5th International Symposium on Three-Dimensional Electromagnetics
Sapporo, Japan, May 7–9, 2013

Parallel Finite-Element Method for 3-D Electromagnetic Modelling in Geophysics

Jelena Koldan[1], Vladimir Puzyrev[1] and José María Cela[1]
[1]Barcelona Supercomputing Center, Computer Applications in Science & Engineering (CASE) Department

SUMMARY

We have developed a parallel nodal finite-element solver for three-dimensional electromagnetic numerical modelling in geophysics. The method can be used for modelling different controlled-source and magnetotelluric problems in anisotropic media. Due to the fact that it supports completely unstructured tetrahedral meshes as well as mesh refinement, it is possible to represent complex geological structures very precisely and thus improve the solution accuracy. In addition, the parallel implementation of the method has proved to be highly scalable – the achieved speed-up is close to linear for up to thousands of processors. Thanks to this, the code is able to deal with extremely large problems, which may have tens of millions of degrees of freedom, in a very efficient way. Considering that efficiency of the forward-problem solver is critical for its future use inside of a 3-D inversion algorithm, we have developed different preconditioning schemes for Krylov subspace methods, namely approximate inverse based and algebraic multigrid based preconditioners. Tests for various problems with different conductivity structures and characteristics have shown that, our preconditioners greatly improve the convergence of different iterative solvers and significantly reduce the total execution time of the program.

Keywords: forward modelling, finite element, preconditioning, algebraic multigrid

INTRODUCTION

The growing use and technical development of electromagnetic (EM) methods in geophysics, e.g. controlled-source (CSEM) and magnetotelluric (MT) techniques, have led to the increasing need for reliable and fast methods of interpretation of three-dimensional EM data acquired in complex geological environments. In order to obtain realistic subsurface images of large Earth areas, industrial large-scale surveys need to collect vast amounts of data, which makes the solution of the 3-D EM inverse problem computationally immensely challenging. One of the reasons for these enormous computational demands is the expensive solution of the 3-D EM forward problem, which is, in addition, solved many times inside an inversion algorithm. 3-D EM modelling is demanding due to the fact that the grids designed to approximate huge realistic 3-D geologies are usually enormous in order to correctly represent complex structures, e.g. bathymetric variations (Sasaki & Meju, 2009). Clearly, the ability to solve the forward problem as efficiently as possible is essential to the strategies for performing 3-D EM inversion.

3-D EM modelling involves the numerical solution of Maxwell's equations in electrically conductive anisotropic media. The most commonly used approaches are finite-difference (FD) and finite-element (FE) methods. No matter which approach is used, the initial forward problem is always reduced to the solution of a system of linear equations. Usually, this resultant linear system, which is very large and sparse, is solved using iterative methods, such as Krylov subspace methods (Saad, 2003). However, Krylov methods most often converge extremely slowly, if they converge at all, since resultant linear systems are generally poorly conditioned (Key & Ovall, 2011). In order to overcome this problem, a variety of preconditioners have been designed and applied, such as the standard SSOR and ILU (Hou, Mallan, & Torres-Verdín, 2006), the low induction number (LIN) preconditioner (Weiss & Newman, 2003) and multigrid preconditioners (Aruliah & Ascher, 2002; Haber, 2004; Mulder, 2006). Algebraic multigrid methods (AMG), originally designed for creating standalone solvers, can be very good preconditioners, as well (Trottenberg, Oosterlee, & Schüller, 2001). Very efficient way to deal with huge computational requirements coming from EM field simulations on large realistic Earth models is to design an algorithm to run on parallel computers. Nowadays, it is well-known that such algorithms are capable of handling large-scale problems that cannot be easily treated otherwise (Commer & Newman, 2004).

We present an efficient parallel nodal finite-element solver for 3-D EM numerical modelling in geophysics. By designing the algorithm from the beginning to run on massively parallel computing platforms, we have created a highly scalable and very fast code that can cope with extremely large problems in an efficient way. Scalability tests on massively parallel computers have shown that our code achieves significant speed-ups for up to thousands of processors. In addition, we present a preconditioning technique based on AMG which improves convergence of

Krylov subspace methods and thus increases efficiency of the whole numerical scheme to a great extent, as proved by tests on various models.

NUMERICAL MODELLING

The physical problem has been formulated in terms of the secondary Coulomb-gauged EM potentials (Badea, Everett, Newman, & Biro, 2001). By employing the EM potentials, the diffusive Maxwell's equations have been transformed into the following equations, which we solve numerically:

$$\nabla^2 \mathbf{A}_s + k\tilde{\sigma}(\mathbf{A}_s + \nabla\Psi_s) = -k\Delta\tilde{\sigma}(\mathbf{A}_p + \nabla\Psi_p) \quad (1)$$

$$\nabla \cdot [k\tilde{\sigma}(\mathbf{A}_s + \nabla\Psi_s)] = -\nabla \cdot [k\Delta\tilde{\sigma}(\mathbf{A}_p + \nabla\Psi_p)] \quad (2)$$

where $k = i\omega\mu_0$ and $\Delta\tilde{\sigma} = \tilde{\sigma} - \tilde{\sigma}_p$ is the difference between the conductivity distribution $\tilde{\sigma}(\mathbf{r})$ and the background conductivity, $\tilde{\sigma}_p$, whose response is already known. Our code assumes transverse anisotropy with horizontal, $\sigma_x = \sigma_y = \sigma_h$, and vertical, $\sigma_z = \sigma_v$, conductivity. We remark that different types of anisotropy could be easily modeled, as well.

The boundaries of a domain are considered to be located far away from the transmitter and we impose homogeneous Dirichlet boundary conditions:

$$(\mathbf{A}_s, \Psi_s) = (\mathbf{0}, 0) \text{ on } \Gamma \quad (3)$$

The equations (1) and (2) together with the boundary conditions (3) fully describe the EM induction caused by different types of sources in anisotropic electrically conductive media.

In order to discretise the governing equations, we apply the finite-element method because it supports completely unstructured meshes and thus is able to take into account arbitrary geometries more accurately than other techniques. The formulation of the physical problem in terms of the EM potentials allows us to employ nodal finite elements. Our program supports different types of elements which makes it easy to shape very irregular and complex geometries.

The resultant FE matrix is a $4N \times 4N$ complex block matrix (N is the number of mesh nodes, each of which has 4 degrees of freedom) composed of 4×4 sub-matrices, $\mathbf{K}_{ij}$ $(i, j = 1, \ldots, N)$, of the following form:

$$\mathbf{K}_{ij} = \begin{pmatrix} \gamma_{ij}\mathbf{I}_{33} & k\int_\Omega \sigma\phi_i\nabla\phi_j \,\mathrm{d}v \\ [k\int_\Omega \sigma\phi_j\nabla\phi_i \,\mathrm{d}v]^T & -k\int_\Omega \sigma\nabla\phi_i\cdot\nabla\phi_j \,\mathrm{d}v \end{pmatrix} \quad (4)$$

where $\mathbf{I}_{33}$ is the 3×3 identity matrix and γ_{ij} is the scalar function:

$$\gamma_{ij} = -\int_\Omega \nabla\phi_i \cdot \nabla\phi_j \,\mathrm{d}v + k\int_\Omega \sigma\phi_i\phi_j \,\mathrm{d}v \quad (5)$$

Considering that this resultant matrix is large, sparse, complex and non-Hermitian, we have implemented three different right-preconditioned Krylov subspace methods that can handle a system with such characteristics: the biconjugate gradient stabilised (BiCGStab), quasiminimal residual (QMR) and generalised minimal residual (GMRES).

PARALLEL IMPLEMENTATION

Our code works under the $Alya$ system (Houzeaux, Vázquez, Aubry, & Cela, 2009). The parallelisation strategy is based on the domain decomposition (mesh partitioning) technique using the Message Passing Interface (MPI) programming paradigm for communication among computational units. In addition to this, we use OpenMP for parallelisation inside of each computational unit. In this way, we have created one powerful hybrid parallel scheme which accelerates the execution of the forward-problem code to a great extent.

The original problem domain, which normally has a great number of elements, is divided into smaller sub-domains. Thanks to this partitioning, many computations can be done concurrently, which substantially reduces the total execution time of the program. In order to perform calculations simultaneously, it is necessary to create one execution process for each sub-domain. Having the sub-domains and the corresponding parallel tasks, the parallel execution is performed in the following way:

- Assembly: Assemble the local matrix and RHS in each sub-domain
- Algebraic solver:
  - Exchange and add local RHS contributions on the interface nodes using `MPI_Sendrecv` so that each sub-domain has the global RHS value on the interface nodes.
  - Perform matrix-vector products locally and then exchange and add contributions on the interface nodes using `MPI_Sendrecv`.
  - Perform scalar products locally and then assemble and sum contributions using `MPI_Allreduce`.

ALGEBRAIC MULTIGRID APPLIED AS PRECONDITIONING

AMG techniques can be very good preconditioners due to the fact that they efficiently reduce all error components – from low-frequency to high-frequency ones. Also, they do not need any geometric information and thus can be used as black-box preconditioners with different iterative schemes.

We have implemented our right-preconditioning AMG scheme without explicit knowledge of the preconditioning matrix $\mathbf{M}$. For each matrix-vector multiplication,

$\mathbf{A}'\mathbf{p} = \mathbf{A}\mathbf{M}^{-1}\mathbf{p}$, instead of doing an explicit multiplication $\mathbf{M}^{-1}\mathbf{p} = \mathbf{z}$, we solve a system $\mathbf{A}\mathbf{z} = \mathbf{p}$ using AMG:

- Do p basic relaxations on the original system, $\mathbf{A}\mathbf{z} = \mathbf{p}$, with the initial guess $\mathbf{z}_0 = 0$, and get an approximation $\mathbf{z}_p$.
- Find the residual: $\mathbf{r}_p = \mathbf{p} - \mathbf{A}\mathbf{z}_p$.
  - Project $\mathbf{r}_p$ on a coarse space: $\tilde{\mathbf{r}}_p = \mathbf{W}^T\mathbf{r}_p$.
    - Solve a coarse-space residual system: $\tilde{\mathbf{A}}\tilde{\mathbf{e}}_p = \tilde{\mathbf{r}}_p$.
  - Project back $\tilde{\mathbf{e}}_p$ on the fine space: $\mathbf{e}_p = \mathbf{W}\tilde{\mathbf{e}}_p$.
- Correct the fine-space approximation: $\mathbf{z}_p = \mathbf{z}_p + \mathbf{e}_p$
- Do p basic relaxations on the original system, $\mathbf{A}\mathbf{z} = \mathbf{p}$, with the initial guess $\mathbf{z}_0 = \mathbf{z}_p$, and get the final, i.e. preconditioned, approximation.

As a smoother, we can use one of three basic relaxation methods that we have implemented: Jacobi, Gauss-Seidel and symmetric successive over-relaxation (SSOR). The parameter p is the number of basic iterations, which is usually 1, 2 or 3. For transferring vectors between the fine and a coarse space, we use the projection matrix $\mathbf{W}$. Matrix $\mathbf{W}$ is also employed to build the coarse-system matrix $\tilde{\mathbf{A}} = \mathbf{W}^T\mathbf{A}\mathbf{W}$. Considering that our linear system of equations is mesh-based, the projection matrix $\mathbf{W}$ is defined by using the points of the mesh, which are divided into groups with the wave-front algorithm. Each group of the fine-mesh nodes is represented by one variable in the coarse space. Since the coarse residual system normally has a very small dimension (in the order of $10^2 - 10^3$), we solve it with a direct method based on LU factorisation.

NUMERICAL RESULTS

We present one of the models for which we have tested the accuracy and efficiency of our method. This model is a large synthetic test case that we created to be realistic in order to demonstrate the possibilities of our FE approach. X-Z slice of the model is shown in Fig. 1 (top). Its dimensions are $15 \times 12 \times 6$ km. The subsurface has 5 different conductivity structures with σ_h varying from 0.1 to 1.5 S/m and σ_v varying from 0.12 to 1.0 S/m. The hydrocarbon reservoir is located $\sim$1000 m below the seabed and has the conductivity $\sigma_h = \sigma_v = 0.01$ S/m. The transmitter is a horizontal electric dipole operating at the frequency of 1 Hz. In order to accurately represent the complex geology and sea-floor bathymetry, we have created a mesh which has 3 million elements and 0.5 million nodes.

Simulations with and without the reservoir have been performed. Normalised secondary electric and magnetic field magnitudes along an in-line profile are presented in Fig. 1 (bottom).

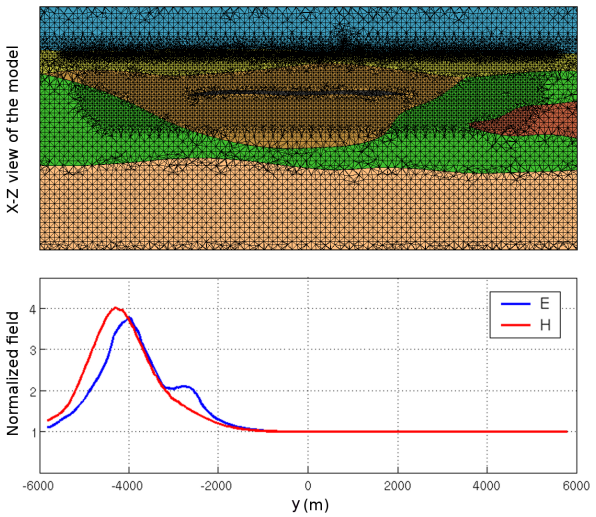

Figure 1. Model and reservoir responses normalised by the background results (without the reservoir).

In all the following experiments, we have used the right-preconditioned BiCGStab method to solve resultant linear systems. The convergence criterion for all BiCGStab iterations is a reduction of the relative residual norm to a value in the order of 10^{-8}. Also, the number of iterations has been limited by the maximum value of 1000. All the simulations have been carried out on the MareNostrum III supercomputer (http://www.bsc.es).

Scalability test

The scalability of the code has been tested by running the same problem for different numbers of CPUs working in parallel. For this test we have used the automatic mesh refinement of the second level that is built in our code to create a very large mesh, which has 191.7 million elements and 32.1 million nodes (128.4 million degrees of freedom). We have measured the total time spent on the construction of the matrix and on solving the system.

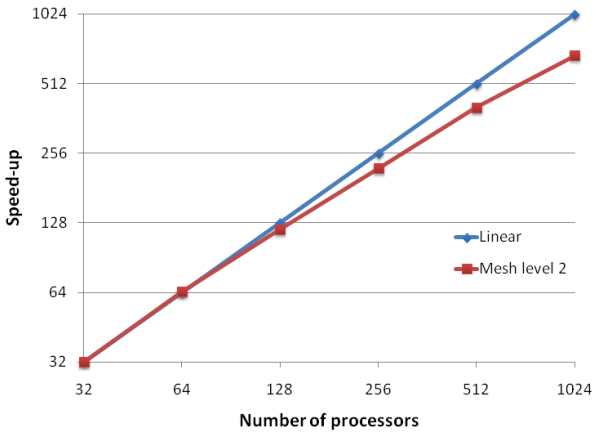

Figure 2. Scalability test for the mesh with two levels of refinement (128.4 million degrees of freedom).

Fig. 2 shows the speed-up obtained for up to 1024 CPUs. The achieved speed-up is close to linear for up to 128 processors. From this number on, the speed-up stops its near-linear growth and slowly begins to saturate since the execution becomes dominated by exchange of messages between the processes. However, the speed-up keeps growing constantly and we observe significant improvements in execution time for more than thousand processors.

AMG efficiency test

To inspect the benefits of our new preconditioning scheme, we have compared the performances of several preconditioning strategies:

- standard Jacobi (diagonal)
- Truncated approximate inverse (TAI), a cheap problem-based preconditioner that converges several times faster than Jacobi
- AMG with Jacobi smoother (AMG-J)
- AMG with SSOR smoother (AMG-SSOR)

The AMG preconditioner has been tested with different parameters:

- number of basic iterations at the beginning and at the end of the AMG procedure: 1, 3
- number of groups, i.e. size of the coarse system: 500, 1000, 5000

In Fig. 3, we compare convergence of the solver for different preconditioning schemes.

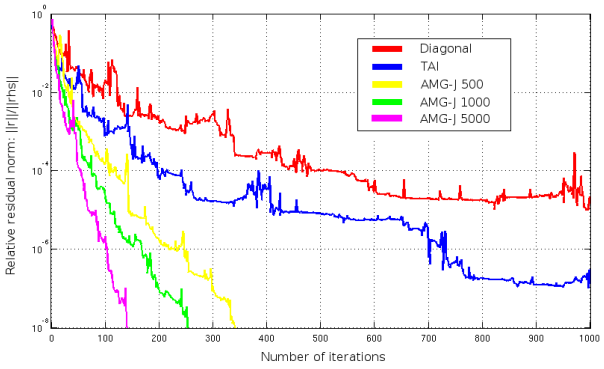

Figure 3. Convergence of BiCGStab with different preconditioners for the original problem (2 million degrees of freedom).

Fig. 3 shows that BiCGStab converges to the desired precision in less than 1000 iterations only when using AMG preconditioning. The results have shown that any reasonable combination of AMG parameters gives good convergence, so, in Fig. 3, we present the results obtained by AMG-J with 1+1 basic iterations, which is computationally the least demanding version of the preconditioner. For this case, the shortest total execution time is 37 seconds on 32 CPUs.

ACKNOWLEDGMENTS

The authors thank Repsol for funding provided for this work.

REFERENCES

Aruliah, D., & Ascher, U. (2002). Multigrid preconditioning for Krylov methods for time-harmonic Maxwell's equations in three dimensions. *SIAM Journal on Scientific Computing*, *24*(2), 702–718.

Badea, E., Everett, M., Newman, G., & Biro, O. (2001). Finite-element analysis of controlled-source electromagnetic induction using Coulomb-gauged potentials. *Geophysics*, *66*(3), 786–799.

Commer, M., & Newman, G. (2004). A parallel finite-difference approach for 3D transient electromagnetic modeling with galvanic sources. *Geophysics*, *69*(5), 1192–1202.

Haber, E. (2004). Quasi-Newton methods for large-scale electromagnetic inverse problems. *Inverse Problems*, *21*(1), 305.

Hou, J., Mallan, R., & Torres-Verdín, C. (2006). Finite-difference simulation of borehole em measurements in 3d anisotropic media using coupled scalar-vector potentials. *Geophysics*, *71*(5), G225–G233.

Houzeaux, G., Vázquez, M., Aubry, R., & Cela, J. (2009). A massively parallel fractional step solver for incompressible flows. *Journal of Computational Physics*, *228*(17), 6316–6332.

Key, K., & Ovall, J. (2011). A parallel goal-oriented adaptive finite element method for 2.5-d electromagnetic modelling. *Geophysical Journal International*, *186*(1), 137–154.

Mulder, W. (2006). A multigrid solver for 3D electromagnetic diffusion. *Geophysical Prospecting*, *54*(5), 633–649.

Saad, Y. (2003). *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics.

Sasaki, Y., & Meju, M. (2009). Useful characteristics of shallow and deep marine csem responses inferred from 3d finite-difference modeling. *Geophysics*, *74*(5), F67–F76.

Trottenberg, U., Oosterlee, C., & Schüller, A. (2001). *Multigrid*. Academic Pr.

Weiss, C., & Newman, G. (2003). Electromagnetic induction in a generalized 3D anisotropic earth, Part 2: The LIN preconditioner. *Geophysics*, *68*(3), 922–930.