Lawrence Berkeley National Laboratory

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

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

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

Publication Date
2008-06-24

Peer reviewed

1 **MASSIVELY-PARALLEL ELECTRICAL-CONDUCTIVITY IMAGING**

2 **OF HYDROCARBONS USING THE BLUE GENE/L SUPERCOMPUTER**

3

4 **M. Commer & G. A. Newman**

5 *Lawrence Berkeley National Laboratory*

6

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

8 *ExxonMobil Upstream Research Company*

9

10 **J. Shiu**

11 *Deep Computing IBM Corporation*

12

13 **ABSTRACT**

14 Large-scale controlled source electromagnetic (CSEM) three-dimensional (3D)

15 geophysical imaging is now receiving considerable attention for electrical conductivity

16 mapping of potential offshore oil and gas reservoirs. To cope with the typically large

17 computational requirements of the 3D CSEM imaging problem, our strategies exploit

18 computational parallelism and optimized finite-difference meshing. We report on an

19 imaging experiment, utilizing 32,768 tasks/processors on the IBM Watson Research Blue

20 Gene/L (BG/L) supercomputer. Over a 24-hour period, we were able to image a large-

21 scale marine CSEM field data set that previously required over four months of computing

22 time on distributed clusters utilizing 1024 tasks on an Infiniband fabric. The total initial

23   data misfit could be decreased by 67 % within 72 completed inversion iterations,

24   indicating an electrically resistive region in the southern survey area below a depth of

25   1500 m below the seafloor. The major part of the residual misfit stems from transmitter-

26   parallel receiver components that have an offset from the transmitter sail line (broadside

27   configuration). Modeling confirms that improved broadside data fits can be achieved by

28   considering anisotropic electrical conductivities. While delivering a satisfactory gross-

29   scale image for the depths of interest, the experiment provides important evidence for the

30   necessity of discriminating between horizontal and vertical conductivities for maximally

31   consistent 3D CSEM inversions.

32

33   **INTRODUCTION**

34   Seismic methods have a long and established history in hydrocarbon, i.e. oil and gas,

35   exploration, and are proven very effective in mapping geologic reservoir formations.

36   However, they are not good at discriminating the different types of reservoir fluids

37   contained in the rock pore space, such as brines, water, oil and gas. This has encouraged

38   the development of new geophysical technologies that can be combined with established

39   seismic methods to directly image fluids. One technique that has recently emerged, with

40   considerable potential, utilizes low frequency electromagnetic (EM) energy to map

41   variations in the subsurface electrical conductivity, σ ($[\sigma]$=S/m), or its reciprocal

42   ($[1/\sigma]$=Ωm), usually called resistivity, of offshore oil and gas prospects [1, 2, 3, 4 and 5].

43   Resistivity is a more meaningful quantity for imaging hydrocarbons. An increase,

44   compared to the surrounding geological strata, may directly indicate potential reservoirs.

45 EM field measurements have been shown to be highly sensitive to changes in the pore

46 fluid types and the location of hydrocarbons, given a sufficient resistivity contrast to fluids

47 like brine or water.

48 With the marine controlled-source electromagnetic (CSEM) measurement technique, a

49 deep-towed electric-dipole transmitter is used to excite a low-frequency (~0.1 to 10 Hz)

50 electromagnetic signal that is measured on the seafloor by electric and magnetic field

51 detectors, where the largest transmitter-detector offsets can exceed 15 km. To cover larger

52 depth ranges, multiple transmitter frequencies are usually employed in a survey. Similar to

53 acoustic wave propagation, the attenuation rate with exploration depth increases with the

54 frequency. Current technologies require low frequency EM signals (< 1 Hz) to interrogate

55 down to reservoir depths as large as 4 km.

56 Exploration with the CSEM technology in the search for hydrocarbons now extends to

57 highly complex and subtle offshore geological environments. The geometries of the

58 reservoirs are inherently 3D and exceedingly difficult to map without recourse to 3D EM

59 imaging experiments, requiring fine model parameterizations, spatially exhaustive survey

60 coverage and multi-component data. The 3D imaging problem, in this paper also referred

61 to as inversion problem, usually has large computational demands, owing to the expensive

62 solution of the forward modeling problem, that is the EM field simulation on a given 3D

63 finite-difference (FD) grid. Moreover, large data volumes require many forward solutions

64 in an iterative inversion scheme. Therefore, we have developed an imaging algorithm that

65 utilizes two levels of parallelization, one over the modeling/imaging volume, and the other

66 over the data volume. The algorithm is designed for arbitrarily large data sets, allowing for

67 an arbitrarily large number of parallel tasks, while the computationally idle message

68 passing is minimized. We have further incorporated an optimal meshing scheme that

69 allows us to separate the imaging/modeling mesh from the simulation mesh. This provides

70 for significant acceleration of the 3D EM field simulation, directly impacting the time to

71 solution for the 3D imaging process.

72 Here, we report an imaging experiment, utilizing 32,768 tasks/processors on the IBM

73 Watson Research BG/L supercomputer. The experiment is a novelty both in terms of

74 computational resources utilized and amount of data inverted. Its main purpose is a

75 feasibility study for the effectiveness of the employed algorithm. Further, the results

76 obtained will improve both important base knowledge for the design of upcoming large-

77 scale CSEM surveys and the automated imaging method for data interpretation.

78

79

80   **PROBLEM FORMULATION**

81 We formulate the inverse problem by finding a model $\mathbf{m}$ with m piecewise constant

82 electrical conductivity parameters that describe the earth model reproducing a given data

83 set. Specifically, the inversion algorithm minimizes the error functional,

84
$$\phi = \tfrac{1}{2}\,\{\mathbf{D}(\mathbf{d}^{p} - \mathbf{d}^{obs})^{\mathbf{T}*}\{\mathbf{D}(\mathbf{d}^{p} - \mathbf{d}^{obs})\} + \tfrac{1}{2}\,\lambda\,\{\mathbf{Wm}\}^{\mathbf{T}}\{\mathbf{Wm}\}, \qquad (1)$$

85 where $\mathbf{T}^{*}$ denotes the Hermitian conjugate operator. In the above expression, the predicted

86 (from a starting model) and observed data vectors are denoted by $\mathbf{d}^{p}$ and $\mathbf{d}^{obs}$, respectively,

87 where each has n complex values. These vectors consist of electric or magnetic field

88   values specified at the measurement points, where the predicted data are determined

89   through solution of the time harmonic 3D Maxwell equations in the diffusive

90   approximation. We have also introduced a diagonal weighting matrix, $\mathbf{D}_{nxn}$, into the error

91   functional to compensate for noisy measurements. To stabilize the minimization of (1) and

92   to reduce model curvature in three dimensions, we introduce a matrix $\mathbf{W}_{mxm}$ based upon a

93   FD approximation to the Laplacian (∇^2) operator applied in Cartesian coordinates. The

94   parameter λ attempts to balance the data error and the model smoothness constraint.

95

96   **The Forward Problem**

97   Within an inversion framework, the forward problem is solved multiple times to simulate

98   the EM field, denoted by the vector $\mathbf{E}$, and thus the data $\mathbf{d}^p$ for a given model $\mathbf{m}$. EM wave

99   propagation is controlled by the vector Helmholtz equation,

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

100   where source vector, free-space magnetic permeability, and angular frequency are denoted

101   by $\mathbf{J}$, μ_0, and ω, respectively (see [6] for specific details). Our solution method is based

102   upon the consideration that the number of model parameters required to simulate realistic

103   3D distributions of the electrical conductivity σ can typically exceed 10^7. FD modeling

104   schemes are ideally suited for this task and can be parallelized to handle large-scale

105   problems that cannot be easily treated otherwise [6]. After approximating equation (**2**) on a

106   staggered grid at a specific angular frequency, using finite differencing and eliminating the

107   magnetic field, we obtain a linear system for the electric field,

108

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

109

110   where $\mathbf{K}$ is a sparse complex symmetric matrix with 13 non-zero entries per row [6]. The

111   diagonal entries of $\mathbf{K}$ depend explicitly on the conductivity parameters that we seek to

112   estimate through the inversion process. Since the electric field, $\mathbf{E}$, also depends upon the

113   conductivity, implicitly, this gives rise to the nonlinearity of the inverse problem. The

114   fields are sourced with a grounded wire or loop embedded within the modeling domain,

115   described by the discrete source vector, $\mathbf{S}$, and includes Dirichlet boundary conditions

116   imposed upon the problem. To help avoid excessive meshing near the source, we favor a

117   scattered-field formulation to the forward modeling problem. In this instance, $\mathbf{E}$ is

118   replaced with $\mathbf{E}_s$ in equation (3). The source term, for a given transmitter, will now depend

119   upon the difference between the 3D conductivity model and a simple background model,

120   weighted by the background electric field $\mathbf{E}_b$, where $\mathbf{E} = \mathbf{E}_b + \mathbf{E}_s$. Simple background

121   models with one-dimensional (1D) conductivity distributions, i.e. σ changes only with

122   depth, are used because fast semi-analytical solutions for $\mathbf{E}_b$ are available. Given the

123   solution of the electric field in equation (3), the magnetic field can be easily determined

124   from a numerical implementation of Faraday's law. An efficient solution process is

125   paramount. We solve equation (3) to a predetermined error level using iterative Krylov

126   subspace methods, using either a biconjugate gradient (BICG) or quasi-minimum residual

127   (QMR) scheme with preconditioning [6].

128

Minimization Procedure

129

130   In large-scale nonlinear inverse problems, as considered here, we minimize (1) using

131   gradient-based optimization techniques because of their minimal storage and

132   computational requirements. We characterize these methods as gradient-based techniques

133   because they employ only first derivative information of the error functional in the

134   minimization process, specifically $-\nabla\phi$. Gradient-based methods include steepest decent,

135   nonlinear conjugate gradient and limited memory quasi-Newton schemes, where the latter

136   usually provide the best inverse solution convergence, however at a larger computational

137   expense. Solution accelerators are discussed in [7], also providing detailed derivation of

138   the gradients and an efficient scheme for their computation. Here, we focus on a non-linear

139   conjugate gradient (NLCG) minimization approach as a tradeoff between inverse solution

140   convergence and computational effort per inversion iteration.

141

Exploitation of Solution Parallelism

143   In order to realistically image the subsurface of large survey areas at a sufficient level of

144   resolution and detail, industrial CSEM data sets can contain up to hundreds of transmitter-

145   receiver arrays, operating at different frequencies, with a spatial covering of more than

146   1000 km^2. This easily requires thousands of solutions to the forward modeling problem for

147   just one imaging experiment. Hence, the computational demands for solving the 3D

148   inverse problem are enormous. To cope with this problem, our algorithm utilizes two

149   levels of parallelization, one over the modeling domain, and the other over the data

150   volume.

151 First, in solving the forward problem on a distributed environment, we split up the FD

152 simulation grid, not the matrix, amongst a Cartesian processor topology, which shall be

153 called local communicator (LC). As the linear system is relaxed during the iterative

154 solution, which involves matrix-vector products on each of the processors, values of the

155 solution vector at the current Krylov iteration not stored on the processor must be passed

156 by neighbors within LC to complete the matrix-vector products. Additional global

157 communication across the LC is needed to complete several dot products at each

158 relaxation step of the Krylov iteration. The solution time increases linearly with the

159 number of parallel tasks, up to a point where the message passing overhead increase

160 dominates. A study of the flop rate versus communicator size for the Intel Paragon

161 architecture is exemplified in [6].

162 To carry out many forward simulations simultaneously, we employ multiple LCs,

163 connected via a group of lead processors, with one lead task assigned to each LC. The

164 topology of this lead group defines the communicator on which the iterative NLCG

165 inversion framework is carried out, here called the global communicator (GC). This

166 distribution of the forward modeling problems, or data decomposition, is highly parallel.

167 Assuming the optimal LC size has been estimated for a given range of mesh sizes, the size

168 of the GC (equals the number of LCs) can be increased linearly with the data volume. The

169 relative amount of communication within the GC remains constant, because

170 communication within the GC is only needed in order to complete several dot products per

171 inversion iteration and to sum up the contributions from each LC to the global gradient

172 vector. The main computational and communication burden occurs with the forward FD

173 solves. As outlined below, we adapt FD mesh sizes according to given transmitter-receiver

174 configurations and minimum spatial sampling requirements. To keep a balanced workload

175 between all LCs, the data decomposition is based on a balanced distribution of the FD

176 grids in terms of grid sizes.

177

178 **Optimal Mesh Considerations**

179 Although our experience using two parallelization levels has been satisfactory, to solve the

180 very large problems of interest requires us to obtain a higher level of efficiency. One

181 promising approach, which we have previously reported in [8], is to design an optimal FD

182 simulation mesh for each source excitation in equation (3). FD meshing for field

183 simulation then only considers part of the total model volume where it can have an

184 appreciable effect in the imaging process. Moreover, minimum spatial grid sampling

185 intervals are dictated by the EM field wavelength, and hence can be optimized according

186 to a specific source excitation frequency. Optimizing both mesh size and spatial sampling,

187 we create a collection of simulation grids, Ω_s, that support the EM field simulation for all

188 different source activations contained in the data set. All simulation grids act upon a

189 common model grid, Ω_m, which defines the imaging volume. Both types of grids are

190 Cartesian with conformal grid axes. Key to the grid separation is an appropriate mapping

191 scheme that transfers the material properties from Ω_m to Ω_s. The imaging process provides

192 piecewise constant estimates of the electrical conductivity, which are defined by the cells

193 of Ω_m. The staggered FD mesh Ω_s, on the other hand, involves edge-based directional

194 conductivities, needed for constructing the stiffness matrix **K** in equation (3) (see also [6]

195   and [9] for details). In the case $\Omega_m = \Omega_s$, an edge conductivity, σ^e, is computed from

196    $\sigma^e = \sum_{i=1}^{4} \sigma_i w_i$, with $w_i = dV_i \Big/ \sum_{j=1}^{4} dV_j$. Here w_i are weights corresponding to volume

197   fractions of the four cells on Ω_m, that share the edge σ^e on Ω_s. Furthermore, the edge

198   conductivity σ^e is simply an arithmetic volume average of the four model cell

199   conductivities. When $\Omega_m \neq \Omega_s$, the conductivity mapping involves parallel/serial circuit

200   analysis resulting in an arithmetic and harmonic conductivity averaging scheme of [8,10].

201   The averaging scheme is exemplified for an x-directed edge conductivity σ_x^e in two

202   dimensions in Figure 1. Here, model and simulation meshes are represented by dashed and

203   solid lines, respectively. The material average is to be specified from the formula

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

204   The inner integration constitutes a point wise parallel conductivity average, while the outer

205   integration provides for the effective conductivity in series, arising over the integrated

206   edge length $(x_{i+1} - x_i)$ of the simulation mesh. The total integration area assigned to σ_x^e is

207   shown by the red rectangle.

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

209   exemplified by

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

210   where ΔX is the edge length of the simulation cell along the *x*-coordinate direction.

211   Similarly, σ_y^e and σ_z^e involve averaging along the *y*- and *z*-coordinates, respectively. Now

212   the averaging along ΔX involves a number of *J* serially connected discrete parallel

213   circuits, P_j, each with a volume V_j. The length of P_j along the edge is Δx_j,

214   where $\sum_{j=1}^{J} \Delta x_j = \Delta X$. Further, I_j is the number of cells on the modeling grid contributing

215   to P_j, with σ_i and dV_i the individual model cell conductivity and volume fraction,

216   respectively.

217   We are also required to specify $\partial\sigma^e / \partial\sigma_k$ which is needed to define the gradient on the

218   modeling grid, because it is linked to the forward modeling problem on the simulation

219   grid(s) (see [9] for details on the equal-grid case). Thus

$$\partial\sigma^e / \partial\sigma_k = \frac{\sigma^{e2}}{\Delta X} \sum_{j=1}^{J} \Delta x_j \left(\frac{1}{V_j} \sum_{i=1}^{I_j} dV_i \sigma_i \right)^{-2} \frac{dV_k}{V_j}, \tag{6}$$

220   where *J* is now the number of discrete parallel circuits with a non-zero contribution

221   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

222   weighting coefficient defined above as $w_k = dV_k \bigg/ \sum_{j=1}^{4} dV_j$.

223

224   **ELECTRICAL-CONDUCTIVITY IMAGING OF HYDROCARBONS USING**

225   **THE BLUE GENE/L SUPERCOMPUTER**

226   CSEM data is usually characterized by a large dynamic range, which can reach more than

227   ten orders of magnitude. This requires the ability to analyze it in a self-consistent manner

228   that incorporates all structures not only on the reservoir scale at tens of meters, but on the

229   geological basin scale at tens of kilometers, and must include salt domes, detail

230   bathymetry, and other 3D peripheral geology structures that can influence the

231   measurements [11, 12]. These complications give rise to the need for an automated 3D

232   conductivity inversion process for successful conductivity imaging of hydrocarbons. Trial-

233   and-error 3D forward modeling is too cumbersome to be effective. Both model size and

234   amount of the required data provides ample justification for utilizing the IBM's massively

235   parallel BG/L supercomputer for the task. Such a platform which can scale up to 131,072

236   processors, allows for the capability to image prospective oil and gas reservoirs at the

237   highest resolution possible, and on time scales acceptable to the exploration process.

238   The 3D imaging experiment we present here demonstrates the above mentioned points.

239   The data were acquired offshore of South America. The sail lines and 23 detector locations

240   on a 40×40 km^2 grid used for subsurface conductivity mapping are shown in Figure 2.

241   Data was collected from nearly 1 million binned transmitter sites along the shown sail

242   lines. Obviously, this amount of data cannot be treated with the current inversion

243   methodology even with a massively parallel implementation. Every source treated by the

244   imaging algorithm requires a forward simulation, an adjoint computation, and two or more

245   additional simulations for step control for each non-linear inversion update. To efficiently

246   deal with this data volume, we employ reciprocity. The positions of the real CSEM

247   transmitter along the sail line become the computational receiver profiles, and the real

248   CSEM detectors on the seafloor become computational sources, referred to as sources in

249   the following.

250   The equivalent reciprocal problem involves 951,423 data points and 207 effective sources,

251   since there are 23 source locations with three polarizations and each operating at the three

252   discrete excitation frequencies 0.125, 0.25, and 0.5 Hz. Each effective transmitter is

253   polarized according to the antenna orientation of its corresponding detector. The exact

254   seafloor detector orientations were determined by analyzing the data polarizations and

255   phase reversals with respect to the source sail lines. Data processing involves binning in

256   time, followed by spectral decomposition and spatial filtering. Timing errors were

257   removed by forcing the data phases to match the frequency-offset scaling behavior

258   appropriate to solutions of Maxwell's equations.

259   The survey layout in Figure 2 contains different transmitter-receiver configurations to be

260   considered, as is illustrated in the upper Figure 2. For the transmitter sail line position with

261   respect to a given detector on the sea bottom, we consider the so-called overflight (a)

262   configuration, where the sail line is directly over the detector. In the broadside

263   configuration (b), the towed transmitter passes at an offset Δy to one side of the detector.

264   Three components are recorded by the detector's receiver antennas: inline horizontal (E_x),

265   perpendicular horizontal (E_y), and vertical (E_z) electric fields.

266   A starting model is necessary to launch the inversion process and resolve some final issues

267   associated with phase components in the data. It is obviously favorable to achieve

268   minimum data misfits with the starting model. Therefore, the model used has been

269   constructed from knowledge of the sea bottom bathymetry, the seawater electrical

270   conductivity-versus-depth profile, and 1D inversion of the amplitude components of the

271   common-receiver gathers, based on the inline overflight measurement configuration (E_x^i).

272   The resulting 1D models were then refined by comparing selected simulation results with

273   field observations. To accommodate all sail lines and detector sites in the model, a large

274   parameterization was required for Ω_m. To model bathymetry, the minimum required

275   spatial grid sampling interval Δ is kept constant with $\Delta=125$ m for the horizontal, x and y,

276   coordinates, while it ranges from 50 to 200 m in z. This amounts to 403 nodes along x and

277   y, and 173 nodes vertically, and thus approximately 27.8 million model cells.

278   To restrict the size of the simulation grid for each source activation, we have assigned each

279   a separate mesh. Both mesh size and spatial grid sampling rate are based on skin depth

280   estimations. The skin depth δ, a commonly used constant in EM applications, is defined as

281   the depth below the surface of a conductor (in our case at the transmitter location) at which

282   the current density decays to $1/e$ (about 0.37) of the surface current density. Using the

283   approximation,

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

285   mesh intervals depend on the source excitation frequency f and the background

286   conductivity σ_b of the employed starting model. Horizontal mesh size is based on ten skin

287   depths from the source midpoint, assuming $\sigma_b=0.5$ S/m; the resulting mesh ranges were of

288   sufficient size to accommodate the specific sail lines of data assigned to the effective

289   sources. The horizontal spatial grid sampling intervals vary with frequency, $\Delta=250$, 200,

290   and 125 m, for the frequencies $f=0.125$, 0.25, and 0.5 Hz, respectively. The vertical

291   meshing was identical to that employed in the modeling mesh in order to honor the

292   bathymetry. With these considerations, we were able to reduce the size of the simulation

293   meshes significantly; the number of x and y grid nodes both ranged from 128 to 162.

294   Solution accuracy was verified against solutions where $\Omega_s = \Omega_m$.

295   A maximum of 256 Mbytes of memory per task was available on BG/L. The largest

296   memory requirement results from temporary storage of the forward solutions within one

297   inversion iteration. To stay within the machine limits each simulation grid was distributed

298   across a local communicator size of 512 processors, relying on the inter-processor

299   bandwidth to support the BiCG/QMR solves. Sixty-four local communicators were then

300   used to distribute the 207 effective sources and its associated data. Thus the total number

301   of tasks employed in the imaging experiment was 32,768. Disk IO and file system

302   performance were minor concerns, as the generated image output was relatively modest,

303   approximately 2.5 Gbytes per inversion update, which was written to disk in parallel using

304   512 tasks. Data output at each inversion iteration consisted of predicted and observed

305   measurements with a total file size of 170 Mbytes. A lead task within the global

306   communicator was assigned to dump the data output after each inversion update.

307   Prior to the actual imaging experiment, performance tests were carried out. Base line

308   evaluation involved an inversion where the large model grid (size 403×403×173 nodes)

309   represented the simulation grid for each source.

310   1) The job performance using 32 MPI tasks completed on BG/L (CPU speed 700

311   MHz) and an Intel (Pentium 4, CPU speed 2.6 GHz) cluster with Gigabit Ethernet

312   fabric was compared. A forward solution used 25 sec per 100 QMR iterations on

313   BG/L, compared to 23 sec on the Intel P4 platform. The computational burden of

314   the QMR solver is dominated by complex double precision matrix-vector

315   multiplications with indexed memory access. BG/L's 64-bit IBM Power

316   architecture is designed for floating point operations achieving an efficient memory

317   access. Profiling shows that for our application the architecture compensates for

318   BG/L's lower processor speed.

319   2) Workload scalability tests revealed a linear QMR solution time decrease up to a

320   number of 4096 tasks.

321   3) A 1024-task job on BG/L showed that the communication averaged to about 25 %

322   of the total solution time per inversion iteration. The distribution of the

323   communication overhead is as follows. Collective communications within GC are

324   mainly global reduction operations, and amount to about 50% with typical message

325   sizes of 16 Bytes. Point-to-point blocking message passing within LC: 20 % with

326   30 Kbytes average message size. Barrier synchronization: 30%.

327

328   The relatively long idle time due to global barrier synchronization, which is done after

329   each inversion iteration, indicates the importance of a balanced workload distribution

330   among all LCs. The QMR solver convergence behavior depends on the condition number

331   of the FD stiffness matrix $\mathbf{K}$ in equation (3), which in turn is governed by the aspect ratio

332   and conductivity contrasts within Ω_s. Because the latter changes dynamically with the

333   model updates during an inversion, a faster barrier synchronization would require an

334   adequate sophisticated scheme for dynamically adapting the LC size.

335   Over a 24-hour period, 72 inversion model updates were realized on BG/L and the relative

336   squared error misfit measure was reduced by nearly 67%. Exemplified in Figure 3, good

337 fits, to within the anticipated noise, were obtained for the horizontal and vertical inline

338 electric field overflight data, E_x^i (a) and E_z^i (b), as well for the horizontal perpendicular and

339 vertical broadside electric fields, E_y^b (c) and E_z^b (d). We observed that the major residual

340 misfits originate from the broadside inline components, E_x^b (e,f).

341 The average resistivity computed over three depth ranges for solution 72 is shown in

342 Figure 4. The sea bottom defines the depth $z=0$. Inspection of the images shows enhanced

343 resistivity in the southern model section for depths below 1500 m. Such is also observed

344 broadside of the sail lines, for the depth range 0-1500 m. Along the sail lines, however,

345 little to no resistivity enhancement is observed and the imaged resistivity volume contains

346 an unacceptable acquisition overprint. A possible explanation for this outcome is the

347 inconsistencies observed in fitting the in-line component of the broadside data compared

348 to other data components. This is particularly true of inline overflight data. Clearly, the

349 overflight data will be most sensitive to resistivity variations along the sail lines, while

350 broadside data are more sensitive to resistivity variations off the sail lines. One possibility

351 for the enhanced resistivity observed off the sail lines arises from the inversion algorithm's

352 attempt to fit the inline broadside data. Enhanced resistivity amplifies the broadside inline

353 model data, reducing the mismatch between observed and predicted data. Nevertheless, it

354 was still not possible to achieve acceptable data fits indicating a systematic bias in the

355 underlying assumptions employed in the inversion processing.

356 One critical assumption in this inversion was that the conductivity is isotropic;

357 conductivity within a cell does not vary with direction. However, it is well known within

358 sedimentary rocks that fine grain bedding planes can induce the rocks to exhibit transverse

359   electrical anisotropy [13 and 14]. In addition, parallel interbedding of rocks with different

360   conductivities can lead to anisotropic behavior. Thus, the conductivity can be expected to

361   depend strongly on directions, parallel and perpendicular to the bedding planes. In the

362   context of marine CSEM, [15] showed that the effects of electrical anisotropy can produce

363   significant anomalies, even as large as target reservoir responses, and a consensus is now

364   emerging that electrical anisotropy plays a bigger factor in influencing marine CSEM

365   measurement than previously believed.

366   Two tests were carried out to verify the importance of anisotropy. First, to test the degree

367   to which electrical anisotropy is affecting the broadside inline data, and to what lesser

368   extent it influences the overflight and broadside perpendicular and vertical data, we

369   repeated the initial stage of the inversion process. This involved an anisotropic model with

370   the vertical conductivity fixed at the conductivity used in the initial isotropic inversion and

371   the horizontal conductivity set to three times the vertical conductivity below the water

372   bottom. A sampling of the results is shown in Figure 5, confirming that the data are very

373   likely significantly more consistent with an anisotropic conductivity model than with an

374   isotropic one. Furthermore, we rerun two inversions with a subset of the data, comprising

375   36 effective transmitters. Using the same isotropic starting model, the inversions differed

376   by using an isotropic and anisotropic model parameterization. After 62 iterations, the

377   anisotropic model achieved a final data fit, which was by 27 % lower, compared to the

378   isotropic result. A complete anisotropic inversion of these data has yet to be carried out.

379

380

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.

406

407

REFERENCES

409   1. T. Eidesmo, S. Ellingsgrud, L. M. MacGregor, S. Constable, M. C. Sinha, S. Johansen,

410   F. N. Kong, and H. Westerdahl, "Sea Bed Logging (SBL), a new method for remote

411   and direct identification of hydrocarbon filled layers in deepwater," *First Break*, **20**, No. 3,

412   144-152 (March 2002).

413

414   2. S. Ellingsrud, T. Eidesmo, S. Johansen, M. C. Sinha, L. M. MacGregor, and S.

415   Constable, "Remote sensing of hydrocarbon layers by seabed logging (SBL): Results from

416   a cruise offshore Angola," *The Leading Edge*, **21**, No. 10, 972-982 (October 2002).

417

418   3. L. M. MacGregor and M. C. Sinha, "Use of marine controlled source electromagnetic

419   sounding for sub-basalt exploration," *Geophysical Prospecting*, **48**, No. 6, 1091-1106

420   (November 2000).

421   4. L. J. Srnka, J. J. Carazzone, M. S. Ephron, and E. A. Eriksen, "Remote reservoir

422   resistivity mapping," *The Leading Edge*, **25**, No. 8, 972-975 (August 2006).

423

424   5. S. Constable, "Marine electromagnetic methods – A new tool for offshore exploration,"

425   *The Leading Edge*, **25**, No. 4, 438-444 (April 2006).

426

427   6. D. L. Alumbaugh, G. A. Newman, L. Prevost, and J. Shadid, "Three-dimensional,

428   wideband electromagnetic modeling on massively parallel computers," *Radio Science*. **31**,

429   No. 1, 1-23 (January-February 1996).

430

431   7. G. A. Newman and P. T. Boggs, "Solution accelerators for large-scale three-

432   dimensional electromagnetic inverse problems," *Inverse Problems*, **20**, doi:10.1088/0266-

433   5611/20/6/S10, S151-S171 (2004).

434   8. M. Commer and G. A. Newman, "Large scale 3D EM inversion using optimized

435   simulation grids non-conformal to the model space," *Exp. Abstr. Soc. Expl. Geophys.*, **25**,

436   760-764 (2006).

437   9. G. A. Newman and D. L. Alumbaugh, "Three-dimensional massively parallel

438   electromagnetic inversion - Part I. Theory," *Geophys. J. Int.*, **128**, No. 2, 345-354

439   (February 1997).

440

441   10. S. Moskow, V. Druskin, T. Habashy, P. Lee, and S. Davdychewa, "A finite difference

442   scheme for elliptic equations with rough coefficients using a Cartesian grid nonconforming

443   to interfaces," *SIAM J. Number. Anal.*, **36**, No. 2, 442-464 (February 1999).

444 11. J. J. Carazzone, O. M. Burtz, K. E. Green, and D. A. Pavlov, "Three dimensional

445 imaging of marine CSEM data," *Exp. Abstr. Soc. Expl. Geophys.*, **24**, 575-578 (2005).

446 12. K. E. Green, O. M. Burtz, L. A. Wahrmund, T. Clee, I. Gallegos, C. Xia, G. Zelewski,

447 A. A. Martinez, M. J. Stiver, C. M. Rodriguez, and J. Zhang, "R3M Case studies:

448 detecting reservoir resistivity in complex settings," *Exp. Abstr. Soc. Expl. Geophys.*, **24**,

449 572-574 (2005).

450 13. J. D. Klein, P. R. Martin, and D. F. Allen, "The petrophysics of electrically anisotropic

451 reservoirs," *Trans. Soc. Petrol. Well-Log Analysts (SPWLA); 36th Ann. Logging Symp.*,

452 (1995).

453

454 14. J. Zhao, D. Zhou, X. Li, R. Chen, and C. Yang, "Laboratory measurements and

455 applications of anisotropy parameters of rocks," *Trans. Soc. Petrol. Well-Log Analysts*

456 *(SPWLA); 35th Ann. Logging Symp.*, (1994).

457 15. G. M. Hoversten, G. A. Newman, N. Geier, and G. Flanagan, G., "3D modeling of a

458 deepwater EM exploration survey," *Geophysics*, **71**, No. 5, G239-G248 (September-

459 October 2006) .

460

461

462

463

464   **Figure captions**

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

466   dimensions.

467   Figure 2. Layout of the sail lines (red and blue) and 23 detector locations (crosses) on the

468   sea bottom for the offshore CSEM survey. Contained survey configurations are illustrated

469   in the upper figure. Bathymetry is given in meters below sea level. The example data

470   shown in this paper corresponds to the transmitter-detector arrays marked in blue.

471   Figure 3. Six selected plots of overflight and broadside electric field data amplitudes

472   (black curves) versus the transmitter offset projected onto the profile lines shown in Figure

473   2. Shown are data fits produced by the starting model (red) and for iteration 72 (blue).

474   Figure 4. Average resistivity computed over three depth ranges for solution 72: a) Water

475   bottom to 500 m below mud line (BML), b) interval 500 to 1500 m BML, c) interval 1500

476   to 2500 m BML. Resistivity is rendered on a base 10 log scale.

477   Figure 5. Six selected plots of overflight and broadside electric field data amplitudes

478   (black curves) versus the transmitter offset projected onto the profile. Shown are data fits

479   produced by a starting model with isotropic (red) and anisotropic (blue) electrical
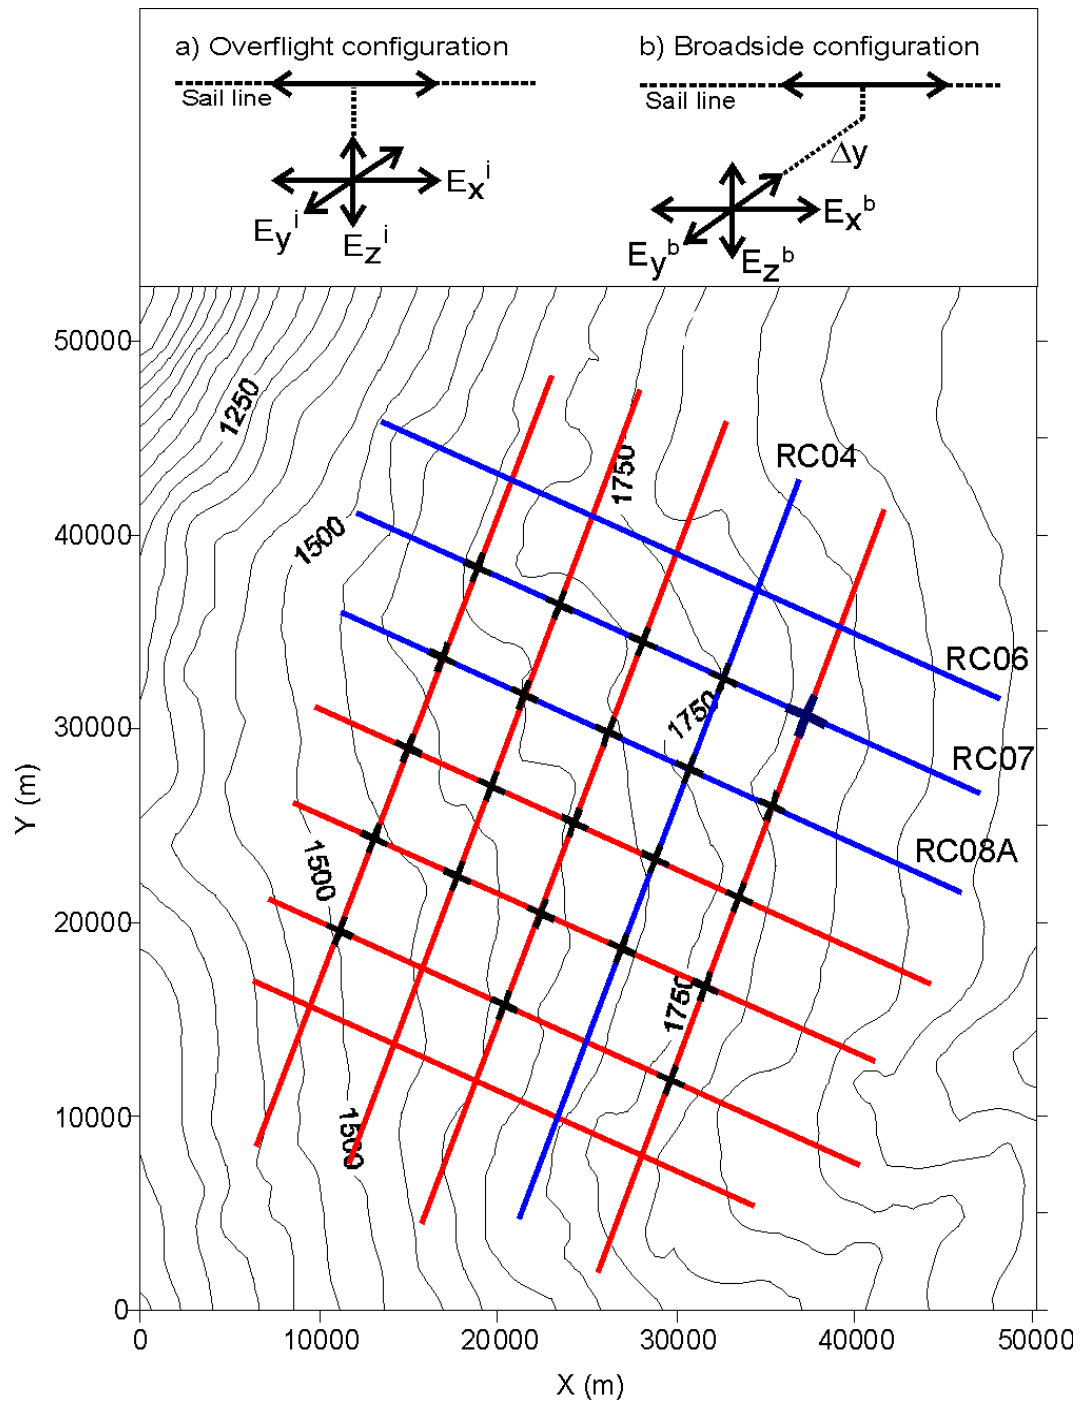
480   conductivity.

481

482

483
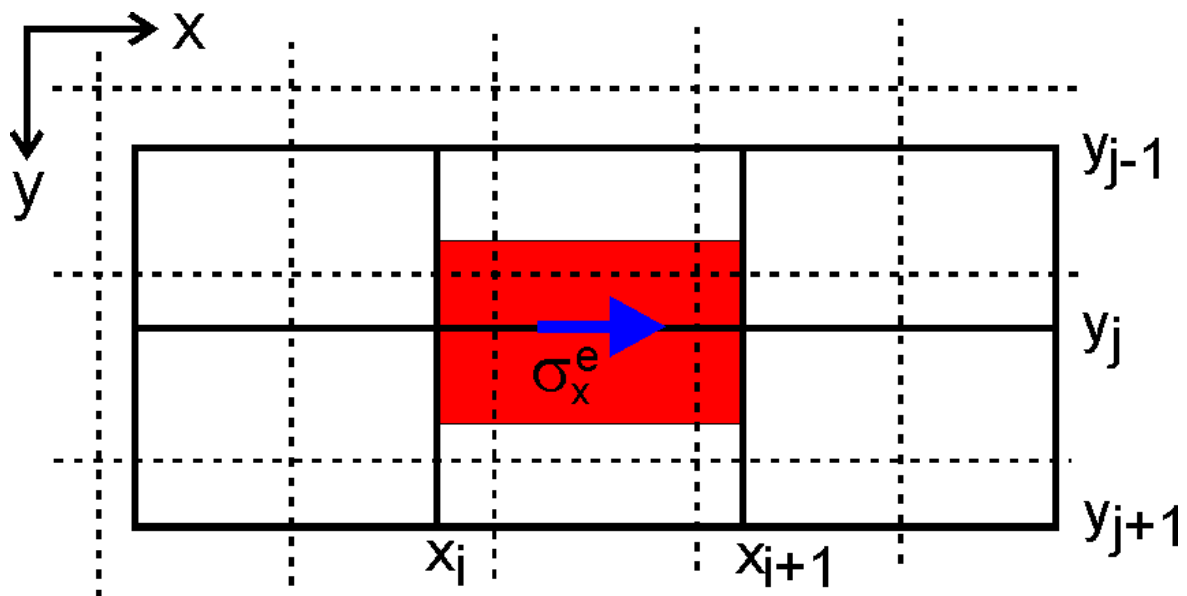
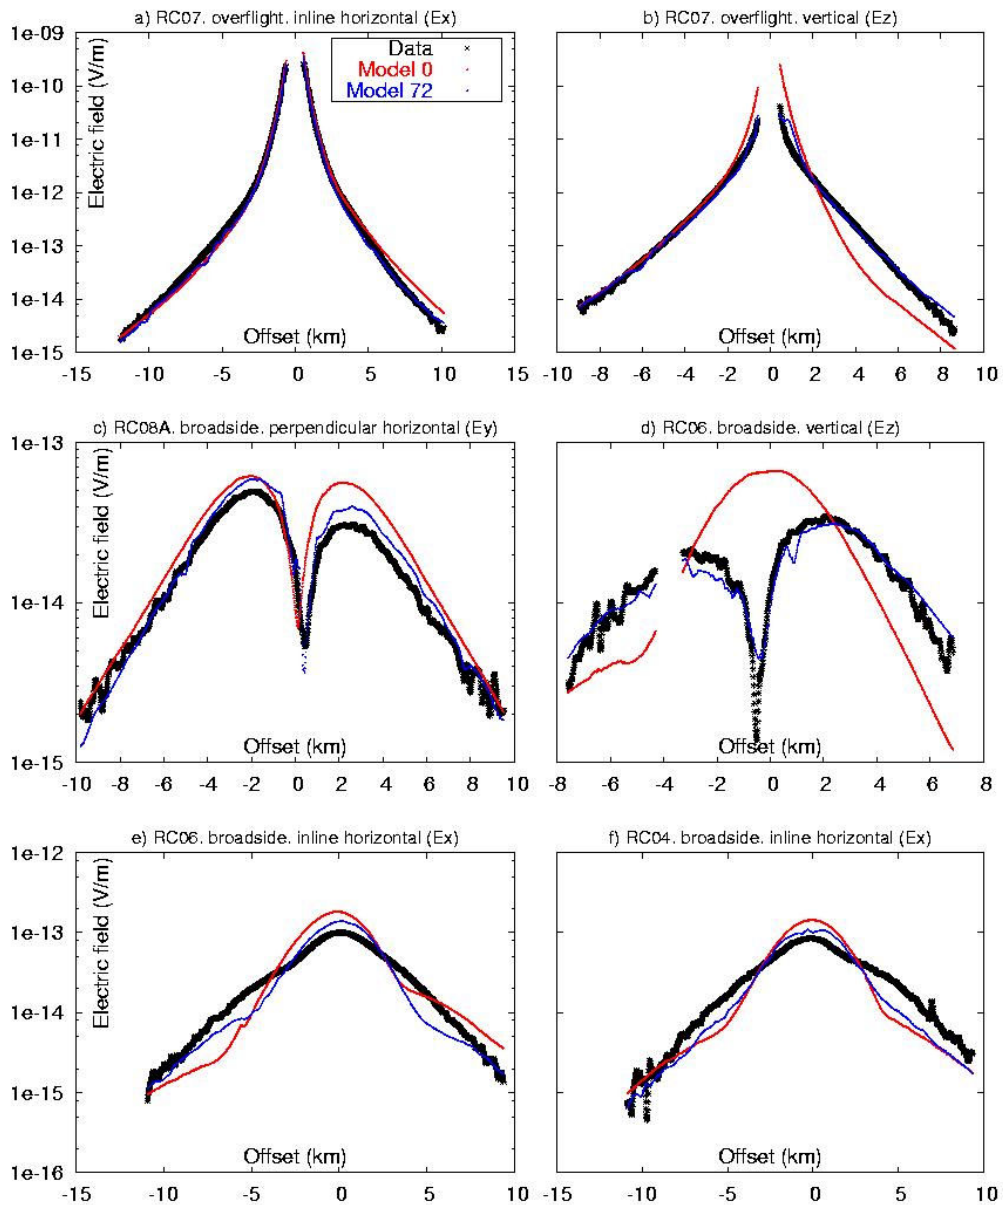

484

485   Fig. 1

486

487   Fig. 2

488

489 Fig. 3

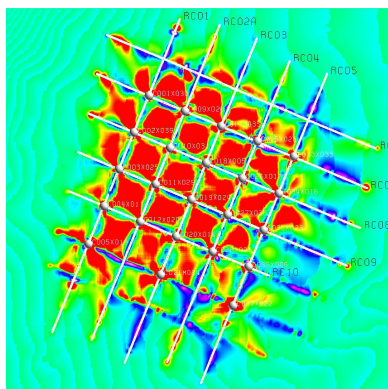490

491

492

493   a)

494

495   b)
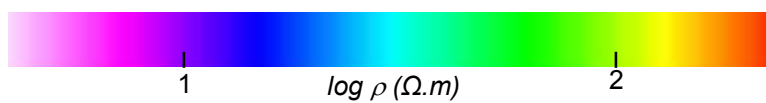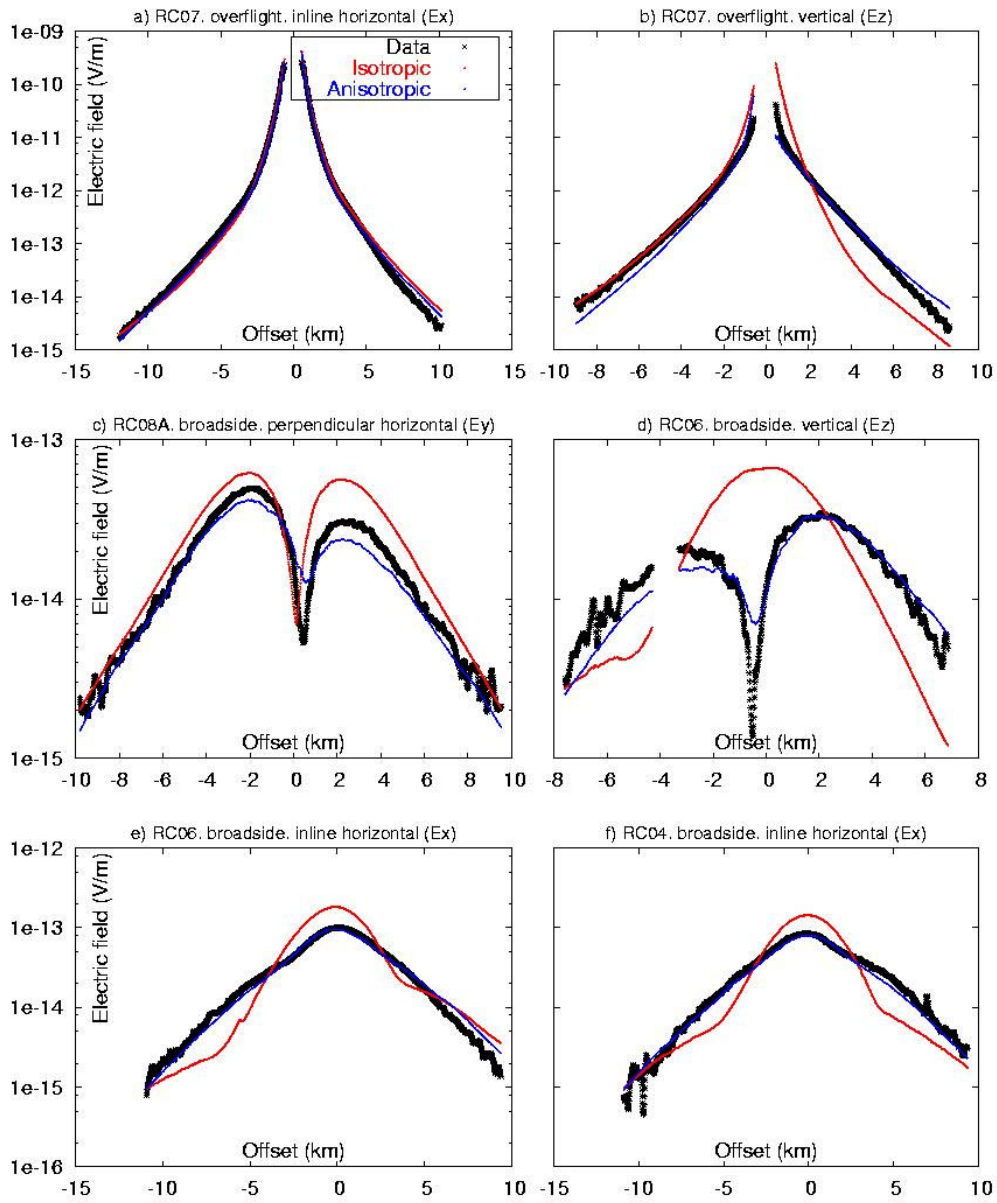

496

497   c)

498

$log\ \rho\ (\Omega.m)$

499   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)
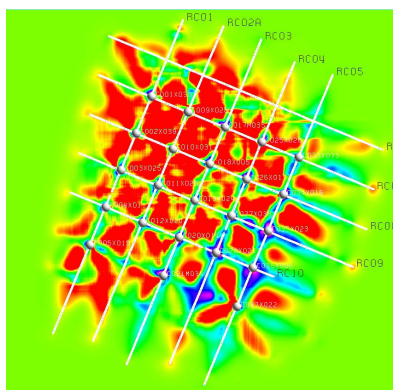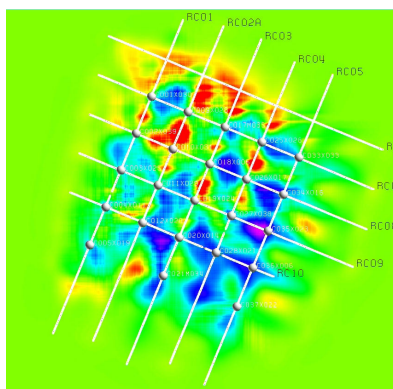
500

501   Fig. 5

502

503

504

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.

548

549

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

570 *ExxonMobil Upstream Research Company, Houston, TX 77252*

571 (denny.e.willen@exxonmobil.com).   Dr. Willen joined Exxon Production Research

572 Company in 1980, after receiving his Ph.D. in Physics from the University of Illinois.

573 Since then he has worked in several areas of exploration geophysics including seismic

574 imaging, near-surface effects, converted waves, well logging, parallel computing, and

575 electromagnetic methods.   He is a member of the Society of Exploration Geophysicists

576 and the International Trumpet Guild.

577

578 **Janet Shiu**

579 *IBM Deep Computing, Two Riverway, Houston, Texas 77056*

580 (jshiu@us.ibm.com). Dr. Shiu received her Ph.D. degree in Physics from University of

581 Pittsburgh. Dr. Shiu was an assistant professor at Old Dominion University, and a

582 principal investigator of research grants from NASA Langley Research Center where her

583 research focus were in the area of high power lasers for space application. Since 1985, Dr.

584 Shiu has been employed with Cray Research Inc., SGI, and Exxon Upstream Technical

585 Company in technical support of High Performance Computing. Dr. Shiu joined IBM in

586 1999, and is a member of IBM Deep Computing Technical Team.

587