

Anelastic sensitivity kernels with parsimonious storage for adjoint tomography and full waveform inversion

Dimitri Komatitsch,¹ Zhinan Xie,^{1,2} Ebru Bozdağ,³ Elliott Sales de Andrade,⁴ Daniel Peter,⁵ Qinya Liu⁴ and Jeroen Tromp⁶

¹LMA, CNRS UPR 7051, Aix-Marseille University, Centrale Marseille, F-13453 Marseille Cedex 13, France. E-mail: komatitsch@lma.cnrs-mrs.fr

²Institute of Engineering Mechanics, China Earthquake Administration, Harbin 150080, China

³Géoazur, University of Nice Sophia Antipolis, 250 rue Albert Einstein, F-06560 Valbonne, France

⁴Department of Physics and Department of Earth Sciences, University of Toronto, Toronto, Ontario, M5S 1A7, Canada

⁵Extreme Computing Research Center, King Abdullah University of Science and Technology (KAUST), 23955-6900 Thuwal, Saudi Arabia

⁶Department of Geosciences and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA

Accepted 2016 June 9. Received 2016 June 8; in original form 2016 March 22

SUMMARY

We introduce a technique to compute exact anelastic sensitivity kernels in the time domain using parsimonious disk storage. The method is based on a reordering of the time loop of time-domain forward/adjoint wave propagation solvers combined with the use of a memory buffer. It avoids instabilities that occur when time-reversing dissipative wave propagation simulations. The total number of required time steps is unchanged compared to usual acoustic or elastic approaches. The cost is reduced by a factor of 4/3 compared to the case in which anelasticity is partially accounted for by accommodating the effects of physical dispersion. We validate our technique by performing a test in which we compare the K_α sensitivity kernel to the exact kernel obtained by saving the entire forward calculation. This benchmark confirms that our approach is also exact. We illustrate the importance of including full attenuation in the calculation of sensitivity kernels by showing significant differences with physical-dispersion-only kernels.

Key words: Numerical solutions; Tomography; Seismic attenuation; Seismic tomography; Computational seismology; Wave propagation.

1 INTRODUCTION

Efficient numerical methods for simulating the propagation of acoustic, elastic, or anelastic waves in the time domain are widely available, for instance, based on finite-difference methods (see e.g. Virieux & Operto 2009, for a review), spectral-element methods (e.g. Komatitsch & Vilotte 1998; Vai *et al.* 1999; Komatitsch & Tromp 1999, 2002), or standard finite-element methods (e.g. Kallivokas *et al.* 2013). Nowadays, these techniques are heavily used for imaging based on full waveform inversion (FWI) or adjoint tomography (e.g. Tromp *et al.* 2005; Plessix 2006; Tromp *et al.* 2008; Virieux & Operto 2009; Fichtner 2010; Monteiller *et al.* 2015). FWI involves fitting bandpass filtered versions of observed seismograms by minimizing least-squared differences between observed and synthetic seismograms. Adjoint tomography generalizes FWI by considering arbitrary measures of misfit, for example, cross-correlation traveltimes, multitaper phase and amplitude anomalies, or instantaneous phase measurements.

In the context of imaging, it is useful to resort to the concept of sensitivity kernels (e.g. Tarantola 1986, 1987, 1988; Tromp *et al.* 2005, 2008; Liu & Tromp 2008; Fichtner 2010; Fichtner & van Driel 2014). Let $\mathbf{s}$ denote the forward displacement wavefield and $\mathbf{s}^\dagger$ the

adjoint wavefield. In an isotropic Earth model, the kernels K_κ and K_μ represent Fréchet derivatives with respect to relative bulk and shear moduli perturbations, respectively. These kernels are given by (e.g. Tromp *et al.* 2008)

$$K_\kappa(\mathbf{x}) = - \int_0^T \kappa(\mathbf{x}) [\nabla \cdot \mathbf{s}^\dagger(\mathbf{x}, T-t)] [\nabla \cdot \mathbf{s}(\mathbf{x}, t)] dt, \quad (1)$$

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

where

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

and

$$\mathbf{D}^\dagger = \frac{1}{2} [\nabla \mathbf{s}^\dagger + (\nabla \mathbf{s}^\dagger)^T] - \frac{1}{3} (\nabla \cdot \mathbf{s}^\dagger) \mathbf{I} \quad (4)$$

denote the traceless strain deviator and its adjoint, respectively, $\mathbf{x}$ is the position vector and κ and μ are the bulk and shear moduli, respectively. Their expression remains valid for elastic perturbations superimposed on an anelastic Earth model (Liu & Tromp 2008) if the regular and adjoint wavefields are computed in that anelastic

reference model. In practical applications, it is often useful to define compressional and shear wave speed sensitivity kernels, namely (Tromp *et al.* 2005, 2008)

$$K_\alpha = 2 \left(\frac{\kappa + \frac{4}{3}\mu}{\kappa} \right) K_\kappa \quad (5)$$

and

$$K_\beta = 2 \left(K_\mu - \frac{4}{3} \frac{\mu}{\kappa} K_\kappa \right). \quad (6)$$

In order to perform the convolution involved in the calculation of the kernels (1) and (2), simultaneous access to the forward wavefield $\mathbf{s}$ at time t and the adjoint wavefield $\mathbf{s}^\dagger$ at time $T - t$ is required (or conversely, since $\int_0^T f(t) f^\dagger(T - t) dt = \int_0^T f(T - t) f^\dagger(t) dt$ when convolving two functions f and $f^\dagger$). Carrying out forward and adjoint simulations simultaneously is insufficient, because in that case both wavefields are only available at a given time t . In addition, to calculate the adjoint wavefield one must prescribe the adjoint source, and that source is computed based on measurements between observed and simulated seismograms, that is, it can only be constructed after the completion of a forward simulation.

A straightforward solution to this dilemma is to store the entire forward simulation to disk and then read it back in reverse order during the adjoint simulation. For 1-D or 2-D models this is feasible (e.g. Pakravan *et al.* 2016), but in 3-D at short periods without lossy compression or significant spatial or temporal subsampling (Fichtner *et al.* 2009; Sun & Fu 2013; Rubio Dalmau *et al.* 2014; Cyr *et al.* 2015) the required amount of disk storage is currently prohibitive. It is worth mentioning that this situation will likely change in the future, but not any time soon. In addition, heavy I/O involved in reading back the forward wavefield can significantly slow down the simulation (Yuan *et al.* 2014).

For non-dissipative acoustic or elastic media, a standard solution for large 3-D problems is to perform three simulations per source (Tromp *et al.* 2008; Peter *et al.* 2011). One performs the forward calculation twice: once to compute the adjoint sources and once again in reverse time simultaneously with the adjoint simulation performed in forward time to correlate the two fields and sum their interaction on the fly over all time steps. Thus, one only needs a small amount of disk storage to store the last time step of the forward run, which is then used as an initial condition to recalculate the forward wavefield backwards in time with a negative time step. Based on this strategy, one has simultaneous access to the adjoint wavefield at time t and the forward wavefield at time $T - t$, which is what is required to perform the convolution involved in the construction of the kernels (1) and (2). In the anelastic case, as shown by Tarantola (1988), the wave equation is no longer self adjoint, leading to exponential growth of energy in the adjoint equation (i.e. ‘anti attenuation’). However, in the calculation of sensitivity kernels this does not matter, because this process involves the *time-reversed* adjoint wave equation rather than the adjoint wave equation. In this respect, it is worth mentioning that Liu & Tromp (2008) have a time-reversed definition of the adjoint state compared to the classical one of, for example, Tarantola (1988) and Virieux & Operto (2009), that is, they use the time-reversed eq. (35) of Plessix (2006) instead of his eq. (32). Thus, ‘anti attenuation’ in forward time becomes attenuation in reverse time, and the time-reversed adjoint wave equation is identical to the classical forward wave equation (Tarantola 1988; Fichtner *et al.* 2006; Liu & Tromp 2008). Consequently, both the forward and the adjoint wavefields are attenuated

during the calculation of the kernel. This makes physical sense, because as the distance between two seismic stations increases, one expects the sensitivity kernel to become weaker and weaker as a result of dissipation. Unfortunately, in the presence of attenuation the process of reconstructing the forward wavefield backwards in time based on the final snapshot and a negative time step is numerically highly unstable (e.g. Liu & Tromp 2006, 2008; Kowar & Scherzer 2011; Ammari *et al.* 2013). This results in numerical instabilities during the calculation of anelastic sensitivity kernels, not because of the adjoint run, but because of the forward run that needs to be performed backwards in time.

An approximate solution to this conundrum involves modifying the wave equation to introduce filtering or other stabilizing terms, or incorporating only certain aspects of anelasticity. Ammari *et al.* (2013) introduced a promising regularized time-reversal imaging technique which corrects attenuation effects to first order. However, their approach involves significant filtering, thereby affecting the quality of the resulting signals, in particular at high frequencies, and it is currently limited to simple models of weak attenuation in homogeneous media. In particular, the approach cannot handle models comprised of standard linear solids. Zhu (2014) and Zhu *et al.* (2014) introduce partial support for attenuation in reverse-time migration by separating amplitude attenuation and phase dispersion operators (Varela *et al.* 1993). They construct attenuation- and dispersion-compensated operators by reversing the sign of the attenuation operator and leaving the sign of the dispersion operator unchanged; they then design a low-pass filter for these operators to stabilize the numerical procedure and avoid amplifying high-frequency noise that would trigger instabilities. Changing the sign of one of the terms in the backward integrations to stabilize the calculations is also classically done in so-called Back-and-Forth Nudging algorithms (Auroux *et al.* 2011, 2013).

An alternative solution consists of resorting to so-called ‘partial checkpointing’ or ‘optimal checkpointing’, that is, using partial storage to disk, but, realizing that storage size limitations or slow-down related to disk storage are important issues (Yuan *et al.* 2014), defining an optimized sequence in which the forward and adjoint time steps are performed, essentially trading storage requirements for longer computation times. This elegant idea was introduced by Restrepo *et al.* (1998) and Griewank & Walther (2000). Restrepo *et al.* (1998) used a recursive strategy to compute the optimized order in which the simulation steps need to be performed and showed that when the schedule is optimized the storage and computational times grow at most logarithmically, and Griewank & Walther (2000) defined an algorithm called ‘Revolve’ that is provably optimal to reduce storage requirements. Charpentier (2001) used it for the meteorological model ‘Meso-NH’ and proposed several variants depending on user preference between CPU time and memory optimization. Akçelik *et al.* (2003), Symes (2007) and Anderson *et al.* (2012) resorted to it for FWI as well as reverse-time migration. Hinze *et al.* (2006) applied it to the instationary Navier–Stokes equations, and more recently Spears *et al.* (2014) included it in a jet-engine noise reduction simulation code. A limitation of the original ‘Revolve’ algorithm of Griewank & Walther (2000) is that it requires *a priori* knowledge of the total number of time steps to be performed, making it incompatible with adaptive time stepping. Wang *et al.* (2009) have addressed that issue by introducing a dynamic checkpointing algorithm that is applicable even when the total number of time steps is *a priori* unknown.

Another—more precise, but also more expensive—approximate solution is to accommodate the effects of physical




Figure 1. Taking into account only the effects of physical dispersion in anelastic simulations, that is, its effect on phase but not on amplitudes (no dissipation), amounts to performing an elastic simulation in a wave speed model shifted to the dominant frequency of the source (red cross) relative to the frequency of the reference model (green cross), based on eq. (8). The logarithmic phase speed in a strictly constant- Q absorption-band model is represented by the dotted line. In practice, in time-domain simulations a constant- Q model is approximated by a small number of standard linear solids (usually Zener solids) in parallel (e.g. Carcione 2014; Blanc *et al.* 2016), which approximate a constant Q inside an absorption band of interest (solid line).

dispersion induced by attenuation, but not the effects of dissipation. To appreciate this, consider the frequency-dependent shear modulus in a constant- Q absorption-band solid, namely (e.g. Liu *et al.* 1976; Dahlen & Tromp 1998; Carcione 2014)

$$\frac{\mu(\omega)}{\mu(\omega_0)} = 1 + \frac{2}{\pi Q_\mu} \ln\left(\frac{\omega}{\omega_0}\right) + \frac{i}{Q_\mu}. \quad (7)$$

Here, ω denotes the angular frequency of interest, ω_0 a chosen reference frequency and Q_μ a frequency-independent shear quality factor. The effects of physical dispersion are captured by the logarithmic term, $2/(\pi Q_\mu) \ln(\omega/\omega_0)$, and the effects of dissipation by the complex part of the modulus, $1/Q_\mu$. Based on these observations, partial support for attenuation is accommodated by performing a total of four simulations per source (e.g. Zhou *et al.* 2011; Zhu *et al.* 2012), instead of three in the non-dissipative case (Tromp *et al.* 2008; Peter *et al.* 2011). In a first stage, one runs a forward simulation twice in the forward direction, once with full attenuation and once with physical dispersion only. The second calculation is therefore a purely elastic calculation, but for a wave speed model that is shifted to the dominant frequency of the source, ω_s , based on the correction

$$\mu(\omega_s) = \mu(\omega_0)[1 + 2/(\pi Q_\mu) \ln(\omega_s/\omega_0)], \quad (8)$$

as illustrated in Fig. 1. The first forward run with full attenuation is used to make the measurements needed in the construction of the adjoint sources for the third run, and the second forward run with physical dispersion only is used to compute and store the final time step to be able to time reverse that calculation in the fourth run. The third and fourth runs are carried out simultaneously to calculate the kernel, and both these runs are purely elastic and use a wave speed model that is shifted to the dominant frequency of the source. The fourth run is thus stable, because it involves time reversal of an elastic simulation. Note, however, that the measurements that are assimilated in the third run are based on synthetics computed with full attenuation.

It is important to recognize that the significance of anelastic effects on sensitivity kernels very much depends on the type of measurement one chooses to make. For example, as demonstrated by Tromp *et al.* (2005), the sensitivity kernel for a cross-correlation traveltime measurement is identical to the so-called ‘banana-doughnut’ kernel first introduced by Dahlen *et al.* (2000). Such kernels may be calculated based on ray theory, and the related expressions are largely unaffected by attenuation. Physically, this reflects the fact that traveltimes are affected by wave speed, and only very marginally by dissipation. More generally, as long as one focuses on measuring phase, for example, frequency-dependent traveltime or instantaneous phase, the corresponding kernels are largely unaffected by attenuation (Zhou *et al.* 2011). However, for inversions involving amplitude measurements, including FWI, attenuation plays a critical role. More generally, attenuation is important on a global scale (e.g. Ruan & Zhou 2010, 2012), in exploration geophysics (e.g. Kurzmann *et al.* 2013; Groos *et al.* 2014) and in near-surface geophysics or for site effects in poorly consolidated sediments (e.g. Askan *et al.* 2007; Assimaki *et al.* 2012). We demonstrate in this paper that global simulations of surface waves at periods less than 40–60 s must accommodate the full effects of attenuation.

To some extent, the accuracy of the gradient is not that critical in the early stages of an inversion, because as part of the iterative inversions scheme, for example, a non-linear conjugate gradient method or an L-BFGS quasi-Newton method, the raw gradient is generally smoothed and pre-conditioned (e.g. Sherali & Ulular 1990; Felgenhauer 1992; Shi 2006). Nevertheless, as the inversion proceeds and the frequency content of the seismograms is increased, details in the gradient matter, and its accurate calculation becomes highly relevant.

2 PARSIMONIOUS STORAGE TECHNIQUE

In view of eqs (1) and (2), for simulations in the time domain consisting of N time steps numbered from 1 to N , the contribution to adjoint-based sensitivity kernels at time step i is obtained by combining information coming from time step i of the adjoint run simultaneously with information coming from time step $N - i + 1$ of the forward run. Two classical approaches can be used to facilitate this, namely, at low to moderate frequencies and/or small to moderate model sizes one can consider storing the entire forward run to disk (Process A; see Fig. 2a), and reading it back from disk in reverse order during the adjoint run. Lossless or lossy compression (Fichtner *et al.* 2009; Sun & Fu 2013; Rubio Dalmau *et al.* 2014; Cyr *et al.* 2015) or spatial or temporal subsampling (Sun & Fu 2013) can be helpful in this context. However, the required amount of storage is currently unaffordable and will remain so for many years to come, and, in addition, heavy I/O will significantly slow down the simulation code (Yuan *et al.* 2014).

Another classical approach (e.g. Liu & Tromp 2006, 2008; Tromp *et al.* 2008) is to first perform the forward run and store its final time step to disk, and in a second stage perform the adjoint run in the forward direction while simultaneously redoing the forward run backwards, reversing time and starting from the final time step (Process B; see Fig. 2b). In the acoustic or elastic case, that is, when total energy is conserved, this process is numerically stable and its only two drawbacks are that the compute time increases by a factor of 3/2 because the forward run needs to be performed twice, and that




Figure 2. For time-domain simulations consisting of N time steps numbered from 1 to N , the contribution to adjoint-based sensitivity kernels at time step i is obtained by combining information coming from time step i of the adjoint run with information coming from time step $N - i + 1$ of the forward run. For low to moderate frequencies and/or small to moderate model sizes, (a, red dashed slices) one can consider storing the entire forward run to disk, reading it back from disk in reverse order while computing the adjoint wavefield. However, for high frequencies and/or large model sizes, the amount of storage needed is currently unaffordable. (b) Another classical approach is to perform the forward run first and store its final time step to disk, and in a second stage perform the adjoint run while simultaneously redoing the forward run backwards, reversing time and starting from the stored final time step. In the acoustic or elastic cases, that process is stable, but not in the anelastic case. However, on computers that have a significant amount of memory per compute node, (c) a third process can be designed, which is stable even in the presence of energy loss. During the first stage, one saves checkpointing/restart files to disk every few hundred or thousand time steps. During the second stage, one still performs two runs simultaneously, but instead of performing the forward run backward from the stored final time step one performs it in chunks in reverse order but in the *forward* direction, in each case starting from the previous restart file read back from disk and storing that subpart of the run in memory (the green region, stored in a memory buffer). Since the run is conducted forward rather than backward in time, this process is always stable, even in the presence of attenuation.

the required memory size increases by a factor of two because during the second stage two runs need to be performed simultaneously in memory. Let us note that the first drawback is not that serious compared to Process A (Fig. 2a), because heavy I/O slows the latter down considerably. Unfortunately, as mentioned above, Process B cannot be used—without heavy filtering and resulting significant loss of accuracy (e.g. Ammari *et al.* 2013)—in the anelastic case or in the presence of any kind of energy loss, because time-reversing energy decay is unstable from a numerical point of view (e.g. Liu & Tromp 2006, 2008; Kowar & Scherzer 2011; Ammari *et al.* 2013,

and references therein). The reason is that while amplifying the fields to restore energy when going backwards, numerical schemes will also amplify numerical noise and thus very quickly become unstable.

Even if full attenuation cannot be taken into account in Process B, it can be partially accommodated by performing two runs instead of one during Stage 1, as we will see in more details in Section 4: one with full attenuation to make measurements to be used in the calculation of the adjoint sources for Stage 2, and another one with physical dispersion only in order to compute and store the final

time step and be able to time reverse that calculation in Stage 2. In the literature, it is often mentioned (e.g. Liu & Tromp 2006, 2008) that during Stage 2 one needs to reverse time and perform the forward run backwards, but, more precisely, the only requirement is to have access to time step $N - i + 1$ of the forward run, regardless of whether it is computed backward or forward in time. Thus, on computers that have a significant amount of memory—which is always the case on modern compute clusters—a third process can be designed, which is stable even in the presence of energy loss (Process C; see Fig. 2c). In this approach, during the first stage one saves a small number of evenly spaced checkpointing/restart files of the three components of the displacement field to disk, typically one every few hundred or thousand time steps; during the second stage one still performs two simulations simultaneously, one adjoint run and one forward run, but instead of performing the forward run *backward* from the stored final time step one performs it in chunks, in reverse order, but in the *forward* direction inside each chunk. In each instance, one starts from the previous restart file of displacement read back from disk, storing only that subpart of the run in memory. Since the run is conducted forward rather than backward in time, this process is always stable, even in the presence of attenuation. It is also exact, since no filtering is involved. Process C is computationally meaningful compared to Process A only if the number of time steps between two checkpoints is sufficiently large, say a few hundred to a few thousand, that is, if the total memory available as a storage buffer is large enough. Fortunately, in practice that is almost always the case on modern compute clusters.

Compared to the more involved ‘Revolve’ algorithm of Griewank & Walther (2000) discussed in the introduction, which is provably optimal in terms of minimizing the number of time steps to store in memory (see also, e.g. Hinze *et al.* 2006), our choice is different and rather makes optimal use of the entire computer memory, that is, it maximizes memory usage instead of trying to minimize it. The rationale for this is that monitoring of typical large wave propagation simulations, for instance, in seismology or in the oil industry, shows that—considering the large compute clusters or supercomputers which are nowadays readily available—users use only a small portion of the memory available per compute node, typically between 5 and 30 per cent, because they generally harness a relatively large number of processor cores to keep the calculation relatively fast. Thus, leaving 5 per cent for the operating system of the machine, between 65 and 90 per cent of the total memory is available and can be used as a memory buffer. This enables one to store at least hundreds of time steps of the displacement vector, sometimes even a few thousand. The number of time steps that can be stored is readily computed in an exact fashion once and for all before the time loop by dividing the size of the available free computer memory by the (constant) size of the array that contains the three components of the displacement vector at a given time step. Note that sensitivity kernel calculations often require the strain (e.g. Tromp *et al.* 2008; Liu & Tromp 2008, as well as eqs 1 and 2), but to reduce storage to disk, which is both disk-space consuming and slow, we usually recompute the strain from the stored displacement instead of storing it.

In terms of implementation, adding this approach to an existing code is easy because it consists mainly of restructuring the time loop and implementing a simple memory buffer system. Note that the cost does not increase compared to Process B, because we perform the same total number of operations, simply in a different order. In fact, if attenuation is partially taken into account in Process B, four runs are needed instead of three, as mentioned above,

and in such a case Process C is cheaper by a factor 4/3. The writing of checkpointing files to disk during Stage 1 of the algorithm may be non-blocking (if technically feasible on the file system), thereby allowing for overlapping of disk writes with calculations, because these restart files are reused much later in the algorithm, during Stage 2. The reason why a memory buffer is needed during Stage 2 is that in order to gain access to time step $N - i + 1$ of the forward run, this buffer will be filled in forward order from the previous restart file, but will then be accessed backward, that is, in reverse order from its end, a policy often called ‘Last In, First Out’ (LIFO).

Interestingly, even in the case of purely acoustic or elastic sensitivity kernels, that is, in the absence of attenuation, Process C is a little more accurate than Process B in terms of numerical errors, because in C one computes the exact same forward run twice, the second time from intermediate restart files, thus resetting numerical errors and getting cumulated numerical dispersion for a total of N time steps only, while in B one performs the forward simulation and then a second forward simulation backwards from the saved snapshot of the final time step, thus getting cumulated numerical dispersion for a total of $2N$ time steps.

3 VALIDATION BENCHMARK

In this section, we validate our approach to compute anelastic sensitivity kernels (Process C shown in Fig. 2c) by comparing its results to those obtained with the exact approach (Process A shown in Fig. 2a). In order to illustrate the effects of full attenuation on sensitivity kernels, we also compare our kernels to an approximate kernel in which only physical dispersion is taken into account, that is, Process B shown in Fig. 2(b), but with a total of four runs performed instead of three, as discussed in Section 1. We calculate kernels for a cross-correlation traveltime measurement, that is, we use time-reversed particle velocity as the adjoint source, as explained in Tromp *et al.* (2005).

In order to perform the benchmark, we resort to the spectral-element method (e.g. Komatitsch & Tromp 1999, 2002). The mesh of hexahedra used in the 3-D simulations is designed to honor all first-order discontinuities in the Preliminary Reference Earth Model (PREM; Dziewoński & Anderson 1981), which are the Moho at a depth of 24.4 km, the upper-mantle discontinuities at depths of 220, 400 and 670 km, the core–mantle boundary and the inner-core boundary; it also honours second-order discontinuities at 600, 771 km, and at the top of the D'' layer. The mesh is doubled in size once below the Moho, a second time below the 670 km discontinuity and a third time in the middle of the outer core (Komatitsch & Tromp 2002). Each of the six chunks that comprise the so-called ‘cubed sphere’ that the spectral-element technique uses to mesh the Earth has 256×256 elements along the free surface and, as a result of the three doublings, 32×32 elements along the inner-core boundary, leading to a total of 4 352 000 spectral elements to mesh the entire globe. The radial density and velocity profiles of the model are determined by PREM. The 3-km-thick water layer of PREM has been replaced with the PREM upper crust. PREM has a transversely isotropic asthenosphere between 24.4 and 220 km, which is also incorporated in our simulations. Based on the size of the mesh cells, the simulations presented in this section are accurate for periods greater than about 17 s. To ensure stability and accuracy of the calculations, we use a time step $\Delta t = 0.19$ s. We simulate a total duration of 5400 s, that is, 28 600 time steps.




Figure 3. (a) Reference exact anelastic K_α sensitivity kernel defined by eq. (5) in the 100–200 s period range at an epicentral distance of 120° obtained by saving the entire forward simulation to disk. (Process A in Fig. 2a.) The source is indicated by the red star and the receiver by the green triangle. (b) Approximate anelastic K_α sensitivity kernel obtained by taking into account physical dispersion only, as explained in Section 1. (Process A in Fig. 2a.) This kernel matches the exact result in terms of its pattern, but differs on average by 30 per cent in magnitude. (c) Anelastic sensitivity kernel computed with our new method. (Process C in Fig. 2c.) This kernel matches the exact result in both pattern and magnitude, with differences of less than 0.01 per cent. (d) The difference between Processes B and A. (e) The difference between Processes C and A, enhanced by a factor of 10^4 . (f)–(h) Vertical-component synthetic seismograms filtered in the 100–200 s passband. The red rectangle indicates the surface wave signal used to create the adjoint source. The seismograms obtained by the exact and new methods are identical since the forward simulation is the same in both cases. The seismogram obtained using physical dispersion only shows the surface wave arriving 30 s early on average, with double the peak amplitude.

We resort to parallel computing using a total of 384 processor cores.

A source with a source time function with a half duration of 11.2 s, strike of 174° , dip of 30° and rake of 67° is located at latitude -16.08° and longitude 168.31° , at a depth of 15 km (corresponding to event 112699G in the global CMT catalogue). A receiver is located on the surface of the Earth at latitude 25.10° and longitude 52.37° , at an epicentral distance of 120° , and records the three components of the displacement vector.

Fig. 3 shows the K_α sensitivity kernel defined by eq. (5) in the 100–200 s period range obtained based on Processes A, B and C. The difference between Processes A and C is shown in Fig. 3(e) and confirms that our technique works well and is exact. When using physical dispersion in Process B, the difference in the resulting kernel is on average 30 per cent of the exact result, as shown in Fig. 3(d). The non-negligible differences that appear when only physical dispersion is taken into account rather than full attenuation highlight the importance of including full attenuation in the calculation of sensitivity kernels, as discussed further in the next section. Note that the convention for traveltimes kernels throughout

this paper is such that $\Delta T = T^{\text{obs}} - T^{\text{syn}}$, where T^{obs} and T^{syn} are the traveltimes of observed and synthetic data (e.g. Marquering *et al.* 1999; Dahlen & Baig 2002).

We have chosen to represent the K_α kernel because, as shown in eq. (5), it requires saving a single scalar to disk during Stage 1 of Process A, namely the trace of the strain tensor, $\nabla \cdot \mathbf{s}$, as a function of time and space. Saving that scalar required 16 TB of disk space, thus illustrating that Process A is currently inconvenient and cannot be routinely used when conducting seismic imaging at relatively high frequencies. In comparison, computing the exact K_β kernel given by eq. (6) via Process A would require 50 TB of additional disk storage.

4 IMPORTANCE OF FULL ATTENUATION IN KERNEL CALCULATIONS

In this section, we compare sensitivity kernels calculated based on physical dispersion only (Process B) with exact kernels calculated




Figure 4. Multitaper traveltime shear wave speed sensitivity kernels K_{β_v} for 40–60 s vertical-component R1 and R2 waves at depths of (A) 30 km and (B) 125 km. The minor-arc epicentral distance is 60° . Traveltime measurements are set to -1 in the computations (i.e. $\Delta T(\omega) = -1$). The locations of the source and receiver are indicated by the red star and green triangle, respectively. The white star and triangle denote the source and receiver antipodes, respectively. S40RTS with Crust2.0 is used as the 3-D model to compute forward and adjoint simulations. Associated seismograms and adjoint sources are as shown in Fig. 5.

based on our new parsimonious storage technique (Process C). As discussed in detail by Zhou *et al.* (2011), for body-wave traveltime measurements attenuation can be safely ignored, as long as the kernels are calculated in models with the appropriate wave speed, that is, taking into account the effects of physical dispersion defined in eq. (8) and illustrated in Fig. 1. Zhou *et al.* (2011) also show that for intermediate-period surface waves at regional distances the physical-dispersion-only approach is perfectly valid. Thus, the use of Process B to accommodate the effects of attenuation in regional-scale studies (e.g. Zhu *et al.* 2012; Zhu & Tromp 2013; Chen *et al.* 2015) is well justified.

Process B may be summarized as follows:

- (i) Compute two sets of synthetic seismograms, (1) using full attenuation, and (2) using physical dispersion only.
- (ii) Make measurements between observed and synthetic seismograms with full attenuation.
- (iii) Calculate sensitivity kernels for this measurement using physical-dispersion-only forward and adjoint wavefields.
- (iv) Compute the gradient by weighting the kernels obtained in step (iii) with the measurements from step (ii).

Note that Process B is suitable for cross-correlation or frequency-dependent (e.g. multitaper) traveltime and amplitude anomaly measurements, but not for FWI. In contrast, the new approach based on Process C, taking into account full attenuation, may be summarized as follows:

- (i) Compute synthetic seismograms using full attenuation.
- (ii) Make measurements between observed and synthetic seismograms with full attenuation.
- (iii) Calculate sensitivity kernels for this measurement using forward and adjoint wavefields with full attenuation based on Process C.
- (iv) Compute the gradient by weighting the kernels obtained in step (iii) with the measurements from step (ii).

The number of simulations required for Process C is reduced by a factor 4/3 compared to Process B, although the extra cost is partially offset by relatively cheaper physical dispersion-only simulations in Process B.

In Fig. 4, we present horizontal cross-sections of multitaper traveltime shear wave speed sensitivity kernels defined by eq. (6) at depths of 30 and 125 km for 40–60 s vertical-component Rayleigh waves using physical-dispersion-only and full attenuation, respectively. We used 3-D mantle model S40RTS (Ritsema *et al.* 2011) together with 3-D crustal model Crust2.0 (Bassin *et al.* 2000) as a background model during forward and adjoint simulations. Confirming observations by Zhou *et al.* (2011), the two sets of kernels are in good agreement. The corresponding R1 and R2 seismograms together with their adjoint sources computed based on cross-correlation and multitaper measurements are shown in Fig. 5. The difference between physical-dispersion-only and full-attenuation kernels is mainly in amplitude, although the former exhibit slight differences in shape compared to the latter. The success of physical-dispersion-only kernels strongly depends on the choice of measurement and the bandpass. As long as the physical-dispersion-only and full attenuation waveforms are similar in shape, the resulting kernels will also be similar, as is clearly shown for 40 s Rayleigh waves.

In Fig. 6, we present horizontal cross-sections of Love-wave multitaper traveltime shear wave speed sensitivity kernels defined by eq. (6) at 30 and 125 km depths. Shown are 40–60 s transverse-component Love-wave kernels using physical-dispersion-only and full attenuation, respectively. Again, S40RTS (Ritsema *et al.* 2011) together with Crust2.0 (Bassin *et al.* 2000) is used during the numerical simulations. The corresponding G1 and G2 seismograms together with their adjoint sources computed based on cross-correlation and multitaper measurements are shown in Fig. 7. For these shorter period Love waves, we see that the physical dispersion-only sensitivity kernels are beginning to breakdown, especially along the major arc, mainly due to their stronger sensitivity to the 3-D crustal heterogeneity.




Figure 5. (A) Vertical-component R1 seismograms computed with physical-dispersion-only and full attenuation (top row) and their associated adjoint sources (bottom row). (B) Vertical-component R2 seismograms computed with physical-dispersion-only and full attenuation (top row) and their associated adjoint sources, where the measurements (i.e. ΔT) are set to -1 (bottom row). CC and MT denote cross-correlation traveltime and multitaper measurements, respectively. The multitaper adjoint sources are used to compute the K_{β_v} kernels presented in Fig. 4. The epicentral distance is 60° , and seismograms were filtered between 40 and 60 s. S40RTS with Crust2.0 is used as the 3-D model to compute the seismograms. Note that, due to the relatively narrow-band signals, the physical-dispersion-only and full-attenuation waveforms are similar. Note also that, due to the non-dispersive behaviour of the wave trains, the CC and MT adjoint sources are very similar.




Figure 6. Multitaper traveltime shear wave speed sensitivity kernels $K_{\beta h}$ for 40–60 s transverse-component G1 and G2 waves at depths of (A) 30 km and (B) 125 km. The minor-arc epicentral distance is 60° . Traveltime measurements are set to -1 in the computations (i.e. $\Delta T(\omega) = -1$). The locations of the source and receiver are shown by the red star and the green triangle, respectively. The white star and triangle denote the source and receiver antipodes, respectively. S40RTS with Crust2.0 is used as the 3-D model to compute forward and adjoint simulations. Associated seismograms and adjoint sources are as shown in Fig. 7.

5 CONCLUSIONS AND FUTURE WORK

We have introduced a method of computing exact anelastic sensitivity kernels in the time domain using parsimonious disk storage and a simple reordering of the time loop, combined with the use of a ‘LIFO’ memory buffer. The total number of time steps required is unaffected compared to usual approaches for the acoustic or elastic (non-dissipative) cases. We reduced the computational cost by a factor 4/3 compared to a commonly used approach in which only the effects of physical dispersion associated with anelasticity are taken into account.

We performed a benchmark in which we compared the compressional wave speed sensitivity kernel obtained based on our approach to the exact kernel obtained by saving the entire forward calculation to disk; the difference was zero, confirming that our approach is also exact. For shorter period surface waves, we discovered non-negligible kernel differences, thus illustrating the importance of including full attenuation in sensitivity kernel calculations for dispersive waves.

The technique applies without modification to problems in reverse-time migration, which may be viewed as a particular case of a sensitivity kernel calculation (e.g. Virieux & Operto 2009; Douma *et al.* 2010), time-reversal seismological source studies (e.g. Larmat *et al.* 2006) and time reversal as used in medical imaging or non-destructive testing (e.g. Fink *et al.* 2000; Tanter & Fink 2014). It would work for Maxwell’s equations as well, since they can be written as a hyperbolic system and are also self-adjoint in the absence of dissipation. The technique works particularly well on GPU-accelerated machines (e.g. Komatitsch *et al.* 2010; Komatitsch 2011), because the entire memory of the CPU is largely unused and thus available as a huge memory buffer.

Our SPECFEM open source spectral-element software package is freely available via the Computational Infrastructure for

Geodynamics (geodynamics.org), including the new developments presented in this paper.

ACKNOWLEDGEMENTS

We thank Mark Asch, Didier Auroux, Cédric Bellis, Élie Bretin, Andreas Fichtner, Josselin Garnier, Thomas Guillet, Ioannis G. Kevrekidis, Bruno Lombard, Vadim Monteiller and William W. Symes for fruitful discussion, and the Computational Infrastructure for Geodynamics (CIG) and Marie Cournille for support. We thank Heiner Igel and an anonymous reviewer for useful comments that improved the manuscript. Part of this work was funded by the Simone and Cino del Duca/Institut de France/French Academy of Sciences Foundation under grant no. 095164, by the European Union Horizon 2020 Marie Curie Action no. 641943 project ‘WAVES’ of call H2020-MSCA-ITN-2014, by U.S. NSF grant 1112906 and by China NSFC grant 51378479. ZX thanks the China Scholarship Council for financial support during his stay at LMA CNRS, and the continuous support from Prof Liao Zhenpeng. ES and QL were supported by the NSERC G8 Research Councils Initiative on Multilateral Research grant no. 490919 and Discovery grant no. 487237. This work was granted access to the European Partnership for Advanced Computing in Europe (PRACE) under allocation TGCC CURIE no. ra2410, to the French HPC resources of TGCC under allocation no. 2015-gen7165 made by GENCI and of the Aix-Marseille Supercomputing Mesocenter under allocations nos 14b013 and 15b034, to the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory, USA, which is supported by the Office of Science of the U.S. Department of Energy under contract no. DE-AC05-00OR22725, and to the Sandybridge cluster at the SciNet HPC Consortium funded by the Canada Foundation for Innovation, the Ontario Research Fund and the University of Toronto Startup Fund. Part of this work was




Figure 7. (A) Transverse-component G1 seismograms computed with physical-dispersion-only and full attenuation (top row) and their associated adjoint sources (bottom row). (B) Transverse-component G2 seismograms computed with physical-dispersion-only and full attenuation (top row) and their associated adjoint sources, where the measurements (i.e. ΔT) are set to -1 (bottom row). CC and MT denote cross-correlation traveltime and multitaper measurements, respectively. The multitaper adjoint sources are used to compute the $K_{\beta h}$ kernels presented in Fig. 6. The epicentral distance is 60° , and seismograms were filtered between 40 and 60 s. S40RTS with Crust2.0 is used as the 3-D model to compute the seismograms. Note that we start observing the effect of full attenuation more for Love waves due to their higher sensitivity to crustal variations.

presented at the GPU'2014 Conference in Roma, Italy, in September 2014.

REFERENCES

- Akçelik, V. *et al.*, 2003. High-resolution forward and inverse earthquake modeling on terascale computers, in *Proceedings of the ACM/IEEE Conference on Supercomputing (SC'03)*, pp. 52–72, Phoenix, Arizona, USA.
- Ammari, H., Bretin, E., Garnier, J. & Wahab, A., 2013. Time-reversal algorithms in viscoelastic media, *Eur. J. Appl. Math.*, **24**(4), 565–600.
- Anderson, J.E., Tan, L. & Wang, D., 2012. Time-reversal checkpointing methods for RTM and FWI, *Geophysics*, **77**(4), S93–S103.
- 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.
- Assimaki, D., Kallivokas, L.F., Kang, J.W., Li, W. & Kucukcoban, S., 2012. Time-domain forward and inverse modeling of lossy soils with frequency-independent Q for near-surface applications, *Soil Dyn. Earthq. Eng.*, **43**, 139–159.
- Auroux, D., Blum, J. & Nodet, M., 2011. Diffusive back and forth nudging algorithm for data assimilation, *C. R. Acad. Sci. Paris, Ser. I*, **349**(15–16), 849–854.
- Auroux, D., Bansart, P. & Blum, J., 2013. An evolution of the back and forth nudging for geophysical data assimilation: application to Burgers equation and comparisons, *Inverse Probl. Sci. Eng.*, **21**(3), 399–419.
- Bassin, C., Laske, G. & Masters, G., 2000. The current limits of resolution for surface wave tomography in North America, *EOS, Trans. Am. geophys. Un.*, **81**, F897.
- Blanc, E., Komatitsch, D., Chaljub, E., Lombard, B. & Xie, Z., 2016. Highly accurate stability-preserving optimization of the Zener viscoelastic model, with application to wave propagation in the presence of strong attenuation, *Geophys. J. Int.*, **205**(1), 427–439.
- Carcione, J.M., 2014. *Wave Fields in Real Media: Wave Propagation in Anisotropic, Anelastic, Porous and Electromagnetic Media*, 3rd edn, Elsevier Science.
- Charpentier, I., 2001. Checkpointing schemes for adjoint codes: application to the meteorological model Meso-NH, *SIAM J. Sci. Comput.*, **22**(6), 2135–2151.
- Chen, M., Niu, F., Liu, Q., Tromp, J. & Zheng, X., 2015. Multi parameter adjoint tomography of the crust and upper mantle beneath East Asia: 1. Model construction and comparisons, *J. geophys. Res.*, **120**, 1762–1786.
- Cyr, E.C., Shadid, J.N. & Wildey, T., 2015. Towards efficient backward-in-time adjoint computations using data compression techniques, *Comput. Methods Appl. Mech. Eng.*, **288**, 24–44.
- Dahlen, F.A. & Baig, A.M., 2002. Fréchet kernels for body-wave amplitudes, *Geophys. J. Int.*, **150**, 440–466.
- Dahlen, F.A. & Tromp, J., 1998. *Theoretical Global Seismology*, Princeton Univ. Press, 944 pp.
- Dahlen, F.A., Hung, S.-H. & Nolet, G., 2000. Fréchet kernels for finite-frequency traveltimes—I. Theory, *Geophys. J. Int.*, **141**(1), 157–174.
- Douma, H., Yingst, D., Vasconcelos, I. & Tromp, J., 2010. On the connection between artifact filtering in reverse-time migration and adjoint tomography, *Geophysics*, **75**(6), S219–S223.
- Dziwowski, A.M. & Anderson, D.L., 1981. Preliminary reference Earth model, *Phys. Earth planet. Inter.*, **25**(4), 297–356.
- Felgenhauer, U., 1992. Quasi-Newton descent methods with inexact gradients, in *Operations Research '91*, pp. 83–86, eds Gritzmann, P., Hettich, R., Horst, R. & Sachs, E., Physica-Verlag, Heidelberg, Germany.
- Fichtner, A., 2010. *Full Seismic Waveform Modelling and Inversion*, Advances in Geophysical and Environmental Mechanics and Mathematics, Springer-Verlag, Berlin, Germany, 343 pp.
- Fichtner, A. & van Driel, M., 2014. Models and Fréchet kernels for frequency-(in)dependent Q , *Geophys. J. Int.*, **198**(3), 1878–1889.
- Fichtner, A., Bunge, H.-P. & Igel, H., 2006. The adjoint method in seismology: I. Theory, *Phys. Earth planet. Inter.*, **157**(1–2), 86–104.
- Fichtner, A., Kennett, B.L.N., Igel, H. & Bunge, H.P., 2009. Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, *Geophys. J. Int.*, **179**(3), 1703–1725.
- Fink, M., Cassereau, D., Derode, A., Prada, C., Roux, P., Tanter, M., Thomas, J.-L. & Wu, F., 2000. Time-reversed acoustics, *Rep. Prog. Phys.*, **63**(12), 1933–1995.
- Griewank, A. & Walther, A., 2000. Algorithm 799: Revolve, an implementation of checkpointing for the reverse or adjoint mode of computational differentiation, *ACM Trans. Math. Softw.*, **26**(1), 19–45.
- Groos, L., Schäfer, M., Forbriger, T. & Bohlen, T., 2014. The role of attenuation in 2D full-waveform inversion of shallow-seismic body and Rayleigh waves, *Geophysics*, **79**(6), R247–R261.
- Hinze, M., Walther, A. & Sternberg, J., 2006. An optimal memory-reduced procedure for calculating adjoints of the instantaneous Navier-Stokes equations, *Optim. Control Appl. Methods*, **27**(1), 19–40.
- Kallivokas, L.F., Fathi, A., Kucukcoban, S., Stokoe, K.H., II, Bielak, J. & Ghattas, O., 2013. Site characterization using full waveform inversion, *Soil Dyn. Earthq. Eng.*, **47**, 62–82.
- Komatitsch, D., 2011. Fluid-solid coupling on a cluster of GPU graphics cards for seismic wave propagation, *C. R. Acad. Sci., Ser. IIb Mec.*, **339**, 125–135.
- 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., 2002. Spectral-element simulations of global seismic wave propagation—I. Validation, *Geophys. J. Int.*, **149**(2), 390–412.
- 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., Erlebacher, G., Göddeke, D. & Michéa, D., 2010. High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster, *J. Comput. Phys.*, **229**(20), 7692–7714.
- Kowar, R. & Scherzer, O., 2011. Photoacoustic imaging taking into account attenuation, in *Mathematical Modeling in Biomedical Imaging II, Lecture Notes in Mathematics*, Vol. 2035, pp. 85–130, Springer-Verlag, Berlin, Germany.
- Kurzmann, A., Przebindowska, A., Köhn, D. & Bohlen, T., 2013. Acoustic full waveform tomography in the presence of attenuation: a sensitivity analysis, *Geophys. J. Int.*, **195**(2), 985–1000.
- Larmat, C., Montagner, J.-P., Fink, M., Capdeville, Y., Tourin, A. & Clévéde, E., 2006. Time-reversal imaging of seismic sources and application to the great Sumatra earthquake, *Geophys. Res. Lett.*, **33**(19), L19312, doi:10.1029/2006GL026336.
- 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, 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.
- Marquering, H., Dahlen, F.A. & Nolet, G., 1999. Three-dimensional sensitivity kernels for finite-frequency traveltimes: the banana-doughnut paradox, *Geophys. J. Int.*, **137**(3), 805–815.
- Monteiller, V., Chevrot, S., Komatitsch, D. & Wang, Y., 2015. Three-dimensional full waveform inversion of short-period teleseismic wavefields based upon the SEM-DSM hybrid method, *Geophys. J. Int.*, **202**(2), 811–827.
- Pakravan, A., Kang, J.W. & Newton, C.M., 2016. A Gauss-Newton full-waveform inversion for material profile reconstruction in viscoelastic semi-infinite solid media, *Inverse Probl. Sci. Eng.*, **24**(3), 393–421.
- Peter, D. *et al.*, 2011. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes, *Geophys. J. Int.*, **186**(2), 721–739.
- Plessix, R.E., 2006. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications, *Geophys. J. Int.*, **167**(2), 495–503.

- Restrepo, J.M., Leaf, G.K. & Griewank, A., 1998. Circumventing storage limitations in variational data assimilation studies, *SIAM J. Sci. Comput.*, **19**(5), 1586–1605.
- 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 traveltimes and normal-mode splitting function measurements, *Geophys. J. Int.*, **184**(3), 1223–1236.
- Ruan, Y. & Zhou, Y., 2010. The effects of 3-D anelasticity (Q) structure on surface wave phase delays, *Geophys. J. Int.*, **181**(1), 479–492.
- Ruan, Y. & Zhou, Y., 2012. The effects of 3-D anelasticity (Q) structure on surface wave amplitudes, *Geophys. J. Int.*, **189**(2), 967–983.
- Rubio Dalmau, F., Hanzich, M., de la Puente, J. & Gutiérrez, N., 2014. Lossy data compression with DCT transforms, in *Proceedings of the EAGE Workshop on High Performance Computing for Upstream*, p. HPC30, Chania, Crete, Greece.
- Sherali, H.D. & Ulular, O., 1990. Conjugate gradient methods using quasi-Newton updates with inexact line searches, *J. Math. Anal. Appl.*, **150**(2), 359–377.
- Shi, Z.-J., 2006. Convergence of quasi-Newton method with new inexact line search, *J. Math. Anal. Appl.*, **315**(1), 120–131.
- Spears, Z., Corrigan, A.T. & Kailasanath, K., 2014. Checkpointing methods for adjoint-based supersonic jet noise reduction, in *Proceedings of the 20th AIAA/CEAS Aeroacoustics Conference*, Vol. 1, pp. 694–699, American Institute of Aeronautics and Astronautics, Atlanta, Georgia, USA, AIAA paper 2014-2472.
- Sun, W. & Fu, L.-Y., 2013. Two effective approaches to reduce data storage in reverse time migration, *Comput. Geosci.*, **56**, 69–75.
- Symes, W.W., 2007. Reverse time migration with optimal checkpointing, *Geophysics*, **72**(5), SM213–SM221.
- Tanter, M. & Fink, M., 2014. Ultrafast imaging in biomedical ultrasound, *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, **61**(1), 102–119.
- Tarantola, A., 1986. A strategy for non linear inversion of seismic reflection data, *Geophysics*, **51**(10), 1893–1903.
- Tarantola, A., 1987. *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation*, Elsevier Science.
- Tarantola, A., 1988. Theoretical background for the inversion of seismic waveforms, including elasticity and attenuation, *Pure appl. Geophys.*, **128**, 365–399.
- 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, *Commun. Comput. Phys.*, **3**(1), 1–32.
- Vai, R., Castillo-Covarrubias, J.M., Sánchez-Sesma, F.J., Komatitsch, D. & Vilotte, J.P., 1999. Elastic wave propagation in an irregularly layered medium, *Soil Dyn. Earthq. Eng.*, **18**(1), 11–18.
- Varela, C.L., Rosa, A.L.R. & Ulrych, T.J., 1993. Modeling of attenuation and dispersion, *Geophysics*, **58**(8), 1167–1173.
- Virieux, J. & Operto, S., 2009. An overview of full-waveform inversion in exploration geophysics, *Geophysics*, **74**(6), WCC1–WCC26.
- Wang, Q., Moin, P. & Iaccarino, G., 2009. Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation, *SIAM J. Sci. Comput.*, **31**(4), 2549–2567.
- Yuan, S., Wen, S., Li, H., Zhang, X. & Liu, Q., 2014. An optimization framework for adjoint-based climate simulations: a case study of the Zebiak-Cane model, *Int. J. High Perform. Comput. Appl.*, **28**(2), 174–182.
- Zhou, Y., Liu, Q. & Tromp, J., 2011. Surface wave sensitivity: mode summation versus adjoint SEM, *Geophys. J. Int.*, **187**(3), 1560–1576.
- Zhu, T., 2014. Time-reverse modelling of acoustic wave propagation in attenuating media, *Geophys. J. Int.*, **197**(1), 483–494.
- Zhu, H. & Tromp, J., 2013. Mapping tectonic deformation in the crust and upper mantle beneath Europe and the North Atlantic ocean, *Science*, **341**(6148), 871–875.
- Zhu, H., Bozdağ, E., Peter, D. & Tromp, J., 2012. Structure of the European upper mantle revealed by adjoint tomography, *Nature Geosci.*, **5**(7), 493–498.
- Zhu, T., Harris, J.M. & Biondi, B., 2014. Q -compensated reverse-time migration, *Geophysics*, **79**(3), S77–S87.