IPTC 16881

Performances of 3D Frequency-Domain Full-Waveform Inversion Based on Frequency-Domain Direct-Solver and Time-Domain Modeling: Application to 3D OBC Data from the Valhall Field

Romain Brossier(1), Vincent Etienne(2), Guanghui Hu (1-2), Stéphane Operto(2) and Jean Virieux(1)
(1) ISTerre-Université Joseph Fourier Grenoble & CNRS
(2) Géoazur-Université de Nice Sophia Antipolis & CNRS

This paper was prepared for presentation at the International Petroleum Technology Conference held in Beijing, China, 26–28 March 2013.

Abstract

This study addresses the performances of different approaches to tackle 3D frequency-domain Full Waveform Inversion (FWI). FWI is becoming an appealing method to build high resolution velocity models from wide-azimuth seismic data. Three-dimensional frequency-domain solutions of the wave-equation for frequency-domain FWI can be computed with several methods. In the frequency domain, direct or iterative solvers can be used to solve the linear system resulting from the discretization of the wave equation. Alternatively, frequency responses can be computed by time-domain modeling coupled with discrete Fourier transform or phase-sensitive detection method. For a large number of seismic sources, the computational time of all of the methods scale along similar lines to the size of the problem, although the performances of the direct-solver approach is hampered by the limited scalability and the memory demand of the lower-upper decomposition of the forward-problem operator.

We implement the direct-solver and the time-domain modeling approaches in 3D acoustic frequency-domain FWI, which is applied to real 3D wide-azimuth OBC data from the Valhall field, North Sea, at low frequencies (< 5 Hz). We show, for this case study, that even if less flexible for high-performance computing, the direct-solver approach is one order of magnitude more efficient in computing time/resources that the time-domain approach, if the signal-to-noise ratio in the data is sufficiently-high to limit the inversion to few frequencies. We finally discuss the performances, pros and cons of both methods, which depend on the wave physics (acoustic versus visco-acoustic, anisotropy, extension toward elastic or visco-elastic), the acquisition geometry (streamer-like versus fixed-spread acquisitions) and the computing architectures (small versus large amount of memory per computing node).

Introduction

Since the last decade, acoustic full waveform inversion (FWI) have proved to be a powerful method to develop accurate velocity models, which can be used as background models for reverse-time migration (RTM). Three-dimensional applications have shown the dramatic improvement in resolution of the velocity models in comparison to the ones built by traveltime tomography (Plessix, 2009, Sirgue et al., 2010, Vigh et al. 2010). With the continuously increasing size of 3D wide-azimuth data-sets, FWI appears to be a challenging computational problem.

While 3D time-domain FWI requires a careful memory and i/o management of both the increasing size of seismic traces and the incident and adjoint wavefields (such as check-pointing method, Symes, 2007), the frequency-domain (FD) FWI relies on the hierarchical inversion of few discrete frequencies and, hence allows for the management of smaller data volumes. The choice of the modeling tool for FD FWI can be driven by the size of the computational mesh, the acquisition design (fixed-spread versus streamer, number of sources) and the architecture of the computational platform (Virieux et al., 2009). FD wave-equation solutions can be directly obtained from the resolution of the linear system resulting from the discretized Helmholtz equation. The linear system can be solved either by iterative solver with appropriate preconditioning techniques (Plessix, 2009) or with direct solvers, requiring a significant amount of core-memory (BenHadjAli et al., 2008). The FD solutions can also be obtained from a time-domain explicit scheme coupled with an extraction of the frequency solution by Fourier transform (Nihei and Li, 2007, Sirgue et al., 2007).

In this study, we first review the time and memory complexities of the three modeling methods for FD FWI. We then present two 3D FWI algorithms based on FD direct solver and time-domain solver plus Fourier transform, respectively. The two algorithms are applied at low frequencies (< 5 Hz) to 3D Ocean Bottom Cable (OBC) data acquired on the Valhall field in Norway. We finally compare the performances of the two algorithms and discuss the pros and cons of each method.

Computing 3D FD wave-equation solutions

Considering FWI in the frequency-domain, the most natural approach to compute the frequency seismic response consists in solving the generalized Helmholtz equation (Marfurt, 1984). This equation can be discretized with various numerical methods, but all of them rely on the resolution of a linear system per seismic source and per frequency. Spatially compact numerical schemes, such as finite-difference stencils, lead to sparse matrices. The frequency approach allows an easy implementation of attenuation without extra-cost through the use of complex velocities (Toksöz and Johnston, 1981).

The sparse linear system can be solved with iterative solvers (Erlangga and Nabben, 2008, Plessix, 2009). The main advantage of iterative solvers is the small memory requirement that allows for an efficient distribution of sources over processors. The main drawback is that convergence is not guaranteed as the matrix is indefinite and ill-conditioned. The number of floating-point operations generally scales to N^4 for 3D N^3 computational grids. Recent preconditioners should however decrease this time complexity to $O(N^3)$ (Erlangga and Nabben, 2008).

The matrix of the linear system can also be factorized once with a direct solver before backward and forward substitutions for each seismic source (Marfurt, 1984, Operto et al., 2007, BenHadjAli et al., 2008). The main advantage of the direct solver is that, once the matrix is factorized, the substitution phase is quite efficient for multiple sources. The main drawbacks of the direct solvers are the limited scalability for high-performance computing (HPC) and the large memory demand of the LU factorization.

An alternative to these approaches is to solve the wave-equation in the time domain. Time-domain modeling must be coupled with a Fourier transform of the time-domain wavefield to extract the frequency response, with two possible implementations based on phase-sensitive detection (Nihei and Li, 2007) or discrete Fourier transform (DFT) computed on the fly (Sirgue et al., 2007). The main advantage of this approach is the high scalability of time-domain solvers for HPC, often based on the same tools than RTM, and the ability to extract several frequency components for FWI during a single time-domain simulation. One of the main limitations is the implementation of the attenuation through the use of memory variables, which increase the required memory and computation time.

Time complexities of the three approaches for N^2 seismic sources, representative of 3D fixed-spread acquisitions, show that all of the approaches are theoretically equivalent (Table 1). However, the direct solver requires a significant extra amount of core-memory to perform the source-independent LU factorization.

From this theoretical analysis, it is difficult to anticipate the most efficient strategy for realistic FWI applications

Table 1 : Time & memory theoretical complexities of the three modeling approaches, describing how the time & memory cost scale to the problem size (without considering parallelism)

Method	Iterative	Time+FD	Direct
Time complexity (rhs-independant tasks)	-	-	$O(N^6)$
Time complexity for 1 rhs	$O(N^4)$	$O(N^4)$	$O(N^4)$
Time complexity for N^2 rhs	$O(N^6)$	$O(N^6)$	$O(N^6)$
Memory complexity	$O(N^3)$	$O(N^3)$	$O(N^4)$

FD FWI algorithms

We describe two FD FWI algorithms, developed in the framework of the SEISCOPE consortium, which rely on the L_2 norm minimization of the misfit function and the adjoint-state method to efficiently compute the gradient of the misfit function.

GeoInv3D code

The *GeoInv3D* code is a FD FWI algorithm based on time-domain modeling coupled with a discrete Fourier transform of the wavefields (Etienne et al., 2010, Castellanos et al., 2011, Etienne et al., 2012). In this study, a $O(\Delta x^4)$ staggered-grid finite-difference scheme is used to discretize the first-order velocity-pressure wave equation (Levander, 1988). This scheme allows one to discretize the computational domain with 4 to 5 grid points per minimum wavelength. The time derivative is computed through a $O(\Delta t^2)$ leap-frog integration scheme. Convolutional perfectly-matched layers (CPML) are used for absorbing boundary conditions (Komatitsch and Martin, 2007). Frequency-domain wavefields required by the FD inversion are computed on the fly through a DFT embedded in the time-marching loop (Sirgue et al., 2007).

The optimization is based on a conjugate-gradient method, preconditioned by the diagonal elements of the pseudo-Hessian (Shin et al., 2001). The line-search is based on a local parabola interpolation of the misfit function.

Two nested levels of parallelism by shot distribution over computing cores and domain decomposition of the computational mesh can be combined by using two message-passing-interface (MPI) communicators, which is useful to optimize the perfomances of the code according to the number of cores, the available memory per node, the size of the subsurface grid and the number of shots.

More detailed of this algorithm are given in Castellanos et al. (2011) and Etienne et al. (2012).

TOY3DAC code

The *TOY3DAC* code is a FD FWI algorithm based on a FD direct solver. An optimized anti-lumped-mass mixed grid stencil provides the accuracy of a $O(\Delta x^4)$ scheme, although it relies on a linear combination of compact $O(\Delta x^2)$ stencils (Operto et al., 2007, Brossier et al., 2010). This compact stencil reduces the fill-in of the matrix during factorization and, hence the memory requirement, while using discretization rule of 4 grid points per minimum wavelength. The linear system is solved with the massively-parallel MUMPS solver (Amestoy et al., 2006), based on a multi-frontal approach. An in-house nested-dissection algorithm takes advantage of the regular pattern of the finite-difference matrix during the fill-reducing matrix ordering before factorization, and outperforms generic reordering tools such as METIS or SCOTCH.

The gradient direction is preconditioned with smoothing constraints implemented with anisotropic Laplace functions to avoid high frequency artifacts. The optimization is based on a quasi-Newton L-BFGS-B scheme (Byrd et al., 1995). Both optimization and line search are performed with the corrected version of algorithm 778 (Zhu et al., 1997). Frequency-domain wavefields are computed in parallel by MUMPS using MPI. Multi-threaded BLAS (Basic Linear Algebra Subroutines) speeds up the LU factorization and the multi-source solution steps when a limited number of MPI processes is used to account for the limited scalability of the LU factorization. The inversion (gradient building, optimization) is centralized on the host MPI process, and is parallelized with shared-memory OpenMP threads.

Application to the 3D OBC Valhall data

The Valhall field (North Sea, Norway) is a shallow-water environment (water depth of 70 m). The target is composed of tertiary sediments with gas trapped layers above a Jurassic (fractured) anticline creating the top of the oil reservoir. The fixed-spread permanent 4-components OBC acquisition is composed of 2302 sensors positions each 50 m along 12 parallel cables each composed of 3 geophones and 1 hydrophone (Sirgue et al. 2009). In this study, only the hydrophone component is used. The 49950 shots positions generated with a grid interval of 50 m in *x* and *y* are used for FWI.

The spatial reciprocity of Green functions is used to process shots as receivers and vice-versa, and hence, reduce the number of seismic modeling during inversion. In both codes, free-surface effects are modeled and hence multiples have not been eliminated from the data before FWI. Data preprocessing consists in band-pass filtering and resampling before Fourier transform for FWI. A constant density of 1000 $kg.m^{-3}$ is considered in both codes. A constant attenuation with Q_P=150 defined by the AVO of the first arrival (Prieux et al., 2011) is considered in *TOY3DAC*, unlike in *GeoInv3D*.

FWI is performed at low frequencies, between 3.5 and 5 Hz. *GeoInv3D* takes advantage of the multi-frequency extraction to invert a first group of 5 frequencies within the [3.5-4] Hz band before inversion of a second group of 9 frequencies within the [4-5] Hz band. *TOY3DAC* performs sequential inversions of the 3.8 Hz and the 4.5 Hz frequencies. The source signature is estimated for each receiver gather during inversion through the resolution of a linear inverse problem in both cases (Pratt, 1999).

The initial model (courtesy of BP) was built by VTI reflection traveltime tomography, as presence of shale leads to seismic anisotropy up to 16% (Prieux et al., 2011). For isotropic FWI, the anisotropic models have been converted into NMO velocity and smoothed for FWI application. This starting model prevents, however, to fit first-arrival traveltimes. A grid interval of 70 m leads to a finite-difference grid of 77 x 130 x 230 points for both codes.

Vertical and inline sections of the final FWI models obtained with *GeoInv3D* and *TOY3DAC* are shown in Figures 1 and 2. Despite the low frequency components used in FWI, an impressive resolution of the FWI models is obtained compared to the one of the starting model. Both codes show similar structures such as the complex network of channels in the shallow sediments below the OBC arrays, which can be compared with those shown by Sirgue et al. (2009, 2010). The imaging of the channels is more significantly hampered by the acquisition footprint in the *GeoInv3D* results because the full offset range has been used in the inversion, while the offsets smaller than 500 m were muted to perform FWI with *TOY3DAC*. A low-velocity anomaly at 1.44 km in depth is associated with gas accumulation and fracturing in the sediment layers. Deep reflector below the anticline level can also be seen on both FWI models.

The misfit function reduction achieved by both methods is comparable: *GeoInv3D* decreases the misfit function by 18% and 60% for the first and second frequency bands, respectively, while *TOY3DAC* reduces the misfit by 10% and 45% for frequencies 3.8 and 4.5Hz, respectively. The difference between these misfit function reductions can be related to the different frequencies processed by the two codes, the noise level being much higher for frequencies smaller than 4 Hz.

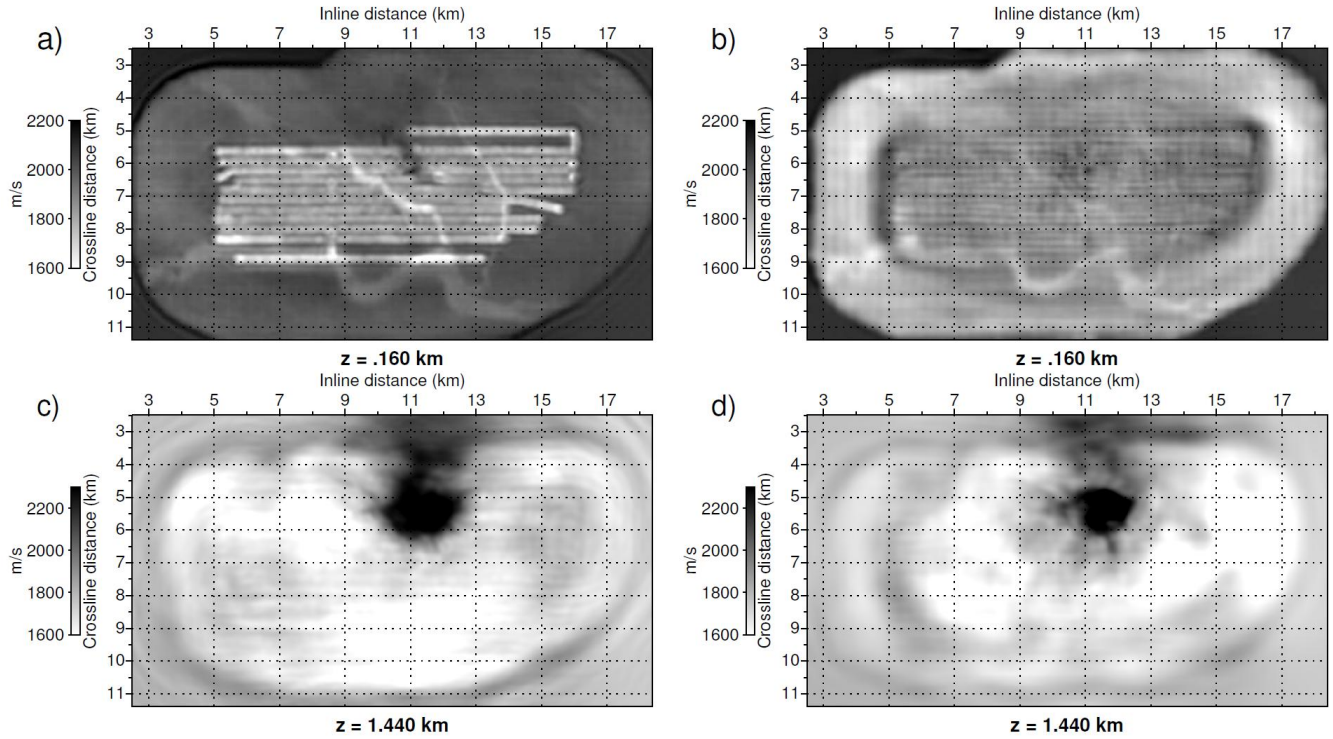

Figure 1 : Horizontal sections across the channels (160 m depth) and the top of the gas cloud (1440m depth). (a)-(c) *GeoInv3D*, (b)-(d) *TOY3DAC*.

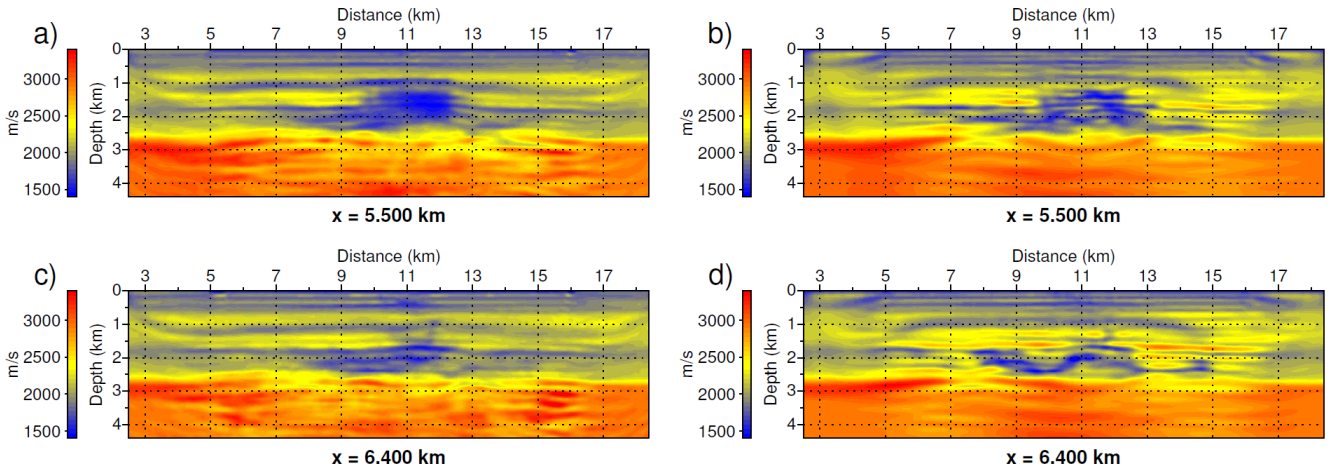

Figure 2 : Inline section across the gas cloud (x=5500 m) and the south end of the gas cloud (x=6400 m). (a)-(c) GeoInv3D. (b)-(d) TOY3DAC. Note the presence of a gas-filled fracture between 1.5 and 1.8 km in depth at around 12 km in distance (see also Sirgue et al. (2010) for comparison).

Modeling of first-arrival traveltimes in the initial FWI model (not shown here) shows a significant delay of the computed traveltimes compared to the recorded ones. Moreover, comparison between a sonic log located at (x=6.8 km, y=9.5 km) with the corresponding FWI logs (Figure 3) shows the possible imprint of anisotropy on isotropic FWI (Prieux et al., 2011): FWI reconstructs higher wave-speeds than the vertical velocity recorded by the sonic log between 1.25 and 1.75 km in depth, where anisotropy reaches a value of 16 %. The reconstructed wave-speeds are close to the horizontal velocities, which are required to match the early arrivals at long offsets, which travel almost horizontally in the first 1.5 km of the subsurface.

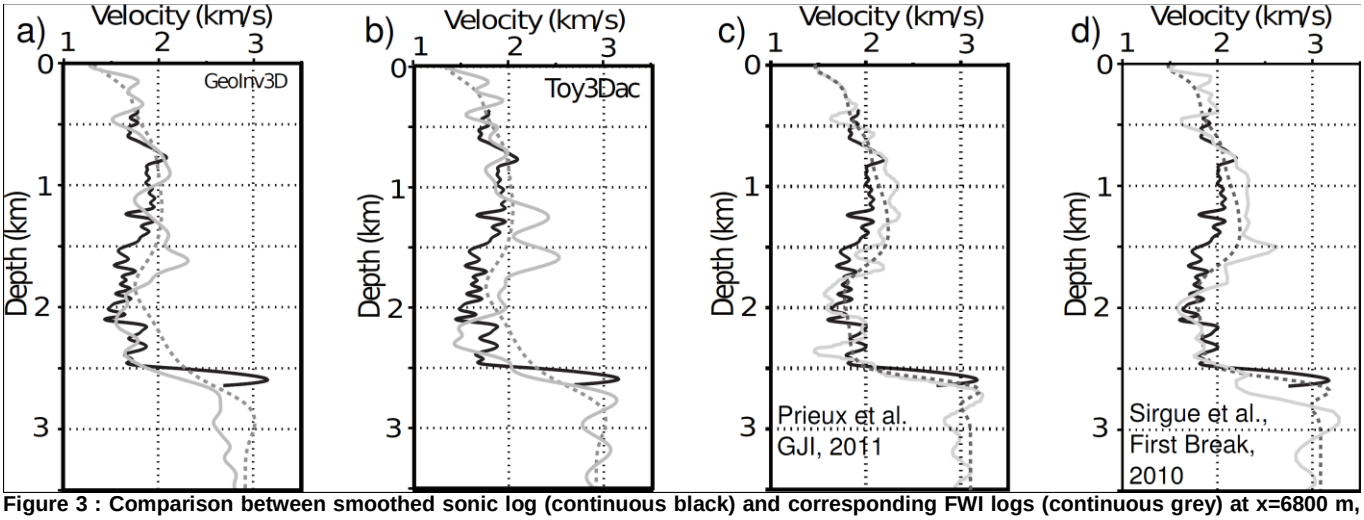

Figure 3 : Comparison between smoothed sonic log (continuous black) and corresponding FWI logs (continuous grey) at x=6800 m, y=9500 m. (a) *GeoInv3D*, (b) *TOY3DAC*, (c) 3D FWI model (Sirgue et al., 2010), (d) 2D FWI model (Prieux et al., 2011). Starting models used by the FWI applications are plotted with a discontinuous grey line. Note that GeoInv3D and TOY3DAC used an initial model, which is smoother than the one used by Prieux et al. (2011).

Performances of GeoInv3D and TOY3DAC

Direct solvers and time-marching algorithms require significant different computing architectures to perform realistic 3D applications with a large number of seismic sources.

Direct solvers do not scale well on a large number of computing cores and hence require a large amount of core memory on a limited number of MPI processes to perform the factorization of the impedance matrix. This *TOY3DAC* application have been run on a cluster equipped with Intel Westmere bi-processor nodes, connected to a high speed Mellanox QDR 40 Gb/s Infiniband network. Each node is composed of 12 computing cores at 2.26 GHz, which share 72 Gb of core memory.

The time-domain approaches are highly scalable for multi-source problems and require much less memory per MPI process. This application has been run on a IBM Blue Gene P, equipped with single-processor Power PC 450 nodes, interconnected through the IBM 3D toroidal network. Each node is composed of 4 computing cores at 850 MHz with 2 Gb of shared memory. *GeoInv3D* has been launched with the two-parallelism levels with 4 sub-domains and 512 seismic sources processed in parallel.

The performances of *GeoInv3D* and *TOY3DAC* are summarized in Table 2, and show that, when using appropriate computing architecture with sufficient core memory, the direct solver approach is 18 times more efficient than the time-domain approach, in terms of theoretical floating point operations. The time-domain approach allows however to tackle multi-frequency FWI without extra computing cost, which can be helpful for sparse acquisitions or noisy data to ensure redundancy during FWI. For this application, both approaches would be roughly equivalent in floating point operations, if groups of 15 to 20 frequencies would have been processed by *TOY3DAC*.

The total amount of distributed memory used by both codes (around 500 Gb) are similar. The parallelism implemented in *GeoInv3D* implies the in-core storage of the distributed multi-frequency wavefields. The corresponding memory demand scales to N^5 (for N^2 sources in parallel) while the memory demand of the LU decomposition scales to N^4.

This application clearly shows that, if appropriate computing architectures are available and if data quality and coverage are sufficient, the direct-solver modeling approach remains the method of choice to tackle 3D FD FWI of fixed-spread data at low frequency. This method is however less suitable for streamer acquisition, where the computational domain may change with the shot positions: in this case, the LU factorization must be performed for each domain for a limited number of sources, a non favorable scenario. Attenuation can however been taken into account without extra cost.

Time-domain approaches are more flexible regarding the computer architectures and the acquisition design thanks to multiple levels of parallelism by shot distribution and domain decomposition. Data preconditioning by time windowing can be also more easily implemented in time-domain modeling. Extension toward elastic FWI seems also much more tractable than with direct solver. Moreover, multi-frequency extraction for FWI can be used without extra cost. With our implementation, the method remains however one order of magnitude less efficient that the direct solver on equivalent grid size.

Table 2 : Performances of TOY3DAC and GeoInv3D on the 3D Valhall application. Times have been averaged over the non-linear FWI iterations of the second frequency band for GeoInv3D and the second frequency for TOY3DAC

Code	*TOY3DAC*	*GeoInv3D*
number of frequencies	1	9
cluster architecture	Intel cluster	IBM Blue Gene P
number of computing cores used	72	2048
elapsed time per FWI iteration (h)	1.33	2.30
total elapsed time per FWI iteration (h) (number of core x elapsed time)	96	4710
theoretical pic floating point operations per FWI iteration	3.1E15	5.8E16

Conclusions

We present two algorithms to perform 3D FD FWI, based either on frequency-domain direct solver or time-domain modeling. We apply these algorithms to real wide-aperture OBC data and discuss the performances, pros and cons of both methods, which are case and computer architecture dependent.

Acknowledgments

This study was funded by the SEISCOPE consortium (http://seiscope.oca.eu), sponsored by BP, CGG-VERITAS, ENI, EXXON-MOBIL, PETROBRAS, SAUDI ARAMCO, SHELL, STATOIL and TOTAL.
This study was granted access to the HPC facilities of CIMENT (Université Joseph Fourier Grenoble), SIGAMM (Observatoire de la Côte d'Azur) and [CINES/IDRIS] under project [gao2280] of GENCI (Grand Equipement National de Calcul Intensif). We thank BP Norge AS and Hess Norge AS for providing us the 3D Valhall data set.
We thank V. Prieux (GeoAzur), O.I. Barkved (BP), J. Kommedal (BP), R.-E. Plessix (Shell), and L. Sirgue (Total) for fruitful discussions.

References

Amestoy, P.R., A. Guermouche, J.Y. L'Excellent, and S. Pralet, 2006, Hybrid scheduling for the parallel solution of linear systems: Parallel Computing, **32**, 136-156.

Ben-Hadj-Ali, H., S. Operto, and J. Virieux, 2008, Velocity model building by 3D frequency-domain, full-waveform inversion of wide-aperture seismic data: Geophysics, **73**, WE101-WE117.

Brossier, R., V. Etienne, S. Operto, and J. Virieux, 2010, Frequency-domain numerical modelling of visco-acoustic waves based on finite-difference and finite-element discontinuous galerkin methods, *ed.* Dissanayake, D.W., *in.*, Acoustic Waves, 125-158, SCIYO.

Byrd, R., P. Lu, and J. Nocedal, 1995, A limited memory algorithm for bound constrained optimization: SIAM Journal on Scientific and Statistical Computing, **16**, 1190-1208.

Castellanos, C., V. Etienne, G. Hu, S. Operto, R. Brossier, and J. Virieux, 2011, Algorithmic and methodological developments towards full waveform inversion in 3D elastic media: SEG Technical Program Expanded Abstracts, **30**, 2793-2798.

Erlangga, Y.A. and R. Nabben, 2008, Multilevel projection-based nested Krylov iteration for boundary value problems: SIAM Journal on Scientific Computing, **30**, 1572-1595.

Etienne, V., S. Operto, J. Virieux, and Y. Jia, 2010, Computational issues and strategies related to full waveform inversion in 3D elastic media: methodological developments: SEG Technical Program Expanded Abstracts, **29**, 1050-1054.

Etienne, V., G. Hu, S. Operto, J. Virieux, O.I. Barkved and J. Kommedal, 2012, Three-dimensional acoustic full waveform inversion: algorithm and application to Valhall : Expanded Abstracts, 74th Annual EAGE Conference & Exhibition, Copenhagen, P343.

Komatitsch, D. and R. Martin, 2007, An unsplit convolutional perfectly matched layer improved at grazing incidence for the seismic wave equation: Geophysics, **72**, SM155-SM167.

Levander, A.R., 1988, Fourth-order finite-difference P-SV seismograms: Geophysics, **53**, 1425-1436.

Marfurt, K., 1984, Accuracy of finite-difference and finite-element modeling of the scalar and elastic wave equations: Geophysics, **49**, 533-549.

Nihei, K.T. and X. Li, 2007, Frequency response modelling of seismic waves using finite difference time domain with phase sensitive detection (TD-PSD): Geophysical Journal International, **169**, 1069-1078.

Operto, S., J. Virieux, P. Amestoy, J.-Y. L'Excellent, L. Giraud, and H. BenHadjAli, 2007, 3D finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study: Geophysics, **72**, SM195-SM211.

Plessix, R.E., 2009, Three-dimensional frequency-domain full-waveform inversion with an iterative solver: Geophysics, **74**, WCC53-WCC61.

Pratt, R.G., 1999, Seismic waveform inversion in the frequency domain, part I : theory and verification in a physic scale model: Geophysics, **64**, 888-901.

Prieux, V., R. Brossier, Y. Gholami, S. Operto, J. Virieux, O.I. Barkved, and J.H. Kommedal, 2011, On the footprint of anisotropy on isotropic full waveform inversion: the Valhall case study: Geophysical Journal International, **187**, 1495-1515.

Shin, C., S. Jang, and D.J. Min, 2001, Improved amplitude preservation for prestack depth migration by inverse scattering theory: Geophysical Prospecting, **49**, 592-606.

Sirgue, L., O.I. Barkved, J. Dellinger, J. Etgen, U. Albertin, and J.H. Kommedal, 2010, Full waveform inversion: the next leap forward in imaging at Valhall: First Break, **28**, 65-70.

Sirgue, L., O.I. Barkved, J.P.V. Gestel, O.J. Askim, and J.H. Kommedal, 2009, 3D waveform inversion on Valhall wide-azimuth OBC: Presented at the 71th Annual International Meeting, EAGE, Expanded Abstracts, U038.

Sirgue, L., T.J. Etgen, U. Albertin, and S. Brandsberg-Dahl, 2007, System and method for 3D frequency-domain waveform inversion based on 3D time-domain forward modeling: US Patent Application Publication, US2007/0282535A1.

Symes, W.W., 2007, Reverse time migration with optimal checkpointing: Geophysics, **72**, SM213-SM221.

Toksöz, M.N. and D.H. Johnston, 1981, Geophysics reprint series, no. 2: Seismic wave attenuation: Society of exploration geophysicists.

Vigh, D., J. Kapoor, and H. Li, 2011, Full-waveform inversion application in different geological settings: SEG Technical Program Expanded Abstracts, **30**, 2374-2378.

Virieux, J., S. Operto, H. BenHadjAli, R. Brossier, V. Etienne, F. Sourbier, L. Giraud, and A. Haidar, 2009, Seismic wave modeling for seismic imaging: The Leading Edge, **28**, 538-544.

Zhu, C., R.H. Byrd, and J. Nocedal, 1997, L-BFGS-B: Algorithm 778: L-BFGS-B, Fortran routines for large scale bound constrained optimization: ACM Transactions on Mathematical Software, **23**, 550-560.