I} + \mathbf{F}(\tau, \mathbf{r}), \tag{12}$$

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

$$\mathbf{J}^{-1}(\tau, \mathbf{r}) = \frac{1}{h} \begin{pmatrix} \hat{\mathbf{r}} & \hat{\boldsymbol{\theta}} & \hat{\boldsymbol{\phi}} \end{pmatrix} \begin{pmatrix} 1 + \dfrac{\tau}{r} & -\dfrac{\partial_\theta \tau}{r} & -\dfrac{\partial_\phi \tau}{r\sin\theta} \\ 0 & 1 + \partial_r \tau & 0 \\ 0 & 0 & 1 + \partial_r \tau \end{pmatrix} \begin{pmatrix} \hat{\mathbf{r}} \\ \hat{\boldsymbol{\theta}} \\ \hat{\boldsymbol{\phi}} \end{pmatrix}, \tag{13}$$

where $h = (r + \tau)(1 + \partial_r \tau)/r$. Note that $\mathbf{J}^{-1}$ shares the same sparsity with $\mathbf{F}$, which can speed-up the computation of the stiffness term. This is the reason why we constrain the geometric mapping, eq. (5), in the radial direction.

The generalized weak form, eq. (7), is formally identical to the original one, eq. (1), except for the parameters. Thus, we may reduce it to an algebraic system of 2-D weak forms following the same approach in Section 2.1, and the result should be formally identical to eq. (3), but with the Fourier coefficients replaced by those of the transformed parameters ($\tilde{\rho}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{f}}$). However, additional attention must be paid to the solid–fluid discontinuities, as addressed in Appendix A. The general theory of Al-Attar & Crawford (2016) has been extended to solid–fluid bodies by Al-Attar *et al.* (2018). Also, the closed form of the source terms must be modified, as given in Appendix B.

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. (10). Because the major symmetry ($\tilde{C}_{ijkl} = \tilde{C}_{klij}$) still holds, $\tilde{C}_{ijkl}$ contains 45 ostensibly independent elements, in contrast to five for transverse isotropy and two for isotropy, the two most

frequently used types in seismology. Therefore, the double contractions in eq. (3), $\nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}}$, will become very expensive if implemented as is. In practice, we do not compute $\tilde{\mathbf{C}}$ by eq. (10). 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 double contractions $\nabla \tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla \tilde{\mathbf{w}}$ by means of $\left(\mathbf{J}^{-\mathrm{T}} \cdot \nabla \tilde{\mathbf{u}}\right) : \mathbf{C} : \left(\mathbf{J}^{-\mathrm{T}} \cdot \nabla \tilde{\mathbf{w}}\right) \det \mathbf{J}$. The implementation will be elaborated at the end of the following subsection.

The particle relabelling transformation can also be understood in analogy to the piecewise transformation used in finite-element-based approaches (including SEM), that is, the irregular physical domain of an element is mapped onto a regular reference domain, such as a square in 2-D (e.g. Leng *et al.* 2016) and a cube in 3-D (e.g. Komatitsch & Tromp 2002a). The mathematics underlying the two transformations are similar both applying a change of variables and the chain rule for the calculation of derivatives and quadratures. Whether one wishes to see these two transformations as being equivalent comes, in part, down to a point of view on 'discretize-then-transform' versus 'transform-then-discretize'. In finite element, the equations of motion are established in the physical configuration, while the transformation, which comes after discretization, serves just as a tool for calculating derivatives and quadratures. In the particle relabelling transformation, however, the equations of motion have been transformed onto the reference configuration before discretization, with their functional forms kept unchanged; a numerical method is then developed to solve these new equations of motion. This point of view is particularly sensible here because AxiSEM3D can only have a spherical computational domain due the Fourier parametrization of solution, which necessitates transformation before discretization.

2.3 FFT-based mode coupling

This subsection aims at accelerating the procedure of Fourier-mode coupling, the most time-consuming part in a 3-D simulation with AxiSEM3D.

As shown by Leng *et al.* (2016), the 2-D weak forms, eq. (3), can be recast into the following strong forms,

$$\sum_{\substack{|\alpha| \leq n_u \\ |\beta - \alpha| \leq n_c}} \mathscr{L}_{ik}^{\beta - \alpha} u_k^\alpha = f_i^\beta, \quad \forall \beta \in \{-n_u, \ldots, n_u\}, \tag{14}$$

where $\mathscr{L}_{ik}^{\beta - \alpha}$ denotes the reduced wave operator in D, determined by the $(\beta - \alpha)$-th order Fourier coefficients of ρ and $\mathbf{C}$,

$$\mathscr{L}_{ik}^{\beta - \alpha} u_k^\alpha = \rho^{\beta - \alpha} \delta_{ik} \partial_t^2 u_k^\alpha - \left(C_{ijkl}^{\beta - \alpha} u_{k;l}^\alpha \right)_{;j}, \tag{15}$$

and n_c the number of terms required in the Fourier series of either ρ or $\mathbf{C}$ for an accurate model description. For a given 3-D model, n_c is constant and equals to the maximum degree ($l_{\max}$) of the spherical harmonics expansion of the model, for example, $n_c = 40$ for the global mantle model S40RTS (Ritsema *et al.* 2011). Eq. (14) indicates how lateral heterogeneities couple the Fourier modes of wavefields: a 3-D ρ or $\mathbf{C}$ will give rise to a non-diagonal operator matrix $\mathscr{L}_{ik}^{\beta - \alpha}$ with a bandwidth of $2n_c + 1$.

Leng *et al.* (2016) implemented eq. (14) straightforwardly by a summation over α, which proved efficient for the state-of-the-art tomographic mantle models such as S40RTS ($n_c = 40$), S20RTS ($n_c = 20$, Ritsema & Van Heijst 1999) and S362ANI ($n_c = 22$, Kustowski *et al.* 2008), based on the fact that the non-diagonal operator matrix is narrowly banded by a small n_c. However, this direct summation approach will cease to work efficiently if more
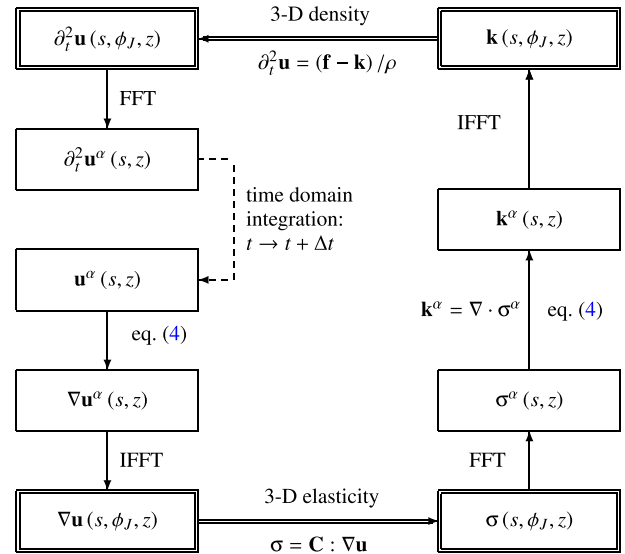
Figure 3. Workflow of one time step in AxiSEM3D. Two additional field variables are introduced: the stress tensor $\boldsymbol{\sigma}$ and the stiffness force $\mathbf{k}$ (the body force corresponding to $\boldsymbol{\sigma}$). FFT indicates the transform from physical to Fourier space and IFFT from Fourier to physical space. All 3-D operations are conducted in the physical space (double-lined) on the cardinal gridpoints anchored at $\phi_J = 2\pi J/(2n_u + 1)$, $J = 0, 1, \ldots, 2n_u$. For aspherical earth models, the parameters ρ, $\mathbf{C}$ and $\mathbf{f}$ will be replaced by $\tilde{\rho}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{f}}$, and the stress tensor $\boldsymbol{\sigma}$ will become asymmetric, $\sigma_{ij} \neq \sigma_{ji}$.

complicated structures are involved, such as the global crustal model Crust 1.0 ($n_c \leq 180$, Laske *et al.* 2013).

Alternatively, we may implement the summation by means of fast Fourier transform (FFT), commonplace in pseudospectral methods for partial differential equations with spatially varying parameters (chapter 9, Boyd 2001). 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 gridpoints evenly distributed along the ϕ-direction. FFT is responsible for efficient transitions between the Fourier space and the physical space. A typical time step is illustrated in Fig. 3.

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 $O(n_u \log n_u)$.[2] Therefore, the performance will be greatly promoted for complex models with a large n_c. In addition, because the cost no longer depends *explicitly* on the model, we may optimize the azimuthal resolution of solution by inspecting the Fourier power spectra of wavefields, as explained later in Section 4.2.

In the remainder of this subsection, we elaborate the implementation of the particle relabelling transformation in the FFT-based mode coupling scheme. In short, all operations related to 3-D model parameters, including the particle relabelling transformation, are conducted in the physical space (double-lined in Fig. 3) on a collection of 3-D gridpoints. These 3-D gridpoints, as termed the cardinal gridpoints (chapter 9, Boyd 2001), are generated by rotating the Gauss–Lobatto–Legendre (GLL) points in the 2-D spectral element

---

[2] In modern FFT libraries such as FFTW3 (Frigo & Johnson 2005), $O(N \log N)$ 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.

mesh along the azimuthal direction, anchored at $\phi_J = 2\pi J/(2n_u + 1)$, $J = 0, 1, \ldots, 2n_u$.

Before time marching, given the geometric mapping $\xi : \tilde{\Omega} \to \Omega$ in eq. (5), we pre-compute the Jacobian **J** and its inverse $\mathbf{J}^{-1}$, respectively, by eqs (12) and (13) on the 3-D cardinal gridpoints $(s_{\text{GLL}}, \phi_J, z_{\text{GLL}})$. The in-plane derivatives, $\partial_r \tau$ and $\partial_\theta \tau$, are computed on the 2-D spectral elements (eqs 66–67, Leng *et al.* 2016) at each ϕ_J. Note that **J** will become singular if $\partial_r \tau = -1$, which corresponds to the case where the thickness of a layer becomes zero after the radial stretching, for example, a point on the Moho stretched to the surface. The azimuthal derivative, $\partial_\phi \tau$, is computed by FFT with the following steps: given the series $\tau(\phi_J)$, $J = 0, 1, \ldots, 2n_u$ at each GLL-point $(s_{\text{GLL}}, z_{\text{GLL}})$, we compute its discrete Fourier transform τ^α, $\alpha = 0, 1, \ldots, n_u$ and then the series $\mathring{\imath}\alpha\tau^\alpha$, $\alpha = 0, 1, \ldots, n_u$, whose inverse discrete Fourier transform yields the azimuthal derivative $\partial_\phi \tau(\phi_J)$, $J = 0, 1, \ldots, 2n_u$.

During time marching, we consider the left-hand branch of Fig. 3 within an undulated range. The process is reciprocal for the right-hand branch. Starting from the displacement Fourier coefficients $\tilde{\mathbf{u}}^\alpha (s_{\text{GLL}}, z_{\text{GLL}})$ defined on the 2-D GLL points in the reference spherical configuration, we first compute the deformation gradient in the Fourier space $\nabla \tilde{\mathbf{u}}^\alpha (s_{\text{GLL}}, z_{\text{GLL}})$ by eq. (4), which is then transformed onto the 3-D cardinal points by IFFT, that is, $\nabla \tilde{\mathbf{u}} (s_{\text{GLL}}, \phi_J, z_{\text{GLL}})$. Then we are able to perform the contraction $\mathbf{C} : (\mathbf{J}^{-\text{T}} \cdot \nabla \tilde{\mathbf{u}})$ on each cardinal gridpoint $(s_{\text{GLL}}, \phi_J, z_{\text{GLL}})$, which yields the Cauchy stress tensor in the physically undulating configuration. In summary, the particle relabelling transformation does not interact with the Fourier series parametrization of solution, as it is conducted point-wise on the 3-D cardinal gridpoints.

3 VERIFICATION

In this section, we verify our approach for handling undulating interfaces by comparing AxiSEM3D to an independent 3-D SEM, SPECFEM3D_GLOBE (SPECFEM, Komatitsch & Tromp 2002a,b). SPECFEM and AxiSEM3D have been cross-verified for spherically shaped earth models with 3-D material perturbations (Leng *et al.* 2016).

We consider three representative aspherical Earth features in our 3-D earth model:

(i) the Earth's ellipticity, which undulates the surface and all the internal interfaces (including solid–fluid ones such as CMB);

(ii) surface topography and bathymetry, which introduce a larger undulation on the surface as well as a 3-D ocean;

(iii) the tomographic mantle model S362ANI, which incorporates topography on the transition zone discontinuities at $\sim$410 and $\sim$650 km.

Full crustal structure (such as Crust 1.0, Laske *et al.* 2013) has been implemented in both SPECFEM and AxiSEM3D, but a discrepancy persists because AxiSEM3D cannot follow the approximation SPECFEM has made to handle the undulating Moho. The main reason is that the undulation on the Moho exhibits sharp gradients beneath the ocean–continent plate boundaries, which causes great difficulty to model discretization while having a strong impact on computational cost and numerical stability. The details will be discussed in Section 6.1.

The simulation parameters used in this section are summarized in Table 1. The nominal simulation period is 5 s, following the SPECFEM convention of simulation periods. This convention corresponds to a one-element-per-wavelength resolution with

Table 1. Simulation parameters for cross-verification between SPECFEM and AxiSEM3D.

Parameter	Value	Notes
1-D reference	PREM	Anisotropic; viscoelastic
ellipticity	WGS 84	Both geographic–geocentric and full correction
Topography	Smoothed ETOPO1	Elevation varies between −7.7 and 5.5 km
Ocean	Smoothed ETOPO1	Approximated as load
Mantle	S362ANI	Anisotropic; $l_{\max} = 22$; undulation on 410 km varies between −13.4 and 13.2 km and that on 650 km between −14.3 and 19.1 km
Period	5 s	The average grid size on the surface is $\sim$2.8 km in SPECFEM and AxiSEM3D meshes
Record length	1 hr	Covering most teleseismic phases of interest
Earthquake	Virginia	Depth 12 km; 37.91° N, 77.93° W
stations	Fig. 4	17×18 imaginary stations
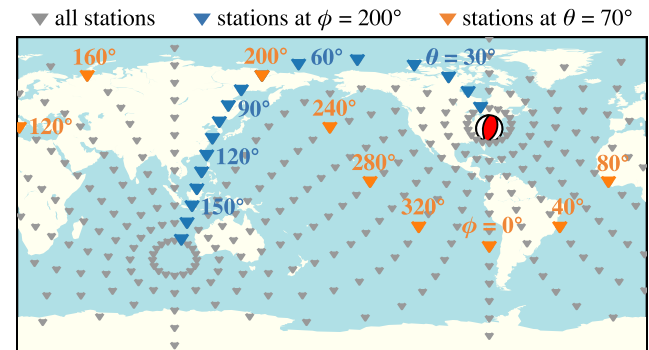
Figure 4. Source and stations for cross-verification between SPECFEM and AxiSEM3D. We compare synthetic seismograms recorded by 17×18 imaginary stations evenly distributed on the surface, spaced by 10° in epicentral distance (θ) and by 20° in azimuth (ϕ). The seismograms recorded by the stations at $\phi = 200°$ (blue) will be shown later in Figs 5 and 7, and those at $\theta = 70°$ (orange) in Fig. 6.

a polynomial order of 4. Therefore, the nominal simulation periods should be the minimum corner periods. In an attempt to benchmark SPECFEM and AxiSEM3D for spherical earth models, Leng *et al.* (2016) noted that such a convention could be insufficient for surface wave modelling. Hence, when comparing the synthetic seismograms, we adopt a bandpass filter with its minimum cut-off period fixed at 1.5 times of the nominal simulation period. We use an error function as the source–time function, with its half-duration equal to the simulation period.

We compare synthetic seismograms, respectively, computed by SPECFEM and AxiSEM3D on an imaginary network of stations, as shown in Fig. 4. Such global coverage on the surface should be able to reflect the goodness-of-fit of wavefields in the interior. Following Chaljub *et al.* (2015), we quantify the difference between two waveforms by the scores of goodness-of-fit, that is, for amplitude and phase, respectively,

$$G_e = 10 \exp(-|m_e|) \quad \text{and} \quad G_p = 10 \exp(-|m_p|), \qquad (16)$$

where m_e and m_p are the globally normalized time–frequency misfits (TF-misfits, Kristeková *et al.* 2009) for amplitude and phase, respectively. The scores of goodness-of-fit range from 0 to 10, with

10 indicating a perfect match. When the scores are above 9 (TF-misfits smaller than 0.105) for all the stations, we regard the match to be 'excellent'.

3.1 Ellipticity

The Earth's ellipticity represents an integrated test of the particle relabelling transformation, as the entire earth model has to be flattened from core to surface, undulating all the external and internal discontinuities. In practice, there are two levels of ellipticity correction, the geographic–geocentric (colatitude-θ) correction and the full ellipticity correction. In the former, the Earth geometry is kept spherical, but the locations of source and stations are corrected. In the latter, the entire model has to be flattened to satisfy the hydrostatic equilibrium condition governed by the Clairaut's equation (chapter 13, Nolet 2008), so the particle relabelling transformation will be activated throughout the entire computational domain.

Excellent agreement has been achieved between SPECFEM and AxiSEM3D. The global averages of G_e and G_p turn out to be 9.69 and 9.89, respectively, and the minimum local values 9.41 and 9.69, respectively, computed on the 1-hr full-length window. To visualize how small the misfits are, we compare the waveforms in Fig. 5, recorded by a meridional section where the maximum phase misfit occurs (blue in Fig. 4). From the zoomed-in windows in Fig. 5, one can observe the effects of the two levels of ellipticity correction. In brief, ellipticity mainly causes phase shifts to body waves, but alters surface wave waveforms to some degree. The geographic–geocentric correction can correctly predict the directions of the induced phase shifts but not their magnitudes.

3.2 Topography and bathymetry

The 3-D geometric perturbation introduced by ellipticity varies very smoothly in space, as reflected by the small phase shifts in Fig. 5. Therefore, we conduct a benchmark solution for topography and bathymetry. Sampled at a 5 s period, the undulation on the surface varies from about −7.7 km (Mariana Trench) to 5.5 km (Tibet), capable of generating obvious waveform anomalies on surface waves. The bottom of the low-velocity zone (depth = 220 km in PREM) is selected to be the 'lower zero boundary' in Fig. 2, at which the undulation vanishes. Between the surface and 220 km, the undulation decreases linearly with depth. In AxiSEM3D, any depth below −7.7 km (not just the radial discontinuities) can be chosen as the 'lower zero boundary'. Without observational constraints, the selection comes down to a performance issue, that is, a trade-off between a larger time step and a smaller number of stretched elements for the particle relabelling transformation. Here we use 220 km because this parameter is hardcoded in SPECFEM.

Now we discuss the spatial sampling of the surface. At a period of 5 s, the average grid size for surface sampling is ∼2.8 km in SPECFEM, that is, 3584 3-D elements (NEX = 896) on a great circle, each containing four gridpoints in one dimension. In AxiSEM3D, we create the 2-D mesh such that the surface sampling is equivalent in the meridian domain, that is, 1792 2-D elements on ∂D in Fig. 1. The azimuthal sampling becomes more complicated. The Fourier expansion order, n_u in eq. (2), should be sufficiently large to sample both the model and the resultant wavefield. Here we first consider model sampling, leaving wavefield sampling to Section 4.2. Among the three aspherical features we are considering, surface topography has the largest spatial variability, for which we use the 1-min gridded global relief data ETOPO1 (Amante &

Eakins 2009). If expanded in spherical harmonics, this model has a l_{max} of 5400 in theory. In this section, we use $n_u = 1200$, which means we downsample ETOPO1 in terms of truncated spherical harmonics. Note that the spectral power of ETOPO1 normalized to degree zero has decreased to ∼3×10^{-8} at $l = 1200$. In SPECFEM, we use one of its built-in topography models, which also originates from ETOPO1 but smoothed in some way. The two methods slightly differ in surface sampling, but, as we will show it later, such difference is insufficient to generate significant waveform differences at 5 s. The l_{max}'s of ellipticity and S362ANI are, respectively, 2 and 22, so they are accurately sampled in the computational models.

Ocean is another fundamental issue. Physical realization of a fluid ocean is essential to surface wave modelling, especially at the high frequencies where the load approximation becomes inaccurate (Komatitsch & Tromp 2002b). It also provides a means to study waves propagating within the ocean, such as the T-phases (Stevens *et al.* 2001), tsunami (Kowalik *et al.* 2005) and seismic noise excitation (Tanimoto 2010). The ocean floor is a multiply connected solid–fluid discontinuity, for which we can no longer rely on the particle relabelling transformation. Therefore, a 3-D fluid ocean is currently not available in AxiSEM3D. In this paper, we approximate the ocean as a hydrodynamic load on the ocean floor, following Komatitsch & Tromp (2002b). The formulations are given in Appendix C. The load approximation is accurate when the wavelengths of interest are larger than the ocean depth. It is suggested by Komatitsch & Tromp (2002b) that such an approximation should be acceptable when the period is above ∼5 s. This assumption, however, needs to be examined by numerically comparing a load-approximated ocean to a fluid ocean, considering the regional variation of ocean depth.

Similar to ellipticity, the agreement between SPECFEM and AxiSEM3D turns out excellent for topography and bathymetry. Here we compute the scores of goodness-of-fit on a surface wave window between the arrival times of '4.1 kmps' and '2.5 kmps', as shown in Fig. 6, because the body waves are barely affected. The global averages of G_e and G_p are 9.53 and 9.70, respectively, and the minimum local values 9.26 and 9.46, respectively. Such goodness-of-fit is lower than that for ellipticity because of surface sampling. Fig. 6 shows the seismograms recorded by the stations located 70° in epicentral distance. Because topography irregularly alters the surface waves, we emphasize the impact of ocean in Fig. 6: the ocean greatly strengthens surface wave dispersion and causes remarkable phase delays over the oceanic paths, for example, nearly a 10 min delay to the high-frequency surface waves at the azimuth of 80°.

3.3 Undulating transition zone

Here we consider the global tomographic mantle model S362ANI, which incorporates topography on the transition zone discontinuities at 410 and 650 km. The undulation ranges between −13.4 and 13.2 km on 410 km and between −14.3 and 19.1 km on 650 km. As constrained by the SPECFEM implementation, we choose the PREM discontinuities 220 and 771 km as the upper and lower zero boundaries on which the undulation vanishes. The undulation varies linearly between two adjacent controlling depths. Besides a comparison to SPECFEM, we intend to verify the uniqueness of solution upon different selections of the spherical reference configuration for the particle relabelling transformation. For this purpose, we create two different 2-D meshes in AxiSEM3D both based on PREM. The only difference between the two meshes is the depth of the two transition zone discontinuities: in the first one, we use 400 and 670 km
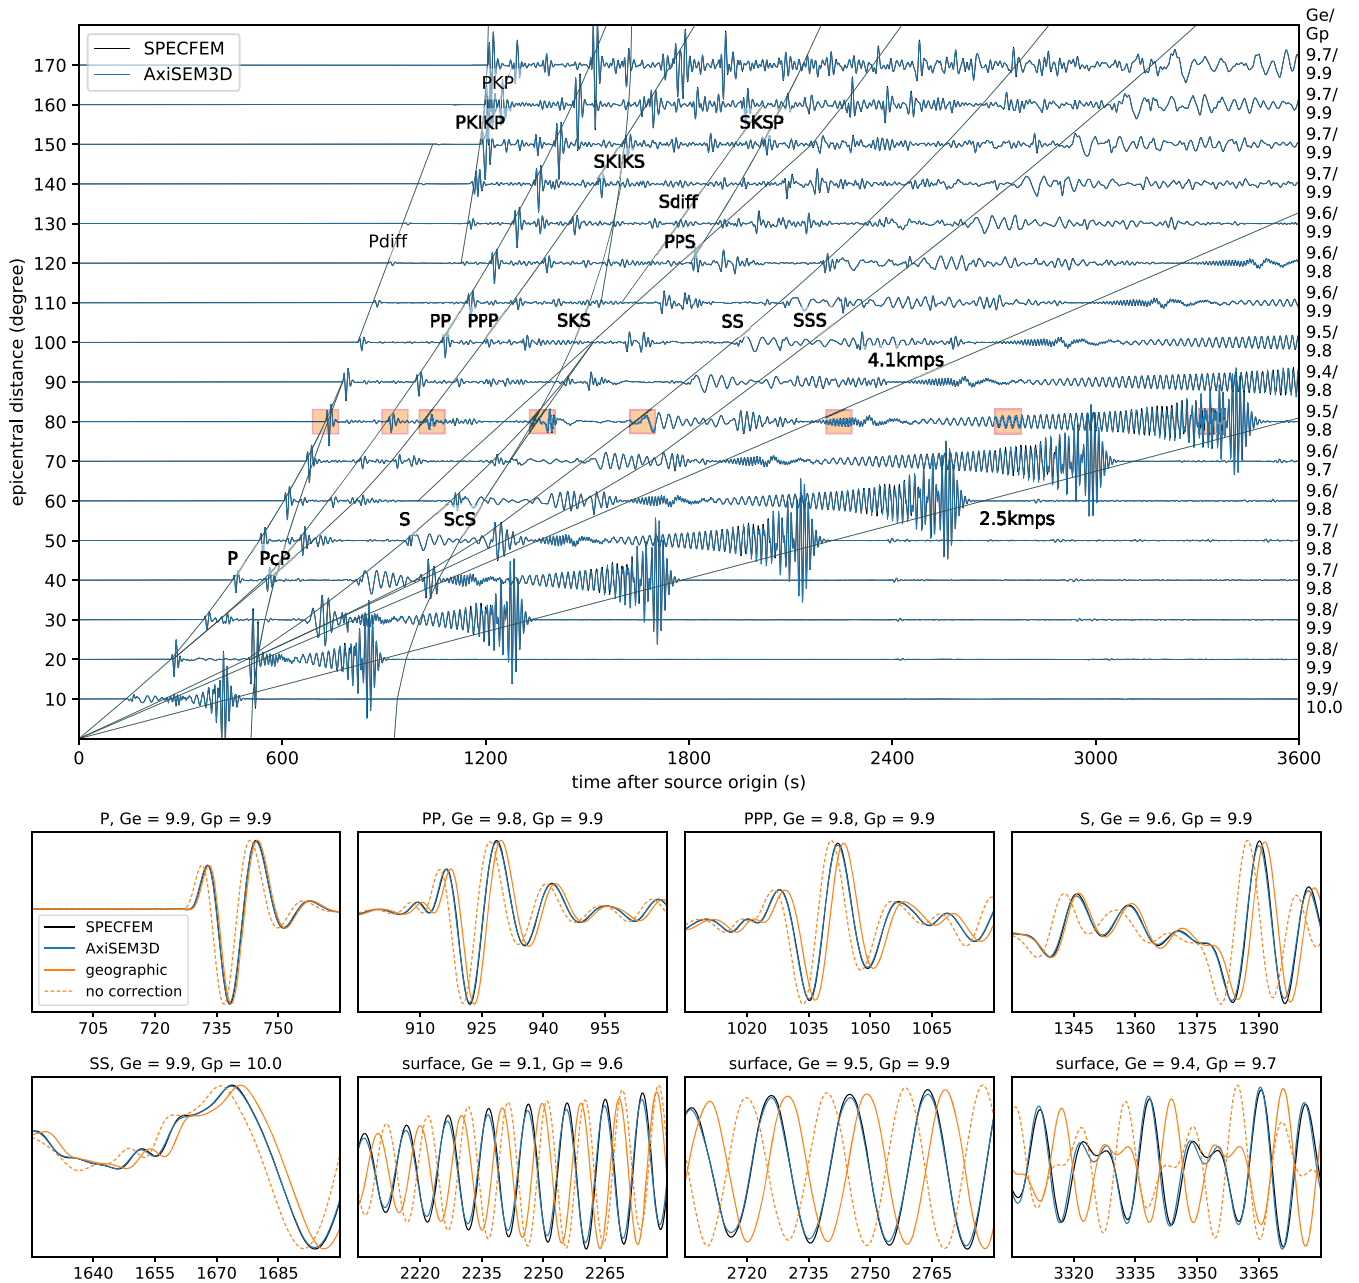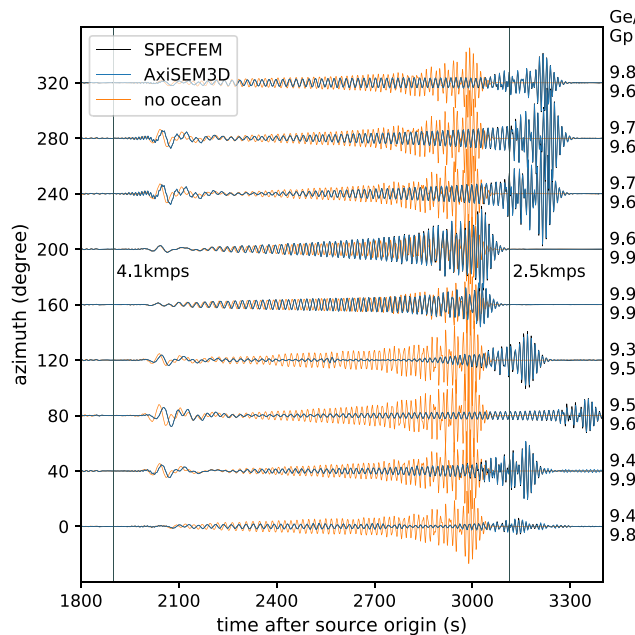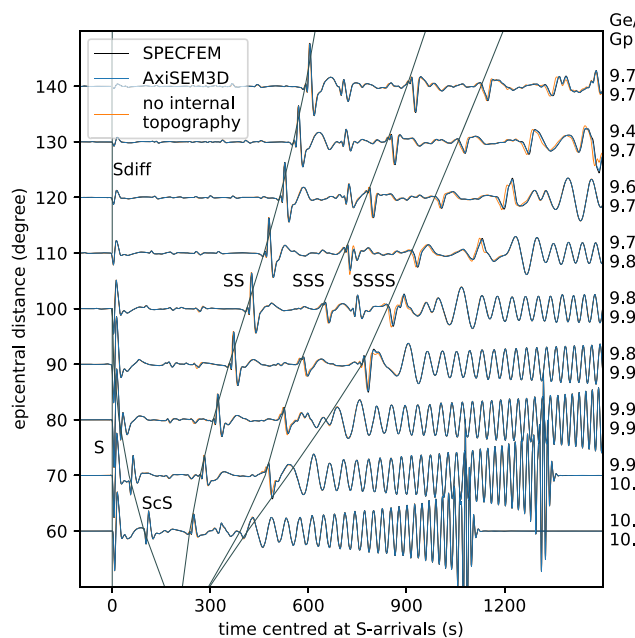
Figure 5. Comparisons of synthetic seismograms (vertical component) for ellipticity. The simulation parameters are summarized in Table 1. The stations are located on the meridian $\phi = 200°$ (blue in Fig. 4). The amplitudes have been scaled with epicentral distance to identify most teleseismic phases. To make the body waves more visible, we also reduce the amplitude of the surface waves by a factor of 5. The SPECFEM traces (black) are mostly invisible because of small waveform differences. To show the effects of ellipticity corrections, we zoom in a few windows from the 80° trace, adding two other traces computed on a spherical geometry, one with geographic–geocentric correction (solid orange) and the other without (dashed orange). The traveltime curves are computed for PREM using OBSPY.TAUP (Krischer *et al.* 2015).

based on PREM and, in the second one, 410 and 650 km based on STW105 (Kustowski *et al.* 2008), the original 1-D reference model of S362ANI. The undulations on these two discontinuities (τ in eq. 5) are determined such that their physical positions remain the same in the two computational models. In other words, the two models created from the two different meshes present the same physical earth model, so the solutions should also be identical. Based on the scaling used in the inversion (Kustowski *et al.* 2008), we use 55 per cent of dv_{sh} and dv_{sv}, respectively, for dv_{ph} and dv_{pv}.

Between SPECFEM and AxiSEM3D, the global averages of G_e and G_p are 9.64 and 9.80, respectively, and the minimum local values 9.36 and 9.65, respectively, computed on a half-hour window after the *S*-arrival where the impact of the undulating transition zone is most outstanding. Fig. 7 shows the seismograms for *SH* waves, emphasizing the time delays over the multiple *S* arrivals caused by the undulating transition zone. Compared to ellipticity and topography, the goodness-of-fit for S362ANI lies in between. The goodness-of-fit for the two equivalent AxiSEM3D models is almost perfect, with the minimum G_e and G_p being 9.97 and 9.99.

Figure 6. Comparisons of synthetic seismograms (radial component) for topography and bathymetry. The simulation parameters are summarized in Table 1. The stations are located at the distance $\theta = 70°$ (orange in Fig. 4). Here we focus on the surface wave window because the body waves are much less affected by topography. The traces labelled 'SPECFEM' (black) and 'AxiSEM3D' (blue) are computed with both topography and ocean incorporated, between which the appended G_e and G_p are calculated. The SPECFEM traces (black) are mostly invisible because of small waveform differences. The traces labelled 'no ocean' (orange) exclude the ocean, computed by SPECFEM (similar by AxiSEM3D).

Figure 7. Comparisons of synthetic seismograms (transverse component) for S362ANI with undulating transition zone boundaries. The simulation parameters are summarized in Table 1. The traces labelled 'SPECFEM' (black) and 'AxiSEM3D' (blue) are computed with topography on 410 and 650 km, between which the appended G_e and G_p are calculated. The SPECFEM traces (black) are mostly invisible because of small waveform differences. The traces labelled 'no internal topography' (orange) exclude transition zone topography but preserve the material heterogeneity of S362ANI. The stations and the amplitude scaling are the same as those used in Fig. 5, except that we align the seismograms by the S arrivals.

4 COMPUTATIONAL EFFICIENCY

In this section, we discuss the computational efficiency of AxiSEM3D. We first summarize its model-adapted cost characteristic, as realized by the spatial variability of the Fourier expansion order of solution, that is, $n_u = n_u(s, z)$ in eq. (2). Next, we introduce the technique of 'wavefield scanning', which can determine an efficient $n_u(s, z)$ for an arbitrary 3-D earth model. Based on this technique, we compare the computational efficiency of SPECFEM and AxiSEM3D, considering a variety of 3-D Earth features and the observable frequency band of global seismic data. Finally, we investigate the scattering effects and cost impact of localized small-scale heterogeneities and show a simulation at a 1 Hz corner frequency.

4.1 Model-adapted computational cost

A model-dependent cost characteristic distinguishes AxiSEM3D from discretized 3-D methods. Generally, the stronger the 3-D structure, the more expensive the simulation becomes. Such a structure-adapted cost characteristic is embodied by the Fourier expansion order of solution, or n_u in eq. (2), which can be adapted *locally* to the azimuthal complexity of wavefields in space and time, that is, $n_u = n_u(s, z; t)$.

For the simplest case of PREM-like 1-D earth models, the wavefield due to an earthquake takes on order two throughout space and time, that is, $n_u(s, z; t) = 2$ (Nissen-Meyer *et al.* 2007a; Leng *et al.* 2016), for which AxiSEM3D performs as efficient as AxiSEM (Nissen-Meyer *et al.* 2014). Let ω denote the dominant frequency. At this 1-D limit, AxiSEM3D has a computational cost of $O(\omega^3)$, that is, the two in-plane dimensions (s, z) and one time dimension t, whereas a discretized 3-D method has a computational cost of $O(\omega^4)$ (three spatial and one time dimension). Therefore, the speed-up of AxiSEM3D scales with frequency for 1-D models.

For a 3-D earth model, $n_u(s, z; t)$ varies with both space and time. Because changing n_u with time is technically impractical and uneconomic (as it requires redistribution of elements across processors), we only consider the spatial dependence of n_u, that is, $n_u = n_u(s, z)$, which should cover the maximum order required throughout time. The theoretical upper limit of $n_u(s, z)$ occurs when the minimum wavelength can be resolved over the azimuth dimension. In other words, $n_u(s, z)$ is determined such that the azimuthal spacing between the cardinal gridpoints is equivalent to the in-plane spacing between the GLL points in the 2-D mesh. Similar to a 3-D mesh, such a $n_u(s, z)$ ignores the azimuthal wavefield complexity and realizes a uniform wavefield sampling over all the three dimensions. At this upper limit, the computational cost of AxiSEM3D will be higher than that of a discretized 3-D method by a factor of $O(\log_2 \omega)$, because the FFT-based mode coupling on the azimuthal dimension will cost $O(\omega \log_2 \omega)$. This extreme case, however, seems irrelevant because a global wavefield with comparable azimuthal and in-plane complexity seems unimaginable; if such a wavefield could exist, it should be solved by a discretized 3-D method such as SPECFEM. For most applications at a global scale, $n_u(s, z)$ is much smaller than this upper limit because of the azimuthal smoothness of global wavefields.

For a given 3-D earth model, a reasonable choice of $n_u(s, z)$ is critical to both solution quality and computational efficiency. However, this can be technically difficult from a user's perspective, because the wavefield complexity, which can be influenced by a variety of factors such as model complexity, source location and depth, propagation path and record length, is difficult to be estimated *a priori*.

To handle this difficulty, we provide two approaches for the determination of $n_u(s, z)$: an empirical equation mainly for tomographic mantle models (Leng *et al.* 2016) and the technique of 'wavefield scanning' for arbitrary 3-D earth models.

For typical tomographic mantle models such as S40RTS and S362ANI, Leng *et al.* (2016) brought forward an efficient empirical equation based on massive trial simulations, taking into account some first-order propagation effects:

$$n_u(s, z) = n_u^{ref} F_s(s) F_\theta(\theta(s, z)) F_d(d(s, z)), \qquad (17)$$

where n_u^{ref} is a constant reference value and F_s, F_θ and F_d three factors determining the pattern of $n_u(s, z)$. $F_s(s)$ alters n_u by s, the distance to the axis, $F_\theta(\theta)$ by the epicentral distance θ, $\theta = \arctan(s/z)$ and $F_d(d)$ locally enlarges n_u near the surface to enhance surface wave resolution, determined by the depth d, $d = R_{earth} - \sqrt{s^2 + z^2}$. Using the direct summation approach for mode coupling and ignoring parallel slowdown, Leng *et al.* (2016) showed that such an empirical equation could lead to a speed-up of 1 to 2 orders of magnitude at a 11 s period compared to a 3-D SEM. It was also observed that such a speed-up basically scaled with frequency across a period range from 68 down to 11 s (Leng *et al.* 2016), a scaling behaviour attractive to high-frequency applications. Based on a number of tomographic mantle models, however, this empirical equation cannot be applied universally, so we bring forward the technique of 'wavefield scanning'.

4.2 Wavefield scanning

For an arbitrary 3-D earth model, 'wavefield scanning' can determine a near-optimal $n_u(s, z)$ that guarantees solution convergence. The idea is straightforward. We run a trial simulation with a starting n_u field, denoted $n_u^{start}(s, z)$, which is sufficiently large to compute the wavefield accurately throughout space and time. For every point in D, the solver monitors the energy of each Fourier mode to identify the unexcited higher-order modes at each time step. At the end of the record length, the highest order ever excited at each point is saved to a file, which can be used repeatedly afterwards for problems with similar parameters, such as those with smoother earth models, nearby events and shorter record lengths. We call this process 'wavefield scanning' and denote the obtained n_u field by $n_u^{scan}(s, z)$, inasmuch as the wavefield automatically optimizes $n_u^{start}(s, z)$ by 'scanning' the model over time.[3] It must be emphasized that the solution converges if $n_u^{scan}(s, z)$ turns out smaller than $n_u^{start}(s, z)$ point-wise, making AxiSEM3D a stand-alone method without a reference solution from a full 3-D method.

The computational cost of a scanning simulation may vary a lot, mainly depending on the choice of $n_u^{start}(s, z)$ and the frequency at which the scanning is performed. We will show it later in this section that $n_u^{scan}(s, z)$ obtained at a lower frequency may still be applicable to higher frequencies. Even at the same frequency, a scanning simulation in AxiSEM3D can be less expensive than the corresponding 3-D simulation unless $n_u^{start}(s, z)$ is unwisely large.

Two points must be addressed for the usage of wavefield scanning. First, in forward modelling, we usually use a Heaviside step function or a δ function as the source–time function, so that we can convolve

the resultant seismograms (the Green's functions) with any source–time functions. However, for a scanning simulation, one should use a source–time function without frequency contents beyond the mesh resolution; otherwise, the scanning result will be dominated by numerical noise. Second, the obtained n_u^{scan} may turn out to be smaller than n_c at some places; in this case, one should increase n_u^{scan} to properly sample the model.

4.3 Global tomographic models

In this subsection, we compare the computational efficiency of SPECFEM and AxiSEM3D for the state-of-the-art global tomographic models. The actual computational cost is measured for the following three 3-D earth models: (i) S362ANI without internal topography, (ii) S40RTS and (iii) S362ANI with internal topography + ellipticity + surface topography + ocean. The model parameters are summarized in Table 1 except for S40RTS, a global tomographic mantle model with $l_{max} = 40$. A comparison between (i) and (ii) can manifest the cost impact of long-wavelength volumetric heterogeneities, and that between (i) and (iii) the cost impact of geometric complexities. Take surface topography for example. It increases the computation cost in three aspects: first, $n_u(s, z)$ will be increased near the surface because surface waves will become more complicated in the presence of topography; second, the global time step will be decreased because the spherical elements in the oceanic crust will become thinner after being stretched by topography; finally, the particle relabelling transformation within the undulated range (between surface and 220 km) demands extra arithmetic operations. These three models should be sufficient to constrain other computational models with different constituents, except for a full 3-D crust, as will be discussed in Section 6.1. The simulation period (T_{simu}) ranges from 50 down to 1 s, reaching the highest observable frequency of global seismic data, while utilizing hardware at all scales from laptop to large-scale supercomputer.

We first show the results of wavefield scanning. The scanning simulations start with a large constant n_u field, that is, $n_u^{start}(s, z) = 1200$. The AxiSEM3D solution has converged with this n_u^{start}, as can be verified by the benchmark solutions in Section 3. Fig. 8 shows the obtained $n_u^{scan}(s, z)$ for our three earth models. The sum of $n_u^{scan}(s, z)$ over all the GLL points, as denoted by Σn_u in Fig. 8, can be a estimation of the total computation cost. The average of $n_u^{scan}(s, z)$, as denoted $n_{u, ave}$ in Fig. 8, represents the computational cost consumed by the azimuthal dimension, as it further excludes the in-plane contributions (the number of GLL points). The maximum of $n_u^{scan}(s, z)$, as denoted $n_{u, max}$ in Fig. 8, indicates the strongest azimuthal scattering strength, which is less significant if such large values are distributed only locally. The following remarks can be made based on Fig. 8:

(i) *general pattern*: $n_u^{scan}(s, z)$ takes on a general pattern for the considered global tomographic models. The highest values are found in the crust and the uppermost mantle (even without surface topography), as the surface waves sample the 3-D structures more extensively than the body waves do. For the majority of the mantle regions far from the surface, n_u^{scan} remains small (less than 200 in our cases), which roughly increases with epicentral distance. Also, n_u^{scan} is very small in the outer core;

(ii) *model-adaptivity*: $n_u^{scan}(s, z)$ increases with model complexity. A comparison between Figs 8(a) and (b) shows that, in the mid and lower mantle, the azimuthal scattering effects of S40RTS ($l_{max} = 40$) are stronger than that of S362ANI ($l_{max} = 22$), as reflected by $n_{u, ave}$; however, S362ANI results in a larger $n_{u, max}$ in the crust,

---

[3] Essentially, this is a learning process for parameter optimization, adopted by many modern numerical packages such as FFTW3 (Frigo & Johnson 2005) and EIGEN3 (Guennebaud *et al.* 2010). In our code, a $n_u^{scan}(s, z)$ is called a 'wisdom', following the spirit of FFTW3.

n_u^{scan} | 2 | 50 | 100 | 150 | ≥200

(a) S362ANI without internal topography

(b) S40RTS
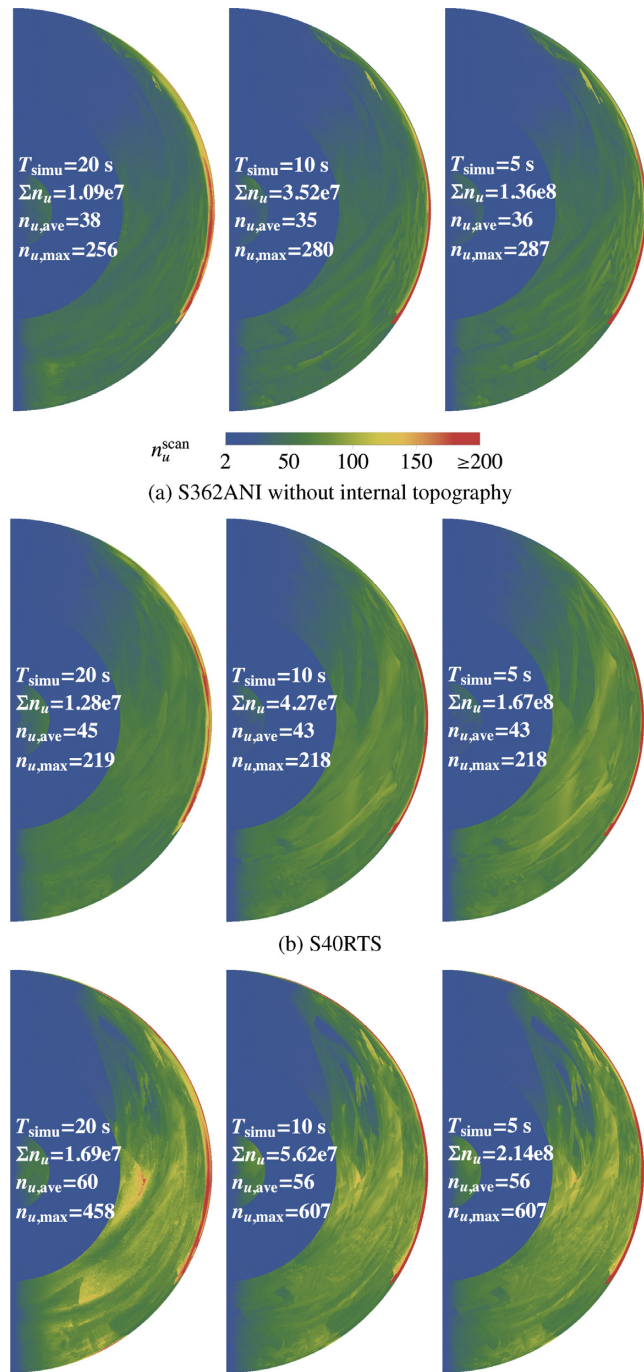
Figure 8. Fields of $n_u^{scan}(s, z)$ obtained by wavefield scanning for different global tomographic models. The model parameters are summarized in Table 1. All the scanning simulations start from a constant n_u field, $n_u^{start}(s, z) = 1200$. In each plot, Σn_u, $n_{u, ave}$ and $n_{u, max}$ denote, respectively, the sum, average and maximum of $n_u^{scan}(s, z)$ over all the GLL points. All the plots use the same colour scale displayed in (a), with saturated colouring for $n_u^{scan}(s, z) \geq 200$.

because it is more complicated than S40RTS in some parts of the uppermost mantle (see e.g. fig. 3 in Leng *et al.* 2016). Comparing Figs 8(a) with (c), we find that surface topography greatly increases n_u^{scan} near the surface, implying its strong scattering effects on the surface waves. The undulating transition zone and ellipticity do not substantially increase n_u^{scan} in the mantle, as they vary smoothly in space.
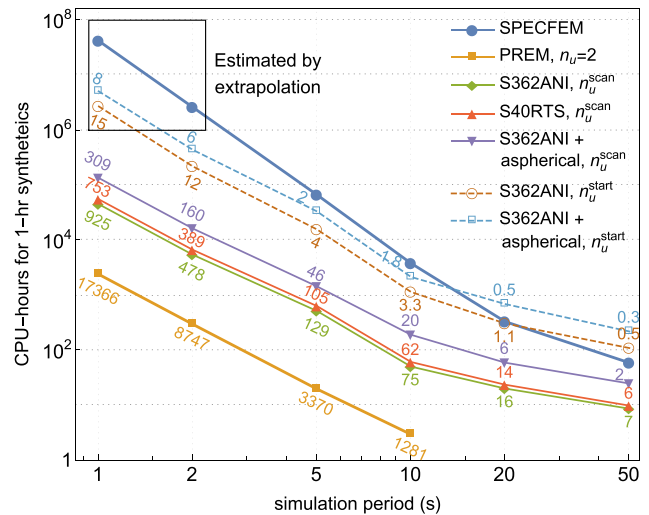
Figure 9. Total CPU-hr required by AxiSEM3D and SPECFEM to compute 1-hr length seismograms. For SPECFEM, we only show the CPU-hr required for PREM (thick blue), which remains the same for S362ANI and S40RTS and will slightly increase when topography is turned on. For AxiSEM3D, we show the CPU-hr for PREM (thick orange) and the three 3-D earth models indicated in Fig. 8. We use $n_u^{start} = 1200$ for the scanning simulations (dashed). Note that we do not show the scanning cost for S40RTS, as it is identical to that for S362ANI. The numbers near the data points show the speed-up of AxiSEM3D with respect to SPECFEM. The framed points in the top-left corner are estimated by extrapolation because of unaffordable solutions. The CPU-hr are measured on a Cray XC30 supercomputer (Archer, UK), including parallelization overheads.

(iii) *weak frequency-dependence*: for all the three models, $n_u^{scan}(s, z)$ basically remains unchanged as the period drops from 10 to 5 s, except for a small fraction of elements near the surface. It is implied that, for global tomographic models with finite power spectra, there exists a frequency above which the azimuthal complexity of wavefields may cease to develop for body waves, so that the computational cost of AxiSEM3D will start to scale with $\omega^{3+\epsilon}$, with ϵ accounting for surface waves, significantly smaller than 1. Such an observation supports the weak frequency-dependence of n_u observed by Leng *et al.* (2016), which benefits high-frequency applications.

To find a near-optimal $n_u(s, z)$ such as those shown in Fig. 8, one needs expensive scanning simulations at the target period and for all different models of interest. In practice, however, one does not always need $n_u(s, z)$ to be optimal, as long as it is correct and efficient. For instance, if one aims at propagating waves in a series of models, wavefield scanning can be performed only on the one with the highest complexity. Also, based on the weak frequency-dependence of n_u^{scan}, one may perform scanning at some lower frequency and enlarge the obtained n_u^{scan} near the surface before applying it to the target frequency.

The CPU-hr required to compute a 1-hr record length are plotted in Fig. 9 for both the scanning simulations using $n_u^{start}(s, z) = 1200$ and the repeated ones reusing $n_u^{scan}(s, z)$. For the considered 3-D earth models and at a period ranging between 5 and 1 s, the speed-up of AxiSEM3D with respect to a discretized 3-D method can reach 2 to 3 orders of magnitude, using the optimized n_u^{scan}. Such a speed-up basically scales with frequency. Therefore, AxiSEM3D cannot only promote the computational efficiency for the period range already achieved by 3-D methods (practically above 5 s), but also cut a path to the highest observable frequency range (up to

1 Hz) in global seismology. This is true for body wave studies of the deep interior, to say the least, because this section does not involve a 3-D crust. For the scanning simulations, AxiSEM3D remains faster at 10 s and below. Fig. 9 again displays the model adaptivity of AxiSEM3D: as the model evolves from PREM to a complex 3-D model, it bridges the gap between fast dimension-reduction methods and slow discretized 3-D methods by locally adapting the Fourier expansion order.

In addition to model complexity, the computational cost of AxiSEM3D can be influenced by many other factors, among which the most significant two are source depth and record length. For deep events without surface wave excitation, n_u will not be locally increased near the surface; for example, the CPU-hr can be reduced by ~15 per cent for (a) and (b) and ~25 per cent for (c). An example will be provided in the next subsection, using a deep earthquake and S40RTS. The computational cost increases with record length in a super-linear style, as the aggregated azimuthal complexity of wavefields always evolves with time, that is, a longer recorder length leads to a higher $n_u(s, z)$; for instance, halving the record length to 30 min will nearly quarter the CPU-hr for those simulations using n_u^{scan}; doubling the record length to 2 hr, however, will only multiply the CPU-hr by a factor of ~2.5, because the body waves have well sampled the interior by 1 hr. Both the above two effects suggest a phase-dependent cost characteristic, the essence of which is the adaptivity of n_u over the time dimension. Such a characteristic greatly benefits studies based on early body wave phases.

4.4 Small-scale heterogeneities

Pseudospectral methods are generally believed to be less competitive in handling localized features because of their white power spectra. Therefore, we investigate in this subsection how a small-scale heterogeneity affects the azimuthal complexity of wavefields and thus the efficiency of AxiSEM3D.

For this purpose, we place a spherical slow blob in the mid-mantle of PREM, as described in Fig. 10. The centre of the blob is 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 per cent, and the density by 20 per cent. Such impedance contrast should be sufficiently strong to exceed by far the most extreme applications in the deep Earth. Given the material perturbations, the scattering effects of the blob depend on both its size and the propagated wavelength. Here we consider three simulation periods (20, 10 and 5 s) and four blob radii (20, 40, 80 and 160 km), with 12 simulations in total. We define the structure–wavelength ratio ζ as

$$\zeta = \frac{2R_{blob}}{\lambda_s} = \frac{2R_{blob}}{v_s T_{simu}}, \tag{18}$$

where R_{blob} is the blob radius and λ_s the S wavelength of the background medium. In our parameter space, ζ increases from ~0.3 to ~10, shifting the scattering regime from point scattering to multiple scattering within a heterogeneous medium.

We use wavefield scanning to investigate how the blob affects the azimuth complexity of the wavefield. Again, we start from $n_u^{start}(s, z) = 1200$. In addition to solution convergence, we have to make sure that the blob is properly sampled along the azimuthal direction by such a n_u field. Let us first consider the in-plane sampling by the 2-D GLL points. The shortest period of interest is 5 s, so the shortest S wavelength is ~32.5 km (with $v_s \approx 6.5$ km s^{-1} at the blob centre). Adopting the one-element-per-wavelength convention with polynomial order of 4, the average spacing between the GLL points (in both SPECFEM and AxiSEM3D meshes) will
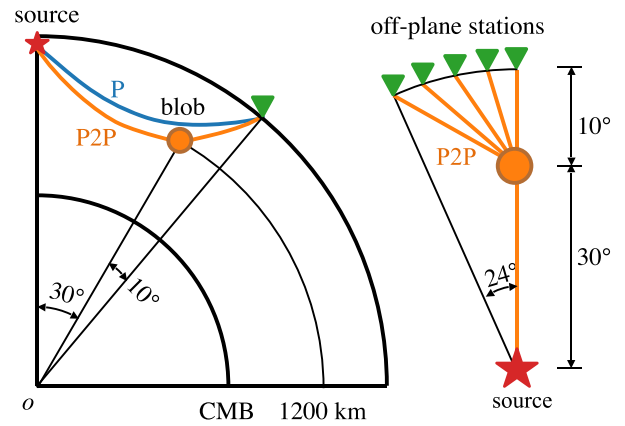

Figure 10. Wave scattering by a spherical slow blob in the mid-mantle of PREM. We show the source–structure plane on the left and the off-plane view on the right. The blob is located 30° in distance and 1200 km in depth. Within the blob, the S- and P-wave velocities are, respectively, reduced by 50 and 30 per cent, and the density by 20 per cent. The perturbation is constant inside the blob and fade in a Gaussian style outside the blob, halved at 1.1R_{blob} away from the centre. The path labelled 'P' corresponds to the direct P arrival and 'P2P' the P-to-P conversion at the blob. Not concerned with surface waves, we use a deep earthquake at depth 647.1 km. The off-plane stations are located at 40° in distance and distributed between 0° and 24° in azimuth; seismograms recorded by these stations will be shown in Fig. 12.

be ~8 km. Now we consider the azimuthal sampling. The distance from the blob centre to the axis is 2585.5 km. Rotating the blob centre around the axis by 2π generates a ring, whose circumference is ~16 245 km. There will be 2401 (=$2n_u + 1$) cardinal gridpoints evenly distributed on this ring. Hence, the spacing between any two cardinal gridpoints will be about 16 245/2401 ≈ 6.8 km, independent of T_{simu}. Therefore, for $T_{simu} \geq 5$ s, the azimuthal sampling is finer than the in-plane sampling. The smallest blob we consider has a diameter of 40 km, which will be sampled by about six GLL points at 5 s and might be totally missed at 20 s. Therefore, for a better model sampling, we double the perturbation range by making the material perturbations fade in a Gaussian style outside the blob. Note that the above difficulty in sampling sub-wavelength structures exists not only in AxiSEM3D but also in any discretized 3-D methods without local refinement.

Fig. 11 shows the results of wavefield scanning. Here we introduce a cost factor to quantify the impact of the blob, as denoted by f_{S40} in Fig. 11, which is given by Σn_u required by the blob relative to that by S40RTS at the same period. In other words, instead of using absolute computing time, we measure the cost impact of the blob with respect to that of a typical tomographic model, so as to exclude any factors other than structure-induced wavefield complexity. In general, Fig. 11 confirms the competency of AxiSEM3D in handling such localized small-scale heterogeneities, as the maximum f_{S40} is only 1.13, with $R_{blob} = 160$ km and $T_{simu} = 20$ s. Looking into the patterns of $n_u^{scan}(s, z)$, we reach the following remarks:

(i) Reading Fig. 11 row-wise, we find that the cost impact of the blob quickly increases with its size, as f_{S40} is enlarged by ~20 times as R_{blob} increases from 20 to 160 km. When the blob is small ($R_{blob} \leq 40$ km), it has localized influence on the azimuthal wave complexity, with n_u^{scan} highly concentrated in a smeared neighbourhood around it, taking on a comet-like pattern. When the blob becomes
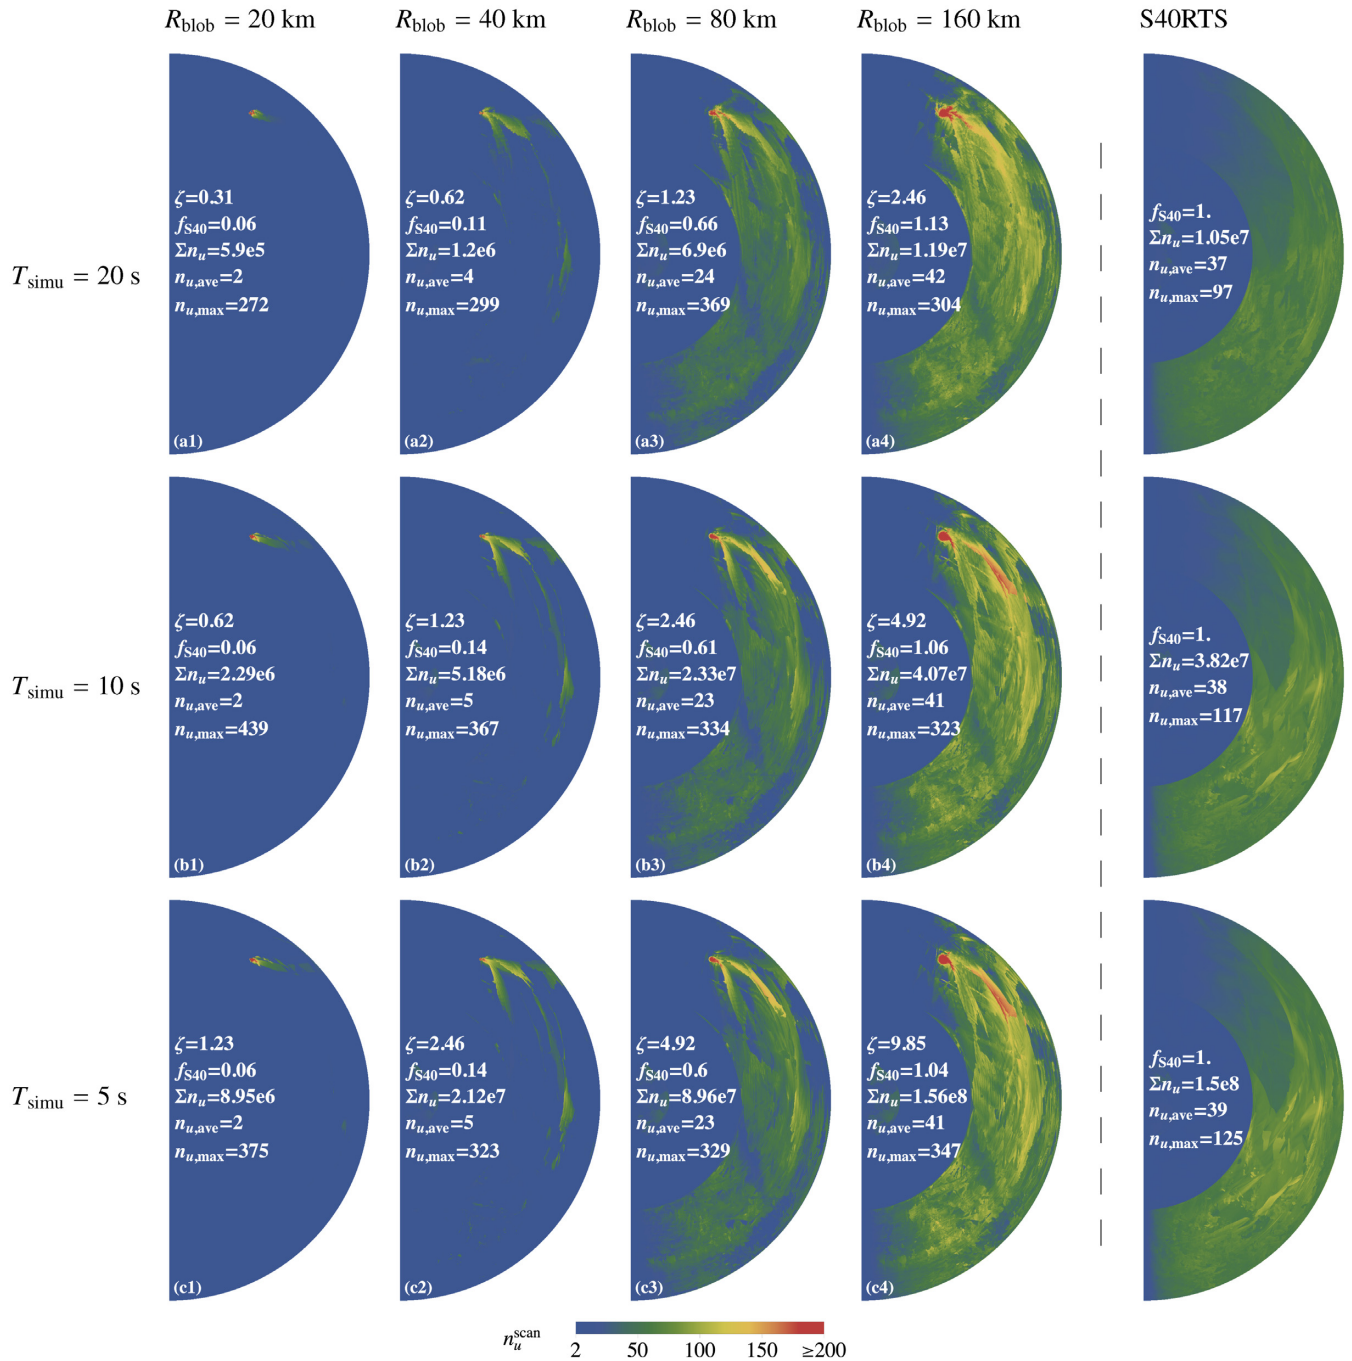
Figure 11. Fields of $n_u^{scan}(s, z)$ obtained by wavefield scanning for a spherical slow blob in the mid-mantle of PREM. The problem is described in Fig. 10. The simulation period T_{simu} varies column-wise and the blob radius R_{blob} varies row-wise. As a reference for cost assessment, $n_u^{scan}(s, z)$ obtained with S40RTS is appended at the end of each row (note that they are different from those shown in Fig. 8(b) because of different event depths). The factor f_{S40} is the ratio between Σn_u for the blob and that for S40RTS, which indicates the azimuthal scattering strength of the blob relative to S40RTS.

increasingly large ($R_{blob} \geq 80$ km), its azimuthal scattering effects extend to the entire domain, reaching a cost impact close to that of a tomographic model when $R_{blob} = 160$ km.

(ii) The cost impact of the blob turns out insensitive to frequency, as f_{S40} remains mostly unchanged column wise. For a fixed blob size, n_u^{scan} grows with frequency only in the near field, as can be verified by all the four columns. Such weak frequency-dependence again benefits high-frequency applications. Technically, we can perform wavefield scanning at a lower frequency and apply the obtained

$n_u^{scan}(s, z)$ to a higher frequency after enhancing it in the near field of a small-scale heterogeneity.

(iii) The maximum value of $n_u^{scan}(s, z)$, $n_{u, max}$, always appears inside the blob. It is shown that the maximum of $n_{u, max}$ does not occur at the largest blob size or at the shortest wavelength; instead, it occurs when the structure–wavelength ratio is close to 1. This is because a smaller blob has a higher azimuthal expansion order (n_c) and is thus supposed to have stronger azimuthal scattering effects; however, as explained at the beginning of this subsection,

(a) $T_{simu} = 5$ s

(b) $T_{simu} = 1$ s
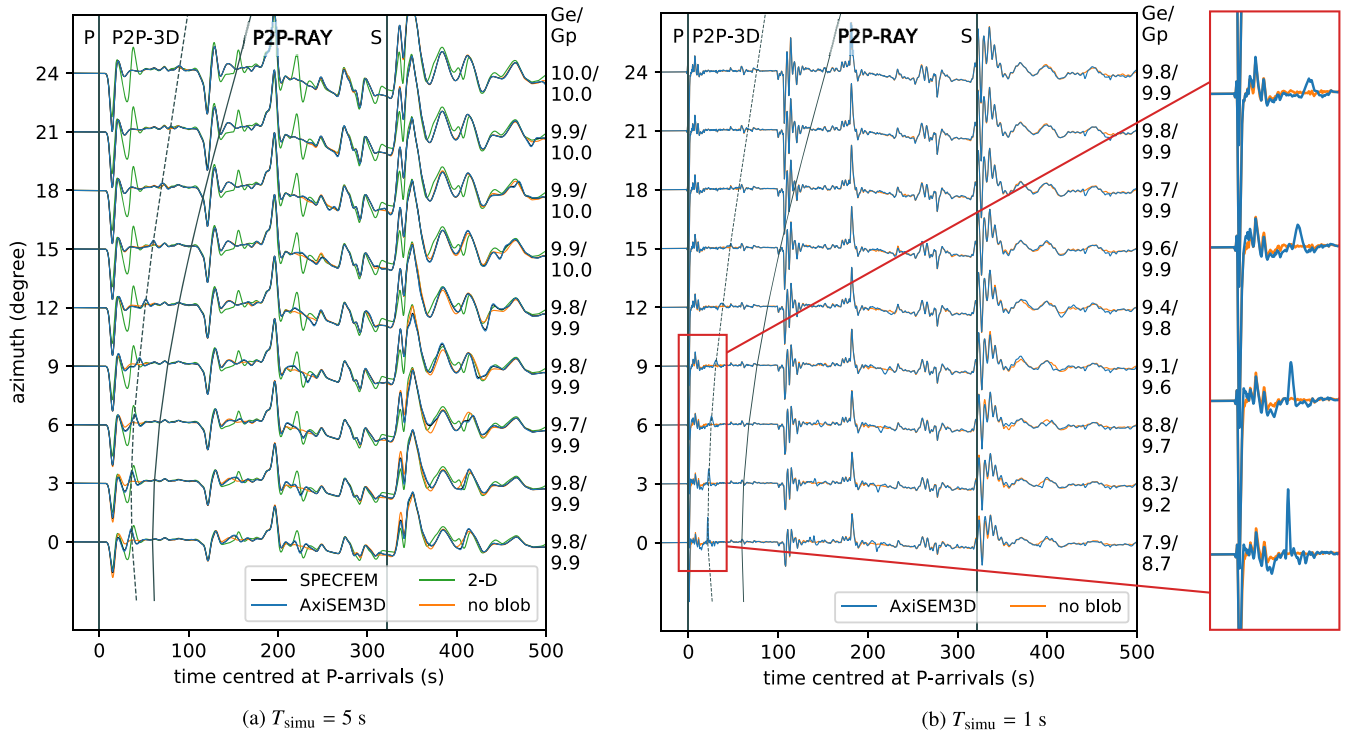
Figure 12. Comparisons of synthetic seismograms (vertical component) for a spherical slow blob in the mid-mantle of PREM. The problem is described in Fig. 10, with $R_{blob} = 160$ km. For the high-frequency simulation in (b), we use four times of the $n_u^{scan}(s, z)$ scanned at 20 s, as shown in (a4) of Fig. 11. The traces labelled 'SPECFEM' (black) and 'AxiSEM3D' (blue) show the full 3-D solutions, those labelled 'no blob' (orange) the 1-D solution for PREM without the blob and those labelled '2-D' (green) the 2-D axisymmetric solution for which the source–blob plane is rotated by 2π around the axis to generate the computational model (so the blob becomes a torus). The SPECFEM traces (black) in (a) are mostly invisible because of small waveform differences. The 2-D traces change with azimuth only because of the radiation pattern from the source. The traveltime curve labelled 'P2P-RAY' is computed by ray tracing, taking the blob centre as its representative location, whereas the one labelled 'P2P-3D' (dashed) simply connects the peaks of P2P in the 3-D solution. In (a), G_e and G_p are calculated using the 'SPECFEM' and 'AxiSEM3D' traces, computed for the same model but by the two different methods; in (b), G_e and G_p are calculated for the two different models (PREM with and without the blob).

(a) element classes

(b) $n_u(s, z)$

(c) measured elemental cost

(d) domain decomposition
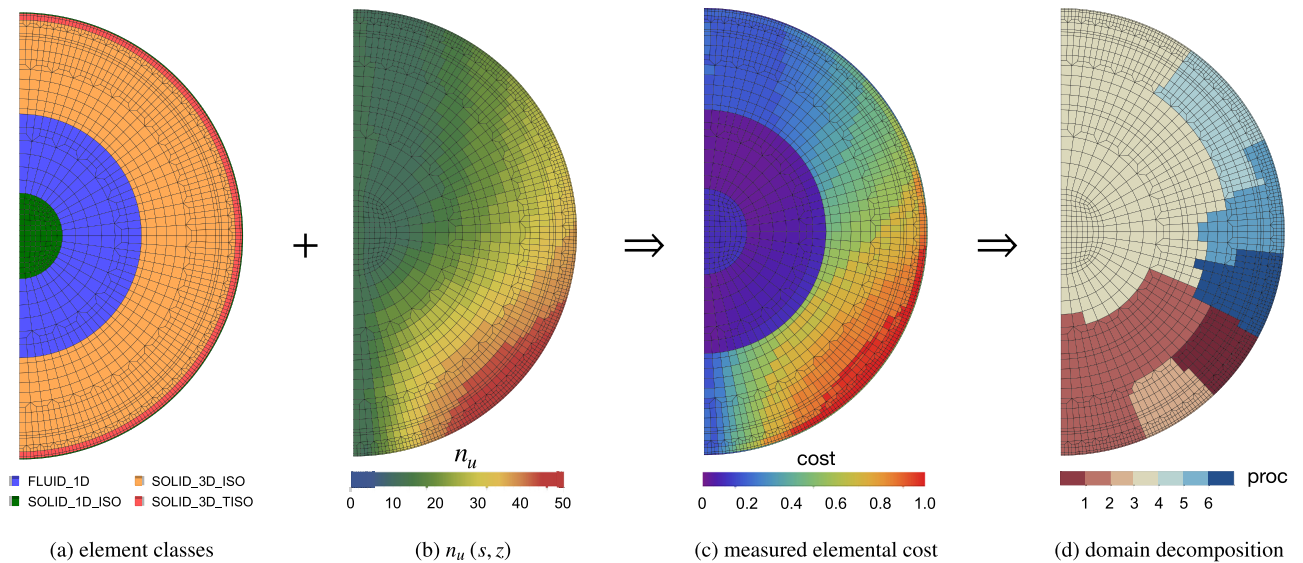
Figure 13. An example of domain decomposition in AxiSEM3D. The global mantle model S40RTS is considered at a 50 s period. (a) Shows the 2-D spectral element mesh, which contains four element classes (in the legend, ISO means isotropic and TISO transversely isotropic). (b) Shows a $n_u(s, z)$ field computed by the empirical equation, eq. (17). 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 cost at runtime. Finally, (d) shows the heterogeneous domain decomposition produced by METIS, weighted by the measured elemental cost in (c). The varying numbers of elements on the different processors clearly reflect the effect of cost weighting.

a sub-wavelength blob may be smoothed by the mesh because of an incomplete spatial sampling. Also, it is shown that $n_{u, max}$ can be smaller than n_c, implying that a structure seen by the wave may be smoother than it actually is; this is consistent with the homogenization theory (Capdeville *et al.* 2010).

Fig. 11 also exhibits a healing effect of azimuthal wave complexity: for a small-scale heterogeneity, n_u^{scan} may have readily recovered on the surface; in this sense, constraining the transverse dimension of a small-scale heterogeneity may be more difficult than constraining the in-plane ones. Such a healing effect takes place over the range of influence of the blob, which mainly depends on the blob size. Note that the overall cost impact actually depends on the proportion such a range of influence makes up of the entire computational domain, and that is how the propagation distance (Aki & Richards 1980; Igel 2017) comes into play. In this study, we do not change the propagation distance, but, for instance, if we switch from the Earth to the Moon, the blob sizes should also be decrease accordingly for the above remarks to hold.

In Fig. 12, we display the seismograms recorded by the off-plane stations shown in Fig. 10, with $R_{blob} = 160$ km. We choose this record section because the off-plane scattering effects can only be modelled with a full 3-D simulation. The seismograms simulated at 5 s are shown in Fig. 12(a), from which the following remarks can be made:

(i) An excellent agreement between SPECFEM (black) and AxiSEM3D (blue) is achieved for this problem.

(ii) 3-D scattering effects can be clearly observed from the difference between the AxiSEM3D (blue) and the 1-D (orange) solution: as the station moves away from the central azimuth, the backscattered energy (amplitude of P2P) decays while the differential traveltime between P and P2P enlarges.

(iii) The difference between the AxiSEM3D (blue) and the 2-D (green) solution manifests the insufficiency of the axisymmetric approximation: first, it overestimates the amplitudes of the scattered phases at the central azimuth; second, it totally misses out the off-plane scattering effects, which can be useful in constraining the lateral scale of the scatterer.

(iv) The difference between the two traveltime curves, respectively, labelled 'P2P-3D' and 'P2P-RAY' disqualifies the centre of blob as a representative location to predict the arrival times of the scattered phases by ray tracing, as the blob size has far exceeded the regime of point scattering.

In addition, we perform a high-frequency simulation at 1 Hz and display the seismograms in Fig. 12(b). All the above remarks hold at 1 Hz except for the first one, because a 1 Hz simulation with SPECFEM is unaffordable. The window around P2P (zoomed-in) contains multiple wiggles at 0° and 3° (note that the blob radius spans ~3.5°), which should originate from multiple encounters on the boundary of the blob. Computationally, we use four times of the $n_u^{scan}(s, z)$ scanned at 20 s, as shown in (a4) of Fig. 11. This choice can be rather conservative, as we may only need to increase n_u^{scan} in the near field, based on the weak frequency-dependence of n_u^{scan} observed in Fig. 11. Nevertheless, even with such a conservative choice, the computational cost is only ~500 000 CPU-hr for a 1-hr record length. This should be the first time that a 1 Hz frequency is achieved with practical resources for global wave propagation in a 3-D earth model.

Finally, it is noted that we have only examined a single scatterer with constant perturbation strength. Concerning the perturbation

(a) strong scaling   (b) weak scaling
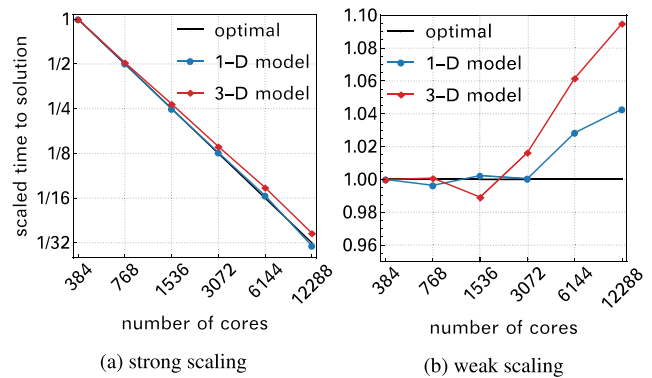
Figure 14. Scaling behaviours of AxiSEM3D on a Cray XC30 supercomputer (Archer, UK). Two earth models are tested, the 1-D model PREM with $n_u(s, z) = 2$, and the 3-D model S40RTS with $n_u(s, z)$ determined by the empirical equation, eq. (17). For strong scaling, the problem (total workload) is fixed, 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 normalized wall-clock times measured for 2000 time steps.
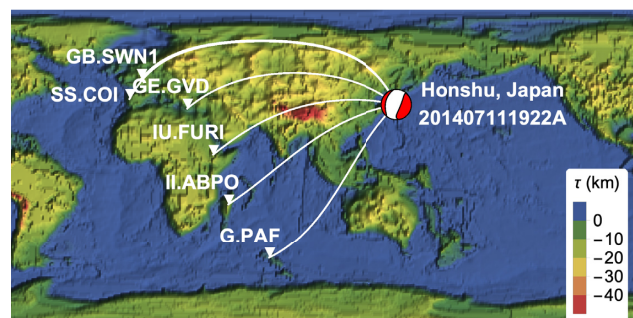
(a) Moho in Crust 1.0
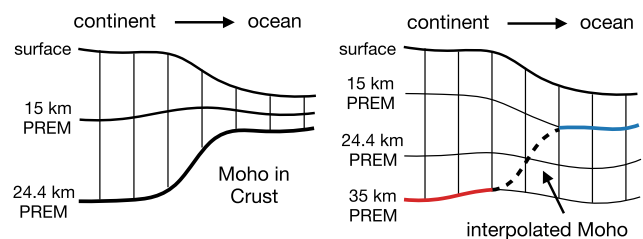
(b) implementations in AxiSEM3D and SPECFEM

Figure 15. Moho discontinuity in Crust 1.0 and its implementation. (a) shows the Moho topography in Crust 1.0 with respect to the PREM Moho (24.4 km). It also includes an earthquake event and several seismic stations for which the seismograms will be displayed in Fig. 16. (b) shows the different strategies adopted in AxiSEM3D (left) and SPECFEM (right) for Moho implementation. AxiSEM3D always stretches the spherical Moho in the 1-D reference model (24.4 km in PREM) to match the undulating 3-D Moho. In SPECFEM, the mid-crust discontinuity (15 km in PREM) is stretched upward to match the oceanic part, while an artificial discontinuity (35 km in PREM) is stretched downwards to match the continental part; for the rest beneath the ocean–continent plate boundaries, the 3-D Moho cuts through the spectral elements and is thus smoothed by means of material interpolation.

strength, we are on the safe side, but it is still interesting to investigate how the wavefield complexity evolves with it. In the case of multiple scatterers, the field of $n_u(s, z)$ has to be large wherever a scatterer presents. For the extreme case where the scatterers are distributed throughout the mantle, a large constant $n_u(s, z)$ may become necessary, which may not be readily reduced by wavefield scanning. It then depends on the model complexity and the period whether and to what extent AxiSEM3D can maintain its computational advantage over a discretized 3-D method, as will be discussed in Section 6.1.

5 PARALLELIZATION

AxiSEM3D is parallelized via the MPI using non-blocking asynchronous communication. Load-balancing across processors is the most important issue around parallelization. In ordinary SEMs such as AxiSEM and SPECFEM, since the computational cost of an individual element is constant (not to consider 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 (Karypis & Kumar 1998) to partition the 2-D computational domain, with a weight assigned to each element. The weights are determined by measuring the elemental cost during runtime.[4] A complete procedure of domain decomposition is described with an example in Fig. 13.

The scaling behaviour of such a measure-and-decompose scheme is tested on the supercomputer Archer (https://www.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, for which $n_u(s, z) = 2$ across the domain, so the decomposition is unweighted. Perfect strong scaling is observed up to 12 288 cores for a 1 Hz simulation, as shown in Fig. 14(a). Weak scaling turns out slightly suboptimal, with a 4 per cent parallelization overhead at 12 288 cores, as shown in Fig. 14(b). 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 cost (Fig. 13c) fluctuates; second, extra workload emerges on the processor boundaries (Fig. 13d) after domain decomposition, which cannot be balanced *a priori*. The 3-D scaling behaviours shown in Fig. 14 are observed with S40RTS, using an empirical $n_u(s, z)$ sufficiently large for solution convergence. For strong and weak scaling, the overheads are, respectively, 6 and 10 per cent at 12 288 cores. Compared to the speed-ups shown in Fig. 9, these overheads are ignorable. It is noted that the above scalabilities are evaluated in reference to the simulations with 384 cores. Compared to serial simulations, AxiSEM3D still suffers a parallelization overhead around 30 per cent for 3-D earth models, mainly due to unstable cost measurement. This is an issue left for future optimization.

The high-performance code is available open-source at https://github.com/kuangdai/AxiSEM3D. The code features an easy implementation of user-defined models. The 2-D spectral element

---

[4]Most of the computing time in AxiSEM3D is consumed by FFT and matrix arithmetics, for which we rely on the high performance packages FFTW3 (Frigo & Johnson 2005) and EIGEN3 (Guennebaud *et al.* 2010), respectively. The element cost becomes unpredictable due to the low-level optimizations in these packages and can thus be obtained only through runtime measurement.

meshes used in AxiSEM3D are generated by the Salvus mesher (Afanasiev *et al.* 2017), which can mesh any 1-D spherically symmetric models not limited to Earth (e.g. Mars, Bozdağ *et al.* 2016b). We provide two approaches to implement a user-defined 3-D model. First, AxiSEM3D accepts discrete models defined on a structured grid, supporting several formats such as the IRIS-EMC database (http://ds.iris.edu/ds/products/emc/) and the SubMachine database (http://submachine.earth.ox.ac.uk/). Second, written in modern C++ with an object-oriented framework, AxiSEM3D allows users to define their models via class inheritance, which maximizes the flexibility for model description. In addition, we attach great importance to wavefield visualization, which zooms in wave scattering effects at all scales. Some animations of global wave propagation can be found on the Youtube channel 'Seismology Oxford'.

6 DISCUSSIONS

Global wave propagation in 3-D earth models is fundamental to seismological probings of the Earth's interior. Though a number of comprehensive 3-D numerical methods have been developed, the insufficiency of computing power remains outstanding, as the multiscale nature of Earth structure and the observable frequency band of global seismic data both require high-frequency waveform modelling.

To exploit the inherent azimuthal smoothness of a global wavefield, Leng *et al.* (2016) parametrized its azimuthal dependence in Fourier series, which substantially reduces the number of degrees of freedom required by an accurate description of solution. Leng *et al.* (2016) developed a pseudospectral/spectral-element hybrid method, AxiSEM3D, to solve the dimension-reduced system. This paper is an essential extension of Leng *et al.* (2016). First, we generalize AxiSEM3D to aspherical earth models with undulating discontinuities using the particle relabelling transformation, which preserves the spherical geometry of the computational domain and thus the functional form of the dimension-reduced system as well as its SEM implementation. We also verify the robustness of AxiSEM3D to localized small-scale structures with strong perturbation strength. In addition, we enhance the usability and performance of AxiSEM3D by developing the technique of wavefield scanning, an FFT-based mode coupling approach and a heterogeneous domain decomposition for load-balanced parallelization.
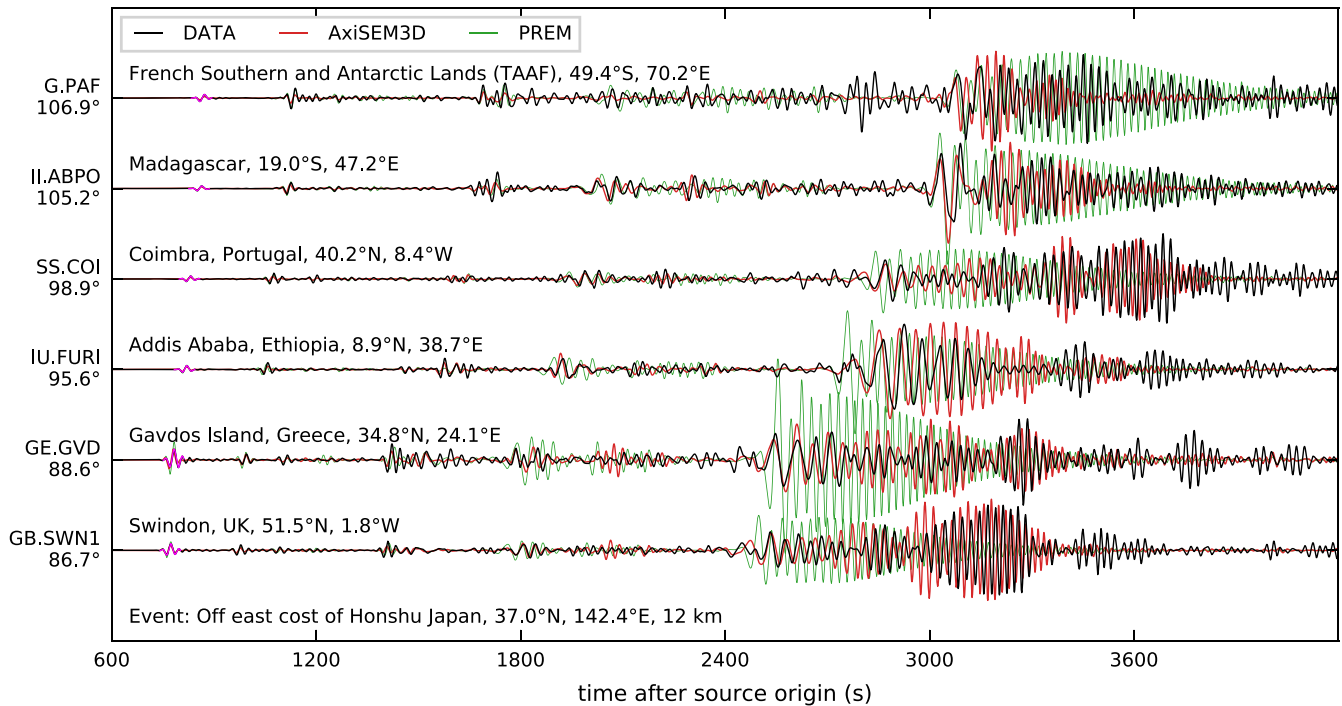
All these efforts make AxiSEM3D a stand-alone numerical method for global wave propagation in 3-D earth models characterized by the following properties:

(i) *Comprehensive*: AxiSEM3D allows for material heterogeneities, such as velocity, density, anisotropy and attenuation, as well as 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. Ocean is modelled as a hydrodynamic load.

(ii) *Accurate*: AxiSEM3D is cross-verified with an independent discretized 3-D SEM, SPECFEM3D_GLOBE, for both material heterogeneity and interface undulation. 3-D crustal structures such as Crust 1.0 have been implemented in AxiSEM3D; in the absence of a reference solution, however, our implementation cannot be verified at this point, as discussed in the following subsection.

(iii) *Efficient*: Model-adaptivity is the central keyword for AxiSEM3D. By *locally* adapting the Fourier expansion order of solution to the azimuthal complexity of a wavefield, AxiSEM3D bridges the gap between the fast dimension-reduction methods (for

Figure 16. Comparisons between synthetic seismograms (vertical component) computed by AxiSEM3D and real seismic data. To construct the 3-D earth model, we use S40RTS for the mantle and Crust 1.0 for the crust, also incorporating ellipticity, topography and ocean. The AxiSEM3D traces (red) are computed using this model. The simulation period is 10 s and we filter the seismograms by a bandpass between 20 and 100 s. The traces are shifted by cross-correlating the direct *P*-wave window (pink) and plotted in their original amplitudes (no normalization or correction). The synthetic traces for PREM (green) are retrieved from Syngine (Krischer *et al.* 2017). The real seismic data (black) are retrieved using obspyDMT (Hosseini & Sigloch 2017).

1-D earth models) and the slow discretized 3-D methods. Excluding a 3-D crust, for the state-of-the-art global tomographic models with undulating discontinuities, the speed-up of AxiSEM3D with respect to a 3-D SEM may reach 2 to 3 orders of magnitude for a seismic period ranging from 5 to 1 s. Given the earth model, such speed-up basically scales with frequency. The observable frequency band (up to 1 Hz) of global seismic data has been achieved for 3-D mantle models with reasonable computing resources.

6.1 3-D Crust

A realistic 3-D crustal structure is essential for surface wave modelling. A global crustal model such as Crust 1.0 involves an undulating Moho discontinuity and material heterogeneity between the surface and the Moho. Crust 1.0 has been implemented in AxiSEM3D, including ice sheets and sediments. However, we cannot fully verify our implementation with SPECFEM due to a persistent inconsistency. In this subsection, we briefly discuss the complication around crust implementation and verification, and compare some synthetic seismograms computed by AxiSEM3D to real seismic data.

Compared to long-wavelength mantle structures, the Earth's crust varies much more drastically in both amplitude and wavelength. The most outstanding challenge comes from a severely varying crustal thickness, characterizing the Moho discontinuity with large undulations and sharp slopes beneath the ocean–continent plate boundaries. For instance, Fig. 15(a) shows the undulating Moho in Crust 1.0 defined on a 1×1 degree grid in which the crustal thickness varies from ∼8 to ∼75 km. Fig. 15(b) explains the different implementations of the Moho discontinuity in AxiSEM3D and SPECFEM. In AxiSEM3D, the 3-D Moho is completely honoured

by the mesh, since the particle relabelling transformation requires homeomorphism of any radial discontinuity to a sphere. In contrast, SPECFEM honours part of the Moho by stretching the radial interfaces and the rest by material interpolation through the interior of spectral elements. The interpolated part takes up around one fourth of the entire Moho area and mainly overlaps with the seismically active regions along the plate boundaries (where the crustal thickness varies dramatically).

Concerning verification, AxiSEM3D cannot follow the way the undulating Moho has been approximated in SPECFEM due to the requirement of homeomorphism. In principle, as a full 3-D method, SPECFEM could be modified such that an integrated Moho is honoured; however, modifying the crustal module in SPECFEM seems technically immense due to its hardcoded nature. Besides, it would be unjustified to substantially modify such a complex code without validation, turning the exercise upside down in that we would always need an independent reference solution.

Computationally, the SPECFEM implementation can benefit from a smaller global time step, as controlled by the number of element layers within the thin oceanic crust. For Crust 1.0, the time step of AxiSEM3D is about half of that of SPECFEM. Because n_u will also be increased near the surface, a 3-D crust could decrease the speed-up of AxiSEM3D by one order of magnitude. Still, AxiSEM3D could remain faster at high frequencies, as can be estimated from Fig. 9.

Because of its non-standard, hardcoded treatment of the Moho discontinuity, SPECFEM could no longer serve as a reference solution for surface wave propagation in the Earth's 3-D crust. We leave the verification to a future study using Salvus (Afanasiev *et al.* 2017), a recently developed 3-D SEM that can honour the Moho in a continuous manner as AxiSEM3D does. In the remainder of this

section, we briefly compare AxiSEM3D synthetics to real seismic data to illustrate the strong scattering effects of a 3-D crust.

Our 3-D earth model is a combination of S40RTS and Crust 1.0. The event is a shallow earthquake in Japan. We show six stations located between 85° and 110° in distance, as shown in Fig. 15(a), with the paths going through complex crustal structures such as the sharp plate boundaries and the thickest part of the Eurasian plate. The seismograms are shown in Fig. 16. The 3-D earth model remarkably enhances the solution quality, compared to PREM, especially for the body waves and the low-frequency surface waves. Some high-frequency surface waves are not captured by the synthetics, for which we presume two possible reasons: first, the resolution of Crust 1.0 might be insufficient to model those high-frequency reverberations; second, as we do not use a 3-D attenuation model, the globally averaged attenuation strength in PREM might be too strong for the considered paths.

In addition to a global 3-D crust, an earth model may contain other types of strong, short-wavelength heterogeneities, such as randomly distributed scatterers and irregularly shaped interfaces. Such heterogeneities have wave scattering effects increasing with frequency and decreasing with propagation distance (problem scale), and are thus of higher interest at regional to local scales. As a convergent method, AxiSEM3D can handle any non-smooth heterogeneities as long as the 3-D model is topologically equivalent to some layered reference model (which is spherical at a global scale). For solution convergence, a more complicated model generally demands a higher Fourier expansion order and thus higher computational cost. Because of the super-linearity of FFT, there exists in theory a limit model complexity at which AxiSEM3D can lose its speed-up with respect to a discretized 3-D method. Such a limit tends to occur at a long period.

In this paper, as focused on a global scale without benchmark for a 3-D crust, AxiSEM3D works in a desirable regime and exhibits remarkable computational advantage. Nevertheless, its cost characteristic in the complete scattering regime remains incompletely revealed, particularly the limit model complexity as a function of frequency and propagation distance (we know that such a limit should increase with frequency and decrease with propagation distance). Extremely varying 3-D heterogeneities are more relevant at a local or regional scale and will be tested in a future study.

ACKNOWLEDGEMENTS

REFERENCES

Afanasiev, M., Boehm, C., van Driel, M., Krischer, L., May, D., Rietmann, M. & Fichtner, A., 2017. Salvus: a flexible high-performance and open-source package for waveform modelling and inversion from laboratory to global scales, in *Geophys. Res. Abstr., EGU2017–9456,* Vol. **19,** Am. Geophys. Un., Washington DC.

Aki, K. & Richards, P.G., 1980. *Quantitative Seismology: Theory and Methods,* 1st edn, W.H. Fredman and Company.

Al-Attar, D. & Crawford, O., 2016. Particle relabelling transformations in elastodynamics, *Geophys. J. Int.,* **205**(1), 575–593.

Al-Attar, D. & Woodhouse, J.H., 2008. Calculation of seismic displacement fields in self-gravitating earth models—applications of minors vectors and symplectic structure, *Geophys. J. Int.,* **175**(3), 1176–1208.

Al-Attar, D., Crawford, O., Valentine, A.P. & Trampert, J., 2018. Hamilton's principle and normal mode coupling in an aspherical planet with a fluid core, *Geophys. J. Int.,* **214,** 485–507.

Amante, C. & Eakins, B., 2009. ETOPO1 1 arc-minute global relief model: procedures, data sources and analysis, in *NOAA Technical Memorandum NESDIS NGDC,* Vol. **24,** NOAA, Silver Spring, MA.

Astiz, L., Earle, P. & Shearer, P., 1996. Global stacking of broadband seismograms, *Seismol. Res. Lett.,* **67**(4), 8–18.

Auer, L., Boschi, L., Becker, T., Nissen-Meyer, T. & Giardini, D., 2014. Savani: a variable resolution whole-mantle model of anisotropic shear velocity variations based on multiple data sets, *J. geophys. Res.,* **119**(4), 3006–3034.

Basini, P., Nissen-Meyer, T., Boschi, L., Casarotti, E., Verbeke, J., Schenk, O. & Giardini, D., 2013. The influence of nonuniform ambient noise on crustal tomography in Europe, *Geochem. Geophys. Geosyst.,* **14**(5), 1471–1492.

Becker, T.W. & Boschi, L., 2002. A comparison of tomographic and geodynamic mantle models, *Geochem. Geophys. Geosyst.,* **3**(1), doi:10.1029/2001GC000168.

Boyd, J.P., 2001. *Chebyshev and Fourier spectral methods,* Courier Corporation.

Bozdağ, E., Peter, D., Lefebvre, M., Komatitsch, D., Tromp, J., Hill, J., Podhorszki, N. & Pugmire, D., 2016a. Global adjoint tomography: first-generation model, *Mon. Not. R. astr. Soc.,* **207**(3), 1739–1766.

Bozdağ, E. *et al.*, 2016b. Simulations of seismic wave propagation on Mars, *Space Sci. Rev.,* **211,** 1–24.

Capdeville, Y., To, A. & Romanowicz, B., 2003. Coupling spectral elements and modes in a spherical Earth: an extension to the 'sandwich' case, *Geophys. J. Int.,* **154**(1), 44–57.

Capdeville, Y., Guillot, L. & Marigo, J.-J., 2010. 2-D non-periodic homogenization to upscale elastic media for P–SV waves, *Geophys. J. Int.,* **182**(2), 903–922.

Capdeville, Y., Stutzmann, É., Wang, N. & Montagner, J.-P., 2013. Residual homogenization for seismic forward and inverse problems in layered media, *Geophys. J. Int.,* **194**(1), 470–487.

Carrington, L., Komatitsch, D., Laurenzano, M., Tikir, M.M., Michéa, D., Le Goff, N., Snavely, A. & Tromp, J., 2008. High-frequency simulations of global seismic wave propagation using SPECFEM3D_GLOBE on 62k processors, in *Proceedings of the 2008 ACM/IEEE conference on Supercomputing,* p. 60, IEEE Press.

Chaljub, E. & Tarantola, A., 1997. Sensitivity of SS precursors to topography on the upper-mantle 660-km discontinuity, *Geophys. Res. Lett.,* **24**(21), 2613–2616.

Chaljub, E. & Valette, B., 2004. Spectral element modelling of three-dimensional wave propagation in a self-gravitating Earth with an arbitrarily stratified outer core, *Geophys. J. Int.,* **158**(1), 131–141.

Chaljub, E., Capdeville, Y. & Vilotte, J.-P., 2003. Solving elastodynamics in a fluid–solid heterogeneous sphere: a parallel spectral element approximation on non-conforming grids, *J. Comput. Phys.,* **187**(2), 457–491.

Chaljub, E. *et al.*, 2015. 3-d numerical simulations of earthquake ground motion in sedimentary basins: testing accuracy through stringent models, *Geophys. J. Int.,* **201**(1), 90–111.

Colombi, A., Nissen-Meyer, T., Boschi, L. & Giardini, D., 2012. Seismic waveform sensitivity to global boundary topography, *Geophys. J. Int.,* **191**(2), 832–848.

Colombi, A., Nissen-Meyer, T., Boschi, L. & Giardini, D., 2014. Seismic waveform inversion for core–mantle boundary topography, *Geophys. J. Int.,* **198**, 55–71.

Cottaar, S. & Romanowicz, B., 2012. An unsually large ULVZ at the base of the mantle near Hawaii, *Earth planet. Sci. Lett.,* **355**, 213–222.

Cupillard, P., Delavaud, E., Burgos, G., Festa, G., Vilotte, J.-P., Capdeville, Y. & Montagner, J.-P., 2012. RegSEM: a versatile code based on the spectral element method to compute seismic wave propagation at the regional scale, *Geophys. J. Int.,* **188**(3), 1203–1220.

Dahlen, F. & Tromp, J., 1998. *Theoretical Global Seismology,* Princeton Univ. Press.

Deuss, A., Redfern, S.A., Chambers, K. & Woodhouse, J.H., 2006. The nature of the 660-kilometer discontinuity in Earth's mantle from global seismic observations of PP precursors, *Science,* **311**(5758), 198–201.

Dumbser, M. & Käser, M., 2006. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – II. the three-dimensional isotropic case, *Geophys. J. Int.,* **167**(1), 319–336.

Dziewonski, A.M. & Anderson, D.L., 1981. Preliminary reference earth model, *Phys. Earth planet. Inter.,* **25**(4), 297–356.

Etienne, V., Chaljub, E., Virieux, J. & Glinsky, N., 2010. An hp-adaptive discontinuous Galerkin finite-element method for 3-D elastic wave modelling, *Geophys. J. Int.,* **183**(2), 941–962.

Friederich, W. & Dalkolmo, J., 1995. Complete synthetic seismograms for a spherically symmetric earth by a numerical computation of the Green's function in the frequency domain, *Geophys. J. Int.,* **122**(2), 537–550.

Frigo, M. & Johnson, S.G., 2005. The design and implementation of FFTW3, *Proc. IEEE,* **93**(2), 216–231 Special issue on "Program Generation, Optimization, and Platform Adaptation".

Gokhberg, A. & Fichtner, A., 2016. Full-waveform inversion on heterogeneous HPC systems, *Comput. Geosci.,* **89**, 260–268.

Gualtieri, L., Stutzmann, É., Farra, V., Capdeville, Y., Schimmel, M., Ardhuin, F. & Morelli, A., 2014. Modelling the ocean site effect on seismic noise body waves, *Geophys. J. Int.,* **197**(2), 1096–1106.

Guennebaud, G. *et al.*, 2010. *Eigen v3,* Available at: http://eigen.tuxfamily.org, last accessed 11 December 2018.

Haugland, S.M., Ritsema, J., Kaneshima, S. & Thorne, M.S., 2017. Estimate of the rigidity of eclogite in the lower mantle from waveform modeling of broadband *S*-to-*P* wave conversions, *Geophys. Res. Lett.,* **44**, 11 778–11 784

Hestholm, S. & Ruud, B., 1994. 2D finite-difference elastic wave modelling including surface topography, *Geophys. Prospect.,* **42**(5), 371–390.

Hosseini, K. & Sigloch, K., 2015. Multifrequency measurements of core-diffracted *P* waves (Pdiff) for global waveform tomography, *Geophys. J. Int.,* **203**(1), 506–521.

Hosseini, K. & Sigloch, K., 2017. ObspyDMT: a Python toolbox for retrieving and processing large seismological data sets, *Solid Earth,* **8**(5), 1047–1070.

Hosseini, K., Matthews, K.J., Sigloch, K., Shephard, G.E., Domeier, M. & Tsekhmistrenko, M., 2018. SubMachine: web-based tools for exploring seismic tomography and other models of Earth's deep interior, *Geochem. Geophys. Geosyst.,* **19**, 1464–1483.

Hwang, Y.K., Ritsema, J., van Keken, P.E., Goes, S. & Styles, E., 2011. Wavefront healing renders deep plumes seismically invisible, *Geophys. J. Int.,* **187**(1), 273–277.

Igel, H., 2017. *Computational Seismology: A Practical Introduction,* Oxford Univ. Press.

Igel, H. & Weber, M., 1995. SH-wave propagation in the whole mantle using high-order finite differences, *Geophys. Res. Lett.,* **22**(6), 731–734.

Igel, H., Nissen-Meyer, T. & Jahnke, G., 2002. Wave propagation in 3D spherical sections: effects of subduction zones, *Physics Earth planet. Inter.,* **132**(1), 219–234.

Ishii, M. & Tromp, J., 2004. Constraining large-scale mantle heterogeneity using mantle and inner-core sensitive normal modes, *Phys. Earth planet. Inter.,* **146**(1), 113–124.

Jenkins, J., Deuss, A. & Cottaar, S., 2017. Converted phases from sharp 1000 km depth mid-mantle heterogeneity beneath Western Europe, *Earth planet. Sci. Lett.,* **459**, 196–207.

Karypis, G. & Kumar, V., 1998. A fast and high quality multilevel scheme for partitioning irregular graphs, *SIAM J. Sci. Comput.,* **20**(1), 359–392.

Kawai, K., Takeuchi, N. & Geller, R.J., 2006. Complete synthetic seismograms up to 2 Hz for transversely isotropic spherically symmetric media, *Geophys. J. Int.,* **164**(2), 411–424.

Komatitsch, D. & Tromp, J., 2002a. Spectral-element simulations of global seismic wave propagation – I. validation, *Geophys. J. Int.,* **149**(2), 390–412.

Komatitsch, D. & Tromp, J., 2002b. Spectral-element simulations of global seismic wave propagation – II. three-dimensional models, oceans, rotation and self-gravitation, *Geophys. J. Int.,* **150**(1), 303–318.

Komatitsch, D., Coutel, F. & Mora, P., 1996. Tensorial formulation of the wave equation for modelling curved interfaces, *Geophys. J. Int.,* **127**(1), 156–168.

Komatitsch, D., Erlebacher, G., Göddeke, D. & Michéa, D., 2010. High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster, *J. Comput. Phys.,* **229**(20), 7692–7714.

Korenaga, J., 2015. Constraining the geometries of small-scale heterogeneities: a case study from the Mariana region, *J. geophys. Res.,* **120**(11), 7830–7851.

Kowalik, Z., Knight, W., Logan, T. & Whitmore, P., 2005. Numerical modeling of the global tsunami: Indonesian tsunami of 26 December 2004, *Sci. Tsunami Hazards,* **23**(1), 40–56.

Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C. & Wassermann, J., 2015. ObsPy: a bridge for seismology into the scientific python ecosystem, *Comput. Sci. Discovery,* **8**(1), 014003.

Krischer, L., Hutko, A.R., van Driel, M., Stähler, S., Bahavar, M., Trabant, C. & Nissen-Meyer, T., 2017. On-demand custom broadband synthetic seismograms, *Seismol. Res. Lett.,* **88**, 1127–1140.

Kristek, J. & Moczo, P., 2003. Seismic-wave propagation in viscoelastic media with material discontinuities: a 3D fourth-order staggered-grid finite-difference modeling, *Bull. seism. Soc. Am.,* **93**(5), 2273–2280.

Kristeková, M., Kristek, J. & Moczo, P., 2009. Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals, *Geophys. J. Int.,* **178**(2), 813–825.

Kustowski, B., Ekström, G. & Dziewoński, A., 2008. Anisotropic shear-wave velocity structure of the Earth's mantle: a global model, *J. geophys. Res.,* **113**(B6), doi:10.1029/2007JB005169.

Laske, G., Masters, G., Ma, Z. & Pasyanos, M., 2013. Update on CRUST 1.0 - a 1-degree global model of Earth's crust, in *Geophysical Research Abstracts, EGU General Assembly,* Vol. **15**, Eur. Geophys. Un., Vienna, Austria, p. 2658.

Leng, K., Nissen-Meyer, T. & van Driel, M., 2016. Efficient global wave propagation adapted to 3-D structural complexity: a pseudospectral/spectral-element approach, *Geophys. J. Int.,* **207**(3), 1700–1721.

Maguire, R., Ritsema, J., van Keken, P.E., Fichtner, A. & Goes, S., 2016. *P*-and *S*-wave delays caused by thermal plumes, *Geophys. J. Int.,* **206**(2), 1169–1178.

Masson, Y., Cupillard, P., Capdeville, Y. & Romanowicz, B., 2014. On the numerical implementation of time-reversal mirrors for tomographic imaging, *Geophys. J. Int.,* **196**, 1580–1599.

Moczo, P., Kristek, J., Galis, M., Chaljub, E. & Etienne, V., 2011. 3-D finite-difference, finite-element, discontinuous-galerkin and spectral-element schemes analysed for their accuracy with respect to P-wave to S-wave speed ratio, *Geophys. J. Int.,* **187**(3), 1645–1667.

Monteiller, V., Chevrot, S., Komatitsch, D. & Wang, Y., 2015. Three-dimensional full waveform inversion of short-period teleseismic wavefields based upon the SEM–DSM hybrid method, *Geophys. J. Int.,* **202**(2), 811–827.

Nissen-Meyer, T., Dahlen, F. & Fournier, A., 2007a. Spherical-earth fréchet sensitivity kernels, *Geophys. J. Int.,* **168**(3), 1051–1066.

Nissen-Meyer, T., Fournier, A. & Dahlen, F.A., 2007b. A two-dimensional spectral-element method for computing spherical-earth seismograms – I. moment-tensor source, *Geophys. J. Int.,* **168**(3), 1067–1092.

Nissen-Meyer, T., Fournier, A. & Dahlen, F., 2008. A 2-d spectral-element method for computing spherical-earth seismograms – II. waves in solid–fluid media, *Geophys. J. Int.,* **174**(3), 873–888.

Nissen-Meyer, T., van Driel, M., Stähler, S., Hosseini, K., Hempel, S., Auer, L., Colombi, A. & Fournier, A., 2014. AxiSEM: broadband 3-D seismic wavefields in axisymmetric media, *Solid Earth,* **5**(1), 425–445.

Nolet, G., 2008. A breviary of seismic tomography, *Imaging Inter,* 234, chapter 13.

Pelties, C., Puente, J., Ampuero, J.-P., Brietzke, G.B. & Käser, M., 2012. Three-dimensional dynamic rupture simulation with a high-order discontinuous Galerkin method on unstructured tetrahedral meshes, *J. geophys. Res.,* **117**(B2), doi:10.1029/2011JB008857.

Peter, D. *et al.*, 2011. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes, *Geophys. J. Int.,* **186**(2), 721–739.

Rietmann, M. *et al.*, 2012. Forward and adjoint simulations of seismic wave propagation on emerging large-scale GPU architectures, in *Proc. Int. Conf. High Perform. Computing, Networking, Storage and Analysis,* p. 38, IEEE Computer Society Press, Washington DC.

Rietmann, M., Peter, D., Schenk, O., Uçar, B. & Grote, M., 2015. Load-balanced local time stepping for large-scale wave propagation, in *Parallel and Distributed Processing Symposium (IPDPS), 2015 IEEE International,* pp. 925–935, IEEE, Washington DC.

Ritsema, J. & Van Heijst, H., 1999. Seismic imaging of structural heterogeneity in Earth's mantle: evidence for large-scale mantle flow., *Sci. Progr.,* **83**, 243–259.

Ritsema, J., Deuss, A., Van Heijst, H. & Woodhouse, J., 2011. S40RTS: a degree-40 shear-velocity model for the mantle from new rayleigh wave dispersion, teleseismic traveltime and normal-mode splitting function measurements, *Geophys. J. Int.,* **184**(3), 1223–1236.

Romanowicz, B., 2003. 3D structure of the Earth's lower mantle, *C. R. Geosci.,* **335**(1), 23–35.

Soldati, G., Koelemeijer, P., Boschi, L. & Deuss, A., 2013. Constraints on core-mantle boundary topography from normal mode splitting, *Geochem. Geophys. Geosyst.,* **14**(5), 1333–1342.

Stähler, S.C., Sigloch, K. & Nissen-Meyer, T., 2012. Triplicated *P*-wave measurements for waveform tomography of the mantle transition zone, *Solid Earth,* **3**(2), 339–354.

Stevens, J.L., Baker, G.E., Cook, R.W., Gerald, L., Berger, L.P. & Day, S.M., 2001. Empirical and numerical modeling of T-phase propagation from ocean to land, in *Monitoring the Comprehensive Nuclear-Test-Ban Treaty: Hydroacoustics,* pp. 531–565, Springer.

Su, W.-J. & Dziewonski, A.M., 1991. Predominance of long-wavelength heterogeneity in the mantle, *Nature,* **352**(6331), 121–126.

Takeuchi, N., 2005. Finite boundary perturbation theory for the elastic equation of motion, *Geophys. J. Int.,* **160**(3), 1044–1058.

Tanimoto, T., 2010. Equivalent forces for colliding ocean waves, *Geophys. J. Int.,* **181**(1), 468–478.

Tape, C., Liu, Q., Maggi, A. & Tromp, J., 2009. Adjoint tomography of the southern California crust, *Science,* **325**(5943), 988–992.

Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation,* SIAM.

Tarrass, I., Giraud, L. & Thore, P., 2011. New curvilinear scheme for elastic wave propagation in presence of curved topography, *Geophys. Prospect.,* **59**(5), 889–906.

Thorne, M.S., Garnero, E.J., Jahnke, G., Igel, H. & McNamara, A.K., 2013. Mega ultra low velocity zone and mantle flow, *Earth planet. Sci. Lett.,* **364**, 59–67.

Toyokuni, G. & Takenaka, H., 2006. FDM computation of seismic wavefield for an axisymmetric earth with a moment tensor point source, *Earth Planets Space,* **58**(8), e29–e32.

Trabant, C., Hutko, A.R., Bahavar, M., Karstens, R., Ahern, T. & Aster, R., 2012. Data products at the IRIS DMC: stepping stones for research and other applications, *Seismol. Res. Lett.,* **83**(5), 846–854.

Tromp, J., Komattisch, D. & Liu, Q., 2008. Spectral-element and adjoint methods in seismology, *Commun. Comput. Phys.,* **3**(1), 1–32.

Tsuboi, S., Ando, K., Miyoshi, T., Peter, D., Komatitsch, D. & Tromp, J., 2016. A 1.8 trillion degrees-of-freedom, 1.24 petaflops global seismic wave simulation on the K computer, *Int. J. High Perform. Comput. Appl.,* **30**(4), 411–422.

van Driel, M. & Nissen-Meyer, T., 2014. Optimized viscoelastic wave propagation for weakly dissipative media, *Geophys. J. Int.,* **199**(2), 1078–1093.

van Driel, M., Krischer, L., Stähler, S., Hosseini, K. & Nissen-Meyer, T., 2015. Instaseis: instant global seismograms based on a broadband waveform database, *Solid Earth Discuss.,* **7**(1), 957–1005.

Woodhouse, J., 1976. On Rayleigh's principle, *Geophys. J. Int.,* **46**(1), 11–22.

Woodhouse, J.H. & Dziewonski, A.M., 1984. Mapping the upper mantle: three-dimensional modeling of Earth structure by inversion of seismic waveforms, *J. geophys. Res.,* **89**(B7), 5953–5986.

Zaroli, C., Lambotte, S. & Lévêque, J.-J., 2015. Joint inversion of normal-mode and finite-frequency s-wave data using an irregular tomographic grid, *Mon. Not. R. astr. Soc.,* **203**(3), 1665–1681.

Zhang, W. & Chen, X., 2006. Traction image method for irregular free surface boundaries in finite difference seismic wave simulation, *Geophys. J. Int.,* **167**(1), 337–353.

Zhang, W., Zhang, Z. & Chen, X., 2012. Three-dimensional elastic wave numerical modelling in the presence of surface topography by a collocated-grid finite-difference method on curvilinear grids, *Geophys. J. Int.,* **190**(1), 358–378.

Zhou, Y., Dahlen, F., Nolet, G. & Laske, G., 2005. Finite-frequency effects in global surface-wave tomography, *Geophys. J. Int.,* **163**(3), 1087–1111.

Zhu, H., Bozdağ, E. & Tromp, J., 2015. Seismic structure of the European upper mantle based on adjoint tomography, *Geophys. J. Int.,* **201**(1), 18–52.

APPENDIX A: ASPHERICAL FLUID DOMAIN

An earth model may contain various solid–fluid discontinuities, collectively denoted Σ, as shown in Fig. 1. Following Komatitsch & Tromp (2002a) and Chaljub & Valette (2004), 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

$$
\int_{\Omega_s} \rho \, \partial_t^2 \mathbf{u} \cdot \mathbf{w} \, d^3\mathbf{r} + \int_{\Omega_s} \nabla\mathbf{u} : \mathbf{C} : \nabla\mathbf{w} \, d^3\mathbf{r} \\
= \int_{\Omega_s} \mathbf{f} \cdot \mathbf{w} \, d^3\mathbf{r} - \int_{\Sigma} \partial_t^2 \chi \, \hat{\mathbf{n}} \cdot \mathbf{w} \, d^2\mathbf{r}, \tag{A1}
$$

in the solid domain Ω_s, and

$$
\int_{\Omega_f} \frac{1}{\kappa} \partial_t^2 \chi \, w \, d^3\mathbf{r} + \int_{\Omega_f} \frac{1}{\rho} \nabla\chi \cdot \mathbf{I} \cdot \nabla w \, d^3\mathbf{r} = \int_{\Sigma} \mathbf{u} \cdot w \hat{\mathbf{n}} \, d^2\mathbf{r}, \tag{A2}
$$

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. (A2) 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. (7). For the

solid domain, eq. (A1) becomes

$$\int_{\tilde{\Omega}_s} \tilde{\rho}\partial_t^2 \tilde{\mathbf{u}} \cdot \tilde{\mathbf{w}}\, d^3\mathbf{r} + \int_{\tilde{\Omega}_s} \nabla\tilde{\mathbf{u}} : \tilde{\mathbf{C}} : \nabla\tilde{\mathbf{w}}\, d^3\mathbf{r}$$
$$= \int_{\tilde{\Omega}_s} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}}\, d^3\mathbf{r} - \int_{\tilde{\Sigma}} \partial_t^2 \tilde{\chi} \tilde{\mathbf{n}} \cdot \tilde{\mathbf{w}}\, d^2\mathbf{r}, \tag{A3}$$

where $\tilde{\rho}$, $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{f}}$ are computed by eqs (9)–(11), 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

$$\tilde{\mathbf{n}}(\mathbf{r}) = \mathbf{J}^{-T}(\tau, \mathbf{r}) \cdot \hat{\mathbf{n}}(\xi(\mathbf{r})) \det \mathbf{J}(\tau, \mathbf{r}). \tag{A4}$$

Similarly, for the fluid domain, eq. (A2) becomes

$$\int_{\tilde{\Omega}_f} \frac{1}{\tilde{\kappa}} \partial_t^2 \tilde{\chi} \tilde{w}\, d^3\mathbf{r} + \int_{\tilde{\Omega}_f} \frac{1}{\tilde{\rho}} \nabla\tilde{\chi} \cdot \tilde{\mathbf{I}} \cdot \nabla\tilde{w}\, d^3\mathbf{r} = \int_{\tilde{\Sigma}} \tilde{\mathbf{u}} \cdot \tilde{w}\tilde{\mathbf{n}}\, d^2\mathbf{r}, \tag{A5}$$

where

$$\tilde{\kappa}(\mathbf{r}) = \kappa(\xi(\mathbf{r})) \det{}^{-1}\mathbf{J}(\tau, \mathbf{r}), \tag{A6}$$

$$\tilde{\rho}(\mathbf{r}) = \rho(\xi(\mathbf{r})) \det{}^{-1}\mathbf{J}(\tau, \mathbf{r}), \tag{A7}$$

$$\tilde{\mathbf{I}}(\mathbf{r}) = \mathbf{J}^{-1}(\tau, \mathbf{r}) \cdot \mathbf{J}^{-T}(\tau, \mathbf{r}). \tag{A8}$$

Again, the 3-D weak forms for an aspherical solid–fluid Earth, eqs (A3) and (A5), are formally identical to those for a spherical solid–fluid Earth, eqs (A1) and (A2), 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.

APPENDIX B: SOURCE NEAR UNDULATING INTERFACES

For a point source moment tensor, denoted $\mathbf{M}$, the right-hand side of the 3-D weak form, eq. (1), 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), can be written as (Leng *et al.* 2016)

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

$$2\pi \int_D f_i^0 w_i^0 s\, ds dz = \left(M_{xx} + M_{yy}\right) w_{s,s}^0 + M_{zz} w_{z,z}^0, \tag{B1}$$

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

$$2\pi \int_D f_i^{-1} w_i^1 s\, ds dz = \frac{M_{xz} - \text{î} M_{yz}}{2} \left(w_{s,z}^1 + 2w_{z,s}^1 + \text{î} w_{\phi,z}^1\right), \tag{B2}$$

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

$$2\pi \int_D f_i^{-2} w_i^2 s\, ds dz = \left(\frac{M_{yy} - M_{xx}}{2} + \text{î} M_{xy}\right) \left(w_{s,s}^2 + \text{î} w_{\phi,s}^2\right), \tag{B3}$$

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. (7) will become

$$\int_{\tilde{\Omega}} \tilde{\mathbf{f}} \cdot \tilde{\mathbf{w}}\, d^3\mathbf{r} = \mathbf{M} : \left(\mathbf{J}^{-T}(\tau_s, \mathbf{r}_s) \cdot \nabla\tilde{\mathbf{w}}(\mathbf{r}_s)\right). \tag{B4}$$

Accordingly, the above closed forms, eqs (B1)–(B3), should be, respectively, generalized as

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

$$2\pi \int_{\tilde{D}} \tilde{f}_i^0 \tilde{w}_i^0 s\, ds dz = \left(1 + \frac{\tau_s}{z_s}\right) \left(M_{xx} + M_{yy}\right) \tilde{w}_{s,s}^0 + M_{zz} \tilde{w}_{z,z}^0$$
$$+ \frac{\partial\tau}{\partial x} M_{xz} \tilde{w}_{z,z}^0 + \frac{\partial\tau}{\partial y} M_{yz} \tilde{w}_{z,z}^0 + \frac{\partial\tau}{\partial z} M_{zz} \tilde{w}_{z,z}^0, \tag{B5}$$

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

$$2\pi \int_{\tilde{D}} \tilde{f}_i^{-1} \tilde{w}_i^1 s\, ds dz = \left(1 + \frac{\tau_s}{z_s}\right) \left(M_{xz} - \text{î} M_{yz}\right) \tilde{w}_{z,s}^1$$
$$+ \frac{M_{xz} - \text{î} M_{yz}}{2} \left(\tilde{w}_{s,z}^1 + \text{î}\tilde{w}_{\phi,z}^1\right)$$
$$+ \frac{\partial\tau}{\partial x} \frac{M_{xx} - \text{î} M_{xy}}{2} \left(\tilde{w}_{s,z}^1 + \text{î}\tilde{w}_{\phi,z}^1\right) \tag{B6}$$
$$+ \frac{\partial\tau}{\partial y} \frac{M_{xy} - \text{î} M_{yy}}{2} \left(\tilde{w}_{s,z}^1 + \text{î}\tilde{w}_{\phi,z}^1\right)$$
$$+ \frac{\partial\tau}{\partial z} \frac{M_{xz} - \text{î} M_{yz}}{2} \left(\tilde{w}_{s,z}^1 + \text{î}\tilde{w}_{\phi,z}^1\right),$$

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

$$2\pi \int_{\tilde{D}} \tilde{f}_i^{-2} \tilde{w}_i^2 s\, ds dz = \left(1 + \frac{\tau_s}{z_s}\right) \left(\frac{M_{yy} - M_{xx}}{2} + \text{î} M_{xy}\right)$$
$$\times \left(\tilde{w}_{s,s}^2 + \text{î}\tilde{w}_{\phi,s}^2\right), \tag{B7}$$

where the derivatives of τ are also evaluated at $\mathbf{r}_s$. The above closed forms degenerate to eqs (B1)–(B3) 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.

APPENDIX C: OCEAN LOAD

When approximated as load, the ocean introduces an additional term (Komatitsch & Tromp 2002b) to the 3-D weak form, eq. (1),

$$\int_\Omega \rho\partial_t^2 \mathbf{u} \cdot \mathbf{w}\, d^3\mathbf{r} + \int_\Omega \nabla\mathbf{u} : \mathbf{C} : \nabla\mathbf{w}\, d^3\mathbf{r}$$
$$= \int_\Omega \mathbf{f} \cdot \mathbf{w}\, d^3\mathbf{r} - \int_{\partial\Omega} \rho_w h_w \left(\partial_t^2 \mathbf{u} \cdot \hat{\mathbf{n}}\right) \left(\mathbf{w} \cdot \hat{\mathbf{n}}\right) d^2\mathbf{r}, \tag{C1}$$

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. (A4), 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 parametrization to h_w, the above weak form reduces to the following algebraic system of 2-D weak forms in D, as a generalization of eq. (3),

$$\sum_{|\alpha| \leq n_u} \left(\int_D \rho^{-(\alpha+\beta)} \partial_t^2 u_i^\alpha w_i^\beta s\, ds dz + \int_D u_{i;j}^\alpha C_{ijkl}^{-(\alpha+\beta)} w_{k;l}^\beta s\, ds dz\right)$$
$$= \int_D f_i^{-\beta} w_i^\beta s\, ds dz - \rho_w \sum_{|\alpha| \leq n_u} \int_{\partial D} h_w^{-(\alpha+\beta)} \partial_t^2 u_i^\alpha w_j^\beta \hat{n}_i \hat{n}_j s\, ds dz. \tag{C2}$$