

Efficient 3D frequency-domain full-waveform inversion of ocean-bottom cable data with sparse block low-rank direct solver: a real data case study from the North Sea

P. Amestoy*, R. Brossier^{||}, A. Buttari[‡], J.-Y. L'Excellent[§], T. Mary**, L. Métivier^{††}, A. Miniussi[¶], S. Operto^{¶,†}, A. Ribodetti[¶], J. Virieux^{||}, C. Weisbecker*

*INPT(ENSEEIH)-IRIT; ‡CNRS-IRIT; §INRIA-LIP; ¶Geoazur-CNRS-UNSA-OCA; ||ISTerre-UJF; **UPS-IRIT; ††ISTerre-LJK-CNRS; †Speaker

SUMMARY

We present an application of 3D frequency-domain full waveform inversion (FWI) on ocean-bottom cable data from the North Sea. Frequency-domain seismic modeling is performed in the visco-acoustic VTI approximation with a sparse direct solver based on the multifrontal method. The computational cost of the multifrontal LU factorization is efficiently reduced with a block-low rank (BLR) approximation of the dense frontal matrices. A multiscale frequency-domain FWI is applied by successive inversions of 11 discrete frequencies in the 3.5Hz-10Hz frequency band. The velocity model built by FWI reveals short-scale features such as channels, scrapes left by drifting icebergs on the paleo-seafloor, fractures and deep reflectors below the reservoir level, although the presence of gas in the overburden. The quality of the FWI results is controlled by time-domain modeling and source wavelet estimation. Next step is the application of multi-parameter FWI with second-order optimization algorithms, which can be efficiently implemented with our frequency-domain modeling engine.

INTRODUCTION

Full waveform inversion (FWI) has known a renew of interest in the oil industry since high performance computing allows for 3D applications. However, the high computational demand still requires the designing of computationally efficient FWI algorithms. The most widespread implementation of 3D FWI is performed in the time domain (Tarantola, 1984), the main issue being to perform efficiently seismic modeling for thousands of sources. An obvious speedup strategy is to distribute the sources over the processors of parallel computers. Ideally, the source distribution requires at least as many processors as sources, and hence requires huge parallel computers. This source parallelism can be complemented by subsurface domain decomposition and multithreading for loop parallelism. Reduction of data dimensionality, either by source encoding (Krebs et al., 2009) or subsampling (Warner et al., 2013), allows to further reduce this computational burden.

Alternatively, full waveform modeling and inversion can be performed in the frequency domain (Pratt, 1999). The forward problem is a boundary value problem which requires the resolution of a large, sparse system of linear equations relating the seismic source to the monochromatic wavefield. This approach is worth of interest when wide-azimuth long-offset acquisitions allow performing FWI with a limited number of discrete frequencies (e.g. Sirgue and Pratt, 2004). Among the possible approaches to solve the linear system, Gauss elimina-

tion methods based on sparse direct solver have the advantage of being efficient to process a large number of right-hand sides (i.e., sources). The efficiency of this approach has been extensively demonstrated with 2D case studies, while the memory demand of the LU decomposition of the system matrix has been considered for a long time as a major bottle neck.

This study illustrates with a real data case study the efficiency of 3D frequency-domain FWI based on sparse direct solver to process ocean bottom seismic data data in the visco-acoustic VTI approximation with limited computational resources.

METHOD

A conventional data-domain frequency-domain FWI is performed by iterative minimization of the least-squares misfit between the recorded and modeled pressure data. A multiscale FWI proceeds successively from low frequencies to the higher ones to mitigate the nonlinearity of the inversion. Monochromatic pressure wavefields are modeled in the visco-acoustic VTI approximation using the finite-difference frequency-domain method described in Operto et al. (2014). The dedicated finite-difference stencil reconciles high-order accuracy and compact support that allow one to minimize both the dimension of the matrix and its numerical bandwidth. The linear system resulting from the discretization of the time-harmonic wave equation is solved with the massively-parallel sparse direct solver MUMPS, which is based on the multifrontal method (MUMPS-team, 2011). To further decrease the computational cost of the LU factorization, we use block-low rank (BLR) approximation of the dense frontal matrices (see Weisbecker et al. (2013), Amestoy et al. (2015b) and a companion abstract, Amestoy et al. (2015a), for more details). The gradient of the misfit function with respect to the vertical wavespeed is computed with the adjoint-state method. A preconditioned steepest-descent method provides the descent direction at each iteration, while the amount of descent is estimated by a conventional line search. The preconditioning is built with the diagonal terms of the so-called pseudo Hessian. No regularization is used in this study. The source signature estimation is alternated with the subsurface update during each nonlinear FWI iteration.

Application to OBC data from Valhall

We apply the 3D frequency-domain FWI to the hydrophone component of a 4C OBC data set collected in the Valhall oil field, a shallow-water giant field in the North Sea characterized by gas-charged sediments in the overburden, forming locally a gas cloud between 1km and 1.5km depth (Sirgue et al., 2010; Barkved et al., 2010). The subsurface is also characterized by a significant anisotropy, which reaches a maximum

value of 16%. The reservoir is located at 2.5km depth. We used the same dataset as that processed by frequency-domain FWI by Sirgue et al. (2010), whose FWI velocity model can be used as a benchmark to assess our results. The main difference between the FWI formulation of Sirgue et al. (2010) and ours lies in the isotropic time-domain versus VTI frequency-domain formulation of the forward problem. The targeted area covers a surface of 145 km^2 . A recording layout of 12 cables (2302 receivers) recorded 49,954 shots, located 5-m below the sea surface (Fig. 1). The nominal distance between cables is 300m, with the two outer cables at 600m. The inline spacing between two consecutive shots and receivers is 50m. The seis-

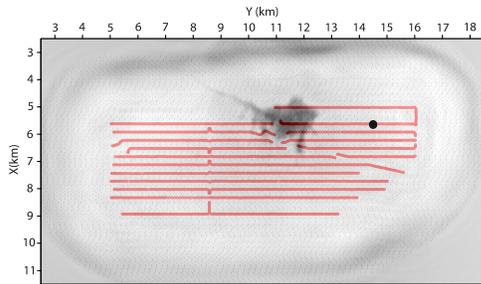


Figure 1: Acquisition layout. The red lines show the cables, the black dot pattern the shot positions. The receiver, whose records are shown in Fig. 2, is located at the position of the black circle. A horizontal slice across the gas cloud at 1km depth (Fig. 3b) shows the zone of influence of the gas.

mic wavefield is dominated by the diving waves that propagate above the gas zone ((Fig. 2a, solid black arrow), the reflection from the top of the gas layers (Fig. 2a, red arrow) and the reflection from the top of the reservoir (Fig. 2a, blue arrows).

We sequentially invert 11 frequencies in the 3.5Hz-10Hz frequency band. The grid interval is adapted to the frequency (70m, 50m, 35m for the 3.5Hz-5Hz, 5.2Hz-7Hz, 8Hz-10Hz frequency bands, respectively). All the 49,954 shots and all of the 2302 hydrophones, that are processed in a reciprocal way during seismic modeling, are used at each iteration of the inversion. Seismic modeling is performed with a free surface, meaning that free surface multiples are accounted for during FWI. We only update the vertical wavefield (V_0) during inversion, while the density ρ , the quality factor Q_P and the Thomsen's parameters δ and ϵ are kept to their initial values.

The initial models for the V_0 , δ and ϵ were built by reflection traveltome tomography (courtesy of BP). The background density model was inferred from the initial V_0 model with the Gardner law while we used a constant $Q_P = 200$ below the sea floor. In this initial model, the gas cloud appears as a smooth low-velocity blob (Fig. 3a,c), while the reservoir level is delineated by a quite sharp interface at 2.5km depth (Fig. 3c). This sharp interface allows yet for accurate modeling of the critical reflection from the reservoir (Fig. 2a, dash blue arrows). In contrast, the smooth representation of the gas cloud leads to inaccurate modeling of short-spread reflection from the top and the bottom of the gas between 1.5s and 3s two-way traveltimes (Fig. 2, solid red and blue arrows) and inaccurate position-

ing in offset of shallow critical reflections (Fig. 2, dash black arrows).

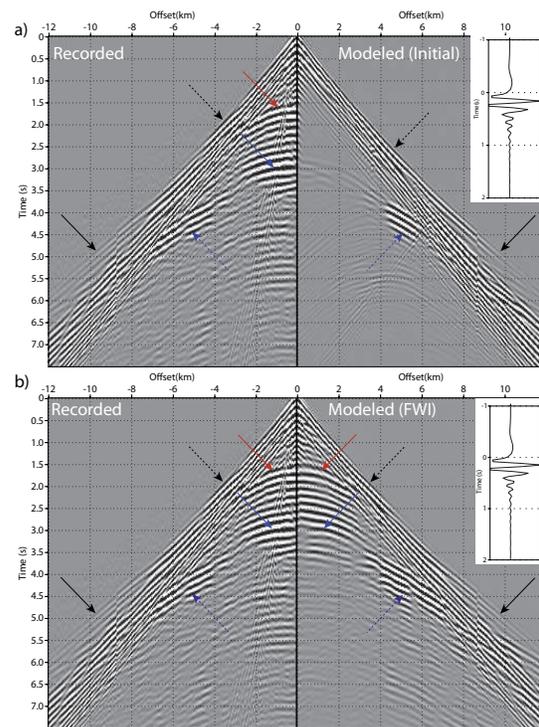


Figure 2: Receiver gather. (a) Recorded data (left) versus synthetic data (right) computed in the initial model. The inset shows the source wavelet estimated by matching the seismograms computed in the initial model with the real data. See text for the interpretation of the main arrivals. (b) Same as (a) for synthetic data computed in the final FWI model.

FWI results

An horizontal slice and a vertical section across the gas cloud of the final FWI model are shown in Fig. 3b,d. In Fig. 3d, we superimpose on the velocity section its vertical derivative to highlight the reflectors. The geometry of the gas cloud has been nicely refined. Low velocity anomalies that radiate out from the gas cloud in the horizontal slice might be interpreted as fractures (Fig. 3b). In the vertical section, FWI provides a fairly clear image of the base cretaceous reflector below the reservoir level at around 3.5km depth (Barkved et al., 2010). The relevance of the main structures can be further assessed by tying the zero-offset reflections of a receiver gather with the reflectors highlighted in the FWI model after depth-to-time conversion (Fig. 3d). In particular, this tying exercise shows that intra-gas reflectors (for example, at 2.3km depth) are clearly correlated with major reflections in the data.

All the features are synthesized in Fig. 4, which shows two expositions of the interior of the FWI V_0 volume. The first shows the horizontal slice across the gas cloud at 1km depth together with a vertical section on the periphery of the gas cloud to emphasize the 3D geometry of several sub-vertical fractures. The second shows the interior of the volume off the gas cloud. On the two panels of Figure 4, quite focused migrated-like images

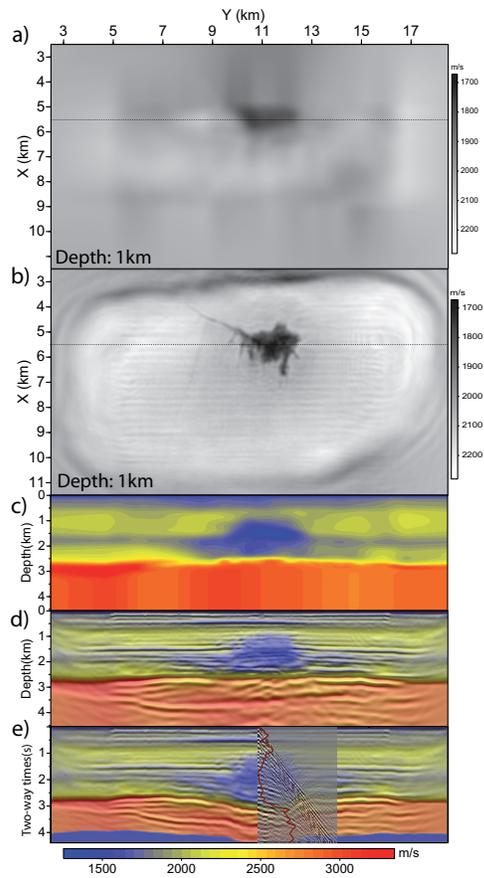


Figure 3: Initial versus final FWI models. (a-b) Horizontal slice of the initial (a) and final (b) FWI models across the gas cloud at 1 km depth. (c-d) Vertical section of the initial (a) and final (b) FWI models across the gas cloud ($X=5.5\text{km}$). In (d), the vertical derivative of the velocities are superimposed on the velocity model to highlight discontinuities. (e) Same as (d) after depth-to-time conversion. A receiver gather as well as a vertical profile of the velocity model (red line) are superimposed (see the text for interpretation).

of deep reflectors below the reservoir level show how the geometry of these reflectors evolves as we move away from the zone of influence of the gas. Other sections of the 3D FWI model, in particular in the shallow part, are shown in the companion abstract (Amestoy et al., 2015a).

Quality control

We assess the relevance of the FWI model by seismic modeling in the time and frequency domains. We achieve an excellent match between the recorded data and the modeled data in the FWI model, despite the use of approximations of physical (visco-acoustic VTI) and numerical (BLR approximation) natures (see Amestoy et al. (2015a) for mode details). Time-domain modeling in the final FWI model allows us to reproduce the main arrivals over the full range of offset and traveltimes (Fig. 2b). In particular, improved delineation of the gas cloud leads to an improved match of the short-spread reflections from the top and bottom of the gas as well as an improved

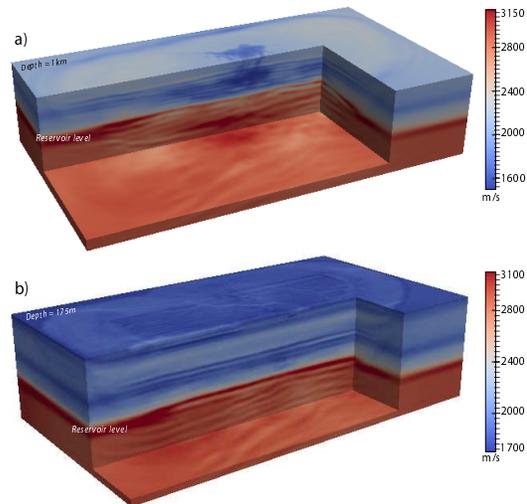


Figure 4: 3D FWI model. (a) Depth slice of the gas cloud (1km depth) and vertical section at the periphery of the gas cloud. (b) Vertical sections off the gas cloud. Note how the geometry of the reflectors below the reservoir level evolves from (a) to (b).

match in offset of shallow critical reflection .

Source wavelet estimation is another tool for quality control of the FWI models (Malinowski et al., 2011; Prioux et al., 2011). We compute wavelets for a representative number of receiver gathers in which a linear gain with offset is applied on the seismograms to increase the sensitivity of the wavelet estimation to the accuracy of the velocity model. We show that the consistency of the wavelets from one receiver gather to the next is greatly improved when the wavelet estimation is performed in the FWI model instead of the initial model (Fig. 5).

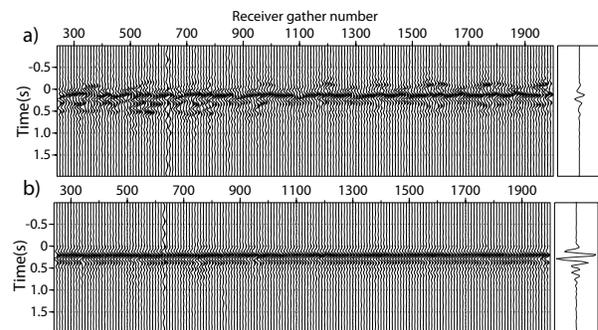


Figure 5: Source wavelet estimation in the initial (a) and final (b) FWI models for a subset of receiver gathers. The inset on the right is the average wavelet. Each wavelet is normalized by its maximum amplitude for sake of readability.

Computational cost

The inversion was performed on computer nodes composed of two 2.5GHz Intel Xeon IvyBridge E5-2670v2 processors with 10 cores per processor. The memory per node is 64Gb and the connecting network is infiniband FDR at 56 Gb/s. We submit

two MPI processes per node and use 10 threads per MPI process. We used 12, 16 and 34 nodes to process the 70m, 50m and 35m grids, respectively. The number of unknowns in the forward-problem linear system is 2.94(70m grid), 7.17(50m grid) and 17.4(35m grid) millions. The memory required by the LU factorization is tractable at these frequencies: 84Gb(70m grid), 288Gb(50m grid), 1.645Tb(35m grid) with the full-rank direct solver. Use of the BLR solver allows us to reduce the number of LU factors by 51%, 61% and 74% on the 70m, 50m and 35m grids, respectively. The elapsed time required to perform the full-rank LU factorization is 76s(70m), 327s(50m) and 1093s(35m). These times are reduced to 38s-41s(70m), 123s-137s(50m) and 351s-394s(35m) with the BLR factorization using a compression threshold ϵ of 10^{-4} - 10^{-5} , respectively (see Amestoy et al. (2015a) for more details).

Although the LU factorization and the multi right-hand sides substitution have the same time complexity ($\mathcal{O}(n^6)$ for a n^3 finite-difference grid), the cost of the full-rank LU factorization is only 11.5%(70m grid), 22%(50m grid), 25%(35m grid) of the cost of the substitution steps. This percentage is further reduced if the BLR factorization is used making the cost of the LU factorization almost negligible. Once the LU factorization is performed, the elapsed time to compute one wavefield solution is only 0.16s(70m grid), 0.37s(50m grid) and 1.1s(35m). Note that, in the current implementation of MUMPS, the substitution step does not exploit the BLR compression, unlike the factorization step. Furthermore, the performances of the solve phase, which inherits the data mapping from the factorization is not optimal and should be revisited; we should also better exploit the sparsity of the right-hand sides in a parallel environment. Therefore, further significant computational saving are expected when these developments will be achieved.

The elapsed time to compute one gradient is 35mn(70m grid), 60mn(50m grid) and 175mn(35m grid). In average, we perform 29, 36, 16 gradients per mono-frequency inversion on the 70m, 50m and 35m grids, respectively. The total time to perform the FWI is $T_{fwi} = \sum_{k=1}^3 N_g^k \times N_f^k \times T_g^k$ where k denotes the three grid intervals and N_g , N_f and T_g are the number of computed gradients, the number of inverted frequencies and the elapsed time to compute one gradient for each grid, respectively. This gives $T_{fwi} = 29 \times 4 \times 35mn + 36 \times 4 \times 60mn + 16 \times 3 \times 175mn = 15$ days. We stop the iterations empirically. Optimized stopping criterion of iteration should contribute to reduce N_g without significantly hampering the model quality.

Time versus frequency-domain FWI

We review some pros and cons of time-domain and frequency-domain FWI based on sparse direct solver to process OBC data. The case study performed by Warner et al. (2013) provides a reference benchmark of time-domain mono-parameter FWI in the 3Hz-6.5Hz frequency band. Their dataset contains 1440 reciprocal sources and 10,000 reciprocal receivers after keeping every third shots and every four receivers. At a second level, they split the reciprocal sources in subset of 80 sources and invert a different subset during each FWI iteration such that the number of iterations times the number of sources in a subset equals to the total number of sources. This subsampling strategy allows them to perform FWI on 40 computer

nodes with 6 cores per node in 2.6 days. With this subsampling strategy, we conclude that the cost of the time-domain and frequency-domain FWI is of the same order of magnitude, a more precise assessment being unrealistic considering the differences between the setup of the two applications.

In both studies, the data were sub sampled to perform FWI. In time-domain FWI, the subsampling was applied in the source and receiver domains, while in this study, the subsampling is applied on the frequency domain. We involve the full set of sources and receivers during each FWI iteration, thanks to the computational efficiency of the substitution step. We believe that the high fold resulting from the fine spatial sampling of the shots and receivers is beneficial to increase the signal-to-noise ratio of the reconstructed subsurface models. We could have applied a similar subsampling strategy over iterations to that adopted by Warner et al. (2013) with the advantage that source subsets of arbitrary size can be chosen. This results because the parallelism in the direct-solver approach does not require to tie the number of sources in a subset to the number of available processors, unlike in time-domain modeling.

One advantage of the frequency domain deals with implementation of attenuation in both seismic modeling and inversion (Malinowski et al., 2011), that will not generate a significant computational burden unlike in the time domain. In relation with the previous item which refers to multiparameter FWI, another advantage concerns the use of more sophisticated optimization algorithms such as the truncated Newton method, which will allow to take into account more accurately the Hessian in the inversion (Métivier et al., 2013, 2014; Castellanos et al., 2015). In the truncated Newton method, the linear system relating the descent direction to the gradient through a Hessian-vector product is solved with an iterative solver. At each iteration of the iterative resolution, two additional modelings need to be performed in the initial model of the current nonlinear iteration. These modelings will be performed efficiently by substitution considering that the LU factors generated for the gradient computation will still be available.

CONCLUSIONS

We show with a real data case study the high computational efficiency of frequency-domain FWI based on sparse direct solver to process OBC data at low frequencies (3.5Hz-10Hz) with limited computational resources. The computational efficiency can still be improved by exploiting the BLR compression during the substitution phase, by taking better advantage of the sparsity of the multiple right-hand sides and by tuning the parallel implementation and data mapping of the solve phase, which has receiver less attention than the factorization phase in the past. Next work deals with multi-parameter FWI with second-order optimization methods.

ACKNOWLEDGMENTS

This study was partly funded by the sponsors of the SEISCOPE consortium. This study was granted access to the HPC resources of SIGAMM and CIMENT computer centers and CINES/IDRIS under the allocation 046091 made by GENCI. We thank BP Norge AS and their Valhall partner Hess Norge AS for allowing access to the Valhall dataset.

EDITED REFERENCES

Note: This reference list is a copyedited version of the reference list submitted by the author. Reference lists for the 2015 SEG Technical Program Expanded Abstracts have been copyedited so that references provided with the online metadata for each paper will achieve a high degree of linking to cited sources that appear on the Web.

REFERENCES

- Amestoy, P.-R., C. Ashcraft, O. Boiteau, A. Buttari, J. Y. L'Excellent, and C. Weisbecker, 2015a, Improving multifrontal methods by means of block low-rank representations: *SIAM Journal on Scientific Computing*, **37**, no. 3, A1451–A1474.
- Amestoy, P., O. Boiteau, R. Brossier, A. Buttari, J. Y. L'Excellent, and T. Mary, L. Métivier, A. Miniussi, S. Operto, C. Puglisi, J. Virieux, and C. Weisbecker, 2015b, 3D frequency-domain seismic modeling with a parallel BIR multifrontal direct solver: 85th Annual International Meeting, SEG, Expanded Abstracts.
- Barkved, O., U. Albertin, P. Heavey, J. Kommedal, J.-P. van Gestel, R. Synnøve, H. Pettersen, and C. Kent, 2010, Business impact of full waveform inversion at Valhall: 80th Annual International Meeting, SEG, Expanded Abstracts, 925–929.
- Castellanos, C., L. Métivier, S. Operto, R. Brossier, and J. Virieux, 2015, Fast full waveform inversion with source encoding and second-order optimization methods: *Geophysical Journal International*, **200**, no. 2, 718–742.
- Krebs, J., J. Anderson, D. Hinkley, R. Neelamani, S. Lee, A. Baumstein, and M.-D. Lacasse, 2009, Fast full-wavefield seismic inversion using encoded sources: *Geophysics*, **74**, no. 6, WCC105–WCC116.
- Malinowski, M., S. Operto, and A. Ribodetti, 2011, High-resolution seismic attenuation imaging from wide-aperture onshore data by visco-acoustic frequency-domain full waveform inversion: *Geophysical Journal International*, **186**, no. 3, 1179–1204.
- Métivier, L., F. Breteau, R. Brossier, S. Operto, and J. Virieux, 2014, Full waveform inversion and the truncated Newton method: Quantitative imaging of complex subsurface structures: *Geophysical Prospecting*, **62**, no. 6, 1353–1375.
- Métivier, L., R. Brossier, J. Virieux, and S. Operto, 2013, Full waveform inversion and the truncated Newton method: *SIAM Journal on Scientific Computing*, **35**, no. 2, B401–B437.
- MUMPS Team, 2011, MULTifrontal Massively Parallel Solver users' guide, version 4.10.0: https://openmodelica.org/svn/OpenModelicaExternal/trunk/3rdParty/Ipopt-3.11.6/ThirdParty/Mumps/MUMPS/doc/userguide_4.10.0.ps, accessed 12 June 2015.
- Operto, S., R. Brossier, L. Combe, L. Métivier, A. Ribodetti, and J. Virieux, 2014, Computationally efficient three-dimensional acoustic finite-difference frequency-domain seismic modeling in vertical transversely isotropic media with sparse direct solver: *Geophysics*, **79**, no. 5, T257–T275.
- Pratt, R. G., 1999, Seismic waveform inversion in the frequency domain, Part 1: Theory and verification in a physical scale model: *Geophysics*, **64**, 888–901.
- Prieux, V., R. Brossier, Y. Gholami, S. Operto, J. Virieux, O. Barkved, and J. Kommedal, 2011, On the footprint of anisotropy on isotropic full waveform inversion: The Valhall case study: *Geophysical Journal International*, **187**, no. 3, 1495–1515.
- 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**, no. 4, 65–70.

- Sirgue, L., and R. G. Pratt, 2004, Efficient waveform inversion and imaging: A strategy for selecting temporal frequencies: *Geophysics*, **69**, 231–248.
- Tarantola, A., 1984, Inversion of seismic reflection data in the acoustic approximation: *Geophysics*, **49**, 1259–1266.
- Warner, M., A. Ratcliffe, T. Nangoo, J. Morgan, A. Umpleby, N. Shah, V. Vinje, I. Stekl, L. Guasch, C. Win, G. Conroy, and A. Bertrand, 2013, Anisotropic 3D full-waveform inversion: *Geophysics*, **78**, no. 2, R59–R80.
- Weisbecker, C., P. Amestoy, O. Boiteau, R. Brossier, A. Buttari, J. Y. L'Excellent, S. Operto, and J. Virieux, 2013, 3D frequency-domain seismic modeling with a block low-rank algebraic multifrontal direct solver: 83rd Annual International Meeting, SEG, Expanded Abstracts, 3411–3416.