



HAL
open science

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling using spectral element method with parallel direct linear solvers

Yang Li, Romain Brossier, Ludovic Métivier

► To cite this version:

Yang Li, Romain Brossier, Ludovic Métivier. Feasibility study on 3D frequency-domain anisotropic elastic wave modeling using spectral element method with parallel direct linear solvers. SEG 2019 annual meeting, Sep 2019, San Antonio, United States. pp.3770-3774, 10.1190/segam2019-3215003.1 . hal-02325586

HAL Id: hal-02325586

<https://hal.science/hal-02325586>

Submitted on 24 Nov 2020

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.

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling using spectral element method with parallel direct linear solvers

Yang Li^{*1}, Romain Brossier¹, Ludovic Métivier^{1,2}

¹Univ. Grenoble Alpes, ISTERre, F-38058 Grenoble, France

²Univ. Grenoble Alpes, CNRS, LJK, F-38058 Grenoble, France

SUMMARY

A feasibility study on 3D frequency-domain anisotropic elastic wave modeling is conducted. The spectral element method is applied to discretize the 3D frequency-domain anisotropic elastic wave equation and the linear system is solved by parallel direct solvers, MUMPS and WSMP. A hybrid implementation of MPI and OpenMP for MUMPS is shown to be more efficient in flops and memory cost during the factorization. The influence of complex topography on MUMPS performance is negligible. With available resources, the largest scale modeling, 30 wavelengths in each dimension, is achieved. Using the block low-rank feature of MUMPS leads to computational gains compared with the full-rank version. Limitation of MUMPS scalability for large number of MPI processes prompts us to investigate the performance of an alternative linear solver, WSMP. Preliminary comparison on small scale modelings shows a better scalability of WSMP while being more computational demanding.

INTRODUCTION

Onshore seismic exploration, such as the elastic full waveform inversion (FWI), is challenging due to the complex topography, free surface boundary condition (FSBC), visco-elastic and anisotropic effects. The main computational cost comes from the repeated solution of wave equations to build the model updates. As performing the anisotropic visco-elastic wave modeling is necessary, the majority of current FWI applications relies on the time-domain forward modeling (Warner et al., 2013; Trinh et al., 2019). Nevertheless, frequency-domain FWI possesses several advantages: First, the restriction of Courant-Friedrichs-Lewy (CFL) stability condition no longer exists in the frequency domain; Second, the seismic attenuation (viscosity) can be incorporated easily in the frequency domain by using complex-valued elastic modulus (Carcione, 2015); Third, while inverting realistic scale dataset, time-domain FWI generally relies on source subsampling techniques, frequency-domain FWI invert for only a limited number of discrete frequencies and can account for all the sources provided a direct solver is used to solve the linear system (Duff and Reid, 1983). All these advantages and the recent successful applications of 3D frequency-domain FWI for visco-acoustic VTI medium in offshore environment (Operto et al., 2015; Operto and Miniussi, 2018) prompt us to investigate the feasibility of 3D frequency-domain anisotropic elastic wave modeling.

The spectral element method (SEM) has been investigated particularly in seismology and seismic exploration (Komatitsch and Vilotte, 1998; Komatitsch and Tromp, 1999; Trinh et al., 2019). As a particular instance of finite element method, SEM incorporates FSBC naturally in the weak form of wave equation and uses an adaptive design of the mesh which simplifies the realization of complex topography and acknowledge the variations in the media. The specific character of SEM con-

sists in using a mesh of hexahedra in 3D and choosing the Gauss-Lobatto-Legendre (GLL) points for the integration and Lagrange interpolation. Using high-order Lagrange polynomials and Gauss quadrature on GLL points enables spectral convergence when solving smooth problems. Note that anisotropy can be considered without making an extra effort, unlike the conventional finite-difference method (FDM).

In this study, SEM is applied to 3D frequency-domain elastic wave modeling, incorporating the heterogeneity, anisotropy, complex topography and FSBC. A Cartesian-based fully deformed mesh is used. To solve the linear system generated from the discretization of the frequency-domain wave equation, we resort to direct solvers due to their stability. MUMPS (MUMPS-team, 2017) and WSMP (Gupta et al., 2009) are used. Both the full-rank (FR) and block low-rank (BLR) version of MUMPS are investigated. Performance of FR MUMPS and the computational gain from the BLR MUMPS in terms of time and memory complexity is investigated. The free version of WSMP is tested on small scale modelings. A comparison on the scalability of WSMP and MUMPS is given.

DISCRETIZATION THROUGH SPECTRAL ELEMENT METHOD

The 3D frequency-domain elastic wave equation reads

$$\rho\omega^2 u_j + \frac{\partial}{\partial x_i} \left(c_{ijkl} \frac{\partial u_k}{\partial x_l} \right) + f_j(\omega, \mathbf{r}_s) = 0, \quad i, j, k, l = 1, 2, 3, \quad (1)$$

where ρ is the density, ω is the angular frequency, u_j is the displacement vector, c_{ijkl} is the elastic modulus tensor and $f_j(\omega, \mathbf{r}_s)$ is the point source force vector located at \mathbf{r}_s . Einstein convention is used here. The seismic attenuation in viscoelastic media is incorporated by using complex-valued elastic moduli. The weak form of equation 1 is obtained by multiplying a test function ϕ and integrating over the physical volume Ω . Using the integration by parts and incorporating FSBC and absorbing boundary conditions, the weak form rewrites as

$$\omega^2 \int_{\Omega} \rho u_j \phi \, d\mathbf{x} + \int_{\Omega} c_{ijkl} \frac{\partial u_k}{\partial x_l} \frac{\partial \phi}{\partial x_i} \, d\mathbf{x} + \int_{\Omega} f_j(\omega, \mathbf{r}_s) \phi \, d\mathbf{x} = 0. \quad (2)$$

The volume Ω is then divided into a set of non-overlapping hexahedral elements. An example hexahedral 5th-order SEM element is shown in Figure 1. A mapping is defined to transform a unitary cube $[-1, 1] \otimes [-1, 1] \otimes [-1, 1]$ into a single element Ω^e . The unitary cube is discretized by GLL points. With respect to these GLL points, a scalar function could be approximated by corresponding Lagrange polynomials and an integral could be evaluated by Gauss quadrature. Taking the basis functions as the test function and incorporating the interpolation and GLL quadrature, we obtain the discretized linear system as follows

$$\mathbf{A}\mathbf{u} = \mathbf{f}, \quad (3)$$

where $A = \omega^2 M + K$ is the impedance matrix, M is the mass matrix and K is the stiffness matrix. Vector \mathbf{u} is the discretized

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling with SEM

Figure 1: Point distribution in one hexahedral element for a 5th-order SEM.

#e/dim	N_{PML}	$ e $ (m)	V_P (m/s)	V_S (m/s)	ρ (g/cm ³)	f (Hz)	DOFs / λ
20	2	100	5000	2500	1	25	5

Table 1: Parameter settings for wave modeling.

displacement vector and \mathbf{f} represents the discretized source vector. Note that the combination of the Gauss quadrature and the Lagrange interpolation at GLL points leads to a diagonal mass matrix M for the discretization of the wave equation. This is a huge advantage for the time-domain wave modeling if an explicit time marching scheme is used, no expensive matrix inversion being required.

For the boundary conditions, we adopt the anisotropic PML method (Zhou and Greenhalgh, 2011; Shi et al., 2016) on lateral and bottom sides of the 3D model. With appropriate arrangement, the complex coordinate stretching technique leads to complex-valued elastic parameters and density, while keeping the original wave equation unchanged. The new parameters are defined as follows

$$\tilde{\rho} = \rho s_{x_1} s_{x_2} s_{x_3}, \quad \tilde{c}_{ijkl} = c_{ijkl} \frac{s_{x_1} s_{x_2} s_{x_3}}{s_{x_i} s_{x_k}}, \quad (4)$$

where s_* corresponds to the complex coordinate stretching in each dimension, defined as

$$s_x(x) = \frac{1}{1 + i \frac{\gamma(x)}{\omega}}, \quad \gamma(x) = c_{\text{PML}} \left(1 - \cos \left(\frac{\pi}{2} \frac{x}{L_{\text{PML}}} \right) \right), \quad (5)$$

where c_{PML} is an empirical parameter and L_{PML} is the width of the PML.

NUMERICAL RESULTS FOR DIFFERENT PARALLEL DIRECT SOLVERS

FULL-RANK MUMPS

In this section, MUMPS 5.1.2 (full-rank version) is used to solve the linear system 3. It is based on a multifrontal method (Duff and Reid, 1983), which recasts the original matrix into multiple frontal matrices and computes the LU decomposition of these smaller matrices to save memory and computational cost. We use MUMPS with a hybrid implementation of MPI and OpenMP (MUMPS-team, 2017). The numerical settings are summarized in Table 1. The number of DOFs per wavelength is 5 as we use a 5th-order SEM (De Basabe and Sen, 2007). Thus the number of elements in each dimension is the same as the number of propagated wavelengths.

The total number of MPI and OpenMP is increased from 96 to 256 and the number of OpenMP threads varies from 1 to 8 in order to fit our hardware settings (2 Intel E5-2670 processors per node, 8 cores per processor). Figure 2 presents the corresponding factorization time and memory cost of MUMPS with

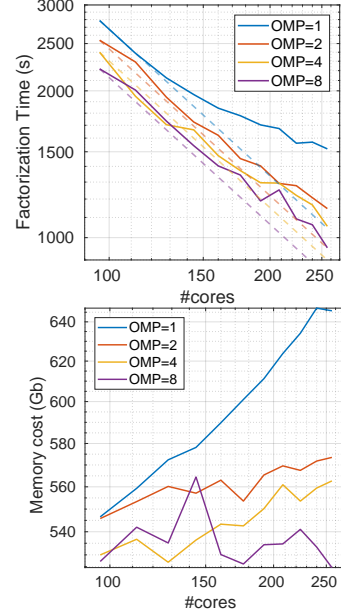


Figure 2: Factorization time and memory cost of MUMPS with different number of OpenMP threads for fixed number of total cores (MPI + OpenMP).

different number of OpenMP threads. The dashed lines indicate the ideal scalability and the solid curves are real computing time. With a fixed total number of MPI and OpenMP, the more OpenMP threads we use, the better MUMPS scales. For a target modeling, fewer MPI leads to larger block matrix, where the BLAS3 part in MUMPS can fully benefit from the matrix-matrix operations. The multithreaded part of MUMPS also improves the overall performance with more OpenMP threads. Both of these allow a better usage of memory when larger number of OpenMP is used, which is also illustrated in Figure 2. Although the memory usage of 8 threads is larger for smaller number of total MPI + OpenMP, the trend agrees well with the expectation when the number of total MPI + OpenMP increases.

To investigate the growth trend of MUMPS memory cost and flops, we increase the model size from $10 \times 10 \times 10$ to $20 \times 20 \times 20$ elements. The parameters settings are the same as in Table 1 except for #e/dim. The number of cores is 64, 96, 160, 192, 224, 256 respectively with #OMP = 8. Free surface boundary condition is taken into account. Two sets of experiments, one with Cartesian non-deformed mesh and the other with vertically deformed mesh, are conducted to test the influence of complex topography. The factorization time and memory cost are presented in Figure 3. It is promising to see that the deformed mesh does not introduce great increment of the factorization time and memory cost because the matrix structure in each case is similar. A direct solver is thus not affected by this modification. Conversely, the surface waves generated from the complex topography and FSBC may lead to drastic convergence delay for iterative solvers (Li et al., 2015).

Table 2 summarizes the largest scale modeling we have achieved so far both with Cartesian and vertically deformed mesh (30 and 28 wavelengths respectively in each dimension). The size of the linear system reaches tens of millions and the number

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling with SEM




Figure 3: Factorization time (upper) and total memory (lower) of wave modeling for different model size, with Cartesian and deformed mesh.

	# λ	DOF	N_{NNZ}	#core	M_T (GB)	T_F (s)	T_S (s)
Cartesian	30	1.4 e7	2.1 e9	384	1913.6	2085.9	5.1
Deformed	28	1.2 e7	5.4 e9	320	1434.6	1828.2	4.6

Table 2: Current largest modeling with MUMPS.

of nonzeros in the matrix surpasses 10^9 . As seen, these results could be obtained using moderate scale of computing resources. Time for solving each RHS is trivial, which is appealing for seismic exploration applications with thousands of sources. The current bottleneck comes from the matrix ordering before the factorization. We use the sequential METIS (Karypis, 2013) which is quite efficient for the subsequent factorization in terms of both time and memory. However, the memory cost of METIS reaches the limit of one cluster node as the model size increases to about $30 \times 30 \times 30$ elements.

PERFORMANCE IMPROVEMENT FROM THE BLOCK LOW-RANK FEATURE OF MUMPS

A Block Low-Rank (BLR) multifrontal solver consists in representing the fronts with low-rank sub-blocks based on the so-called BLR format (Amestoy et al., 2015). Experiments are conducted to assess the performance gain from the BLR feature of MUMPS. Version 5.2.0-beta-pre1 is provided by courtesy of MUMPS team (MUMPS-team, 2019). The model size is $2.5 \times 2.5 \times 2.5$ km. The physical parameters are the same as in Table 1 except for $\#e/\text{dim}$ and f (Hz). The frequencies are 5 Hz, 10 Hz and 20 Hz, corresponding to elements of size 500 m, 250 m and 125 m, and meshes consisted of $5 \times 5 \times 5$, $10 \times 10 \times 10$ and $20 \times 20 \times 20$ elements respectively. The BLR solutions are computed with three values of threshold $\epsilon = 10^{-3}$, 10^{-4} and 10^{-5} and assessed with analytical solutions. The reduction of factorization time, flops and memory cost from the BLR approximation are summarized in Table 3. The computational gain on factorization time provided by the BLR approximation increases with the frequencies. The same trend is shown for

f (Hz)	Factorization time (FR)	Factorization time (%) (BLR)		
		$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$
5	22.207	91.1%	101.3%	108.8%
10	178.193	67.3%	76.7%	88.3%
20	647.065	59.2%	66.6%	77.6%
f (Hz)	Flop count LU (FR)	Flop count LU (%) (BLR)		
		$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$
5	3.666e12	24.5%	37.6%	53.8%
10	4.609e13	15.5%	25.6%	39.9%
20	1.157e15	16.2%	23.5%	35.7%
f (Hz)	Mem LU (FR, GB)	Mem LU (%) (BLR, GB)		
		$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$	$\epsilon = 10^{-5}$
5	8	69.1%	79.9%	88.8%
10	46	60.5%	70.9%	81.9%
20	517	72.9%	80.2%	89.2%

Table 3: Statistics of the full-rank (FR) and block low-rank (BLR) results. f (Hz): the frequency. Flop count LU: number of flops during the LU factorization. Mem LU: memory cost for LU factors in GB. The metrics of BLR for different values of threshold ϵ are given as the percentage of those required by FR factorization.

f (Hz)	FR	BLR (10^{-5})	BLR (10^{-4})	BLR (10^{-3})
5	0.275%	0.276%	0.281%	1.76%
10	0.133%	0.133%	0.193%	14.6%
20	0.181%	0.163%	0.392%	50.5%

Table 4: Relative errors of FR and BLR solutions u_y for different frequencies.

f (Hz)	FR	BLR (10^{-5})	BLR (10^{-4})	BLR (10^{-3})
5	3.07 e-08	3.29 e-07	3.69 e-06	4.93 e-05
10	1.24 e-06	8.35 e-07	8.64 e-06	4.63 e-04
20	2.85 e-05	4.01 e-06	1.35 e-05	3.73 e-04

Table 5: Scaled residuals defined as $\|\mathbf{A}\mathbf{u} - \mathbf{f}\|_{\infty} / \|\mathbf{A}\|_{\infty} \|\mathbf{u}\|_{\infty}$ of FR and BLR solutions for different frequencies.

the flops, but it is not the case for the memory cost. Figure 4 presents the analytical solutions and the difference between analytical solutions and FR, BLR solutions. Table 4 gives the relative errors of u_y wavefields $\|u_y^{\text{ana}} - u_y^{\text{num}}\| / \|u_y^{\text{ana}}\|$. The scaled residuals defined as $\|\mathbf{A}\mathbf{u} - \mathbf{f}\|_{\infty} / \|\mathbf{A}\|_{\infty} \|\mathbf{u}\|_{\infty}$ provided by MUMPS outputs are summarized in Table 5. BLR solutions with $\epsilon = 10^{-3}$ is shown to be inaccurate, especially for high frequency modelings. BLR solutions with $\epsilon = 10^{-5}$ is almost the same as the FR solutions. For 20 Hz, BLR solutions with $\epsilon = 10^{-5}$ is even more accurate than that of FR, which is indicated both from the relative errors and the scaled residuals. For $\epsilon = 10^{-4}$ where the solutions are accurate enough, 60–70% of flops and 20–30% memory cost can be saved compared with the FR version. Performance gain on larger scale modelings could be expected in the future research.

COMPARISON BETWEEN MUMPS AND WSMP

MUMPS is unable to maintain a satisfactory scalability when using more than about one hundred MPI processes (Amestoy et al., 2001; Mary, 2017). Although resorting to a hybrid implementation of MPI and OpenMP extend the number of cores usable for MUMPS, the number of threads has a limit due to different design of CPUs. For large scale modeling where thousands, or tens of thousands of cores have to be used,

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling with SEM




Figure 4: Wavefields u_y of 5 Hz (1st column), 10 Hz (2nd column) and 20 Hz (3rd column), obtained from analytical solutions (1st row). Difference between analytical solutions and BLR solutions with $\epsilon = 10^{-3}$ (2nd row), $\epsilon = 10^{-4}$ (3rd row), $\epsilon = 10^{-5}$ (4th row) and FR solutions (5th row).

MUMPS seems not to be an appropriate choice. In this frame, we investigate the performance of the linear solver WSMP as an alternative because of its scalability up to tens of thousands of cores (Puzryev et al., 2016; Gupta, 2018). WSMP uses a modified multifrontal algorithm with a MPI/OpenMP parallelization. Preliminary tests are conducted using the free version with a limit of 128 cores. As this version is in double complex precision, the tests with MUMPS are changed accordingly (ZMUMPS). Experiments are conducted on three models, with $8 \times 8 \times 8$, $9 \times 9 \times 9$ and $10 \times 10 \times 10$ elements. Pure MPI implementation is set for both solvers. The number of MPI processes increases from 16 to 128. The number of factors and flops with 128 MPIs are summarized in Table 6. Scaling curves of the factorization time are shown in Figure 5, where the ideal scaling is shown in dotted line. We do not show the data for $N_e = 9, N_{MPI} = 16$ and $N_e = 10, N_{MPI} = 16, 32$ because WSMP can not run in such settings. As shown, MUMPS is more efficient for such small scale modelings, in terms of both flops and time. The fewer number of factors also indicates smaller memory consumption of MUMPS. This can be due to the quality of matrix reordering. The graph-partitioning based ordering algorithms used in WSMP may not be efficient as METIS used in MUMPS. However, as shown in Figure 5, MUMPS deviates more from the ideal scaling lines as the number of MPI increases. WSMP stays close to the ideal lines, indicating a better scalability, although longer time and more memory are required. This preliminary test prompt us to investigate the performance of WSMP on larger modelings with more powerful clusters in future research.

		N_e	8	9	10
Factors	MUMPS		2.27e9	3.09e9	4.21e9
	WSMP		4.19e9	5.86e9	7.94e9
Flops	MUMPS		1.29e13	2.04e13	3.27e13
	WSMP		9.71e13	1.61e14	2.54e14

Table 6: Number of factors and flops of MUMPS and WSMP.

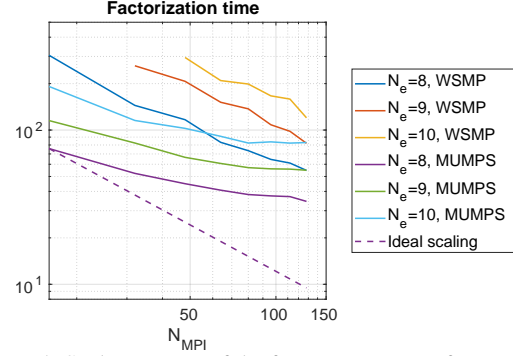


Figure 5: Scaling curves of the factorization time for MUMPS and WSMP, with 10 elements in each dimension.

CONCLUSIONS AND PERSPECTIVES

SEM is applied to 3D frequency-domain anisotropic elastic wave modeling, taking into account complex topography, deformed mesh and FSBC. The performance of FR and BLR MUMPS and WSMP are investigated to solve the generated linear system. MUMPS with a hybrid implementation of MPI and OpenMP presents a satisfactory performance in terms of scalability, flops and memory cost. Deformed mesh only introduce slight extra flops and memory cost compared with Cartesian mesh. It is promising for future onshore applications where complex topography has to be considered. With limited computing resources, a moderate scale modeling (30 wavelengths in each dimension) is achieved. Tests on BLR MUMPS indicates that both the flops and memory cost are reduced due to the low-rank feature of the matrix. To avoid current bottleneck on matrix reordering in MUMPS, alternative methods will be studied. However, the limitation of MUMPS scalability still inhibits its application on larger modelings. As an alternative, WSMP is compared with MUMPS on small scale modelings. WSMP shows better scalability but is more demanding in memory and flops. Further investigation on the performance of WSMP for larger scale modelings using more powerful clusters is thus worthwhile. Other parallel direct solvers will also be future investigation aspects.

ACKNOWLEDGEMENTS

This study was partially funded by the SEISCOPE consortium (<https://seiscope2.osug.fr>), sponsored by AKERBP, CGG, CHEVRON, EQUINOR, EXXON-MOBIL, JGI, PETROBRAS, SCHLUMBERGER, SHELL, SINOPEC and TOTAL. This study was granted access to the HPC resources of CIMENT infrastructure (<https://ciment.ujf-grenoble.fr>) and CINES/IDRIS/TGCC under the allocation 046091 made by GENCI. The authors thank the MUMPS team for providing MUMPS package and appreciate the discussions with Théo Mary, Alfredo Buttari, Jean-Yves L'Écellent and Patrick Amestoy, which improves greatly on the usage of MUMPS.

Feasibility study on 3D frequency-domain anisotropic elastic wave modeling with SEM

REFERENCES

- Amestoy, P. R., C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, and C. Weisbecker, 2015, Improving multifrontal methods by means of block low-rank representations: *SIAM Journal on Scientific Computing*, **37**, 1451–1474.
- Amestoy, P. R., I. S. Duff, J. Koster, and J. Y. L'Excellent, 2001, A fully asynchronous multifrontal solver using distributed dynamic scheduling: *SIAM Journal of Matrix Analysis and Applications*, **23**, 15–41.
- Carcione, J. M., 2015, *Wave fields in real media, wave propagation in anisotropic, anelastic, porous and electromagnetic media*, third edition ed.: Elsevier.
- De Basabe, J. D., and M. K. Sen, 2007, Grid dispersion and stability criteria of some common finite-element methods for acoustic and elastic wave equations: *Geophysics*, **72**, T81–T95.
- Duff, I. S., and J. K. Reid, 1983, The multifrontal solution of indefinite sparse symmetric linear systems: *ACM Transactions on Mathematical Software*, **9**, 302–325.
- Gupta, A., 2018, *WSMP: Watson Sparse Matrix Package Part II - direct solution of general systems - version 18.06*. IBM T. J. Watson Research Center.
- Gupta, A., S. Koric, and T. George, 2009, Sparse matrix factorization on massively parallel computers: *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, ACM, 1:1–1:12.
- Karypis, G., 2013, *METIS - A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices - version 5.1.0*. University of Minnesota.
- Komatitsch, D., and J. Tromp, 1999, Introduction to the spectral element method for three-dimensional seismic wave propagation: *Geophysical Journal International*, **139**, 806–822.
- Komatitsch, D., and J. P. Vilotte, 1998, The spectral element method: an efficient tool to simulate the seismic response of 2D and 3D geological structures: *Bulletin of the Seismological Society of America*, **88**, 368–392.
- Li, Y., L. Métivier, R. Brossier, B. Han, and J. Virieux, 2015, 2D and 3D frequency-domain elastic wave modeling in complex media with a parallel iterative solver: *Geophysics*, **80**(3), T101–T118.
- Mary, T., 2017, *Block Low-Rank multifrontal solvers: complexity, performance and scalability*: PhD thesis, PhD thesis, Université de Toulouse.
- MUMPS-team, 2017, *MUMPS - MULTifrontal Massively Parallel Solver users' guide - version 5.1.2 (october, 2017)*. ENSEEIHT-ENS Lyon, <http://www.enseeiht.fr/apo/MUMPS/> or <http://graal.ens-lyon.fr/MUMPS>.
- , 2019, *MUMPS - MULTifrontal Massively Parallel Solver (MUMPS 5.2.0-beta-pre1) Users' guide (february, 2019)*. ENSEEIHT-ENS Lyon, <http://www.enseeiht.fr/apo/MUMPS/> or <http://graal.ens-lyon.fr/MUMPS>.
- Operto, S., and A. Miniussi, 2018, On the role of density and attenuation in 3D multi-parameter visco-acoustic VTI frequency-domain FWI: an OBC case study from the North Sea: *Geophysical Journal International*, **213**, 2037–2059.
- Operto, S., A. Miniussi, R. Brossier, L. Combe, L. Métivier, V. Monteiller, A. Ribodetti, and J. Virieux, 2015, Efficient 3-D frequency-domain mono-parameter full-waveform inversion of ocean-bottom cable data: application to Valhall in the visco-acoustic vertical transverse isotropic approximation: *Geophysical Journal International*, **202**, 1362–1391.
- Puzyrev, V., S. Koric, and S. Wilkin, 2016, Evaluation of parallel direct sparse linear solvers in electromagnetic geophysical problems: *Computers & Geosciences*, **89**, 79–87.
- Shi, L., Y. Zhou, J.-M. Wang, M. Zhuang, N. Liu, and Q. H. Liu, 2016, Spectral element method for elastic and acoustic waves in frequency domain: *Journal of Computational Physics*, **327**, 19 – 38.
- Trinh, P. T., R. Brossier, L. Métivier, L. Tvard, and J. Virieux, 2019, Efficient 3D time-domain elastic and viscoelastic Full Waveform Inversion using a spectral-element method on flexible Cartesian-based mesh: *Geophysics*, **84**, R75–R97.
- 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**, R59–R80.
- Zhou, B., and S. A. Greenhalgh, 2011, 3-d frequency-domain seismic wave modelling in heterogeneous, anisotropic media using a gaussian quadrature grid approach: *Geophysical Journal International*, **184**, 507–526.