Surv Geophys (2014) 35:85–100
DOI 10.1007/s10712-013-9260-0

A Review of High-Performance Computational Strategies for Modeling and Imaging of Electromagnetic Induction Data

Gregory A. Newman

Received: 24 September 2012 / Accepted: 23 October 2013 / Published online: 29 November 2013

Abstract Many geoscientific applications exploit electrostatic and electromagnetic fields to interrogate and map subsurface electrical resistivity—an important geophysical attribute for characterizing mineral, energy, and water resources. In complex three-dimensional geologies, where many of these resources remain to be found, resistivity mapping requires large-scale modeling and imaging capabilities, as well as the ability to treat significant data volumes, which can easily overwhelm single-core and modest multicore computing hardware. To treat such problems requires large-scale parallel computational resources, necessary for reducing the time to solution to a time frame acceptable to the exploration process. The recognition that significant parallel computing processes must be brought to bear on these problems gives rise to choices that must be made in parallel computing hardware and software. In this review, some of these choices are presented, along with the resulting trade-offs. We also discuss future trends in high-performance computing and the anticipated impact on electromagnetic (EM) geophysics. Topics discussed in this review article include a survey of parallel computing platforms, graphics processing units to multicore CPUs with a fast interconnect, along with effective parallel solvers and associated solver libraries effective for inductive EM modeling and imaging.

Keywords Three-dimensional electromagnetic modeling and inversion · Magnetotelluric soundings · Electromagnetic induction · High-performance computing · Parallel solvers · Resistivity imaging for the Earth

G. A. Newman (✉)
Lawrence Berkeley National Laboratory, MS 90-1116, Berkeley, CA 94720, USA
e-mail: ganewman@lbl.gov

1 Introduction

At the dawn of the high-performance computing (HPC) era,[1] (late 1980s to the mid-1990s), fast approximate methods were favored among electromagnetic (EM) modelers, such as the extended Born technique of Habashy et al. (1993) for modeling and extended by Torres-Verdin and Habashy (1994, 1995) for inversion. Other variants of these techniques included the quasi-linear approximation for EM modeling (Zhdanov and Fang 1999) and approximate sensitivities for inversion (Farquharson and Oldenburg 1996), along with rapid relaxation imaging (RRI) for magnetotelluric (MT) data (Smith and Booker 1991). While the aforementioned techniques were fast and often produced answers in close agreement with more rigorous integral equation (IE) and finite difference (FD) schemes, they could also produce erroneous results. This resulted in an ultimate lack of confidence in these methods within the EM community and their eventual abandonment.

Finite difference (FD) modeling schemes made their appearance in MT applications in the early to mid-1990s (Mackie and Madden 1993; Smith 1992, 1996a, b). FD controlled source time and frequency domain EM solutions were also developed during this time (Wang and Hohmann 1993; Newman and Alumbaugh 1995). These schemes allowed for accurate simulations of significantly greater model complexity. Such models could not be effectively simulated with integral equation (IE) techniques (cf. Hohmann 1975; Wannamaker et al. 1984; Newman et al. 1986) or with the approximate methods already mentioned, due to accuracy issues. Nevertheless, the size of models that could be treated using FD modeling techniques was still too limiting on single CPU machines, due to excessive computation times and memory restrictions. Newman and Alumbaugh (1995) reported that simulating an airborne EM profile for 29 source positions over a simple three-dimensional (3D) dyke model required nearly 14 h of CPU time and 120 MB of memory on a single-processer, top-of-the-line, IBM RS 6000 Workstation. This represented a significant demand on machine resources at that time. Recognition of such limitations provided significant motivation to migrate FD EM modeling and imaging schemes to massively parallel (MP) computing platforms.

Application of HPC systems to the solution of large-scale EM geophysical studies first appeared with the pioneering works of Alumbaugh et al. (1996), Alumbaugh and Newman (1997) and Newman and Alumbaugh (1997, 1999). Exploitation of massively computing resources showed the possibilities for solving EM problems at scales with realistic 3D geologies, within days. Realizing that accurate solutions of the FD formulations of the EM field equations, using HPC resources, could be obtained, within a day, for realistic, large-scale, earth models and 3D data acquisition geometries, modelers had increasingly less use for approximate schemes (as well as integral equation schemes) over the intervening years (to the present). Within the last several years, parallel finite element (FE) solutions in two-dimensions (2D) and (3D) have started to appear in the scientific literature; we refer the interested reader to the recent works of Franke et al. (2007), Börner et al. (2008), Schwarzbach et al. (2011), Key and Ovall (2011), Vieira da Silva et al. (2012), Schwarzbach and Haber (2013), and Puzyrev et al. (2013). The appeal of FE solutions compared to FD solutions is more accurate treatment of topography and bathymetry, because unstructured FE solutions formulated on unstructured grids can be refined to conform to geological interfaces. Nevertheless, one significant problem for 3D FE schemes is that of robust grid

[1] I consider the beginning of the HPC era circa 1988, with the development of multiple instructions multiple data (MIMD) asynchronous computing architectures (cf. Fox 1988).

generation. This is still an open research problem (cf. Um et al. 2010), along with the need fast and robust iterative solvers (cf. Um et al. 2013).

The application of large-scale 3D EM and MT imaging experiments continues to grow. Driven in large part by the worldwide hunt for energy resources, EM and MT integrated with seismic data are being used to better define hydrocarbon resources. Marine controlled-source electromagnetic (CSEM) imaging has become a mature industry in its own right. The CSEM method is sensitive to reservoir fluids and saturation, and thus may indicate hydrocarbons directly, which makes it a source of data complementary to seismic data (cf. Constable 2006; Eidesmo et al. 2002; Ellingsrud et al. 2002; MacGregor et al. 2006). Another very important observation is that two common, problematic seismic lithologies in hydrocarbon exploration (using seismic methods) are salt and basalt. Both materials have fast seismic velocities and scatter the seismic energy, often resulting in poor seismic data resolution beneath them (cf. Hoversten et al. 2000; Maresh and White 2005). Energy exploration companies have realized the potential of EM and MT as additional data sources to aid in determining the geological structure beneath these problematic units (cf. Jegen et al. 2009; Colombo et al. 2012).

With these new industrial scale applications, successfully extracting and processing of information from CSEM and MT data has proven to be a formidable computational problem. To provide geologists with maximally consistent electromagnetic data, interpretation requires dense survey coverage and multicomponent data volumes. As a consequence, large-scale 3D imaging is receiving considerable attention in the interpretation of CSEM and MT data (Commer et al. 2008; Commer and Newman 2008, 2009; Carazzone et al. 2008, 2005; Gribenko and Zhdanov 2007; Newman and Commer 2009; Plessix and Mulder 2008; Plessix and van der Sman 2007, 2008; Zach et al. 2008). Here, exploitation of HPC systems is essential to effectively treat the resulting large-scale 3D imaging problems that arise in industrial applications on an acceptable time scale of a few days to a week (Fig. 1).

The recent arrival of new HPC technologies is enabling new approaches to parallelization of EM modeling and imaging schemes. Modern graphics processing units (GPUs), designed for efficiently manipulating computer graphics, have a high parallel architecture (hundreds of independent computational threads on each GPU) that also makes them suitable for computationally intensive scientific applications. Recent inquiries into GPU performances for EM applications can be found in the works of Commer et al. (2011) and Weiss and Schultz (2011). Later in this review, we discuss some performance comparisons between multicore CPU with message–passage-interface (MPI) and GPUs. Clearly, the field of HPC is rapidly evolving. In a decade, the HPC architectures we employ for our problem-solving may well be vastly different, thus demanding new program and parallelization paradigms. To better understand where HPC is headed for our problems, it is important to discuss platforms and software currently in use for EM modeling and imaging. It is also important to discuss the parallel computational approaches for our applications, including multiple levels of parallelization, and the choices one makes regarding software and solver libraries.

2 Effective 3D Solvers for Electrical Resistivity Modeling

Before discussing HPC implementations for 3D EM and DC resistivity modeling, it is useful to review the solvers for the Maxwell's equations in use today. Such solvers form the computational kernel of HPC implementation. They apply to a large class of

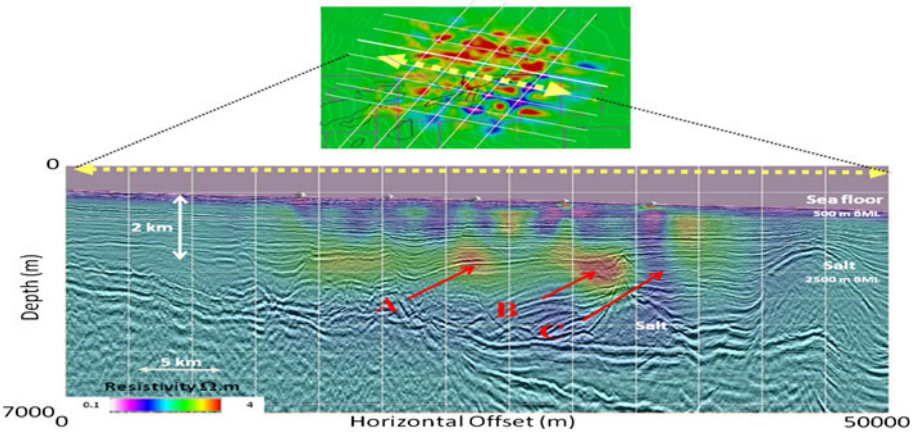

Fig. 1 Illustrated is an industrial size example of large-scale CSEM imaging experiment, processed with over 1,000 compute cores, offshore Campos Basin, Brazil. Rendered at the top is the average vertical resistivity map from 500 to 2,500 m below the seafloor, superimposed with sail lines used to acquire the Campos Basin data. The cross-section at the *bottom* shows the vertical resistivity image along the indicated transect. The Campos Basin experiment demonstrated the necessity to incorporate electrical anisotropy into the imaging processes for accurate results. The EM image is shown together with seismic reflection horizons. Anomaly A is related to enhanced resistivity due to a known oil field. Anomaly B is enhanced resistivity and may indicate a possible hydrocarbon trap above a large salt body. Enhanced conductivity at C is likely to be related to conductive brines originating from salt below. Results presented by Carazzone et al. (2008) and Newman et al. (2011)

geoscientific applications that involve boundary value problems for simulating EM and MT fields for geophysical prospecting and subsurface imaging of electrical resistivity. Inductive EM modeling can be done in either the frequency or time domain. In the frequency domain, FD and FE implementations for the CSEM and MT modeling give rise to a large, sparse linear system resulting from the discrete versions of the Maxwell equations for complex 3D geological media; IE implementation also gives rise to a linear system, but it is dense and not that practical for complex 3D modeling investigations. Specifically, the linear system of N rows and columns can be written as

$$\boldsymbol{A}_{N\times N}\boldsymbol{x} = \boldsymbol{b}. \tag{1}$$

In the time domain, explicit time-stepping solutions to the Maxwell equations work well. The explicit discrete forms of the equations in the time domain also hide an underlying sparse matrix operator, which, at every time step is equivalent to a matrix–vector product involving the electric or magnetic fields to be advanced to the next time step. In (1), the matrix $\boldsymbol{A} \in \Re^{N\times N}$ is regular (sparse or dense), and the solution and right-hand side vectors are $\boldsymbol{x}, \boldsymbol{b} \in \Re^{N}$, respectively. Direct methods can be used to factor the matrix (Cholesky and LU factorizations for symmetric and non-symmetric matrices; in resistivity applications, the matrices are symmetric [DC resistivity methods] or complex symmetric [EM and MT soundings]). Upon factorization, solution to multiple right-hand sides can be rapidly calculated. However, beyond a certain problem size, depending upon the given computational resources and acceptable time to solution, usage of direct solvers becomes prohibitive. Thus, in these instances, iterative Krylov subspace techniques are commonly used for solution.

Krylov subspace methods are defined as projection (Galerkin) or generalized projection (Petrov–Galerkin) methods for the solution of the linear system (1). The solution involves constructing the Krylov subspace $\boldsymbol{K_m}$,

$$\boldsymbol{K_m} = \boldsymbol{K_m}(\boldsymbol{A}, \boldsymbol{r}_0) = \mathrm{span}\left\{\boldsymbol{r}_0, \boldsymbol{A}\boldsymbol{r}_0, \ldots, \boldsymbol{A}^{m-1}\boldsymbol{r}_0\right\} \tag{2}$$

Starting with the residual vector, $\boldsymbol{r}_0 = \boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_0$, Krylov methods compute the approximation $\boldsymbol{x_m} \in \boldsymbol{x}_0 + \boldsymbol{K_m}$ to the solution of (1) in an iterative manner, where at each iteration the dimension m of $\boldsymbol{K}$ is updated; in exact arithmetic and for symmetric positive definite systems, the iteration will produce an exact solution to $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$. However, because of round-off errors in machine computation, this property is not realized in practice. For nonsymmetric and complex symmetric linear systems, there is no guarantee of convergence with the Krylov iteration, but in practice, reasonable convergence has been observed for EM problems (cf. Newman and Alumbaugh 1995, 2000). Krylov methods are named after the Russian applied mathematician and naval engineer Krylov (1931), whose paper formed the basis of all Krylov methods later developed (http://en.wikipedia.org/wiki/Krylov_subspace).

A common near-surface application is the DC resistivity method, where a DC (or a very low frequency) current is introduced as a means of studying earth electrical resistivity, for example in groundwater mapping. Here, depth of interrogation is controlled by source-receiver offsets of the electrodes and earth resistivity. Readers are referred to the works of Spitzer (1995) and Spitzer and Wurmstich (1999) for further details on DC resistivity modeling.

Time-harmonic EM prospecting methods, including magnetotellurics, are methods in which depth of exploration and spatial resolution can be further controlled with source frequency. The Krylov solvers of interest for these applications are designed to handle linear systems where $\boldsymbol{A}$ is real symmetric or complex symmetric. For the real symmetric case, the conjugate gradient (CG) method is optimal (Hestenes and Stiefel 1952). CG is designed to be applied to any Hermitian (complex) linear system that is symmetric positive definite; the DC resistivity modeling problem is an example where such systems may arise. For the complex symmetric case, where the matrix is non-Hermitian, the bi-conjugate gradient (BiCG) method first proposed by Lanczos (1952) and quasi-minimum residual (QMR), more recently proposed by Freund and Nachtigal (1991) and Freund (1992), are effective and cost-efficient solvers. Another good choice for complex symmetric as well as nonsymmetric problems is BiCGSTAB, which was introduced by van der Vorst (1992). The solvers BiCG, QMR, and BiCGSTAB are appropriate for simulating complex EM and MT fields.

Krylov iteration also provides the basis for more sophisticated solver techniques that can be applied to the solution of 3D geo-electromagnetic modeling problems. The spectral Lanczos decomposition method (SLDM) (cf. Druskin and Knizhnerman 1994; Druskin et al. 1999) and Padé-type approximations used by Börner et al. (2008) allow for rapid solution of 3D EM fields in either the time or frequency domain, based on model reduction. The essential idea with reduced order modeling is that only a modest number of eigenvalues and eigenvectors are necessary to accurately describe the simulated field over a range of frequencies and decay times. These singular values and vectors can be estimated quickly with a Krylov subspace projection. The larger the Krylov subspace, the greater the range of frequencies and decay times that can be efficiently simulated. Finally, multigrid methods that exploit Krylov iteration can be extremely efficient in solving linear systems (Greenbaum 1997). In this context, the multigrid method acts as a preconditioner for the

Krylov subspace methods. Application of multigrid methods for DC modeling problems can be found in the works of Lu et al. (2010) and Moucha and Bailey (2004).

3 Parallel Computing Platforms

3.1 The Multicore CPU Platform

Massively parallel (MP), multiple instruction multiple data (MIMD) machines have been the standard platform for HPC computation for nearly the last 24 years. These machines have dedicated access to tens to ten thousands of compute cores. The smallest MIMD machines are typically clusters of a few tens to several hundred compute cores. The largest machines are to be found in supercomputing centers around the world, such as the National Energy Research Scientific Computing (NERSC) center, which is a division of Lawrence Berkeley National Laboratory, located in Berkeley California. As an example of a top-tier machine, consider NERSC's first peta-flop system, Hopper, which is a Cray XE6. It has peak performance of 1.28×10^{15} floating point operations per second,[2] 153,216 compute cores, 217 TB of memory, and 2 PB of disk. Large and small MIMD machines rely on a dedicated backbone for communication among the computing cores. The standard message passing interface (MPI) (Gropp et al. 1996) is used for interprocessor communication. MPI provides portability so that MP software can be run across a range of MIMD platforms, including dedicated distributed machines and/or a distributed network of machines. The works of Alumbaugh et al. (1996), Alumbaugh and Newman (1997), Newman and Alumbaugh (1997, 1999), and Newman and Alumbaugh (2000) were some of the first to exploit the MIMD-MPI computing architectures for 3D EM modeling and inversion, with some of the more recent works including those of Chen et al. (2012a, b), Oldenburg et al. (2013) and Puzyrev et al. (2013).

Fast performance on MIMD machines depends upon the speed at which messages can be delivered through the backbone fabric that connects all the compute cores. One of the largest MP applications for EM modeling is the 3D CSEM imaging experiment reported by Commer et al. (2008), in which 32,768 compute cores were employed. While this application required extensive resources to execute, it was a demonstration that 3D EM imaging problems could be solved in days, rather than weeks, on more modest size clusters. It offered important verification that industrial size 3D EM imaging problems could exploit fine- and coarse-grain parallelism to achieve solutions on a time scale acceptable to energy exploration companies. At the fine-grain level of parallelization, a domain decomposition is employed for each source and frequency over a fixed number of compute cores; while at the coarse-grained level, multiple copies of the domain decomposition are employed for the different sources and frequencies, which is perfectly scalable, resulting in an embarrassingly parallel computation.

3.2 GPU Processors

To provide access to the multithreaded GPU computational resources and their associated memory bandwidth graphics, hardware manufacturers have introduced new application programming interfaces such as CUDA and Open CL, thus enabling numerical calculations in a fashion similar to the MIMD parallel computing paradigms discussed above. Since

[2] One petaflop is equivalent to 10^{15} operations per second.

2006, GPU technology has evolved rapidly for scientific applications (http://en.wikipedia.org/wiki/Graphics_Processing_Unit), and the technology is now used extensively in processing and imaging of seismic data in oil and gas exploration. For EM induction problems, Commer et al. (2011) and Weiss and Schultz (2011) provide studies on the feasibility of using a single GPU for large-scale EM field simulations. (I will discuss these comparisons between GPUs and MIMD-MPI computing platforms at length, later in this review.) When GPUs come with significant memory access (3–5 GB), it is possible to simulate reasonable-sized EM field calculations on a single GPU, but memory can be an issue with larger calculations. An obvious strategy for treating larger problems is to distribute the calculations across multiple GPUs. Unfortunately, a current bottleneck with the GPU architecture is the performance hit observed with distributing calculations across multiple GPUs. Such data transfers are not nearly as straightforward as using a pure MPI-CPU interconnect. This is an area of active research; the interested reader is referred to the works of Yang et al. (2011), Lawlor (2009), and Stuart and Owens (2009).

A new trend that has emerged recently is the pairing of CPUs with MPI interconnects, with GPU accelerators, avoiding direct message passing between GPU accelerators. Instead, parallelization across multicore CPUs is combined with fine-grain parallelization of the GPU on the CPU for much faster scientific computations (approximately 5- to 10-fold). As an example, consider a domain decomposition that is distributed across the CPUs, with internal calculations on each CPU passed to the GPU for accelerated computation without significant communication overhead between CPU and GPU. Since communication or message passing is carried out extensively on the coarse-grain level between CPUs and infrequently on the GPUs, much faster times to solutions can be realized with careful load balancing. In late 2012, utilizing this type of strategy, the Titan supercomputer at Oak Ridge National Laboratory in the US—consisting of 18,688 Nvida k20x GPUs and 18,688 AMD Opetron CPUs with 560,640 processor cores—reported a performance of nearly 18 petaflops on the standard Linpack benchmark, making it the world's fastest supercomputer. Peak performance for Titan is 20 petaflops or a quadrillion calculations per second.

4 Parallel EM Applications

Let us now discuss the range of applications in which HPC computing has impacted EM investigations. The two applications of EM where HPC has made an impact and will continue to make impacts are in large-scale 3D modeling and inversion studies.

4.1 Large-Scale 3D Modeling and Inversion

The size and complexity, and thus the realism, of 3D models that can be simulated on traditional serial computers are limited by memory and flop-rate of the processor. However, with HPC, the rate at which the simulations can proceed is dramatically increased, because thousands to tens of thousands of processors can operate on the problem simultaneously. Examples of EM modeling problems using finite difference techniques that exploit such resources can be found in Newman and Alumbaugh (1999) for airborne 3D EM simulations, Newman and Commer (2005) for 3D transient EM (TEM) modeling and inversion, and Commer et al. (2008) and Newman et al. (2010) for 3D inverse modeling of a deep-water EM exploration survey in the Campos Basin, offshore Brazil. The airborne EM and marine CSEM examples can be quantified as requiring hundreds of sources in the simulation experiment along with large-scale model parameterizations. Other examples can be found in the literature (cf. Plessix and Mulder 2008 and Oldenburg et al. 2013). Typically, the 3D imaging solutions use a Cartesian mesh for the model parameterization. Recently, however, Schwarzbach and Haber (2013) completed an unstructured finite element solution to the 3D CSEM imaging problem, which could allow for more effective imaging in the presence of bathymetry and topography.

Calculations that attempt to include all the physics in the simulation are highly desirable, for several reasons. First, they provide benchmarks on faster approximate solution methods that explicitly ignore some of the salient physics of EM field simulation and inverse modeling (the Born Approximation, for example), indicating their range of validity and applicability. Note that all numerical methods involve some sort of approximations, but these errors can be adjusted to acceptable level by mesh refinement. What I am speaking about here are approximations whose errors cannot be reduced by higher-order meshing and derivative and quadrature approximations. Second, the computational end-member numerical solutions I speak of are the gold standard in EM modeling. They include model complexities, along with a complete description of the physics of the EM field that cannot be simulated otherwise, especially since analytical solutions do not exist for these problems. Provided that adequate HPC resources are available, such end-member calculations should always be considered.

4.2 Stochastic Imaging

Bayesian stochastic framework and Markov Chain Monte Carlo sampling methods may offer (in some cases) a clear alternative to deterministic inversion methods for estimating the resistivity of the subsurface along with the associated uncertainties. The papers of Chen and Dickens (2009) and Chen et al. (2007) provide a good overview of the methods for reservoir estimation problems, and a more recent work by Chen et al. (2012) applies the stochastic inversion to MT data using a sharp 2D boundary parameterization to a geothermal field site in Indonesia. For DC problems, Schwarzbach et al. (2005) applied a genetic algorithm (also considered a stochastic imaging approach) to 2D resistivity data. Liu et al. (2012) reported some initial results using a modified genetic algorithm allowing for mutation in the sampling strategy of resistivity models with modest size (<15,000 parameters), relevant to small-scale 3D hydrological investigations. Reported solution times were much slower than the traditional deterministic imaging in order to eliminate dependence on the initial resistivity model.

With stochastic sampling, the likelihood function is the link between the resistivity and EM/MT/DC data. It is clearly an embarrassingly parallel application that can be used to construct the joint posterior probability function defined by the Baysesian stochastic framework. Moreover, because each sample of the likelihood function requires the solution of a specific EM or MT boundary value problem, the solution of such a problem can be efficiently obtained by distributing it over a number of parallel computing tasks (cf. Chen et al. 2012a, b). Hence, multiple levels of parallelization can be exploited. Nevertheless, the advantages of stochastic estimation come at considerable computational expense—hence the need for HPC resources. Currently, stochastic sampling methods are viable for simplified layered and 2D geologies. However, the case for stochastic parameter estimation for realistically parameterized three-dimension problems is a different matter. Global-stochastic estimation methods are still too time and memory intensive to be of practical use in the foreseeable future.

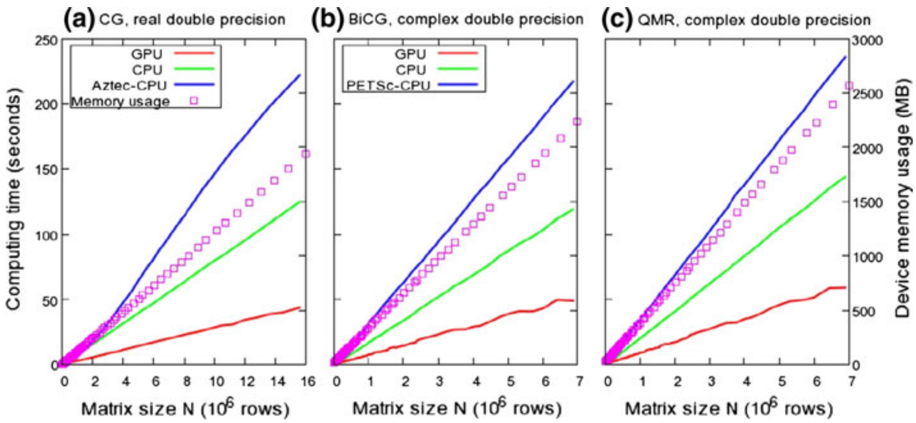

Fig. 2 Solution times for three different iterative Krylov solvers, CG (**a**), BiCG (**b**), and QMR (**c**); figure replicated from Commer et al. (2011). Each GPU implementation is compared against its original parallel FORTRAN90 version, as well as a parallel solver from an external library. All CPU solvers were run on eight processor cores. The solid lines correspond to the *left axis* and show computing times in seconds required for 1,000 iterations. Memory usage is shown on the *right y axis* and denote the requirements for storing all Krylov solver data structures in memory. The *curves* denoted as Aztec-CPU (CG) and PETSc-CPU (BiCG and QMR) are reported solution times using the Aztrec and PETSc external parallel solver libraries. GPU simulations were carried out on NVIDIA Tesla C2050 (Fermi) GPU, which includes three gigabytes of memory and 448 parallel CUDA processor cores. CPU only computations were preformed on two Intel Nehalem 5530-2.4 GHz CPU with eight processor node. Each node consists of two CPUs

4.3 Parallel Computational Approaches for EM Modeling and Inverse Problems

As has already been remarked, computational approaches to 3D EM and MT modeling now favor finite difference and finite element implementations, because of their flexibility in modeling complex geologies. Another consideration is the size of the data volume to be analyzed, which can include hundreds of sources and thousands of detectors per source. Successful approaches in dealing with large-scale model parameterization and data volumes exploit multiple levels of parallelization. For example, Commer et al. (2008, 2009) use a domain decomposition of the model across a subset of computational tasks, called a “local processor group” and distribute the calculation across multiple copies of these groups, called “data planes.” Because the data calculations are independent on each data plane, results are embarrassingly parallel and scale with the number of data planes employed. Implicit solver methods take advantage of this processor topology, where iterative Krylov solvers are ideal. Here, each source calculation is independent of other sources, and copies of the solvers can be distributed across the different data planes; a similar strategy can be used for explicit time-stepping TEM modeling approaches for multiple sources. Recently, there has been interest in multifrontal direct solvers in large-scale EM simulations (cf. Vieira da Silva et al. 2012; Streich 2009). The appeal of these solvers is that, once the linear system operator is factored, solutions for multiple right-hand sides can be computed quickly. Popular, parallel multifront solvers libraries include MUMPS (Amestoy et al. 2001, 2006), Super LU (Li and Demmel 2003), and PARDISO (Schenk and Gärtner 2004, 2006), among others. While there is much appeal in exploiting these parallel solvers, large model parameterizations and corresponding meshes that arise in CSEM applications limit their applicability to modest-sized problems. Airborne EM modeling is another limited application because of the size of the grids that can be treated.

Even though there could be thousands of sources, it is not practical to use a single mesh for all of them. Like Krylov solvers, copies of multifrontal direct solvers can also be implemented independently on different data planes, if sufficient memory is available. Thus, the fields arising from groups of sources on each data plane can be solved efficiently.

Practical solutions of the 3D EM/TEM/MT inverse problems use some variant of Newton, Gauss–Newton, and nonlinear/steepest decent techniques to arrive at a solution. It is best to avoid direct formulation of the Hessian and Jacobian and its transpose. These linear systems, which result from discretization of these operators, are dense, and the cost of their explicit formulation is considerable, even on a distributed computational platform. Instead, it is advisable to compute only the action of these discrete operators on a vector, which correspond to a matrix–vector multiplication. The net result is that two to three forward solves are necessary to complete the operation per source/frequency (cf. Newman and Hoversten 2000). As a consequence, the computational bottleneck in 3D imaging problems is the solution of the forward modeling problem. Therefore, the use of HPC resources, effective solver strategies, and multiple levels of parallelism are essential for large-scale problems.

5 Parallel Solver Libraries

Parallel solver libraries are an essential resource that non-computer scientists can take advantage of when developing their own parallel modeling and imaging algorithms. The alternative is to develop one's own parallel solver. In certain cases, this can be achieved, depending upon the solver type. Iterative Krylov solvers are not too hard to implement, and substantial benefits in performance can be obtained with them, compared to external parallel libraries (approaching a speedup of nearly a factor of two, as demonstrated below). However, it will be necessary to write the MPI communication routines required by the iterative solver. With explicit time-stepping methods, which arise in TEM modeling and inverse problems, it is not too difficult to develop similar communicators to forward and back-propagate EM fields. For direct solvers, however, this is not advisable, because of the level of complexity and sophisticated programming structures needed to achieve the desired efficiency (cf. Wang et al. 2012). Popular iterative Krylov solver libraries include the PETSE (Balay et al. 2010) and the TRILINOS-AZTEC (Heroux et al. 2003, 2005; Tuminaro et al. 1999) packages. These libraries were developed at United States Department of Energy Laboratories, with the explicit goal of helping software developers in implementing their applications on HPC machines. They come with the necessary communicators to carry out the required message passing, relieving the developer of the task of development, which is considered one of the hardest aspects of parallel programming. For direct solver libraries, as have already been mentioned, MUMPS, Super LU, and PARDISO are popular choices. It is outside the scope of this review to provide specific details on the relative merits of these direct solver libraries. However, selection of a specific solver will depend upon the available computing resources, particularly memory, algorithm efficiency, and the size of the problem to be solved. Streich (2009), Oldenburg et al. (2013), Swarzbach and Haber (2013), and Grayver et al. (2013) have reported success using MUMPS or PARDISO for 3D CSEM and time domain EM modeling and inversion. For extreme large-scale problems, typical in simulating seismic wave propagation, a new variant of Super LU (cf. Wang et al. 2011, 2012), is worth mentioning.

6 GPU-Multicore CPU Comparison Study

Commer et al. (2011) recently implemented efficient Krylov subspace methods for the iterative solution of sparse linear systems on GPUs. The solvers studied were suitable for electrical and electromagnetic simulation problems that arise in geophysical resistivity prospecting. Timing comparisons clearly indicated the increasing efficiency of the GPU solvers for increasingly larger matrix sizes. For the largest problems examined (16 and 7 million unknowns, respectively, for real and complex systems), the GPU performance is equivalent to 23 CPU compute cores of the CG iteration and 19 cores for both the BiCG and QMR iterations; the number of compute cores on a CPU typically varies anywhere from 4 to 12. The comparisons are for the faster CPU solvers that are custom designed and do not use external libraries (Fig. 2). A comparison of solution times for the Krylov solvers (CG, BiCG, and QMR), using only CPU compute cores with an MPI interconnect, clearly shows the computational efficiencies that can be obtained when carefully implementing solver communicators oneself, compared to those supplied by the external solver libraries (Azterc and PETSc).

While the GPU performance in these tests are clearly impressive, it is important to remember that one still needs to distribute the solution of a single forward problem across the order of 100 CPU cores or more, to address both computing and memory needs of systems arising from typical large-scale CSEM exploration applications—along with acceptable processing times. Even with considerations for grid optimization schemes to reduce the computational burden (cf. Commer and Newman 2008; Plessix and Mulder 2008), hundreds to thousands of CPU compute cores are a necessity for realistic 3D CSEM imaging experiments (cf. Newman et al. 2010). While Commer et al. (2011) have not yet fully explored a corresponding GPU approach, i.e., solving one system on multiple GPUs, currently there are constraints imposed by the amount of GPU memory, limiting the usage of GPU solvers to maximum grid sizes as reported here. Nevertheless, one can expect hardware and message improvements with GPUs, so the situation is likely to improve.

7 Future Trends: The Cloud

Computing is ubiquitous. Large server farms serviced by companies such as Amazon, Microsoft, and Google allow for the possibilities of carrying out modeling and inversion work in the Cloud, alleviating the need for the EM geophysicist to support in-house computational clusters and corresponding operating systems software. In a recent work, Mudge et al. (2011) discussed the evolution of inversion methods, focusing on MT, with Cloud computing. This work shows us new possibilities for distributed geophysical modeling and inversion computations. Nevertheless, there are some issues that need to be addressed before this approach to computation can really take off. First is security. If data cannot be made secure on remote server farms, industrial applications will not drive the shift away from in-house computing to that in the Cloud. Second, large-scale EM field simulation and inversion requires a dedicated interconnect between different processor cores. Mudge et al. (2011) report that MPI so far does not map well to the cloud architecture. Because MPI has become the most popular programming model for HPC, it will be necessary to improve its efficiency in the Cloud, particularly for computations that are tightly coupled, requiring fast and efficient message passing.

8 Conclusions

This review has discussed the importance of HPC on 3D modeling and imaging of earth resistivity properties, be it DC, EM, or MT resistivity soundings. In the 1980s, the most complex models that we could simulate were simple prismatic bodies in layered geological media. As software and hardware have progressed, over nearly 30 years, we have witnessed the maturation of EM modeling and inversion methods, which has allowed their application to complex 3D geologies once inconceivable. HPC resources have played a significant role in this advancement of EM modeling and imaging technologies. What will the next 30 years bring? With distributed cloud computing on the near horizon, assuming that the present limitations described above are overcome, the possibility exists for freeing EM geophysicists from supporting in-house computer clusters, associated software maintenance, and cumbersome supercomputer centers. The simple idea of uploading data and input into web-based 3D modeling and imaging applications, anywhere there is connectivity, has significant appeal, provided security and interconnect issues can be resolved. With calculations done remotely, securely, and quickly in the Cloud, solution times will contract, resulting in new breakthroughs in our EM modeling and imaging of the earth's subsurface, with important implications in the exploration of energy resources, environmental site characterization, and geological hazard studies.

High-performance computing (HPC) and scientific computation are inexorably linked. New algorithms and computational approaches are constantly being developed to allow for faster and more accurate large-scale EM modeling. Initially, these numerical approaches are developed using convenient scientific software framework, such as MATLAB, on a single-core/CPU processor for quick and easy proof of concept testing. On the other hand, efficient implementation on a distributed computing system requires careful considerations of algorithm scalability, along with memory considerations and message passing, along with the optimal approach for utilizing computing resources.

In closing, I believe the forefront HPC EM modeling and associated imaging research now focus on finite element solutions to Maxwell's equations and the associated DC equations on unstructured meshes. The linear systems that result from unstructured meshes are large, sparse, and highly ill-conditioned. Schwarzbach and Haber (2013) proposed using parallel direct solvers on these irregular meshes because they can be substantially smaller than the matrices produced by finite difference discretization on regular meshes. Nevertheless, efficient and scalable iterative solutions to these types of matrix systems are an area of active research (cf. Um et al. 2013, Puzyrev et al. 2013). While Krylov iterative solvers can be effective for these types of problems, efficient preconditioning is necessary. (Scalability of the solver is also an ongoing concern.)

Finally, it is interesting to note that algebraic multigrid (AMG) solvers have not received much attention for these types of problems. AMG is reportedly the optimal iterative solution method (Greenbaum 1997) that can be applied to sparse linear systems. While AMG (as well as geometric multigrid solvers) has been successfully developed for real systems (DC type problems—for example, Lu et al. 2010; Moucha and Bailey 2004), application to complex, and complex symmetric systems, have proven more difficult. One possible reason may be attempts to use AMG libraries designed for real systems on complex systems, which are expressed in terms of equivalent real forms (ERF). Interestingly, Freund (1992) warned that Krylov solvers designed for real systems, such as GMRES, were not very effective when applied to complex symmetric linear systems expressed in ERF. Freund claimed that it is better to solve such systems directly in the complex domain and developed the QMR and the transpose-free QMR methods for that task. One reason for poor performance on ERF is the resulting complex eigenvalue distribution, which is folded about the real axis, and this folding could cause problems with the Krylov iteration. From my own experience, a similar problem is observed when AMG, designed for real systems, is applied to complex symmetric systems. Development of AMG solvers formulated specifically for complex symmetric systems thus appears warranted.

Acknowledgments I wish to thank Yasuo Ogawa, Graham Heinson, and other members of the 21st EM Workshop Program Committee for the invitation and opportunity to write this review article. Input from the two referees, Klaus Spitzer and Chester Weiss, also improved the content of the review. Finally, I also wish to acknowledge my employer, Lawrence Berkeley Laboratory, and the U.S. Department of Energy Office of Science for funding, under contract number DE-AC02-05CH11231.

References

Alumbaugh DL, Newman GA (1997) 3-D massively parallel electromagnetic inversion—part II. Anal Crosswell Exp Geophys J Int 128:355–363

Alumbaugh DL, Newman GA, Prevost L, Shadid JN (1996) Three dimensional, wideband electromagnetic modeling on massively parallel computers. Radio Sci 31:1–23

Amestoy PR, Duff IS, Koster J, L'Excellent JY (2001) A fully asynchronous multifrontal solver using distributed dynamic scheduling. SIAM J Matrix Anal Appl 23(1):15–41

Amestoy PR, Guermouche A, L'Excellent JY, Pralet S (2006) Hybrid scheduling for the parallel solution of linear systems. Parallel Comput 32(2):136–156

Balay S, Brown J, Buschelman K, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Smith BF, Zhang H (2010) PETSc users manual. Tech. Rep. Number ANL-95/11—revision 3.1, Argonne National Laboratory

Börner R-U, Ernst OG, Spitzer K (2008) Fast 3-D simulation of transient electromagnetic fields by model reduction in the frequency domain using Krylov subspace projection. Geophys J Int 173:766–780. doi:10.1111/j.1365-246x.2008.03750.x

Carazzone JJ, Burtz OM, Green KE, Pavlov DA, Xia C (2005) Three-dimensional imaging of marine CSEM data. In: 75th Annual international meeting, SEG, expanded abstracts, pp 575–578

Carazzone JJ, Dickens TA, Green KE, Jing C, Wahrmund LA, Willen DE, Commer M, Newman GA (2008) Inversion study of a large marine CSEM survey. In: 78th Annual international meeting, SEG, expanded abstracts, pp 644–647

Chen J, Dickens T (2009) Effects of uncertainty in rock-physics models on reservoir parameter estimation using seismic amplitude variation with angle and controlled-source electromagnetics data. Geophys Prospect 57:61–74

Chen J, Hoversten GM, Vasco D, Rubin Y, Zhou Z (2007) A Bayesian model for gas saturation estimation using marine seismic AVA and CSEM data. Geophysics 72:WA85–WA95

Chen J, Tompkins M, Zhang P, Wilt M, Mackie R (2012) Frequency-domain EM modeling of 3D anisotropic magnetic permeability and analytical analysis. In: 82nd Annual international meeting, SEG extended abstracts, pp 1–5. doi:10.1190/segam2012-0308.1

Chen J, Hoversten GM, Key K, Nordquest G, Cumming W (2012b) Stochastic inversion of magnetotelluric data using a sharp boundary parameterization and application to a geothermal site. Geophysics 77(4):E265–E279. doi:10.1190/geo2011-0430.1

Colombo D, Keho T, McNeice G (2012) Integrated seismic-electromagnetic workflow for sub-basalt exploration in northwest Saudi Arabia. Lead Edge 31:42–52

Commer M, Newman GA (2008) New advances in three-dimensional controlled-source electromagnetic inversion. Geophys J Int 172:513–535

Commer M, Newman GA (2009) Three-dimensional controlled-source electromagnetic and magnetotelluric joint inversion. Geophys J Int 178:1305–1316

Commer M, Newman GA, Carazzone JJ, Dickens TA, Green KE, Wahrmund LA, Willen DE, Shiu J (2008) Massively parallel electrical-conductivity imaging of hydrocarbons using the IBM Blue Gene/L supercomputer. IBM J Res Dev 52(½):93–103

Commer M, Maia FRN, Newman GA (2011) Iterative Krylov solution methods for geophysical electromagnetic simulations on throughput-oriented processing units. Int J High Perform Comput Appl 26(4):378–385. doi:10.1177/1094342011428145

Constable S (2006) Marine electromagnetic methods—a new tool for offshore exploration. Lead Edge 25:438–444

Druskin V, Knizhnerman L (1994) Spectral approach to solving three-dimensional Maxwell's diffusion equations in the time and frequency domains. Radio Sci 29(4):937–953

Druskin V, Knizhnerman LA, Ping L (1999) New spectral Lanczos decomposition method for induction modeling in arbitrary 3-D geometry. Geophysics 64(3):701–706

Eidesmo T, Ellingsrud S, MacGregor LM, Constable S, Sinha MC, Johansen S, Kong S, Westerdahl FN (2002) Sea Bed Logging (SBL), a new method for remote and direct identification of hydrocarbon filled layers in deepwater areas. First Break 20(3):144–152

Ellingsrud S, Eidesmo T, Johansen S, Sinha MC, MacGregor LM, Constable S (2002) Remote sensing of hydrocarbon layers by seabed logging (SBL): results from a cruise offshore Angola. Lead Edge 21:972–982

Farquharson CG, Oldenburg DW (1996) Approximate sensitivities for the electromagnetic inverse problem. Geophys J Int 126:235–252

Fox GC (1988) Solving problems on concurrent processors. Prentice Hall, Old Tappan, NJ

Franke A, Börner R-U, Spitzer K (2007) Adaptive unstructured grid finite element simulation of two-dimensional magnetotelluric fields for arbitrary surface and seafloor topography. Geophys J Int 171:71–86

Freund R (1992) Conjugate gradient type methods for linear systems with complex symmetric coefficient matrices. SIAM J Sci Stat Comput 13:425–448

Freund R, Nachtigal N (1991) QMR: a quasi-minimal residual method for non-hermitian linear systems. Numer Math 60:315–339

Grayver AV, Streich R, Ritter O (2013) Three-dimensional parallel distributed inversion of CSEM data using a direct forward solver. Geophys J Int 193(3):1432–1446. doi:10.1093/gji/ggt055

Greenbaum A (1997) Iterative methods for solving linear systems. SIAM, Philadelphia, PA

Gribenko A, Zhdanov MS (2007) Rigorous 3D inversion of marine CSEM data based on the integral equation method. Geophysics 72:73–84

Gropp W, Lusk E, Doss N, Skjellum A (1996) A high-performance, portable implementation of the MPI message passing interface standard. Parallel Comput 22:789–828

Habashy TM, Groom RW, Spies BR (1993) Beyond the Born and Rytov approximations: a nonlinear approach to electromagnetic scattering. J Geophy Res 98:1759–1775. doi:10.1029/92JB02324

Heroux MA, Willenbring JM, Heaphy R (2003) Trilinos developers guide part II: ASCI software quality engineering practices version 1.0. Tech. Rep. SAND2003-1899, Sandia National Laboratories

Heroux MA, Salinger AG, Bartlett RA, Thornquist HK, Howle VE, Tuminaro RS, Hoekstra RJ, Willenbring JM, Hu JJ, Willina SA, Kolda T, Lehoucq RB, Long KR, Pawlowski RP, Philipps ET, Stanley KS (2005) An overview of the Trilinos project. ACM Trans Math Softw 31:397–423

Hestenes MR, Stiefel E (1952) Methods of conjugate directions for solving linear systems. J Res Natl Bureau Stand 49:409–435

Hohmann GW (1975) Three-dimensional induced polarization and electromagnetic modeling. Geophysics 40:309–324

Hoversten G, Constable S, Morrison H (2000) Marine magnetotellurics for base-of-salt mapping: Gulf of Mexico field test at the Gemini structure. Geophysics 65:1476–1488

Jegen MD, Hobbs R, Tarits P, Chave A (2009) Joint inversion of marine magnetotelluric and gravity data incorporating seismic constraints: preliminary results of sub-basalt imaging off the Faroe Shelf. Earth Planet Sci Lett 282:47–55

Key K, Ovall J (2011) A parallel goal-oriented adaptive finite element method for 2.5-D electromagnetic modeling. Geophys J Int 186(1):137–154. doi:10.1111/j.1365-246X2011.05025.x

Krylov A (1931) On the numerical solution of the equation by which in technical questions frequencies of small oscillations of material systems are determined. Izv. Akad. Nauk SSSR 7:491–539 (in Russian)

Lanczos C (1952) Solution of systems of linear equations by minimized iterations. J Res Natl Bureau Stand 49:33–53

Lawlor OS (2009) Message passing for GPGPU clusters. In: CudaMPI: IEEE international conference on cluster computing and workshops, 2009. CLUSTER '09, pp 1–8

Li XS, Demmel JW (2003) SuperLU_DIST: a scalable distributed-memory sparse direct solver for unsymmetric linear systems. ACM Trans Math Softw 29(2):110–140

Liu B, Li SC, Nie LC, Wang J, Nie LC, Wang J, Zhang QS (2012) 3D resistivity inversion using an improved Genetic Algorithm based on control method of mutation direction. J Appl Geophys 87:1–8. doi:10.1016/j.jappgeo.2012.08.002

Lu JJ, Wu XP, Spitzer K (2010) Algebraic multigrid methods for 3D DC resistivity modeling. Chin J Geophys. Special issue of the 19th international workshop on electromagnetic induction in the Earth, Beijing, Oct 23–29, 2008, vol 53, pp 700–707

MacGregor L, Andeis D, Tomlinson T, Barker N (2006) Controlled-source electromagnetic imaging of the Nuggets-1 reservoir. Lead Edge 25:984–992

Mackie RL, Madden TR (1993) Conjugate direction relaxation solutions for 3-D magnetotelluric modeling. Geophysics 58:1052–1057
Maresh J, White RS (2005) Seeing through a glass, darkly: strategies for imaging through basalt. First Break 23:27–33
Moucha R, Bailey RC (2004) An accurate and robust multi-grid algorithm for 2D resistivity modeling. Geophys Prospect 52:197–212
Mudge JC, Heinson GS, Thiel S (2011) Evolving inversion methods in geophysics with cloud computing—a case study of an eScience collaboration. In: Proceedings of IEEE eScience, pp 119–125
Newman GA, Alumbaugh DL (1995) Frequency domain modeling of airborne electromagnetic responses using staggered finite differences. Geophys Prospect 43:1021–1042
Newman GA, Alumbaugh DL (1997) 3-D massively parallel electromagnetic inversion—part I theory. Geophys J Int 128:345–354
Newman GA, Alumbaugh DL (1999) 3-D electromagnetic modeling and inversion on massively parallel computers. In: Oristaglio MN, Spies BR (eds) Three-dimensional electromagnetics. Society of exploration geophysicists, Geophysical Developments No. **7**, Tulsa OK, pp 299–321
Newman GA, Alumbaugh DL (2000) Three-dimensional magnetotelluric inversion using non-linear conjugate gradients. Geophys J Int 140:410–424
Newman GA, Commer M (2005) New advances in transient electromagnetic inversion. Geophys J Int 160:5–32
Newman GA, Commer M (2009) Massively parallel electrical conductivity imaging of the subsurface. J Phys Conf Ser 180:012063
Newman GA, Hohmann GW, Anderson WL (1986) Transient electromagnetic response of a three-dimensional body in a layered earth. Geophysics 51:1608–1627
Newman GA, Commer M, Carazzone JJ (2010) Imaging CSEM data in the presence of electrical anisotropy. Geophysics 75:51–61
Oldenburg DW, Haber E, Shekhtman R (2013) Three dimensional inversion of multisource time domain electromagnetic data. Geophysics 78:E47–E57
Plessix RE, Mulder WA (2008) Resistivity imaging with controlled-source electromagnetic data: depth and data weighting. Inverse Prob 24:1–22
Plessix RE, van der Sman P (2007) 3D CSEM modeling and inversion in complex geological settings. In: 77th Annual international meeting, SEG, expanded abstracts, pp 589–593
Plessix RE, van der Sman P (2008) Regularized and blocky 3D controlled source electromagnetic inversion. In: 24th Progress in electromagnetic research symposium, abstracts, pp 755–760
Puzyrev V, Koldan J, de la Puente J, Houzeaux G, Vazquez M, Cele J (2013) A parallel finite-element method for 3D controlled-source electromagnetic forward modeling. Geophys J Int 193:678–693. doi:10.1093/gji/ggt027
Schenk O, Gärtner K (2004) Solving unsymmetric sparse systems of linear equations with PARDISO. J Future Gen Comput Syst 20(3):475–487
Schenk O, Gärtner K (2006) On fast factorization pivoting methods for symmetric indefinite systems. Elec Trans Numer Anal 23:158–179
Schwarzbach C, Haber E (2013) Finite-element based inversion for time-harmonic electromagnetic problems. Geophys J Int 193:615–634. doi:10.1093/gji/ggt006
Schwarzbach C, Börner R-U, Spitzer K (2005) 2D inversion of direct current resistivity data using a parallel, multi-objective genetic algorithm. Geophys J Int 162:685–695
Schwarzbach C, Börner R-U, Spitzer K (2011) Three-dimensional adaptive higher order finite element simulation for geo-electromagnetics—a marine CSEM example. Geophys J Int 187:63–74
Smith JT (1992) Conservative modeling of 3-D electromagnetic fields. Paper presented at the 11th workshop on electromagnetic induction in the earth. International association of geomagnetism and aeronomy. Wellington, New Zealand, Aug 26–Sept 2
Smith JT (1996a) Conservative modeling of 3-D electromagnetic fields, part I: properties and error analysis. Geophysics 61:1308–1318
Smith JT (1996b) Conservative modeling of 3-D electromagnetic fields, part II: biconjugate gradient solution and an accelerator. Geophysics 61:1319–1324
Smith JT, Booker JR (1991) Rapid inversion of two- and three dimensional magnetotelluric data. J Geophys Res 96:3905–3922
Spitzer K (1995) A 3D finite difference algorithm for DC resistivity modeling using conjugate gradient methods. Geophys J Int 123:903–914
Spitzer K, Wurmstich B (1999) Speed and accuracy in 3D resistivity modeling. In: Oristaglio ML, Spies BR (eds) Three-dimensional electromagnetics, SEG book series "Geophysical Developments", No. **7**, Society of exploration geophysicists, pp 161–176, Tulsa, OK

Streich R (2009) 3D finite-difference frequency-domain modeling of controlled-source electromagnetic data: direct solution and optimization for high accuracy. Geophysics 74(5):F95–F105. doi:10.1190/1.3196241

Stuart JA, Owens JD (2009) Message passing on data-parallel architectures. In: Proceedings of the 23rd IEEE international parallel and distributed processing symposium

Torres-Verdin C, Habashy TM (1994) Rapid 2.5-dimensional forward modeling and inversion via a new nonlinear scattering approximation. Radio Sci 29:1051–1079

Torres-Verdin C, Habashy TM (1995) A two step linear inversion of two dimensional electrical conductivity. IEEE Trans Antenna Propag 43:405–415

Tuminaro RS, Heroux M, Hutchinson SA, Shadid JN (1999) Official Aztec user's guide: version 2.1, sand report SAND99-8801J, Sandia National Laboratories

Um E, Harris JM, Alumbaugh DL (2010) 3D time-domain simulation of electromagnetic diffusion phenomena: a finite-element electric-field approach. Geophysics 75:115–126

Um E, Commer M, Newman GA (2013) Efficient pre-conditioned iterative solution strategies for the electromagnetic diffusion in the Earth: finite-element frequency-domain approach. Geophys J Int. doi:10.1093/gji/ggt071

van der Vorst H (1992) Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of non-symmetric linear systems. SIAM J Sci Statist Comput 13:631–644

Vieira da Silva N, Morgan JV, Macgregor L, Warner M (2012) A finite element multifrontal method for 3D CSEM modeling in the frequency domain. Geophysics 77:101–115

Wang T, Hohmann GW (1993) A finite-difference, time-domain solution for three-dimensional electromagnetic modeling. Geophysics 58:797–809

Wang S, de Hoop MV, Xia J (1011) 2011, On 3D modeling of seismic wave propagation via a structured parallel multifrontal direct Helmholtz solver. Geophys Prospect 59:857–873. doi:11/j.1365-2478.2011.00982.x

Wang S, de Hoop MV, Xia J, Li XS (2012) Massively parallel structured multifrontal solver for time-harmonic elastic waves in 3-D anisotropic media. Geophys J Int 191(1):346–366. doi:10.1111/j.1365.246X.2012.05634.x

Wannamaker PE, Hohmann GW, Ward SH (1984) Magnetotelluric responses of three-dimensional bodies in layered earths. Geophysics 49:1517–1533

Weiss CJ, Schultz A (2011) An evaluation of parallelization strategies for low-frequency electromagnetic induction simulators using staggered grid discretizations. In: American geophysical union fall meeting conference proceedings, informatics session, San Francisco

Yang C, Huang C, Lin C (2011) Hybrid CUDA, OpenMP, and MPI parallel programming on multicore GPU clusters. Comput Phys Commun 182(1):266–269

Zach JJ, Bjørke AK, Støren T, Maaø F (2008) 3D inversion of marine CSEM data using a fast finite-difference time-domain forward code and approximate Hessian-based optimization. In: 78th Annual international meeting, SEG, expanded abstracts, pp 614–618

Zhdanov MS, Fang S (1999) 3D electromagnetic inversion based on the quasi-linear approximation. In: Three-dimensional electromagnetics, SEG book series "Geophysical Developments", No. **7**, Society of exploration geophysicist, pp 233–255. Society of Exploration Geophysicists, Tulsa, OK