Lawrence Berkeley National Laboratory
Energy Geosciences

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

Permalink
https://escholarship.org/uc/item/3mp5q7pp

Journal
Lawrence Berkeley National Laboratory

Authors
Commer, M.
Newman, G.A.
Carazzone, J.J.
et al.

Publication Date
2008-06-24

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

MASSIVELY-PARALLEL ELECTRICAL-CONDUCTIVITY IMAGING OF HYDROCARBONS USING THE BLUE GENE/L SUPERCOMPUTER

M. Commer & G. A. Newman

Lawrence Berkeley National Laboratory

J. J. Carazzone, T. A. Dickens, K. E. Green, L. A. Wahrmund & D. E. Willen

ExxonMobil Upstream Research Company

J. Shiu

Deep Computing IBM Corporation

ABSTRACT

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/processors on the IBM Watson Research Blue Gene/L (BG/L) supercomputer. Over a 24-hour period, we were able to image a large-scale marine CSEM field data set that previously required over four months of computing time on distributed clusters utilizing 1024 tasks on an Infiniband fabric. The total initial

data misfit could be decreased by 67 % within 72 completed inversion iterations, indicating an electrically resistive region in the southern survey area below a depth of 1500 m below 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 proven very effective in mapping geologic reservoir formations. However, they 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 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 ($[\sigma]$=S/m), or its reciprocal ($[1/\sigma]=\Omega$m), usually called resistivity, of offshore oil and gas prospects [1, 2, 3, 4 and 5]. Resistivity is a more meaningful quantity for imaging hydrocarbons. An 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 to fluids like brine or water.

With the marine controlled-source electromagnetic (CSEM) measurement technique, a deep-towed electric-dipole transmitter is used to excite a low-frequency (~0.1 to 10 Hz) electromagnetic signal that is measured on the seafloor by electric and magnetic field detectors, where the largest transmitter-detector offsets can exceed 15 km. To cover larger depth ranges, multiple transmitter frequencies are usually employed in a survey. Similar to acoustic wave propagation, the attenuation rate with exploration depth increases with the frequency. Current technologies require low frequency EM signals (< 1 Hz) to interrogate down to reservoir depths as large 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 3D and exceedingly difficult to map without recourse to 3D EM imaging experiments, requiring fine model parameterizations, spatially exhaustive survey coverage and multi-component data. The 3D imaging problem, in this paper also referred to as inversion problem, usually has large computational demands, owing to the 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 over the modeling/imaging volume, and the other over the data volume. The algorithm is designed for arbitrarily large data sets, 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/modeling mesh from the simulation mesh. This provides for significant acceleration of the 3D EM field simulation, directly impacting the time to solution for the 3D imaging process.

Here, we report an imaging experiment, utilizing 32,768 tasks/processors on the IBM Watson Research BG/L supercomputer. The experiment is a novelty both in terms of computational resources utilized and amount of data inverted. Its main purpose is a feasibility study for the effectiveness of the employed algorithm. Further, the results obtained will improve both important base knowledge for the design of upcoming large-scale CSEM surveys and the automated imaging method for data interpretation.

PROBLEM FORMULATION

We formulate the inverse problem by finding a model **m** with *m* piecewise constant electrical conductivity parameters that describe the earth model reproducing a given data set. Specifically, the inversion algorithm minimizes the error functional,

$$\phi = \tfrac{1}{2} \{\mathbf{D}(\mathbf{d}^{p} - \mathbf{d}^{obs})^{\mathbf{T}*} \{\mathbf{D}(\mathbf{d}^{p} - \mathbf{d}^{obs})\} + \tfrac{1}{2} \lambda \{\mathbf{Wm}\}^{\mathbf{T}} \{\mathbf{Wm}\}, \qquad (1)$$

where $\mathbf{T}^{*}$ denotes the Hermitian conjugate operator. In the above expression, the predicted (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 are determined through solution of the time harmonic 3D Maxwell equations in the diffusive approximation. We have also introduced a diagonal weighting matrix, $\mathbf{D}_{nxn}$, into the error functional to compensate for noisy measurements. To stabilize the minimization of (1) and to reduce model curvature in three dimensions, we introduce a matrix $\mathbf{W}_{mxm}$ based upon a 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 **E**, and thus the data $\mathbf{d}^p$ for a given model **m**. EM wave propagation is controlled by the vector Helmholtz equation,

$$\nabla \times \nabla \times \mathrm{E} + i\omega\mu_0\sigma\mathrm{E} = -i\omega\mu_0 J \tag{2}$$

where source vector, free-space magnetic permeability, and angular frequency are denoted by **J**, μ_0, and ω, respectively (see [6] for specific details). Our solution method is based upon 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,

$$\mathbf{KE}=\mathbf{S} \tag{3}$$

where **K** is a sparse complex symmetric matrix with 13 non-zero entries per row [6]. The diagonal entries of **K** depend explicitly on the conductivity parameters that we seek to estimate through the inversion process. Since the electric field, **E**, also depends upon the conductivity, implicitly, this gives rise to the nonlinearity of the inverse problem. The fields are sourced with a grounded wire or loop embedded within the modeling domain, described by the discrete source vector, **S**, and includes 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, **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. 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 (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 decent, nonlinear conjugate gradient and limited memory quasi-Newton schemes, where the latter usually provide the best inverse solution convergence, however at a larger computational expense. Solution accelerators are discussed in [7], also providing detailed derivation of the gradients and an efficient scheme for their computation. Here, we focus on a non-linear conjugate gradient (NLCG) minimization approach as a tradeoff 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 data sets can contain up to hundreds of transmitter-receiver arrays, operating at different frequencies, with a spatial covering of more than 1000 km^2. 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 up the FD simulation grid, not the matrix, amongst 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 where the message passing overhead increase dominates. A study of the flop 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 (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 only needed in order to complete several dot products per inversion iteration and to sum up the contributions from each LC to the global gradient vector. The main computational and communication burden occurs with the forward FD

solves. As outlined below, we adapt FD mesh sizes according to given transmitter-receiver configurations and minimum spatial sampling requirements. To keep a balanced 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, to solve the very large problems of interest requires us to obtain a higher level of efficiency. One promising approach, which we have previously reported in [8], is to design an optimal FD simulation mesh for each source excitation in equation (3). FD meshing for field simulation then only considers 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 source activations contained in the data set. 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 **K** in equation (3) (see also [6]

and [9] for details). In the case $\Omega_m = \Omega_s$, an edge conductivity, σ^e, is computed from $\sigma^e = \sum_{i=1}^{4} \sigma_i w_i$, with $w_i = dV_i \Big/ \sum_{j=1}^{4} dV_j$. Here w_i are 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/serial circuit analysis resulting in an arithmetic and harmonic conductivity averaging scheme of [8,10]. The averaging scheme is exemplified for an *x*-directed edge conductivity σ_x^e in two dimensions in Figure 1. Here, model and simulation meshes are represented by dashed and solid lines, respectively. The material average is to be specified from the formula

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

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.

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

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

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. Now 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 conductivity and volume fraction, 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

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

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 $w_k = dV_k \Big/ \sum_{j=1}^{4} dV_j$.

ELECTRICAL-CONDUCTIVITY IMAGING OF HYDROCARBONS USING THE BLUE GENE/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 it 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 give rise to 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 provides ample justification for utilizing the IBM's 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 time scales acceptable to the exploration process.

The 3D imaging experiment we present here demonstrates the above mentioned points. The data were acquired offshore of South America. The sail lines and 23 detector locations on a 40×40 km^2 grid used for subsurface conductivity mapping are shown in Figure 2. Data was collected from nearly 1 million binned transmitter sites along the shown sail lines. 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, and two or more additional simulations for step control for each non-linear inversion update. To efficiently deal with this data volume, we employ reciprocity. The positions of the real CSEM transmitter along the sail line become the computational receiver profiles, and the real 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 0.125, 0.25, 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 is illustrated in the upper Figure 2. For the transmitter sail line position with respect to a given detector on the sea bottom, we consider the so-called overflight (a) configuration, where the sail line is directly over the detector. In the broadside configuration (b), the towed transmitter passes at an offset Δy to one side of the detector. Three components are recorded by the detector's receiver antennas: inline horizontal (E_x), perpendicular horizontal (E_y), and vertical (E_z) electric fields.

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 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. To accommodate all sail lines and detector sites in the model, a large parameterization was required for Ω_m. To model bathymetry, the minimum required spatial grid sampling interval Δ is kept constant with Δ=125 m for the horizontal, *x* and *y*, coordinates, while it ranges from 50 to 200 m in *z*. This amounts to 403 nodes along *x* and *y*, and 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 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 (about 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 σ_b=0.5 S/m; 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 vary with frequency, Δ=250, 200, and 125 m, for the frequencies *f*=0.125, 0.25, and 0.5 Hz, respectively. The vertical meshing was identical to that employed in the modeling mesh in order to honor the bathymetry. 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 where $\Omega_s=\Omega_m$.

A maximum of 256 Mbytes of memory per task was available on BG/L. 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 a local communicator size of 512 processors, relying on the inter-processor bandwidth to support the BiCG/QMR solves. Sixty-four local communicators were then used to distribute the 207 effective sources and its associated data. Thus the total number of tasks employed in the imaging experiment was 32,768. Disk IO and file system performance were minor concerns, as the generated image output was relatively modest, approximately 2.5 Gbytes 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 Mbytes. 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. Base line evaluation involved an inversion where the large model grid (size 403×403×173 nodes) represented the simulation grid for each source.

1) The job performance using 32 MPI tasks completed on BG/L (CPU speed 700 MHz) and an Intel (Pentium 4, CPU speed 2.6 GHz) cluster with Gigabit Ethernet fabric was compared. A forward solution used 25 sec per 100 QMR iterations on BG/L, compared to 23 sec on the Intel P4 platform. The computational burden of the QMR solver is dominated by complex double precision matrix-vector

multiplications with indexed memory access. BG/L's 64-bit IBM Power architecture is designed for floating point operations achieving an efficient memory access. Profiling shows that for our application the architecture compensates for BG/L's lower processor speed.

2) Workload scalability tests revealed a linear QMR solution time decrease up to a number of 4096 tasks.

3) A 1024-task job on BG/L showed that the communication averaged to 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% with typical message sizes of 16 Bytes. Point-to-point blocking message passing within LC: 20 % with 30 Kbytes average message size. Barrier synchronization: 30%.

The relatively long idle time 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 the condition number of the FD stiffness matrix **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, a faster barrier synchronization would require an adequate sophisticated scheme for dynamically adapting the LC size.

Over a 24-hour period, 72 inversion model updates were realized on BG/L and the relative squared error misfit measure was reduced by nearly 67%. 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 (a) and E_z^i (b), as well for the horizontal perpendicular and vertical broadside electric fields, E_y^b (c) and E_z^b (d). We observed that the major residual misfits originate from the broadside inline components, E_x^b (e,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 enhanced resistivity in the southern model section for depths below 1500 m. Such is also observed broadside of the sail lines, for the depth range 0-1500 m. Along the sail lines, however, little to no resistivity enhancement is observed and the imaged resistivity volume contains an unacceptable acquisition overprint. A possible explanation for this outcome is the inconsistencies observed in fitting the in-line 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 are more sensitive to resistivity variations off the sail lines. One possibility for the enhanced resistivity observed off the sail lines arises from the inversion algorithm's attempt 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; 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 and 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, [15] showed 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 plays a bigger factor in influencing marine CSEM measurement than previously believed.

Two tests were carried out to verify the importance of anisotropy. First, 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, we repeated the initial stage of the inversion process. This involved 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 are very likely significantly more consistent with an anisotropic conductivity model than with an isotropic one. Furthermore, we rerun 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, which was by 27 % lower, compared to the isotropic result. A complete anisotropic inversion of these 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 to solve large 3D imaging problems that cannot be addressed otherwise in a timely manner.

Results of the Blue Gene/L experiment for this offshore data showed 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 donation of Blue Gene/L computing resources by the IBM Corporation. Base funding for this work was provided by the ExxonMobil Corporation and the United States Department of Energy, Office of Basic Energy Sciences, under contract DE-AC02-05CH11231. We also wish to thank the German Alexander-von-Humboldt Foundation for support of Michael Commer through a Feodor-Lynen research fellowship. We wish to 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.

REFERENCES

1. T. Eidesmo, S. Ellingsgrud, 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 (March 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 (October 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 (November 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 (August 2006).

5. S. Constable, "Marine electromagnetic methods – A new tool for offshore exploration," *The Leading Edge*, **25**, No. 4, 438-444 (April 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 (January-February 1996).

7. G. A. Newman and P. T. Boggs, "Solution accelerators for large-scale three-dimensional electromagnetic inverse problems," *Inverse Problems*, **20**, doi:10.1088/0266-5611/20/6/S10, S151-S171 (2004).

8. M. Commer and G. A. Newman, "Large scale 3D EM inversion using optimized simulation grids non-conformal 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 (February 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. Number. Anal.*, **36**, No. 2, 442-464 (February 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. Wahrmund, T. Clee, I. Gallegos, C. Xia, G. Zelewski, A. A. Martinez, M. J. Stiver, C. M. Rodriguez, and J. Zhang, "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 Ann. Logging Symp.*, (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 Ann. Logging Symp.*, (1994).

15. G. M. Hoversten, G. A. Newman, N. Geier, and G. Flanagan, G., "3D modeling of a deepwater EM exploration survey," *Geophysics*, **71**, No. 5, G239-G248 (September-October 2006) .

Figure captions

Figure 1. Illustration of the conductivity averaging scheme of equation (4) in two dimensions.

Figure 2. Layout of the sail lines (red and blue) and 23 detector locations (crosses) on the sea bottom for the offshore CSEM survey. Contained survey configurations are illustrated in the upper figure. Bathymetry is given in meters below sea level. The example data shown in this paper corresponds to the transmitter-detector arrays marked in blue.

Figure 3. Six selected plots of overflight and broadside electric field data amplitudes (black curves) versus the transmitter offset projected onto the profile lines shown in Figure 2. Shown are data fits produced by the starting model (red) and for iteration 72 (blue).

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 1500 m BML, c) interval 1500 to 2500 m BML. Resistivity is rendered on a base 10 log scale.

Figure 5. Six selected plots of overflight and broadside electric field data amplitudes (black curves) versus the transmitter offset projected onto the profile. Shown are data fits produced by a starting model with isotropic (red) and anisotropic (blue) electrical conductivity.

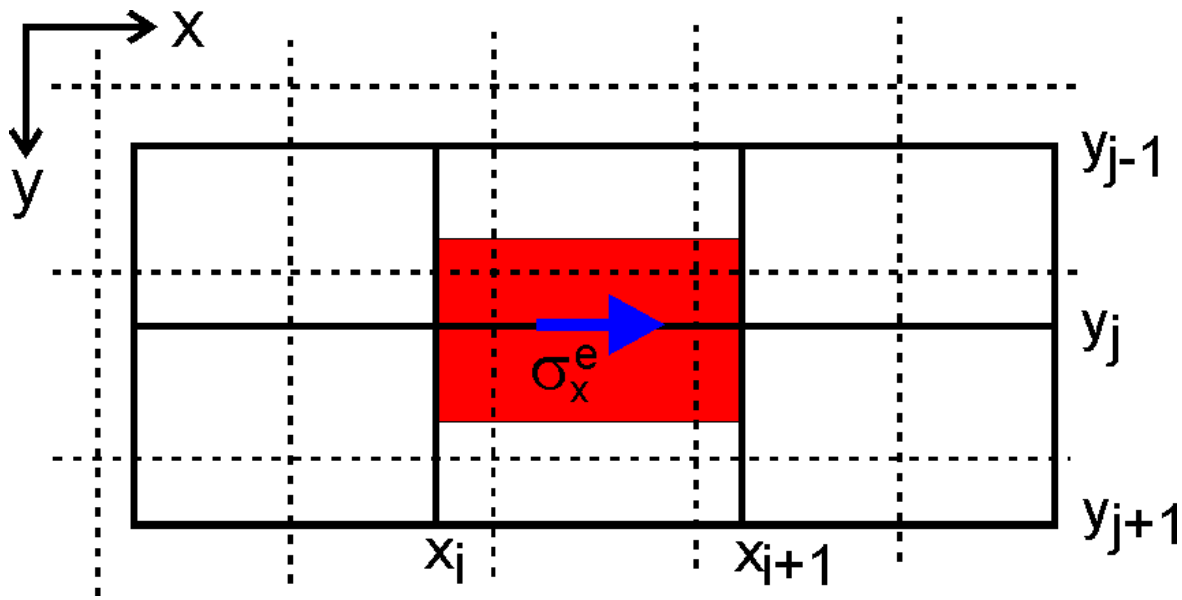

Fig. 1

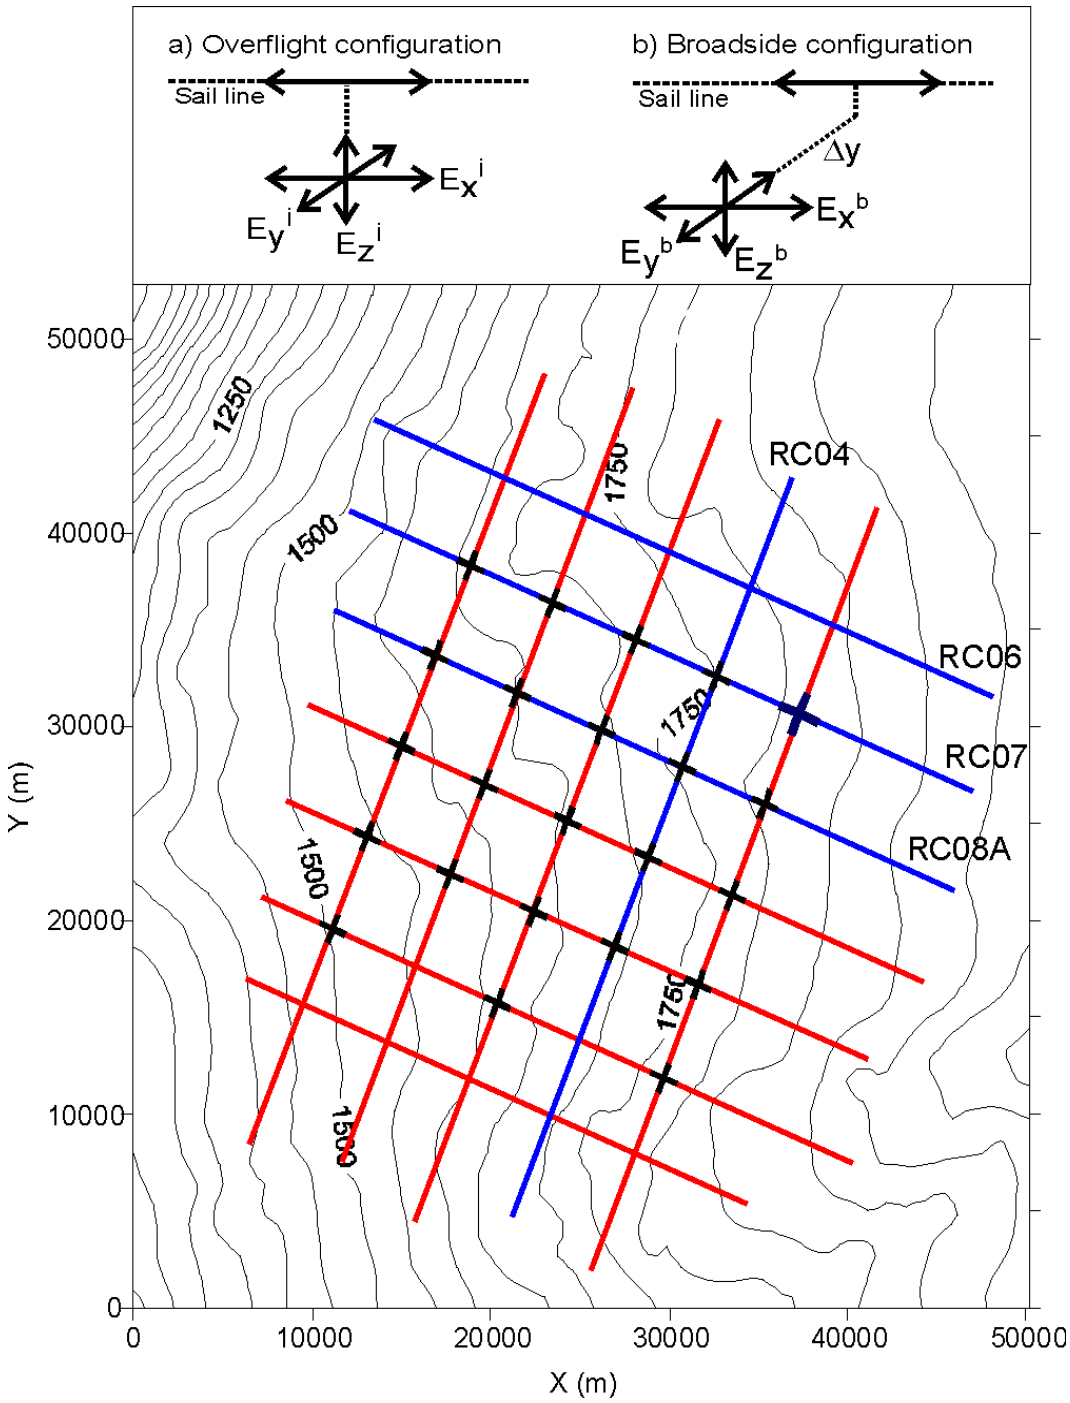
a) Overflight configuration
Sail line
E_x^i
E_y^i
E_z^i
b) Broadside configuration
Sail line
Δy
E_x^b
E_y^b
E_z^b
1250
1500
1750
RC04
RC06
RC07
RC08A
Y (m)
X (m)
0
10000
20000
30000
40000
50000

Fig. 2

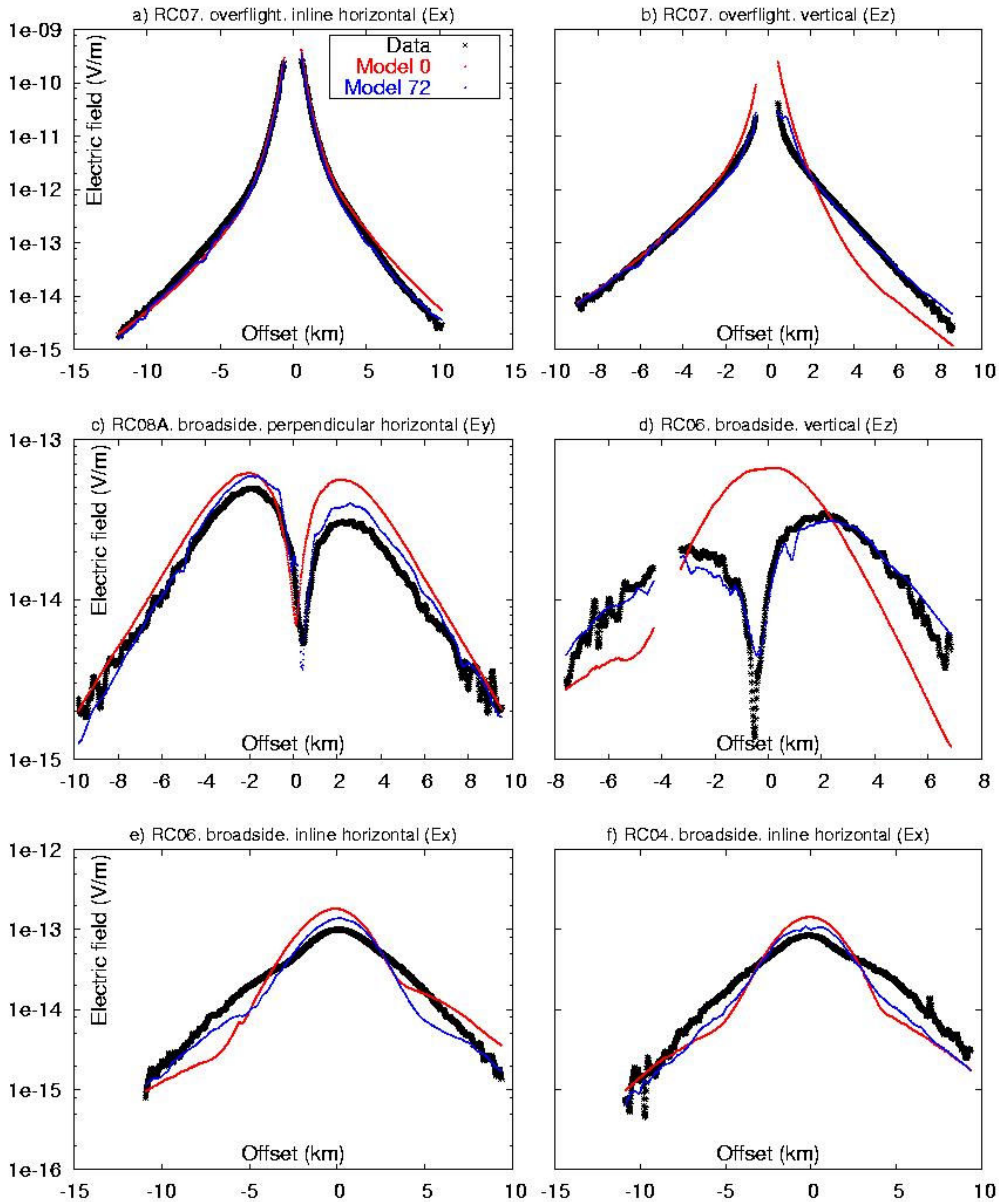
a) RC07. overflight. inline horizontal (Ex)
b) RC07. overflight. vertical (Ez)
c) RC08A. broadside. perpendicular horizontal (Ey)
d) RC06. broadside. vertical (Ez)
e) RC06. broadside. inline horizontal (Ex)
f) RC04. broadside. inline horizontal (Ex)
Data
Model 0
Model 72
Electric field (V/m)
Offset (km)

Fig. 3

a)

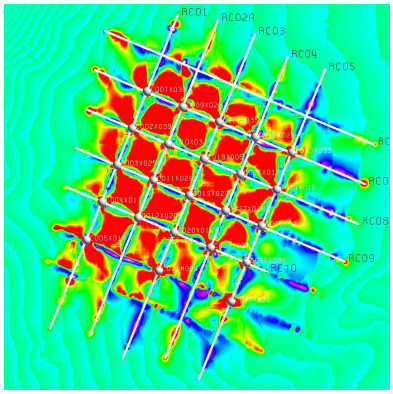
RC01
RC02A
RC03
RC04
RC05
RC08
RC09
RC10

b)

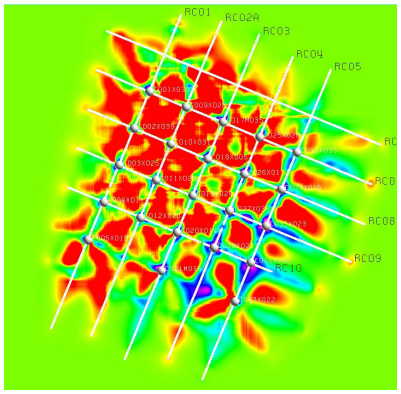
RC01
RC02A
RC03
RC04
RC05
RC08
RC09
RC10

c)

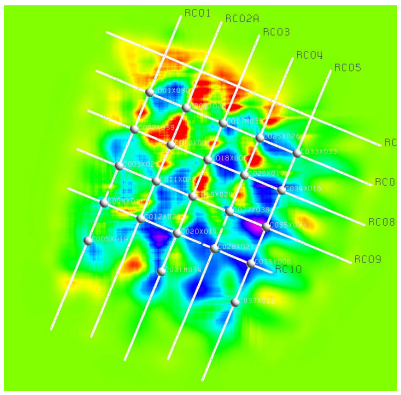
RC01
RC02A
RC03
RC04
RC05
RC08
RC09
RC10

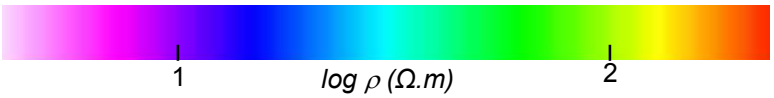
1
log ρ (Ω.m)
2

Fig. 4

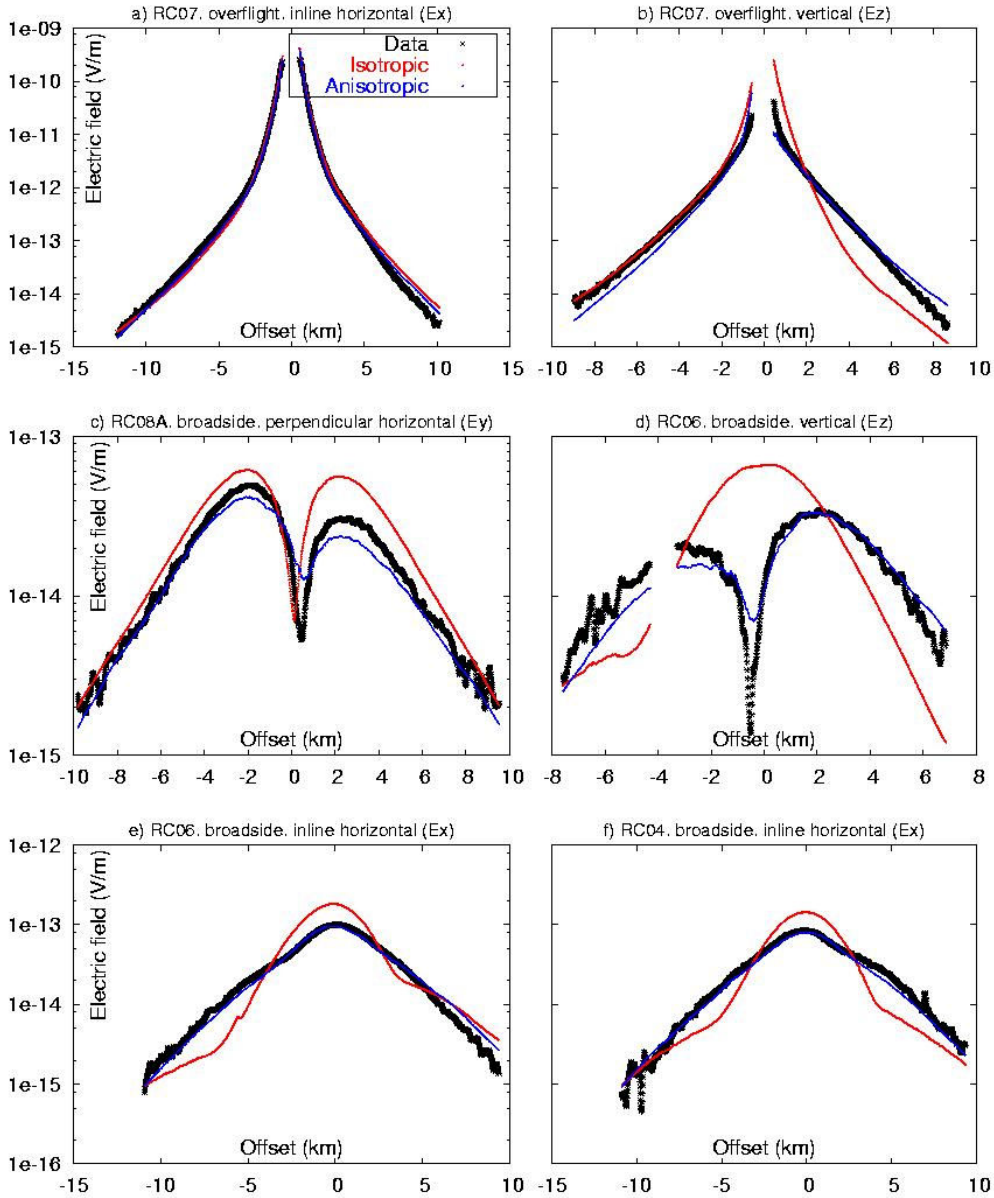
a) RC07. overflight. inline horizontal (Ex)
b) RC07. overflight. vertical (Ez)
c) RC08A. broadside. perpendicular horizontal (Ey)
d) RC06. broadside. vertical (Ez)
e) RC06. broadside. inline horizontal (Ex)
f) RC04. broadside. inline horizontal (Ex)
Data
Isotropic
Anisotropic
Electric field (V/m)
Offset (km)

Fig. 5

Michael Commer

Lawrence Berkeley National Laboratories, Berkeley, CA 94720 (MCommer@lbl.gov).

M. 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 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, CA 94720 (ganewman@lbl.gov).

Gregory 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 Lawrence Berkeley National Laboratory. Previously he has been associated with Sandia National Laboratories and Institute of Geophysics and Meteorology, Cologne Germany.

James J. Carazzone

ExxonMobil Upstream Research Company, Houston, TX 77252

(jim.j.carazzone@exxonmobil.com). James Carazzone received his Ph.D. degree in Physics from Harvard University in 1975 and had two post-doctoral 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 Exxon's upstream research organization 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 of the European Association of Geophysicists and Engineers.

Thomas A. Dickens

ExxonMobil Upstream Research Company, Houston, TX 77252

(tom.a.dickens@exxonmobil.com). Thomas Dickens received a B.S. in Physics from the University of Virginia (1981) and a Ph.D. in Physics from Princeton University in 1987, where he was awarded the Joseph Henry Prize. From 1987-1990 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, and 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 SEG, APS, and SIAM.

Kenneth E. Green

ExxonMobil Upstream Research Company, Houston, TX 77252

(ken.e.green@exxonmobil.com). Ken 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 function at Exxon Production Research Company 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. Ken is a member of the Society of Exploration Geophysicists.

Leslie A. Wahrmund

ExxonMobil Upstream Research Company, Houston, TX 77252

(leslie.a.wahrmund@exxonmobil.com). Leslie Wahrmund received her B.A. degree in Geological Sciences from the University of California at Santa Barbara, and her Ph.D 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, TX 77252

(denny.e.willen@exxonmobil.com). Dr. Willen joined Exxon Production Research Company in 1980, after receiving his Ph.D. 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, Two Riverway, Houston, Texas 77056

(jshiu@us.ibm.com). Dr. Shiu received her Ph.D. degree in Physics from University of Pittsburgh. Dr. Shiu was an assistant professor at Old Dominion University, and a principal investigator of research grants from NASA Langley Research Center where her research focus were in the area of high power lasers for space application. Since 1985, Dr. Shiu has been employed with Cray Research Inc., SGI, and Exxon Upstream Technical Company in technical support of High Performance Computing. Dr. Shiu joined IBM in 1999, and is a member of IBM Deep Computing Technical Team.