

Massively parallel electrical-conductivity imaging of hydrocarbons using the IBM Blue Gene/L supercomputer

M. Commer
G. A. Newman
J. J. Carazzone
T. A. Dickens
K. E. Green
L. A. Wahrmond
D. E. Willen
J. Shiu

Large-scale controlled-source electromagnetic (CSEM) three-dimensional (3D) geophysical imaging is now receiving considerable attention for electrical-conductivity mapping of potential offshore oil and gas reservoirs. To cope with the typically large computational requirements of the 3D CSEM imaging problem, our strategies exploit computational parallelism and optimized finite-difference meshing. We report on an imaging experiment utilizing 32,768 tasks (and processors) on the IBM Blue Gene/L™ (BG/L) supercomputer at the IBM T. J. Watson Research Center. Over a 24-hour period, we were able to image a large-scale marine CSEM field dataset that previously required more than 4 months of computing time on distributed clusters utilizing 1,024 tasks on an InfiniBand® fabric. The total initial data-fitting errors (i.e., “misfits”) could be decreased by 67% within 72 completed inversion iterations, indicating the existence of an electrically resistive region in the southern survey area below a depth of 1,500 m underneath the seafloor. The major part of the residual misfit stems from transmitter-parallel receiver components that have an offset from the transmitter sail line (broadside configuration). Modeling confirms that improved broadside data fits can be achieved by considering anisotropic electrical conductivities. While delivering a satisfactory gross-scale image for the depths of interest, the experiment provides important evidence for the necessity of discriminating between horizontal and vertical conductivities for maximally consistent 3D CSEM inversions.

Introduction

Seismic methods have a long and established history in hydrocarbon (i.e., oil and gas) exploration and are very effective in mapping geologic reservoir formations. However, these methods are not good at discriminating the different types of reservoir fluids contained in the rock pore space, such as brines, water, oil, and gas. This discrimination challenge has encouraged the development of new geophysical technologies that can be combined with established seismic methods to directly image fluids. One technique that has recently emerged, with considerable potential, utilizes low-frequency electromagnetic (EM) energy to map variations in the subsurface electrical conductivity, σ (in units of Siemens per meter), or its

reciprocal ($1/\sigma$ in units of ohm-meters), usually called *resistivity*, of offshore oil and gas prospects [1–5]. Resistivity is a more meaningful quantity for imaging hydrocarbons. A resistivity increase, compared to the surrounding geological strata, may directly indicate potential reservoirs. EM field measurements have been shown to be highly sensitive to changes in the pore fluid types and the location of hydrocarbons, given a sufficient resistivity contrast between hydrocarbons and fluids such as brine or water.

With the marine controlled-source EM (CSEM) measurement technique, a deep-towed electric-dipole transmitter is used to excite a continuous low-frequency (~ 0.1 to 10 Hz) EM signal that is measured on the

©Copyright 2008 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied by any means or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/08/\$5.00 © 2008 IBM

seafloor by electric and magnetic field detectors, where the largest transmitter-detector offsets can exceed 15 km. The data is initially collected in the time domain.

Conversion to the frequency domain involves application of a window function, with its characteristics defined by the fundamental signal frequency, which is 0.125 Hz in the present survey. To cover larger depth ranges, multiple transmitter frequencies are usually employed in a survey. Similar to acoustic wave propagation, the attenuation rate as a function of exploration depth increases with the wave frequency. Current technologies require low-frequency EM signals (<1 Hz) to interrogate to reservoir depths as deep as 4 km.

Exploration with the CSEM technology in the search for hydrocarbons now extends to highly complex and subtle offshore geological environments. The geometries of the reservoirs are inherently three-dimensional (3D) and exceedingly difficult to map without recourse to 3D EM imaging experiments, requiring fine model parameterizations, spatially exhaustive survey coverage, and multicomponent data. The 3D imaging problem, which in this paper is also referred to as the *inversion problem*, usually has large computational demands, due to the computationally expensive solution of the forward modeling problem, that is, the EM field simulation on a given 3D finite-difference (FD) grid. Moreover, large data volumes require many forward solutions in an iterative inversion scheme. Therefore, we have developed an imaging algorithm that utilizes two levels of parallelization, one applied to the modeling (or imaging) volume and the other applied to the data volume. The algorithm is designed for arbitrarily large datasets, allowing for an arbitrarily large number of parallel tasks, while the computationally idle message passing is minimized. We have further incorporated an optimal meshing scheme that allows us to separate the imaging or modeling mesh from the simulation mesh. This provides for significant acceleration of the 3D EM field simulation, having a direct impact on the time to solution for the 3D imaging process.

Here, we report an imaging experiment, utilizing 32,768 tasks (and processors) on the IBM Blue Gene/L* (BG/L) supercomputer located at the IBM T. J. Watson Research Center. The experiment is novel in terms of both computational resources utilized and amount of data inverted. Our main purpose is to provide a feasibility study for the effectiveness of the employed algorithm. Further, the results obtained will improve important base knowledge for the design of upcoming large-scale CSEM surveys and improve the automated imaging method for data interpretation.

Problem formulation

We formulate the inverse problem, mentioned in the introduction, by finding a model \mathbf{m} with m piecewise

constant electrical-conductivity parameters that describe the earth model reproducing a given dataset. Specifically, the inversion algorithm minimizes the error functional,

$$\Phi = \frac{1}{2} \left[\mathbf{D}(\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{p}}) \right]^{\text{T}*} \left[\mathbf{D}(\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{p}}) \right] + \frac{1}{2} \lambda (\mathbf{W}\mathbf{m})^{\text{T}} (\mathbf{W}\mathbf{m}), \quad (1)$$

where T denotes the transpose operator, and T^* denotes the Hermitian conjugate operator. In the above expression, the predicted data vectors (from a starting model) and observed data vectors are denoted by \mathbf{d}^{p} and \mathbf{d}^{obs} , respectively, where each has n complex values. These vectors consist of electric or magnetic field values specified at the measurement points, where the predicted data is determined through solution of the time harmonic 3D Maxwell's equations in the diffusive approximation. We have also introduced a diagonal weighting matrix, $\mathbf{D}_{n \times n}$, into the error functional to compensate for noisy measurements. To stabilize the minimization of Equation (1), and to reduce model curvature in three dimensions, we introduce a matrix $\mathbf{W}_{m \times m}$ based upon an FD approximation to the Laplacian (∇^2) operator applied in Cartesian coordinates. The parameter λ attempts to balance the data error and the model smoothness constraint.

The forward problem

Within an inversion framework, the forward problem is solved multiple times to simulate the EM field, denoted by the vector \mathbf{E} , and thus, the data \mathbf{d}^{p} for a given model \mathbf{m} . EM wave propagation is controlled by the vector Helmholtz equation,

$$\nabla \times \nabla \times \mathbf{E} + i\omega\mu_0\sigma\mathbf{E} = -i\omega\mu_0\mathbf{J}, \quad (2)$$

where source vector, free-space magnetic permeability, and angular frequency are denoted by \mathbf{J} , μ_0 , and ω , respectively (see [6] for specific details). Our solution method is based on the consideration that the number of model parameters required to simulate realistic 3D distributions of the electrical conductivity σ can typically exceed 10^7 . FD modeling schemes are ideally suited for this task and can be parallelized to handle large-scale problems that cannot be easily treated otherwise [6]. After approximating Equation (2) on a staggered grid at a specific angular frequency, using finite differencing and eliminating the magnetic field, we obtain a linear system for the electric field,

$$\mathbf{K}\mathbf{E} = \mathbf{S}, \quad (3)$$

where \mathbf{K} is a sparse complex symmetric matrix with 13 non-zero entries per row [6]. The diagonal entries of \mathbf{K} depend explicitly on the conductivity parameters that we seek to estimate through the inversion process. Since the

electric field, \mathbf{E} , also depends upon the conductivity, implicitly, this gives rise to the nonlinearity of the inverse problem. The fields are sourced (i.e., generated) with a grounded wire or loop embedded within the modeling domain, described by the discrete source vector, \mathbf{S} , and include Dirichlet boundary conditions imposed upon the problem. To help avoid excessive meshing near the source, we favor a scattered-field formulation to the forward modeling problem. In this instance, \mathbf{E} is replaced with \mathbf{E}_s in Equation (3). The source term, for a given transmitter, will now depend upon the difference between the 3D conductivity model and a simple background model, weighted by the background electric field \mathbf{E}_b , where $\mathbf{E} = \mathbf{E}_b + \mathbf{E}_s$. Simple background models with one-dimensional (1D) conductivity distributions, i.e., models in which σ changes only with depth, are used because fast semi-analytical solutions for \mathbf{E}_b are available. Given the solution of the electric field in Equation (3), the magnetic field can be easily determined from a numerical implementation of Faraday's law. An efficient solution process is paramount. We solve Equation (3) to a predetermined error level using iterative Krylov subspace methods, using either a biconjugate gradient (BICG) or quasi-minimum residual (QMR) scheme with preconditioning [6].

Minimization procedure

In large-scale nonlinear inverse problems, as considered here, we minimize Equation (1) using gradient-based optimization techniques because of their minimal storage and computational requirements. We characterize these methods as gradient-based techniques because they employ only first-derivative information of the error functional in the minimization process, specifically $-\nabla\phi$. Gradient-based methods include steepest descent, nonlinear conjugate gradient, and limited memory quasi-Newton schemes, with the latter usually providing the best inverse solution convergence, however, at a larger computational expense. Solution accelerators are discussed in [7], which also provides detailed derivation of the gradients and an efficient scheme for their computation. Here, we focus on a nonlinear conjugate gradient (NLCG) minimization approach as a trade-off between inverse solution convergence and computational effort per inversion iteration.

Exploitation of solution parallelism

In order to realistically image the subsurface of large survey areas at a sufficient level of resolution and detail, industrial CSEM datasets can contain up to hundreds of transmitter–receiver arrays operating at different frequencies, with a spatial covering of more than 1,000 km². This easily requires thousands of solutions to the forward modeling problem for just one imaging

experiment. Hence, the computational demands for solving the 3D inverse problem are enormous. To cope with this problem, our algorithm utilizes two levels of parallelization, one over the modeling domain and the other over the data volume.

First, in solving the forward problem on a distributed environment, we split the FD simulation grid, not the matrix, among a Cartesian processor topology, which shall be called *local communicator* (LC). As the linear system is relaxed during the iterative solution, which involves matrix–vector products on each of the processors, values of the solution vector at the current Krylov iteration not stored on the processor must be passed by neighbors within LC to complete the matrix–vector products. Additional global communication across the LC is needed to complete several dot products at each relaxation step of the Krylov iteration. The solution time increases linearly with the number of parallel tasks, up to a point at which the increase in message-passing overhead dominates. A study of the FLOP (floating-point operation) rate versus communicator size for the Intel Paragon** architecture is exemplified in [6].

To carry out many forward simulations simultaneously, we employ multiple LCs, connected via a group of lead processors, with one lead task assigned to each LC. The topology of this lead group defines the communicator on which the iterative NLCG inversion framework is carried out, here called the *global communicator* (GC). This distribution of the forward modeling problems, or data decomposition, is highly parallel. Assuming the optimal LC size has been estimated for a given range of mesh sizes, the size of the GC (which equals the number of LCs) can be increased linearly with the data volume. The relative amount of communication within the GC remains constant because communication within the GC is needed only in order to complete several dot products per inversion iteration and to sum the contributions from each LC to the global gradient vector. The main computational and communication burden occurs with the forward FD solving. As outlined below, we adapt FD mesh sizes according to given transmitter–receiver configurations and minimum spatial sampling requirements. To keep a balanced computational and communication workload between all LCs, the data decomposition is based on a balanced distribution of the FD grids in terms of grid sizes.

Optimal mesh considerations

Although our experience using two parallelization levels has been satisfactory, in order to solve the very large problems of interest, we must obtain a higher level of efficiency. One promising approach, which we have previously reported [8], is to design an optimal FD




Figure 1

Illustration of the conductivity averaging scheme of Equation (4) in two dimensions.

simulation mesh for each source excitation in Equation (3). FD meshing for field simulation then considers only part of the total model volume where it can have an appreciable effect in the imaging process. Moreover, minimum spatial grid sampling intervals are dictated by the EM field wavelength and, hence, can be optimized according to a specific source excitation frequency. Optimizing both mesh size and spatial sampling, we create a collection of simulation grids, Ω_s , that support the EM field simulation for all different sources contained in the dataset. All simulation grids act upon a common model grid, Ω_m , which defines the imaging volume. Both types of grids are Cartesian with conformal grid axes. Key to the grid separation is an appropriate mapping scheme that transfers the material properties from Ω_m to Ω_s . The imaging process provides piecewise constant estimates of the electrical conductivity, which are defined by the cells of Ω_m . The staggered FD mesh Ω_s , on the other hand, involves edge-based directional conductivities, needed for constructing the stiffness matrix \mathbf{K} in Equation (3) (see also [6] and [9] for details). In the case in which $\Omega_m = \Omega_s$, an edge conductivity, σ^e , is computed from $\sigma^e = \sum_{i=1}^4 \sigma_i w_i$, with $w_i = dV_i / \sum_{j=1}^4 dV_j$. Here, w_i represent weights corresponding to volume fractions of the four cells on Ω_m that share the edge σ^e on Ω_s . Furthermore, the edge conductivity σ^e is simply an arithmetic volume average of the four model cell conductivities. When $\Omega_m \neq \Omega_s$, the conductivity mapping involves parallel and serial circuit analysis, resulting in an arithmetic and harmonic conductivity averaging scheme [8, 10]. The averaging scheme is exemplified in **Figure 1** for an x -directed edge conductivity σ_x^e in two dimensions. Here, model and simulation meshes are represented by dashed and solid lines, respectively. The material average is to be specified from the formula

$$\sigma_x^e = \left[\int_{x_i}^{x_{i+1}} \left(\int_{y_{j-1/2}}^{y_{j+1/2}} \sigma(x, y) dy \right) dx \right]^{-1}. \quad (4)$$

The inner integration constitutes a point-wise parallel conductivity average, while the outer integration provides for the effective conductivity in series, arising over the integrated edge length $(x_{i+1} - x_i)$ of the simulation mesh. The total integration area assigned to σ_x^e is shown by the red rectangle in Figure 1.

Extension to the full 3D case is straightforward, with the discrete representation exemplified by

$$\sigma_x^e = \sum_{j=1}^J \left[\left(\frac{1}{V_j} \sum_{i=1}^{I_j} dV_i \sigma_i \right)^{-1} \Delta x_j \right]^{-1} \Delta X, \quad (5)$$

where ΔX is the edge length of the simulation cell along the x -coordinate direction. Similarly, σ_y^e and σ_z^e involve averaging along the y - and z -coordinates, respectively. In Equation (5), the averaging along ΔX involves a number of J serially connected discrete parallel circuits, P_j , each with a volume V_j . The length of P_j along the edge is Δx_j , where $\sum_{j=1}^J \Delta x_j = \Delta X$. Further, I_j is the number of cells on the modeling grid contributing to P_j , with σ_i and dV_i the individual model cell conductivities and volume fractions, respectively.

We are also required to specify $\partial \sigma^e / \partial \sigma_k$, which is needed to define the gradient on the modeling grid because it is linked to the forward modeling problem on the simulation grid(s) (see [9] for details on the equal-grid case). Thus,

$$\partial \sigma^e / \partial \sigma_k = \frac{\sigma^e{}^2}{\Delta X} \sum_{j=1}^J \Delta x_j \left(\frac{1}{V_j} \sum_{i=1}^{I_j} dV_i \sigma_i \right)^{-2} \frac{dV_k}{V_j}, \quad (6)$$

where J is now the number of discrete parallel circuits with a non-zero contribution from σ_k . When $\Omega_m = \Omega_s$, we have $J = 1$, $\Delta x_j = \Delta X$ and $\partial \sigma^e / \partial \sigma_k = w_k$, which is the weighting coefficient defined above as $w_k = dV_k / \sum_{j=1}^4 dV_j$.

Electrical-conductivity imaging of hydrocarbons using the BG/L supercomputer

CSEM data is usually characterized by a large dynamic range, which can reach more than ten orders of magnitude. This requires the ability to analyze data in a self-consistent manner that incorporates all structures not only on the reservoir scale at tens of meters, but on the geological basin scale at tens of kilometers, and must include salt domes, detail bathymetry, and other 3D peripheral geology structures that can influence the measurements [11, 12]. These complications motivate the need for an automated 3D conductivity inversion process for successful conductivity imaging of hydrocarbons.

Trial-and-error 3D forward modeling is too cumbersome to be effective. Both model size and amount of the required data provide ample justification for utilizing the massively parallel BG/L supercomputer for the task. Such a platform, which can scale up to 131,072 processors, allows for the capability to image prospective oil and gas reservoirs at the highest resolution possible and on timescales acceptable to the exploration process.

The 3D imaging experiment we present here demonstrates the points mentioned above. The data was acquired offshore of South America. The sail lines and 23 detector locations on a $40 \times 40 \text{ km}^2$ grid used for subsurface conductivity mapping are shown in **Figure 2**. Data was collected from nearly 1 million transmitter sites along the sail lines shown. Obviously, this amount of data cannot be treated with the current inversion methodology, even with a massively parallel implementation. Every source treated by the imaging algorithm requires a forward simulation, an adjoint computation, as well as two or more additional simulations for step control for each nonlinear inversion update. To efficiently deal with this data volume, we use a general reciprocity principle that involves the interchange of transmitter and receiver points. Hence, the positions of the actual CSEM transmitter along the sail line become the computational receiver profiles, and the actual CSEM detectors on the seafloor become computational sources, referred to as *sources* in the following.

The equivalent reciprocal problem involves 951,423 data points and 207 effective sources, since there are 23 source locations with three polarizations and each operating at the three discrete excitation frequencies of 0.125 Hz, 0.25 Hz, and 0.5 Hz. Each effective transmitter is polarized according to the antenna orientation of its corresponding detector. The exact seafloor detector orientations were determined by analyzing the data polarizations and phase reversals with respect to the source sail lines. Data processing involves binning in time, followed by spectral decomposition and spatial filtering. Timing errors were removed by forcing the data phases to match the frequency-offset scaling behavior appropriate to solutions of Maxwell's equations.

The survey layout in Figure 2 contains different transmitter–receiver configurations to be considered, as illustrated in Figure 2(a). For the transmitter sail line position with respect to a given detector on the sea bottom, we consider the so-called *overflight* (left) configuration, in which the sail line is directly over the detector. In the broadside configuration (right), the towed transmitter passes at an offset Δy to one side of the detector. Three components are recorded by the receiver antennas of the detector: inline horizontal (E_x), perpendicular horizontal (E_y), and vertical (E_z) electric fields.

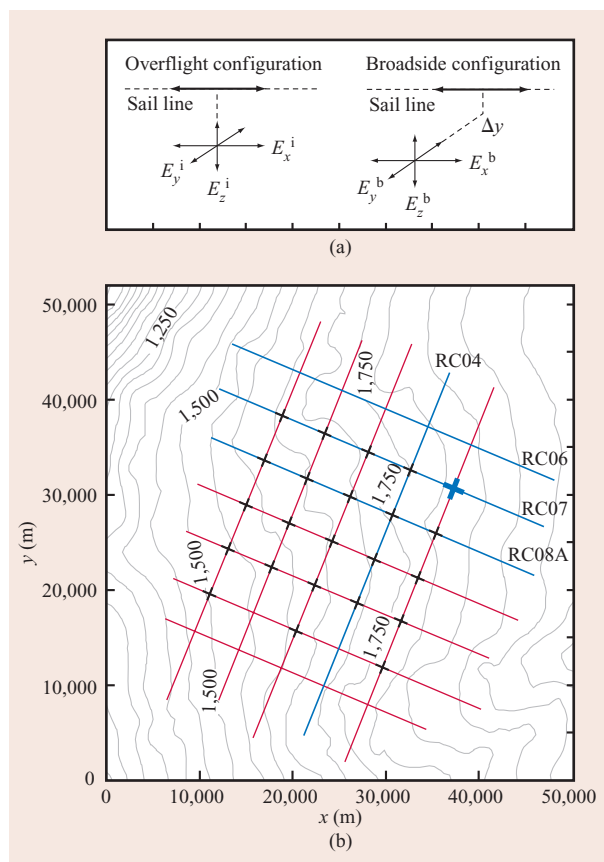


Figure 2

Layout of the sail lines (red and blue) and 23 detector locations (crosses) on the sea bottom for the offshore controlled-source electromagnetic (CSEM) survey. Part (a) illustrates configuration geometries within the dataset. Part (b) shows bathymetry values (seafloor elevation values on contour lines) in meters below sea level. The example data shown in this paper corresponds to the reciprocal treatment of the detector and the sail lines, respectively, marked by the blue cross and the blue lines. (Designators such as RC04 and RC08A are labels for sail lines assigned by ExxonMobil.)

A starting model is necessary to launch the inversion process and resolve some final issues associated with phase components in the data. It is obviously favorable to achieve minimum data misfits (i.e., data-fitting errors) with the starting model. Therefore, the model used has been constructed from knowledge of the sea bottom bathymetry, the seawater electrical-conductivity-versus-depth profile, and 1D inversion of the amplitude components of the common-receiver gathers, based on the inline overflight measurement configuration (E_x^i). The resulting 1D models were then refined by comparing selected simulation results with field observations. In order to accommodate all sail lines and detector sites in the model, a large parameterization was required for Ω_m .

In order to model bathymetry, the minimum required spatial grid sampling interval Δ is kept constant with $\Delta = 125$ m for the horizontal (x and y) coordinates, while it ranges from 50 to 200 m in z . This corresponds to 403 nodes along x and y , 173 nodes vertically, and thus approximately 27.8 million model cells.

To restrict the size of the simulation grid for each source activation, we have assigned each source a separate mesh. Both mesh size and spatial grid sampling rate are based on skin depth estimations. The skin depth δ , a commonly used constant in EM applications, is defined as the depth below the surface of a conductor (in our case, at the transmitter location) at which the current density decays to $1/e$ (~ 0.37) of the surface current density. Using the approximation

$$\delta = 503 / \sqrt{\sigma_b f},$$

mesh intervals depend on the source excitation frequency, f , and the background conductivity, σ_b , of the employed starting model. Horizontal mesh size is based on ten skin depths from the source midpoint, assuming $\sigma_b = 0.5$ S/m (Siemens/meter); the resulting mesh ranges were of sufficient size to accommodate the specific sail lines of data assigned to the effective sources. The horizontal spatial grid sampling intervals, $\Delta = 250$ m, 200 m, and 125 m, vary with frequencies $f = 0.125$ Hz, 0.25 Hz, and 0.5 Hz, respectively. The vertical meshing was identical to that employed in the modeling mesh in order to account for an accurate representation of the seafloor elevation values. With these considerations, we were able to reduce the size of the simulation meshes significantly; the number of x and y grid nodes both ranged from 128 to 162. Solution accuracy was verified against solutions in which $\Omega_s = \Omega_m$.

A maximum of 256 MB of memory per task was available on the BG/L supercomputer. The largest memory requirement results from temporary storage of the forward solutions within one inversion iteration. To stay within the machine limits, each simulation grid was distributed across an LC size of 512 processors, relying on the interprocessor bandwidth to support the BICG/QMR solvers. Sixty-four LCs were then used to distribute the 207 effective sources and their associated data. Thus, the total number of tasks employed in the imaging experiment was 32,768. Disk I/O and file system performance were minor concerns, as the generated image output was relatively modest, ~ 2.5 GB per inversion update, which was written to disk in parallel using 512 tasks. Data output at each inversion iteration consisted of predicted and observed measurements with a total file size of 170 MB. A lead task within the global communicator was assigned to dump the data output after each inversion update.

Prior to the actual imaging experiment, performance tests were carried out. Baseline evaluation involved an

inversion in which the large model grid (size $403 \times 403 \times 173$ nodes) represented the simulation grid for each source. Here are some findings:

1. We compared the job performance using 32 message-passing interface (MPI) tasks completed on the BG/L platform (with a CPU speed of 700 MHz) versus an Intel (Pentium** 4, with a CPU speed of 2.6 GHz) cluster with Gigabit Ethernet fabric. A forward solution used 25 seconds per 100 QMR iterations on the BG/L platform, compared to 23 seconds on the Intel Pentium 4 platform. The computational burden of the QMR solver is dominated by complex double-precision matrix-vector multiplications with indexed memory access. The 64-bit IBM Power Architecture* of the BG/L platform is designed for FLOPs, achieving an efficient memory access. Profiling shows that for our application, the architecture compensates for the lower processor speed of the BG/L platform.
2. Using an LC size of up to 4,096, workload scalability tests revealed that the QMR solution time decreases linearly with the number of parallel tasks. The linear behavior starts to flatten for a larger number of tasks.
3. A 1,024-task job on the BG/L platform showed that the average communication time was about 25% of the total solution time per inversion iteration. The distribution of the communication overhead is as follows. Collective communications within GC are mainly global reduction operations and amount to about 50% of the overhead with typical message sizes of 16 bytes. Point-to-point blocking message passing within LC amounted to 20% overhead, with 30 KB for the average message size. Barrier synchronization provided 30% of the overhead.

The relatively long idle time that is due to global barrier synchronization, which is done after each inversion iteration, indicates the importance of a balanced workload distribution among all LCs. The QMR solver convergence behavior depends on how numerically well-posed the problem is. Numerically, this can be expressed in terms of the *condition number* of the FD stiffness matrix \mathbf{K} in Equation (3), which in turn is governed by the aspect ratio and conductivity contrasts within Ω_s . Because the latter changes dynamically with the model updates during an inversion, faster barrier synchronization would require an adequately sophisticated scheme for dynamically adapting the LC size.

Over a 24-hour period, 72 inversion-model updates were realized on the BG/L platform, and the relative squared error misfit measure was reduced by nearly 67%.

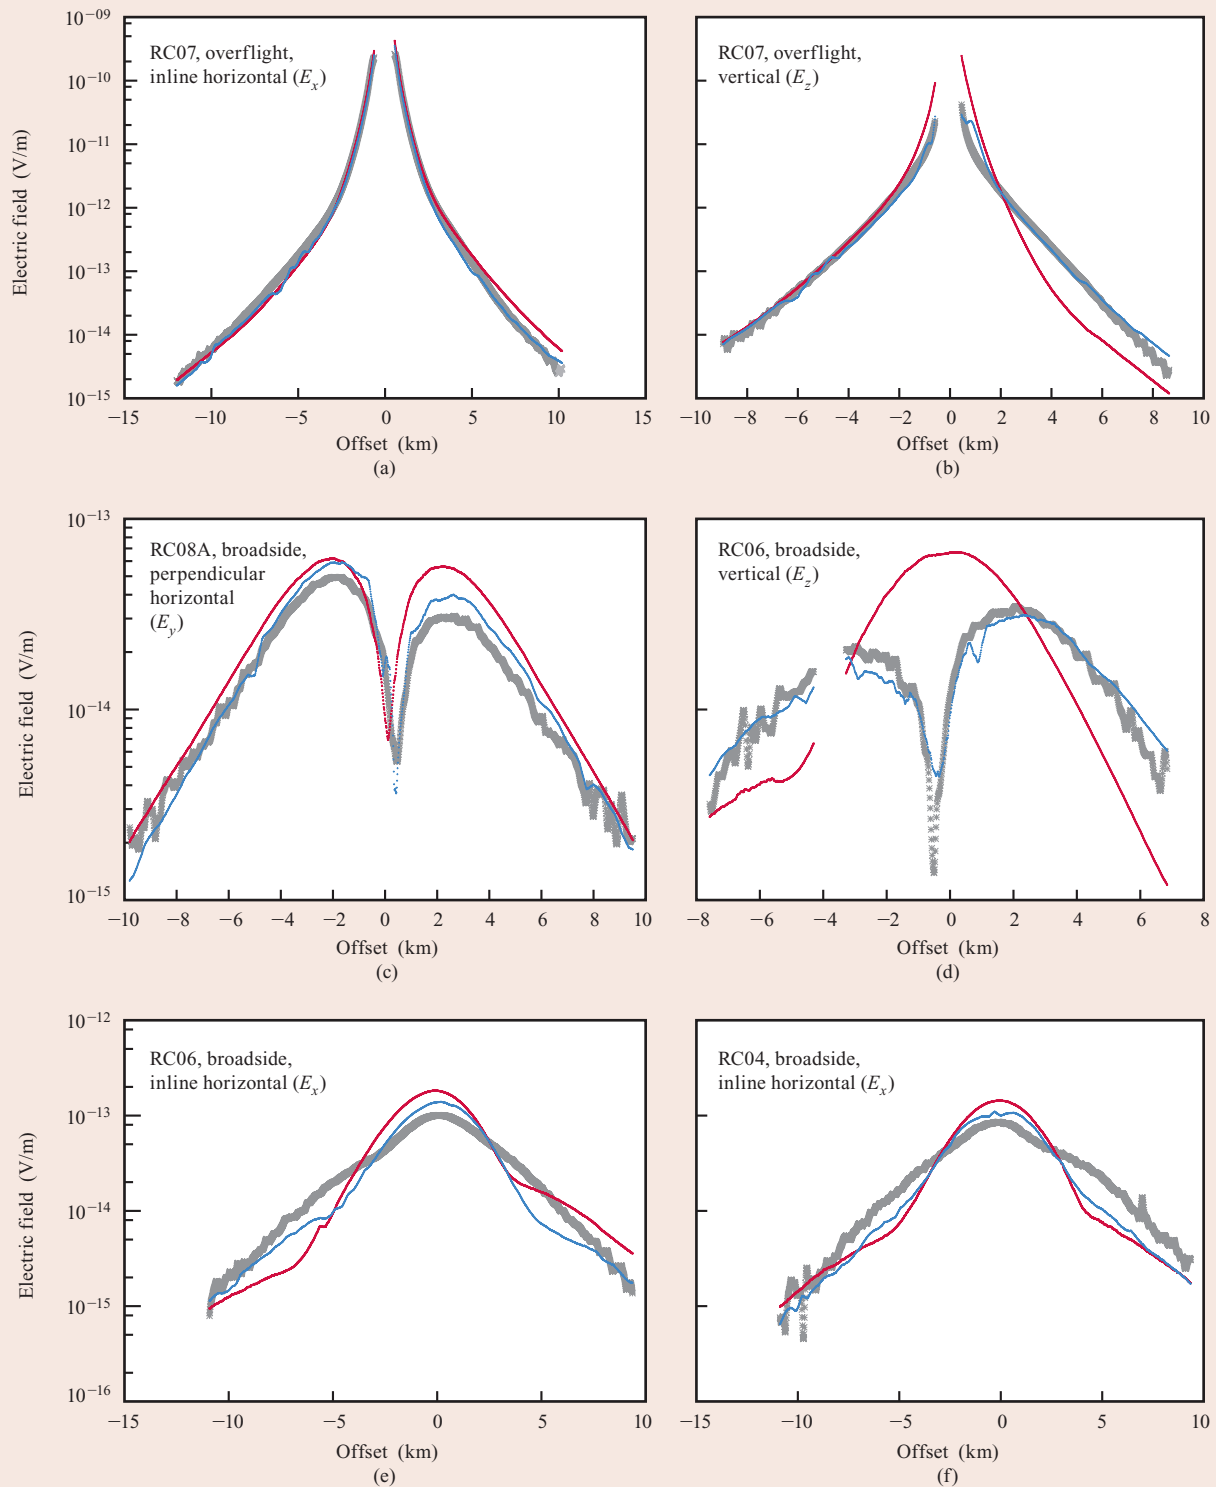


Figure 3

Six selected plots of overflight and broadside electric field data amplitudes (gray) vs. the transmitter offset projected onto the receiver profile lines (RC04, RC06, RC07, RC08A) shown in Figure 2. Shown here are data fits produced by the starting model (red) and for iteration 72 (blue). See text for details.

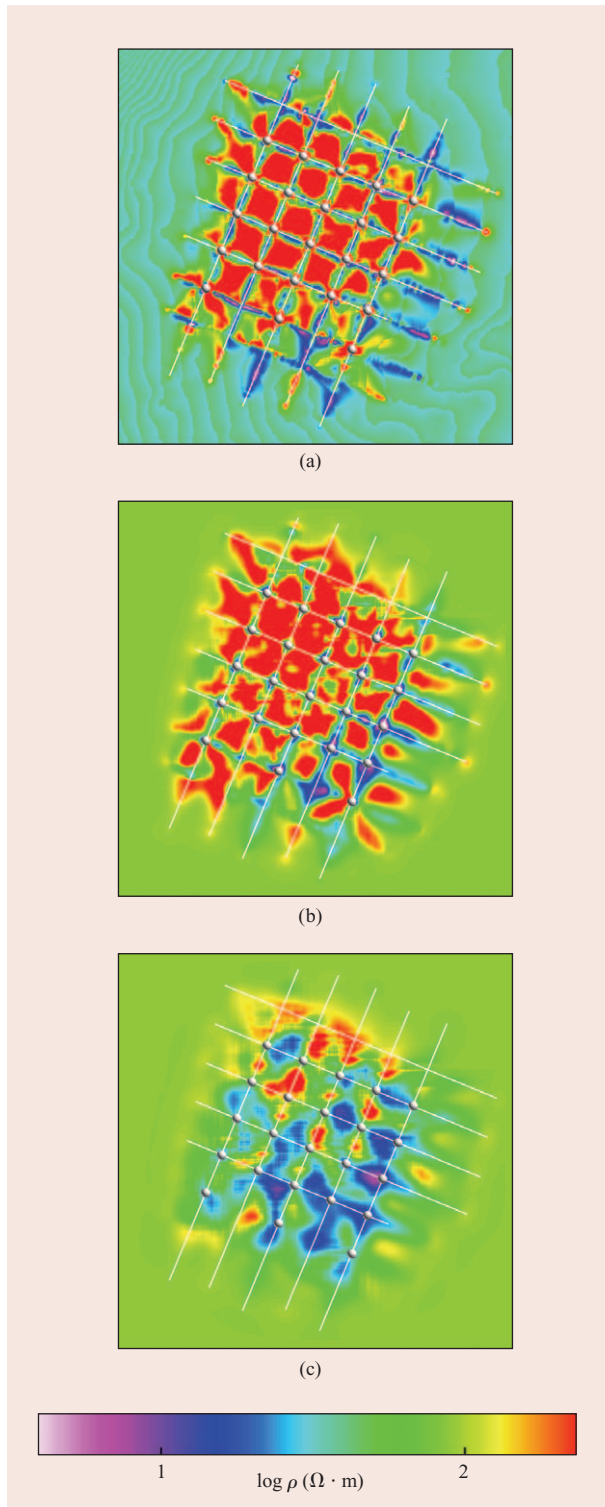


Figure 4

Average resistivity computed over three depth ranges for solution 72: (a) water bottom to 500 m below mud line (BML); (b) interval 500 to 1,500 m BML; (c) interval 1,500 to 2,500 m BML. Resistivity is rendered on a base 10 log scale.

As exemplified in **Figure 3**, good fits, to within the anticipated noise, were obtained for the horizontal and vertical inline electric field overflight data, E_x^i [Figure 3(a)] and E_z^i [Figure 3(b)], as well as for the horizontal perpendicular and vertical broadside electric fields, E_y^b [Figure 3(c)] and E_z^b [Figure 3(d)]. We observed that the major residual misfits originate from the broadside inline components, E_x^b [Figures 3(e) and 3(f)].

The average resistivity computed over three depth ranges for solution 72 is shown in **Figure 4**. The sea bottom defines the depth $z = 0$. Inspection of the images shows increased resistivity in the southern model section for depths below 1,500 m. Such is also observed broadside of the sail lines, for the depth range 0–1,500 m. Along the sail lines, however, little to no resistivity increase is observed. Here, the image has distortions, which are likely caused by the data acquisition method.

A possible explanation for this outcome is the inconsistencies observed in fitting the inline component of the broadside data compared to other data components. This is particularly true of inline overflight data. Clearly, the overflight data will be most sensitive to resistivity variations along the sail lines, while broadside data will be more sensitive to resistivity variations off the sail lines. This enhanced resistivity observed off the sail lines may arise from the attempt of the inversion algorithm to fit the inline broadside data. Enhanced resistivity amplifies the broadside inline model data, reducing the mismatch between observed and predicted data. Nevertheless, it was still not possible to achieve acceptable data fits, indicating a systematic bias in the underlying assumptions employed in the inversion processing.

One critical assumption in this inversion was that the conductivity is isotropic; that is, conductivity within a cell does not vary with direction. However, it is well known within sedimentary rocks that fine-grain bedding planes can induce the rocks to exhibit transverse electrical anisotropy [13, 14]. In addition, parallel interbedding of rocks with different conductivities can lead to anisotropic behavior. Thus, the conductivity can be expected to depend strongly on directions, parallel and perpendicular to the bedding planes. In the context of marine CSEM, reference [15] shows that the effects of electrical anisotropy can produce significant anomalies, even as large as target reservoir responses, and a consensus is now emerging that electrical anisotropy is a stronger factor in influencing marine CSEM measurement than previously believed.

Two tests were carried out to verify the importance of anisotropy. First, we repeated the initial stage of the inversion process to test the degree to which electrical anisotropy is affecting the broadside inline data, and to what lesser extent it influences the overflight and broadside perpendicular and vertical data. This involved

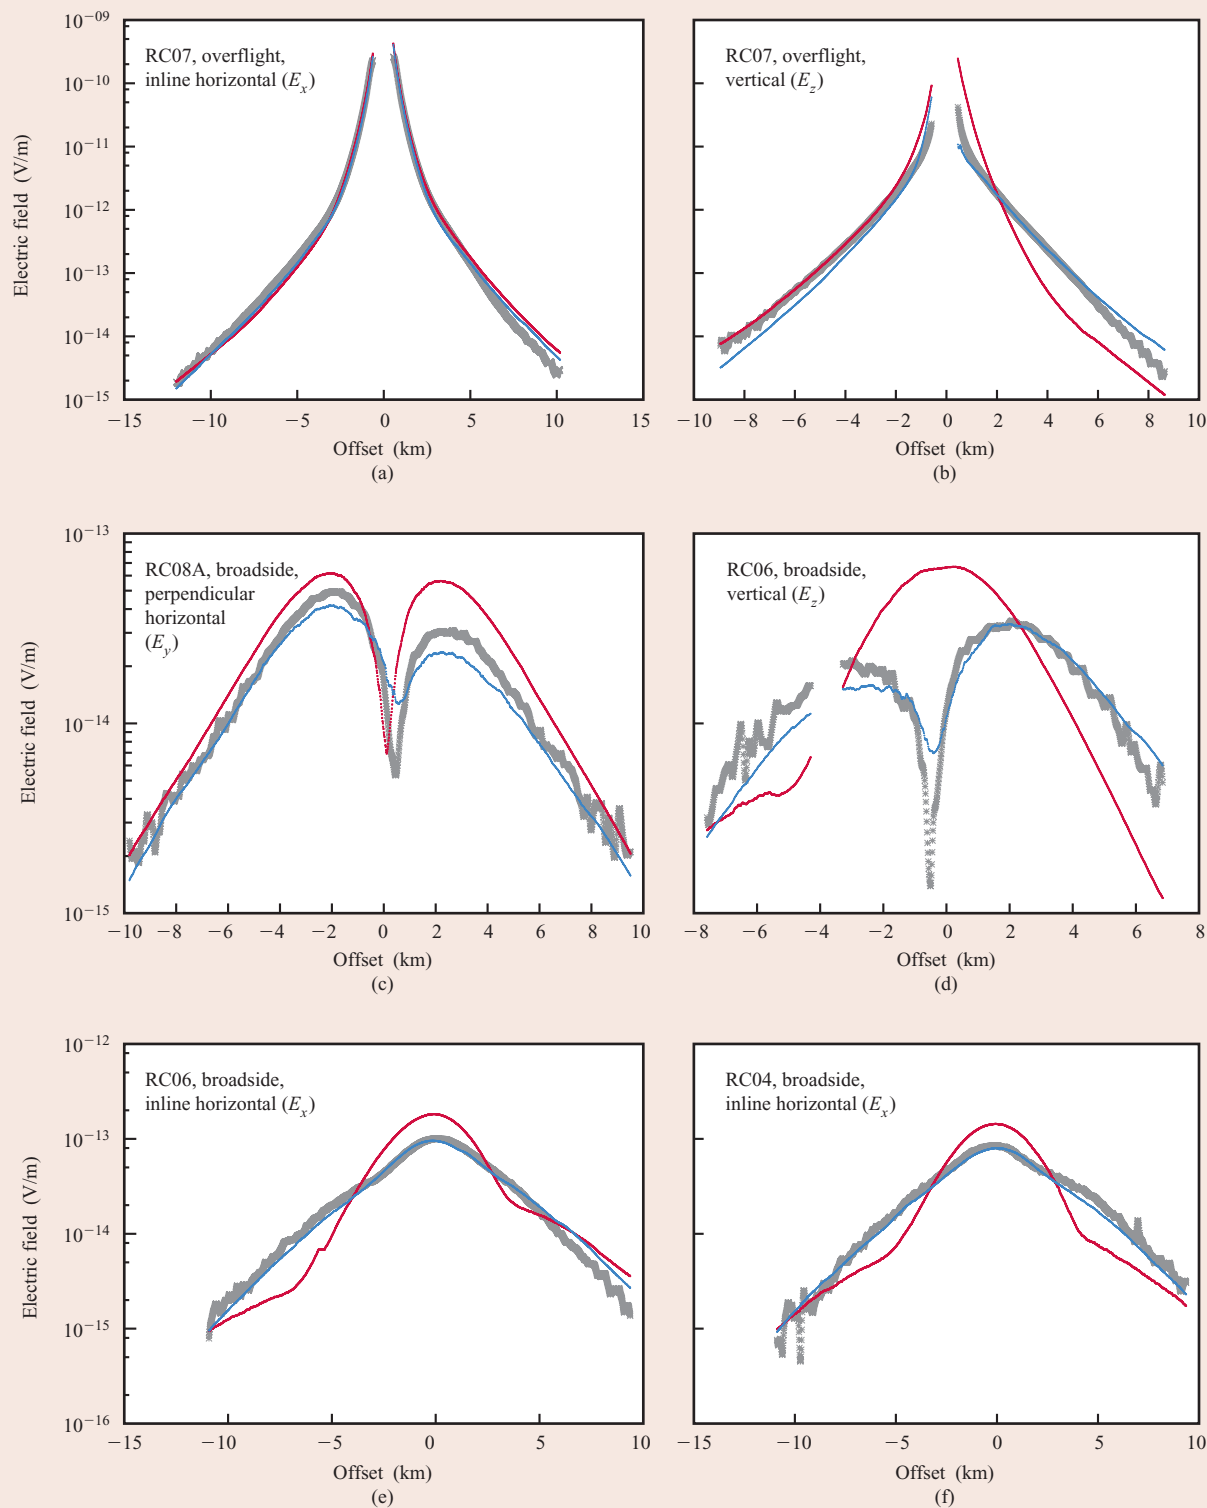


Figure 5

Six selected plots of overflight and broadside electric field data amplitudes (gray) vs. the transmitter offset projected onto the profile. Shown are data fits produced by a starting model with isotropic (red) and anisotropic (blue) electrical conductivity.

an anisotropic model with the vertical conductivity fixed at the conductivity used in the initial isotropic inversion and the horizontal conductivity set to three times the vertical conductivity below the water bottom. A sampling of the results is shown in **Figure 5**, confirming that the data is very likely to be significantly more consistent with an anisotropic conductivity model than with an isotropic one. Furthermore, we reran two inversions with a subset of the data, comprising 36 effective transmitters. Using the same isotropic starting model, the inversions differed by using an isotropic and anisotropic model parameterization. After 62 iterations, the anisotropic model achieved a final data fit that was 27% lower than the isotropic result. A complete anisotropic inversion of this data has yet to be carried out.

Conclusions

We have made significant progress in reducing the computational demands of large-scale 3D EM imaging problems. Exploiting multiple levels of parallelism over the data and model spaces and utilizing different meshing for field simulation and imaging provides a capability for solving large 3D imaging problems that cannot be addressed otherwise in a timely manner.

Results of the BG/L platform experiment for this offshore data show that the broadside inline component data displays a systematic bias that is most likely attributable to conductivity anisotropy between the vertical and horizontal directions. The other field components were satisfactorily fit by an isotropic model, showing that these field components are significantly less sensitive to this kind of anisotropy. The speed at which the Blue Gene/L supercomputer delivered this result is essential to the time frame in which the exploration process is conducted. This work provides motivation to extend the 3D conductivity imaging methodology to the anisotropic situation.

Acknowledgments

The authors gratefully acknowledge the donation of Blue Gene/L computing resources by the IBM Corporation. Base funding for this work was provided by the ExxonMobil Corporation and the U.S. Department of Energy, Office of Basic Energy Sciences, under contract DE-AC02-05CH11231. We also thank the German Alexander-von-Humboldt Foundation for support of Michael Commer through a Feodor-Lynen research fellowship. We acknowledge the contributions of our colleague Dr. Xinyou Lu, who provided the 1D inversion code and the contributions of our colleagues Dr. Dmitriy A. Pavlov and Dr. Charlie Jing of ExxonMobil who contributed many useful insights into the behavior of CSEM data in anisotropic conductivity models.

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of Intel Corporation in the United States, other countries, or both.

References

1. T. Eidesmo, S. Ellingsrud, L. M. Macgregor, S. Constable, M. C. Sinha, S. Johansen, F. N. Kong, and H. Westerdahl, "Sea Bed Logging (SBL), A New Method for Remote and Direct Identification of Hydrocarbon Filled Layers in Deepwater," *First Break* **20**, No. 3, 144–152 (2002).
2. S. Ellingsrud, T. Eidesmo, S. Johansen, M. C. Sinha, L. M. Macgregor, and S. Constable, "Remote Sensing of Hydrocarbon Layers by Seabed Logging (SBL): Results from a Cruise Offshore Angola," *The Leading Edge* **21**, No. 10, 972–982 (2002).
3. L. M. Macgregor and M. C. Sinha, "Use of Marine Controlled Source Electromagnetic Sounding for Sub-Basalt Exploration," *Geophysical Prospecting* **48**, No. 6, 1091–1106 (2000).
4. L. J. Srnka, J. J. Carazzone, M. S. Ephron, and E. A. Eriksen, "Remote Reservoir Resistivity Mapping," *The Leading Edge* **25**, No. 8, 972–975 (2006).
5. S. Constable, "Marine Electromagnetic Methods—A New Tool for Offshore Exploration," *The Leading Edge* **25**, No. 4, 438–444 (2006).
6. D. L. Alumbaugh, G. A. Newman, L. Prevost, and J. Shadid, "Three-Dimensional, Wideband Electromagnetic Modeling on Massively Parallel Computers," *Radio Science* **31**, No. 1, 1–23 (1996).
7. G. A. Newman and P. T. Boggs, "Solution Accelerators for Large-Scale Three-Dimensional Electromagnetic Inverse Problems," *Inverse Problems* **20**, No. 6, S151–S170 (2004).
8. M. Commer and G. A. Newman, "Large Scale 3D EM Inversion Using Optimized Simulation Grids Nonconforming to the Model Space," *Exp. Abstr. Soc. Expl. Geophys.* **25**, 760–764 (2006).
9. G. A. Newman and D. L. Alumbaugh, "Three-Dimensional Massively Parallel Electromagnetic Inversion—Part I. Theory," *Geophys. J. Int.* **128**, No. 2, 345–354 (1997).
10. S. Moskow, V. Druskin, T. Habashy, P. Lee, and S. Davdychewa, "A Finite Difference Scheme for Elliptic Equations with Rough Coefficients Using a Cartesian Grid Nonconforming to Interfaces," *SIAM J. Numer. Anal.* **36**, No. 2, 442–464 (1999).
11. J. J. Carazzone, O. M. Burtz, K. E. Green, and D. A. Pavlov, "Three Dimensional Imaging of Marine CSEM Data," *Exp. Abstr. Soc. Expl. Geophys.* **24**, 575–578 (2005).
12. K. E. Green, O. M. Burtz, L. A. Wahrmond, T. Clee, I. Gallegos, C. Xia, G. Zelewski, et al., "R3M Case Studies: Detecting Reservoir Resistivity in Complex Settings," *Exp. Abstr. Soc. Expl. Geophys.* **24**, 572–574 (2005).
13. J. D. Klein, P. R. Martin, and D. F. Allen, "The Petrophysics of Electrically Anisotropic Reservoirs," *Trans. Soc. Petrol. Well-Log Analysts (SPWLA); 36th Annual Logging Symposium* (1995).
14. J. Zhao, D. Zhou, X. Li, R. Chen, and C. Yang, "Laboratory Measurements and Applications of Anisotropy Parameters of Rocks," *Trans. Soc. Petrol. Well-Log Analysts (SPWLA); 35th Annual Logging Symposium* (1994).
15. G. M. Hoversten, G. A. Newman, N. Geier, and G. Flanagan, "3D Modeling of a Deepwater EM Exploration Survey," *Geophysics* **71**, No. 5, G239–G248 (2006).

Received March 14, 2007; accepted for publication April 25, 2007; Internet publication December 14, 2007

Michael Commer *Lawrence Berkeley National Laboratories, Berkeley, California 94720 (MCommer@lbl.gov)*. Dr. Commer received his Ph.D. degree in geophysics from the University of Cologne, Germany. His research focused on transient electromagnetic modeling and inversion and was awarded the German Klaus Liebrecht prize for outstanding dissertations. Current research areas include large-scale time- and frequency-domain modeling and data inversion using massively parallel computers. Since 2004, Dr. Commer has been employed with Lawrence Berkeley National Laboratories (LBNL), where he initially started as a post-doctoral fellow supported by the German Alexander von Humboldt Foundation.

Gregory A. Newman *Lawrence Berkeley National Laboratories, Berkeley, California 94720 (ganewman@lbl.gov)*. Dr. Newman received his Ph.D. degree in geophysics from the University of Utah in 1987. He is an expert in large-scale electromagnetic field modeling and inversion, with an emphasis on massively parallel implementations. In 2004, Dr. Newman accepted a Senior Scientist appointment within the Earth Sciences Division at LBNL. Previously, he was associated with Sandia National Laboratories and the Institute of Geophysics and Meteorology in Cologne, Germany.

James J. Carazzone *ExxonMobil Upstream Research Company, Houston, Texas 77252 (jim.j.carazzone@exxonmobil.com)*. Dr. Carazzone received his Ph.D. degree in physics from Harvard University in 1975 and had two postdoctoral appointments in the areas of elementary particle physics and quantum field theory at the Fermi National Accelerator Laboratory and at the Institute for Advanced Study at Princeton. In 1978, he joined the upstream research organization of Exxon, where he has worked ever since in the areas of seismic and electromagnetic modeling and inversion applied to hydrocarbon exploration. Dr. Carazzone is a member of the Society of Exploration Geophysicists and the European Association of Geophysicists and Engineers.

Thomas A. Dickens *ExxonMobil Upstream Research Company, Houston, Texas 77252 (tom.a.dickens@exxonmobil.com)*. Dr. Dickens received a B.S. degree in physics from the University of Virginia (1981) and a Ph.D. degree in physics from Princeton University (1987), where he was awarded the Joseph Henry Prize. From 1987 to 1990, he was employed by MIT Lincoln Laboratory, where he worked on synthetic aperture radar and infrared imaging techniques. He joined Exxon Production Research (now ExxonMobil Upstream Research) in 1990 and has performed research in the areas of tomography, parallel computing, imaging of complex structures, anisotropic depth migration, and electromagnetic inversion. His current research interests include seismic and electromagnetic imaging, signal processing, and modeling. He is a member of the Society of Exploration Geophysicists, the American Physical Society (APS), and the Society for Industrial and Applied Mathematics (SIAM).

Kenneth E. Green *ExxonMobil Upstream Research Company, Houston, Texas 77252 (ken.e.green@exxonmobil.com)*. Dr. Green received his undergraduate degree in geophysics from MIT in 1974 and a doctorate in oceanography from the MIT-Woods Hole Oceanographic Institution in 1980. He joined the Geoscience Group at Exxon Production Research Company (now ExxonMobil Upstream Research) in 1980. His work at ExxonMobil has covered many aspects of exploration in basin hydrocarbon systems, including the subsurface visualization of earth resistivity volumes applied to oil and gas migration,

entrapment, and production. Dr. Green is a member of the Society of Exploration Geophysicists.

Leslie A. Wahrmond *ExxonMobil Upstream Research Company, Houston, Texas 77252 (leslie.a.wahrmond@exxonmobil.com)*. Dr. Wahrmond received her B.A. degree in geological sciences from the University of California at Santa Barbara and her Ph.D. degree in geology from the University of Texas at Austin. She has worked at Exxon Production Research (now ExxonMobil Upstream Research) since 1991, primarily in seismic interpretation and seismic attribute analysis. She has worked on integration and interpretation of CSEM data since 2003.

Dennis E. Willen *ExxonMobil Upstream Research Company, Houston, Texas 77252 (denny.e.willen@exxonmobil.com)*. Dr. Willen joined Exxon Production Research Company in 1980, after receiving his Ph.D. degree in physics from the University of Illinois. Since then he has worked in several areas of exploration geophysics including seismic imaging, near-surface effects, converted waves, well logging, parallel computing, and electromagnetic methods. He is a member of the Society of Exploration Geophysicists and the International Trumpet Guild.

Janet Shiu *IBM Deep Computing Group, Two Riverway, Houston, Texas 77056 (jshiu@us.ibm.com)*. Dr. Shiu received her Ph.D. degree in physics from University of Pittsburgh. She was an assistant professor at Old Dominion University and a principal investigator of research grants from the NASA Langley Research Center, where her research focus was in the area of high-power lasers for space application. Since 1985, Dr. Shiu has been employed with Cray Research, Inc., SGI, and the Exxon Upstream Research Company in technical support of high-performance computing. Dr. Shiu joined IBM in 1999 and is a member of IBM Deep Computing technical team.