

Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes

Daniel Peter,¹ Dimitri Komatitsch,^{2,3} Yang Luo,¹ Roland Martin,² Nicolas Le Goff,² Emanuele Casarotti,⁴ Pieyre Le Loher,² Federica Magnoni,⁴ Qinya Liu,⁵ Céline Blitz,² Tarje Nissen-Meyer,⁶ Piero Basini⁶ and Jeroen Tromp^{1,7}

¹Princeton University, Department of Geosciences, 318 Guyot Hall, Princeton, NJ 08544, USA. E-mail: dpeter@princeton.edu

²Université de Pau et des Pays de l'Adour, CNRS & INRIA Magique-3D, Laboratoire de Modélisation et d'Imagerie en Géosciences UMR 5212, Avenue de l'Université, 64013 Pau Cedex, France

³Institut universitaire de France, 103 boulevard Saint-Michel, 75005 Paris, France

⁴Istituto Nazionale di Geofisica e Vulcanologia, Via di Vigna Murata 605, 00143, Rome, Italy

⁵Department of Physics, University of Toronto, Ontario, Canada

⁶Institute of Geophysics, ETH Zurich, Sonneggstr. 5, CH-8092 Zurich, Switzerland

⁷Princeton University, Program in Applied & Computational Mathematics, Princeton, NJ 08544, USA

Accepted 2011 April 12. Received 2011 April 8; in original form 2011 February 2

SUMMARY

We present forward and adjoint spectral-element simulations of coupled acoustic and (an)elastic seismic wave propagation on fully unstructured hexahedral meshes. Simulations benefit from recent advances in hexahedral meshing, load balancing and software optimization. Meshing may be accomplished using a mesh generation tool kit such as CUBIT, and load balancing is facilitated by graph partitioning based on the SCOTCH library. Coupling between fluid and solid regions is incorporated in a straightforward fashion using domain decomposition. Topography, bathymetry and Moho undulations may be readily included in the mesh, and physical dispersion and attenuation associated with anelasticity are accounted for using a series of standard linear solids. Finite-frequency Fréchet derivatives are calculated using adjoint methods in both fluid and solid domains. The software is benchmarked for a layercake model. We present various examples of fully unstructured meshes, snapshots of wavefields and finite-frequency kernels generated by Version 2.0 ‘Sesame’ of our widely used open source spectral-element package SPECFEM3D.

Key words: Tomography; Interferometry; Computational seismology; Wave propagation.

1 INTRODUCTION

We present a new software package, SPECFEM3D Version 2.0 ‘Sesame’, capable of simulating forward and adjoint seismic wave propagation on fully unstructured hexahedral meshes of arbitrary shaped model domains. In view of unrelenting growth in computational power, it has become more and more important to develop software capable of harnessing powerful computers to address a broad range of seismological forward and inverse problems. A well-established numerical technique for solving such problems in a fast and highly accurate manner is the spectral-element method (SEM). The SEM was originally developed in computational fluid dynamics (Patera 1984; Maday & Patera 1989) and has been successfully adapted to address problems in seismic wave propagation. Early seismic wave propagation applications of the SEM, utilizing Legendre basis functions and a perfectly diagonal mass matrix, include Cohen *et al.* (1993), Komatitsch (1997), Faccioli *et al.* (1997), Casadei & Gabellini (1997), Komatitsch & Vilotte (1998) and

Komatitsch & Tromp (1999), whereas applications involving Chebyshev basis functions and a non-diagonal mass matrix include Seriani & Priolo (1994), Priolo *et al.* (1994) and Seriani *et al.* (1995).

The SEM is a continuous Galerkin technique, which may be made discontinuous (Bernardi *et al.* 1994; Chaljub 2000; Kopriva *et al.* 2002; Chaljub *et al.* 2003; Legay *et al.* 2005; Kopriva 2006; Wilcox *et al.* 2010; Acosta Minolia & Kopriva 2011); it is then close to a particular case of the discontinuous Galerkin technique (Reed & Hill 1973; Arnold 1982; Falk & Richter 1999; Hu *et al.* 1999; Cockburn *et al.* 2000; Giraldo *et al.* 2002; Rivière & Wheeler 2003; Monk & Richter 2005; Grote *et al.* 2006; Ainsworth *et al.* 2006; Bernacki *et al.* 2006; Dumbser & Käser 2006; De Basabe *et al.* 2008; de la Puente *et al.* 2009; Wilcox *et al.* 2010; De Basabe & Sen 2010; Etienne *et al.* 2010), with optimized efficiency because of its tensorized basis functions (Wilcox *et al.* 2010; Acosta Minolia & Kopriva 2011).

An important feature of the SEM is that it can accurately handle very distorted mesh elements (Oliveira & Seriani 2011), and

thus conforming non-structured mesh doubling bricks can efficiently accommodate mesh size variations (Komatitsch & Tromp 2002a, 2004; Lee *et al.* 2008, 2009a,b). The method has very good accuracy and convergence properties, such as a spectral rate of convergence (Canuto *et al.* 1988; Maday & Patera 1989; Seriani & Priolo 1994; Deville *et al.* 2002; Cohen 2002; De Basabe & Sen 2007; Seriani & Oliveira 2008). In this sense the SEM is close to the family of pseudo-spectral methods (see e.g. Canuto *et al.* 1988; Carcione *et al.* 1988a, 1992; Carcione & Wang 1993; Komatitsch *et al.* 1996), but combined with the flexibility of finite elements, in particular in terms of mesh design. For reviews of the SEM in seismology, see for example, Komatitsch *et al.* (2005), Chaljub *et al.* (2007), Tromp *et al.* (2008) and Fichtner (2010).

The SEM is well suited to parallel implementations on very large supercomputers (Komatitsch & Tromp 2002a; Komatitsch *et al.* 2003; Tsuboi *et al.* 2003; Komatitsch *et al.* 2008; Carrington *et al.* 2008; Komatitsch *et al.* 2010b) as well as on clusters of GPU accelerating graphics cards (Komatitsch *et al.* 2009, 2010a, Komatitsch 2011). Tensor products inside each element may be optimized to reach very high efficiency (Deville *et al.* 2002), and mesh point and element numbering may be optimized to reduce processor cache misses and improve cache reuse (Komatitsch *et al.* 2008). The SEM can handle triangular (in 2-D) or tetrahedral (in 3-D) elements (Wingate & Boyd 1996; Taylor & Wingate 2000; Komatitsch *et al.* 2001; Cohen 2002; Mercerat *et al.* 2006), as well as mixed meshes, although with increased cost and reduced accuracy in these non-tensorized elements, as in the discontinuous Galerkin method.

In many cases of practical seismological interest, using a conforming mesh and a continuous formulation are sufficient, because in most geological models material property contrasts are not too dramatic. When this ceases to be true, requiring a discontinuous formulation, one can either turn to a discontinuous version of the SEM (Bernardi *et al.* 1994; Chaljub 2000; Kopriva *et al.* 2002; Chaljub *et al.* 2003; Legay *et al.* 2005; Kopriva 2006; Wilcox *et al.* 2010; Acosta Minolia & Kopriva 2011) or to a discontinuous Galerkin technique. A discontinuous formulation is particularly suitable for dynamic rupture simulations, because high frequencies or super-shear rupture need to be accommodated near the fault, where a significantly denser mesh and a more sophisticated (upwind) time scheme are required, thereby suppressing the amplification of unstable modes (see e.g. Benjema *et al.* 2007, 2009; de la Puente *et al.* 2009; Tago *et al.* 2010). Another example that may require a discontinuous formulation involves the resolution of a shallow geotechnical layer, in which seismic shear wave speeds may be reduced by an order of magnitude.

For seismological applications, the SEM has been successfully implemented for 3-D global- and regional-scale simulations (Komatitsch & Vilotte 1998; Paolucci *et al.* 1999; Chaljub 2000; Komatitsch & Tromp 2002a,b; Capdeville *et al.* 2003; Chaljub & Valette 2004; Fichtner *et al.* 2009a), as well as local-scale simulations in complex and/or densely populated regions, for example in southern California, USA (Komatitsch *et al.* 2004; Tape *et al.* 2009, 2010), Taipei, Taiwan (Lee *et al.* 2008, 2009a,b), Caracas, Venezuela (Delavaud *et al.* 2006) and Grenoble, France (Chaljub *et al.* 2005; Stupazzini *et al.* 2009; Chaljub *et al.* 2010). The SEM may also be used to study elastic wave propagation on smaller scales, for instance the propagation of ultrasonic waves in crystals (van Wijk *et al.* 2004).

Two complementary SEM software packages—namely, SPEC-FEM3D_GLOBE for global and regional simulations,

and SPEC-FEM3D for local simulations—are feature-rich, well benchmarked and documented implementations. Data parallelism in the SEM is efficiently exploited using the Message-Passing Interface (MPI) standard, crucial for modern high-performance computing. These open source packages are freely available via the Computational Infrastructure for Geodynamics (CIG) and widely used by the seismological community.

To extend the range of local-scale applications, easing the task of mesh generation is paramount. The two community software packages separate a simulation into two distinct steps: first, creation of a hexahedral mesh, and second, solution of the seismic wave equation. This separation avoids the overhead of remeshing when running multiple simulations for the same region, for example, repeated simulations at the same resolution. Focussing on local-scale simulations, previous versions of SPEC-FEM3D used an internal mesher which was explicitly tied to the specific purposes of the package: all geological models were based on a layercake model. Consequently, the solver was restricted by its internal mesher. It was impossible to run spectral-element simulations on more complex 3-D models without significant recoding, nor was it possible to run such simulations in regions of interest for on- and off-shore exploration seismology, because acoustic wave propagation in fluids was not supported by the package.

The purpose of this paper is to present forward and adjoint simulations in various 3-D models using the new software package, SPEC-FEM3D Version 2.0 ‘Sesame’, thereby illustrating its current capabilities. The original SPEC-FEM3D package for local simulations was extended, improved and optimized in various ways. The Version 2.0 ‘Sesame’ release includes a more flexible internal mesher and accommodates more powerful external meshers, such as CUBIT (Blacker *et al.* 1994; White *et al.* 1995; Mitchell 1996; Casarotti *et al.* 2008). Adding such external meshers into the workflow greatly increases flexibility for high-performance applications, as illustrated by the GeoELSE software package (Casadei & Gabellini 1997; Stupazzini *et al.* 2009; Chaljub *et al.* 2010). Advantages of GeoELSE include the accommodation of viscoplastic and non-linear rheologies, whereas benefits of SPEC-FEM3D include coupled fluid-solid domains and adjoint capabilities; the latter enable one to address seismological inverse problems. Load balancing parallel simulations in SPEC-FEM3D is accomplished based on the graph partitioning software package SCOTCH (Pellegrini & Roman 1996; Chevalier & Pellegrini 2008). The new package facilitates coupled forward and adjoint acoustic/(an)elastic simulations, which are especially interesting for problems in exploration seismology, ocean acoustics and medical tomography. The new software is freely available under the GNU GPL Version 2 license via CIG.

2 GOVERNING EQUATIONS

Let us briefly summarize the equations governing seismic wave propagation implemented in SPEC-FEM3D. For more technical details, the reader is referred to Komatitsch & Tromp (1999). SPEC-FEM3D Version 2.0 ‘Sesame’ implements wave propagation in coupled (an)elastic and acoustic materials on local scales. We may thus safely neglect additional effects that would arise from self-gravitation and rotation (Komatitsch & Tromp 2002b, 2005; Chaljub *et al.* 2007), which are important at longer periods. In the following, we first discuss (an)elastic wave propagation and subsequently consider acoustic waves.

2.1 Elastic domain

For elastic materials, the displacement wavefield $\mathbf{s}(\mathbf{x}, t)$ is governed by

$$\rho \partial_t^2 \mathbf{s} = \nabla \cdot \mathbf{T} + \mathbf{f}, \quad (1)$$

where ρ denotes mass density, \mathbf{T} the stress tensor and \mathbf{f} the seismic source. On free surfaces, the traction vector must vanish, that is,

$$\hat{\mathbf{n}} \cdot \mathbf{T} = \mathbf{0}, \quad (2)$$

where $\hat{\mathbf{n}}$ denotes the unit outward normal on the surface. On boundaries between different elastic materials, both traction $\hat{\mathbf{n}} \cdot \mathbf{T}$ and displacement \mathbf{s} need to be continuous. On boundaries between elastic and acoustic domains, traction $\hat{\mathbf{n}} \cdot \mathbf{T}$ and the normal component of displacement $\hat{\mathbf{n}} \cdot \mathbf{s}$ need to be continuous. The initial conditions are

$$\mathbf{s}(\mathbf{x}, 0) = \mathbf{0}, \quad \partial_t \mathbf{s}(\mathbf{x}, 0) = \mathbf{0}. \quad (3)$$

We thus initiate the simulation in a medium at rest. To accommodate simulations under pre-stressed conditions, these initial conditions may be modified in an appropriate manner.

For elastic materials, the force \mathbf{f} in eq. (1) represents the earthquake, which for a simple point source may be written as

$$\mathbf{f} = -\mathbf{M} \cdot \nabla \delta(\mathbf{x} - \mathbf{x}_s) S(t), \quad (4)$$

where \mathbf{M} denotes the moment tensor, \mathbf{x}_s the source location, $\delta(\mathbf{x} - \mathbf{x}_s)$ the Dirac delta distribution located at \mathbf{x}_s and $S(t)$ the source-time function. The software also accommodates kinematic rupture simulations, which may be captured by prescribing a moment-density tensor field.

The stress tensor \mathbf{T} is linearly related to the strain via the constitutive relationship

$$\mathbf{T} = \mathbf{c} : \nabla \mathbf{s}, \quad (5)$$

where \mathbf{c} denotes the stiffness tensor that describes the elastic properties of the medium. The implementation is general and can handle a fully anisotropic tensor with 21 independent parameters (Chen & Tromp 2007; Sieminski *et al.* 2007a,b). Using a linear constitutive relationship is valid under the assumption that perturbations to the reference state are small. Note that non-linear effects are sometimes observed, for example, non-linear soil amplification, and non-linear constitutive relationships become important for studying such effects, for example, for risk mitigation (Xu *et al.* 2003; Dupros *et al.* 2010).

In an anelastic medium, we approximate an absorption-band solid using a series of L standard linear solids (Liu *et al.* 1976), and model the time evolution of the isotropic shear modulus μ by

$$\mu(t) = \mu_R \left[1 - \sum_{l=1}^L \left(1 - \frac{\tau_l^\epsilon}{\tau_l^\sigma} \right) e^{-t/\tau_l^\sigma} \right] H(t), \quad (6)$$

where μ_R denotes the relaxed modulus, $H(t)$ the Heaviside function and τ_l^ϵ & τ_l^σ the stress and strain relaxation times of the l th standard linear solid. Experience shows that three solids generally suffice for simulating an absorption band (Emmerich & Korn 1987). For further details, see Carcione *et al.* (1988b), Robertsson (1996), Day & Bradley (2001), Moczo & Kristek (2005), Komatitsch *et al.* (2005), Carcione (2007) and Savage *et al.* (2010). Simulations of seismic wave propagation in laboratory-scale rock samples or in the context of medical tomography involve very high frequencies (in the kHz or even MHz range), and strong attenuation must be taken into account.

The SEM solves the equations of motion in the weak form, which is obtained by dotting the momentum eq. (1) with an arbitrary test

vector \mathbf{w} and integrating by parts over the model volume Ω . We focus on elastic domains and consider coupling interfaces with acoustic domains. Thus, we obtain

$$\int_{\Omega} \rho \mathbf{w} \cdot \partial_t^2 \mathbf{s} d^3 \mathbf{x} = \int_{\partial\Omega} \hat{\mathbf{n}} \cdot \mathbf{T} \cdot \mathbf{w} d^2 \mathbf{x} - \int_{\Omega} \nabla \mathbf{w} : \mathbf{T} d^3 \mathbf{x} + \mathbf{M} : \nabla \mathbf{w}(\mathbf{x}_s) S(t). \quad (7)$$

Note that in this formulation the traction-free surface condition is implicitly accounted for by setting the contribution from the free surface to zero.

When and where necessary, we use Clayton–Engquist–Stacey absorbing conditions (Clayton & Engquist 1977; Stacey 1988; Quarteroni *et al.* 1998) to absorb outgoing waves on fictitious boundaries of the mesh, thereby representing a semi-infinite domain. It would be more efficient to use a Perfectly Matched Layer (PML) (see e.g. Komatitsch & Martin 2007, 2008c; Martin & Komatitsch 2009), but a parallel implementation with good load-balancing properties is challenging because additional equations need to be solved. This issue becomes important when high-order time marching is required to reduce numerical dispersion in difficult case studies that involve complex media with poroelastic or viscoelastic rheologies (Martin *et al.* 2008b, 2010) or Newtonian compressible fluids (Martin & Couder-Castaneda 2010). Consequently, additional computations need to be performed in PML layers, in particular in corners, where contributions along several directions are summed (Komatitsch & Martin 2007).

At a solid–fluid boundary, the interface integral over the coupling surface $\partial\Omega$ is used to exchange pressure from the fluid p_{fluid} to the solid: $\hat{\mathbf{n}} \cdot \mathbf{T} = -p_{\text{fluid}} \hat{\mathbf{n}}$.

2.2 Acoustic domain

We define a scalar potential ϕ such that the displacement \mathbf{s} may be written as

$$\mathbf{s} = \rho^{-1} \nabla \phi. \quad (8)$$

The equation of motion in terms of the potential ϕ becomes

$$\kappa^{-1} \partial_t^2 \phi = \nabla \cdot (\rho^{-1} \nabla \phi) + f, \quad (9)$$

where κ denotes the bulk modulus. It follows that velocity \mathbf{v} and pressure p may be expressed as

$$\mathbf{v} = \rho^{-1} \nabla \partial_t \phi, \quad (10)$$

$$p = -\kappa (\nabla \cdot \mathbf{s}) = -\partial_t^2 \phi. \quad (11)$$

The resulting formulation for pressure p is the reason why we choose to define the potential ϕ as in eq. (8). Since pressure is continuous across first-order discontinuities, it follows that $\partial_t^2 \phi$ and thus ϕ must be continuous, a requirement which is honoured automatically by the basis functions of the SEM. The source f may be expressed in terms of pressure P acting at location \mathbf{x}_s .

$$f = -\kappa^{-1} P(t) \delta(\mathbf{x} - \mathbf{x}_s). \quad (12)$$

Note that the source is multiplied by a factor κ^{-1} due to the formulation used in eq. (9).

Using Gauss' theorem and a scalar test function w , the weak form becomes

$$\int_{\Omega} \kappa^{-1} w \partial_t^2 \phi d^3 \mathbf{x} = \int_{\partial\Omega} \rho^{-1} w \hat{\mathbf{n}} \cdot \nabla \phi d^2 \mathbf{x} - \int_{\Omega} \rho^{-1} \nabla w \cdot \nabla \phi d^3 \mathbf{x} - \kappa^{-1} P(t) w(\mathbf{x}_s). \quad (13)$$

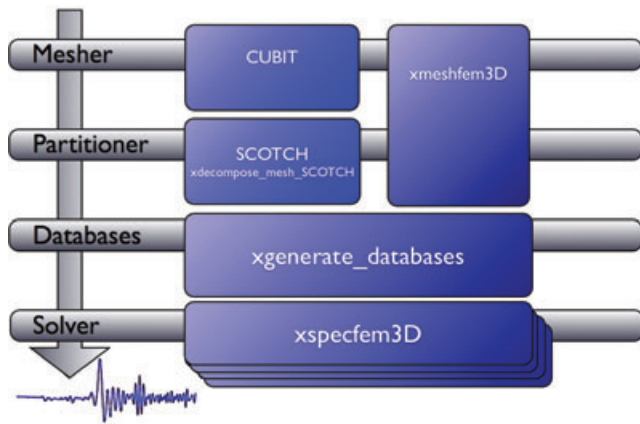


Figure 1. Workflow for running spectral-element simulations with SPECSEM3D Version 2.0 ‘Sesame’.

At the free surface $\partial\Omega$ we set the pressure $p = -\partial_t^2\phi = 0$, thereby enforcing $\phi = 0$, $\partial_t\phi = 0$ and $\partial_t^2\phi = 0$, that is, we implement a Dirichlet boundary condition along the surface. At a fluid–solid boundary, the interface coupling integral may be used to exchange the normal component of displacement between fluid and solid: $\rho^{-1}\hat{\mathbf{n}} \cdot \nabla\phi = \hat{\mathbf{n}} \cdot \mathbf{s}_{\text{solid}}$.

3 MESHING, MESH PARTITIONING AND LOAD BALANCING

The first step in a SEM consists of constructing a high-quality mesh for the region of interest. In this section, we outline the key issues based on various 3-D examples. Fig. 1 draws the schematic workflow from meshing and partitioning to finally running spectral-element simulations. We discuss each phase separately, focussing on the use of an external mesher, in our case CUBIT (Blacker *et al.* 1994).

3.1 Hexahedral meshing

We subdivide the model volume Ω into a set of non-overlapping, hexahedral elements. We impose that the discretization creates a conforming mesh, that is, elements match on a full face or edge, and the mesh cannot be discontinuous. Using the SEM with hexahedral elements leads to computational benefits over tetrahedral finite elements (Komatitsch *et al.* 2001; Mercier *et al.* 2006; Vos *et al.* 2010). Especially for parallel implementations, taking advantage of the diagonal mass matrix and optimized tensor products is critical in terms of computational speed (Komatitsch *et al.* 2003; Carrington *et al.* 2008; Vos *et al.* 2010). Hexahedral meshing is also attractive for the SEM because it benefits from reduced errors and generally smaller element counts compared to tetrahedral meshing (Hesthaven & Teng 2000; Komatitsch *et al.* 2001; Vos *et al.* 2010).

Unfortunately, automatic 3-D hexahedral mesh generation is more demanding than unstructured tetrahedral meshing (Shepherd & Johnson 2008; Staten *et al.* 2010). To construct hexahedral meshes, our examples make use of an external hexahedral mesher, such as CUBIT (Blacker *et al.* 1994). We focus on this particular mesh generation tool kit because it is a well-documented and feature-rich package, on which most of our own experience is based. One may readily use other meshing tools, such as Abaqus (SIMULIA 2008), ANSYS (ANSYS 2011), GOCAD (Mallet 1992; Caumon *et al.* 2009), GiD (Gardia-Donoro *et al.* 2010; Ribó *et al.* 2011), Gmsh (Geuzaine & Remacle 2009), TrueGrid (Noble &

Nuss 2004; Rainsberger 2006) or Salome (Ribes & Caremoli 2007; Bergeaud *et al.* 2010).

Fig. 2 shows several examples of fully unstructured hexahedral meshes. In the Mount St Helens region, the mesh employs a mesh tripling layer to increase resolution at the topographic surface. Tripling is the default refinement in CUBIT for subdividing hexahedral elements in a conforming fashion. Surface topography is imported using Shuttle Radar Topographic Mission (SRTM) data, converted to Universal Transverse Mercator (UTM) coordinates with an original resolution of 90 m (Jarvis *et al.* 2008). Meshing is performed automatically by CUBIT using a sweep algorithm. The resolution of the mesh enables seismic wave simulations with frequencies up to ~ 1.5 Hz. The Mesh for the L’Aquila region, Italy, consists of ~ 7 M hexahedra with an element size of ~ 90 m at the top surface. This mesh facilitates simulations of seismic wave propagation up to ~ 5 Hz. For the exploration geophysics model, the hexahedral mesh honours a salt dome body inside a 3-D model capped by a water layer. The mesh for asteroid 433-Eros with a close-bound surface has a resolution of roughly 300 m. Finally, the filled coffee cup model discretized into hexahedra couples an elastic domain for the cup with an acoustic domain for the coffee inside the cup.

To ensure compatibility with previous versions of SPECSEM3D (see e.g. Komatitsch *et al.* 2004; Liu *et al.* 2004), the in-house mesher based on analytical linear interpolation from the top to the bottom of the mesh has been adapted to the new code structure. It facilitates the design of simpler, alternative meshes for layercake models.

3.2 Partitioning and load balancing

Balancing the computational load and distributing the mesh on a large number of cores is crucial for optimized high-performance simulations (Martin *et al.* 2008a). To do so, we make use of an external partitioner, namely SCOTCH (Pellegrini & Roman 1996; Chevalier & Pellegrini 2008), which we use to balance spectral-element computations on an arbitrary number of cores. An alternative partitioner able to fulfill these tasks is METIS (Karypis & Kumar 1998), but SCOTCH is more actively maintained (Chevalier & Pellegrini 2008) and performs better in many cases that we have tested.

Especially for simulations involving coupled elastic and acoustic domains, balancing the mesh becomes paramount. Most of the computation time is spent resolving the divergence of the stress tensor in each element. The computational cost for an elastic element is approximately four times larger than for an acoustic element, which may be established by running simulations for one domain at a time. During partitioning, we therefore weight each element according to its associated domain type and computational cost to balance the overall numerical cost rather than simply the number of elements between partitions. The major improvement in SPECSEM3D code performance focuses on these tensor products, using highly efficient algorithms developed by Deville *et al.* (2002) and optimizing cache usage. Another key aspect of mesh partitioning is minimization of the number of edge cuts, because this reduces the amount of MPI communications between processor cores (an edge cut occurs when two contiguous elements are assigned to distinct cores). On machines comprising a very large number of cores, it is crucial to resort to non-blocking communications between compute nodes, for instance using non-blocking MPI message passing, to obtain good performance scaling (Danielson & Namburu 1998; Komatitsch *et al.* 2008; Martin *et al.* 2008a).




Figure 2. Mesh examples: (a) Mount St Helens meshed by hexahedral elements. The mesh honours surface topography and includes a mesh tripling layer in the middle of the model. The smallest element size is approximately 280 m. (b) L'Aquila, Italy, region discretized for high-frequency simulations. The mesh honours surface and Moho topography and includes two mesh tripling layers. The yellow and blue volumes denote slower and faster than average wave speeds, respectively. (c) Salt dome body meshed inside an exploration model for a SEG/EAGE benchmark test. (d) 3-D hexahedral mesh of the asteroid 433 Eros. (e) Arbitrarily shaped mesh for coupled solid–fluid simulations involving a coffee cup.

Fig. 3 presents a simple example of partitioning and load balancing the mesh around Mount St Helens, as shown in Fig. 2. For illustrative purposes, we decompose the mesh onto four cores using the SCOTCH library. The total number of spectral-elements is $\sim 24\,000$, such that each partition contains ~ 6000 elements after decomposition. Partitioning and load balancing equally distributes the elements over the different cores, since the whole domain is purely elastic. A partitioner such as SCOTCH can also load balance computationally more complex meshes, for example, containing PML elements along absorbing boundaries of the model; this is the subject of future research.

In a final, separate step we generate mesh databases for each partition needed for the spectral-element solver. These databases contain Gauss–Lobatto–Legendre (GLL) points for all spectral elements. Material properties are assigned to these GLL points, and thus sampling resolution of a geological model not only depends on element size but also on polynomial degree. Furthermore, the generation of mesh databases automatically detects interfaces between elastic and acoustic domains, needed for coupling seismic waves from one domain to another. Load-balancing of the simulation persists, because we keep the polynomial degree fixed for all spectral elements. Note that this final step of generating mesh




Figure 3. Mount St Helens mesh partitioned and load balanced to run in parallel on four cores. The four partitions are indicated by different colours.




Figure 4. Outer (highlighted) and inner (transparent colours) elements for the mesh shown in Fig. 3. Outer elements have at least one point in common with an element from another slice and must therefore be computed first, before initiating non-blocking MPI communications. While MPI messages are travelling across the computer network, simultaneous computations are performed on inner elements. Non-blocking MPI communication is crucial to obtain good scaling results for simulations running on a large number of parallel cores.

databases provides additional freedom in assigning and changing wave speeds, which is important for seismic inversion procedures.

3.3 Overlapping computation and communication

The elements that compose the mesh slices shown in Figs 2 and 3 are in contact through a common face, edge or point. To allow for overlap of communication between compute nodes with computations within each mesh slice—thereby speeding up the simulation—a list of all elements in contact with any other mesh slice through a common face, edge or point is created. Members of this list are termed ‘outer’ elements, and all other elements are termed ‘inner’ elements, as illustrated in Fig. 4.

Once the outer elements have been identified following a standard procedure (see e.g. Danielson & Namburu 1998; Martin *et al.* 2008a; Micikevicius 2009; Michéa & Komatitsch 2010; Komatitsch 2010a, 2011), MPI buffers are filled and a non-blocking MPI call is issued, which initiates communication and returns immediately. While MPI messages are travelling across the network, computations are performed on inner elements. Achieving effective overlap requires that the ratio of the number of inner to outer elements is sufficiently large, a condition which is satisfied for suitably large mesh slices. Under these circumstances, MPI data transfer will generally finish before the completion of computations on inner elements.

4 SAMPLE SIMULATIONS

In this section, we present various simulations with increasing complexity to highlight the flexibility of our new spectral-element

package. We start with a layercake model and finish with an example of an arbitrarily shaped model.

4.1 Validation example: two-layer model

The SEM has been well benchmarked against discrete wavenumber methods for layercake models by Komatitsch & Tromp (1999). Here we compare their two-layer model solution (Fig. 8, left-hand side) against the solution obtained by the new code. The model has a horizontal size of $134 \text{ km} \times 134 \text{ km}$, with a depth of 60 km. We discretize the model into 70 200 elements, using an approximate element size of 1000 m at the top and 4500 m at the bottom. A mesh tripling layer is placed below the upper layer, between 3 km and 10 km, with the wave speed properties of the lower layer. We use SCOTCH to partition the model onto six cores, each with 11 700 elements. The final mesh is generated using GLL points for a polynomial degree $N = 5$, which results in 9 025 941 global mesh points.

A vertical force is placed at a depth of 25.05 km in the middle of the model. The source–time function is a Ricker wavelet with a dominant frequency of 0.4 Hz. The simulation uses a time step of 6.5 ms and propagates for 6000 steps. We compare our solution with seismograms obtained by Komatitsch & Tromp (1999, Fig. 9). The mesh and seismograms are shown in Fig. 5. The seismograms match very closely with the reference solutions, exhibiting almost identical displacements. Maximum waveform differences reach ~ 0.3 per cent, arising from differences in mesh geometry and source implementation.

The performance of the code is summarized in Fig. 6, using simulations with the optimized routines by Deville *et al.* (2002)

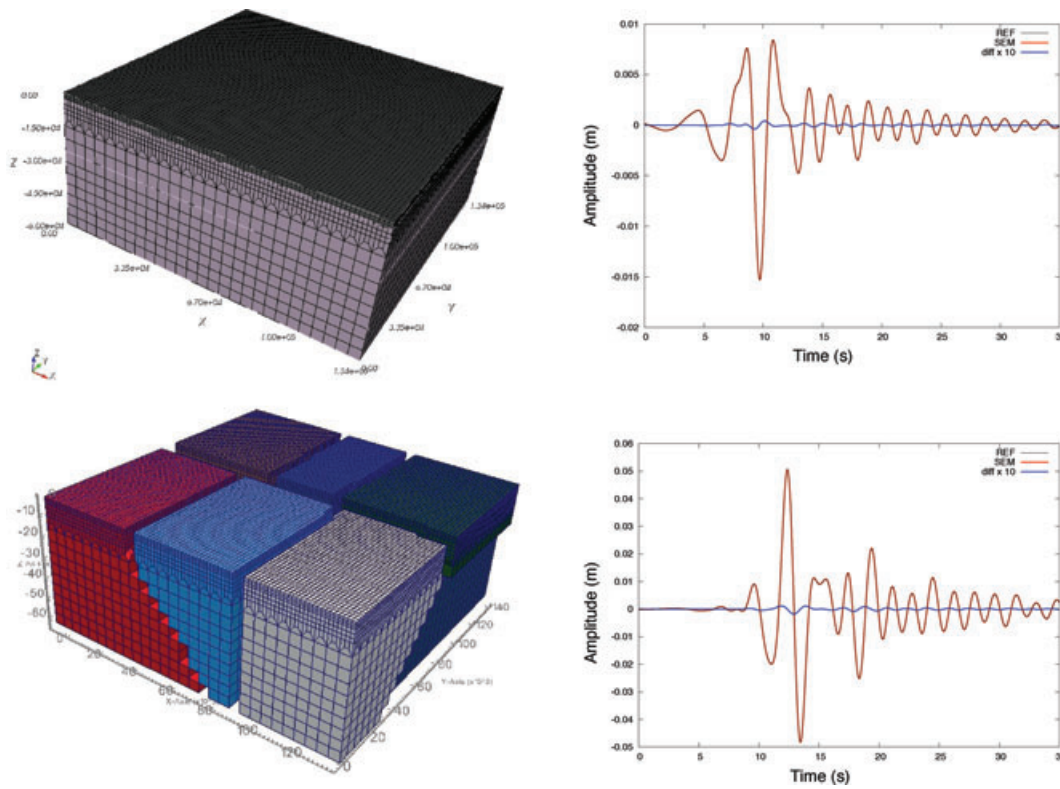


Figure 5. Validation for a two-layer mesh (left, top panel), using six partitions (left, bottom panel), and seismograms recorded at the surface at horizontal distances of 2.39 km (right, top panel) and 31.11 km (right, bottom panel). Plotted are radial displacements (SEM, red) against reference solutions (REF, black) from Komatitsch & Tromp (1999), as well as their exaggerated differences (blue).

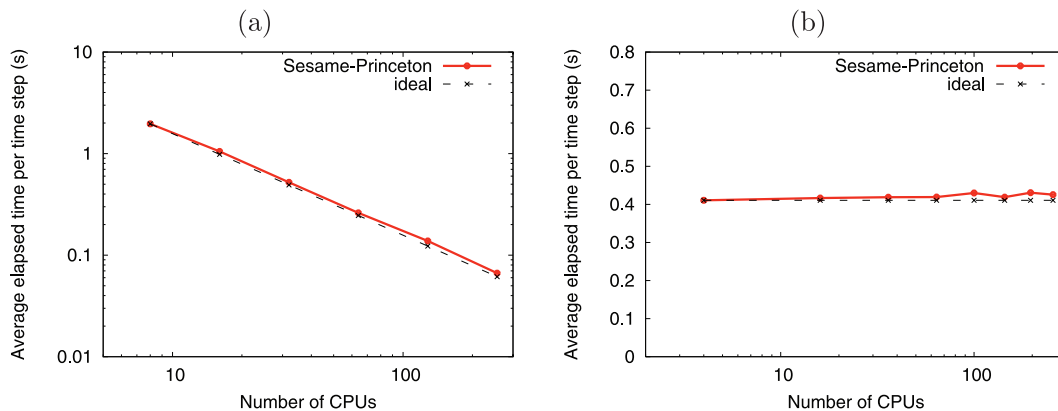


Figure 6. CPU scaling results for the model shown in Fig. 5, (a) using a fixed total problem size (strong scaling) and (b) a fixed problem size per processor (weak scaling) for up to 256 cores. Perfect weak scaling deviates slightly from a straight line, because a larger number of processors involves more MPI buffers and therefore more computational overhead.

and a polynomial degree $N = 4$. We are interested in how the code behaves when the number of calculations is decreased linearly with the number of CPU cores (strong scaling), and how performance varies when the number of calculations on each core is kept constant while increasing the total number of CPU cores (weak scaling). To assess strong scaling, we fix the total mesh size but vary the number of CPU cores used for the simulation. We run the simulation for a duration of 4000 time steps and show the corresponding average elapsed time per time step in Fig. 6(a). More interesting for high-performance applications, we assess weak scaling by fixing the problem size per processor and varying the number of CPU cores. This leads to higher mesh resolutions for an increasing

number of CPU cores but should keep the average elapsed time per time step constant. We summarize the simulation times in Fig. 6(b). The computations were performed on a high-performance cluster with compute nodes consisting of two Intel Nehalem quad-core processors; each core has 3 GB of RAM. The code scales linearly within ~ 90 per cent up to 256 CPU cores for both strong and weak scaling, achieving excellent performance on this parallel system. Note that for the strong scaling examples shown here, simulations using more than 64 CPUs see a performance decrease since communications no longer overlap, thus they no longer profit from the default non-blocking MPI scheme (Martin *et al.* 2008a).




Figure 7. Wavefield snapshots around Mount St Helens. Plotted are vertical displacements (up/down coloured red/blue, respectively) at the free surface of the model.




Figure 8. Wavefield snapshots for the 2009 April 6, L'Aquila earthquake, taken after 6 s, 11 s, 16 s and 21 s. Plotted are vertical displacements (up/down as red/blue).

4.2 Mount St Helens example: layercake model with surface topography

To include surface topography, we import SRTM data with an original resolution of 90 m (Jarvis *et al.* 2008) and convert it to UTM coordinates for the corresponding UTM zone. We read in this data set using CUBIT and create a surface honouring these data points. A 3-D volume is built manually with topography on top.

The simulation uses an explosive source at a depth of 5 km. In Fig. 7, we show the vertical displacement field at the free surface at consecutive times. Note that once the wavefield hits the model boundary, it gets absorbed by the Clayton–Engquist–Stacey absorbing boundary conditions.

4.3 L'Aquila example: layercake model honouring surface and Moho topography

The purpose of this example is to show that additional surfaces may be honoured by the mesh, for example the Moho. We import not only surface topography, but also create a Moho surface that is

honoured by the boundaries of the spectral elements. The mesh for the L'Aquila region was built using an additional 'Python' library that semi-automates the mesh creation process with CUBIT (Casarotti *et al.* 2008). Once these mesh files are constructed, the default partitioning and database generation process may be used to create fully load-balanced spectral-element simulations on an arbitrary number of parallel processors.

Fig. 8 shows several snapshots of the seismic wavefield at consecutive times for an anelastic material, using a kinematic source description for the 2009 April 6, L'Aquila earthquake. Simulations are accurate up to 5 Hz and may be used to discriminate between different wave speed models and/or kinematic source solutions. These high-frequency simulations may be used to assess the response of engineered structures and may guide the development of better seismic building codes for the L'Aquila region.

4.4 SEG/EAGE salt dome example: exploration model

Our new spectral-element package can combine acoustic and (an)elastic simulations by coupling these distinct domains. In this




Figure 9. Wavefield snapshots for an exploration geophysics simulation taken after 5 s, 10 s, 15 s and 20 s. Plotted are vertical velocities at the free surface of the water layer.

example, we generate acoustic waves in the top water layer and propagate them down through a salt dome body included in the lower, anelastic domain. The mesh honours the surface of the salt dome and the fluid–solid boundary, that is, the bathymetry.

Fig. 9 shows the acoustic wavefield at the free surface of the water layer at different times. The source is a pressure source, located slightly below the free surface in the water layer, with a Ricker source–time function. The wavefield is reflected and refracted by the salt dome in the anelastic domain below the water layer. Note how these reflected/refracted waves, which include *P*-to-*S* converted waves, are recorded in the water layer.

4.5 Asteroid example: arbitrarily shaped model

This final example shows that our new software package may be used for simulating wave propagation in arbitrarily shaped models, such as asteroid Eros, which was imaged by the NEAR spacecraft in 2000–2001. This silicated asteroid is 34 km long with a peanut-like shape and is thought to be covered with a regolith layer, corresponding to a blanket of loose material crushed by impacts (Richardson *et al.* 2005). We meshed the asteroid with 5 797 440 hexahedral elements having an approximate resolution of 70 m. To simulate a thin, 70 m regolith layer superimposed on strong bedrock, as suggested by Robinson *et al.* (2002), we assigned a low wave speed material to the elements touching the free surface and a high wave speed material to elements inside the asteroid, representing solid bedrock.

We simulated the propagation of seismic waves from a source represented by a point force normal to the surface. The source–time function corresponds to a Dirac pulse low-pass filtered up to a cut-off frequency of 5 Hz. Fig. 10 displays wavefield snapshots for the first ~10 s of the simulation. It shows the propagation of *P*, *S* and surface waves with a refocusing effect on the opposite side. The regolith

layer strongly increases physical dispersion of surface waves. Peak ground accelerations are plotted in Fig. 11 for a simulation without a regolith layer, showing that refocussing occurs on the asteroid.

5 ADJOINT SENSITIVITY KERNELS

An important goal in seismology is to use differences between observed and simulated seismograms to improve Earth and source models, that is, we are interested in the inverse problem. An elegant way to address this issue is to take advantage of adjoint methods (Tarantola 1984; Tromp *et al.* 2005) to calculate Fréchet derivatives for a pre-defined objective function. These derivatives may then be used in a conjugate-gradient approach to minimize differences between data and synthetics. The key ingredients of such an adjoint approach are sensitivity kernels. Following Tromp *et al.* (2005), Liu & Tromp (2006, 2008) and Tromp *et al.* (2010), we show examples of sensitivity kernels for various models using our new software package.

5.1 Elastic sensitivity kernels

Following Tromp *et al.* (2005), we may write the variation of a misfit function χ as

$$\delta\chi = \int_V (K_\rho \delta \ln \rho + K_{c_{jklm}} \delta c_{jklm}) d^3\mathbf{x}, \quad (14)$$

where $\delta \ln \rho = \delta\rho/\rho$ denotes relative perturbations in density and δc_{jklm} denotes perturbations in the elastic tensor. The misfit kernels are given by

$$K_\rho = -\rho \int_0^T \mathbf{s}^\dagger(T-t) \cdot \partial_t^2 \mathbf{s}(t) dt, \quad (15)$$




Figure 10. Wavefield snapshots for an asteroid simulation taken after 3 s, 4.5 s, 6.5 s and 10.5 s. Plotted is the norm of the velocity at the free surface of the asteroid.




Figure 11. ShakeMap views for an asteroid simulation. Plotted are different views of the peak ground accelerations at the free surface of the asteroid.

$$K_{c_{jklm}} = - \int_0^T \epsilon_{jk}^\dagger(T-t) \epsilon_{lm}(t) dt, \quad (16)$$

where ϵ_{lm} and ϵ_{jk}^\dagger denote elements of the strain and adjoint strain tensors, and where we have suppressed the spatial dependence to avoid clutter.

In an isotropic model, we have $c_{jklm} = (\kappa - 2\mu/3) \delta_{jk} \delta_{lm} + \mu (\delta_{jl} \delta_{km} + \delta_{jm} \delta_{kl})$, and thus eq. (14) may be rewritten as

$$\delta\chi = \int_V (K_\rho \delta \ln \rho + K_\mu \delta \ln \mu + K_\kappa \delta \ln \kappa) d^3 \mathbf{x}. \quad (17)$$

The isotropic misfit kernels K_μ and K_κ are defined by

$$K_\mu = -2\mu \int_0^T \mathbf{D}^\dagger(T-t) : \mathbf{D}(t) dt, \quad (18)$$

$$K_\kappa = -\kappa \int_0^T [\nabla \cdot \mathbf{s}^\dagger(T-t)] [\nabla \cdot \mathbf{s}(t)] dt, \quad (19)$$

where $\mathbf{D} = \frac{1}{2}[\nabla \mathbf{s} + (\nabla \mathbf{s})^T] - \frac{1}{3}(\nabla \cdot \mathbf{s}) \mathbf{I}$ and $\mathbf{D}^\dagger = \frac{1}{2}[\nabla \mathbf{s}^\dagger + (\nabla \mathbf{s}^\dagger)^T] - \frac{1}{3}(\nabla \cdot \mathbf{s}^\dagger) \mathbf{I}$ are the traceless strain deviator and its adjoint, respectively. In terms of a parametrization involving compressional wave speed α , shear wave speed β and density ρ , the corresponding kernels are given by a linear combination of these primary kernels

(Tromp *et al.* 2005, eq. 20):

$$\begin{aligned} K_\alpha &= 2 \left(\frac{\kappa + \frac{4}{3}\mu}{\kappa} \right) K_\kappa, \\ K_\beta &= 2 \left(K_\mu - \frac{4}{3} \frac{\mu}{\kappa} K_\kappa \right), \quad K'_\rho = K_\kappa + K_\mu + K_\rho. \end{aligned} \quad (20)$$

Note that a suitable parametrization for isotropic inversions is to use bulk sound wave speed $\Phi = \sqrt{\kappa/\rho}$, shear wave speed β and density ρ (Tarantola 1987). Bulk sound and shear wave speeds are independent combinations of the bulk and shear moduli κ and μ . The corresponding kernels are given by

$$K_\Phi = 2K_\kappa, \quad K'_\beta = 2K_\mu, \quad K'_\rho = K_\kappa + K_\mu + K_\rho. \quad (21)$$

We place an explosive source at a depth of 7 km and a horizontal distance of 16 km from the receiver in a homogeneous model with topography around Mount St Helens. The P wave at the receiver is used to construct a traveltime adjoint source for the kernel simulation. Fig. 12(a) shows the isotropic kernels K_κ , K_μ and K_ρ , and Fig. 12(b) the isotropic kernels K_α , K_β and K'_ρ for the same model and source–receiver geometry.

Note that although we construct the adjoint source using the P wave, significant non-zero S -wave sensitivity is visible in the K_β and K_μ kernels. We interpret these areas of high sensitivity as




Figure 12. Traveltime sensitivity to elastic structure. Fréchet derivatives for isotropic parametrizations (a) K_κ , K_μ & K_ρ and (b) K_α , K_β & K'_ρ are compared in a model of Mount St Helens using traveltime adjoint sources for the P wave. Shown are vertical cross-sections through the source–receiver line and perpendicular to this line.

P -to- S scattering locations, which affect the signal within the chosen time window. As may be observed in Fig. 12, such scattering sensitivity is especially strong at the free surface close to the receiver.

5.2 Acoustic sensitivity kernels

Liu & Tromp (2008) calculated global sensitivity kernels, which include sensitivity to the liquid outer core. In this section, we present acoustic sensitivity kernels for general local- or regional-scale models. Such kernels may be used, for example, in ocean acoustics, non-destructive testing and medical tomography.

For acoustic simulations, the kernels are given by

$$K_\rho = \rho^{-1} \int_0^T [\nabla \partial_t \phi^\dagger(T-t)] \cdot [\nabla \partial_t \phi(t)] dt, \quad (22)$$

$$K_\kappa = -\kappa^{-1} \int_0^T \partial_t^2 \phi^\dagger(T-t) \partial_t^2 \phi(t) dt, \quad (23)$$

where ϕ and ϕ^\dagger denote the acoustic scalar potential and adjoint potential, respectively. To illustrate these kernels, we use a model with acoustic and elastic regions. The model combines a water layer on top of a homogeneous elastic layer, separated by a bathymetric surface. The dimensions of the model volume are approximately $2 \text{ km} \times 2 \text{ km}$ horizontally and 1 km in depth. Bathymetry is taken from a location in front of Pearl Harbor (Hawaii, USA), with a resolution of $\sim 11 \text{ m}$. For the forward simulation, we use a pressure source in the form of an explosion with a Gaussian source–time function, and record pressure variations at the receiver. Both source and receiver are in the acoustic domain at a depth of 10 m and $\sim 1.1 \text{ km}$ apart from each other. We use the simulated pressure

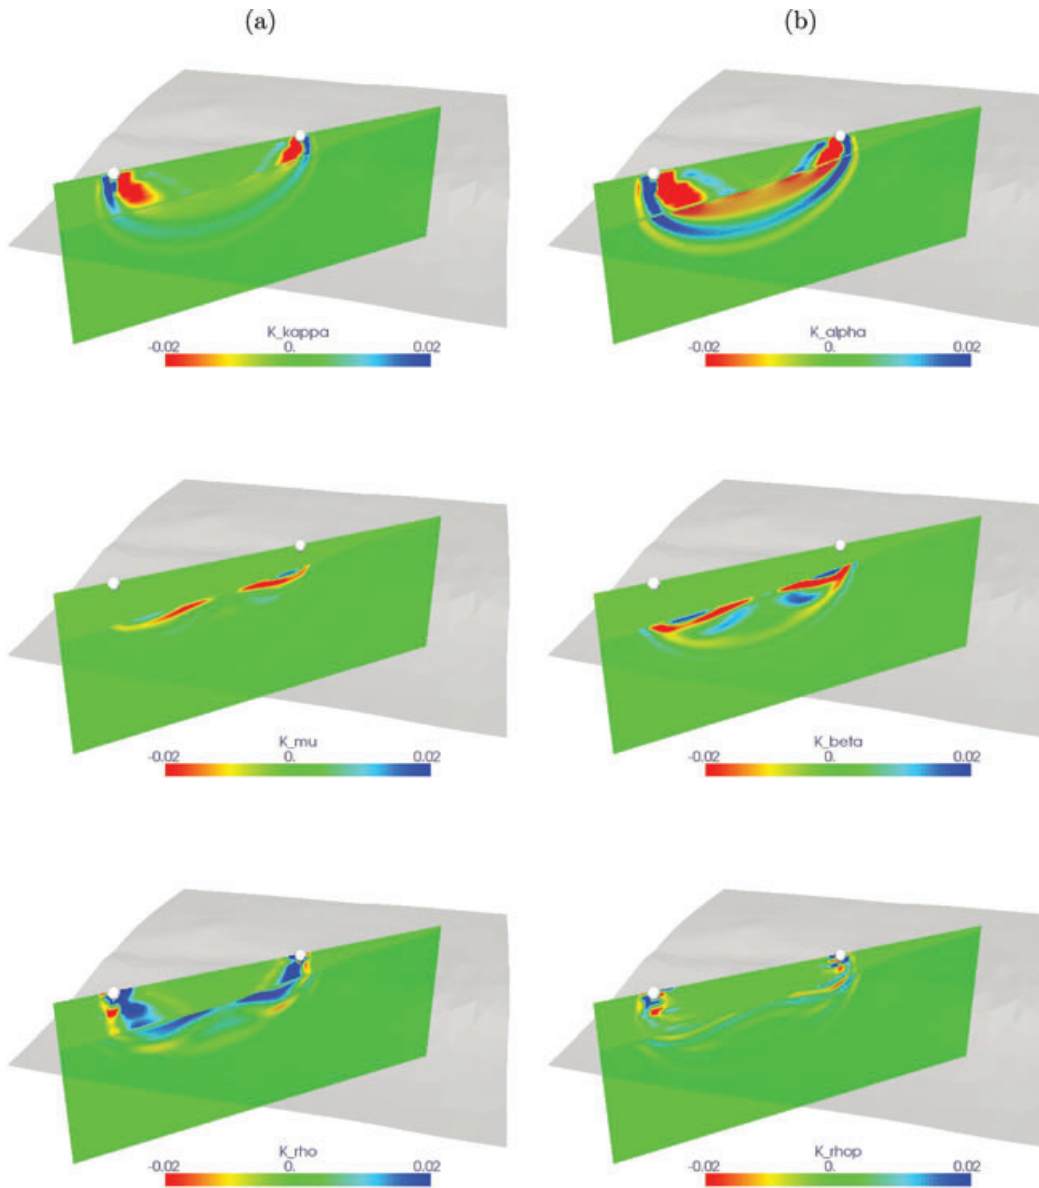


Figure 13. Waveform sensitivity to acoustic and elastic structure in a coupled fluid–solid simulation. The bathymetric surface of the Pearl Harbor model, separating the two media, is shown in grey together with a vertical cross-section through source (right-hand side) and station (left-hand side). Plotted are combined acoustic and elastic kernels using a parametrization (a) K_κ , K_μ & K_ρ and (b) K_α , K_β & K'_ρ .

variation within the measurement window as the pressure misfit for the adjoint source, as explained in Section A1 of the Appendix.

Fig. 13 shows the corresponding combined acoustic and elastic kernels. The kernels highlight how the pressure waveform in the chosen measurement window is affected by a head wave (a Scholte wave) travelling along the seafloor. Since the acoustic region does not support shear waves, the kernels K_μ and K_β are zero in this upper domain. However, they do exhibit non-zero sensitivity in the elastic domain, due to P -to- S coupling along the seafloor.

5.3 Noise sensitivity kernels

As demonstrated by Tromp *et al.* (2010), noise cross-correlation sensitivity kernels may also be calculated based on an adjoint method, and the new package has the necessary capabilities to perform such calculations. Consider two receivers located at \mathbf{x}^α and \mathbf{x}^β . In seismic interferometry, ensemble sensitivity kernels are

calculated in terms of interactions between an ensemble forward wavefield Φ^α , generated at location \mathbf{x}^α , and an ensemble adjoint wavefield $\Phi^{\dagger\alpha\beta}$, generated at \mathbf{x}^β and triggered by the differences between simulated and observed ensemble-averaged cross correlations at \mathbf{x}^α and \mathbf{x}^β . The isotropic ensemble sensitivity kernels are given by

$$\langle K_\rho \rangle = - \int \rho [\Phi^{\dagger\alpha\beta}(-t) \cdot \partial_t^2 \Phi^\alpha(t) + \Phi^{\dagger\beta\alpha}(-t) \cdot \partial_t^2 \Phi^\beta(t)] dt, \quad (24)$$

$$\langle K_\mu \rangle = - \int 2\mu [\mathbf{D}^{\dagger\alpha\beta}(-t) : \mathbf{D}^\alpha(t) + \mathbf{D}^{\dagger\beta\alpha}(-t) : \mathbf{D}^\beta(t)] dt, \quad (25)$$

$$\langle K_\kappa \rangle = - \int \kappa [\nabla \cdot \Phi^{\dagger\alpha\beta}(-t) \nabla \cdot \Phi^\alpha(t) + \nabla \cdot \Phi^{\dagger\beta\alpha}(-t) \nabla \cdot \Phi^\beta(t)] dt, \quad (26)$$




Figure 14. Noise cross-correlation sensitivity to elastic structure. Shown are (a) first, (b) second and (c) summed contributions to the $\langle K_\beta \rangle$ Fréchet derivative in a homogeneous isotropic model. Plotted are vertical and horizontal cross-sections through the line connecting the two receivers (white dots) and perpendicular to this line.

where

$$\mathbf{D}^\alpha = \frac{1}{2} [\nabla \Phi^\alpha + (\nabla \Phi^\alpha)^T] - \frac{1}{3} (\nabla \cdot \Phi^\alpha) \mathbf{I}, \quad (27)$$

$$\mathbf{D}^{\dagger\alpha\beta} = \frac{1}{2} [\nabla \Phi^{\dagger\alpha\beta} + (\nabla \Phi^{\dagger\alpha\beta})^T] - \frac{1}{3} (\nabla \cdot \Phi^{\dagger\alpha\beta}) \mathbf{I}, \quad (28)$$

denote the traceless ensemble strain deviator and corresponding adjoint.

Fig. 14 shows the isotropic kernel $\langle K_\beta \rangle$ calculated according to eq. (20) using the primary isotropic ensemble sensitivity kernels given. Plotted are the two contributions from ensemble wavefields generated at the first receiver location, \mathbf{x}^α , and the second receiver location, \mathbf{x}^β , as well as the combined ensemble sensitivity kernel $\langle K_\beta \rangle$, which is the sum of the two contributions. The two receivers are placed at a distance of 65 km from each other on top of a homogeneous block model with dimensions of 134 km \times 134 km horizontally and 60 km in depth. We smooth the kernel contributions using a 3-D Gaussian filter with a standard deviation of 2 km in the horizontal and vertical directions. Note that these noise sensitivity kernels exhibit strong 3-D variability. Depth sensitivity is controlled by the period range (5–100 s in this example).

6 CONCLUSIONS AND FUTURE WORK

We have taken advantage of recent advances in high-performance computing, fully unstructured hexahedral meshing, load balancing and mesh partitioning to facilitate forward and adjoint simulations of seismic wave propagation in coupled fluid and solid domains. Our new open source software package, SPECFEM3D Version 2.0 ‘Sesame’, performs acoustic and (an)elastic simulations of seismic wave propagation in complex geological models. Hexahedral meshes may be generated based on packages such as CUBIT, Abaqus, ANSYS, GOCAD, GiD, Gmsh, TrueGrid or

Salome, but the simple in-house mesher used in previous versions of SPECFEM3D remains available for backcompatibility.

Partitioning and load balancing meshes may be accomplished based on graph partitioning software, such as SCOTCH. By coupling acoustic and (an)elastic wave propagation, we are able to calculate related sensitivity kernels, which are useful for waveform inversions in off-shore exploration seismology, ocean acoustics, non-destructive testing and medical tomography. Attenuation is important on all scales of seismic wave propagation and is accommodated based on a series of standard linear solids. In particular for simulations in medical tomography, strong attenuation and related dispersion play a dominant role. Finally, the new package can be used to calculate finite-frequency noise cross-correlation sensitivity kernels, which may be used for seismic interferometry. In future work, we will add C-PML and GPU support to the package. Viscoplastic and non-linear elastic rheologies (e.g. Xu *et al.* 2003; di Prisco *et al.* 2007) are accommodated by the GeoELSE software package (Stupazzini *et al.* 2009; Chaljub *et al.* 2010), and we will consider such non-linear constitutive relationships in future releases of SPECFEM3D.

The next grand challenge involves the development of 3-D seismic imaging and inversion tools for the characterization of earthquakes, Earth ‘noise’ and mapping of the Earth’s interior on all scales, that is, to address the seismological inverse problem. The goal is to harness the power of forward and adjoint modelling tools, such as SPECFEM3D, together with modern computers to enhance the quality of images of the Earth’s interior and the earthquake rupture process. Most traditional tomographic methods utilize traveltimes or phase information measured by comparing data with simulations, and interpret such measurements based on ray theory or other approximate methods. Because of the limitations of these approximate theories, only parts of seismograms can be used, and initial models are generally restricted to be spherically symmetric. With the new generation of modelling tools we can go beyond

classical tomography by using fully 3-D initial models (e.g. Akçelik *et al.* 2002, 2003; Askan *et al.* 2007; Chen *et al.* 2007; Fichtner *et al.* 2009b; Fichtner 2010), and utilizing as much information contained in seismograms as possible (e.g. Maggi *et al.* 2009; Valentine & Woodhouse 2010). Our approach will be to minimize frequency-dependent phase and amplitude differences between simulated and observed seismograms based on adjoint techniques in combination with conjugate gradient methods, an approach we refer to as ‘adjoint tomography’ (Tape *et al.* 2009, 2010). The development of such capabilities will affect the fields of exploration geophysics, regional & global seismology, ocean acoustics, non-destructive testing, medical tomography and helioseismology.

ACKNOWLEDGMENTS

We thank two anonymous reviewers for comments which improved the manuscript. All simulations were performed on a Dell cluster built and maintained by the Princeton Institute for Computational Science & Engineering (USA). The open source spectral-element software package SPECFEM3D Version 2.0 ‘Sesame’ used for this paper is available for download via CIG. 3-D graphics were produced with the open source visualization application Paraview. We acknowledge support by the US National Science Foundation under grant EAR-0711177.

REFERENCES

- Acosta Minolia, C.A. & Kopriva, D.A., 2011. Discontinuous Galerkin spectral element approximations on moving meshes, *J. Comput. Phys.*, **230**(5), 1876–1902.
- Ainsworth, M., Monk, P. & Muniz, W., 2006. Dispersive and dissipative properties of discontinuous Galerkin finite element methods for the second-order wave equation, *J. Sci. Comput.*, **27**(1), 5–40.
- Akçelik, V., Biros, G. & Ghattas, O., 2002. Parallel multiscale Gauss–Newton–Krylov methods for inverse wave propagation, *Proc. ACM/IEEE Supercomputing SC’2002 Conference*, CD-ROM. Available at www.sc-conference.org/sc2002 (last accessed 2011 May 2).
- Akçelik, V. *et al.*, 2003. High resolution forward and inverse earthquake modeling on terascale computers, *Proc. ACM/IEEE Supercomputing SC’2003 conference*, CD-ROM. Available at www.sc-conference.org/sc2003.
- ANSYS, 2011. *ANSYS Academic Research Workbench Platform*, release 12.1 edn, ANSYS, Inc., Canonsburg, PA.
- Arnold, D.N., 1982. An interior penalty finite element method with discontinuous elements, *SIAM J. Numer. Anal.*, **19**(4), 742–760.
- Askan, A., Akçelik, V., Bielak, J. & Ghattas, O., 2007. Full waveform inversion for seismic velocity and anelastic losses in heterogeneous structures, *Bull. seism. Soc. Am.* **97**(6), 1990–2008.
- Benjema, M., Glinisky-Olivier, N., Cruz-Atienza, V.M., Virieux, J. & Piperno, S., 2007. Dynamic non-planar crack rupture by a finite volume method, *Geophys. J. Int.*, **171**(1), 271–285.
- Benjema, M., Glinisky-Olivier, N., Cruz-Atienza, V.M. & Virieux, J., 2009. 3D dynamic rupture simulation by a finite volume method, *Geophys. J. Int.*, **178**(1), 541–560.
- Bergeaud, V., Lefebvre, V. & Rossignon, E., 2010. *SALOME The Open Source Integration Platform for Numerical Simulation*, Open Cascade, Guyancourt, France.
- Bernacki, M., Lanteri, S. & Piperno, S., 2006. Time-domain parallel simulation of heterogeneous wave propagation on unstructured grids using explicit, nondiffusive, discontinuous Galerkin methods, *J. Comput. Acoust.*, **14**(1), 57–81.
- Bernardi, C., Maday, Y. & Patera, A.T., 1994. A new nonconforming approach to domain decomposition: the Mortar element method, in *Nonlinear Partial Differential Equations and their Applications*, Séminaires du Collège de France, pp. 13–51, Pitman, Paris.
- Blacker, T. *et al.*, 1994. *CUBIT Mesh Generation Environment Users Manual*, Vol. 1, Sandia National Laboratories, Albuquerque, NM.
- Canuto, C., Hussaini, M.Y., Quarteroni, A. & Zang, T.A., 1988. *Spectral Methods in Fluid Dynamics*, Springer-Verlag, New York, NY.
- Capdeville, Y., Chaljub, E., Vilotte, J.P. & Montagner, J.P., 2003. Coupling the spectral element method with a modal solution for elastic wave propagation in global Earth models, *Geophys. J. Int.*, **152**, 34–67.
- Carcione, J.M., 2007. *Wave Fields in Real Media: Theory and Numerical Simulation of Wave Propagation in Anisotropic, Anelastic, Porous and Electromagnetic Media*, 2nd edn, Elsevier Science, Amsterdam.
- Carcione, J.M. & Wang, P.J., 1993. A Chebyshev collocation method for the wave equation in generalized coordinates, *Comp. Fluid Dyn. J.*, **2**, 269–290.
- Carcione, J.M., Kosloff, D. & Kosloff, R., 1988a. Wave propagation simulation in an elastic anisotropic (transversely isotropic) solid, *Q. J. Mech. appl. Math.*, **41**(3), 319–345.
- Carcione, J.M., Kosloff, D. & Kosloff, R., 1988b. Wave propagation simulation in a linear viscoelastic medium, *Geophys. J. Int.*, **95**, 597–611.
- Carcione, J.M., Kosloff, D., Behle, A. & Seriani, G., 1992. A spectral scheme for wave propagation simulation in 3-D elastic-anisotropic media, *Geophysics*, **57**(12), 1593–1607.
- Carrington, L., Komatitsch, D., Laurenzano, M., Tikir, M., Michéa, D., Le Goff, N., Snively, A. & Tromp, J., 2008. High-frequency simulations of global seismic wave propagation using SPECFEM3D_GLOBE on 62 thousand processor cores, in *Proceedings of the ACM/IEEE Supercomputing SC’2008 Conference*, Austin, TX, pp. 1–11, Article #60.
- Casadei, F. & Gabellini, E., 1997. Implementation of a 3D coupled spectral element solver for wave propagation and soil-structure interaction simulations, Tech. rep., *European Commission Joint Research Center Report EUR17730EN*, Ispra, Italy.
- Casarotti, E., Stupazzini, M., Lee, S., Komatitsch, D., Piersanti, A. & Tromp, J., 2008. Cubit and seismic wave propagation based upon the spectral-element method: An advanced unstructured mesher for complex 3D geological media, in *Proceedings of the 16th International Meshing Roundtable*, Seattle, WA, Vol. 5B.4, pp. 579–597.
- Cauman, G., Collon-Drouaillet, P., le Carlier de Veslund, C., Sausse, J. & Viseur, S., 2009. Surface-based 3D modeling of geological structures, *Math. Geosci.*, **41**, 927–945.
- Chaljub, E., 2000. Modélisation numérique de la propagation d’ondes sismiques en géométrie sphérique : application à la sismologie globale (Numerical modeling of the propagation of seismic waves in spherical geometry: application to global seismology), *PhD thesis*, Université Paris VII Denis Diderot, Paris.
- Chaljub, E. & Valette, B., 2004. Spectral element modelling of three-dimensional wave propagation in a self-gravitating Earth with an arbitrarily stratified outer core, *Geophys. J. Int.*, **158**, 131–141.
- Chaljub, E., Capdeville, Y. & Vilotte, J.P., 2003. Solving elastodynamics in a fluid-solid heterogeneous sphere: a parallel spectral-element approximation on non-conforming grids, *J. Comput. Phys.*, **187**(2), 457–491.
- Chaljub, E., Cornou, C., Guéguen, P., Causse, M. & Komatitsch, D., 2005. Spectral-element modeling of 3D wave propagation in the alpine valley of Grenoble, France, *Geophys. Res. Abstr.*, **7**, 05225.
- Chaljub, E., Komatitsch, D., Vilotte, J.P., Capdeville, Y., Valette, B. & Festa, G., 2007. Spectral element analysis in seismology, in *Advances in Wave Propagation in Heterogeneous Media*, Vol. 48: Advances in Geophysics, pp. 365–419, eds Wu, R.-S. & Maupin, V., Elsevier/Academic Press, London.
- Chaljub, E., Moczo, P., Tsuno, S., Bard, P.-Y., Kristek, J., Käser, M., Stupazzini, M. & Kristekova, M., 2010. Quantitative comparison of four numerical predictions of 3D ground motion in the Grenoble valley, France, *Bull. seism. Soc. Am.*, **100**(4), 1427–1455.
- Chen, M. & Tromp, J., 2007. Theoretical and numerical investigations of global and regional seismic wave propagation in weakly anisotropic earth models, *Geophys. J. Int.*, **168**(3), 1130–1152.

- Chen, P., Zhao, L. & Jordan, T.H., 2007. Full 3D tomography for the crustal structure of the Los Angeles region, *Bull. seism. Soc. Am.*, **97**(4), 1094–1120.
- Chevalier, C. & Pellegrini, F., 2008. PT-SCOTCH: a tool for efficient parallel graph ordering, *Parallel Comput.*, **34**(6–8), 318–331.
- Clayton, R. & Engquist, B., 1977. Absorbing boundary conditions for acoustic and elastic wave equations, *Bull. seism. Soc. Am.*, **67**, 1529–1540.
- Cockburn, B., Karniadakis, G.E. & Shu, C.-W., 2000. *Discontinuous Galerkin Methods: Theory, Computation and Applications*, Springer, Heidelberg.
- Cohen, G., 2002. *Higher-Order Numerical Methods for Transient Wave Equations*, Springer-Verlag, Berlin.
- Cohen, G., Joly, P. & Tordjman, N., 1993. Construction and analysis of higher-order finite elements with mass lumping for the wave equation, in *Proceedings of the Second International Conference on Mathematical and Numerical Aspects of Wave Propagation*, pp. 152–160, SIAM, Philadelphia, PA.
- Danielson, K.T. & Namburu, R.R., 1998. Nonlinear dynamic finite element analysis on parallel computers using Fortran90 and MPI, *Adv. Eng. Softw.*, **29**(3–6), 179–186.
- Day, S.M. & Bradley, C., 2001. Memory-efficient simulation of anelastic wave propagation, *Bull. seism. Soc. Am.*, **91**, 520–531.
- De Basabe, J.D. & Sen, M.K., 2007. Grid dispersion and stability criteria of some common finite-element methods for acoustic and elastic wave equations, *Geophysics*, **72**(6), T81–T95.
- De Basabe, J.D. & Sen, M.K., 2010. Stability of the high-order finite elements for acoustic or elastic wave propagation with high-order time stepping, *Geophys. J. Int.*, **181**(1), 577–590.
- De Basabe, J.D., Sen, M.K. & Wheeler, M.F., 2008. The interior penalty discontinuous Galerkin method for elastic wave propagation: grid dispersion, *Geophys. J. Int.*, **175**(1), 83–93.
- Delavaud, E., Cupillard, P., Festa, G. & Vilotte, J.-P., 2006. 3D spectral element method simulations of the seismic response in the Caracas basin, in *Proc. of the Third International Symposium on the Effects of Surface Geology on Seismic Motion*, Vol. 1, pp. 512–522, Grenoble, France.
- Deville, M.O., Fischer, P.F. & Mund, E.H., 2002. *High-Order Methods for Incompressible Fluid Flow*, Cambridge University Press, Cambridge.
- Dumbser, M. & Käser, M., 2006. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes-II. The three-dimensional isotropic case, *Geophys. J. Int.*, **167**(1), 319–336.
- Dupros, F., de Martin, F., Foerster, E., Komatitsch, D. & Roman, J., 2010. High-performance finite-element simulations of seismic wave propagation in three-dimensional non linear inelastic geological media, *Parallel Comput.*, **36**(5–6), 308–325.
- Emmerich, H. & Korn, M., 1987. Incorporation of attenuation into time-domain computations of seismic wave fields, *Geophysics*, **52**, 1252–1264.
- Étienne, V., Chaljub, E., Virieux, J. & Glinsky, N., 2010. An hp-adaptive discontinuous Galerkin finite-element method for 3-D elastic wave modelling, *Geophys. J. Int.*, **183**(2), 941–962.
- Faccioli, E., Maggio, F., Paolucci, R. & Quarteroni, A., 1997. 2D and 3D elastic wave propagation by a pseudo-spectral domain decomposition method, *J. Seismol.*, **1**, 237–251.
- Falk, R.S. & Richter, G.R., 1999. Explicit finite element methods for symmetric hyperbolic equations, *SIAM J. Numer. Anal.*, **36**(3), 935–952.
- Fichtner, A., 2010. *Full Seismic Waveform Modelling and Inversion*, Springer-Verlag, Heidelberg.
- Fichtner, A., Igel, H., Bunge, H.-P. & Kennett, B.L.N., 2009a. Simulation and inversion of seismic wave propagation on continental scales based on a spectral-element method, *J. Numer. Anal., Ind. Appl. Math.*, **4**(1–2), 11–22.
- Fichtner, A., Kennett, B.L.N., Igel, H. & Bunge, H.-P., 2009b. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, *Geophys. J. Int.*, **179**, 1703–1725.
- Gardia-Donoro, D., Garcia-Castillo, L.E. & Gómez-Revuelto, I., 2010. An interface between an hp-adaptive finite element package and the pre- and post-processor GiD, *Finite Elem. Anal. Des.*, **46**(4), 328–338.
- Geuzaine, C. & Remacle, J.F., 2009. Gmsh: A three-dimensional finite element mesh generator with built-in pre- and post-processing facilities, *Int. J. Numer. Meth. Eng.*, **79**(11), 1309–1331.
- Giraldo, F.X., Hesthaven, J.S. & Warburton, T., 2002. Nodal high-order discontinuous Galerkin methods for the spherical shallow water equations, *J. Comput. Phys.*, **181**(2), 499–525.
- Grote, M.J., Schneebeli, A. & Schötzau, D., 2006. Discontinuous Galerkin finite element method for the wave equation, *SIAM J. Numer. Anal.*, **44**(6), 2408–2431.
- Hesthaven, J.S. & Teng, C.H., 2000. Stable spectral methods on tetrahedral elements, *SIAM J. Sci. Comput.*, **21**, 2352–2380.
- Hu, F.Q., Hussaini, M.Y. & Rasetarinera, P., 1999. An analysis of the discontinuous Galerkin method for wave propagation problems, *J. Comput. Phys.*, **151**(2), 921–946.
- Hung, S.-H., Dahlen, F.A. & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes—II. Examples, *Geophys. J. Int.*, **141**, 175–203.
- Jarvis, A., Reuter, H., Nelson, A. & Guevara, E., 2008. Hole-filled seamless SRTM data V4. Tech. rep., International Centre for Tropical Agriculture (CIAT). Available at: <http://srtm.csi.cgiar.org> (last accessed 2011 May 2).
- Karypis, G. & Kumar, V., 1998. A parallel algorithm for multilevel graph partitioning and sparse matrix ordering, *J. Parallel Distrib. Comput.*, **48**, 71–85.
- Komatitsch, D., 1997. Méthodes spectrales et éléments spectraux pour l'équation de l'élastodynamique 2D et 3D en milieu hétérogène (Spectral and spectral-element methods for the 2D and 3D elastodynamics equations in heterogeneous media), *PhD thesis*, Institut de Physique du Globe, Paris.
- Komatitsch, D., 2011. Fluid-solid coupling on a cluster of GPU graphics cards for seismic wave propagation, *Comptes Rendus de l'Académie des Sciences - Mécanique*, **339**, 125–135, doi:10.1016/j.crme.2010.11.007.
- Komatitsch, D. & Martin, R., 2007. An unsplit convolutional Perfectly Matched Layer improved at grazing incidence for the seismic wave equation, *Geophysics*, **72**(5), SM155–SM167.
- Komatitsch, D. & Tromp, J., 1999. Introduction to the spectral-element method for 3-D seismic wave propagation, *Geophys. J. Int.*, **139**(3), 806–822.
- Komatitsch, D. & Tromp, J., 2002a. Spectral-element simulations of global seismic wave propagation—I. Validation, *Geophys. J. Int.*, **149**(2), 390–412.
- Komatitsch, D. & Tromp, J., 2002b. Spectral-element simulations of global seismic wave propagation—II. 3-D models, oceans, rotation, and self-gravitation, *Geophys. J. Int.*, **150**(1), 303–318.
- Komatitsch, D. & Vilotte, J.P., 1998. The spectral-element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures, *Bull. seism. Soc. Am.*, **88**(2), 368–392.
- Komatitsch, D., Coutel, F. & Mora, P., 1996. Tensorial formulation of the wave equation for modelling curved interfaces, *Geophys. J. Int.*, **127**(1), 156–168.
- Komatitsch, D., Martin, R., Tromp, J., Taylor, M.A. & Wingate, B.A., 2001. Wave propagation in 2-D elastic media using a spectral element method with triangles and quadrangles, *J. Comput. Acoust.*, **9**(2), 703–718.
- Komatitsch, D., Tsuboi, S., Ji, C. & Tromp, J., 2003. A 14.6 billion degrees of freedom, 5 teraflops, 2.5 terabyte earthquake simulation on the Earth Simulator, *Proceedings of the ACM/IEEE Supercomputing SC'2003 Conference*, Phoenix, AZ, pp. 4–11.
- Komatitsch, D., Liu, Q., Tromp, J., Süss, P., Stidham, C. & Shaw, J.H., 2004. Simulations of ground motion in the Los Angeles basin based upon the spectral-element method, *Bull. seism. Soc. Am.*, **94**(1), 187–206.
- Komatitsch, D., Tsuboi, S. & Tromp, J., 2005. The spectral-element method in seismology, *Geophys. Monogr.*, **157**, 205–228.
- Komatitsch, D., Labarta, J. & Michéa, D., 2008. A simulation of seismic wave propagation at high resolution in the inner core of the Earth on 2166 processors of MareNostrum, *Lecture Notes in Computer Science*, **5336**, 364–377.
- Komatitsch, D., Michéa, D. & Erlebacher, G., 2009. Porting a high-order finite-element earthquake modeling application to NVIDIA graphics cards using CUDA, *J. Parallel Distrib. Comput.*, **69**(5), 451–460.

- Komatitsch, D., Erlebacher, G., Göddeke, D. & Michéa, D., 2010a. High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster, *J. Comput. Phys.*, **229**(20), 7692–7714.
- Komatitsch, D., Vinnik, L.P. & Chevrot, S., 2010b. SHdiff/SVdiff splitting in an isotropic Earth, *J. geophys. Res.*, **115**(B7), B07312, doi:10.1029/2009JB006795.
- Kopriva, D.A., 2006. Metric identities and the discontinuous spectral element method on curvilinear meshes, *J. Sci. Comput.*, **26**(3), 301–327.
- Kopriva, D.A., Woodruff, S.L. & Hussaini, M.Y., 2002. Computation of electromagnetic scattering with a non-conforming discontinuous spectral element method, *Int. J. Numer. Meth. Eng.*, **53**(1), 105–122.
- Lee, S.J., Chen, H.W., Liu, Q., Komatitsch, D., Huang, B.S. & Tromp, J., 2008. Three-dimensional simulations of seismic wave propagation in the Taipei basin with realistic topography based upon the spectral-element method, *Bull. seism. Soc. Am.*, **98**(1), 253–264.
- Lee, S.J., Chan, Y.C., Komatitsch, D., Huang, B.S. & Tromp, J., 2009a. Effects of realistic surface topography on seismic ground motion in the Yangminshan region of Taiwan based upon the spectral-element method and LiDAR DTM, *Bull. seism. Soc. Am.*, **99**(2A), 681–693.
- Lee, S.J., Komatitsch, D., Huang, B.S. & Tromp, J., 2009b. Effects of topography on seismic wave propagation: an example from northern Taiwan, *Bull. seism. Soc. Am.*, **99**(1), 314–325.
- Legay, A., Wang, H.W. & Belytschko, T., 2005. Strong and weak arbitrary discontinuities in spectral finite elements, *Int. J. Numer. Meth. Eng.*, **64**(8), 991–1008.
- Liu, H.P., Anderson, D.L. & Kanamori, H., 1976. Velocity dispersion due to anelasticity: implications for seismology and mantle composition, *Geophys. J. R. astr. Soc.*, **47**, 41–58.
- Liu, Q. & Tromp, J., 2006. Finite-frequency kernels based on adjoint methods, *Bull. seism. Soc. Am.*, **96**(6), 2383–2397.
- Liu, Q. & Tromp, J., 2008. Finite-frequency sensitivity kernels for global seismic wave propagation based upon adjoint methods, *Geophys. J. Int.*, **174**(1), 265–286.
- Liu, Q., Polet, J., Komatitsch, D. & Tromp, J., 2004. Spectral-element moment tensor inversions for earthquakes in Southern California, *Bull. seism. Soc. Am.*, **94**(5), 1748–1761.
- Maday, Y. & Patera, A.T., 1989. Spectral-element methods for the incompressible Navier-Stokes equations, in *State of the Art Survey in Computational Mechanics*, pp. 71–143, eds Noor, A.K. & Oden, J.T., Am. Soc. Mech. Eng., New York, NY.
- Maggi, A., Tape, C., Chen, M., Chao, D. & Tromp, J., 2009. An automated time window selection algorithm for seismic tomography, *Geophys. J. Int.*, **178**, 257–281.
- Mallet, J.L., 1992. GOCAD: a computer aided design program for geological applications, in *Three-Dimensional Modeling with Geoscientific Information Systems*, pp. 123–141, Kluwer Academic Publishers, Dordrecht.
- Martin, R. & Couder-Castaneda, C., 2010. An improved unsplit and convolutional perfectly matched layer (CPML) absorbing technique for the Navier-Stokes equations using cut-off frequency shift, *Comput. Model. Eng. Sci.*, **63**(1), 47–77.
- Martin, R. & Komatitsch, D., 2009. An unsplit convolutional perfectly matched layer technique improved at grazing incidence for the viscoelastic wave equation, *Geophys. J. Int.*, **179**(1), 333–344.
- Martin, R., Komatitsch, D., Blitz, C. & Le Goff, N., 2008a. Simulation of seismic wave propagation in an asteroid based upon an unstructured MPI spectral-element method: blocking and non-blocking communication strategies, *Lecture Notes in Computer Science*, **5336**, 350–363.
- Martin, R., Komatitsch, D. & Ezziani, A., 2008b. An unsplit convolutional perfectly matched layer improved at grazing incidence for seismic wave equation in poroelastic media, *Geophysics*, **73**(4), T51–T61.
- Martin, R., Komatitsch, D. & Gedney, S.D., 2008c. A variational formulation of a stabilized unsplit convolutional perfectly matched layer for the isotropic or anisotropic seismic wave equation, *Comput. Model. Eng. Sci.*, **37**(3), 274–304.
- Martin, R., Komatitsch, D., Gedney, S.D. & Bruthiaux, E., 2010. A high-order time and space formulation of the unsplit perfectly matched layer for the seismic wave equation using Auxiliary Differential Equations (ADE-PML), *Comput. Model. Eng. Sci.*, **56**(1), 17–42.
- Mercerat, E.D., Vilotte, J.P. & Sánchez-Sesma, F.J., 2006. Triangular spectral-element simulation of two-dimensional elastic wave propagation using unstructured triangular grids, *Geophys. J. Int.*, **166**(2), 679–698.
- Michéa, D. & Komatitsch, D., 2010. Accelerating a 3D finite-difference wave propagation code using GPU graphics cards, *Geophys. J. Int.*, **182**(1), 389–402.
- Micikevicius, P., 2009. 3D finite-difference computation on GPUs using CUDA, in *GP-GPU-2: Proceedings of the 2nd Workshop on General Purpose Processing on Graphics Processing Units*, pp. 79–84, Washington, DC.
- Mitchell, S., 1996. A characterization of the quadrilateral meshes of a surface which admit a compatible hexahedral mesh of the enclosed volume, in *STACS 96, Lecture Notes in Computer Science*, Vol. 1046, pp. 465–476, eds Puech, C. & Reischuk, R., Springer, Berlin.
- Moczo, P. & Kristek, J., 2005. On the rheological models used for time-domain methods of seismic wave propagation, *Geophys. Res. Lett.*, **32**, L01306.
- Monk, P. & Richter, G.R., 2005. A discontinuous Galerkin method for linear symmetric hyperbolic systems in inhomogeneous media, *J. Sci. Comput.*, **22-23**(1–3), 443–477.
- Noble, C.R. & Nuss, L.K., 2004. Nonlinear seismic analysis of morrow point dam, in *13th World Conference on Earthquake Engineering*, Vancouver, Canada.
- Oliveira, S.P. & Seriani, G., 2011. Effect of element distortion on the numerical dispersion of spectral element methods, *Comm. Comput. Phys.*, **9**(4), 937–958.
- Paolucci, R., Faccioli, E. & Maggio, F., 1999. 3D response analysis of an instrumented hill at Matsuzaki, Japan, by a spectral method, *J. Seismol.*, **3**, 191–209.
- Patera, A.T., 1984. A spectral element method for fluid dynamics: laminar flow in a channel expansion, *J. Comput. Phys.*, **54**, 468–488.
- Pellegrini, F. & Roman, J., 1996. SCOTCH: A software package for static mapping by dual recursive bipartitioning of process and architecture graphs, *Lecture Notes in Computer Science*, **1067**, 493–498.
- Priolo, E., Carcione, J.M. & Seriani, G., 1994. Numerical simulation of interface waves by high-order spectral modeling techniques, *J. acoust. Soc. Am.*, **95**(2), 681–693.
- di Prisco, C., Stupazzini, M. & Zambelli, C., 2007. Nonlinear SEM numerical analyses of dry dense sand specimens under rapid and dynamic loading, *Int. J. Numer. Anal. Meth. Geomech.*, **31**, 757–788.
- de la Puente, J., Ampuero, J.P. & Käser, M., 2009. Dynamic rupture modeling on unstructured meshes using a discontinuous Galerkin method, *J. geophys. Res.*, **114**, B10302, doi:10.1029/2008JB006271.
- Quarteroni, A., Tagliani, A. & Zampieri, E., 1998. Generalized Galerkin approximations of elastic waves with absorbing boundary conditions, *Comput. Meth. Appl. Mech. Eng.*, **163**, 323–341.
- Rainsberger, R., 2006. *TrueGrid User's Manual*, version 2.3.0, XYZ Scientific Applications, Inc., Livermore, CA.
- Reed, W.H. & Hill, T.R., 1973. Triangular mesh methods for the neutron transport equation, Tech. Rep. LA-UR-73-479, Los Alamos Scientific Laboratory, Los Alamos, NM.
- Ribes, A. & Caremoli, C., 2007. Salome platform component model for numerical simulation, in *COMPSAC '07: Proceedings of the 31st Annual International Computer Software and Applications Conference*, vol. 2, pp. 553–564, IEEE Computer Society.
- Ribó, R., de Riera Pasenau, M.A. & Escolano, E., 2011. *GiD Reference Manual, Pre and Post Processing System for F.E.M. Calculations*, International Center For Numerical Methods In Engineering (CIMNE).
- Richardson, J.E., Melosh, H.J., Greenberg, R.J. & O'Brien, D.P., 2005. The global effects of impact-induced seismic activity on fractured asteroid surface morphology, *Icarus*, **179**, 325–349.
- Rivière, B. & Wheeler, M.F., 2003. Discontinuous finite element methods for acoustic and elastic wave problems, *Contemporary Mathematics*, **329**, 271–282.
- Robertsson, J.O.A., 1996. A numerical free-surface condition for elastic/viscoelastic finite-difference modeling in the presence of topography, *Geophysics*, **61**, 1921–1934.

- Robinson, M.S., Thomas, P.C., Veverka, J., Murchie, S.L. & Wilcox, B.B., 2002. The geology of 433 Eros, *Meteorit. planet. Sci.*, **37**, 1651–1684.
- Savage, B., Komatitsch, D. & Tromp, J., 2010. Effects of 3D attenuation on seismic wave amplitude and phase measurements, *Bull. seism. Soc. Am.*, **100**(3), 1241–1251.
- Seriani, G. & Oliveira, S.P., 2008. Dispersion analysis of spectral-element methods for elastic wave propagation, *Wave Motion*, **45**, 729–744.
- Seriani, G. & Priolo, E., 1994. A spectral element method for acoustic wave simulation in heterogeneous media, *Finite Elem. Anal. Des.*, **16**, 337–348.
- Seriani, G., Priolo, E. & Pregarz, A., 1995. Modelling waves in anisotropic media by a spectral element method, in *Proceedings of the Third International Conference on Mathematical and Numerical Aspects of Wave Propagation*, pp. 289–298, SIAM, Philadelphia, PA.
- Shepherd, J.F. & Johnson, C.R., 2008. Hexahedral mesh generation constraints, *Eng. Comput.*, **24**(3), 195–213.
- Sieminski, A., Liu, Q., Trampert, J. & Tromp, J., 2007a. Finite-frequency sensitivity of surface waves to anisotropy based upon adjoint methods, *Geophys. J. Int.*, **168**(3), 1153–1174.
- Sieminski, A., Liu, Q., Trampert, J. & Tromp, J., 2007b. Finite-frequency sensitivity of body waves to anisotropy based upon adjoint methods, *Geophys. J. Int.*, **171**, 368–389.
- SIMULIA, 2008. *Abaqus User’s Manual*, version 6.7 edn, Dassault Systems, Providence, RI.
- Stacey, R., 1988. Improved transparent boundary formulations for the elastic wave equation, *Bull. seism. Soc. Am.*, **78**(6), 2089–2097.
- Staten, M.L., Kerr, R.A., Owen, S.J., Blacker, T.D., Stupazzini, M. & Shimada, K., 2010. Unconstrained plastering—hexahedral mesh generation via advancing-front geometry decomposition, *Int. J. Numer. Methods Eng.*, **81**, 135–171.
- Stupazzini, M., Paolucci, R. & Igel, H., 2009. Near-fault earthquake ground-motion simulation in the Grenoble valley by a high-performance spectral element code, *Bull. seism. Soc. Am.*, **99**(1), 286–301.
- Tago, J., Cruz-Atienza, V.M., Étienne, V., Virieux, J., Benjema, M. & Sánchez-Sesma, F.J., 2010. 3D dynamic rupture with anelastic wave propagation using an hp-adaptive discontinuous Galerkin method, in *Abstract S51A-1915 Presented at 2010 AGU Fall Meeting*, San Francisco, CA, www.agu.org/meetings/fm10/waisfm10.html.
- Tape, C., Liu, Q., Maggi, A. & Tromp, J., 2009. Adjoint tomography of the southern California crust, *Science*, **325**, 988–992.
- Tape, C., Liu, Q., Maggi, A. & Tromp, J., 2010. Seismic tomography of the southern California crust based on spectral-element and adjoint methods, *Geophys. J. Int.*, **180**, 433–462.
- Tarantola, A., 1984. Inversion of seismic reflection data in the acoustic approximation, *Geophysics*, **49**, 1259–1266.
- Tarantola, A., 1987. *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation*, Elsevier Science Publishers, Amsterdam.
- Taylor, M.A. & Wingate, B.A., 2000. A generalized diagonal mass matrix spectral element method for non-quadrilateral elements, *Appl. Num. Math.*, **33**, 259–265.
- Tromp, J., Tape, C. & Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels, *Geophys. J. Int.*, **160**(1), 195–216.
- Tromp, J., Komatitsch, D. & Liu, Q., 2008. Spectral-element and adjoint methods in seismology, *Comm. Comput. Phys.*, **3**(1), 1–32.
- Tromp, J., Luo, Y., Hanasoge, S. & Peter, D., 2010. Noise cross-correlation sensitivity kernels, *Geophys. J. Int.*, **183**, 791–819.
- Tsuij, S., Komatitsch, D., Ji, C. & Tromp, J., 2003. Spectral-element simulations of the November 3, 2002, Denali, Alaska earthquake on the Earth Simulator, *Phys. Earth planet. Int.*, **139**(3–4), 305–313.
- Valentine, A. & Woodhouse, J., 2010. Approaches to automated data selection for global seismic tomography, *Geophys. J. Int.*, **182**, 1001–1012.
- Vos, P.E.J., Sherwin, S.J. & Kirby, R.M., 2010. From h to p efficiently: implementing finite and spectral/hp element methods to achieve optimal performance for low- and high-order discretisations, *J. Comput. Phys.*, **229**, 5161–5181.
- White, D.R., Mingwu, L., Benzley, S.E. & Sjaardema, G.D., 1995. Automated hexahedral mesh generation by virtual decomposition, in *Proceedings of the 4th International Meshing Roundtable*, Albuquerque, NM, pp. 165–176.
- van Wijk, K., Komatitsch, D., Scales, J.A. & Tromp, J., 2004. Analysis of strong scattering at the micro-scale, *J. acoust. Soc. Am.*, **115**(3), 1006–1011.
- Wilcox, L.C., Stadler, G., Burstedde, C. & Ghattas, O., 2010. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media, *J. Comput. Phys.*, **229**(24), 9373–9396.
- Wingate, B.A. & Boyd, J.P., 1996. Spectral element methods on triangles for geophysical fluid dynamics problems, in *Proceedings of the Third International Conference on Spectral and High-Order Methods*, pp. 305–314, Houston J. Mathematics, Houston, TX.
- Xu, J., Bielak, J., Ghattas, O. & Wang, J., 2003. Three-dimensional nonlinear seismic ground motion modeling in inelastic basins, *Phys. Earth planet. Int.*, **137**(1–4), 81–95.

APPENDIX A: ACOUSTIC ADJOINT EQUATIONS

A1 Pressure waveform misfit kernels

For acoustic tomographic studies, it is convenient to define a pressure misfit function

$$\chi = \frac{1}{2} \sum_i \int ||p_i^{\text{syn}}(\mathbf{m}) - p_i^{\text{obs}}||^2 dt, \quad (\text{A1})$$

where p_i^{syn} is the numerically computed pressure and p_i^{obs} the observed pressure at location \mathbf{x}_i . The variation in pressure may be written in terms of the variation in the potential ϕ as

$$\delta p = -\partial_t^2 \delta \phi, \quad (\text{A2})$$

which follows from the definition of the scalar potential ϕ in eq. (11). The corresponding action in the acoustic case is given by

$$\chi = \frac{1}{2} \sum_i \int ||p_i^{\text{syn}} - p_i^{\text{obs}}||^2 dt - \int \int_{\Omega} \lambda [\kappa^{-1} \partial_t^2 \phi - \nabla \cdot (\rho^{-1} \nabla \phi) - f] d^3 \mathbf{x} dt, \quad (\text{A3})$$

where λ denotes a scalar Lagrange multiplier. Setting $\Delta p_i = p_i^{\text{syn}} - p_i^{\text{obs}}$ and taking the variation of the action, we obtain

$$\begin{aligned} \delta \chi = & \sum_i \int \Delta p_i \delta p_i dt - \int \int_{\Omega} [\delta \kappa^{-1} \lambda \partial_t^2 \phi - \nabla \cdot (\delta \rho^{-1} \lambda \nabla \phi) - \lambda \delta f] d^3 \mathbf{x} dt \\ & - \int \int_{\Omega} [\kappa^{-1} \partial_t^2 \lambda - \nabla \cdot (\rho^{-1} \nabla \lambda)] \delta \phi d^3 \mathbf{x} dt - \int \int_{\partial \Omega} \hat{\mathbf{n}} \cdot (\rho^{-1} \nabla \lambda) \delta \phi d^2 \mathbf{x} dt. \end{aligned} \quad (\text{A4})$$

Since eq. (A4) must be stationary when no model perturbations are present, that is, $\delta\rho = 0$, $\delta\kappa = 0$ and $\delta f = 0$, we obtain

$$\begin{aligned}
& \int \int_{\Omega} [\kappa^{-1} \partial_t^2 \lambda - \nabla \cdot (\rho^{-1} \nabla \lambda)] \delta\phi \, d^3 \mathbf{x} \, dt \\
&= \int \int_{\Omega} \sum_i \Delta p_i \delta(\mathbf{x} - \mathbf{x}_i) \delta p \, d^3 \mathbf{x} \, dt - \int \int_{\partial\Omega} \hat{\mathbf{n}} \cdot (\rho^{-1} \nabla \lambda) \delta\phi \, d^2 \mathbf{x} \, dt \\
&= - \int \int_{\Omega} \sum_i \Delta p_i \delta(\mathbf{x} - \mathbf{x}_i) \partial_t^2 \delta\phi \, d^3 \mathbf{x} \, dt - \int \int_{\partial\Omega} \rho^{-1} \hat{\mathbf{n}} \cdot \nabla \lambda \delta\phi \, d^2 \mathbf{x} \, dt \\
&= - \int \int_{\Omega} \sum_i \partial_t^2 \Delta p_i \delta(\mathbf{x} - \mathbf{x}_i) \delta\phi \, d^3 \mathbf{x} \, dt - \int \int_{\partial\Omega} \rho^{-1} \hat{\mathbf{n}} \cdot \nabla \lambda \delta\phi \, d^2 \mathbf{x} \, dt,
\end{aligned} \tag{A5}$$

where \mathbf{x}_i is the station location of the corresponding i th measurement. Note that the last integration by parts of the first term is valid under the assumption that Δp_i and $\partial_t \Delta p_i$ vanish at the limits of the time integration, that is, for a given measurement window $[0, T]$, $\Delta p_i(\mathbf{x}, 0) = \Delta p_i(\mathbf{x}, T) = 0$ and $\partial_t \Delta p_i(\mathbf{x}, 0) = \partial_t \Delta p_i(\mathbf{x}, T) = 0$. This is generally true because we taper the ends of the misfit window.

Let us define the adjoint scalar potential as

$$\phi^\dagger(\mathbf{x}, t) \equiv \lambda(\mathbf{x}, T - t). \tag{A6}$$

It follows from (A5) that ϕ^\dagger must satisfy the adjoint wave equation

$$\kappa^{-1} \partial_t^2 \phi^\dagger - \nabla \cdot (\rho^{-1} \nabla \phi^\dagger) = f^\dagger, \tag{A7}$$

where the adjoint source is given by

$$f^\dagger(\mathbf{x}, t) = - \sum_i \partial_t^2 \Delta p_i(T - t) \delta(\mathbf{x} - \mathbf{x}_i). \tag{A8}$$

The initial conditions for the adjoint potential must satisfy $\phi^\dagger(T) = 0$ and $\partial_t \phi^\dagger(T) = 0$. The corresponding fluid–solid boundary conditions involve terms with $\rho^{-1} \hat{\mathbf{n}} \cdot \nabla \phi^\dagger$.

For acoustic simulations, there is no shear contribution and we may set $K_\mu = 0$. Using

$$\nabla \cdot \mathbf{s} = -\kappa^{-1} p = \kappa^{-1} \partial_t^2 \phi, \tag{A9}$$

$$\nabla \cdot \mathbf{s}^\dagger = \kappa^{-1} \partial_t^2 \phi^\dagger, \tag{A10}$$

the kernel K_κ given in eq. (18) becomes

$$K_\kappa = - \int_0^T \kappa^{-1} \partial_t^2 \phi^\dagger(T - t) \partial_t^2 \phi(t) \, dt. \tag{A11}$$

It is this last kernel expression that is actually implemented, since the values for $\partial_t^2 \phi$ and $\partial_t^2 \phi^\dagger$ are obtained at each time step in the Newark time scheme used to propagate acoustic waves.

A2 Pressure traveltime adjoint sources

Instead of measuring waveform misfits, one may also define a traveltime misfit for pressure signals, that is,

$$\chi = \frac{1}{2} \sum_i \|T_i^{\text{syn}}(\mathbf{m}) - T_i^{\text{obs}}\|^2, \tag{A12}$$

where $T_i^{\text{syn}}(\mathbf{m})$ denotes the arrival time in the synthetic pressure records computed for model \mathbf{m} , and T_i^{obs} the arrival time of the observed pressure wave. The variation in traveltime δT may be written to first order in terms of perturbations in pressure as δp (Hung *et al.* 2000)

$$\delta T = \frac{1}{N} \int \partial_t p \delta p \, dt, \tag{A13}$$

where $N = \int p \partial_t^2 p \, dt$. Using $\delta p = -\partial_t^2 \delta\phi$, this leads to

$$\delta T = -\frac{1}{N} \int \partial_t p \partial_t^2 \delta\phi \, dt. \tag{A14}$$

Defining $\Delta T_i \equiv T_i^{\text{syn}}(\mathbf{m}) - T_i^{\text{obs}}$, the variation of the action becomes

$$\sum_i \Delta T_i \delta T_i = - \sum_i \frac{1}{N} \Delta T_i \int \partial_t p \partial_t^2 \delta\phi \, dt \tag{A15}$$

$$= - \int \int_{\Omega} \sum_i \frac{1}{N} \Delta T_i \delta(\mathbf{x} - \mathbf{x}_i) \partial_t p \partial_t^2 \delta\phi \, d^3 \mathbf{x} \, dt. \tag{A16}$$

Under the assumption that $\partial_t p(0) = \partial_t p(T) = 0$ and $\partial_t^2 p(0) = \partial_t^2 p(T) = 0$ (which can be achieved by carefully selecting and tapering the measurement time windows), we find after some further manipulation that the adjoint source for a travelt ime misfit between observed and simulated pressure signals is given by

$$f^\dagger(\mathbf{x}, t) = - \sum_i \frac{1}{N} \Delta T_i \partial_t^3 p(\mathbf{x}, T - t) \delta(\mathbf{x} - \mathbf{x}_i). \quad (\text{A17})$$