



HAL
open science

Large-scale frequency-domain seismic wave modeling on h-adaptive tetrahedral meshes with iterative solver and multi-level domain-decomposition preconditioners

Victorita Dolean, Pierre Jolivet, Stéphane Operto, Pierre-Henri H Tournier

► To cite this version:

Victorita Dolean, Pierre Jolivet, Stéphane Operto, Pierre-Henri H Tournier. Large-scale frequency-domain seismic wave modeling on h-adaptive tetrahedral meshes with iterative solver and multi-level domain-decomposition preconditioners. SEG 2020 - Annual Meeting of the Society of Exploration Geophysicists, Society of Exploration Geophysicists, Oct 2020, Houston, United States. pp.1949-4645, 10.1190/segam2020-3427414.1 . hal-03136868

HAL Id: hal-03136868

<https://hal.science/hal-03136868>

Submitted on 23 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Large-scale frequency-domain seismic wave modeling on h -adaptive tetrahedral meshes with iterative solver and multi-level domain-decomposition preconditioners

V. Dolean^{*}, P. Jolivet[†], S. Operto[‡], P.H. Tournier[§]

^{*} Univ. of Strathclyde, United Kingdom/UCA, CNRS, LJAD, France;

[†] IRIT-CNRS, University of Toulouse, France;

[‡] UCA, CNRS, Geoazur, France;

[§] Sorbonne University, CNRS, France;

SUMMARY

Frequency-domain full-waveform inversion (FWI) is suitable for long-offset stationary-recording acquisition, since reliable subsurface models can be reconstructed with a few frequencies and attenuation is easily implemented without computational overhead. In the frequency domain, wave modeling is a Helmholtz-type boundary-value problem which requires to solve a large and sparse system of linear equations per frequency with multiple right-hand sides (sources). This system can be solved with direct or iterative methods. While the former are suitable for FWI application on 3D dense OBC acquisitions covering spatial domains of moderate size, the later should be the approach of choice for sparse node acquisitions covering large domains (more than 50 millions of unknowns). Fast convergence of iterative solvers for Helmholtz problems remains however challenging in high frequency regime due to the non definiteness of the Helmholtz operator, on one side and on the discretization constraints in order to minimize the dispersion error for a given frequency, on the other side, hence requiring efficient preconditioners. In this study, we use the Krylov subspace GMRES iterative solver combined with a two-level domain-decomposition preconditioner. Discretization relies on continuous Lagrange finite elements of order 3 on unstructured tetrahedral meshes to comply with complex geometries and adapt the size of the elements to the local wavelength (h -adaptivity). We assess the accuracy, the convergence and the scalability of our method with the acoustic 3D SEG/EAGE Overthrust model up to a frequency of 20 Hz.

INTRODUCTION

The ocean bottom node (OBN) acquisition is emerging for deep-offshore seismic imaging by full waveform inversion (FWI) (Beaudoin and Ross, 2007). These stationary-recording acquisitions have the versatility to design ultra-long offset surveys, which provide a wide angular illumination of the subsurface amenable to broadband velocity models. This wide-angle illumination allows for efficient frequency-domain (FD) FWI by decimating the multi-fold wavenumber coverage through a coarse frequency sampling (Pratt, 1999). This frequency subsampling makes FD modeling competitive with time-marching methods and leads to compact datasets (Plessix, 2017). Moreover, attenuation is easily implemented in FWI without computational overheads, even improving the spectral properties (which condition the behavior of solvers) of Helmholtz operators. In this context, we present a new solver for 3D FD wave simulation as a forward engine for FWI. Note that a similar solver has been successfully used in solving a medical imag-

ing problem (Tournier et al., 2019). FD seismic modeling is a boundary-value problem, that after discretization by a finite element method, for example, leads to a sparse linear system whose unknown is the wavefield, the right-hand side (RHS) the seismic source and the coefficients embed the subsurface properties. To solve such a linear system, one can use either a *sparse direct solver* (Duff et al., 2017) or an iterative solver (Saad, 2003). While a direct solver is efficient when processing multiple RHSs for problems of moderate size ($< 50.10^6$ unknowns) (Amestoy et al., 2016), the memory overhead generated by the storage of the LU factors and the limited scalability of the LU decomposition, makes application on large scale problems challenging. The second approach relies on *iterative solvers*, whose natural scalability and moderate memory demand make them suitable for large-scale problems. However, one major issue is the convergence speed of iterative solvers when applied to Helmholtz problems and this convergence deteriorates as the frequency increases. The use of an iterative method, depends on an efficient preconditioner with the ultimate goal to make the iteration count independent to frequencies, and the processing of multiple RHSs.

Here, we focus on the second category because we target large computational domains (several hundred of millions of unknowns) with a limited number of reciprocal sources (from few hundreds to few thousands). Our method relies on a finite-element discretization on a tetrahedral mesh with Lagrange elements of order 3 whose dispersion properties are improved compared to the 27-points finite difference scheme from Operto et al. (2007, 2014); Gosselin-Cliche and Giroux (2014), the Krylov subspace GMRES solver (Saad, 2003) and a Schwarz two-level domain decomposition preconditioner (Graham et al., 2017; Bonazzoli et al., 2019). Compared to the well-known preconditioner based upon shifted Laplacian (Erlangga, 2008), it is less sensitive to the shift (added attenuation) and can be used without it. In the following, we briefly review the discretization and then the solution method, before assessing the strong and weak scalability of the solver on the 3D SEG/EAGE Overthrust model.

TETRAHEDRAL FINITE ELEMENT METHOD

The mathematical model of acoustic wave propagation is the Helmholtz equation

$$\left(\Delta + k^2(\mathbf{x})\right)u(\mathbf{x}, \omega) = b(\mathbf{x}, \omega), \text{ in a subsurface domain } \Omega, \quad (1)$$

where u is the monochromatic pressure wavefield, b the monochromatic source, $k(\mathbf{x}, \omega) = \omega/c(\mathbf{x})$, with ω denoting frequency, $c(\mathbf{x})$ the wavespeed (which is complex valued in viscous media) and $\mathbf{x} = (x, y, z) \in \Omega$.

frequency-domain seismic wave solvers

After discretization, eq. (1) can be written in matrix form as

$$\mathbf{A}\mathbf{u} = \mathbf{b}. \quad (2)$$

We implement the above equation with absorbing boundary conditions along the vertical and bottom faces of Ω and a homogeneous Dirichlet condition on the pressure along the top face (i.e., free-surface boundary condition).

We discretize the Helmholtz equation, eq. (1), with Lagrange finite elements (FE) on a tetrahedral mesh Γ of the domain Ω . The rationale behind this choice is multiple: Compared to finite-difference methods on uniform Cartesian grid (Operto et al., 2014; Gosselin-Cliche and Giroux, 2014), the flexibility offered by unstructured meshes to adapt the size of the elements to the local wavelengths (the so-called h -adaptivity) offers a good trade-off between the precision and the number of degrees of freedom (d.o.f) in the mesh. This is particularly true for elastic wave simulation where the shear wavespeeds can reach very low values just below the sea bottom. Also, compared to hexahedral meshes used with the spectral element method (Li et al., 2020), tetrahedral elements are more versatile to conform the mesh to complex known boundaries (topography, sea bottom, salt bodies) and refine the discretization when FWI proceeds over different frequency bands.

We conduct a series of numerical experiments to find the degree of Lagrange polynomial providing a sufficient accuracy for a discretization rule of five grid points per wavelength (ppwl). We solve the Helmholtz system (2) with the preconditioned GMRES solver (see next section for more details) for an infinite homogeneous $20\text{km} \times 20\text{km} \times 4.65\text{km}$ medium and a 500-m wavelength and compare the solution with the analytical solution. Polynomials of degree 3 (P3) and 5 ppwl (Figure 2) achieve the same level of precision than polynomials of degree 2 (P2) and 10 ppwl (Figure 1). Compared to P2, the P3 discretization also implies a reduction from 206 million to 74 million d.o.f. and a reduction of the solution time from 93s to 24s on 768 cores. We therefore choose to use the P3 elements in the next section to assess the performances of the parallel preconditioner. The choice of P3 elements is consistent with the conclusions of Chaumont-Frelet (2016) for Helmholtz problem and Mulder et al. (2019) for time-domain simulations on tetrahedral meshes.

These numerical results are also supported by the dispersion analysis of Ainsworth and Wajid (2010), from which we plot the normalized numerical phase velocity against $1/G$, with G the number of points per wavelength. The dispersion curve shows that the P3 elements provide a better accuracy over a wide range of wavelengths than the 27-point finite-difference stencil of Operto et al. (2014., see their Figure 3a), while the accuracy of the P2 elements is clearly insufficient for $G = 5$. Note however that the number of non-zero coefficients per row in \mathbf{A} ranges between 47 and 217 for P3 elements, while this number is 27 for the finite-difference discretization.

The simulations of Figures 1 and 2 were performed with a stopping criterion $\varepsilon = \|\mathbf{A}\mathbf{u} - \mathbf{b}\|/\|\mathbf{b}\|$ equal to 10^{-3} (Wilkinson, 1963). Figure 4 shows that this minimum tolerance is the optimal one as going beyond doesn't affect the precision of the numerical solution. This value of ε is also consistent with the analysis of Sourbier et al. (2011, Their Figure 3). Therefore, $\varepsilon = 10^{-3}$ will be used in the next section for the preconditioner

assessment.




Figure 1: Monochromatic wavefield in infinite homogeneous medium. (a) Analytical solution. (b) P2 FE solution with 10 ppwl. (c) Difference. (d) Profile across (a), (b) and (c). Only small amplitude discrepancies are shown.

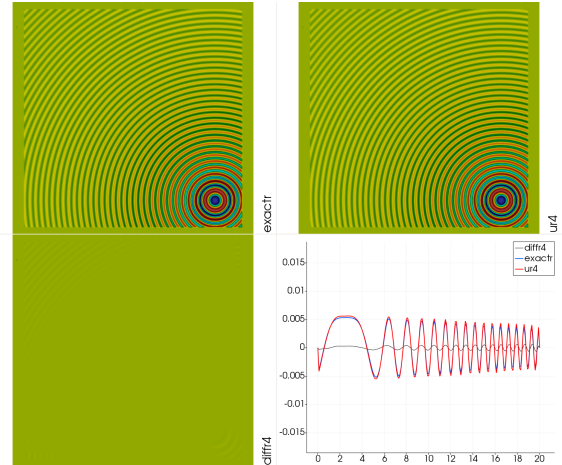


Figure 2: Same as Fig. 1 for P3 elements and 5 ppwl. An accuracy similar to that achieved by P2 and 10ppwl is shown.

DOMAIN DECOMPOSITION PRECONDITIONER

We now review the preconditioner that we use to solve efficiently the linear system (2). A well-known iterative solver for this type of indefinite linear systems is the Krylov subspace Generalized Minimal RESidual Method (GMRES) (Saad, 2003). However, the Helmholtz operator requires efficient preconditioning which can be done by domain decomposition (Dolean et al., 2015, section 2.2.1).

In this study, we solve system (2) with a two-level domain decomposition preconditioner \mathbf{M}^{-1}

$$\mathbf{M}^{-1} = \mathbf{M}_1^{-1}(\mathbf{I} - \mathbf{A}\mathbf{Q}) + \mathbf{Q}, \quad \text{with } \mathbf{Q} = \mathbf{Z}\mathbf{E}^{-1}\mathbf{Z}^T, \quad \mathbf{E} = \mathbf{Z}^T\mathbf{A}\mathbf{Z}, \quad (3)$$

where \mathbf{M}_1^{-1} is the one-level domain decomposition preconditioner called Optimized Restricted Additive Schwarz (ORAS)

frequency-domain seismic wave solvers




Figure 3: Dispersion curves for P2 and P3 spectral and classical finite elements. G denotes the number of grid points per wavelength. The normalized phase velocity is the ratio between the numerical phase velocity and the wavespeed.




Figure 4: Optimal backward error criterion. P3 FE solution with 5 ppwl and different tolerance ($\epsilon = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}$). No significant improvement is shown for ϵ smaller than 10^{-3} .

and \mathbf{Z}^T is the interpolation matrix from the finite element space defined on Γ onto a finite element space defined on a coarse mesh Γ_H . The construction of the domain decomposition preconditioner is described in detail in Bonazzoli et al. (2019). Let $\{\Gamma_i\}_{1 \leq i \leq N_d}$ be an overlapping decomposition of the mesh Γ into N_d subdomains. Let $\{\mathbf{A}_i\}_{1 \leq i \leq N_d}$ denote local Helmholtz operators with absorbing (or transmission) boundary conditions at the subdomain interfaces. The one-level ORAS preconditioner is

$$\mathbf{M}_1^{-1} = \sum_{i=1}^{N_d} \mathbf{R}_i^T \mathbf{D}_i \mathbf{A}_i^{-1} \mathbf{R}_i, \quad (4)$$

where $\{\mathbf{R}_i\}_{1 \leq i \leq N_d}$ are the Boolean restriction matrices from the global to the local finite element spaces and $\{\mathbf{D}_i\}_{1 \leq i \leq N_d}$ are local diagonal matrices representing the partition of unity. The key ingredient of the ORAS method is that the local matrices $\{\mathbf{A}_i\}_{1 \leq i \leq N_d}$ incorporate more efficient boundary conditions (i.e. absorbing boundary conditions) than in the standard RAS preconditioner based on local Dirichlet boundary value problems.

The coarse problem \mathbf{E} in (3) is also solved iteratively by performing 10 GMRES iterations with a one-level ORAS precon-

ditioner. We use the same spatial subdomain partitioning for the coarse and fine meshes. Each computing core is assigned to one spatial subdomain and holds the corresponding coarse and fine local matrices. Each application of the global preconditioner \mathbf{M}^{-1} relies on local concurrent subdomain solves on the coarse and fine levels, which are performed by a direct solver. This hybrid direct/iterative solver requires careful strong scalability analysis to achieve the best compromise between parallel efficiency and memory storage.

NUMERICAL RESULTS

The two-level solver is implemented using the high-performance domain decomposition library HPDDM* (High-Performance unified framework for Domain Decomposition Methods) (Jolivet et al., 2013). We assess the solver on the Irène supercomputer of TGCC[†] with the 3D $20 \times 20 \times 4.65$ km SEG/EAGE Overthrust model (Aminzadeh et al., 1997). We perform wave simulation with P3 finite elements on regular and adaptive tetrahedral meshes (Fig. 5) for the 5 Hz, 10 Hz and 20 Hz frequencies (Tab. 2) in double and single precision. The average length of the element edges is set to 5 nodes per minimum wavelength on the regular tetrahedral mesh, and 5 nodes per local wavelengths in the adaptive tetrahedral mesh (2.5 for the coarser mesh used in the two-level method). We use a ho-




Figure 5: Meshing of Overthrust model. (a) Regular and (b) adaptive tetrahedral meshes.

mogeneous Dirichlet boundary condition at the surface and first-order absorbing boundary conditions along the other five faces of the model. The source is located at $(2.5, 2.5, 0.58)$ km. For weak scalability analysis, we keep #dofs per subdomain roughly constant from one frequency to the next (Tab. 2). The h -adaptivity in the unstructured tetrahedral mesh decreases #dofs relative to the regular mesh by a factor of 2.07. The stopping tolerance ϵ for GMRES is set to 10^{-3} . The consistency between the 10 Hz wavefields computed in the regular and adaptive tetrahedral meshes is shown in Fig. 6.

First, we carry out a set of numerical simulations at 5 Hz on the regular mesh in order to illustrate the benefits of performing computations in single precision arithmetic (versus double precision), as well as using an approximate factorization for the fine local matrices to apply \mathbf{A}_i^{-1} in (4). More precisely,

* <http://github.com/hpddm/hpddm>

[†] <http://www-hpc.cea.fr>

frequency-domain seismic wave solvers




Figure 6: 10 Hz monochromatic wavefields computed in (a) regular and (b) adaptive tetrahedral meshes. (c) Zoom of (b).

we compare incomplete Cholesky factorization (ICC) to complete Cholesky factorization performed by Intel MKL PARDISO. The experiments are performed on 1060 cores with P3 finite elements and 5 ppwl, resulting in 74 million dofs. Results are reported in Tab. 1. First, we can see that performing the whole computation in single precision instead of double precision yields a speedup of about 1.4 for the solution phase. The number of GMRES iterations is the same, there is no loss of accuracy or additional numerical instability. Additionally, the setup phase is drastically reduced (speedup 1.8) when performing Cholesky factorization in single precision. Second, we can see that using an incomplete Cholesky factorization for the fine local matrices yields a speedup of about 1.6 with respect to complete factorization, once again with no effect on the number of GMRES iterations. Moreover, the memory savings are pretty significant: with complete Cholesky factorization we run out of memory with 768 cores, while the simulation runs on 265 cores using ICC. In the rest of this paper, the experiments are performed in single precision and using incomplete Cholesky factorization for the fine local matrices. Timings for the adaptive tetrahedral mesh are around two times smaller than those obtained on the regular mesh (Tab. 2). The simulation at 20 Hz on the adaptive mesh (Fig. 7) involves 2,285 millions of dofs and requires 16,960 cores. The elapsed time achieved by the 2-level preconditioner is 15s and 37s for 10 Hz and 20 Hz respectively (Tab. 2).

CONCLUSION

We have proposed an efficient and accurate forward engine for 3D frequency-domain FWI from ultra-long offset stationary-recording survey. Two key ingredients are the h -adaptive tetrahedral P3 finite-element discretization to optimize the number of d.o.f in the domain and conform the mesh to complex boundaries, and a massively-parallel preconditioned iterative solver for efficient solution of Helmholtz problem. For a comparison with the one-level method and the management of multiple RHS in the case of P2 discretization, see Dolean et al. (2020). Multi-RHS processing can be further improved with the Krylov subspace recycling method GCRO-DR (Parks et al., 2006) and its block variant, which are already implemented in the HPDDM library and have been applied successfully for medical imaging based on a multi-antenna microwave device (Tournier et al., 2019). Extension to visco-elastic media is part of a future work.

Acknowledgments:

This study was granted access to the HPC resources of SIGAMM[‡]

[‡]<http://crimson.oca.eu>

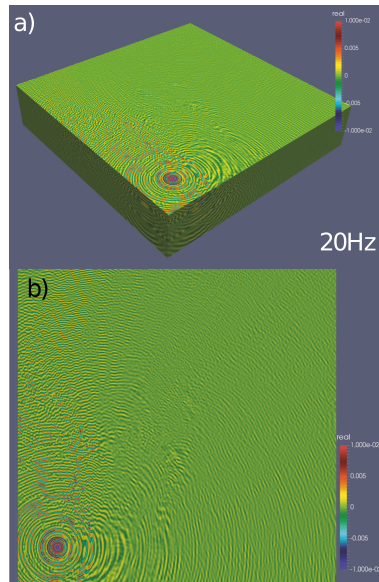


Figure 7: 20 Hz wavefield in the adaptive tetrahedral mesh.

and CINES/IDRIS under the allocation 0596 made by GENCI. This study was partially funded by the WIND consortium[§].

Cartesian grid, $f = 5\text{Hz}$				
precision	fine local solver	#it	setup	GMRES
double	Cholesky	10	92.5s	15.5s
double	ICC	10	30.2s	8.9s
single	Cholesky	10	50.3s	10.3s
single	ICC	10	25.8s	6.3s

Table 1: Comparison between Cholesky and incomplete Cholesky factorization (ICC) of local matrices at the fine level, and single versus double precision arithmetic for the whole computation, at 5Hz with P3 finite elements and 5 ppwl (74M dofs) on 1060 cores. #it: number of iterations. Elapsed time in seconds for the setup phase (assembly and factorization of local matrices) and solution phase with GMRES ($\epsilon = 10^{-3}$).

Regular mesh					
Freq (Hz)	#core	#elts (M)	#dofs (M)	#it	GMRES
5	265	16	74	7	16s
10	2,120	131	575	15	33s
Adaptive mesh					
Freq (Hz)	#core	#elts (M)	#dofs (M)	#it	GMRES
10	2,120	63	286	14	15s
20	16,960	506	2,285	30	37s

Table 2: Statistics of simulation in regular and adaptive tetrahedral meshes. $Freq(Hz)$: frequency; #core: number of cores; #elts: number of elements; #dofs: number of degrees of freedom; #it: iteration count. Elapsed time (seconds) in GMRES.

[§]<https://www.geoazur.fr> WIND sponsored by Chevron, Shell and Total

frequency-domain seismic wave solvers

REFERENCES

- Ainsworth, M., and H. Wajid, 2010, Optimally blended spectral-finite element scheme for wave propagation and nonstandard reduced integration: *SIAM Journal on Numerical Analysis*, **48**, 346–371.
- Amestoy, P., R. Brossier, A. Buttari, J.-Y. L'Excellent, T. Mary, L. Métivier, A. Miniussi, and S. Operto, 2016, Fast 3D frequency-domain full-waveform inversion with a parallel block low-rank multifrontal direct solver: Application to OBC data from the North Sea: *Geophysics*, **81**, R363–R383.
- Aminzadeh, F., J. Brac, and T. Kunz, 1997, 3-D Salt and Overthrust models: SEG/EAGE 3-D Modeling Series No.1.
- Beaudoin, G., and A. Ross, 2007, Field design and operation of a novel deepwater, wide-azimuth node seismic survey: *The Leading Edge*, **26**, 494–503.
- Bonazzoli, M., V. Dolean, I. G. Graham, E. A. Spence, and P.-H. Tournier, 2019, A 2-level domain decomposition preconditioner for the time-harmonic Maxwell's equations: *Math. Comp.*, **88**, 2559–2604.
- Chaumont-Frelet, T., 2016, On high order methods for the heterogeneous Helmholtz equation: *Computers and Mathematics with Applications*, **72**, 2203–2225.
- Dolean, V., P. Jolivet, and F. Nataf, 2015, An introduction to domain decomposition methods: Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA. (Algorithms, theory, and parallel implementation).
- Dolean, V., P. Jolivet, P.-H. Tournier, and S. Operto, 2020, Iterative frequency-domain seismic wave solvers based on multi-level domain-decomposition preconditioners: Presented at the 82th Annual EAGE Meeting (Amsterdam), preprint arXiv:2004.06309.
- Duff, I. S., A. M. Erisman, and J. K. Reid, 2017, Direct methods for sparse matrices, second edition: Oxford University Press, Inc.
- Erlangga, Y., 2008, Advances in iterative methods and preconditioners for the Helmholtz equation: *Archives of Computational Methods in Engineering*, **15**, 37–66.
- Gosselin-Cliche, B., and B. Giroux, 2014, 3D frequency-domain finite-difference viscoelastic-wave modeling using weighted average 27-point operators with optimal coefficients: *Geophysics*, **79**, T169–T188.
- Graham, I. G., E. A. Spence, and E. Vainikko, 2017, Recent results on domain decomposition preconditioning for the high-frequency helmholtz equation using absorption: Lahaye D., Tang J., Vuik K. (eds) *Modern Solvers for Helmholtz Problems*. Geosystems Mathematics. Birkhäuser, Cham, 3–26.
- Jolivet, P., F. Hecht, F. Nataf, and C. Prud'homme, 2013, Scalable domain decomposition preconditioners for heterogeneous elliptic problems: Presented at the Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, ACM.
- Li, Y., R. Brossier, and L. Métivier, 2020, 3d frequency-domain elastic wave modeling with the spectral element method using a massively parallel direct solver: *GEOPHYSICS*, **85**, T71–T88.
- Mulder, W., S. Geewers, and J. Van der Vegt, 2019, Speeding up a Mass-Lumped Tetrahedral Finite-Element Method for Wave Propagation: Presented at the 81th Annual EAGE Meeting (London).
- Operto, S., R. Brossier, L. Combe, L. Métivier, A. Ribodetti, and J. Virieux, 2014, Computationally-efficient three-dimensional visco-acoustic finite-difference frequency-domain seismic modeling in vertical transversely isotropic media with sparse direct solver: *Geophysics*, **79**(5), T257–T275.
- Operto, S., J. Virieux, P. Amestoy, J.-Y. L'Excellent, L. Giraud, and H. Ben Hadj Ali, 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.
- Parks, M., E. de Sturler, G. Mackey, D. Johnson, and S. Maiti, 2006, Recycling krylov subspaces for sequences of linear systems: *SIAM Journal of Scientific Computing*, **28**, 1651–1674.
- Plessix, R.-E., 2017, *in* Some computational aspects of the time and frequency domain formulations of seismic waveform inversion: Springer, 159–187.
- Pratt, R. G., 1999, Seismic waveform inversion in the frequency domain, part I: theory and verification in a physical scale model: *Geophysics*, **64**, 888–901.
- Saad, Y., 2003, *Iterative Methods for Sparse Linear Systems*: SIAM.
- Sourbier, F., A. Haïddar, L. Giraud, H. Ben-Hadj-Ali, S. Operto, and J. Virieux, 2011, Three-dimensional parallel frequency-domain visco-acoustic wave modelling based on a hybrid direct/iterative solver: *Geophysical Prospecting*, **59**, 834–856.
- Tournier, P.-H., I. Aliferis, M. Bonazzoli, M. de Buhan, M. Darbas, V. Dolean, F. Hecht, P. Jolivet, I. E. Kanfoud, C. Migliaccio, F. Nataf, C. Pichot, and S. Semenov, 2019, Microwave tomographic imaging of cerebrovascular accidents by using high-performance computing: *Parallel Computing*, **85**, 88 – 97.
- Wilkinson, J. H., 1963, *Rounding Errors in Algebraic Processes*: Prentice-Halls, Englewood Cliffs, New Jersey.