

# Seismic wave modeling for seismic imaging

JEAN VIRIEUX, *University Joseph Fourier*

STÉPHANE OPERTO, HAFEDH BEN-HADJ-ALI, ROMAIN BROSSIER, VINCENT ETIENNE, and FLORENT SOURBIER, *University of Nice Sophia Antipolis*

LUC GIRAUD and AZZAM HAIDAR, *INPT-ENSEEIH*

Many scientific applications require accurate modeling of seismic wave propagation in complex media. These objectives can include fundamental understanding of seismic wave propagation of the Earth on a global scale, including fluid envelopes, mitigation of seismic risk with better quantitative estimates of seismic hazard, and improved exploitation of the natural resources in the crust of the Earth. Accurate quantification is a continual quest, and benchmark protocols have been designed for model definitions and comparison of solutions. Through this quest, an impressive number of numerical tools have been developed, ranging from efficient finite-difference methods to more sophisticated finite-element methods, and including the so-called pseudospectral methods (see Wu and Maupin for a review). The main motivation behind these permanent developments has been to improve the efficiency and accuracy of forward modeling. To achieve this, one systematic choice for both finite-difference and finite-element methods has involved explicit time-stepping integration to avoid matrix inversion.

This study focuses on seismic modeling as a tool for seismic imaging by full waveform inversion (FWI), where recorded and modeled seismograms should fit (Tarantola, 1987). FWI has been recast as a local iterative optimization problem, where the desired model is searched for in the vicinity of a starting model and in the opposite direction of the gradient of the objective function. Frequency-domain FWI (Pratt et al, 1998) has shown distinct advantages over time-domain FWI (Tarantola, 1984), as it introduces a natural multiscale approach by proceeding from the low frequencies to the higher ones. Another advantage of frequency-domain FWI is the ability to manage and process compact volumes of data by limiting the inversion to a limited number of frequency components of wide-aperture/wide-azimuth data (Sirgue and Pratt, 2004).

Various forward-modeling tools have been designed for 2D and 3D time-domain FWI and frequency-domain FWI. For 2D frequency-domain FWI, the forward problem is naturally implemented in the frequency domain with direct solvers. In the 3D case, the optimal strategy for forward modeling is less obvious, and it can range from time-domain modeling (TDM) combined with discrete Fourier transformation to extract the frequency response (Nihei and Li, 2008; Sirgue et al., 2008), to frequency-domain modeling based on direct and hybrid (Soubier et al., 2008) or iterative solvers (Erlangga and Herrmann, 2008). These solvers are used to solve the linear system resulting from the discretization of the time harmonic wave equation. In the following, we review the advantages and drawbacks of various modeling strategies for frequency-domain FWI.

## Seismic modeling for waveform inversion

For FWI, we have identified several criteria that need to be

taken into account when assessing the merits and drawbacks of forward-modeling tools.

The most obvious criterion is the central processing unit (CPU) efficiency for multiple-source modeling. A 3D acquisition survey typically involves thousands of shots and receivers. The estimation of the gradient of the objective function (i.e., the direction along which the perturbation model is searched) requires  $2 \times \min(N_{shot}, N_{rec})$  modeling, where  $N_{shot}$  and  $N_{rec}$  represent the numbers of shots and receivers, respectively. The estimation of the step length (i.e., the distance in the direction of the gradient) can require 1–2 times this number of forward modelings. These forward modelings must be repeated for each inversion iteration.

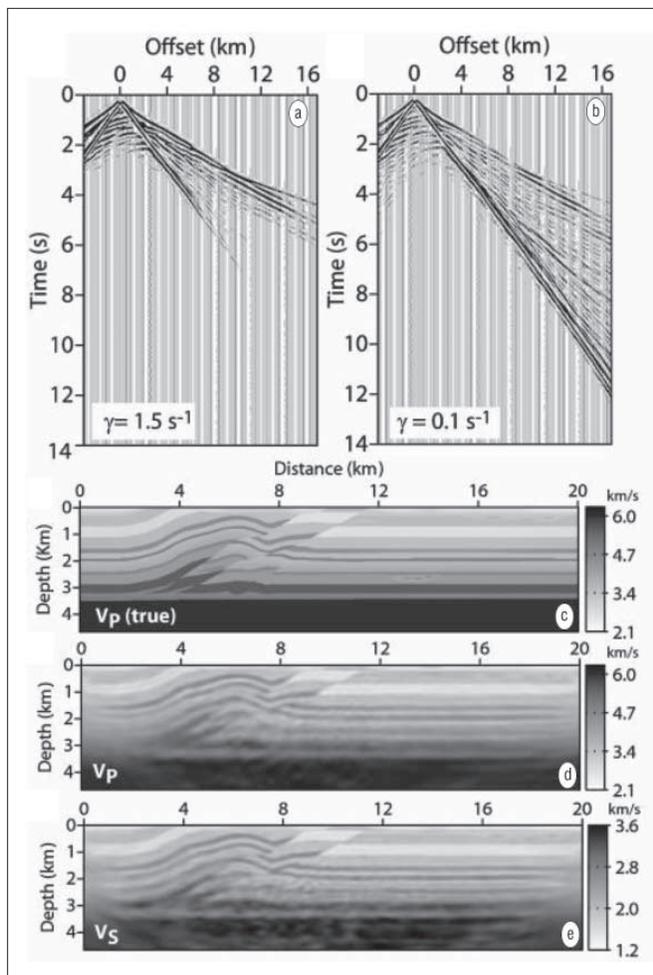
A second criterion is related to the memory requirement of the modeling. We discuss this in detail below when considering various forward-modeling techniques.

A third criterion is related to the scalability of the modeling on a large distributed-memory computational platform, i.e., the ability of a parallel algorithm to use an increasing number of processors. Given the size of the problem (from a few million to a few billion unknowns in the numerical grid), high-performance computing on a large distributed-memory computational platform is unavoidable nowadays for efficient multisource modeling.

A fourth criterion is related to the robustness, which describes the ability of the method to process models of arbitrary levels of heterogeneity.

A fifth criterion is related to the need to reduce the nonlinearity of FWI. This reduction is achieved through a hierarchic multiscale imaging workflow: We proceed from recovering the large wavelengths to small ones. The multiscale character of the imaging can be implemented using various methodologies that operate on different data subsets with increasing resolution power. The data subsets can be designed according to different quantities, such as frequency, offset, aperture angle, and specific arrival-window selection. Frequency-domain FWI is generally performed by successive inversions of single frequencies or of frequency groups of increasingly high-frequency content. Proceeding from the low frequencies to the higher ones helps reduce the risk of cycle-skipping artifacts at the expense of data redundancy. A second regularization level can be introduced by selecting some specific arrivals with respect to time. In the time domain, this selection can be naturally implemented by time windowing. In the frequency domain, complex frequencies are used as a last resort to dampen (in time) the monochromatic wavefields from a given traveltime, typically as the first-arrival traveltime. This is equivalent in spirit to selected aperture angles in the data (Figure 1). This data-driven regularization must be taken into account in assessing the forward method.

A last, sixth criterion is the feasibility of the extension of



**Figure 1.** Elastic FWI of preconditioned onshore data. (a, b) Synthetic seismograms (vertical particle velocities) computed in an elastic version of the overthrust model. A frequency-domain modeling code is used for computation of time-domain seismograms. Complex frequencies that have their imaginary parts,  $\gamma$ , indicated in the figure are used to damp arrivals from the first arrival time. (c) True  $V_p$  model. The  $V_s$  model was inferred from the  $V_p$  model using a constant Poisson's ratio of 0.24. (d, e) Reconstructed  $V_p$  and  $V_s$  models after successive inversion of two overlapping groups of frequencies: 1.7, 2.1, 2.5, and 3.5 Hz and 3.5, 4.1, 4.7, 6, and 7.2 Hz. Without hierarchic inversion of complex-valued frequencies, FWI fails to converge toward acceptable models.

the modeling approach to more realistic physical descriptions of the media. Today, 3D acoustic FWI is feasible and is the field of active research. Extensions to attenuation, elasticity, and anisotropy might need specific improvements.

### Time-domain versus frequency-domain modeling

Numerical methods transform partial differential operators into algebraic operations that can be expressed in matrix notation. Wave equations attempt to estimate the vector  $\vec{f}$  (pressure, solid particle velocities, or fluid/solid particle velocities) through the explicit system in time; for instance:

$$M(x, y, z) d\vec{f} / dt = A(x, y, z) \vec{f} + S(x, y, z, t) \quad (1)$$

where the mass matrix  $M$  is diagonal, and the stiffness matrix

$A$  should operate on the desired solution either in the space domain or by going back and forth in the spectral domain. The source term  $S$  is often a local point source in controlled-source seismology. The corresponding formulation in the frequency domain drives us to a generalization of the Helmholtz equation:

$$B(x, y, z, \omega) \vec{f} = S(x, y, z, \omega) \quad (2)$$

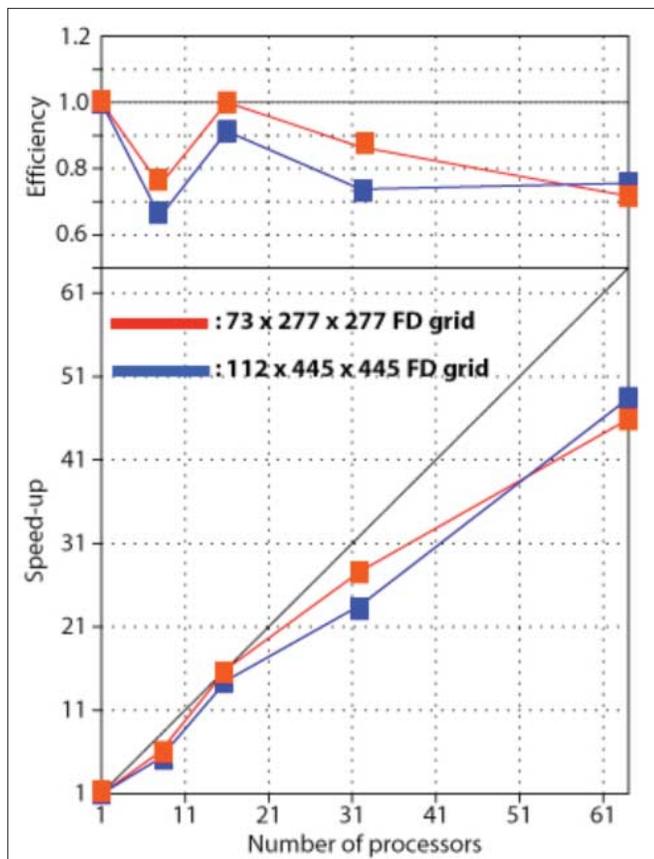
where the impedance matrix  $B$  is complex valued and has a symmetric pattern when a finite discretization is considered. The dimension of the matrix is the number of unknowns in the computational grids. The numerical bandwidth and the number of nonzero coefficients in the matrix depends on the numerical stencil used for the discretization of differential operators embedded in  $A$ . In the frequency domain, wave modeling reduces to the resolution of a large and sparse system of linear equations with multiple right-hand sides (rhs), each rhs corresponding to a seismic source. This linear system can be solved either with direct, hybrid (direct/iterative) or iterative solvers, as will be discussed later.

### Time-domain modeling with explicit integration schemes

TDM is usually performed according to an explicit time-marching algorithm through solving system 1. At each time step, the solution at each spatial grid point is estimated from the solutions at the previous time steps. For frequency-domain FWI, core-memory storage of the full time series is useless as the frequency-domain wavefield is obtained by a discrete Fourier summation over time (Sirgue et al., 2008). The time-domain algorithm can be efficiently parallelized using standard domain-decomposition methods, where the computational domain is divided into subdomains of limited dimension. The efficiency of these algorithms is generally close to 1.0 (Figure 2). In the framework of multiple-source simulations, the low memory requirements of TDM also allow a coarse-grain parallelism over sources that can be combined with domain-decomposition parallelism if the number of processors is significantly greater than the number of sources. In the following discussion, the dimension of a 3D  $N^3$  computational grid is denoted as  $N$ . If only a parallelism over sources is to be implemented, Table 1 gives the memory and time complexities of the TDM. Realistic 3D surveys require a significant amount of memory to store the  $N_{\text{rhs}}$  wavefields distributed over the processors,  $O(N^3 N_{\text{rhs}}) = O(N^5)$ , where the number of sources  $N_{\text{rhs}}$  is over one side of the grid. Propagation in media of arbitrary heterogeneous levels without instabilities is performed accurately by TDM algorithms, provided a sufficiently fine time discretization is used. Extensions are possible for elastic, anisotropic, and attenuation modeling although the memory request increases significantly.

### Frequency-domain modeling based on a direct solver

The limited number of frequencies required by frequency-domain FWI can be efficiently modeled for a large number of sources once the impedance matrix  $B$  has been LU factor-

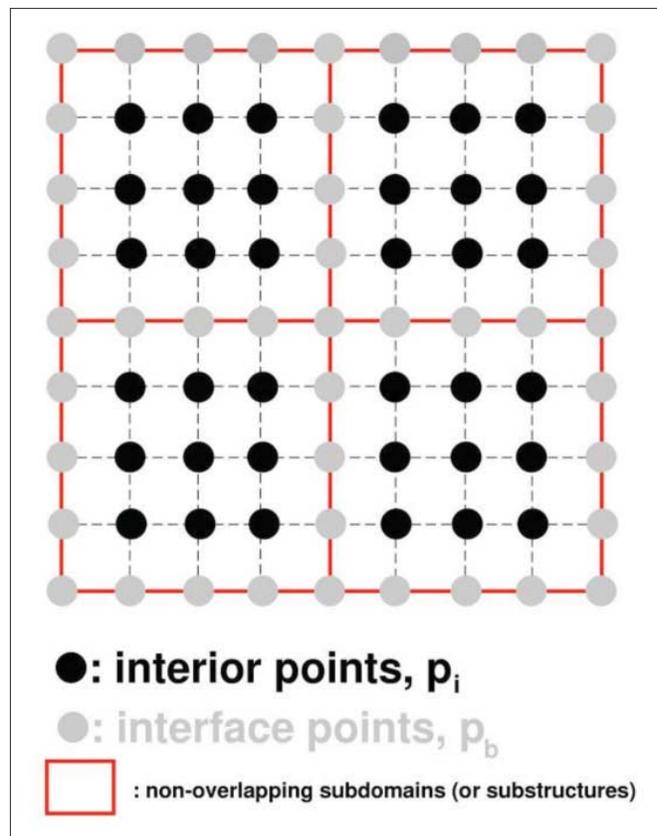


**Figure 2.** Speedup (bottom) and efficiency (top) of TDM. The simulations were performed on the SEG/EAGE overthrust model for two finite-difference grid intervals (46 and 75 m) suitable for a maximum frequency of 7 Hz and 12 Hz, respectively. The dimensions of the two FD grids are provided in the figure. Simulations were performed on the IBM Power 6 at the IDRIS computer center.

ized for each frequency (system 2). This direct-solver method (DSM) has been promoted by Pratt et al. (1998).

Parallelism in DSM is implemented through the use of a massive parallel direct solver, such as MUMPS (<http://graal.ens-lyon.fr/MUMPS/>). After parallel lower-upper triangle matrix factorization (LU), the LU factors remain distributed over processors, allowing us to also perform the solution step (i.e., forward/backward substitution) in parallel. The most limiting factor of direct solvers is their intrinsic limited scalability. Our practical experience suggests that a speedup greater than 15 is difficult to achieve whatever the number of processors used for both 2D and 3D applications. (The speedup and the efficiency are defined by  $T_{ref}/T_n$  and  $T_{ref} \times N_{nref}/T_n \times N_n$ , respectively, where  $N_{ref}$  is the minimum number of processors,  $T_{ref}$  is the related elapsed time, and  $T_n$  is the elapsed time obtained with  $N_n$  processors.)

The memory and time complexities of a direct solver for 2D finite-difference problems are  $O(N^2 \log_2 N)$  and  $O(N^3)$ , respectively, and they increase dramatically in 3D as  $O(N^4)$  and  $O(N^6)$ , respectively. Table 1 provides the memory and time complexities of the LU factorization and of the solution steps in DSM for the 3D geometry. According to the limited memory and time scalability of direct solvers, optimal com-



**Figure 3.** Domain decomposition of a 2D, second-order, finite-difference grid for the Schur complement approach. Four subdomains are considered. Black and gray circles indicate interior and interface unknowns, respectively.

puter platforms for DSM are typically composed of nodes with a large amount of shared memory.

An application of 3D acoustic frequency-domain FWI based on DSM performed on a limited number of processors (60 processors, with 4Gb of memory each) was reported by Ben-Hadj-Ali et al. (2008). For low frequencies, this approach appears quite appealing, in spite of the significant memory request. To date, 3D DSM is restricted to acoustic formulations.

#### Frequency-domain modeling based on an iterative solver

Another approach for frequency-domain modeling is based on an iterative solver, ISM (Ryanti et al., 2006; Plessix, 2007; Erlangga and Herrmann, 2008). The main advantage here with respect to DSM is the small memory requirement, typically  $O(N^3)$  for 3D. The drawback is that the impedance matrix is indefinite (the real eigenvalues of  $B$  change sign), and therefore ill-conditioned. Designing an efficient preconditioner for Equation 2 is an active field of research. A multi-level Krylov method has been recently developed by Erlangga and Herrmann (2008) which leads to a number of iterations that is almost independent of the problem size or frequency. In this ideal scenario, the time complexity of the iterative solver would be  $O(N^3)$  for 3D problems, while previous versions of the method revealed a time complexity of  $O(N^4)$  (il-

lustrating a linear increase of iterations with  $N$ ). For a multisource problem and for coarse-grain parallelization over sources, the time complexity of an iterative solver would be  $O(N^4)$ . Theoretically, the iterative solver should therefore provide the most efficient numerical scheme. Moreover, attenuation effects can be easily introduced. One can, however, question the robustness of the iterative solver in the case of very heterogeneous media, as well as the feasibility of the evolution towards elastic and anisotropic modeling. Since we have no practical experience on iterative solvers, we will not address this approach further.

**Frequency-domain modeling based on a hybrid solver**

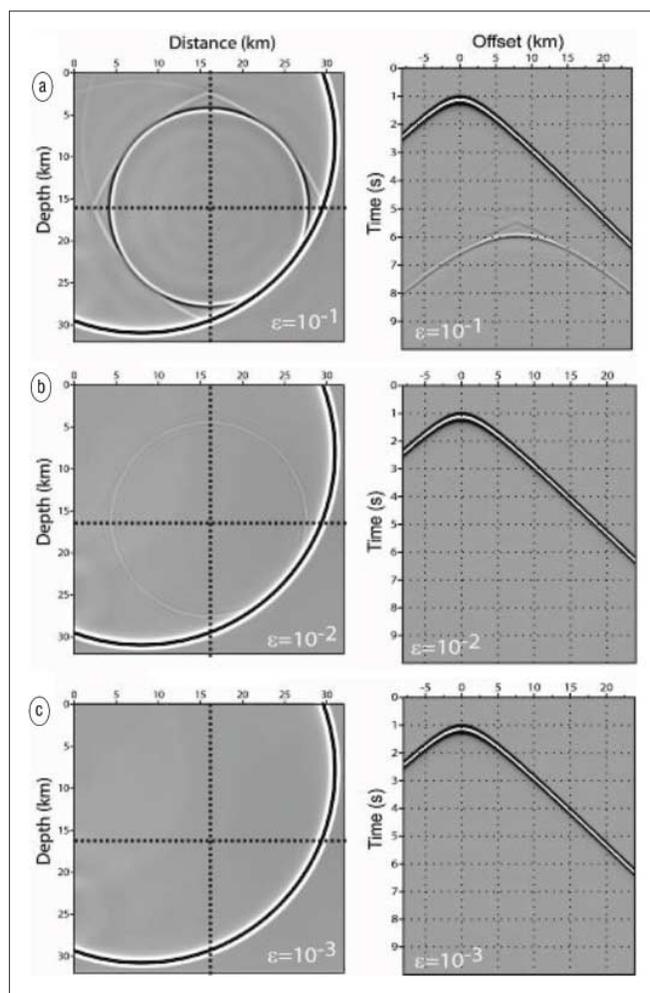
A hybrid frequency-domain modeling approach, referred to as HSM, may provide a good compromise between DSM and ISM in terms of memory requirement and efficiency of multi-rhs simulation. The HSM is based on a domain decomposition method which makes use of a hybrid direct/iterative solver (Saad, 2003).

The governing idea is to split the unknowns in two subsets: the unknowns associated with subdomain interfaces and the unknowns associated with subdomain interiors (Figure 3). One processor is assigned to each subdomain. First, the direct solver is used on each processor to perform sequential LU factorization of local matrices assembled on each subdomain. Second, an iterative solver such as GMRES is used to solve a reduced system, the so-called Schur complement system, which is better conditioned than the full system tackled by ISM. Solutions are the interface unknowns, and the rhs were inferred from the previous factorization steps. Once the interface unknowns have been determined, the interior unknowns can be efficiently computed by forward/backward substitutions on each processor. Table 1 provides the memory and time complexities of the HSM. The memory requirement of the HSM decreases with  $k$ , the number of subdomains in one direction, highlighting the main advantage of the HSM with respect to the DSM. On the other hand, the number of iterations of GMRES increases with  $k$ , and that illustrates the sensitivity of the convergence of GMRES to the domain decomposition. This suggests that the domain decomposition must balance memory saving and computational efficiency in the HSM, and an optimal trade-off should be found through trial and error which is often related to available computer hardware.

The stopping criterion of the GMRES iteration is represented by the parameter  $\epsilon$  (norm of the residual over norm of the rhs). Figure 4 illustrates the sensitivity of the simulation accuracy to  $\epsilon$  for a 2D homogeneous medium subdivided into

Complexity	DSM	HSM	TDM
Memory	$O(N^4) + O(N^3 N_{rhs})$	$O(N^4 / k) + O(N^3 N_{rhs})$	$O(N^3 N_{rhs})$
Time	$\frac{O(N^6)}{S_{LU}} + \frac{O(N^4 N_{rhs} N_{DSM} / N_p)}{S_S}$	$O(N^4 N_{GMRES} N_{rhs} / k / N_p)$	$O(N^3 N_i N_{rhs} / N_p)$
Time ( $N_{rhs} = N_p$ )	$\frac{O(N^6)}{S_{LU}} + \frac{O(N^3 N_{DSM})}{S_S}$	$O(N^4 N_{GMRES} / k)$	$O(N^3 N_i)$

**Table 1.** Memory and time complexities of DSM, HSM, and TDM for multisource problems and two levels of parallelism.  $N$  = dimension of a 3D cubic grid.  $N_p$  = total number of processors.  $N_{rhs}$  = total number of sources.  $N_{DSM}$  = number of MPI processes assigned to LU decomposition in DSM.  $S_{LU}$  = speedup of LU factorization.  $S_S$  = speedup of substitution step in DSM.  $N_{HSM}$  = number of MPI processes assigned to HSM domain decomposition.  $k$  = number of subdomains in one direction ( $N_{HSM} = k^3$ ). NGMRES = number of GMRES iterations. The last row corresponds to  $N_p = N_{rhs}$ . Since  $N_{GMRES} / k$  is roughly constant, and  $Nt = aN$  where  $a$  is a constant factor, the time complexities of HSM and TDM are equal in this specific configuration.



**Figure 4.** Solution accuracy as a function of the stopping criterion of iterations in the HSM approach. The medium is homogeneous. Four subdomains, delineated by dashed lines, are considered. (a)  $\epsilon = 10^{-1}$ , (b)  $\epsilon = 10^{-2}$ , and (c)  $\epsilon = 10^{-3}$ . (See text for details.)

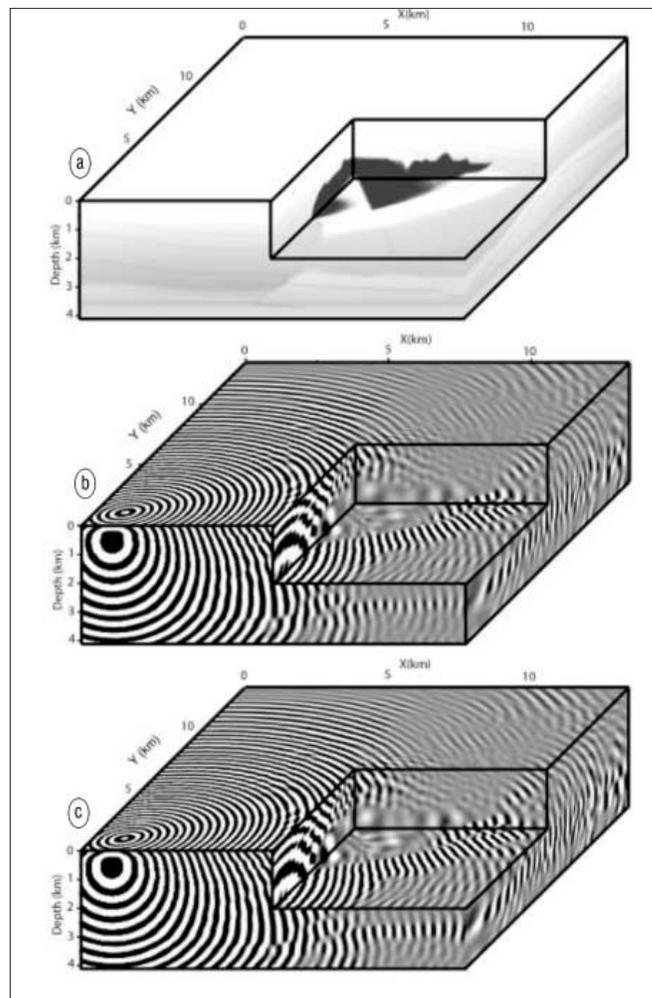
four subdomains. Diffractions and refractions along the subdomain interfaces appear when we select a stopping criterion ( $\epsilon = 10^{-1}$ ) that is too high. A value of  $10^{-3}$  provides artifact-free simulation, while small artifacts are visible for  $\epsilon = 10^{-2}$ . Our

practical experience of FWI suggests that this accuracy is sufficient for imaging applications.

The 7-Hz monochromatic wavefields computed in the 3D SEG/EAGE salt model (Figure 5a) with the HSM and TDM approaches are compared in Figures 5b and 5c. Reasonable qualitative agreement between the two wavefields is seen. For the 7-Hz simulation, the salt model is resampled with a grid interval of 37.5 m. The grid dimensions are  $120 \times 372 \times 372$  which leads to 16.6 million unknowns. The number of processors and subdomains is 576, corresponding to  $4 \times 12 \times 12$  subdomains along the three Cartesian directions  $x$ ,  $y$ , and  $z$ . The dimensions of the subdomains are, on average,  $31 \times 31 \times 31$ , which allows us to perform sequential LU factorization on each processor with 2 Gb of memory. The simulation is performed on the IBM Blue Gene of the IDRIS computer center. The iterations are 113 and 283 for accuracies of  $\epsilon=10^{-2}$  and  $10^{-3}$ , respectively. The total elapsed time is 682 s, and the time dedicated to the GMRES iterative solver is 165 s for  $\epsilon=10^{-3}$ .

### Numerical example

We will assess the TDM, DSM, and HSM approaches with a realistic numerical example. We perform a 7-Hz monofrequency simulation in the SEG/EAGE overthrust model, sampled with a grid interval of 75 m, which roughly corresponds to four grid points per minimum wavelength. The grid dimensions are  $277 \times 277 \times 73$ , leading to 5.6 million unknowns. For the comparison of the three approaches here, we consider two levels of parallelism implemented on  $N_p$  processors: one coarse-grain level over groups of shots, and a finer one related to domain decomposition. For DSM, we assume that one LU factorization performed on  $N_{DSM}$  processors provides the best compromise between CPU efficiency and memory saving. Following the same reasoning for HSM,  $N_{HSM}$  denotes the number of processors dedicated to one domain decomposition. For TDM, we consider only parallelism over shots with an efficiency of 1.0. The theoretical elapsed times for the three approaches are provided in Table 2. The TDM and the HSM simulations are performed on the IBM Power 6 computer of the IDRIS computer center (<http://www.idris.fr/>). The DSM simulation is performed on the cluster of the SIGAMM computer center (<https://sigamm.oca.eu/>). We used 192 MPI processes for the HSM simulation, with 2 Gb of memory per processor. We used 48 MPI processes for the DSM, with 8 Gb of memory and four threads per MPI process. One sequential run of the TDM provides the correct solution by discrete Fourier transform at the selected frequency. The main features of the three simulations are outlined in Table 3. The elapsed time as a function of the total number of shots is shown in Figure 6. The total number of processors decreases the slope of the three curves by the same amount. For our implementation, we observe slower slopes for the TDM than for the HSM, which makes the former approach superior, although any improvements in the iterative preconditioning can question this conclusion. DSM performs better for a large number of sources (above 500) due to the efficiency of the solution step.



**Figure 5.** Simulation in the SEG/EAGE salt model using HSM and TDM. The frequency was 7 Hz. (a) Salt model resampled with a grid interval of 37.5 m, and smoothed accordingly. (b) HSM solution. (c) TDM solution after discrete Fourier transform.

### Discussion and conclusions

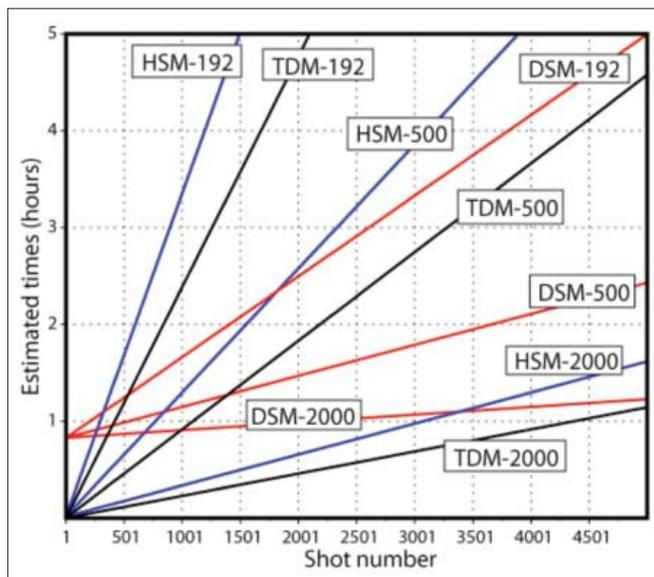
Two-dimensional, frequency-domain FWI is efficiently performed using DSM, even when considering elastic, anisotropic and other extensions. Three-dimensional, frequency-domain FWI requires more subtle distinction. From the previous numerical experiment, we may conclude that 3D DSM remains the most efficient approach to tackle small models (i.e., involving less than 10 million unknowns) if a large number of sources needs to be considered and if a suitable distributed-memory platform composed of a limited number of processors with a significant amount of shared memory per processor is available. Due to the low memory requirement of TDM, coarse-grain parallelization over sources can be viewed on distributed-memory platforms with a large number of processors, typically of the order of the number of sources. If the number of processors significantly exceeds the number of sources, an additional level of parallelism can be easily implemented through classic domain decomposition. TDM also allows us to extract an arbitrary number of frequency components by discrete Fourier transform, with a

	DSM	HSM	TDM
Time	$T_{DSM} = T_{LU} + N_{rhs} \times T_{rhs} \times N_{DSM} / N_p$	$T_{HSM} = T_{LU+M} + N_{rhs} \times T_{GMRES} \times N_{HSM} / N_p$	$T_{TDM} = T_{seq} \times N_{rhs} / N_p$

**Table 2.** Elapsed time for DSM, HSM, and TDM (see text for details).  $N_p$  = total number of processors.  $N_{rhs}$  = total number of sources. For DSM,  $N_{DSM}$  = number of processors dedicated to one LU factorization,  $T_{LU}$  = elapsed time for LU factorization, and  $T_{rhs}$  = elapsed time for one rhs solution. For HSM,  $N_{HSM}$  = number of processors for one domain decomposition,  $T_{LU+M}$  = elapsed time to perform rhs-independent tasks (local factorizations + preconditioner building), and  $T_{GMRES}$  = elapsed time for GMRES. For TDM,  $T_{seq}$  = elapsed time for 1-rhs sequential simulation.

Method	$T_{LU}$ (s)	$T_{rhs}$ (s)
DSM	2980	3
HSM	55	12
TDM	-	1647

**Table 3.** Elapsed times for the SEG/EAGE overthrust model simulations: each method might perform better than the two others depending on computer resources and number of sources/receivers of the seismic experiment. (See text for details.)



**Figure 6.** Elapsed time as a function of total number of sources for DSM (red), HSM (blue), and TDM (black). The curves are computed for three numbers of available processors: 192, 500, and 2000 (indicated in boxes). Groups of 192 processors were assigned to parallel LU factorization for DSM and domain decomposition for HSM. Initial LU factorization for HSM has a negligible value. For TDM, parallelization is performed over shots.

marginal extra computational effort if simultaneous inversion of multiple frequencies needs to be considered in frequency-domain FWI. TDM should also be the most robust with respect to the complexity of the medium, and the most suitable to evolve towards elastic and anisotropic wave modeling. As noted by Pratt and Sirgue during the 2008 EAGE workshop on FWI, a nice feature of TDM is the flexible selection of specific arrivals by time windowing during modeling, while leaving the inversion in the frequency domain. The frequency domain remains the domain of choice for implementation of attenuation effects of arbitrary complexities. The HSM

approach reveals itself so far as less efficient, although it decreases the memory requirement of DSM for large-scale problems. However, conclusions that can be driven for iterative-solver-based approaches are

highly dependent on the preconditioner efficiency, which is still a field of active research. Recent developments of iterative solvers have revealed very promising results for the 3D Helmholtz equation, with an optimal time complexity. Extension to the 3D elastic wave equation remains an ongoing problem.

**Suggested reading.** “Frequency-domain full-waveform modeling using a hybrid direct-iterative solver based on a parallel domain decomposition method: A tool for 3D full-waveform inversion?” by Sourbier et al. (SEG 2008 *Expanded Abstracts*). “Frequency response modeling of seismic waves using finite-difference time domain with phase sensitive detection (TD-PSD)” by Nihei and Li (*Geophysical Journal International*, 2007). “A Helmholtz iterative solver for 3D seismic-imaging problems” by Plessix (GEOPHYSICS, 2007). “A parallel multigrid-based preconditioner for the 3D heterogeneous high-frequency Helmholtz equation” by Riyanti et al. (GEOPHYSICS, 2007). “Gauss-Newton and full Newton methods in frequency-space seismic waveform inversion” by Pratt et al. (*Geophysical Journal International*, 1998). “Inversion of seismic reflection data in the acoustic approximation” by Tarantola (GEOPHYSICS, 1984). *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation* by Tarantola (Elsevier, 1987). *Advances in Wave Propagation in Heterogeneous Earth* edited by Wu and Maupin (Academic Press, 2006). “Efficient waveform inversion and imaging: a strategy for selecting temporal frequencies” by Sirgue and Pratt (GEOPHYSICS, 2004). “3D frequency-domain waveform inversion using time-domain finite-difference methods” by Sirgue et al. (EAGE 2008 *Extended Abstracts*). “Iterative methods for sparse linear systems” by Saad (*SIAM*, 2003). “An iterative multilevel method for computing wavefields in frequency-domain seismic inversion” by Erlangga and Herrmann (SEG 2008 *Expanded Abstracts*). “3D frequency-domain full waveform tomography based on a domain decomposition forward problem” by Ben-Hadj-Ali et al. (SEG 2008 *Expanded Abstracts*). **TLE**

*Acknowledgments:* We thank BP, CGGVeritas, Exxon-Mobil, Shell, and Total for supporting our Seiscope consortium at <http://seiscope.unice.fr>, and the Agence Nationale de la Recherche (ANR) under project ANR-05-NT05-2-42427. We thank P. Amestoy and J.-Y. L'Excellent for providing the MUMPS direct solver, available at <http://graal.ens-lyon.fr/MUMPS/index.html>. The simulations were performed through the facilities provided by the computer centers IDRIS (<http://www.idris.fr>) and SIGAMM (<https://sigamm.oca.eu>).

Corresponding author: Jean.Virieux@obs.ujf-grenoble.fr