

AxiSEM: Broadband 3D seismic wavefields in axisymmetric media

Tarje Nissen-Meyer¹, Martin van Driel², Simon C. Stähler³, Kasra Hosseini³, Stefanie Hempel⁴, Ludwig Auer²,
Andrea Colombi⁵, and Alexandre Fournier⁶

¹Dept. of Earth Sciences, University of Oxford, South Parks Road, Oxford OX1 3AN, UK

²Institute of Geophysics, ETH Zurich, Sonneggstrasse 5, 8092 Zurich, Switzerland

³Dept. of Earth Sciences, Ludwig-Maximilians-Universität München, Theresienstrasse 41, 80333 München, Germany

⁴Institut für Geophysik, Universität Münster, Correnzstrasse 24, 48149 Münster, Germany

⁵ISTerre, Université Joseph Fourier, Grenoble, BP 53 38041 Grenoble CEDEX 9, France

⁶Institut de Physique du Globe, Sorbonne Paris Cité, Université Paris Diderot, UMR 7154 CNRS, F-75005 Paris, France

Correspondence to: Tarje Nissen-Meyer
(tarjen@earth.ox.ac.uk)

Abstract. We present a methodology to compute 3D global seismic wavefields for realistic earthquake sources in viscoelastic anisotropic media, covering applications across the observable seismic frequency band with moderate computational resources. This is accommodated by mandating axisymmetric background models which allow for a multipole expansion such that only a 2D computational domain is needed, whereas the azimuthal third dimension is computed analytically on-the-fly. This dimensional collapse opens doors for storing space-time wavefields on disk which can be used to compute Fréchet sensitivity kernels for waveform tomography. We use the corresponding publicly available open-source spectral-element code AxiSEM (www.axisem.info), demonstrate its excellent scalability on supercomputers, a diverse range of applications ranging from normal modes to small-scale lowermost mantle structures, tomographic models, comparison to observed data, and discuss further avenues to pursue with this methodology.

1 Introduction

Seismology currently enjoys transformative progress upon a simultaneous surge in instrumentation, software, and hardware. The dawn of high-performance computing and sophisticated numerical techniques to address seismic wave propagation in a physically robust and realistic manner has enabled seismologists to capture relevant physics of wave propagation in the seismic far-field, and resolve structures for which direct comparisons to waveform data are feasible. This in turn opens doors to incorporating full waveforms into inver-

sion algorithms, using, for instance, adjoint methods in conjunction with 3D wave propagation (e.g. Tromp et al., 2005). Traditionally, global seismic tomography (consult Rawlinson et al., 2010, for a comprehensive summary) has been based upon ray theory utilizing traveltimes, rather than full waveforms. Indeed, phase delays relate to wavespeed variations in a more robust manner than amplitude information. So why would numerical methods, within the realm of tomography at least, strive to capture the entire waveform? Most modern measurements of “traveltimes” such as cross-correlation (Nolet, 2008), time-frequency phase delays (Fichtner et al., 2008), or instantaneous phase (Bozdag et al., 2011) are based on waveforms, and therefore necessitate full wavefield modeling. Moreover, high-frequency waveform modelling (e.g. Thorne et al., 2007) is often employed to fit smaller-scale heterogeneities. Such studies are subject to significant trade-offs especially if secondary measurements such as traveltimes were used. Thirdly, accurate computation of the gradient of measurements with respect to model variations, often termed Fréchet derivative, requires the convolution of a forward-propagating wavefield with a backward-, or adjoint wavefield, both of which need to be sampled in 3D space and time (Nissen-Meyer et al., 2007a). For the purpose of this paper, let us postulate a desire for a method to deliver

1. **3D wavefields** for realistic sources and structures,
2. across the **observable frequency band**,
3. at a **reasonable computational cost** for tomography.

We will now delve into some of these issues in more detail, and present our compromise for a solution which covers a significant, realistic and relevant fraction of these aspirations.

1.1 Effective Earth models and data

Spherically symmetric models are widely established as a common basis for Earth properties not only in seismology, but also as a bridge to mineral physics and geodynamics. This popularity stems not only from the relative simplicity in modeling 1D structures, but largely from the fact that such laterally averaged models represent and fit a large majority of seismic (traveltime) data, as has been established since the traveltime tables by Jeffreys & Bullen (1940) and subsequent models such as PREM (Dziewonski & Anderson, 1981), IASP91 (Kennett & Engdahl, 1991), and ak135 (Kennett et al., 1995). Our understanding and interpretation of the Earth’s interior has come a long way from the detection of its radial structure, and has been significantly fueled by means of seismic tomography (Rawlinson et al., 2010). At the global scale, 3D tomographic models usually amount to a few percent wavespeed perturbation from spherically symmetric models (Becker & Boschi, 2002), and behave close to linear in seismic traveltimes (Mercerat & Nolet, 2013).

Global tomographic models exhibit considerable agreement up to spherical harmonic degree 5 (Becker & Boschi, 2002; Auer et al., 2013), that is, for very large-scale structures, but often diverge at smaller length scales due to shortcomings such as insufficient data coverage and modeling. In this multi-scale context, it is important to remember that any discrete Earth model used in a numerical method is inherently upscaled (either by its own nature, or by porting to the discrete mesh), and at best a blurred rendition of reality. The challenge lies in tying the background model to both the desired frequency range and type of measurement extracted from the wavefield in a feasible, realistic manner.

Clearly, utilizing a maximal amount of broadband data is as desirable as capturing complexities in structure and wave physics. Even in times of a surge in data acquisition, source-receiver geometries are still largely controlled by continents, tectonic boundaries and the Northern hemisphere. It thus seems desirable to seek a compromise for modeling between broad frequency ranges, realistic effective Earth models, while exploiting a maximal amount of usable data.

1.2 Numerical wave propagation

Unlike disciplines subject to more complex, non-linear physical systems such as fluid dynamics, the availability of mature, comprehensive seismic wave-propagation codes such as SPECSEM3D_GLOBE (e.g. Komatitsch & Tromp, 2002b) seems to have resolved most challenges in capturing the underlying physics of wave propagation. Instead, it appears to shift attention for achieving realistic wave propagation simulations towards

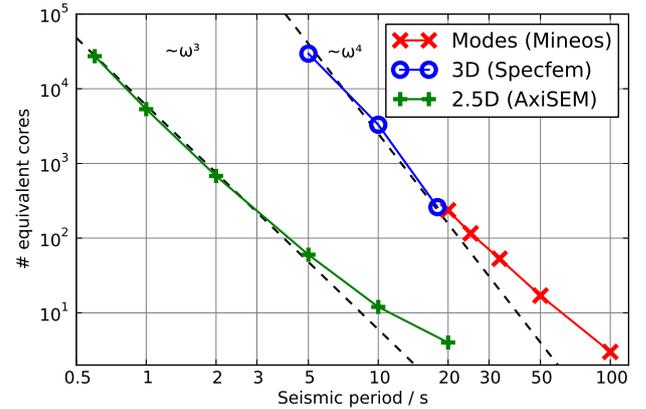


Fig. 1. The cost of global wavefield simulations in “real-time” (i.e. seismogram length equals wallclock simulation time) for different methods. Each data point is based on actual simulation times, and gives as a result the number of processors needed to achieve this, assuming perfect scalability. The cost of normal-mode solutions (mineos, available from geodynamics.org) and wave propagation in 3D domains (SPECSEM3D_GLOBE, geodynamics.org) scales with the seismic period to the fourth power, whereas the axisymmetric method (AxiSEM) scales to the third power. We calculate the cost estimation upon saving 10^6 spatial points, a moderate task to compute wavefields. This is especially noteworthy for the mode solution whose cost scales directly with the number of saved spatial points.

- feasible choices for source and structure,
- usability of the method,
- availability of computational resources.

While all of these issues are common to any numerical method and not easily resolvable, the latter point is especially stringent. Worse still, for global seismology the prohibitive cost will remain a dominant limitation on the maximally resolved frequencies and exponentially increasing number of usable data with millions of recorded waveforms (IRIS annual report 2011, www.iris.edu) for years to come. Full 3D models in spherical geometry can be incorporated by spherical finite-differences (Igel et al., 2002), or spectral-element methods (Komatitsch & Tromp, 2002a; Chaljub et al., 2007). As the computational cost scales with the frequency to the fourth power (three space, one time dimension), such comprehensive methods are still extraordinarily expensive for global-scale wave propagation at high resolution, and certainly more so if large numbers of simulations are needed as in most cases of geophysical interest. Alternative attempts such as high-order expansions of the Born series (Takeuchi et al., 2000) is similarly prohibitive for complex media and high frequencies. This is unfortunate, as especially the domain of seismic periods below 10 seconds (see Fig. 1) offers both a wealth of seismic data and a resolution harboring many open geophysical questions. The plot in Fig. 1 is

an approximate, order-of-magnitude estimation of computational cost, and minor algorithmic optimization does not affect this overarching trend drastically. The slightly increased cost for AxiSEM above 10 seconds represents the fact that the thin crustal layers at these long periods start to dominate the smallest element size and thus increase the relative cost due to this geometric constraint on the global timestep. This is not seen in SPECFEM3D_GLOBE, since intra-crustal layers are not explicitly meshed. Further commentary on Fig. 1) is given in Section 2.7.

Considering desirable features for the inversion such as comprehensive model-sampling, uncertainty analysis, or probabilistic approaches, this represents not only a formidable challenge, but is essentially not computable even with most optimistic estimates of the evolution of computation on a decadal time-scale, especially in 3D. Several strategies of speeding up numerical methods exist, focused on either the physical system or the implementation. Code optimization may exploit dedicated hardware infrastructures such as GPUs (Rietmann et al., 2012), or algorithmic tasks such as tensor-vector products (Nissen-Meyer et al., 2007b), irregular meshing (Zhu et al., 2009) or local time-stepping. These approaches usually lead to a performance speedup of about 2-3 in total CPU time. Physics-based approximations often limit the frequency range either on the high end (as implicitly done due to prohibitive cost in 3D methods) or lower end (ray theory). Additionally, we commonly find cost reductions related to reduced dimensionality (e.g., 2D, Zhu et al. (2009)), rheology (e.g. acoustic wave propagation), or structural complexity by means of homogenization (Capdeville et al., 2013). Such approximations can lead to orders of magnitude faster codes, but need to be chosen carefully depending on each application.

1.3 3D waves in axisymmetric media

Several methods have been developed to effectively accommodate various levels of complexity in background structures. For spherically symmetric Earth models, normal-mode summation (Dahlen & Tromp, 1998) elegantly tackles the grave end of the spectrum including such effects as gravity and rotation (Dahlen, 1968). For higher frequencies, the direct-solution method (Kawai et al., 2006), GEMINI (Friederich & Dalkolmo, 1995), or Yspec (Al-Attar & Woodhouse, 2008) have proven efficient in delivering accurate seismograms. While in principle doable, all of these methods become computationally expensive if an entire wavefield is needed as for sensitivity kernels (Nissen-Meyer et al., 2007a), and do not allow for lateral heterogeneities. Axisymmetric finite difference methods (Toyokuni & Takenaka, 2006; Jahnke et al., 2008) may accommodate this effectively, but suffer various shortcomings such as approximate sources, lack of fluid domains and anisotropy (Jahnke et al., 2008), and high dispersion errors for large propagation distances of interface-sensitive phases such as surface or diffracted

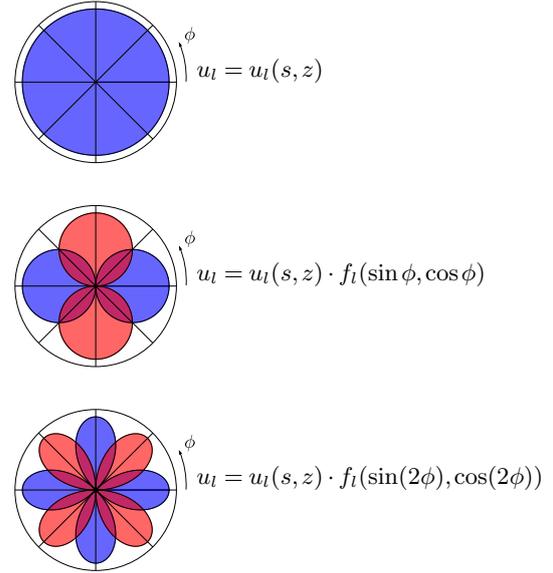


Fig. 2. Radiation patterns for monopole (top), dipole (middle), and quadrupole angular orders of the respective moment tensor elements. The azimuthal radiation patterns encapsulated by f_l depend on multipole order m as well as component l , that is, no summation is implied by the above products.

waves. However, recent advances include a full moment tensor, attenuation, and the Earth's center (Toyokuni & Takenaka, 2012).

The purpose of this paper is to introduce the axisymmetric spectral-element implementation **AxiSEM** as a new and publicly available, production-ready method and code for global wave propagation, which taps into parameter regimes that have been previously unavailable at similar computational cost. We motivate the relevance of these parameter regimes by various examples and present ideas for further extensions and applications. Exploitation of moment-tensor source and single-force radiation patterns allow the computational domain to be collapsed to a 2D semi-disk, and the azimuthal third dimension is computed analytically. Radiation pattern symmetries require all sources to be located along the axis, and lateral heterogeneities are translated into a 2.5-dimensional torus-like structure. Due to the dimensional reduction, global wave propagation at typical seismic periods can be tackled serially on workstations. Novel features in this manuscript with respect to the methodology already described in Nissen-Meyer et al. (2007b, 2008) include 2D parallelization, scalability to > 8000 cores, benchmarks at 1Hz and for normal modes, extensions to visco-elastic anisotropic media, fluid spheres, finite sources, axisymmetric structures, tomographic models, comparison to data, generic post-processing for arbitrary source-receiver settings, sensitivity kernels, and availability as an open-source code.

This paper is organized as follows. A methodological chapter briefly summarizes the mathematical background of our approach, delegating more details to previous publications, and focusing instead on practical matters such as scalability, runtime requirements, I/O, and code availability. Chapter 3 describes those source types that may be simulated with AxiSEM, ranging from moment-tensors, single forces, to finite faults and stochastic sources. Chapter 4 similarly describes a range of background models to be discretized in AxiSEM, including anisotropy, intrinsic attenuation, classical 1D Earth models, solar models, small-scale 2.5D heterogeneities and tomographic cross-sections, random media and explicit mesh representations of the crust and oceans (oceans are not part of the current code as of April 2014). Chapter 5 shows simulation results for a selection of the previously mentioned ranges of applicability, covering the entire seismic frequency spectrum, 3D wavefield visualization, lowermost mantle structures, simulations for 2.5D slices through a tomography model, comparison with observed data from lowermost mantle and core phases, and sensitivity kernels. A concluding chapter discusses the general applicability, limitations and an outlook on future developments.

AxiSEM is a mature methodology and code, able to address a number of intriguing scientific questions. As should be commonly known per any software implementation, the level of automatism for the applications listed here is diverse, and readers should refer to the manual of the most recent release version for up-to-date features of the code.

2 Methodology

The mathematical foundation and validation of spherically symmetric, solid-fluid lower-frequency settings is detailed in previous publications (Nissen-Meyer et al., 2007b, 2008). In this section, we only sketch key methodological concepts, while focusing on new additions and practical matters related to usability, functionality and applicability. Our approach accurately simulates 3D wavefields in axisymmetric Earth models, and distinguishes itself by

1. decreasing the computational costs by orders of magnitude compared to 3D method by running in 2D,
2. taking no limiting assumptions about wave-propagation physics (except for very long-period effects such as rotation (see section 4.7)) or kinematic earthquake radiation.

It therefore falls in between traditional end members that are typically optimized for either end of the frequency spectrum (e.g. ray theory, normal-mode summation) and 3D modeling, by not compromising on essential wave-propagation physics or the coverage of the entire recorded frequency band between 0.001-1Hz. The efficiency gain is grounded upon assuming axisymmetric background models, which reduces the

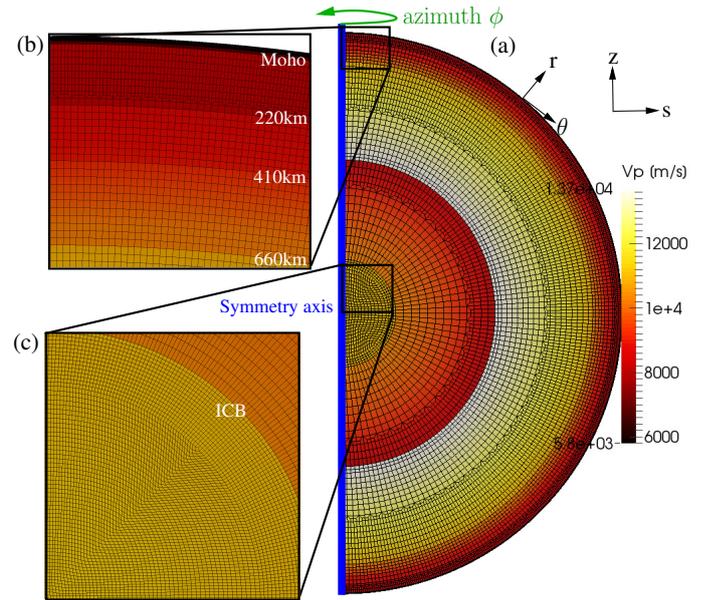


Fig. 3. The 2D computational domain D upon which the collapsed numerical system operates with a symmetry axis (blue). The method solves the three-dimensional equations of motion, but allows for an analytical representation within the azimuthal dimension (green). Sources and structure therefore obey axisymmetry with respect to the axis. Colors denote compressional velocities of the PREM model (Dziewonski & Anderson, 1981), and black lines an elemental mesh for a seismic period of 20s. Zoom panels show a higher-resolution version (5s) with upper-mantle discontinuities honored by the mesh (b), as well as the rotated-coordinate meshing below the inner-core boundary (ICB) (Nissen-Meyer et al., 2008).

numerical cost to a 2D domain whereas the third dimension is tackled analytically. We shall forego detailed treatment of classical spectral-element methods to highlight the peculiarities associated with this axisymmetric setting.

2.1 Equations of motion

The 3D integral (weak-form) elastodynamic equations of motion in the solid Earth \oplus read

$$\underbrace{\int_{\oplus} \rho \mathbf{w} \cdot \partial_t^2 \mathbf{u} d^3 \mathbf{x}}_{\text{mass term: } \mathbf{M}(\mathbf{u})} + \underbrace{\int_{\oplus} \nabla \mathbf{w} : \mathbf{C} : \nabla \mathbf{u} d^3 \mathbf{x}}_{\text{stiffness term: } \mathbf{K}(\mathbf{u})} = \underbrace{\int_{\oplus} \mathbf{w} \cdot \mathbf{f} d^3 \mathbf{x}}_{\text{source term: } \mathbf{F}(\mathbf{u})} \quad (1)$$

where \mathbf{u} is the sought displacement vector, \mathbf{w} a suitably chosen test vector, \mathbf{f} the source term, ρ the mass density, and \mathbf{C} the anisotropic fourth-order elasticity tensor with 21 independent parameters (consult Nissen-Meyer et al. (2007b) for details). It may be time-dependent for intrinsic attenuation in which case the double contraction $:$ implies a convolution.

2.2 Axisymmetric dimensional collapse

As shown in Nissen-Meyer et al. (2007a), one may analytically separate radiation patterns into individual responses to each moment-tensor element M_{ij} factorized in azimuthal functions

$$\mathbf{u}_m(\mathbf{x}) = \begin{pmatrix} u_s(\tilde{\mathbf{x}}) \cos m\phi \\ u_\phi(\tilde{\mathbf{x}}) \sin m\phi \\ u_z(\tilde{\mathbf{x}}) \cos m\phi \end{pmatrix}, \quad (2)$$

where $m = 0, 1, 2$ are monopole, dipole, quadrupole radiation types, respectively (Fig. 2), and $\tilde{\mathbf{x}} = (s, z) = (r, \theta)$ spans a two-dimensional domain (Fig. 3) by cylindrical (s, ϕ, z) or spherical (r, θ, ϕ) coordinates, respectively. This relation is accurate for axisymmetry in source $\mathbf{f} = \mathbf{f}(\tilde{\mathbf{x}})$ and structure $\rho = \rho(\tilde{\mathbf{x}})$, $\mathbf{C} = \mathbf{C}(\tilde{\mathbf{x}})$. After solving the set of 2D problems, seismograms and wavefields at any location (s, ϕ, z) are obtained by multiplication with these azimuthal radiation factors in eq. (2) during the post-processing stage (Section 2.5). Conceptually, 3D integrals in \oplus over any integrand ψ that contains azimuthal dependencies such as in eq. (2) are then collapsed to 2D integrals in D as

$$\int_{\oplus} \psi(\mathbf{x}) d^3\mathbf{x} \Rightarrow \int_D \psi(\tilde{\mathbf{x}}) d^2\tilde{\mathbf{x}}, \quad (3)$$

by evaluating integration over ϕ analytically. This delivers solutions for the 3D displacement vector \mathbf{u} within a 2D computational domain (Fig. 3). Symmetry about the axis (blue in Fig. 3) mandates all structural heterogeneities away from it to adopt a torus-shaped, azimuthally invariant elongation, whereas the point source remains along the axis. Such lateral in-plane heterogeneities may prove useful for parameter studies at sufficiently high frequencies (see Section 4.4).

2.3 Spatial discretization

Finite-element based methods compute derivatives and integration upon reference coordinates. This entails a mapping $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}(\boldsymbol{\xi})$ from a generic reference coordinate frame $\boldsymbol{\xi}$ to the physical domain $\tilde{\mathbf{x}}$, represented by the Jacobian

$$\mathcal{J}(\boldsymbol{\xi}) = \left| \frac{\partial \tilde{\mathbf{x}}}{\partial \boldsymbol{\xi}} \right|, \quad (4)$$

where $|\cdot|$ is the determinant. This mapping is purely analytical for all element types of AxiSEM's automated mesh generator with the exception of the cube at the center of the sphere (Nissen-Meyer et al., 2008). We then expand the wavefield within each elemental integral upon a basis of order N (typically 4, 5, 6) as

$$\mathbf{u}(\boldsymbol{\xi}) \approx \sum_{i,j=0}^N \mathbf{u}_{ij} l_{ij}(\boldsymbol{\xi}), \quad (5)$$

with two-dimensional Lagrange polynomials l_{ij} (Nissen-Meyer et al., 2007b). Partial derivatives $\partial_{\boldsymbol{\xi}} \mathbf{u}(\boldsymbol{\xi})$ are given by

analytically differentiating l_{ij} along $\boldsymbol{\xi}$. These derivative operations are responsible for the bulk of the computational cost in typical spectral-element methods. Having performed these algebraic operations at the level of elements, these elemental contributions are gathered to define the discrete global stiffness $\mathbf{K}\mathbf{u}$ and mass terms $\mathbf{M}\mathbf{u}$. Our formulation with cylindrical coordinates leads to singularities of the type s^{-1} (Fig. 3) in the gradients at the symmetry axis. This is accommodated by a different basis compared to the interior domain, l'Hospital's rule (Fournier et al., 2005), and asymptotic expressions to accommodate boundary conditions (Nissen-Meyer et al., 2007a). By choice of either kind of basis function, the mass matrix is exactly diagonal.

2.4 Temporal discretization

Such a discrete system leads to a set of ordinary differential equations in time, which may be rearranged as

$$\ddot{\mathbf{u}}(t) = \mathbf{M}^{-1}(\mathbf{f}(t) - \mathbf{K}\mathbf{u}(t)), \quad (6)$$

This system is conveniently solved by various explicit time-evolution schemes such as second-order Newmark, or higher-order symplectic schemes (Nissen-Meyer et al., 2008) up to eighth-order accuracy. Note that for the case of solid-fluid domains, the time-stepping becomes a combined system of these two domains, and needs to be iterated appropriately across the solid-fluid interface (Chaljub & Valette, 2004).

2.5 Post-processing: Summation, rotation, filtering

Unlike most seismic wave-propagation codes, AxiSEM requires a crucial sequence of post-processing steps to retrieve the full solution, see eq. (2). While this may seem as an undesirable additional burden, it represents a high level of flexibility, leaving a maximum amount of parameter choices to this post-processing step instead of having them fixed for the bulk simulations. For instance, one does not need to decide on the source mechanism, source-time function, filtering, instrument response, and receiver components at the time of the actual simulation but can defer this to post-processing. The only necessary geophysical choices at the time of the simulation are source depth, receiver distances, maximal frequency, and background model. Fig. 4 depicts an example of an automated output from postprocessing to be read by Google Earth, containing source (red dot) and receiver (yellow pins) locations. Each receiver pin is linked to an image of the corresponding post-processed seismograms, and source information is provided as text (see Fig. 4). Section 5.3 sketches a generalization of this post-processing that fully exploits its flexibility in the context of solving the forward problem once-and-for-all, deferring the choice of the source-receiver geometry to post-processing as well.

Recasting the 3D equations of motion into a suite of 2D problems yields a system of four independent wave equa-

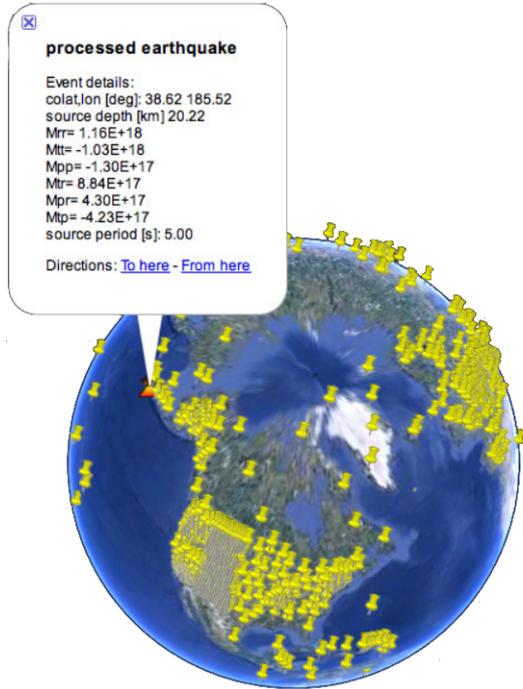


Fig. 4. A Google-Earth rendition of the source-receiver geometry used in an AxiSEM simulation. A generic output of the post-processing embedded within AxiSEM, this *kml*-file contains earthquake parameters (red dot), and actual seismogram images as links at the station locations (yellow pins).

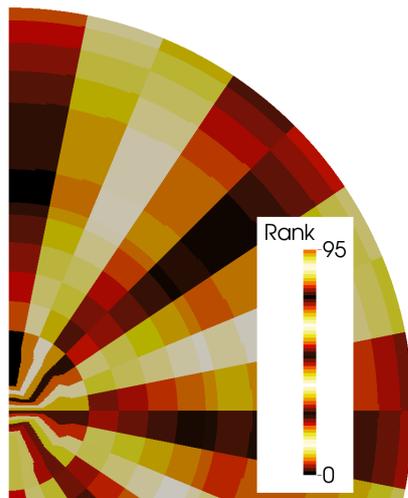


Fig. 5. A typical mesh decomposition for the PREM model running at dominant period of 9 seconds on 96 cores. Load balancing is exact, and the arbitrary permissible multiplication between number of horizontal and radial slices guarantees flexibility for adapting numerical settings to existent hardware infrastructures.

tions to represent all six independent elements of the moment tensor ($M_{rr}, M_{\phi\phi}, M_{\theta\theta}, M_{r\theta}, M_{r\phi}, M_{\theta\phi}$) separately (see Fig. 7). This collapse from six to four independent systems honors azimuthal redundancy between the dipoles $M_{r\theta} \sim M_{r\phi}$ as well as quadrupoles $(M_{\theta\theta} - M_{\phi\phi}) \sim M_{\theta\phi}$ (Nissen-Meyer et al., 2007a). Consequently, AxiSEM simulations are by construction always given for each individual element of the moment tensor. The task to sum to a full moment-tensor is described in Section 3.1. Additional features of post-processing are rotation from the pole-centric to an actual source-receiver geometry, bandpass filtering, convolution with a source-time function, rotation to arbitrary seismogram component systems, choice between displacement and velocity seismogram. A similar set of operations applies to 3D wavefield visualizations. Users may for this visualization case also specify rendering perspectives, wavefield components, 3D cuts, hypersurface extractions within post-processing. In effect, this allows for in-situ visualization and merges seismic trace analysis with visualization on-the-fly.

2.6 Parallelization

At frequencies around 1Hz, the required run-time memory (roughly 20GB) for 1 Million elements exceeds the typical memory of contemporary cluster cores. More importantly, the CPU time-to-solution becomes prohibitively lengthy if the system is simulated on a single core (although possible). We thus incorporated a generic, automated 2D domain decomposition into $2N_{\theta}N_r$ domains, where $N_{\theta}N_r$ represent positive integers for the number of latitudinal and radial slices, respectively. This guarantees the simultaneous realization of the three crucial factors for scalability: 1) a minimal amount of neighboring domains (maximally eight), 2) minimal interfaces size (i.e., length of messages), and 3) exact load balancing. The non-blocking, asynchronous message passing implementation is entirely hidden behind the computation of the stiffness term, which will be seen in the excellent scaling in the next section.

2.7 Performance & scaling

The reduction of 3D wave propagation to a 2D computational domain is reflected by the method's performance compared to 3D methods (Fig. 1). This equally holds true against methods which are extremely efficient and fast in delivering singular seismograms such as normal-mode summation or DSM, but whose computational cost depends on the amount of desired output locations. To compute sensitivity kernels for the inverse problem, one needs to save the entire space-time wavefield everywhere, and hence such dependencies become inefficient especially when moving to higher resolutions. Fig. 1 gives a flavor of the computational task along a typical range of global-scale seismic periods, quantified in terms of required amount of CPU cores to achieve real-time

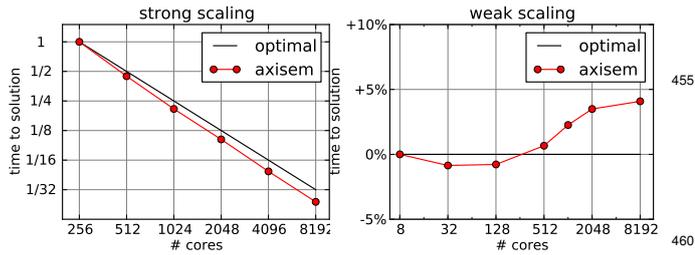


Fig. 6. Scalability of AxiSEM on a Cray XE6 at CSCS (Switzerland). Left: Strong scaling, i.e. fixed global problem size (8 Million 2D elements) as a function of the number of used cores communicating via the message-passing interface, for 12000 time steps. AxiSEM scales super-optimal, which is mostly due to the more efficient usage of run-time memory if less memory is used per core. Right: Weak scaling, i.e. fixed problem size per core (1000 elements) for 1000 time steps. The desired constant time-to-solution is exceeded by 4% for >8000 cores in which case communication is not entirely hidden behind the computation of the stiffness terms.

simulations (seismogram time equals CPU wall-clock time), assuming perfect scalability. Fig. 6 shows the new implementation of 2D parallelization and strong as well as weak scaling results on a Cray XE6 supercomputer installed at CSCS, Switzerland. In both cases, the performance is excellent: Strong scaling (fixed global degrees of freedom) is even super-optimal due to efficient memory usage. Weak scaling shows a slightly sub-optimal behavior at 96 % for > 8000 cores, still a remarkable figure indicating that message passing and parallelization are essentially hidden within the code. It is noteworthy to recognize that AxiSEM has little run-time memory, and applications at the high-frequency end benefit from vast multi-core systems mainly to reduce wall-clock time, unlike 3D seismic methods which are often memory-bound. As in any (2D) time-domain discrete method, it is important to recognize that half the dominant period takes about 8 times longer if seismogram length is fixed: The mesh is about 4 times larger, and the time step about twice as small. Note that monopole source types run faster than dipoles and quadrupoles due to their sparser stiffness terms (Nissen-Meyer et al., 2007b).

The number and meaning of input parameters for AxiSEM are kept at a generic, streamlined minimum, providing a robust basis in an effort to reduce failure. This is amended by a comprehensive number of sanity checks prior to the time loop, including critical tests upon mesh configurations, source, model, receiver settings parallelization, discrete volume and mass of the Earth, accuracy of internal surfaces, numerical quadrature, mass matrix and boundary terms. As a rough estimate, each 2D element occupies 1.5 wavelengths and about 2.5 kB, and for seismic periods of 5 seconds approximately 400,000 elements are needed. The code requires about eight microseconds simulation time per time-step and element.

2.8 Excessive input/output

Spectral-element methods of the kind presented here have excellent scalability properties in general (Fig. 6). The bottleneck, especially when moving to higher resolution and larger parallelization, lies in disk access which is necessary for saving wavefields and seismograms at run-time. For storage of synthetic seismic data, especially for a database of pre-computed waveforms (see sect. 5.3), platform-independence of the data is needed. However, the storage format of Fortran binary files is not even compiler-independent. To ensure true platform independence, AxiSEM fully supports the widely accepted NetCDF4 (<http://www.unidata.ucar.edu/netcdf/>) format to store seismograms and wavefields, but users may also revert to Fortran binary if desired. A NetCDF4 file is a container, in which very large variables (e.g. wavefields) as well as single scalar values (e.g. general simulation information) can be stored. The format allows transparent compression of the variables using the SZIP algorithm (shu Yeh et al., 2002), which saves around 50% of hard drive space for a typical seismic wavefield with respect to generic binary format. The container character of a NetCDF file means that direct access to selected data is possible, i.e. small amounts of data such as time series can be read from a (potentially very large) file, without loading the whole file into memory. The code allows to store all simulation output in *one* self-contained NetCDF4 file, which facilitates handling of a large number of simulation results, for example in parameter studies.

NetCDF4, which is based on the HDF5 format, allows for parallel writing into one file. Since this makes use of parallel file systems very efficiently, it might provide big performance gains on the next generation of supercomputers. On the current generation however, the installation of parallel NetCDF4 is not generally reliable yet. Therefore and to increase compatibility with older machines, AxiSEM uses a serial round Robin-scheme for writing data to disk. All processors buffer their respective wavefield output locally. After a set number of time steps, one instance spawns a new thread and transfers its wavefield buffer to it. This new thread then opens the output file and compresses and writes the buffer to disk, while the original processor continues to simulate the wavefield. This non-blocking IO scheme has been tested to work well up to 224 parallel instances, in that wavefield storage marginally affects CPU time and performance.

2.9 Implementation and availability

The AxiSEM code is written in Fortran2003 combined with MPI message passing, requiring corresponding compilers. Optional additional packages are NetCDF4 (Rew & Davis, 1990) for improved I/O, fftw3 (Frigo & Johnson, 2005) for post-processing in the frequency domain, TauP (Crotwell et al., 1999) for traveltimes picks, paraview for visualization of *vtk*- and *xdmf*-based wavefields, Matlab for visu-

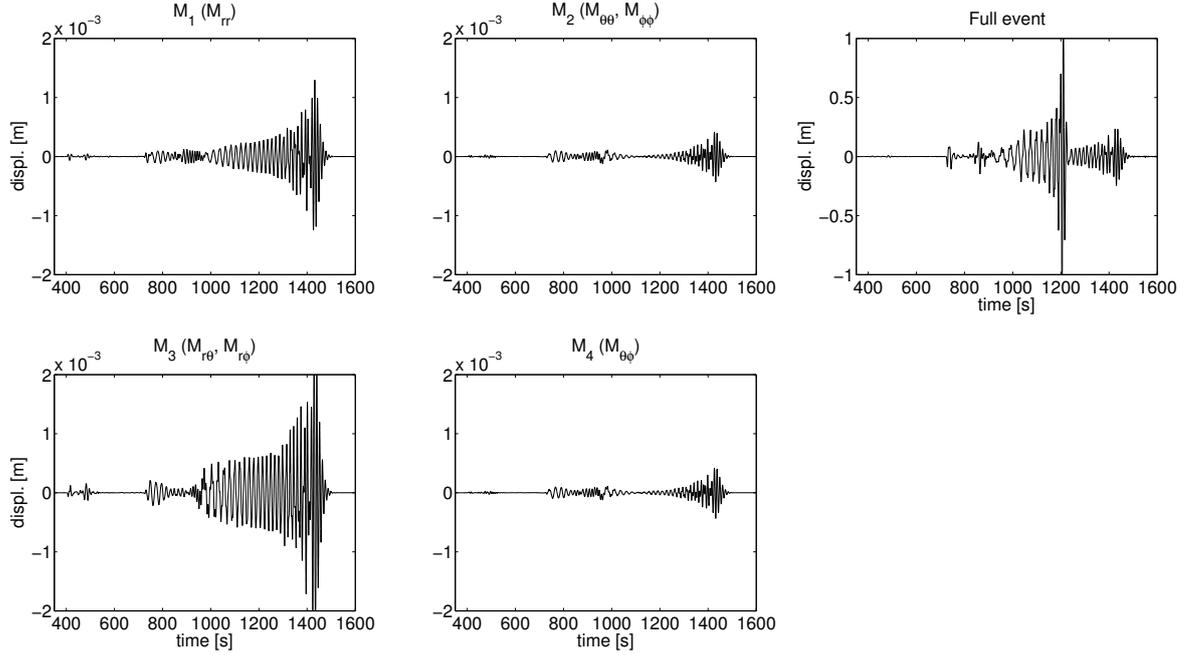


Fig. 7. Left quadrant: The four time series upon the generic moment-tensor types (see eq. (9)–(12)). Right: Summation to the full seismogram for the 2011 M9 Tohoku (point-source) event. Plotted is the displacement in the s -direction (i.e., perpendicular to the symmetry axis, see Fig. 3), with a dominant period of 10 seconds recorded at station BILL (East Siberia) at 33° distance.

alization of record sections, gnuplot for creating seismogram image files, Google Earth for visualization of source-receiver geometries and seismograms on the sphere, python wrappers for streamlined input/output and linkage to Obspy (Beyreuther et al., 2010). The Fortran2003 code is divided into a Mesher utilizing OpenMP, a Solver utilizing the message-passing interface (MPI) for communication between separate domains, and extensive post-processing for ease of visualization, filtering, source-time functions, various receiver component systems, and moment-tensor solutions.

AxiSEM is available through a release version with GPL license from www.axisem.info, and comes with no guarantee of functionality or support, but each version contains a detailed manual, examples, nightly builds and tractable sub-version control system as well as an existent userbase.

3 Seismic sources

In global seismology, it is customary to rely on the point-source approximation and corresponding moment tensors \mathbf{M}^P (not to be confused with mass matrix \mathbf{M} in the last section). The implementation of indigenous earthquake sources or single forces located at \mathbf{x}_p is detailed in Nissen-Meyer et al. (2007b). We use temporal Dirac delta functions acting at time t_p in the simulations, such that a displacement time

series is obtained by convolving a source-time function $S^p(t)$ (incorporated into a time-dependent moment-tensor term as $\mathbf{M}^P(t) = \mathbf{M}^P S^p(t)$), with the Green's tensor solution \mathbf{G}^P

$$\mathbf{u}^p(\mathbf{x}, \omega) = \mathbf{M}^P(\omega) : \nabla_p \mathbf{G}^P(\mathbf{x}, \omega), \quad (7)$$

where we reverted to frequency domain ω for concise notation, and ∇_p denotes spatial differentiation with respect to the source coordinate (no summation implied). Here, we focus on basic necessary post-processing operations to obtain the response to a full moment-tensor, the extension to finite kinematic faults, stochastic sources (as for example in noise seismology or helioseismology), and the problem of handling Dirac delta functions $\delta(\mathbf{x})$ in a discrete world.

3.1 Moment-tensor & single forces

To obtain the response to a full moment tensor (e.g., CMT catalogue, www.globalcmt.org), one applies posterior summation honoring the respective radiation patterns along the azimuth along with the convolution:

$$\mathbf{u}^p(\mathbf{x}, \omega) = \sum_{m=1}^4 \tilde{M}_m^P(\phi, \omega) \tilde{\mathbf{G}}_m^P(\tilde{\mathbf{x}}, \omega), \quad (8)$$

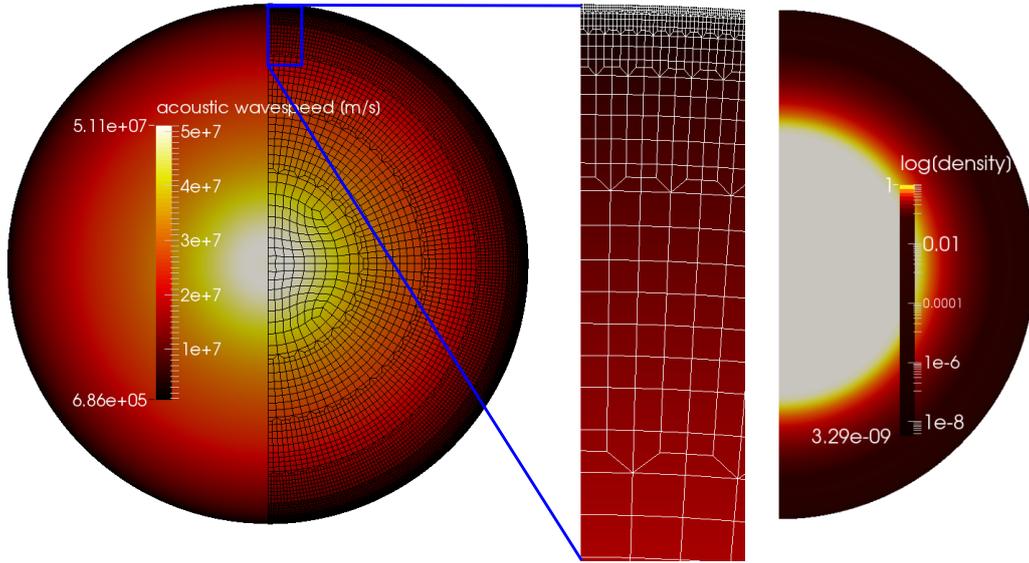


Fig. 8. A mesh for the Sun’s interior which accommodates the radial structure of the Sun for frequencies up to 5 mHz. It honors acoustic wavespeed variations of the Sun across two orders of magnitude (left), leading to seven coarsening (doubling) layers. Density (right) varies by eleven orders of magnitude, but does not affect the meshing process so long as these variations are smooth at the scale of elements. Such a mesh represents the basis for wave propagation and imaging the solar interior utilizing stochastic noise excitation within the framework of time-distance helioseismology.

where $\tilde{\mathbf{G}}_m^p$ represent 2D vectorial Green’s responses to each simulation m for point source p , and \tilde{M}_m read

$$\tilde{M}_1(\phi) = M_{rr}, \quad (9)$$

$$\tilde{M}_2(\phi) = (M_{\theta\theta} + M_{\phi\phi})/2, \quad (10)$$

$$\tilde{M}_3(\phi) = M_{r\theta} \cos \phi + M_{r\phi} \sin \phi, \quad (11)$$

$$\tilde{M}_4(\phi) = (M_{\theta\theta} - M_{\phi\phi}) \cos 2\phi + M_{\phi\theta} \sin 2\phi. \quad (12)$$

Only these four independent types of radiation patterns exist (monopoles \tilde{M}_1, \tilde{M}_2 ; dipole \tilde{M}_3 ; and quadrupole \tilde{M}_4) in this axisymmetric framework. For a full earthquake moment tensor, four independent simulations are thus undertaken to account for the six moment-tensor elements M_{ij} , whereas single forces (as needed for Lamb’s problem, ambient noise, impacts, or adjoint wavefields), require two simulations (monopole vertical force, dipole horizontal force) to account for the three components (Nissen-Meyer et al., 2007b).

Fig. 7 depicts an example for the individual displacement solutions for each \tilde{M}_m (quadrant on the left), and the final sum for a response to the full 2011 M9 Tohoku earthquake, recorded at station BILL (Eastern Siberia) at 33° distance. Note that the summed trace bears little resemblance with any of the generic solutions to radiation patterns eqs (9)–(12).

3.2 Finite faults

Kinematic rupture over a fault plane can be modelled as a discrete sequence of point sources distributed across the fault plane, each of which may have individual moment tensors, magnitudes and source-time functions to mimic time-dependent slip. AxiSEM is well positioned for an effective incorporation of such finite faults: Due to the rotational symmetries outlined above, the number of simulations for an arbitrary fault is simply given by its number of discrete depth points. The solution for finite-fault displacements may be written in terms of the solution to individual point-source solutions \mathbf{u}^p

$$\mathbf{u}(\mathbf{x}, t) = \sum_p \mathbf{u}^p(\mathbf{x}, t). \quad (13)$$

Note that the dependence on the point-source locations \mathbf{x}^p exists for moment-tensor \mathbf{M}^p (by means of radiation pattern and source-time function), and Green tensor \mathbf{G}^p in eq. (7), requiring separate solutions to the wave equation for each location \mathbf{x}^p . A significant shortcut can be made in the case of spherically symmetric media by saving seismograms at “all” distances, and applying rotational properties to the Green tensor. As such, all laterally distributed points \mathbf{x}^p are accommodated within one simulation, and only the discrete depths need to be honored by separate simulations. This is advantageous, as most finite-fault models are mainly distributed

laterally, and only require a few depth samplings. This allows for considerable flexibility should one wish to change certain properties of the fault model without conducting new simulations. In light of the common problem of local minima in (non-linear) source inversions, this offers an efficient engine to perform comprehensive studies on the behavior of different fault models and methodologies (Page et al., 2011). The modeling of finite sources is thereby largely delegated to post-processing (see sect. 5.3), such that existent AxiSEM databases can simply be applied to finite sources as well, and finite faults can be naturally embedded within any application of AxiSEM with little additional computational effort.

3.3 Stochastic sources

The rotational invariance also facilitates applications of spatially distributed stochastic sources such as ambient noise generated by ocean-continent interactions or crustal scattering (e.g. Boué et al., 2013), or random pressure fluctuations in the Sun's interior (Gizon & Birch, 2002). Similar to finite faults, one simulates point sources at the relevant number of depths (for ocean-continent ambient noise, this is one depth) and relegates the spatial distribution and stochastic time-frequency behavior to post-processing for rotations and filtering, respectively. This helps not only for generating a diffuse wavefield for structural imaging, but also for inverting for ambient-noise source locations.

3.4 Discrete Dirac delta distribution

It is desirable to simulate Green's functions, as they offer flexibility with respect to filtering and source-time functions after the simulation. The "source-time function" for the simulation is then a Dirac delta "function", which, from a rigorous perspective, is meaningless in any discretized system. To retain its attractive properties as the "source time function" to generate Green's functions, one instead utilizes a triangle function that obeys integral properties of the Dirac distribution. Should one wish to extract a downsampled time series from a simulation of this kind, then the "width" of the Delta function must be adjusted to guarantee the trade-off between 1) its Delta-function characteristics and 2) sufficient sampling below Nyquist frequency. This is automatically computed in the code, depending on the sampling and period ranges.

4 Structural properties

The definition, discretization and implementation of background models is one of the most critical aspects for accurate wave propagation. Amongst the decision factors are:

- the scale-lengths of structural variations, and their feasible upscaling from a potentially smaller-scale model,
- merging diverse models of source and structure,

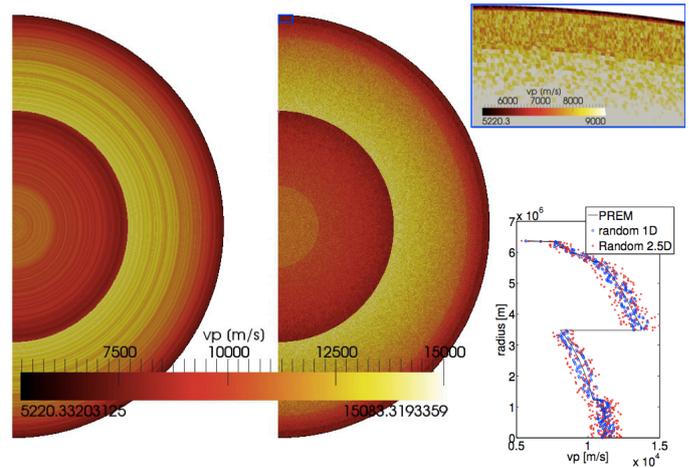


Fig. 9. Two examples of random wavespeed variations superimposed on PREM. Left: variations with depth only. Middle: Variations in 2.5D. Right (top): Zoom into the crust for the 2.5D random medium. Right (bottom): radial profile of these two realizations with respect to PREM.

- sharp versus smooth variations,
- local reliability and resolution, for instance in global tomographic models.

Uncertain choices amongst these points may lead to an entirely wrong model and consequently useless wave-propagation results. Discretization and meshing in finite-element based methods usually strive to replicate all sharp boundaries of the model. Apart from algorithmic limitations in meshing arbitrary hexahedral elements, any failure to mesh a desired interface leads to false solutions.

AxiSEM offers an internal meshing algorithm which optimally honors any discontinuities in spherically symmetric Earth models, as well as arbitrary discrete spheres of any radial distribution of solid and fluid domains. Axisymmetric structures that are invariant in the azimuthal direction may then be superimposed onto the background mesh on-the-fly in the solver. These structures can be of pre-defined shape types, or arbitrarily superimposed by interpolation of discrete external grids (see Section 4.4).

4.1 Spherically symmetric models

Spherically symmetric models such as PREM, IASP91, and ak135 are automatically incorporated in AxiSEM. The code also allows for flexible inclusion of arbitrary 1D structures in the meshing process, such that other Earth models such as those based on mineral physics can be easily accommodated. One may also apply the methodology to other planetary bodies (e.g. Moon, Mars, Europa, not shown here) including purely fluid media to facilitate acoustic wave propagation (a computational shortcut still popular in the exploration indus-

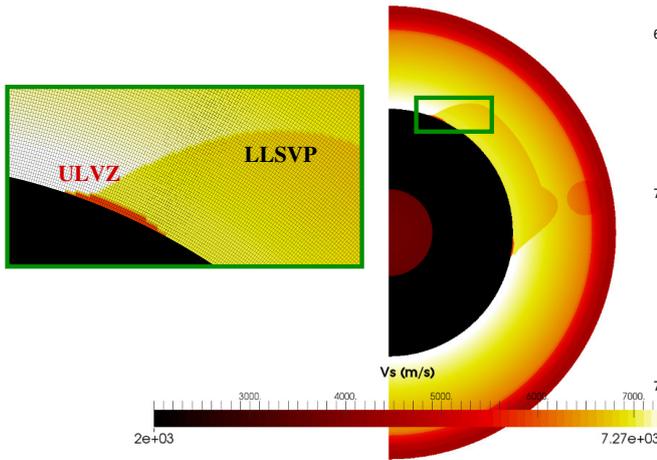


Fig. 10. An example of various lateral heterogeneities, representing realistic deep-mantle structures projected onto the source-receiver plane with azimuthal invariance. The large volume in yellow denotes a Large Low-shear wave Velocity Province (LLSVP), flanked by two exaggerated ultra-low velocity zones (50 km height, 10% P -velocity decrease, 20% S -velocity decrease, 10% density increase), underlying a detached uprising in the mid-mantle. The implementation is done by assigning laterally heterogeneous properties to the coefficients of the basis functions, as commonly done in high-order spectral-element methods (Peter et al., 2011) so long as elements are sufficiently small to capture variations in a smooth manner.

try), which drastically reduces computational cost. As a curious case of extreme medium variations readily discretized by our methodology, we reach out for our central star: The sun is a giant fluid sphere subject to turbulent redistribution of masses, magnetic field variations, and acoustic body waves (Gizon & Birch, 2002). The background structure of the sun covers many orders of magnitude in density due to its huge size and gravitational force, and about two orders of magnitude in acoustic wavespeeds. Only the radial wavespeed gradient matters for meshing, thus it is easily possible to adapt the AxiSEM mesh to the Sun as seen in Fig. 8. Surface boundary conditions for acoustic waves (i.e. (vanishing pressure) pose no technical difficulty, but are not yet included in the first code release and shall be added in a timely future version.

4.2 Crustal variations

Crustal thickness variations from 6–8 km (oceans) to 60–80 km (continental shields) owing to lithospheric composition have a significant imprint on those seismic phases that are sensitive to shallow structure, such as surface waves which traverse the crust to large distances. Additionally, covering 70% of Earth's surface, the oceans are also a contributor to wavefield modifications, even though most seismometers are installed on land.

The computational efficiency of a 2D numerical method allows for sufficiently small elements to explicitly mesh the crust, which is also necessary for wavelengths in the range of crustal thicknesses. This is also true for the actual oceans, which may be discretized by actual fluid-domain elements, instead of resorting to a loading equivalent (Komatitsch & Tromp, 2002b) or a homogenized crust (Fichtner & Igel, 2008). Similar to axisymmetric structures, this divides the sphere into oceanic and continental pole-centric rings. Discretized oceans are not available in the first code release, but will be added in the near future.

4.3 Random media

Spherically symmetric or axisymmetric variations in properties can be as general as desired in the method, including random media variations so long as they are sufficiently smooth and mildly deviatoric. In Fig. 9, we show two types of such random variations, perturbing either radial structure (left), or 2.5D structure by maximally 10% velocity variations. Wave propagation through such complex structures can deliver useful insight into wave effects as a function of spatial scale dependence, scattering and homogenization properties, or the relation between structural heterogeneities and seismic measurements (Baig & Dahlen, 2003).

4.4 Localized heterogeneities

The next level of complexity in structural properties is represented by axisymmetric media, which may have arbitrary variations in the source-receiver plane, but are invariant in the perpendicular azimuthal direction. Especially in high-frequency regimes, 3D edge effects from large-scale structures may become less dominant given the decreased ratio between seismic and structural wavelengths, such that Fresnel zones reside within the azimuthal extent of an anomaly. Axisymmetry can then represent a tangible basis for waveform modelling of unknown arrivals, precursors, undetected arrival delays, oblique reflectors. The only neglected part of the wavefield compared to 3D background models are 3D wave effects from off-plane structures such as 3D elastic lense focusing and off-plane back-scattering. All other scenarios in which structures vary preferentially in the source-receiver direction are respected, as for instance wave propagation through certain configurations for subduction zones, or forward scattering of small-scale lowermost mantle structures. Fig. 10 shows one example of lateral heterogeneities implemented in AxiSEM, including a Large Low-Shear Wave Velocity Province (LLSVP) (Romanowicz & Gung, 2002), an Ultra-Low Velocity Zone (ULVZ) (Rost, 2013), and a disconnected uprising (e.g. Zhao, 2001). These variations are sufficiently smooth to be picked up by the elemental basis functions within the spectral-element mesh (see left panel). The inclusion of such lateral heterogeneities can be done by functional parameterization as shown here, but also

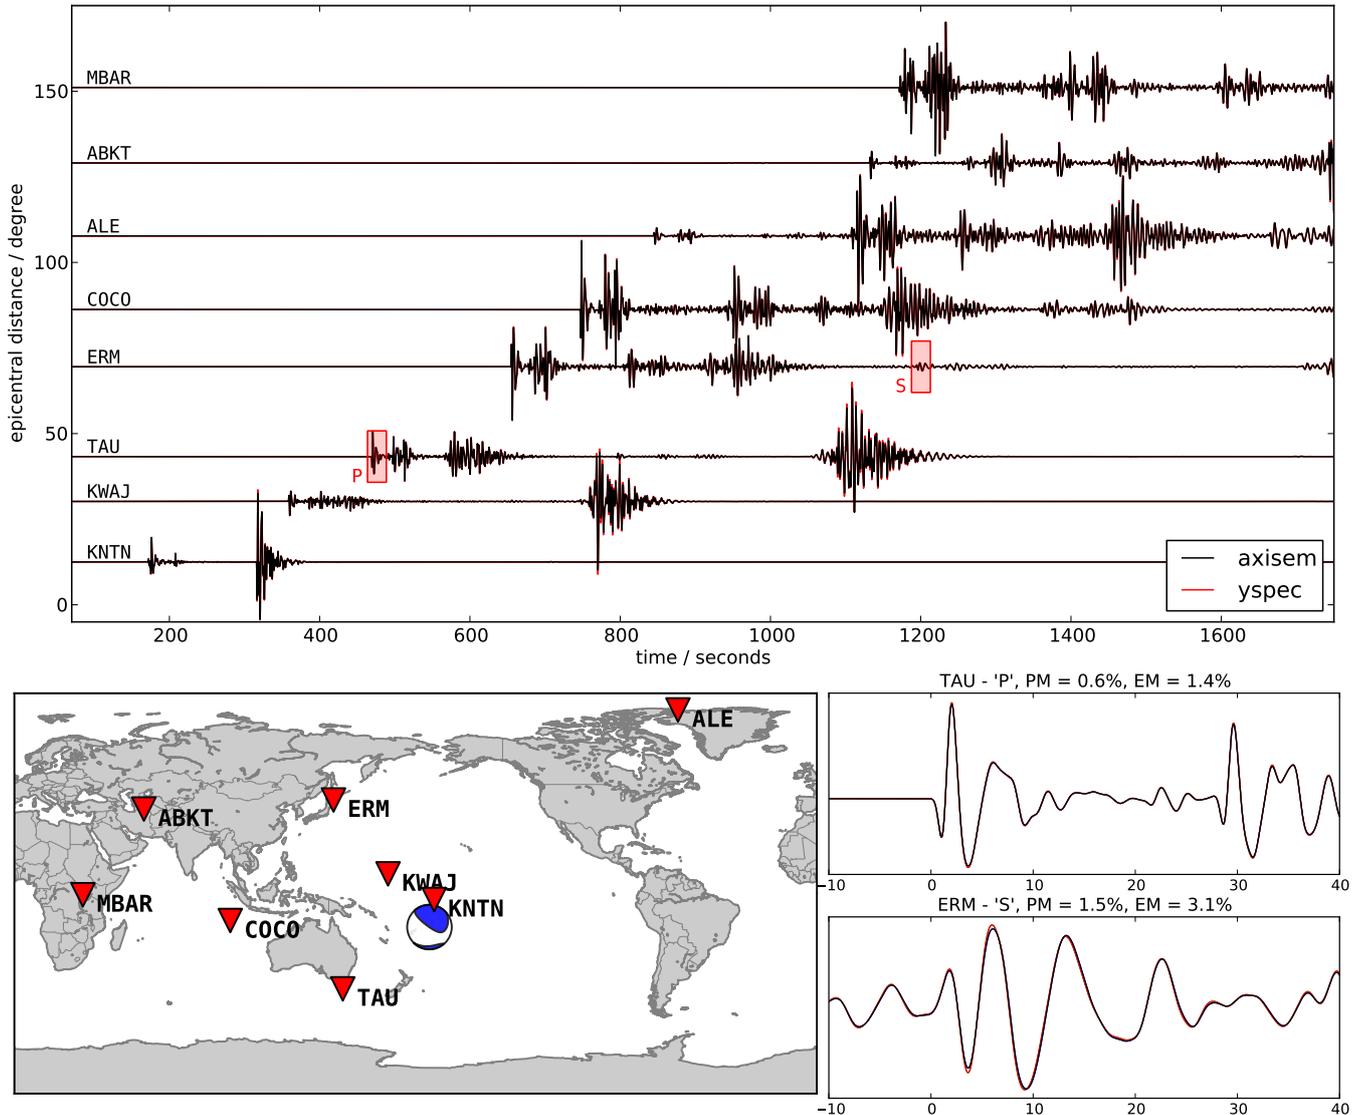


Fig. 11. High-frequency validation (1 Hz dominant frequency) between AxiSEM and YSpec. Top: Record section for vertical displacements of a M4.1 event in Tonga (depth: 126 km), recorded at the stations shown on the map (bottom left) as red triangles. The background model is PREM, including anisotropy and attenuation, and the traces are filtered between seismic frequencies of 0.1-1 Hz, i.e. at the limit of recordable signals in global seismology. The traces from AxiSEM and Yspec are virtually indistinguishable. The zoom sections for individual seismograms (bottom right) on *P* and *S* waves (red boxes) represent phases that traveled 500 and 1200 wavelengths, respectively. Time in these panels is normalized to the ray-theoretical phase arrivals (Crotwell et al., 1999), and includes phase (PM) and envelope misfits (EM) measured following Kristeková et al. (2009).

745 by discrete external models. Such arbitrary models are incorporated via KD-tree (Kennel, 2004) search for nearest neighbors and interpolation, and therefore allows for any shape and complexity. The accuracy of wave propagation through such models is governed by the scale of heterogeneity versus 755
 750 finite-element size, in that strong variations at short spatial scales tend to be smoothed in the discrete model.

4.5 Anisotropy

In a spherically symmetric scenario, the most complex anisotropy is transverse isotropy with 5 independent parameters. In axisymmetry, we may incorporate the full elasticity tensor with triclinic symmetry and 21 independent parameters, so long as the anisotropy does not vary within the azimuthal direction of the Fresnel zone. While logically mandated, this theoretical fact is in itself intriguing: An individual source-receiver wave senses only anisotropic vari-

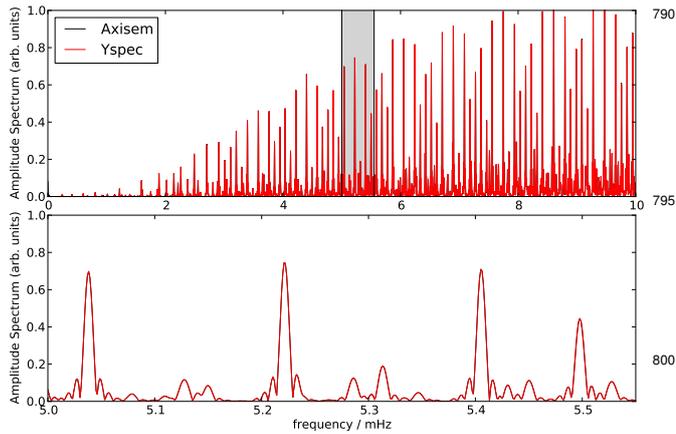


Fig. 12. Amplitude spectra simulated by AxiSEM and YSpec for the PREM model (Dziewonski & Anderson, 1981) for frequencies below 0.02 mHz. The time-domain solution provided by AxiSEM extended over 48-hour time series using 1.7 Million time steps. While amplitude spectra do not exhibit issues related to numerical dispersion, the fit between these two different methods is remarkable.

ations within a sufficiently narrow azimuthal range of sensitivity. If, however, a higher complexity of anisotropy is present but varies at a scale larger than the sensitivity of the traveling wave, then this level of large-scale complexity is not extractable from singular seismograms but instead represents an effective image of the actual structure. Van Driel & Nissen-Meyer (2014a) provide a detailed analysis and implementation strategy for anisotropic wave propagation in axisymmetry, including proofs related to the multipole expansion for the presence of general anisotropy in the axisymmetric environment.

4.6 Attenuation

Intrinsic attenuation or visco-elastic damping is a natural property of bulk real-Earth structure at relevant frequency ranges of seismic wave propagation. Although models for the quality factor Q (inverse damping) of the mantle show little agreement and the origin of damping may be inseparably coupled with elastic small-scale scattering, seismic attenuation on waveforms is a well-detected and significant phenomenon. We implemented an improved methodology based on coarse-grain memory variables with negligible additional computational cost compared to purely elastic wave propagation. This further includes attenuative bulk and shear deformation, 5 relaxation mechanisms, a combined linear/non-linear approach to identify optimal sets of parameters for the range of realistic Q , and analytical time stepping. Details on this new implementation, which is applicable to any higher-order finite-element method, are described in a separate paper (Van Driel & Nissen-Meyer, 2014b).

4.7 Lack of ellipticity & rotation

The Earth's radius differs, depending on latitude, by up to 40 km between poles and equator. For reasons of axial symmetry, AxiSEM does not allow waves to propagate from a non-polar point source through a pole-centric ellipsoid. To account for ellipticity a posteriori, three options are suggested: 1) phase correction, 2) epicentral distance correction, 3) Born perturbation theory. Phase-specific ellipticity corrections may be applied by shifting waveforms according to predicted traveltimes. This is useful only if individual phases are assessed, such as in most cases of tomography, and phase-specific waveform modeling. Epicentral distance correction may be conducted by recalculating receiver coordinates to account for the difference between purely spherical and ellipsoidal geometries, similar to the standard method in traditional tomography (Kennett & Gudmundsson, 1996). Finally, Born theory may be applied by assuming ellipticity (including internal interfaces) to act as a boundary perturbations to the spherical model domain. This way, entire seismograms are accounted for jointly. This approach has not been implemented or tested at this point. Rotation of the polar axis can in principle be incorporated in AxiSEM, but only for a polar source, which clearly is a rather unique case of rotation. At the scale of interest where rotation comes into play (above periods of 100 seconds), one could devise a torus-shaped, off-axis source in case its azimuthal radiation is of lesser significance - as may be the case for free oscillations. This would be a field of further study and implementation. In summary, such effects grow into a visible and recordable first-order concern for rather specific cases of seismic data analysis concerned with pathological body-wave paths and very long-period seismology.

5 Wave-propagation applications

Our methodology and the actual code AxiSEM are production-ready and may be used to tackle a diverse range of applications. Here, we sketch some of these, ranging from basic validation against reference solutions across the frequency spectrum, indefinite solutions to wave propagation, 3D wavefield visualization, lowermost mantle heterogeneities, tomographic models, comparison to recorded data, and sensitivity kernels. All examples are deliberately disconnected as a showcase for the diverse range of applications.

5.1 High-frequency body waves

Previous publications on an early version of this method and implementation showcased the accuracy by comparison against normal-mode summation (Nissen-Meyer et al., 2008). Normal-mode summation is difficult to achieve at high method frequencies due to the computational cost in generating mode catalogues, as well as numerical issues related to determining the eigenfrequencies. We use an al-

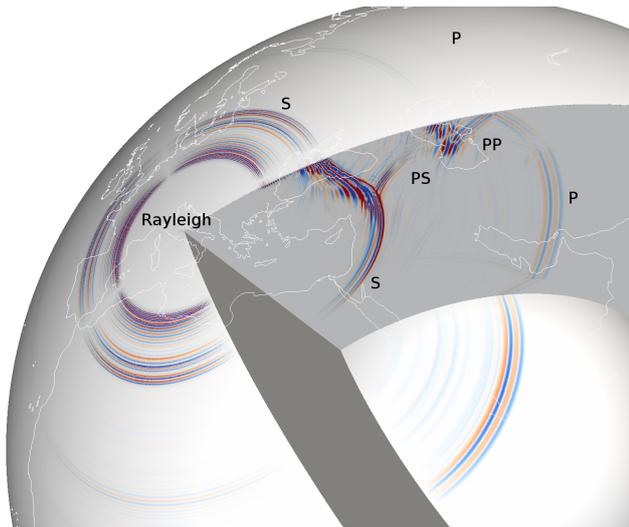


Fig. 13. Snapshot of a 3D wavefield emanating from a strike-slip event in Italy after 400 seconds. The background model is isotropic, anelastic PREM, and the simulation done at a dominant period of 10s. Note the effect of the radiation pattern on the wavefield in 3D. Similar snapshots are automatically generated in the post-processing of AxiSEM. A movie is available as supplementary material.

ternative frequency-domain reference solution YSpec (Al-Attar & Woodhouse, 2008) capable of covering the entire relevant frequency band from 0.001 Hz to 1 Hz. Fig. (11) shows a record section and some details for both AxiSEM and YSpec modeling results in an anisotropic, visco-elastic PREM model for a Tonga event at 126 km depth, simulated at a dominant frequency of 1 Hz, i.e. at the limit of teleseismic detection of body waves. The fit is excellent for all phases and distances; the two solutions are indistinguishable almost everywhere out to 1600 propagated wavelengths. Minor differences are amplitude differences, and most probably due to a cut-off in the summation done in YSpec. To our knowledge, this is the first accurate validation of two completely independent methods for anisotropic, viscoelastic media at such high resolution.

5.2 Free Oscillations

At the grave end of the spectrum, free oscillations dominate and have revealed a great deal about Earth's internal structure, in particular the Earth's density structure (Dahlen & Tromp, 1998). We strive to provide a numerical method applicable to wave propagation across the observable frequency band. We thus compare amplitude spectra stemming from AxiSEM in a simulation over 48 hours, 1.7 Million time steps, against YSpec in Fig. (12). The fit is again excellent, which is not trivial considering that time-domain numerical methods are exposed to steadily growing numerical dispersion errors with increasing numbers of propagated cycles.

To the best of our knowledge, this is the first direct benchmark between time- and frequency-domain methods for free oscillations of the Earth, even if for a spherically symmetric, non-rotating, non-gravitating Earth model. Phase spectra, which may be more informative for actual studies with normal modes, are shown in (Van Driel & Nissen-Meyer, 2014a).

5.3 Instantaneous forward solutions

The reduced dimensionality of AxiSEM opens doors to simulating the entire response due to a given background model once-and-for-all, for all possible source-receiver choices. This seemingly daunting task is rendered tractable by the rotational properties of the displacement vector eq. (2), such that seismograms only need to be stored along the distance range $0 - 180^\circ$ for sources at a range of depths. This is computable. The remaining problem lies in deciding on a discrete sampling for source depth and receiver spacing to mimic continuous coverage. In the case of depths, this may be defined upon depth uncertainties in different earthquake catalogues, and in the case of receivers by choosing the closest or interpolating upon epicentral distance uncertainty levels.

The computation of such a once-and-for-all solution can be conducted by taking into account the reciprocity of the Green's function, resulting in only two simulations: one with a vertical and one with a horizontal single force, upon which the strain tensor needs to be stored for all realistic earthquake depths 0–660 km at all distances. This reciprocity shortcut is fueled by the fact that AxiSEM carries the full 3D wavefield automatically, as opposed to reflectivity, DSM, Yspec or normal-modes solutions for which the number of saved seismogram locations factors into the computational cost. The problem is thereby shifted from CPU-time to hard-drive storage. The permanent storage for the entire parameter space spanning all source-earthquake configurations and several Earth models is feasible (tens of Terabytes). Queries to such databases (such as a record section of arbitrary source-receiver geometries, filters, source-time functions, and a range of spherically symmetric Earth models), can be completed within minutes by means of the same kind of post-processing as done upon the AxiSEM code. This can be tremendously beneficial in studies that need to sample a large range of parameters such as source inversion problems, especially in a probabilistic framework (Stähler & Sigloch, 2013).

5.4 Wavefield visualization

Most of the applications in this section, as well as in the literature, rely upon seismogram analysis. However, one of the major benefits of this method is the availability of the full global 3D space-time wavefield, for both research and teaching purposes (Thorne et al., 2013). This is possible only due to the collapse to 2D at run-time and the convenience of

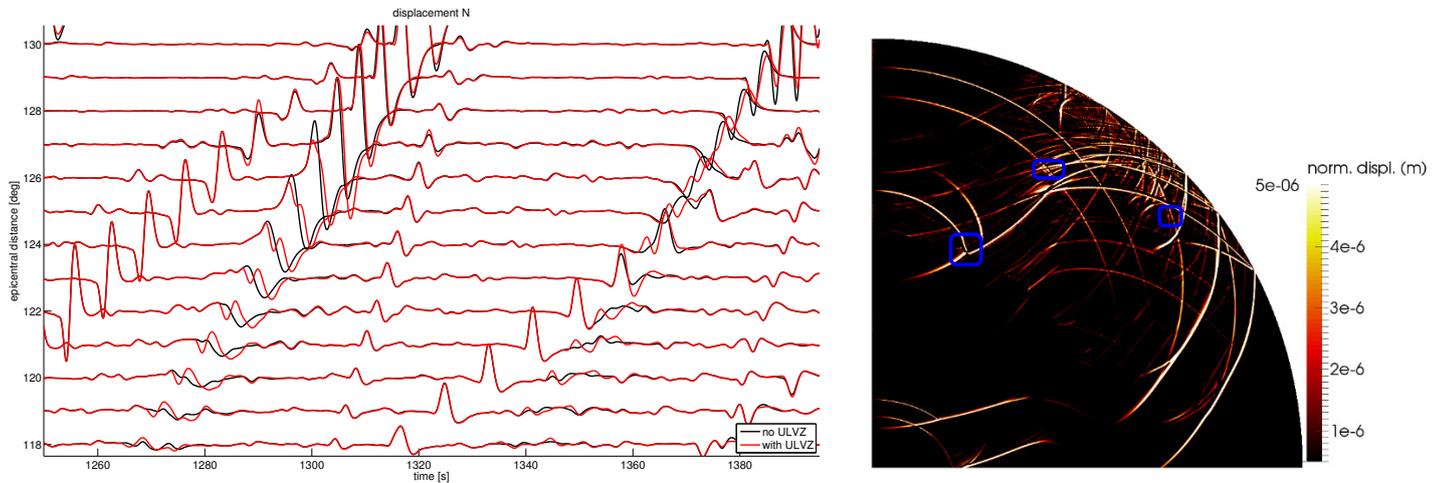


Fig. 14. A forensic application of AxiSEM on the in-plane detectability of an azimuthally invariant representation of two adjacent structures: A “ULVZ” near a “LLSVP” (see model in Fig. 10). Left: Seismograms for a model with a ULVZ as in Fig. 10 (red traces), and one exactly the same but without the ULVZ (black traces). The underlying Earth model is isotropic PREM (Dziewonski & Anderson, 1981), dominant period 2s. The N -displacement record sections are at considerably large epicentral distance ranges. Right: Wavefield snapshot of the same simulation with ULVZ, at time 604 seconds. Blue quadrants denote those parts of the wavefield that are most affected by the presence of the ULVZ (in comparison to a similar plot for the simulation without ULVZ).

on-the-fly extraction of the 3D radiation upon azimuthal fac-⁹⁴⁵
 tors. In practice, this means that one may save the entire 2D
 wavefield in space and time, and then subsequently decide
 on any moment tensor, summation, source-receiver geome-
 try, and rendering choices. Fig. 13 depicts a snapshot of a
 typical simulation of a strike-slip event in Italy with a domi-⁹⁵⁰
 nant period of 10s and isotropic, anelastic PREM background
 model. Note the characteristic dispersion in the surface wave
 train, the large amplitudes in the PP phase, and the 3D radia-
 tion pattern. A movie of this setting is available in the sup-⁹⁵⁵
 plementary material. Such visualization may offer comple-
 mentary insight into complex propagation patterns beyond
 singular trace analysis, in particular as they can be devised as
 differential wavefields for diagnostic purposes in tracing the
 influence of changes in model parameters.⁹⁶⁰

5.5 Wave propagation through in-plane heterogeneities

As mentioned in Section 4.4, we may readily insert lateral
 heterogeneities and compute 3D wavefields upon those torus-⁹⁶⁵
 like structures. To neglect the imprint of the torus-shaped
 azimuthal invariance, seismic frequencies should be chosen
 such that they represent local wavelengths that are smaller
 than the expected azimuthal extent of a 3D heterogeneity
 to be modeled. In the regime of 1 Hz, this is warranted
 for a number of examples, and various lowermost mantle
 structures have been studied and constrained by waveform
 modelling for decades (e.g. Igel & Weber, 1996; Cottaar &
 Romanowicz, 2013). Constraining geometry and structural

composition of these features is crucial for understanding the
 thermo-chemical, dynamical regime of the Earth’s deep in-
 terior. In many previous applications, such structures have been
 modelled as azimuthally invariant with source-receiver-plane
 heterogeneity using approximations at frequencies below
 0.07 Hz to study core-mantle boundary scattering (Thomas
 et al., 2000), D’ layer (Thorne et al., 2007), and LLSVP
 structure (Sun et al., 2007). Fig. 14 displays seismograms and
 wavefield snapshots for the model in Fig. 10. The record sec-
 tions highlight phases and distances at which the existence
 of a ULVZ may be tested, possibly with array methods and
 stacking. The wavefield snapshots represent a complemen-
 tary diagnostic for differential studies, from which the most
 significant imprints can be traced back to the surface, affect-
 ing phases such as *PcP*, *ScS*, and *SPKS*.

Modeling of such lateral heterogeneities taps into a regime
 of wave propagation that offers a grasp of wave effects at
 resolution and computational cost that is difficult to achieve
 with alternative methods. Users should however, as always,
 be cautioned to recognize the structural assumptions imposed
 on such in-plane features: These may either approximate ac-
 tual 3D structures well (if the above-mentioned scale separa-
 tion is warranted), or act as an upper bound of waveform ef-
 fects (by means of azimuthal overestimation), but conversely
 they may also neglect elastic focusing and thus underestimate
 3D effects.

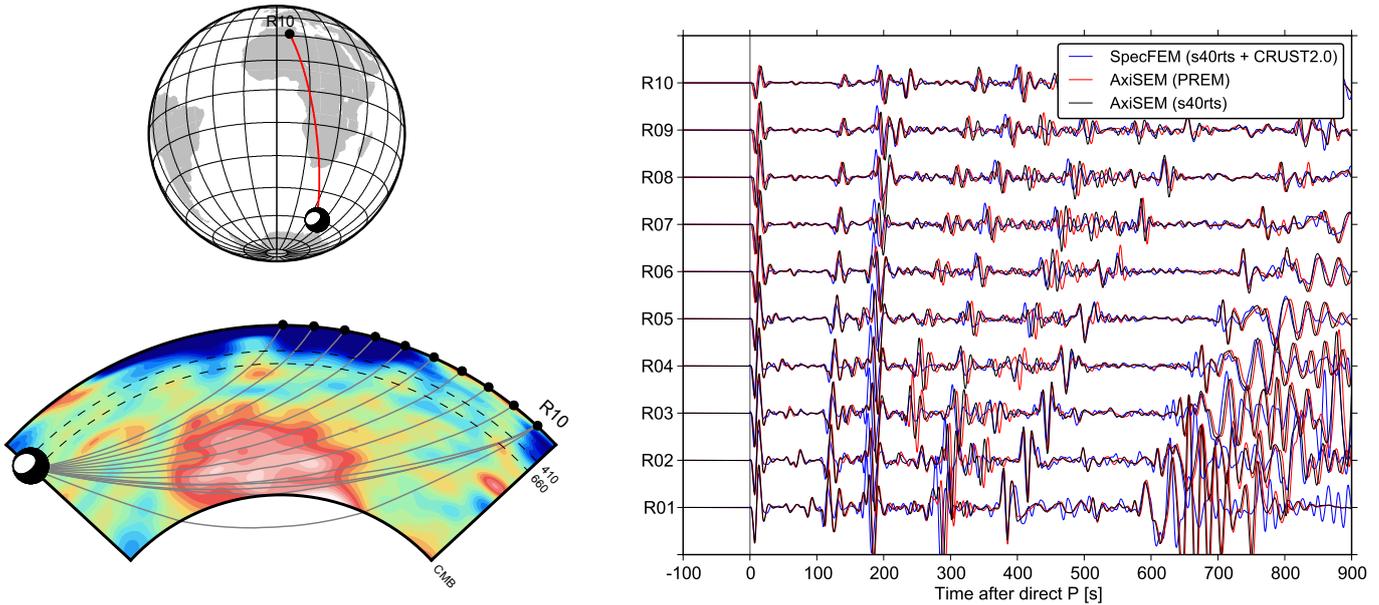


Fig. 15. Modelling 3D wave propagation through a 2.5D version of tomographic model S40RTS (Ritsema et al., 2011) (bottom left) for an event near Antarctica (left top). Right: Seismograms filtered at 10s from the receivers denoted on the cross section (bottom left) for 1D model, 2.5D tomographic model, and SPECFEM3D_GLOBE synthetics through S40RTS and CRUST2.0 (Bassin et al., 2000), aligned with the P -wave arrival time.

5.6 Tomographic models

Global models derived from tomography can also be approximated by an in-plane rendition with AxiSEM, in that they usually deviate only mildly and smoothly from spherically symmetric Earth models. Just as in the previous example, wavefield effects captured by this methodology are those that obey forward scattering, whereas true 3D-medium effects such as off-plane scattering are neglected. Of course, this azimuthal invariance does not represent our nature's dimensionality, but mimics a substantial sub-set of those data that are actually used for waveform modeling or tomography rather well and at a cost many orders of magnitude below that of simulating a 3D domain. This can be seen in Fig. 15, a comparison between synthetic modeling through PREM, an in-plane collapse of tomographic model S40RTS (Ritsema et al., 2011), and SPECFEM3D_GLOBE synthetics for S40RTS and CRUST2.0 (Bassin et al., 2000). As seen by the waveforms in Fig. 15, many phases are largely insensitive to the added complexity in these models, partly due to the smooth nature of tomographic models (as mandated by their inversion technique). Direct body waves and other early arrivals barely notice the different models, whereas later arrivals and surface waves exhibit considerable differences, most of which can be attributed to the crustal layer. The overall imprint of crustal variations overrides that of the tomographic model. The neglected effects such as 3D

back-scattering may indeed not contribute all that significantly to resultant seismograms, but this is subject to further parameter-space studies. In general, this provides an efficient new approach should one wish to validate different tomographic models within a synthetic exercise, or modify local properties for a given source-receiver geometry. The actual incorporation of tomographic models is trivial in AxiSEM for any model that is given by discrete cartesian, spherical grids, or spherical harmonics.

5.7 Relating to data

The ultimate raison d'être of any seismic modeling is its capability to relate to actual observed data, at least in some useful fraction of the generally impenetrable overall parameter space. Here, we showcase a comparison of waveforms at considerably high resolution (5 seconds) to observed data. This resolution is at the cutting edge of supercomputing with 3D methods (see Fig. 1), at a frequency range applicable to tomography, and also interesting for waveform modeling of relatively small-scale features in the lower mantle. The map (top right) shows the event and station locations (red triangles for $PKiKP$, blue for $Pdiff$). Filtering has been applied at 5–15 seconds (top), and 15–45 seconds. In the latter case, we included SPECFEM3D_GLOBE synthetics for the S362ANI tomographic model (Kustowski et al., 2008) and CRUST2.0 (Bassin et al., 2000) which are accurate to about 17 seconds (Tromp et al., 2010). AxiSEM synthetics

are based on an inverted moment tensor and depth, whereas SPEC-FEM3D_GLOBE synthetics are taken from the IRIS database, i.e. calculated for GCMT. All synthetics have been convolved with an inverted source-time function. The phases have been aligned by frequency-dependent cross-correlation, forming the basis for tomographic inversions. The waveform differences between all three traces fall within a feasible range of conducting waveform tomography. The timeshifts based on the 13 Pdiff and 6 PKiKP measurements are:

filter [s]	method	$\Delta t(\text{Pdiff})[s]$	$\Delta t(\text{PKiKP})[s]$
5 – 15	AxiSEM	2.5 ($\sigma^2 = 0.99$)	2.3 ($\sigma^2 = 0.5$)
15 – 45	AxiSEM	4.0 ($\sigma^2 = 0.89$)	3.4 ($\sigma^2 = 0.3$)
15 – 45	SPEC-FEM	3.7 ($\sigma^2 = 0.6$)	6.0 ($\sigma^2 = 0.37$)

Such comparisons include (as per usual) inevitable differences in processing such as event origin time and location, source time function. However, it is noteworthy to recognize the waveform similarity confirming that wave propagation in spherically symmetric Earth models provides an excellent basis for broadband waveform tomography, in particular in the regime of periods below 10 seconds.

5.8 Wavefield sensitivity kernels

As a final example, we present the essential and possibly most intriguing application to time- and frequency dependent sensitivity kernels (Nissen-Meyer et al., 2007a). The ability to store entire space-time wavefields from AxiSEM makes it principally trivial to compute time-dependent sensitivity kernels as a comprehensive basis for mapping seismograms to Earth structure, and thus the model-to-data operator for the tomographic inverse problem (Fuji et al., 2012; Colombi et al., 2014). As such, it logically extends existent ray-based or finite-frequency tomographies (which were based on approximate physics) by incorporating complete seismograms, arbitrary time- and frequency-windows as well as arbitrary wave effects such as triplicated phases from the mantle transition zone (Stähler et al., 2012), core-mantle diffraction (Colombi et al., 2012), or caustics. Fig. (17) shows a sensitivity kernel for cross-correlation traveltimes with respect to compressional wavespeeds. This was computed with wavefields from AxiSEM, by time-integrating the velocity waveform of dominant seismic period 10 s within a 20 s time window around the P -wave arrival time at 90° epicentral distance. The kernel exhibits considerable heterogeneity (partly due to saturated colorscales to highlight its complexity), notably missing the “donut-hole” that is present for pure P -wave kernels (Hung et al., 2000). This stems from the fact that our wavefield-based approach honors the large time window, which obscures the purity of phase-based approaches, but correctly represents the measurement corresponding to the time window. Time-dependent sensitivity (i.e. without integrating over the time window) is useful not only as a basis for tomography, but also in a forensic sense to detect faint signals of sensitivity due to a given region or structure. Note

that the (separate) calculation of sensitivity kernels is not part of the AxiSEM release, but will be added in the future as a separate package.

6 Conclusions

This paper presents a mature method and implementation for global seismic wave propagation across the seismic frequency spectrum by means of a diverse range of applications. It describes crucial extensions with respect to the initial papers (Nissen-Meyer et al., 2007b, 2008) such as the inclusion of anisotropy, attenuation, lateral heterogeneities, finite sources, the basis for sensitivity kernels, and innovative visualization. The method is, to our knowledge, the first time-domain local numerical method successfully benchmarked against independent solutions across the entire frequency band recorded in global seismology, and exhibits excellent scaling on large multi-core systems. The code offers a diverse range of realistic applications in forward and inverse modeling and showcases promising comparisons to recorded data. The moderate computational cost allows for reaching any desirable frequency with moderate resources, and storage of full space-time wavefields for sensitivity kernels.

The presented methodology is most accurate, efficient and useful in parameter regimes which are quite complementary to well-established, mature methods such as normal-mode summation (low-frequency seismology), 3D numerical methods (with local basis functions) such as SPEC-FEM3D_GLOBE (3D Earth models at intermediate frequencies), and asymptotic ray theory (high-frequency regime with potentially complex wavespeeds). As with any method, the realm of validity for AxiSEM is limited, approximative and blurred, and any application must be undertaken with caution despite the excellent and robust validation shown here and with the actual code available from <http://www.axisem.info>. This parameter space promises a diverse range of applications which were previously inconvenient, inexact, or unattainable due to limited computational resources or methodologies. Specifically, the main factor attributed to its efficiency in a 2D computational domain is the availability of space-time wavefields for axisymmetric, viscoelastic anisotropic media and realistic earthquake sources.

We have touched upon a few key applications, far and away from explaining or validating each one of them. Rather, the purpose of this paper is to present a new open-source methodology and scan its usability specifically in those directions that we deem most benefitting from this modeling tool. Details on specific applications and implementation are found in other publications to be submitted, and the state of reliable features in the code should always be consulted in its concurrent manual.

All limitations of the methodology are by construction related to the existence of the symmetry axis, which mainly translates into neglecting true 3D media (effects such as 3D

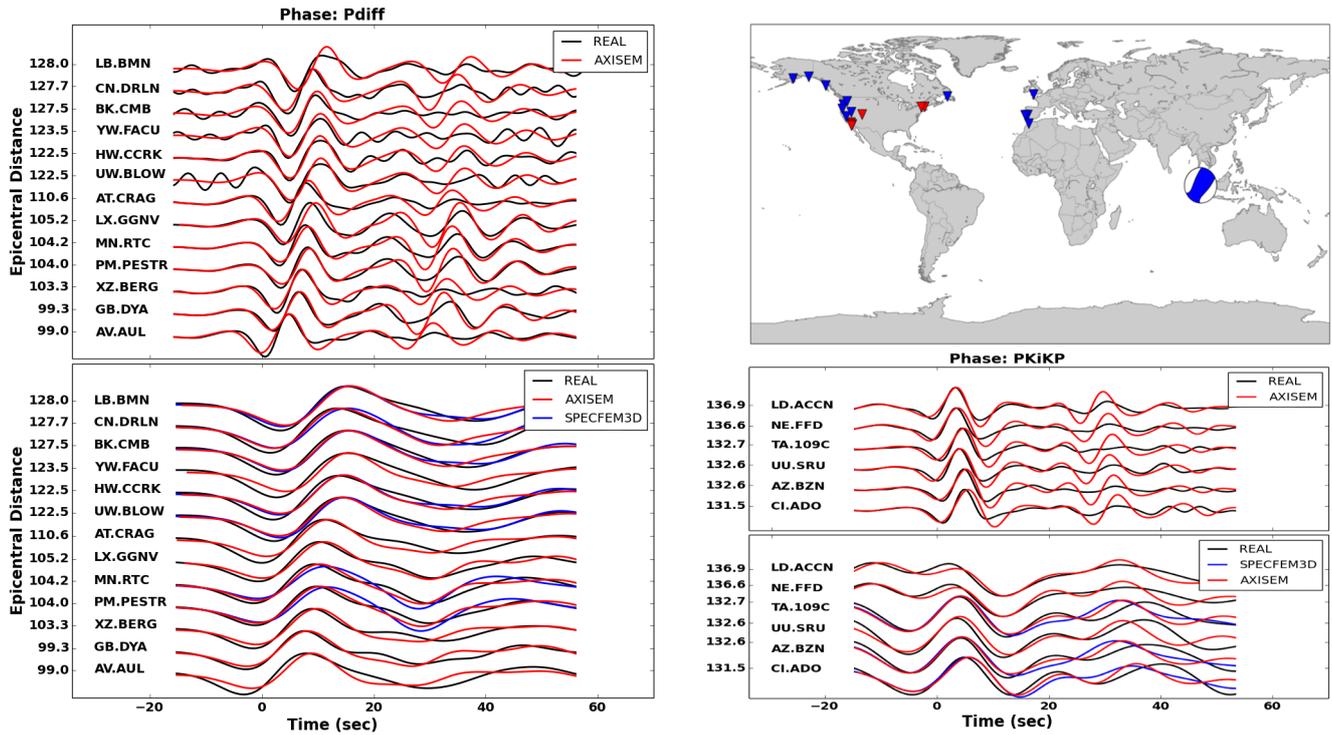


Fig. 16. A comparison of AxiSEM synthetics with recorded data and SPECIFEM3D_GLOBE synthetics (Tromp et al., 2010) for a Mw7.5 2009 event in Southern Sumatra at 82km depth. Top right: Event-station distribution, where red triangles are for core-phase PKiKP, blue for CMB-diffracted phase Pdiff. Left: Pdiff synthetics and observed data filtered at 5-15 seconds (top), and 15–45 seconds (bottom). In the latter case, SPECIFEM3D_GLOBE synthetics are included which are accurate down to 17 seconds. Bottom right: The same for PKiKP. AxiSEM synthetics are simulated through a viscoelastic, anisotropic PREM model, SPECIFEM3D_GLOBE synthetics through the S362ANI tomographic model (Kustowski et al., 2008), and both sets are shifted by cross-correlation traveltimes to align with the respective phases (left: Pdiff, right: PKiKP). Traveltime shifts are about 2-6 seconds (see main text).

off-plane scattering and focusing) and realistic Earth rotation. All other limitations (lack of ocean layer, gravity and topography) mentioned here or in the code reflect the current stage of the algorithm, but pose no fundamental restriction.

6.1 Future additions

Current and future extensions of the presented methodology include low-frequency effects like gravitation (Chaljub & Valette, 2004), internal and external topography, and a local-scale version of the method. Sensitivity kernels upon AxiSEM also deliver the basis for scattering solutions to wave propagation, which may then allow for considering mild effects of 3D (Born) scattering, which can be applied to both 3D volumetric and boundary topography. AxiSEM-generated wavefields may also be injected into a small 3D box of local 3D heterogeneities in a hybrid sense (Tong et al., 2013). This will allow for the consideration of teleseismic wavefields to locally travel through 3D heterogeneities (Masson et al., 2013), for instance beneath a dense seismic array above a tectonically active region such as USArray in Western USA or the Pyrenees (Monteiller et al., 2012). It may

further be useful to attempt a cost-accuracy benefit analysis across various wave-propagation codes which cover a sensible overlapping parameter space. This is a non-trivial task, as efficiency highly depends on the actual problem at hand (frequency range, distance range, number of sources, number of receivers, multi-scale models, source complexity, solid-fluid domains etc).

Acknowledgements. Reviews by Guust Nolet and Nobuaki Fuji improved the manuscript. Financial support has been provided by the EU Framework program 7 (QUEST network), and the HP2C project “Petaquake” funded by the Swiss National Fund. Simulations were done at ETH’s computing facilities and CSCS in Switzerland, as well as HECToR (UK) whose support is appreciated. The contribution of AF is IGP contribution 3514. SCS was funded by the Munich Centre for Advanced Computing (MAC). We are indebted to Jean-Paul Ampuero, Lapo Boschi, Emmanuel Chaljub, Tony Dahlen, Andreas Fichtner, Domenico Giardini, Heiner Igel, Dimitri Komatitsch, Guust Nolet, Karin Sigloch, Frederik Simons, Christine Thomas and Jeroen Tromp for constructive ideas, suggestions, and support which greatly improved AxiSEM over the past years. The online resource for AxiSEM is (including source tar-

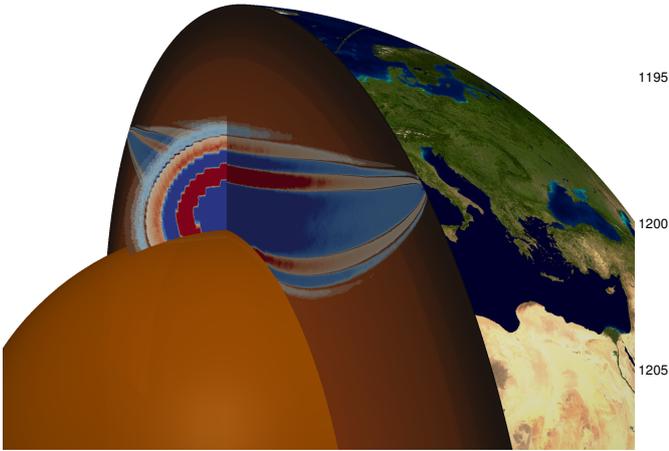


Fig. 17. A sensitivity kernel computed with wavefields from AxiSEM, time-integrated (time window: 20 s) over the arrival of the direct P-wave at 90° distance for a dominant period of 10 s. This denotes the “region of influence” in which this particular time-window in the seismogram may “see” compressional structure which deviates from the background model. Such kernels are not only the basis for waveform tomography, but may also aid in identifying obscure arrivals in the seismogram.

balls, manual etc): <http://www.axisem.info>, and development of the code takes place on GitHub, the latest version can be downloaded under Linux or MacOS with the terminal command:

```
git clone https://github.com/geodynamics/axisem.git
```

References

- 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.
- Auer, L., Boschi, L., Becker, T., Nissen-Meyer, T., & Giardini, D., 2013. Savani: a variable-resolution whole-mantle model of anisotropic shear-velocity variations based on multiple datasets, *J. Geophys. Res.*, under review.
- Baig, A. M. & Dahlen, F. A., 2003. Traveltimes of waves in three-dimensional random media, *Geophysical Journal International*, **153**(2), 467–482.
- Bassin, C., Laske, G., & Masters, G., 2000. The current limits of resolution for surface wave tomography in North America, *Eos*, **81**.
- Becker, T. W. & Boschi, L., 2002. A comparison of tomographic and geodynamic mantle models, *Geochem. Geophys. Geosys.*, **3**, 2001GC000168.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J., 2010. ObsPy: A Python Toolbox for Seismology, *SRL*, **81**(3), 530–533.
- Boué, P., Poli, P., Campillo, M., Pedersen, H., Briand, X., & Roux, P., 2013. Teleseismic correlations of ambient seismic noise for deep global imaging of the earth, *Geophysical Journal International*, **194**, 844–848.
- Bozdag, E., Trampert, J., & Tromp, J., 2011. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements, *Geophys. J. Int.*, **185**, 845–870.
- Capdeville, Y., Stutzmann, E., Wang, N., & Montagner, J.-P., 2013. Residual homogenization for seismic forward and inverse problems in layered media, *Geophysical Journal International*, **194**(1), 470–487.
- Chaljub, E. & Valette, B., 2004. Spectral element modeling of three dimensional wave propagation in a self-gravitating Earth with an arbitrarily stratified outer core, *Geophys. J. Int.*, **158**, 131–141.
- Chaljub, E., Komatitsch, D., Vilotte, J.-P., Capdeville, Y., Valette, B., & Festa, G., 2007. Spectral-element analysis in seismology, *Advances in Geophysics*, **48**, 365–419.
- Colombi, A., Nissen-Meyer, T., Boschi, L., & Giardini, D., 2012. Seismic waveform sensitivity to global boundary topography, *Geophysical Journal International*, **191**(2), 832–848.
- Colombi, A., Nissen-Meyer, T., Boschi, L., & Giardini, D., 2014. Seismic waveform inversion for core-mantle boundary topography, *Geophysical Journal International*, p. 10.1093/gji/ggu112.
- Cottaar, S. & Romanowicz, B., 2013. Observations of changing anisotropy across the southern margin of the african LLSVP, *Geophys. J. Int.*, **195**, 1184–1195.
- Crotwell, H. P., Owens, T. J., & Ritsema, J., 1999. The taup toolkit: Flexible seismic travel-time and ray-path utilities, *Seismological Research Letters*, **70**(2), 154–160.
- Dahlen, F. A., 1968. The Normal Modes of a Rotating, Elliptical Earth, *Geophys. J. Roy. Astr. Soc.*, **16**, 329–367.
- Dahlen, F. A. & Tromp, J., 1998. *Theoretical Global Seismology*, Princeton University Press, Princeton.
- Dziewonski, A. & Anderson, D., 1981. Preliminary Reference Earth Model, *Phys. Earth Planet. Inter.*, **25**, 297–356.
- Fichtner, A. & Igel, H., 2008. Efficient numerical surface wave propagation through the optimization of discrete crustal models - a technique based on non-linear dispersion curve matching (DCM), *Geophys. J. Int.*, **173**(2), 519–533.
- Fichtner, A., Kennett, B. L. N., Igel, H., & Bunge, H.-P., 2008. Theoretical background for continental and global scale full-waveform inversion in the time-frequency domain, *Geophys. J. Int.*, **175**(2), 665–685.
- Fournier, A., Bunge, H.-P., Hollerbach, R., & Vilotte, J.-P., 2005. A Fourier-spectral element algorithm for thermal convection in rotating axisymmetric containers, *J. Comp. Phys.*, **204**, 462–489.
- 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**, 537–550.
- Frigo, M. & Johnson, S. G., 2005. The design and implementation of fftw3, *Proc. IEEE*, **93**, 216–231.
- Fuji, N., Chevrot, S., Zhao, L., Geller, R. J., & Kawai, K., 2012. Finite-frequency structural sensitivities of short-period compressional body waves, *Geophysical Journal International*, **190**(1), 522–540.
- Gizon, L. & Birch, A. C., 2002. Time-distance helioseismology: The forward problem for random distributed sources, *The Astrophysical Journal*, **571**(2), 966–986.
- Hung, S.-H., Dahlen, F. A., & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes—II. Examples, *Geophys. J. Int.*

- 141, 175–203.
- Igel, H. & Weber, M., 1996. P-SV wave propagation in the Earth's mantle using finite differences: Application to heterogeneous lowermost mantle structure, *Geophys. Res. Lett.*, **23**, 415–418.
- 1255 Igel, H., Nissen-Meyer, T., & Jahnke, G., 2002. Wave propagation in 3D spherical sections: Effects of subduction zones, *Phys. Earth Planet. Inter.*, **132**, 219–234.
- Jahnke, G., Thorne, M. S., Cochard, A., & Igel, H., 2008. Global sh-wave propagation using a parallel axisymmetric spherical finite-difference scheme: application to whole mantle scattering, *Geophys. J. Int.*, **173**, 815–826.
- 1260 Jeffreys, H. & Bullen, K. E., 1940. *Seismological Tables*, British Association for the Advancement of Science, London.
- Kawai, K., Takeuchi, N., & Geller, R. J., 2006. Complete synthetic seismograms up to 2 Hz for transversely isotropic spherically symmetric media, *Geophysical Journal International*, **164**(2), 411–424.
- 1265 Kennel, M. B., 2004. KDTree 2: Fortran 95 and c++ software to efficiently search for near neighbors in a multi-dimensional euclidean space, *Arxiv Physics Archive 0408067*.
- 1270 Kennett, B. & Engdahl, E., 1991. Traveltimes for global earthquake location and phase identification, *Geophys. J. Int.*, **105**, 429–465.
- Kennett, B., Engdahl, E., & Buland, R., 1995. Constraints on seismic velocities in the earth from traveltimes, *Geophys. J. Int.*, **122**, 108–124.
- 1275 Kennett, B. L. N. & Gudmundsson, O., 1996. Ellipticity corrections for seismic phases, *Geophysical Journal International*, **127**(1), 40–48.
- Komatitsch, D. & Tromp, J., 2002a. Spectral-element simulations of global seismic wave propagation—I. Validation, *Geophys. J. Int.*, **149**, 390–412.
- 1280 Komatitsch, D. & Tromp, J., 2002b. Spectral-element simulations of global seismic wave propagation—II. 3-D models, oceans, rotation, and self-gravitation, *Geophys. J. Int.*, **150**, 303–318.
- 1285 Kristeková, M., Kristek, J., & Moczo, P., 2009. Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals, *Geophysical Journal International*, **178**(2), 813–825.
- Kustowski, B., Ekström, G., & Dziewonski, A. M., 2008. Anisotropic shear-wave velocity structure of the Earth's mantle: A global model, *Journal of Geophysical Research: Solid Earth*, **113**(B6), 2156–2202.
- 1290 Masson, Y., Cupillard, P., Capdeville, Y., & Romanowicz, B., 2013. On the numerical implementation of time-reversal mirrors for tomographic imaging, *Geophysical Journal International*.
- 1295 Mercerat, E. & Nolet, G., 2013. On the linearity of cross-correlation delay times in finite-frequency tomography, *Geophys. J. Int.*, **192**, 681–687.
- 1300 Monteiller, V., Chevrot, S., Komatitsch, D., & Fuji, N., 2012. A hybrid method to compute short-period synthetic seismograms of teleseismic body waves in a 3-d regional model, *Geophysical Journal International*, **192**, 230–247.
- Nissen-Meyer, T., Dahlen, F. A., & Fournier, A., 2007a. Spherical-earth Fréchet sensitivity kernels, *Geophys. J. Int.*, **168**, 1051–1066.
- 1305 Nissen-Meyer, T., Fournier, A., & Dahlen, F. A., 2007b. A 2-D spectral-element method for computing spherical-earth seismograms—I. Moment-tensor source, *Geophys. J. Int.*, **168**, 1067–1093.
- Nissen-Meyer, T., Fournier, A., & Dahlen, F. A., 2008. A 2-D spectral-element method for computing spherical-earth seismograms—II. Waves in solid-fluid media, *Geophys. J. Int.*, **174**, 873–888.
- Nolet, G., 2008. *A Breviary of Seismic Tomography: Imaging the Interior of the Earth and the Sun*, Cambridge University Press, Cambridge.
- Page, M., Mai, P. M., & Schorlemmer, D., 2011. Testing earthquake source inversion methodologies, *Eos*, **92**, No. 9.
- Peter, D., Komatitsch, D., Luo, Y., Martin, R., Le Goff, N., Casarotti, E., Le Locher, P., Magnoni, F., Liu, Q., Blitz, C., Nissen-Meyer, T., P., B., & Tromp, J., 2011. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes, *Geophys. J. Int.*, **186**, 721–739.
- Rawlinson, N., Pozgaya, S., & Fishwick, S., 2010. Seismic tomography: A window into deep Earth, *Phys. Earth Planet. Inter.*, **178**, 101–135.
- Rew, R. & Davis, G., 1990. NetCDF: an interface for scientific data access, *Computer Graphics and Applications, IEEE*, **10**(4), 76–82.
- Rietmann, M., Messmer, P., Nissen-Meyer, T., Peter, D., Basini, P., Komatitsch, D., Schenk, O., Tromp, J., Boschi, L., & Giardini, D., 2012. Forward and adjoint simulations of seismic wave propagation on emerging large-scale gpu architectures, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '12*, pp. 1–11, IEEE Computer Society Press, Los Alamitos, CA, USA.
- Ritsema, J., Deuss, A., van Heijst, H. J., & Woodhouse, J. H., 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. & Gung, Y., 2002. Superplumes from the core-mantle boundary to the lithosphere: Implications for heat flux, *Science*, **296**(5567), 513–516.
- Rost, S., 2013. Deep earth: Core-mantle boundary landscapes, *Nature Geoscience*, **6**(2), 89–90.
- shu Yeh, P., Xia-serafino, W., Miles, L., Kobler, B., & Menasce, D., 2002. Implementation of ccscs lossless data compression in hdf, in *Earth Science Technology Conference–2002*.
- Stähler, S. C. & Sigloch, K., 2013. Fully probabilistic seismic source inversion – Part 1: Efficient parametrisation, *Solid Earth Discussions*, **5**(2), 1125–1162.
- 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.
- Sun, D., Tan, E., Helmberger, D., & Gurnis, M., 2007. Seismological support for the metastable superplume model, sharp features, and phase changes within the lower mantle, *Proceedings of the National Academy of Sciences*, **104**(22), 9151–9155.
- Takeuchi, N., Geller, R., & Cummins, P., 2000. Complete synthetic seismograms for 3-D heterogeneous Earth models computed using modified DSM operators and their applicability to inversion for Earth structure, *Phys. Earth Planet. Inter.*, **119**, 25–36.
- Thomas, C., Igel, H., Weber, M., & Scherbaum, F., 2000. Acoustic simulation of P-wave propagation in a heterogeneous spherical Earth: Numerical method and application to precursor energy to PKPdf, *Geophys. J. Int.*, **141**, 307–320.

- 1370 Thorne, M., Lay, T., Garnero, E., G., J., & Igel, H., 2007. Seismic
imaging of the laterally varying d'' region beneath the cocos plate,
Geophys. J. Int., **170**, 635–648.
- Thorne, M., Crotwell, H., & Jahnke, G., 2013. An educational re-
source for visualizing the global seismic wavefield, *Seismologi-
cal Research Letters*, **84**(4), 711–717.
- 1375 Tong, P., Chen, C.-w., Komatitsch, D., Basini, P., & Liu, Q., 2013.
High-resolution seismic array imaging based on a sem-fk hybrid
method, *Geophysical Journal International*, in press.
- Toyokuni, G. & Takenaka, H., 2006. FDM computation of seismic
wavefield for an axisymmetric earth with a moment tensor point
source, *Earth Planets Space*, **58**, e29–e32.
- 1380 Toyokuni, G. & Takenaka, H., 2012. Accurate and efficient model-
ing of global seismic wave propagation for an attenuative earth
model including the center, *Phys. Earth Planet. Inter.*, **200-201**,
45–55.
- Tromp, J., Tape, C., & Liu, Q., 2005. Seismic tomography, adjoint
1385 methods, time reversal and banana-doughnut kernels, *Geophys.
J. Int.*, **160**, 195–216.
- Tromp, J., Komatitsch, D., Hjörleifsdóttir, V., Liu, Q., Zhu, H., Pe-
ter, D., Bozdog, E., McRitchie, D., Friberg, P., Trabant, C., &
Hutko, A., 2010. Near real-time simulations of global cmt earth-
1390 quakes, *Geophysical Journal International*, **183**(1), 381–389.
- Van Driel, M. & Nissen-Meyer, T., 2014. Seismic wave propaga-
tion in fully anisotropic axisymmetric media, *Geophys. J. Int.*,
submitted.
- Van Driel, M. & Nissen-Meyer, T., 2014. iefficient visco-elastic
1395 wave propagation with optimized material parameters and coarse
grained memory variables, *Geophys. J. Int.*, to be submitted.
- Zhao, D., 2001. Seismic structure and origin of hotspots and mantle
plumes, *Earth Planet. Sci. Lett.*, **192**(3), 251–265.
- 1400 Zhu, H., Luo, Y., Nissen-Meyer, T., Morency, C., & Tromp, J.,
2009. Elastic imaging and time-lapse migration based upon ad-
joint methods, *Geophysics*, **74**, WCA167–WCA177.