

Frequency-Domain Numerical Modelling of Visco-Acoustic Waves with Finite-Difference and Finite-Element Discontinuous Galerkin Methods

Romain Brossier^{1,2}, Vincent Etienne¹, Stéphane Operto¹ and Jean Virieux²

¹*Geoazur - CNRS - UNSA - IRD - OCA*

²*LGIT - University Joseph Fourier
France*

1. Introduction

Seismic exploration is one of the main geophysical methods to extract quantitative inferences about the Earth's interior at different scales from the recording of seismic waves near the surface. Main applications are civil engineering for cavity detection and landslide characterization, site effect modelling for seismic hazard, CO₂ sequestration and nuclear-waste storage, oil and gas exploration, and fundamental understanding of geodynamical processes. Acoustic or elastic waves are emitted either by controlled sources or natural sources (i.e., earthquakes). Interactions of seismic waves with the heterogeneities of the subsurface provide indirect measurements of the physical properties of the subsurface which govern the propagation of elastic waves (compressional and shear wave speeds, density, attenuation, anisotropy). Quantitative inference of the physical properties of the subsurface from the recordings of seismic waves at receiver positions is the so-called seismic inverse problem that can be recast in the framework of local numerical optimization. The most complete seismic inversion method, the so-called full waveform inversion (Virieux & Operto (2009) for a review), aims to exploit the full information content of seismic data by minimization of the misfit between the full seismic wavefield and the modelled one. The theoretical resolution of full waveform inversion is half the propagated wavelength. In full waveform inversion, the full seismic wavefield is generally modelled with volumetric methods that rely on the discretization of the wave equation (finite difference, finite element, finite volume methods).

In the regime of small deformations associated with seismic wave propagation, the subsurface can be represented by a linear elastic solid parameterized by twenty-one elastic constants and the density in the framework of the constitutive Hooke's law. If the subsurface is assumed isotropic, the elastic constants reduce to two independent parameters, the Lamé parameters, which depend on the compressional (P) and the shear (S) wave speeds. In marine environment, the P wave speed has most of the time a dominant footprint in the seismic wavefield, in particular, on the hydrophone component which records the pressure wavefield. The dominant footprint of the P wave speed on the seismic

Source: Acoustic Waves, Book edited by: Don W. Dissanayake,
ISBN 978-953-307-111-4, pp. 466, September 2010, Sciyo, Croatia, downloaded from SCIYO.COM

wavefield has prompted many authors to develop and apply seismic modelling and inversion under the acoustic approximation, either in the time domain or in the frequency domain.

This study focuses on frequency-domain modelling of acoustic waves as a tool to perform seismic imaging in the acoustic approximation. In the frequency-domain, wave modelling reduces to the resolution of a complex-valued large and sparse system of linear equations for each frequency, the solution of which is the monochromatic wavefield and the right-hand side (r.h.s) is the source. Two key issues in frequency-domain wave modelling concern the linear algebra technique used to solve the linear system and the numerical method used for the discretization of the wave equation. The linear system can be solved with Gauss elimination techniques based on sparse direct solver (e.g., Duff et al.; 1986), Krylov-subspace iterative methods (e.g., Saad; 2003) or hybrid direct/iterative method and domain decomposition techniques (e.g., Smith et al.; 1996). In the framework of seismic imaging applications which involve a large number of seismic sources (i.e., r.h.s), one motivation behind the frequency-domain formulation of acoustic wave modelling has been to develop efficient approaches for multi-r.h.s modelling based on sparse direct solvers (Marfurt; 1984). A sparse direct solver performs first a LU decomposition of the matrix which is independent of the source followed by forward and backward substitutions for each source to get the solution (Duff et al.; 1986). This strategy has been shown to be efficient for 2D applications of acoustic full waveform inversion on realistic synthetic and real data case studies (Virieux & Operto; 2009). Two drawbacks of the direct-solver approach are the memory requirement of the LU decomposition resulting from the fill-in of the matrix during the LU decomposition (namely, the additional non zero coefficients introduced during the elimination process) and the limited scalability of the LU decomposition on large-scale distributed-memory platforms. It has been shown however that large-scale 2D acoustic problems involving several millions of unknowns can be efficiently tackled thanks to recent development of high-performance parallel solvers (e.g., MUMPS team; 2009), while 3D acoustic case studies remain limited to computational domains involving few millions of unknowns (Operto et al.; 2007). An alternative approach to solve the time-harmonic wave equation is based on Krylov-subspace iterative solvers (Riyanti et al.; 2006; Plessix; 2007; Riyanti et al.; 2007). Iterative solvers are significantly less memory demanding than direct solvers but the computational time linearly increases with the number of r.h.s. Moreover, the impedance matrix, which results from the discretization of the wave equation, is indefinite (the real eigenvalues change sign), and therefore ill-conditioned. Designing efficient pre-conditioner for the Helmholtz equation is currently an active field of research (Erlangga & Nabben; 2008). Efficient preconditioners based on one cycle of multigrid applied to the damped wave equation have been developed and leads to a linear increase of the number of iterations with frequency when the grid interval is adapted to the frequency (Erlangga et al.; 2006). This makes the time complexity of the iterative approaches to be $\mathcal{O}(N^4)$, where N denotes the dimension of the 3D N^3 cubic grid. Intermediate approaches between the direct and iterative approaches are based on domain decomposition methods and hybrid direct/iterative solvers. In the hybrid approach, the iterative solver is used to solve a reduced system for interface unknowns shared by adjacent subdomains while the sparse direct solver is used to factorize local impedance matrices assembled on each subdomains during a preprocessing step (Haidar; 2008; Sourbier et al.; 2008). A short review of the time and memory complexities of the direct, iterative and hybrid approaches is provided in Virieux et al. (2009).

The second issue concerns the numerical scheme used to discretize the wave equation. Most of the methods that have been developed for seismic acoustic wave modelling in the frequency domain rely on the finite difference (FD) method. This can be justified by the fact that, in many geological environments such as offshore sedimentary basins, the subsurface of the earth can be viewed as a weakly-contrasted medium at the scale of the seismic wavelengths, for which FD methods on uniform grid provide the best compromise between accuracy and computational efficiency. In the FD time-domain method, high-order accurate stencils are generally designed to achieve the best trade-off between accuracy and computational efficiency (Dablain; 1986). However, direct-solver approaches in frequency-domain modelling prevent the use of such high-order accurate stencils because their large spatial support will lead to a prohibitive fill-in of the matrix during the LU decomposition (Stekl & Pratt; 1998; Hustedt et al.; 2004). Another discretization strategy, referred to as the mixed-grid approach, has been therefore developed to perform frequency-domain modelling with direct solver: it consists of the linear combination of second-order accurate stencils built on different rotated coordinate systems combined with an anti-lumped mass strategy, where the mass term is spatially distributed over the different nodes of the stencil (Jo et al.; 1996). The combination of these two tricks allows one to design both compact and accurate stencils in terms of numerical anisotropy and dispersion.

Sharp boundaries of arbitrary geometry such as the air-solid interface at the free surface are often discretized along staircase boundaries of the FD grid, although embedded boundary representation has been proposed (Lombard & Piraux; 2004; Lombard et al.; 2008; Mattsson et al.; 2009), and require dense grid meshing for accurate representation of the medium. The lack of flexibility to adapt the grid interval to local wavelengths, although some attempts have been performed in this direction (e.g., Pitarka; 1999; Taflove & Hagness; 2000), is another drawback of FD methods. These two limitations have prompted some authors to develop finite-element methods in the time domain for seismic wave modelling on unstructured meshes. The most popular one is the high-order spectral element method (Seriani & Priolo; 1994; Priolo et al.; 1994; Faccioli et al.; 1997) that has been popularized in the field of global scale seismology by Komatitsch and Vilotte (1998); Chaljub et al. (2007). A key feature of the spectral element method is the combined use of Lagrange interpolants and Gauss-Lobatto-Legendre quadrature that makes the mass matrix diagonal and, therefore, the numerical scheme explicit in time-marching algorithms, and allows for spectral convergence with high approximation orders (Komatitsch & Vilotte; 1998). The selected quadrature formulation leads to quadrangle (2D) and hexahedral (3D) meshes, which strongly limit the geometrical flexibility of the discretization. Alternatively, discontinuous form of the finite-element method, the so-called discontinuous Galerkin (DG) method (Hesthaven & Warburton; 2008), popularized in the field of seismology by Kaser, Dumbser and co-workers (e.g., Dumbser & Käser; 2006) has been developed. In the DG method, the numerical scheme is strictly kept local by duplicating variables located at nodes shared by neighboring cells. Consistency between the multiply defined variables is ensured by consistent estimation of numerical fluxes at the interface between two elements. Numerical fluxes at the interface are introduced in the weak form of the wave equation by means of integration by part followed by application of the Gauss's theorem. Key advantages of the DG method compared to the spectral element method is its capacity of considering triangular (2D) and tetrahedral (3D) non-conform meshes. Moreover, the uncoupling of the elements provides a higher level of flexibility to locally adapt the size of

the elements (h adaptivity) and the interpolation orders within each element (p adaptivity) because neighboring cells exchange information across interfaces only. Moreover, the DG method provides a suitable framework to implement any kind of physical boundary conditions involving possible discontinuity at the interface between elements. One example of application which takes fully advantage of the discontinuous nature of the DG method is the modelling of the rupture dynamics (BenJemaa et al.; 2007, 2009; de la Puente et al.; 2009). The dramatic increase of the total number of degrees of freedom compared to standard finite-element methods, that results from the uncoupling of the elements, might prevent an efficient use of DG methods. This is especially penalizing for frequency-domain methods based on sparse direct solver where the computational cost scales with the size of the matrix N in $\mathcal{O}(N^6)$ for 3D problems. The increase of the size of the matrix should however be balanced by the fact the DG schemes are more local and sparser than FEM ones (Hesthaven & Warburton; 2008), which makes smaller the numerical bandwidth of the matrix to be factorized.

When a zero interpolation order is used in cells (piecewise constant solution), the DG method reduces to the finite volume method (LeVeque; 2002). The DG method based on high-interpolation orders has been mainly developed in the time domain for the elastodynamic equations (e.g., Dumbser & Käser; 2006). Implementation of the DG method in the frequency domain has been presented by Dolean et al. (2007, 2008) for the time-harmonic Maxwell equations and a domain decomposition method has been used to solve the linear system resulting from the discretization of the Maxwell equations. A parsimonious finite volume method on equilateral triangular mesh has been presented by Brossier et al. (2008) to solve the 2D P-SV elastodynamic equations in the frequency domain. The finite-volume approach of Brossier et al. (2008) has been extended to low-order DG method on unstructured triangular meshes in Brossier (2009).

We propose a review of these two quite different numerical methods, the mixed-grid FD method with simple regular-grid meshing and the DG method with dense unstructured meshing, when solving frequency-domain visco-acoustic wave propagation with sparse direct solver in different fields of application. After a short review of the time-harmonic visco-acoustic wave equation, we first review the mixed-grid FD method for 3D modelling. We first discuss the accuracy of the scheme which strongly relies on the optimization procedure designed to minimize the numerical dispersion and anisotropy. Some key features of the FD method such as the absorbing and free-surface boundary conditions and the source excitation on coarse FD grids are reviewed. Then, we present updated numerical experiments performed with the last release of the massively-parallel sparse direct solver MUMPS (Amestoy et al.; 2006). We first assess heuristically the memory complexity and the scalability of the LU factorization. Second, we present simulations in two realistic synthetic models representative of oil exploration targets. We assess the accuracy of the solutions and the computational efficiency of the mixed-grid FD frequency-domain method against that of a conventional FD time-domain method. In the second part of the study, we review the DG frequency-domain method applied to the first-order acoustic wave equation for pressure and particle velocities. After a review of the spatial discretization, we discuss the impact of the order of the interpolating Lagrange polynomials on the computational cost of the frequency-domain DG method and we present 2D numerical experiments on unstructured triangular meshes to highlight the fields of application where the DG method should outperform the FD method.

Although the numerical methods presented in this study were originally developed for seismic applications, they can provide a useful framework for other fields of application such as computational ocean acoustics (Jensen et al.; 1994) and electrodynamic (Taflove & Hagness; 2000).

2. Frequency-domain acoustic wave equation

Following standard Fourier transformation convention, the 3D acoustic first-order velocity-pressure system can be written in the frequency domain as

$$\begin{aligned}
 -i\omega p(x, y, z, \omega) &= \kappa(x, y, z) \left(\frac{\partial v_x(x, y, z, \omega)}{\partial x} + \frac{\partial v_y(x, y, z, \omega)}{\partial y} + \frac{\partial v_z(x, y, z, \omega)}{\partial z} \right) \\
 -i\omega v_x(x, y, z, \omega) &= b(x, y, z) \cdot \frac{\partial p(x, y, z, \omega)}{\partial x} + f_x(x, y, z, \omega) \\
 -i\omega v_y(x, y, z, \omega) &= b(x, y, z) \cdot \frac{\partial p(x, y, z, \omega)}{\partial y} + f_y(x, y, z, \omega) \\
 -i\omega v_z(x, y, z, \omega) &= b(x, y, z) \cdot \frac{\partial p(x, y, z, \omega)}{\partial z} + f_z(x, y, z, \omega),
 \end{aligned} \tag{1}$$

where ω is the angular frequency, $\kappa(x, y, z)$ is the bulk modulus, $b(x, y, z)$ is the buoyancy, $p(x, y, z, \omega)$ is the pressure, $v_x(x, y, z, \omega)$, $v_y(x, y, z, \omega)$, $v_z(x, y, z, \omega)$ are the components of the particle velocity vector. $f_x(x, y, z, \omega)$, $f_y(x, y, z, \omega)$, $f_z(x, y, z, \omega)$ are the components of the external forces. The first block row of equation 1 is the time derivative of the Hooke’s law, while the three last block rows are the equation of motion in the frequency domain.

The first-order system can be recast as a second-order equation in pressure after elimination of the particle velocities in equation 1, that leads to a generalization of the Helmholtz equation:

$$\frac{\omega^2}{\kappa(\mathbf{x})} p(\mathbf{x}, \omega) + \frac{\partial}{\partial x} b(\mathbf{x}) \frac{\partial p(\mathbf{x}, \omega)}{\partial x} + \frac{\partial}{\partial y} b(\mathbf{x}) \frac{\partial p(\mathbf{x}, \omega)}{\partial y} + \frac{\partial}{\partial z} b(\mathbf{x}) \frac{\partial p(\mathbf{x}, \omega)}{\partial z} = s(\mathbf{x}, \omega), \tag{2}$$

where $\mathbf{x} = (x, y, z)$ and $s(\mathbf{x}, \omega) = \nabla \cdot \mathbf{f}$ denotes the pressure source. In exploration seismology, the source is generally a local point source corresponding to an explosion or a vertical force. Attenuation effects of arbitrary complexity can be easily implemented in equation 2 using complex-valued wave speeds in the expression of the bulk modulus, thanks to the correspondence theorem transforming time convolution into products in the frequency domain. For example, according to the Kolsky-Futterman model (Kolsky; 1956; Futterman; 1962), the complex wave speed \bar{c} is given by:

$$\bar{c} = c \left[\left(1 + \frac{1}{\pi Q} |\log(\omega / \omega_r)| \right) + i \frac{\text{sgn}(\omega)}{2Q} \right]^{-1}, \tag{3}$$

where the P wave speed is denoted by c , the attenuation factor by Q and a reference frequency by ω_r .

Since the relationship between the wavefields and the source terms is linear in the first-order and second-order wave equations, equations 1 and 2 can be recast in matrix form:

$$[\mathbf{M} + \mathbf{S}]\mathbf{u} = \mathbf{A}\mathbf{u} = \mathbf{b}, \quad (4)$$

where \mathbf{M} is the mass matrix, \mathbf{S} is the complex stiffness/damping matrix. The sparse impedance matrix \mathbf{A} has complex-valued coefficients which depend on medium properties and angular frequency. The wavefield (either the scalar pressure wavefield or the pressure-velocity wavefields) is denoted by the vector \mathbf{u} and the source by \mathbf{b} (Marfurt; 1984). The dimension of the square matrix \mathbf{A} is the number of nodes in the computational domain multiplied by the number of wavefield components. The matrix \mathbf{A} has a symmetric pattern for the FD method and the DG method discussed in this study but is generally not symmetric because of absorbing boundary conditions along the edges of the computational domain. In this study, we shall solve equation 4 by Gaussian elimination using sparse direct solver. A direct solver performs first a LU decomposition of \mathbf{A} followed by forward and backward substitutions for the solutions (Duff et al.; 1986).

$$\mathbf{A}\mathbf{u} = (\mathbf{L}\mathbf{U})\mathbf{u} = \mathbf{b} \quad (5)$$

$$\mathbf{L}\mathbf{y} = \mathbf{b}; \quad \mathbf{U}\mathbf{u} = \mathbf{y} \quad (6)$$

Exploration seismology requires to perform seismic modelling for a large number of sources, typically, up to few thousands for 3D acquisition. Therefore, our motivation behind the use of direct solver is the efficient computation of the solutions of the equation 4 for multiple sources. The LU decomposition of \mathbf{A} is a time and memory demanding task but is independent of the source, and, therefore is performed only once, while the substitution phase provides the solution for multiple sources efficiently. One bottleneck of the direct-solver approach is the memory requirement of the LU decomposition resulting from the fill-in, namely, the creation of additional non-zero coefficients during the elimination process. This fill-in can be minimized by designing compact numerical stencils that allow for the minimization of the numerical bandwidth of the impedance matrix. In the following, we shall review a FD method and a finite-element DG method that allow us to fulfill this requirement.

3. Mixed-grid finite-difference method

3.1 Discretization of the differential operators

In FD methods, high-order accurate stencils are generally designed to achieve the best tradeoff between accuracy and computational efficiency (Dablain; 1986). However, direct-solver methods prevent the use of high-order accurate stencils because their large spatial support will lead to a prohibitive fill-in of the matrix during the LU decomposition (Hustedt et al.; 2004). Alternatively, the mixed-grid method was proposed by Jo et al. (1996) to design both accurate and compact FD stencils. The governing idea is to discretize the differential operators of the stiffness matrix with different second-order accurate stencils and to linearly combine the resulting stiffness matrices with appropriate weighting coefficients. The different stencils are built by discretizing the differential operators along different rotated coordinate systems (\bar{x} , \bar{y} , \bar{z}) such that their axes span as many directions as possible in the FD cell to mitigate numerical anisotropy. In practice, this means that the partial derivatives with respect to x , y and z in equations 1 or 2 are replaced by a linear combination of partial derivatives with respect to \bar{x} , \bar{y} and \bar{z} using the chain rule followed by the

discretization of the differential operators along the axis \bar{x} , \bar{y} and \bar{z} . In 2D, the coordinate systems are the classic Cartesian one and the 45°-rotated one (Saenger et al.; 2000) which lead to the 9-point stencil (Jo et al.; 1996). In 3D, three coordinate systems have been identified (Operto et al.; 2007) (Figure 1): [1] the Cartesian one which leads to the 7-point stencil, [2] three coordinate systems obtained by rotating the Cartesian system around each Cartesian axis x , y , and z . Averaging of the three elementary stencils leads to a 19-point stencil. [3] four coordinate systems defined by the four main diagonals of the cubic cell. Averaging of the four elementary stencils leads to the 27-point stencil. The stiffness matrix associated with the 7-point stencil, the 19-point stencil and the 27-point stencil will be denoted by \mathbf{S}_1 , \mathbf{S}_2 , \mathbf{S}_3 , respectively.

The mixed-grid stiffness matrix \mathbf{S}_{mg} is a linear combination of the stiffness matrices just-mentioned:

$$\mathbf{S}_{mg} = w_1\mathbf{S}_1 + \frac{w_2}{3}\mathbf{S}_2 + \frac{w_3}{4}\mathbf{S}_3, \tag{7}$$

where we have introduced the weighting coefficients w_1 , w_2 and w_3 which satisfy:

$$w_1 + w_2 + w_3 = 1 \tag{8}$$

In the original mixed-grid approach (Jo et al.; 1996), the discretization on the different coordinate systems was directly applied to the second-order wave equation, equation 2, with the second-order accurate stencil of Boore (1972). Alternatively, Hustedt et al. (2004) proposed to discretize first the first-order velocity-pressure system, equation 1, with second-order staggered-grid stencils (Yee; 1966; Virieux; 1986; Saenger et al.; 2000) and, second, to




Fig. 1. Elementary FD stencils of the 3D mixed-grid stencil. Circles are pressure grid points. Squares are positions where buoyancy needs to be interpolated in virtue of the staggered-grid geometry. Gray circles are pressure grid points involved in the stencil. a) Stencil on the classic Cartesian coordinate system. This stencil incorporates 7 coefficients. b) Stencil on the rotated Cartesian coordinate system. Rotation is applied around x on the figure. This stencil incorporates 11 coefficients. Same strategy can be applied by rotation around y and z . Averaging of the 3 resultant stencils defines a 19-coefficient stencil. c) Stencil obtained from 4 coordinate systems, each of them being associated with 3 main diagonals of a cubic cell. This stencil incorporates 27 coefficients (Operto et al.; 2007).

eliminate the auxiliary wavefields (i.e., the velocity wavefields) following a parsimonious staggered-grid method originally developed in the time domain (Luo & Schuster; 1990). The parsimonious staggered-grid strategy allows us to minimize the number of wavefield components involved in the equation 4, and therefore to minimize the size of the system to be solved while taking advantage of the flexibility of the staggered-grid method to discretize first-order difference operators. The parsimonious mixed-grid approach originally proposed by Hustedt et al. (2004) for the 2D acoustic wave equation was extended to the 3D wave equation by Operto et al. (2007) and to a 2D pseudo-acoustic wave equation for transversely isotropic media with tilted symmetry axis by Operto et al. (2009). The staggered-grid method requires interpolation of the buoyancy in the middle of the FD cell which should be performed by volume harmonic averaging (Moczo et al.; 2002).

The pattern of the impedance matrix inferred from the 3D mixed-grid stencil is shown in Figure 2. The bandwidth of the matrix is of the order of N^2 (N denotes the dimension of a 3D cubic N^3 domain) and was kept minimal thanks to the use of low-order accurate stencils.




Fig. 2. Pattern of the square impedance matrix discretized with the 27-point mixed-grid stencil (Operto et al.; 2007). The matrix is band diagonal with fringes. The bandwidth is $\mathcal{O}(2N_1N_2)$ where N_1 and N_2 are the two smallest dimensions of the 3D grid. The number of rows/columns in the matrix is $N_1 \times N_2 \times N_3$. In the figure, $N_1 = N_2 = N_3 = 8$

3.2 Anti-lumped mass

The linear combination of the rotated stencils in the mixed-grid approach is complemented by the distribution of the mass term ω^2/κ in equation 2 over the different nodes of the mixed-grid stencil to mitigate the numerical dispersion:

$$\frac{\omega^2}{\kappa_{000}} p_{000} \Rightarrow \omega^2 \left(w_{m1} \begin{bmatrix} p \\ \kappa \end{bmatrix}_0 + w_{m2} \begin{bmatrix} p \\ \kappa \end{bmatrix}_1 + w_{m3} \begin{bmatrix} p \\ \kappa \end{bmatrix}_2 + w_{m4} \begin{bmatrix} p \\ \kappa \end{bmatrix}_3 \right), \quad (9)$$

where

$$w_{m1} + \frac{w_{m2}}{6} + \frac{w_{m3}}{12} + \frac{w_{m4}}{8} = 1. \tag{10}$$

In equation 9, the different nodes of the 27-point stencils are labelled by indices lmn where $l, m, n \in \{-1, 0, 1\}$ and 000 denotes the grid point in the middle of the stencil.

We used the notations

$$\begin{aligned} \left[\begin{matrix} p \\ \kappa \end{matrix} \right]_0 &= \frac{p_{000}}{\kappa_{000}}, \\ \left[\begin{matrix} p \\ \kappa \end{matrix} \right]_1 &= \frac{p_{100} + p_{010} + p_{001} + p_{-100} + p_{0-10} + p_{00-1}}{\kappa_{100} + \kappa_{010} + \kappa_{001} + \kappa_{-100} + \kappa_{0-10} + \kappa_{00-1}}, \\ \left[\begin{matrix} p \\ \kappa \end{matrix} \right]_2 &= \frac{p_{110} + p_{011} + p_{101} + p_{-110} + p_{0-11} + p_{-101} + p_{1-10} + p_{01-1} + p_{10-1} + p_{-1-10} + p_{0-1-1} + p_{-10-1}}{\kappa_{110} + \kappa_{011} + \kappa_{101} + \kappa_{-110} + \kappa_{0-11} + \kappa_{-101} + \kappa_{1-10} + \kappa_{01-1} + \kappa_{10-1} + \kappa_{-1-10} + \kappa_{0-1-1} + \kappa_{-10-1}}, \\ \left[\begin{matrix} p \\ \kappa \end{matrix} \right]_3 &= \frac{p_{111} + p_{-1-1-1} + p_{-111} + p_{1-11} + p_{11-1} + p_{-1-11} + p_{1-1-1} + p_{-11-1}}{\kappa_{111} + \kappa_{-1-1-1} + \kappa_{-111} + \kappa_{1-11} + \kappa_{11-1} + \kappa_{-1-11} + \kappa_{1-1-1} + \kappa_{-11-1}}. \end{aligned}$$

This anti-lumped mass strategy is opposite to mass lumping used in finite element methods to make the mass matrix diagonal. The anti-lumped mass approach, combined with the averaging of the rotated stencils, allows us to minimize efficiently the numerical dispersion and to achieve an accuracy representative of 4th-order accurate stencil from a linear combination of 2nd-order accurate stencils. The anti-lumped mass strategy introduces four additional weighting coefficients w_{m1} , w_{m2} , w_{m3} and w_{m4} , equations 9 and 10. The coefficients w_1 , w_2 , w_3 , w_{m1} , w_{m2} , w_{m3} and w_{m4} are determined by minimization of the phase-velocity dispersion in infinite homogeneous medium. Alternatives FD methods for designing optimized FD stencils can be found in Holberg (1987); Takeuchi and Geller (2000).

3.3 Numerical dispersion and anisotropy

The dispersion analysis of the 3D mixed-grid stencil was already developed in details in Operto et al. (2007). We focus here on the sensitivity of the accuracy of the mixed-grid stencil to the choice of the weighting coefficients w_1 , w_2 , w_3 , w_{m1} , w_{m2} , w_{m3} . We aim to design an accurate stencil for a discretization criterion of 4 grid points per minimum propagated wavelength. This criterion is driven by the spatial resolution of full waveform inversion, which is half a wavelength. To properly sample subsurface heterogeneities, the size of which is half a wavelength, four grid points per wavelength should be used according to Shannon’s theorem.

Inserting the discrete expression of a plane wave propagating in a 3D infinite homogeneous medium of wave speed c and density equal to 1 in the wave equation discretized with the mixed-grid stencil gives for the normalized phase velocity (Operto et al.; 2007):

$$\tilde{v}_{ph} = \frac{G}{\sqrt{2J}\pi} \sqrt{w_1(3-C) + \frac{w_2}{3}(6-C-B) + \frac{2w_3}{4}(3-3A+B-C)}, \tag{11}$$

where $J = (w_{m1} + 2w_{m2}C + 4w_{m3}B + 8w_{m4}A)$ with

$$\begin{aligned} A &= \cos a \cos b \cos c, \\ B &= \cos a \cos b + \cos a \cos c + \cos b \cos c, \\ C &= \cos a + \cos b + \cos c. \end{aligned}$$

and $a = \frac{2\pi}{G} \cos\phi \cos\theta$; $b = \frac{2\pi}{G} \cos\phi \sin\theta$; $c = \frac{2\pi}{G} \sin\phi$. Here, the normalized phase velocity is the ratio between the numerical phase velocity ω/k and the wave speed c . $G = \frac{\lambda}{h} = \frac{2\pi}{kh}$ is the number of grid points per wavelength λ . ϕ and θ are the incidence angles of the plane wave. We look for the 5 independent parameters w_{m1} , w_{m2} , w_{m3} , w_1 , w_2 which minimize the least-squares norm of the misfit $(1 - \tilde{v}_{ph})$. The two remaining weighting coefficients w_{m4} and w_3 are inferred from equations 8 and 10, respectively. We estimated these coefficients by a global optimization procedure based on a Very Fast Simulating Annealing algorithm (Sen & Stoffa; 1995). We minimize the cost function for 5 angles ϕ and θ spanning between 0 and 45° and for different values of G .

In the following, the number of grid points for which phase velocity dispersion is minimized will be denoted by G_m . The values of the weighting coefficients as a function of G_m are given in Table 1. For high values of G_m , the Cartesian stencil has a dominant contribution (highlighted by the value of w_1), while the first rotated stencil has the dominant contribution for low values of G_m as shown by the value of w_2 . The dominant contribution of the Cartesian stencil for large values of G_m is consistent with the fact that it has a smaller spatial support (i.e., $2 \times h$) than the rotated stencils and a good accuracy for G greater than 10 (Virieux; 1986). The error on the phase velocity is plotted in polar coordinates for four values of G (4, 6, 8, 10) and for $G_m=4$ in Figure 3a. We first show that the phase velocity dispersion is negligible for $G=4$, that shows the efficiency of the optimization. However, more significant error (0.4 %) is obtained for intermediate values of G (for example, $G=6$ in Figure 3a). This highlights the fact that the weighting coefficients were optimally designed to minimize the dispersion for one grid interval in homogeneous media. We show also the good isotropy properties of the stencil, shown by the rather constant phase-velocity error whatever the direction of propagation. The significant phase-velocity error for values of G greater than G_m prompt us to simultaneously minimize the phase-velocity dispersion for four values of G : $G_m=4,6,8,10$ (Figure 3b). We show that the phase-velocity error is now more uniform over the values of G and that the maximum phase-velocity-error was reduced (0.25 % against 0.4 %). However, the nice isotropic property of the mixed-grid stencil was degraded and the phase-velocity dispersion was significantly increased for $G=4$. We conclude that the range of wavelengths propagated in a given medium should drive the discretization criterion used to infer the weighting coefficients of the mixed grid stencil and that a suitable trade-off should be found between the need to manage the heterogeneity of the medium and the need to minimize the error for a particular wavelength. Of note, an optimal strategy might consist of adapting locally the values of the weighting coefficients to the local wave speed during the assembling of the impedance matrix. This strategy was not investigated yet.

G_m	4,6,8,10	4	8	10	20	40
w_{m1}	0.4966390	0.5915900	0.5750648	0.7489436	0.7948160	0.6244839
w_{m2}	7.51233E-02	4.96534E-02	5.76759E-02	1.39044E-02	3.71392E-03	5.06646E-02
w_{m3}	4.38464E-03	5.10851E-03	5.56914E-03	6.38921E-03	5.54043E-03	1.42369E-03
w_{m4}	6.76140E-07	6.14837E-03	1.50627E-03	1.13699E-02	1.45519E-02	6.8055E-03
w_1	5.02480E-05	8.8075E-02	0.133953	0.163825	0.546804	0.479173
w_2	0.8900359	0.8266806	0.7772883	0.7665769	0.1784437	0.2779923
w_3	0.1099138	8.524394E-02	8.87589E-02	6.95979E-02	0.2747527	0.2428351

Table 1. Coefficients of the mixed-grid stencil as a function of the discretization criterion G_m for the minimization of the phase velocity dispersion.

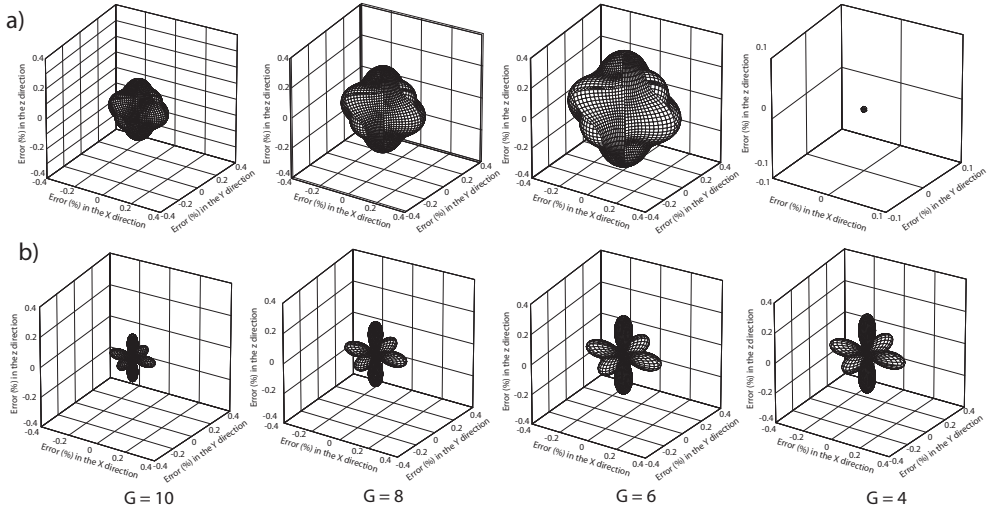


Fig. 3. Phase-velocity dispersion shown in spherical coordinates for four values of G . (a) The phase-velocity dispersion was minimized for $G = 4$. (b) the phase-velocity dispersion was minimized for 4 values of G : 4, 6, 8 and 10.

Comparison between numerical and analytical pressure monochromatic wavefields computed in a homogeneous medium of wave speed 1.5 km/s and density 1000 kg/m³ confirms the former theoretical analysis (Figure 4). The frequency is 3.75 Hz corresponding to a propagated wavelength of 400 m. The grid interval for the simulation is 100 m corresponding to $G = 4$. Simulations were performed when the weighting coefficients of the mixed-grid stencils are computed for $G_m = 4$ and $G_m = \{4, 6, 8, 10\}$. The best agreement is obtained for the weighting coefficients associated with $G_m = 4$ as expected from the dispersion analysis.

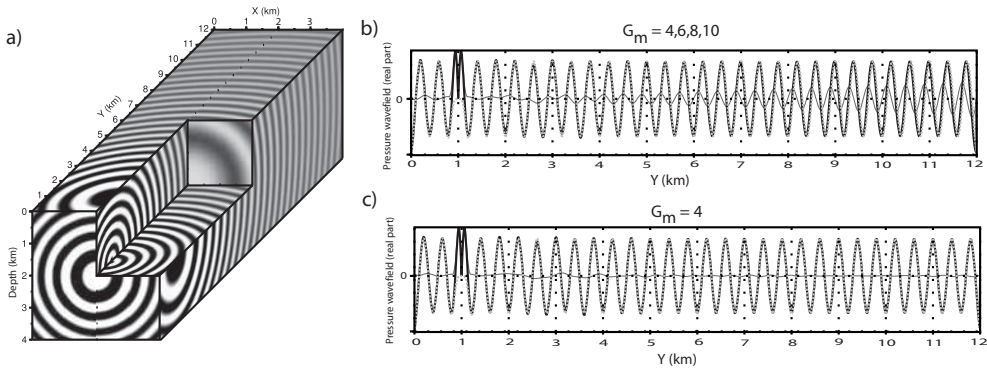


Fig. 4. (a) Real part of a 3.75-Hz monochromatic wavefield computed with the mixed-grid stencil in a 3D infinite homogeneous medium. The explosive point source is at $x=2$ km, $y=1$ km, $z=2$ km. (b-c) Comparison between the analytical (gray) and the numerical solution (black) for a receiver line oriented in the Y direction across the source position. The thin black line is the difference. The amplitudes were corrected for 3D geometrical spreading. (b) $G_m = 4, 6, 8, 10$. (c) $G_m = 4$.

3.4 Boundary conditions

In seismic exploration, two boundary conditions are implemented for wave modelling: absorbing boundary conditions to mimic an infinite medium and free surface conditions on the top side of the computational domain to represent the air-solid or air-water interfaces.

3.4.1 PML absorbing boundary conditions

We use Perfectly-Matched Layers (PML) absorbing boundary conditions (Berenger; 1994) to mimic an infinite medium. In the frequency domain, implementation of PMLs consists of applying in the wave equation a new system of complex-valued coordinates \tilde{x} defined by (e.g., Chew & Weedon; 1994):

$$\frac{\partial}{\partial \tilde{x}} = \frac{1}{\xi_x(x)} \frac{\partial}{\partial x} \quad (12)$$

In the PML layers, the damped wave equation writes:

$$\left[\frac{\omega^2}{\kappa(\mathbf{x})} + \frac{1}{\xi_x(x)} \frac{\partial}{\partial x} \frac{b(\mathbf{x})}{\xi_x(x)} \frac{\partial}{\partial x} + \frac{1}{\xi_y(y)} \frac{\partial}{\partial y} \frac{b(\mathbf{x})}{\xi_y(y)} \frac{\partial}{\partial y} + \frac{1}{\xi_z(z)} \frac{\partial}{\partial z} \frac{b(\mathbf{x})}{\xi_z(z)} \frac{\partial}{\partial z} \right] p(\mathbf{x}, \omega) = -s(\mathbf{x}, \omega), \quad (13)$$

where $\xi_i(x) = 1 + i\gamma_i(x)/\omega$ and $\gamma_i(x)$ is a 1D damping function which defines the PML damping behavior in the PML layers. These functions differ from zero only inside the PML layers. In the PML layers, we used $\gamma(x) = c_{pml} \left(1 - \cos\left(\frac{\pi}{2} \frac{L-x}{L}\right) \right)$ where L denotes the width of the PML layer and x is a local coordinate in the PML layer whose origin is located at the outer edges of the model. The scalar c_{pml} is defined by trial and error depending on the width of the PML layer. The procedure to derive the unsplitted second-order wave equation with PML conditions, equation 13, from the first-order damped wave equation is given in Operto et al. (2007).

The absorption of the PML layers at grazing incidence can be improved by using convolutional PML (C-PML) (Kuzuoglu & Mittra; 1996; Roden & Gedney; 2000; Komatitsch & Martin; 2007). In the C-PML layers, the damping function $\xi_i(x)$ becomes:

$$\xi_x(x) = \kappa_x + i \frac{d_x}{\alpha_x + i\omega}, \quad (14)$$

where d_x and α_x are generally quadratic and linear functions, respectively. Suitable expression for κ_x , d_x and α_x are discussed in Kuzuoglu & Mittra (1996); Collino & Monk (1998); Roden & Gedney (2000); Collino & Tsogka (2001); Komatitsch & Martin (2007); Drossaert & Giannopoulos (2007).

3.4.2 Free surface boundary conditions

Planar free surface boundary conditions can be simply implemented in the frequency domain with two approaches. In the first approach, the free surface matches the top side of the FD grid and the pressure is forced to zero on the free surface by using a diagonal impedance matrix for rows associated with collocation grid points located on the top side of the FD grid. Alternatively, the method of image can be used to implement the free surface along a virtual plane located half a grid interval above the topside of the FD grid (Virieux;

1986). The pressure is forced to vanish at the free surface by using a fictitious plane located half a grid interval above the free surface where the pressure is forced to have opposite values to that located just below the free surface.

From a computer implementation point of view, an impedance matrix is typically built row per row. One row of the linear system can be written as:

$$\sum_{i_3=-1,1} \sum_{i_2=-1,1} \sum_{i_1=-1,1} a_{i_1 i_2 i_3} p_{i_1 i_2 i_3} = s_{000} \tag{15}$$

where $a_{i_1 i_2 i_3}$ are the coefficients of the 27-point mixed grid stencil and 000 denote the indices of the collocation coefficient located in the middle of the stencil in a local coordinate system. The free surface boundary conditions writes:

$$p_{-1i_2 i_3} = -p_{0i_2 i_3}, \tag{16}$$

for $i_2 = \{-1, 0, 1\}$ and $i_3 = \{-1, 0, 1\}$. The indices $i_1=-1$ and $i_1=0$ denotes here the grid points just above and below the free surface, respectively.

For a grid point located on the top side of the computational domain (i.e., half a grid interval below free surface), equation 15 becomes:

$$\sum_{i_3=-1,1} \sum_{i_2=-1,1} a_{1i_2 i_3} p_{1i_2 i_3} + \sum_{i_3=-1,1} \sum_{i_2=-1,1} (a_{0i_2 i_3} - a_{-1i_2 i_3}) p_{0i_2 i_3} = s_{000}, \tag{17}$$

where $p_{-1i_2 i_3}$ has been replaced by the opposite value of $p_{0i_2 i_3}$ according to equation 16. Our practical experience is that both implementation of free surface boundary conditions give results of comparable accuracy. Of note, rigid boundary conditions (zero displacement perpendicular to the boundary) or periodic boundary conditions (Ben-Hadj-Ali et al.; 2008) can be easily implemented with the method of image following the same principle than for the free surface condition.

3.5 Source implementation on coarse grids

Seismic imaging by full waveform inversion is initiated at frequency as small as possible to mitigate the non linearity of the inverse problem. The starting frequency for modelling can be as small as 2 Hz which can lead to grid intervals as large as 200 m. In this framework, accurate implementation of point source at arbitrary position in a coarse grid is critical. One method has been proposed by Hicks (2002) where the point source is approximated by a windowed Sinc function. The Sinc function is defined by

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}, \tag{18}$$

where $x = (x_g - x_s)$, x_g denotes the position of the grid nodes and x_s denotes the position of the source. The Sinc function is tapered with a Kaiser function to limit its spatial support. For multidimensional simulations, the interpolation function is built by tensor product construction of 1D windowed Sinc functions. If the source positions matches the position of one grid node, the Sinc function reduces to a Dirac function at the source position and no approximation is used for the source positioning. If the spatial support of the Sinc function

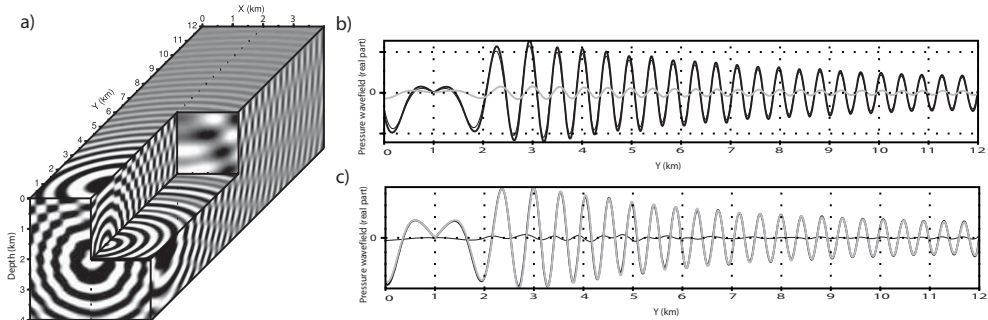


Fig. 5. a) Real part of a 3.75-Hz monochromatic wavefield in a homogeneous half space. (b) Comparison between numerical (black) and analytical (gray) solutions at receiver positions. The Sinc interpolation with 4 coefficients was used for both the source implementation and the extraction of the solution at the receiver positions on a coarse FD grid.

intersects a free surface, part of the Sinc function located above the free surface is mirrored into the computational domain with a reverse sign following the method of image. Vertical force can be implemented in a straightforward way by replacing the Sinc function by its vertical derivative. The same interpolation function can be used for the extraction of the pressure wavefield at arbitrary receiver positions. The accuracy of the method of Hicks (2002) is illustrated in Figure 5 which shows a 3.5-Hz monochromatic wavefield computed in a homogeneous half space. The wave speed is 1.5 km/s and the density is 1000 kg/m³. The grid interval is 100 m. The free surface is half a grid interval above the top of the FD grid and the method of image is used to implement the free surface boundary condition. The source is in the middle of the FD cell at 2 km depth. The receiver line is oriented in the Y direction. Receivers are in the middle of the FD cell in the horizontal plane and at a depth of 6 m just below the free surface. This setting is representative of a ocean bottom survey where the receiver is on the sea floor and the source is just below the sea surface (in virtue of the spatial reciprocity of the Green functions, sources are processed here as receivers and *vice versa*). Comparison between the numerical and the analytical solutions at the receiver positions are first shown when the source is positioned at the closest grid point and the numerical solutions are extracted at the closest grid point (Figure 5b). The amplitude of the numerical solution is strongly overestimated because the numerical solution is extracted at a depth of 50 m below free surface (where the pressure vanishes) instead of 6 m. Second, a significant phase shift between numerical and analytical solutions results from the approximate positioning of the sources and receivers. In contrast, a good agreement between the numerical and analytical solutions both in terms of amplitude and phase is shown in Figure 5c where the source and receiver positioning were implemented with the windowed Sinc interpolation.

3.6 Resolution with the sparse direct solver MUMPS

To solve the sparse system of linear equations, equation 4, we used the massively parallel direct MUMPS solver designed for distributed memory platforms. The reader is referred to Guermouche et al. (2003); Amestoy et al. (2006); MUMPS team (2009) for an extensive description of the method and their underlying algorithmic aspects. The MUMPS solver is based on a multifrontal method (Duff et al.; 1986; Duff and Reid; 1983; Liu; 1992), where the

resolution of the linear system is subdivided into 3 main tasks. The first one is an analysis phase or symbolic factorization. Reordering of the matrix coefficients is first performed in order to minimize fill-in. We used the METIS algorithm which is based on a hybrid multilevel nested-dissection and multiple minimum degree algorithm (Karypis & Kumar; 1999). Then, the dependency graph which describes the order in which the matrix can be factored is estimated as well as the memory required to perform the subsequent numerical factorization. The second task is the numerical factorization. The third task is the solution phase performed by forward and backward substitutions. During the solution phase, multiple-shot solutions can be computed simultaneously from the LU factors taking advantage of threaded BLAS3 (Basic Linear Algebra Subprograms) library and are either assembled on the host or kept distributed on the processors for subsequent parallel computations.

We performed the factorization and the solutions phases in complex arithmetic single precision. To reduce the condition number of the matrix, a row and column scaling is applied in MUMPS before factorization. The sparsity of the matrix and suitable equilibration have made single precision factorization accurate enough so far for the 2D and 3D problems we tackled. If single precision factorization would be considered not accurate enough for very large problems, an alternative approach to double precision factorization may be the postprocessing of the solution by a simple and fast iterative refinement performed in double precision (Demmel (1997), pages 60-61 and Langou et al. (2006); Kurzak & Dongarra (2006)).

The main two bottlenecks of sparse direct solver is the time and memory complexity and the limited scalability of the LU decomposition. By complexity is meant the increase of the computational cost (either in terms of elapsed time or memory) of an algorithm with the size of the problem, while scalability describes the ability of a given algorithm to use an increasing number of processors. The theoretical memory and time complexity of the LU decomposition for a sparse matrix, the pattern of which is shown in Figure 2, is $\mathcal{O}(N^4)$ and $\mathcal{O}(N^6)$, respectively, where N is the dimension of a 3D cubic N^3 grid.

We estimated the observed memory complexity and scalability of the LU factorization by means of numerical experience. The simulations were performed on the SGI ALTIX ICE supercomputer of the computer center CINES (France). Nodes are composed of two quad-core INTEL processors E5472. Each node has 30 Gbytes of useful memory. We used two MPI process per node and four threads per MPI process. In order to estimate the memory complexity, we performed simulations on cubic models of increasing dimension with PML absorbing boundary conditions along the 6 sides of the model. The medium is homogeneous and the source is on the middle of the grid. Figure 6a shows the memory required to store the complex-valued LU factors as a function of N . Normalization of this curve by the real memory complexity will lead to a horizontal line. We found an observed memory complexity of $\mathcal{O}(\text{Log}_2(N)N^{3.9})$ (Figure 6b) which is consistent with the theoretical one. In order to assess the scalability of the LU factorization, we consider a computational FD grid of dimensions $177 \times 107 \times 62$ corresponding to 1.17 millions of unknowns. The size of the grid corresponds to a real subsurface target for oil exploration at low frequency (3.5 Hz). We computed a series of LU factorization using an increasing number of processors N_p , starting with $N_{p_{ref}} = 2$. The elapsed time of the LU factorization (T_{LU}) and the parallelism efficiency ($T_{LU}(N_{p_{ref}}) \times N_{p_{ref}} / T_{LU}(N_p) \times N_p$) are shown in Figure 6(c-d). The efficiency drops rapidly as the number of processors increased, down to a value of 0.5 for $N_p = 32$ (Figure 6d). This clearly indicates that the most suitable platform for sparse direct solver should be composed

of a limited number of nodes with a large amount of shared memory. The efficiency of the multi-r.h.s solution phase is significantly improved by using multithreaded BLAS3 library.

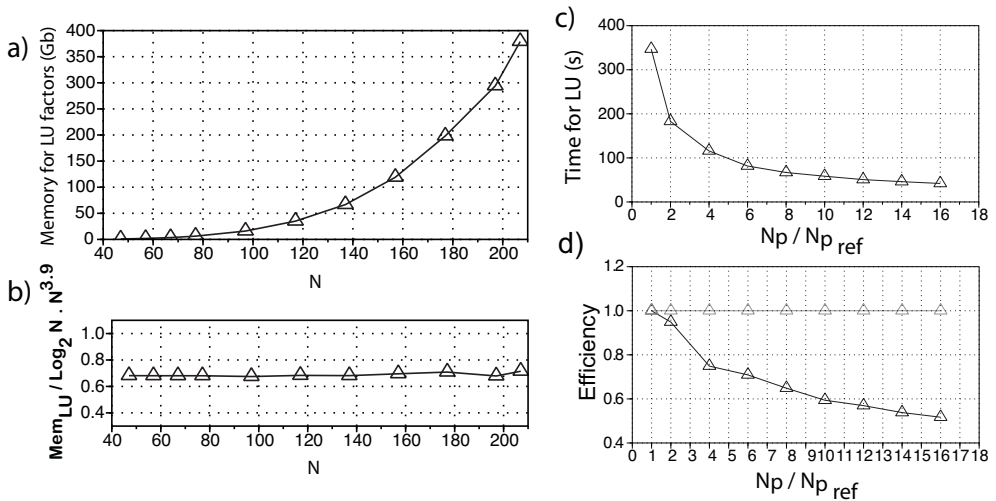


Fig. 6. (a-b) Memory complexity of LU factorization. (a) Memory in Gbytes required for storage of LU factors. (b) Memory required for storage of LU factors divided by $\text{Log}_2 N \cdot N^{3.9}$. N denotes the dimension of a 3D N^3 grid. The largest simulation for $N = 207$ corresponds to 8.87 millions of unknowns. (c-d) Scalability analysis of LU factorization. (c) Elapsed time for LU factorization versus the number of MPI processes. (d) Efficiency.

3.7 Numerical examples

We present acoustic wave modelling in two realistic 3D synthetic velocity models, the SEG/EAGE overthrust and salt models, developed by the oil exploration community to assess seismic modelling and imaging methods (Aminzadeh et al.; 1997). The simulation was performed on the SGI ALTIX ICE supercomputer just described.

3.7.1 3D EAGE/SEG overthrust model

The 3D SEG/EAGE Overthrust model is a constant density onshore acoustic model covering an area of $20 \text{ km} \times 20 \text{ km} \times 4.65 \text{ km}$ (Aminzadeh et al.; 1997) (Figure 7a). From a geological viewpoint, it represents a complex thrust sedimentary succession constructed on top of a structurally decoupled extensional and rift basement block. The overthrust model is discretized with 25 m cubic cells, representing an uniform mesh of $801 \times 801 \times 187$ nodes. The minimum and maximum velocities in the Overthrust model are 2.2 and 6.0 km/s respectively. We present the results of a simulation performed with the mixed-grid FD method (referred to as FDFD in the following) for a frequency of 7 Hz and for a source located at $x=2.4 \text{ km}$, $y=2.4 \text{ km}$ and $z=0.15 \text{ km}$. The model was resampled with a grid interval of 75 m that corresponds to four grid points per minimum wavelength. The size of the resampled FD grid is $266 \times 266 \times 62$. PML layers of 8 grid points were added along the 6 sides of the 3D FD grid. This leads to 6.2 millions of pressure unknowns. For the simulation,

we used the weights of the mixed-grid stencil obtained for $G_m = 4, 6, 8, 10$. These weights provided slightly more accurate results than the weights obtained for $G_m = 4$, in particular for waves recorded at long source-receiver offsets. The 7-Hz monochromatic wavefield computed with the FDFD method is compared with that computed with a classic $\mathcal{O}(\Delta t^2, \Delta x^4)$ staggered-grid FD time-domain (FDTD) method where the monochromatic wavefield is integrated by discrete Fourier transform within the loop over time steps (Sirgue et al.; 2008) (Figure 7).

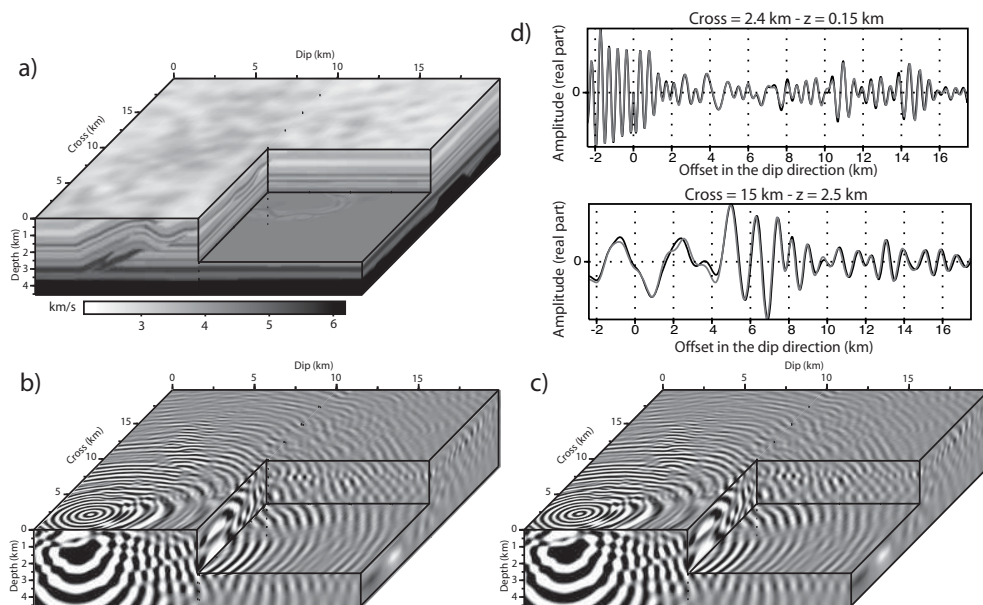


Fig. 7. (a) Overthrust velocity model. (b-c) 7-Hz monochromatic wavefield (real part) computed with the FDFD (b) and FDTD (c) methods. (d) Direct comparison between FDFD (gray) and FDTD (black) solutions. The receiver line in the dip direction is: (top) at 0.15-km depth and at 2.4 km in the cross direction. The amplitudes were corrected for 3D geometrical spreading; (bottom) at 2.5-km depth and at 15 km in the cross direction.

We used the same spatial FD grid for the FDTD and FDFD simulations. The simulation length was 15 s in the FDTD modelling. We obtain a good agreement between the two solutions (Figure 7d). The statistics of the FDFD and FDTD simulations are outlined in Table 2. The FDFD simulation was performed on 32 MPI processes with 2 threads and 15 Gbytes of memory per MPI process. The total memory required by the LU decomposition of the impedance matrix was 260 Gbytes. The elapsed time for LU decomposition was 1822 s and the elapsed time for one r.h.s was 0.97 s. Of note, we processed efficiently groups of 16 sources in parallel during the solution step by taking advantage of the multi-rhs functionality of MUMPS and the threaded BLAS3 library. The elapsed time for the FDTD simulation was 352 s on 4 processors. Of note, C-PML absorbing boundary conditions were implemented in the full model during FDTD modelling to mimic attenuation effects

Thank You for previewing this eBook

You can read the full version of this eBook in different formats:

- HTML (Free /Available to everyone)
- PDF / TXT (Available to V.I.P. members. Free Standard members can access up to 5 PDF/TXT eBooks per month each month)
- Epub & Mobipocket (Exclusive to V.I.P. members)

To download this full book, simply select the format you desire below

