

Combining direct and iterative solvers for improving efficiency of solving wave equations when considering multi-sources problems

by F. Sourbier¹, A. Haidar², L. Giraud³, R. Brossier⁴, S. Operto⁵ and J. Virieux⁶

¹ Géosciences Azur, Sophia-Antipolis, France, E-Mail: sourbier@geoazur.unice.fr

² CERFACS, Toulouse, France, E-Mail: haidar@cerfacs.fr

³ ENSEEIHT-IRIT, Toulouse, France, E-Mail: Luc.Giraud@enseeiht.fr

⁴ Géosciences Azur, Sophia-Antipolis, France, E-Mail: brossier@geoazur.unice.fr

⁵ Géosciences Azur, Villefranche-sur-mer, France, E-Mail: operto@geoazur.obs-vlfr.fr

⁶ LGIT, Grenoble, France, E-Mail: jean.virieux@obs.ujf-grenoble.fr

Introduction

Frequency-domain full-waveform inversion (FWI) has been extensively developed during last decade to build high-resolution velocity models (Pratt, 2004). One advantage of the frequency domain is that inversion of a few frequencies are enough to build velocity models from wide-aperture acquisitions. Multi-source frequency-domain wave modeling requires resolution of a large sparse system of linear equations with multiple right-hand side (RHS). In 3D geometries or for very large 2D problems, the memory requirements of state-of-the-art direct solvers preclude applications involving hundred millions of unknowns. In order to overcome this limitation, we investigate a domain decomposition method based on the Schur complement approach for 2D/3D frequency-domain acoustic wave modeling. The method relies on a hybrid direct-iterative solver. Direct solver is applied to sparse impedance matrices assembled on each subdomain, hence, reducing the memory requirement of the overall simulation. Iterative solver based on a preconditioned Krylov method is used for solving the interface nodes between adjacent domains. A possible drawback of the hybrid approach is that the time complexity of the iterative part linearly increases with the number of RHS, if single-RHS Krylov subspace method is sequentially applied to each RHS. We mention that block-Krylov techniques or deflation techniques can be used in that case to partially overcome this effect. In the following, we introduce the domain decomposition method before illustrating its features with 2D and 3D simulations.

Helmholtz equation

The visco-acoustic wave equation is written in the frequency domain as

$$\frac{\omega^2}{\kappa(\mathbf{x})}p(\mathbf{x}, \omega) + \nabla \left(\frac{1}{\rho(\mathbf{x})} \nabla p(\mathbf{x}, \omega) \right) = -s(\mathbf{x}, \omega) \quad (1)$$

where $\rho(\mathbf{x})$ is density, $\kappa(\mathbf{x})$ is the bulk modulus, ω is angular frequency; $p(\mathbf{x}, \omega)$ and $s(\mathbf{x}, \omega)$ denote the pressure and source respectively. Eq. (1) can be recast in matrix form as $\mathbf{A}\mathbf{p} = \mathbf{s}$, where the complex-valued impedance matrix \mathbf{A} depends on ω , κ and ρ . \mathbf{p} and \mathbf{s} are stored as vectors of dimension equal to the product of the dimensions of the cartesian computational grid. For a 3D geometry, we discretized eq. (1) with the mixed-grid finite-difference (FD) stencil (Operto *et al.*, 2007) which has an accuracy similar to that of 4th-order accurate stencils while minimizing the numerical bandwidth of \mathbf{A} , a key point to mitigate its fill-in during LU factorization.

Schur complement method

The domain decomposition consists of splitting the computational domain into sub-domains without overlap. Each sub-domain shares interface nodes with its adjacent sub-domains. Such decomposition is referred as sub-structuring schemes (Saad, 2003; Smith *et al.*, 1996). After reordering the interior nodes by sub-domain and labeling the interface nodes last, the system $\mathbf{A}\mathbf{p} = \mathbf{s}$ becomes

$$\begin{bmatrix} A_{ii}^1 & & & A_{ib}^1 \\ & A_{ii}^2 & & A_{ib}^2 \\ & & \ddots & \vdots \\ & & & A_{ii}^n & A_{ib}^n \\ A_{bi}^1 & A_{bi}^2 & \cdots & A_{bi}^n & \bar{A}_{bb} \end{bmatrix} \begin{bmatrix} p_i^1 \\ p_i^2 \\ \vdots \\ p_i^n \\ \bar{p}_b \end{bmatrix} = \begin{bmatrix} s_i^1 \\ s_i^2 \\ \vdots \\ s_i^n \\ \bar{s}_b \end{bmatrix} \quad (2)$$

where p_i^j denote unknowns located at interior nodes of sub-domain j and \bar{p}_b denote unknowns located at all interface nodes. Note that indices b and i label interface and interior nodes respectively while the exponent labels sub-domains. The number of sub-domains is denoted n .

The system (2) can be written in compact form as

$$\begin{bmatrix} A_{ii} & A_{ib} \\ A_{bi} & \bar{A}_{bb} \end{bmatrix} \begin{bmatrix} p_i \\ \bar{p}_b \end{bmatrix} = \begin{bmatrix} s_i \\ \bar{s}_b \end{bmatrix} \quad (3)$$

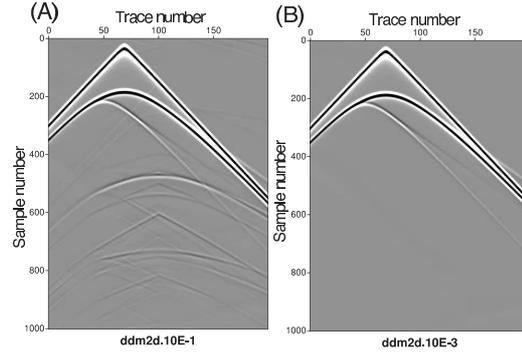


Figure 1: Seismograms over two hundred receivers computed in the corner-edge model with the hybrid solver. The convergence tolerance ϵ is 10^{-1} on the left panel (a) and 10^{-3} on the right panel (b).

Eliminating p_i from the second block row of eq. (3) leads to the following reduced system for \bar{p}_b

$$(\bar{A}_{bb} - A_{bi}A_{ii}^{-1}A_{ib})\bar{p}_b = s_b - A_{bi}A_{ii}^{-1}s_i \quad (4)$$

The matrix $S = \bar{A}_{bb} - A_{bi}A_{ii}^{-1}A_{ib}$ is the Schur complement matrix. In parallel distributed environment, each sub-domain is assigned to one processor so that the solution of system (3) is computed in parallel using the following steps:

- concurrent factorizations are performed on each processor to form the so called “local Schur complement” using the direct solver MUMPS (MUMPS-team, 2006);
- build the RHS of equation (4);
- solve eq. (4) using a parallel distributed preconditioned Krylov solver implemented using the Generalized Minimal Residual iterative Solver (GMRES) (Frayssé *et al.*, 2003). We used both an additive Schwarz preconditioner (M_{as}) (Giraud *et al.*, 2008; Giraud and Tuminaro, 2006) and its variant incorporating a complex shift on the diagonal (M_{asc}) (Erlangga *et al.*, 2004). The preconditioner is derived from the local assembled Schur complements \bar{S}^j , namely, the restriction of the Schur complement to the interface nodes of sub-domain j . These local assembled Schur complements are computed from the local Schur complements S^j by incorporating neighbor contributions through a few point-to-point communications. We have $M_{as} = \sum_{j=1}^n (\bar{S}^j)^{-1}$, where the dense \bar{S}^j are factorized with a Lapack routine.
- Once system (4) was solved, the unknowns p^j associated with interior nodes in eq. (2) are computed concurrently by solving $p_i = A_{ii}^{-1}(s_i - A_{ib}\bar{p}_b)$.

Accuracy of the solution: a 2D illustration of the method

Contrary to direct solvers for which solutions have a high accuracy (machine precision), the hybrid solver requires a criterion for ending iterations. The convergence tolerance ϵ for the backward error on the GMRES algorithm controls the accuracy of the solution. Of course, the number of iterations and the CPU time increase in GMRES when the convergence tolerance ϵ decreases. All simulations were performed in single precision.

We first compare the 2D/3D analytical and hybrid solver solutions in a homogenous media. A value of $\epsilon = 10^{-3}$ seems to provide the best compromise between accuracy and iteration count. Secondly, we compare numerical solutions in 2D heterogeneous media provided by a finite-difference frequency-domain method based on a direct and hybrid solvers respectively. The velocity model is a corner-edge model composed of two homogeneous layers delineated by a horizontal and vertical interfaces forming a corner. The grid dimensions are 801×801 with a grid interval of 40 m. Velocities are $4km/s$ and $6km/s$ in the upper-left and bottom-right layers respectively. Source wavelet is a Ricker with a dominant frequency of 5 Hz. Time-domain seismograms computed with the hybrid solver for the 2 values of ϵ are shown in Figure 1 (the second value leads to a simulation similar to the one obtained through a direct solver procedure). Implementing the hybrid solver into a Finite Difference Frequency Domain Full Waveform Tomography code is the next step to assess precisely which convergence tolerance is needed for imaging applications.

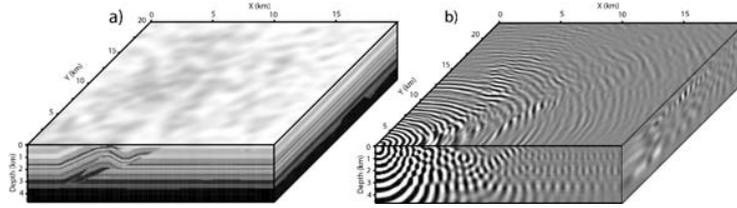


Figure 2: a) 3D SEG/EAGE overthrust model. b) 7-Hz monochromatic wavefield computed with the hybrid solver.

Simulation in the 3D Overthrust model

The 3D SEG/EAGE Overthrust model is a constant-density acoustic model covering an area of $20 \times 20 \times 4.65 \text{ km}^3$. We perform a simulation for the 7-Hz frequency using both the direct and hybrid approaches (Figure 2). The model is resampled with a grid interval of 75 m corresponding to four grid points per minimum wavelength at 7 Hz. This leads to a velocity grid of $277 \times 277 \times 73$ nodes including PML layers (5.6 millions of unknowns). We use processors with 2 Gb of memory. Due to the memory requirement of the direct solver, simulations using a direct and hybrid solvers are performed in 75% of the model (corresponding to 4.25×10^6 unknowns) using 192 processors. For the hybrid solver, we use $\epsilon = 10^{-3}$ and the M_{asc} preconditioner. The statistics of the two simulations are summarized in Table 1. For 192 processors, the hybrid solver requires 2.5 less memory than the direct one and is 8 times faster for one RHS. The direct solver will become faster for a number of RHS greater than 9 suggesting that the direct approach remains the approach of choice for small 3D grids and for 2D grids (except quite large ones).

We also perform a simulation at 7 Hz in the full overthrust model (5.6 millions of unknowns) using the hybrid solver. The minimum number of processors required to tackle this simulation is 128 corresponding to $128 = 8 \times 8 \times 2$ sub-domains (Table 2). The statistics of two simulations performed on 128 and 192 processors are compared in Table 2. Increasing the number of processors from 128 to 192 allowed to reduce the memory requirement from 164 to 137 Gb and the total CPU time from 598 s to 488 s. Note that the memory requirement decreases with the number of processors for the hybrid approach contrary to the direct approach for which an increase of memory is observed because of overheads.

3D	Direct Solver	Hybrid Solver
N_p	192	192
N_u	4.25×10^6	4.25×10^6
Mem (Gb)	235	95.6
$T_f(s)$	2876	14
$T_s(s)$	9	-
$T_i(s)$	-	296
$T_M(s)$	-	46
$T_{total}(s)$	2885	356

Table 1: Computational cost of the 3D-overthrust simulations using a direct (left) and hybrid (right) solvers. N_p : number of processors. N_u : number of unknowns. $Mem(Gb)$: total allocated memory. $T_f(s)$: elapsed time for factorization in the direct and hybrid approaches. $T_s(s)$: elapsed time for resolution in the direct approach; $T_i(s)$: elapsed time for GMRES in the hybrid approach; $T_M(s)$: elapsed time to build the preconditioner in the hybrid approach; $T_{total}(s)$: total elapsed time to achieve the simulation.

Scalability analysis

The hybrid direct-iterative solver allows a significant memory saving compared to direct solvers. This is quantified from the theoretical memory complexity of LU factorization of sparse matrix and the ratio between the size of the full domain and that of the sub-domains (Table 3). This heuristic complexity

N_p	Mem (Gb)	T_f(s)	T_i(s)	T_M(s)	T_total(s)
128	164	68	410	120	598
192	137	28	391	69	488

Table 2: Simulations in the full 3D-overthrust model at $7 Hz$ (5.60×10^6 unknowns) using the hybrid solver for 2 numbers of processors. Notations of Table 1 are used.

estimation through the overhead coefficient ov is specific to the FD discretization we use and to the reordering method (nested dissection) (Operto *et al.*, 2007). The CPU time complexity is more difficult to assess because it depends on several parameters such that the number of sub-domains, their geometries, the convergence criterion, the number of RHSs. Increasing the number of sub-domains reduces the CPU time requirement of the direct solver at the partial expense of the number of iterations in GMRES because of more approximate preconditioner.

	Direct Solver	Hybrid Solver
3D	$ov \times 35 \times n^4$	$107 \times n^4/k$

Table 3: Memory complexity of the direct and hybrid solvers. The dimension size of a 3D n^3 grid is denoted by n while the number of sub-domain along each direction is denoted by k . The memory overhead coefficient is denoted by ov .

We perform a scalability analysis of the hybrid solver using a small target of the 3D SEG/EAGE Overthrust model centred on a channel. The model covers an area of $7 \times 8.8 \times 2.3 km^3$ and is discretized with 50 m cubic cells. This lead to a $141 \times 176 \times 46$ grid. The model is augmented with PML layers spanning over 5 grid points leading to 1.6 millions of unknowns. The frequency is $5 Hz$. We perform a serie of simulations by increasing the number of processors, ie. the number of sub-domains (Table 4). One can note that the total memory may punctually increase with the number of processors because the domain partitioning resulting from some particular number of processors may not be optimal to minimize the number of interface nodes. The optimal subdomain geometry to minimize the number of interface nodes is that of cubes. However, we observe an overall decrease of the memory when the number of processor increases as previously illustrated in Table 2. As expected, the number of iterations performed by GMRES increases with the number of sud-domains due to more approximate preconditioner but we observe so far a decrease of the time of the iterative solving loop and, therefore, of the total CPU time, thanks to the intrinsic scalability of the iterative solver.

N_p	20	32	40	60	168
Mem_f(Mb)	1141	767	592	359	84
Mem_s(Mb)	339	224	250	111	36
Mem (Gb)	36	39	44	35	26
Iter	23	36	38	102	151
T_f(s)	142	77	53	26	4
T_M(s)	-	134	84	46	7
T_i(s)	-	205	118	110	38
T_total(s)	-	416	255	182	49

Table 4: Scalability analysis for the Channel model. Notations of Table 1 are used. $Mem_f(Mb)$: memory allocated per processor for local factorization. $Mem_s(Mb)$: memory used to store each local Schur complement matrix. $Mem(Gb)$: total memory used. $Iter$: number of GMRES iterations to achieve convergence at $\epsilon = 10^{-3}$. '-' means not enough memory.

Conclusions and future works

A domain decomposition method based on a hybrid direct-iterative solver has been developed for 2D/3D frequency-domain acoustic wave modeling. The memory saving provided by the hybrid solver with respect to the direct one is significant, especially for 3D problems, and increases with the number of processors. The hybrid solver is also significantly faster than the direct one for single-source problem but this trend tends to reverse when the number of sources increases due to the efficiency of the solution phase of direct solvers. The ability of block-Krylov methods or deflation techniques to efficiently handle multiple RHSs will be the key issue to design efficient modeling tool for 3D frequency-domain FWI based on hybrid solver. Using data reduction based on source stacking (Capdeville *et al.*, 2005) may provide a last resort to overcome the computational overburden of multishot simulations.

Acknowledgements

This study has been partially funded by the French ANR through grant ANR-05-NT05-242427 and the SEISCOPE consortium (<http://seiscope.unice.fr>) sponsored by BP, CGG-Veritas, EXXON-Mobil, SHELL and TOTAL. The simulations have been computed on the PC Cluster of Mesocentre SIGAMM.

References

- Capdeville, Y., Gung, Y., and Romanowicz, B. [2005] Towards global earth tomography using the spectral element method: a technique based on source stacking. *Geophys. J. Int.* 162, 541–554.
- Erlangga, Y. A., Vuik, C., and Oosterlee, C. W. [2004] On a class of preconditioners for solving the helmholtz equation. *Appl. Numer. Math.* 50(3-4), 409–425.
- Frayssé, V., Giraud, L., Gratton, S., and Langou, J. 2003. *A set of GMRES routines for real and complex arithmetics on high performance computers* Technical Report TR/PA/03/03. CERFACS, Toulouse, France.
- Giraud, L., and Tuminaro, R. S. 2006. *Algebraic domain decomposition preconditioners* Technical Report ENSEEIHT-IRIT RT/APO/06/07. IRIT Preliminary version of a book chapter to appear in *Mesh partitioning techniques and domain decomposition methods*, Civil-Comp Ltd, F. Magoules ed, p. 187-216, 2007.
- Giraud, L., Haidar, A., and Watson, L. T. [2008] Parallel scalability study of three dimensional additive Schwarz preconditioners in non-overlapping domain decomposition. *Parallel computing, in press.*
- MUMPS-team. 2006 (Feb.). *MUMPS - Multifrontal Massively Parallel Solver users' guide - version 4.6.3 (february 2006)* ENSEEIHT-ENS Lyon, <http://www.enseeiht.fr/apo/MUMPS/> or <http://graal.ens-lyon.fr/MUMPS>.
- Operto, S., Virieux, J., Amestoy, P., L'Excellent, J.Y., Giraud, Luc, and Ben-Hadj-Ali, H. [2007] 3D finite-difference frequency-domain modeling of visco-acoustic wave propagation using a massively parallel direct solver: A feasibility study. *Geophysics* 72, 195–211.
- Pratt, R. G. [2004] Velocity models from frequency-domain waveform tomography: past, present and future. *In: Expanded abstracts Eur. Ass. Expl. Geophys.*
- Saad, Y. [2003] *Iterative methods for sparse linear systems* Vol. second edition Philadelphia: SIAM.
- Smith, B. F., Bjørstad, P. E., and Gropp, W. [1996] *Domain decomposition: Parallel multilevel methods for elliptic partial differential equations* Cambridge University Press.