3D Geophysical Imaging of the Subsurface on Multi-Core Machines

Gregory A. Newman
Michael Commer & Filipe Maia
Earth Sciences Division
National Energy Supercomputer Center [NESEA]
Lawrence Berkeley National Laboratory

January 26, 2011

PRESENTATION OVERVIEW

- Motivation for Imaging on Multi-Core Machines
  - Controlled Source EM and Magnetotelluric Data Acquisition
- Formulation of the Imaging Problem
  - Large Scale Modeling Considerations
- Case Studies
  - Offshore Brazil
  - Gulf of Mexico (synthetic example)
  - Coso Geothermal Field, Eastern California
- Seismic Imaging
  - 10 to 100X Larger Computational Demands !!!
- Computing Alternatives
  - GPU
  - FPGA
- Conclusions

Marine CSEM & MT Surveying

CSEM

Deep-towed Electric Dipole transmitter

- ~ 100 Amps
- Water Depth 1 to 7 km
- Alternating current 0.01 to 3 Hz
- 'Flies' 50 m above the sea floor
- Profiles 10's of km in length
- Excites vertical & horizontal currents
- Depth of interrogation ~ 3 to 4 km
- Sensitive to thin resistive beds

MT

Natural Source Fields

- Less than 0.1 Hz
- Measured with CSEM detectors
- Sensitive to horizontal currents
- Depth of interrogation 10's km
- Resolution is frequency dependent
- Sensitive to larger scale geology

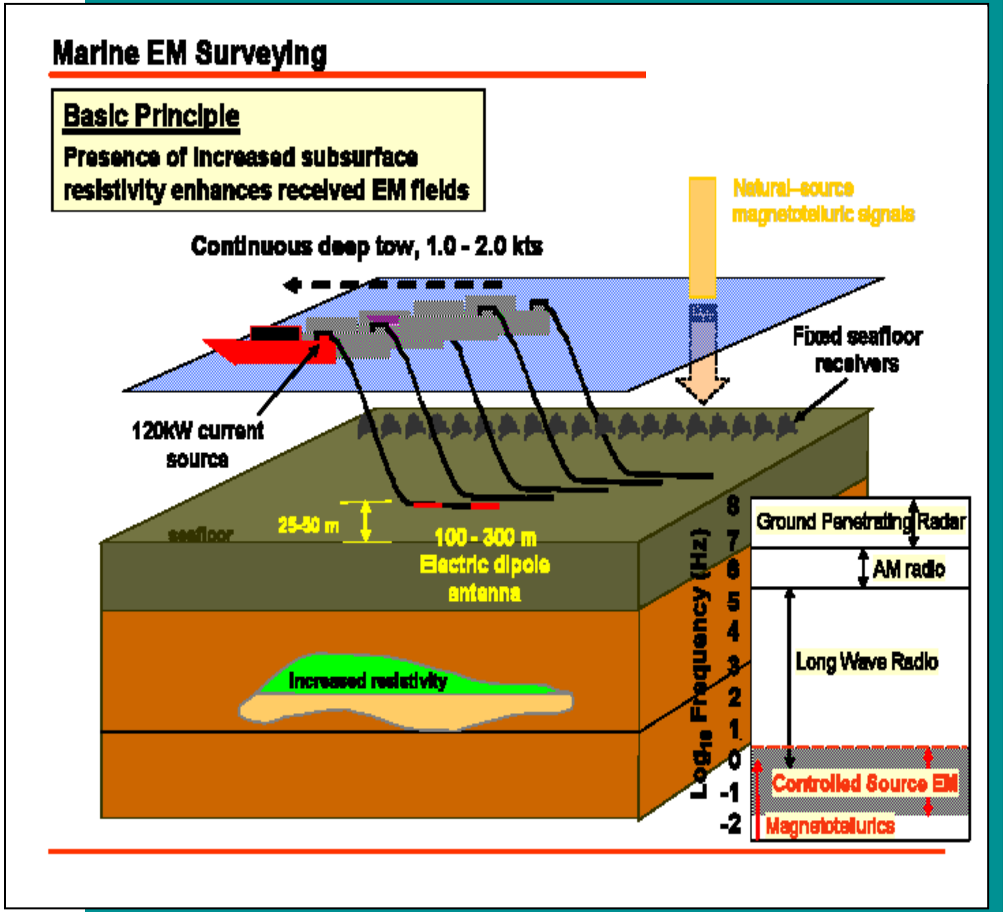

3D INVERSE MODELING

Minimize:

$$\varphi = \alpha \sum_{j=1}^{N} \{(d_j^{obs} - d_j^{p})/\varepsilon_j\}^2 + \beta \sum_{j=1}^{M} \{(Z_j^{obs} - Z_j^{p})/\pi_j\}^2$$

$$+ \lambda_h \mathbf{m}_h \mathbf{W}^T \mathbf{W} \mathbf{m}_h + \lambda_v \mathbf{m}_v \mathbf{W}^T \mathbf{W} \mathbf{m}_v$$

$$\text{s.t.} \quad \mathbf{m}_v \leq \mathbf{m}_h$$

$\mathbf{d}^{obs}$ and $\mathbf{d}^{p}$ are N observed and predicted CSEM data

$\mathbf{Z}^{obs}$ and $\mathbf{Z}^{p}$ are M observed and predicted MT impedance data

ε & π = CSEM and MT data weights

$\mathbf{m}_h$ = horizontal conductivity parameters

$\mathbf{m}_v$ = vertical conductivity parameters

$\mathbf{W} = \nabla^2$ operator; constructs a smooth model

λ_h & λ_v = horizontal & vertical tradeoff parameters

α & β = scaling factors for CSEM and MT data types

LARGE-SCALE 3D MODELING CONSIDERATIONS

- Require Large-Scale Modeling and Imaging Solutions
  - 10's of million's field unknowns (fwd problem)
    - » Solved with finite difference approximations & iterative solvers
  - Imaging grids 400 nodes on a side
    - » Exploit gradient optimization schemes, adjoint state methods

- Parallel Implementation
  - Two levels of parallelization
    - » Model Space (simulation and inversion mesh)
    - » Data Space (each transmitter/MT frequency - receiver set fwd calculation independent)
    - » Installed & tested on multiple distributed computing systems; 10 – 30,000 Processors

- Above procedure satisfactory except for very largest problems
  - To treat such problems requires a higher level of efficiency

- Optimal Grids
  - Separate inversion grid from the simulation/modeling grid
  - Effect: A huge increase in computational efficiency ~ can be orders of magnitude

BERKELEY LAB

Optimal Grids

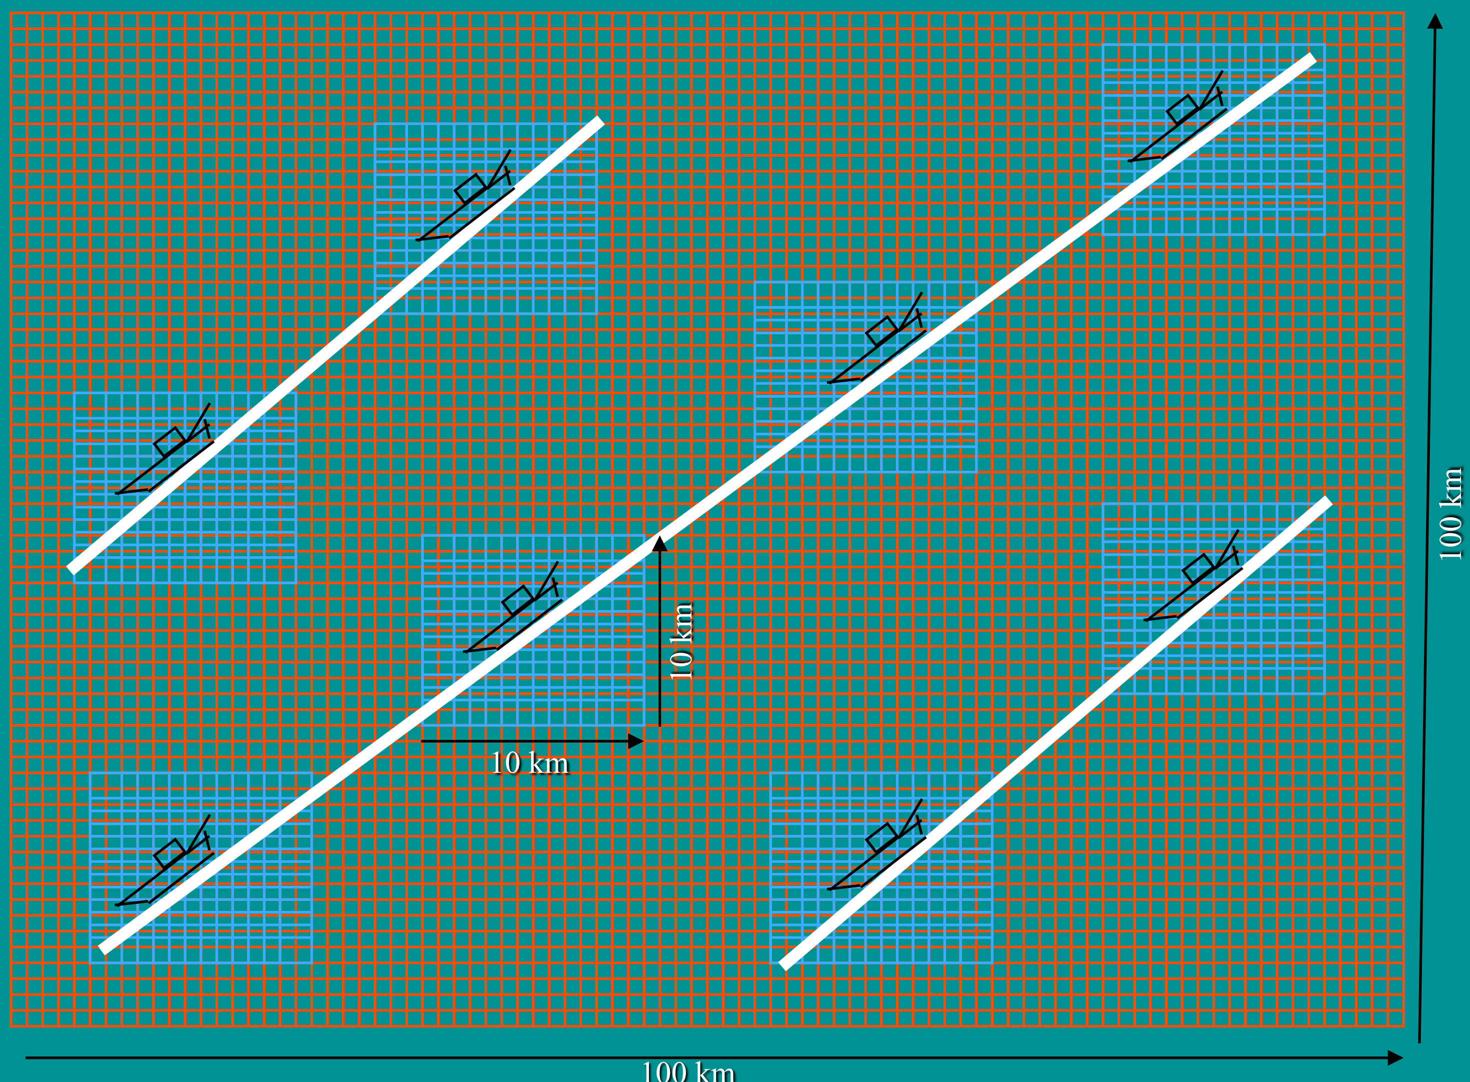

Ω_m imaging grid Ω_s simulation grid sail lines

Campos Basin CSEM Survey

Offshore Brazil

- Study: CSEM Imaging in the presence of electrical anisotropy
- Field Data: 23 detectors, 10 sail lines, 3 frequencies @ 1.25, 0.75, 1.25 Hz
- Image Processing: ~ 1 million data points, 27 million image cells
- Processing Times: 24 hours, 32,768 tasks, IBM Blue Gene (BG/L)
- Conclusions: data cannot be fit using isotropic model, anisotropic model required

Survey layout

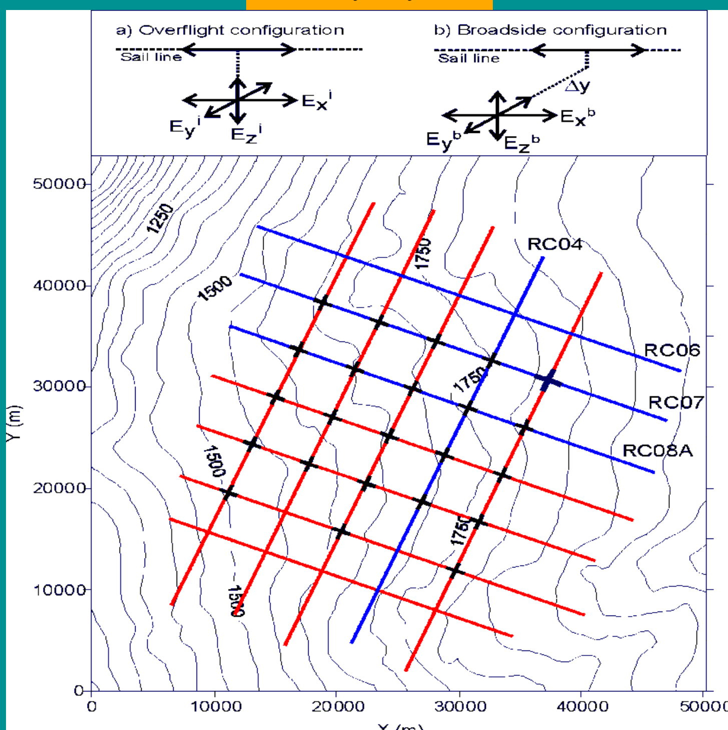

Overflight electric field

Broadside electric field

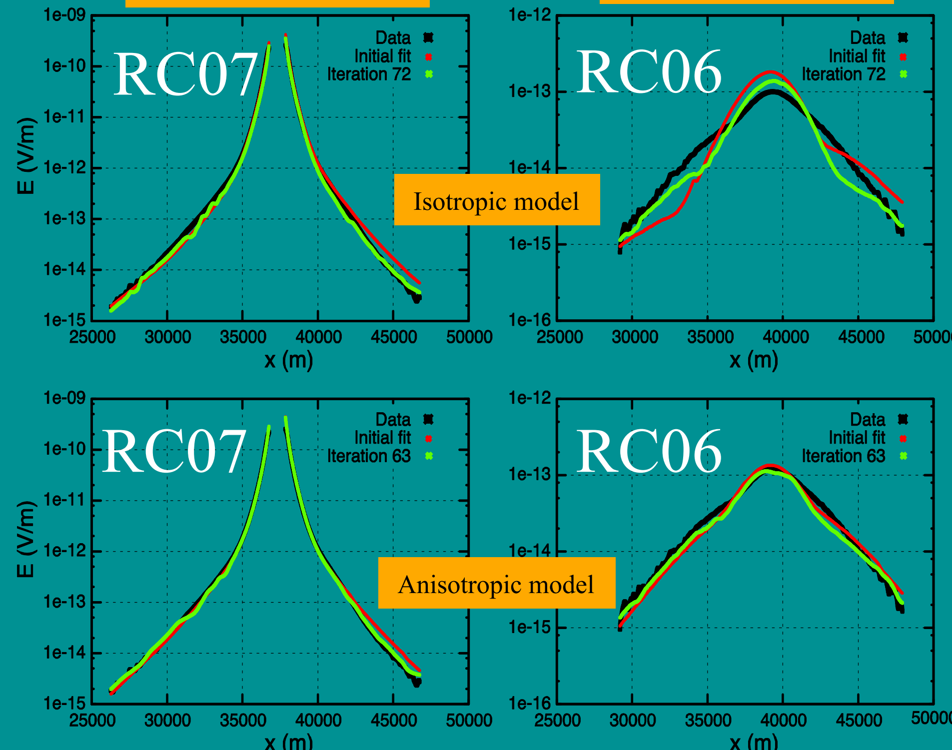

3D Vertical Resistivity Imaging

Offshore Brazil

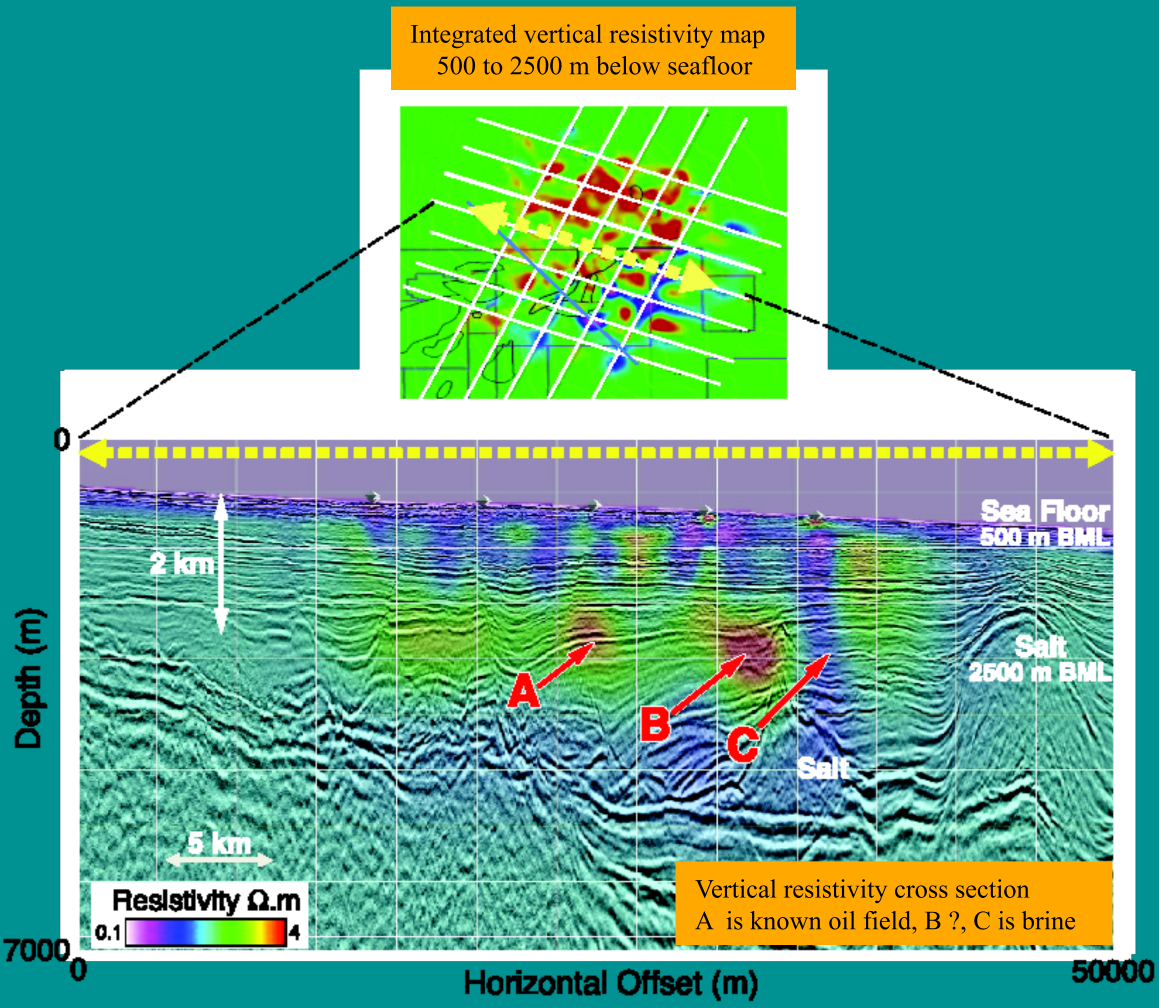

Joint CSEM - MT Imaging

Mahogany Prospect, Gulf of Mexico

- Study: 3D Imaging of oil bearing horizons with complex salt structures present
- Simulated Example: 100 m thick reservoir, 1 km depth, salt below reservoir
- Model: 0.01 S/m salt, 2 S/m seabed, 0.05 S/m reservoir, 3 S/m seawater
- MT Data: 7,436 data points, 143 stations & 13 frequencies 0.0005 to 0.125 Hz
- CSEM Data: 12,396 data points, 126 stations & 2 frequencies 0.25 and 0.75 Hz
- Starting Model: Background Model without reservoir or salt
- Processing Times: 5 to 9 hours, 7,785 tasks, NERSC Franklin Cray XT4 System

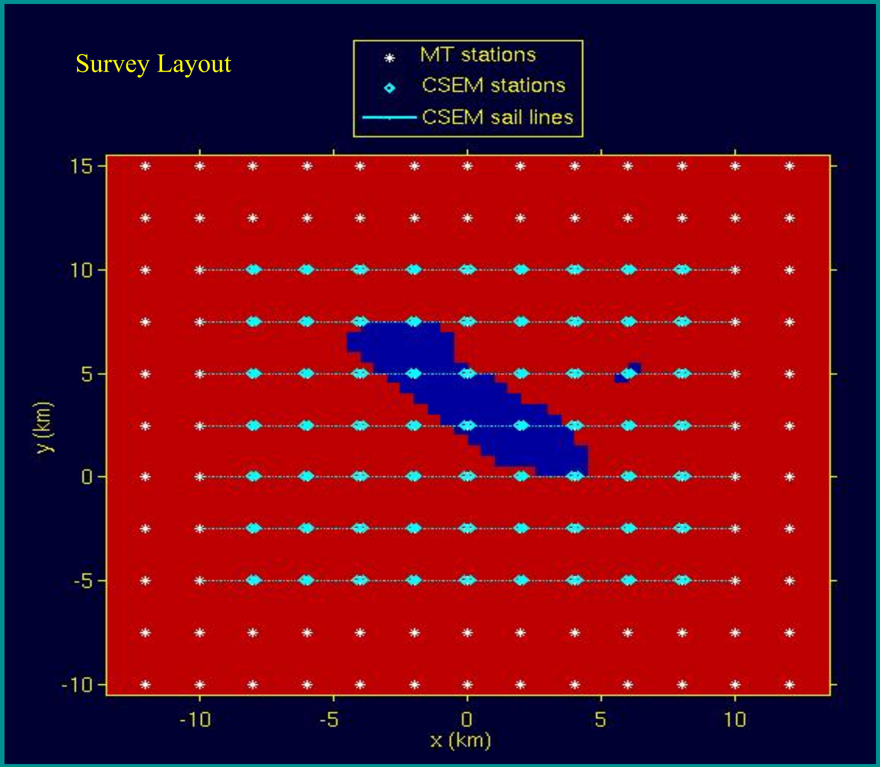

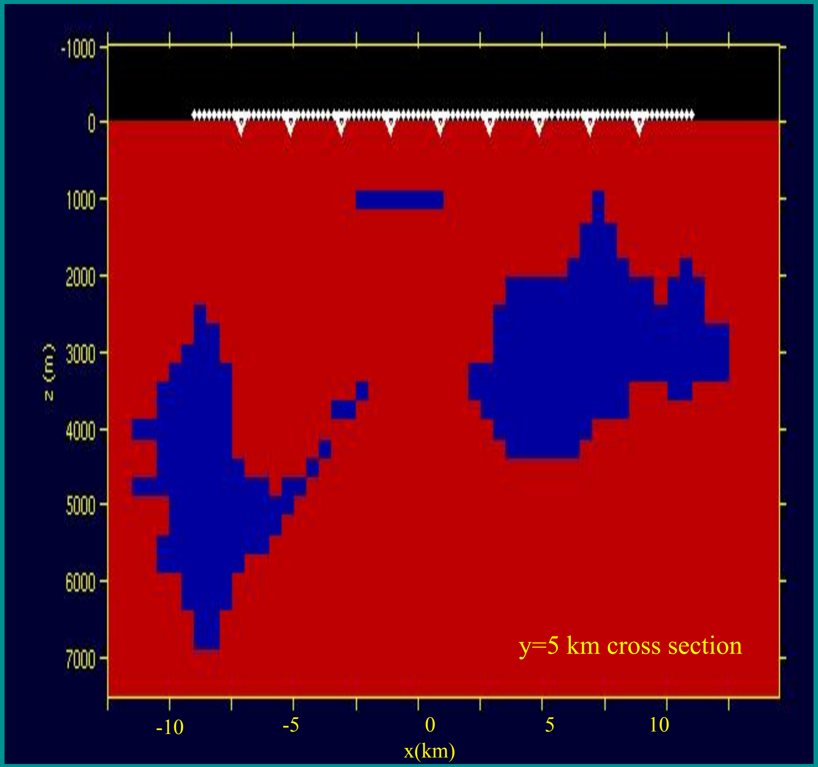

JOINT CSEM-MT IMAGING:

The Benefits

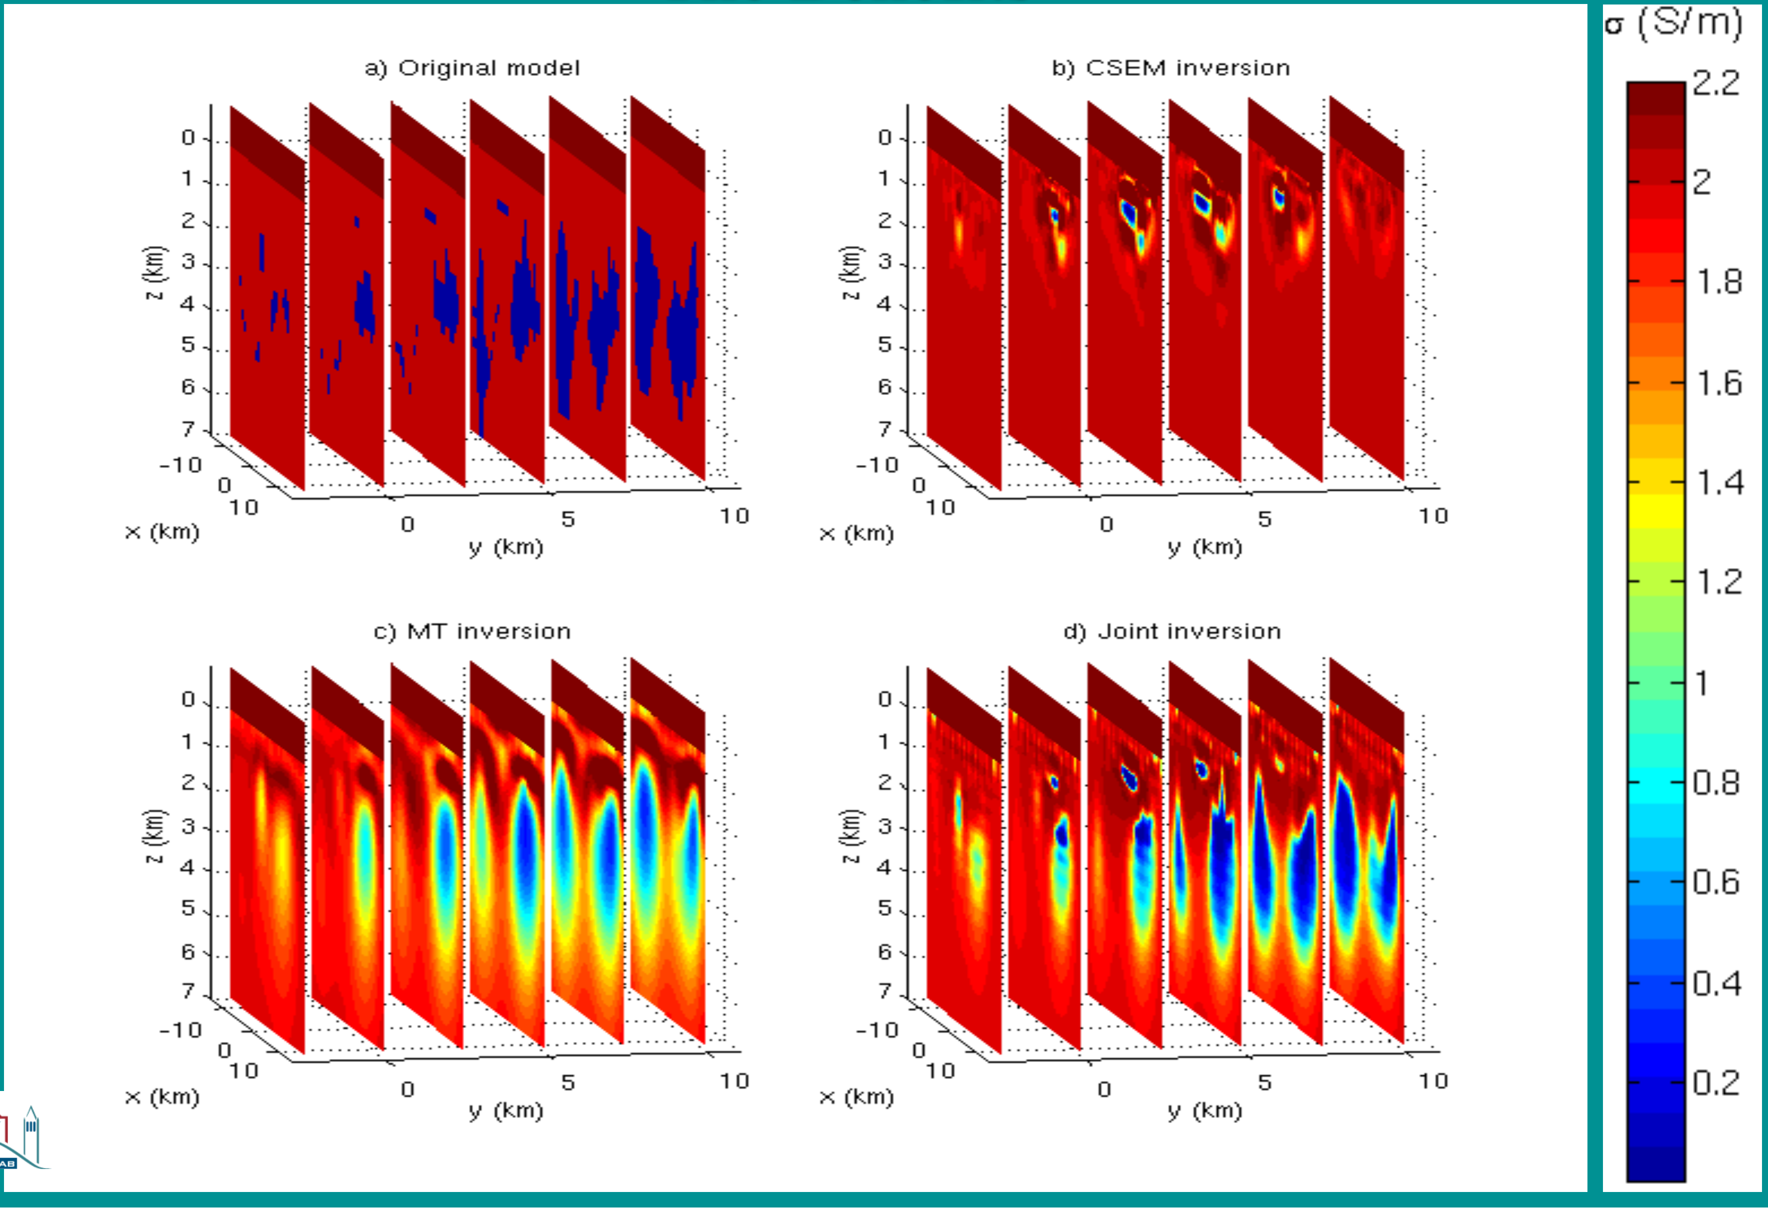

Joint CSEM - MT Imaging
Mahogany Prospect Gulf of Mexico

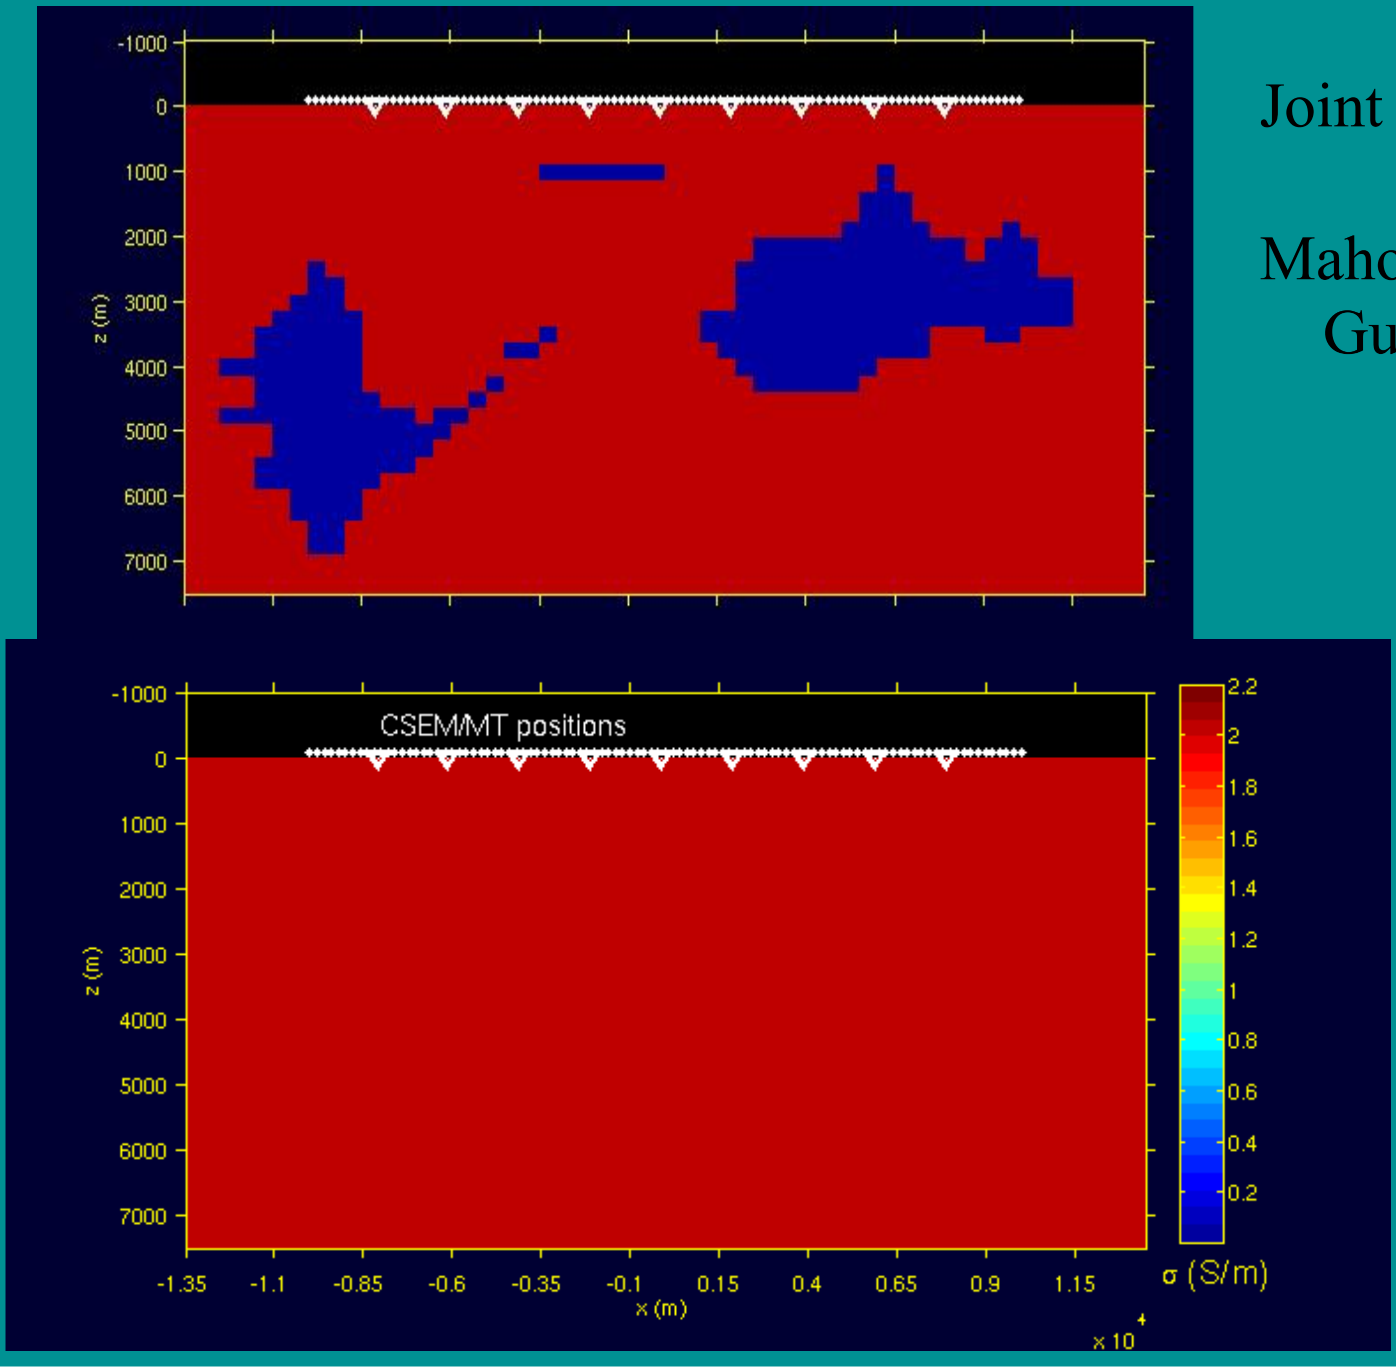

MT Imaging for Geothermal Resources

The Coso Field

- ➢The reservoir is located in Eastern California, Southern End Owens Valley
- ➢Over 120 MT soundings acquired over the eastern flank of the field
- ➢High density profile along the line NA1
- ➢Remote referencing used to suppress noise from western US power grid
- ➢The data span a frequency range from 100 to 0.001 Hz.
- ➢ Run on 512 Cores : NERSC Seaborg Machine IBM SP2 Processors

Survey Layout

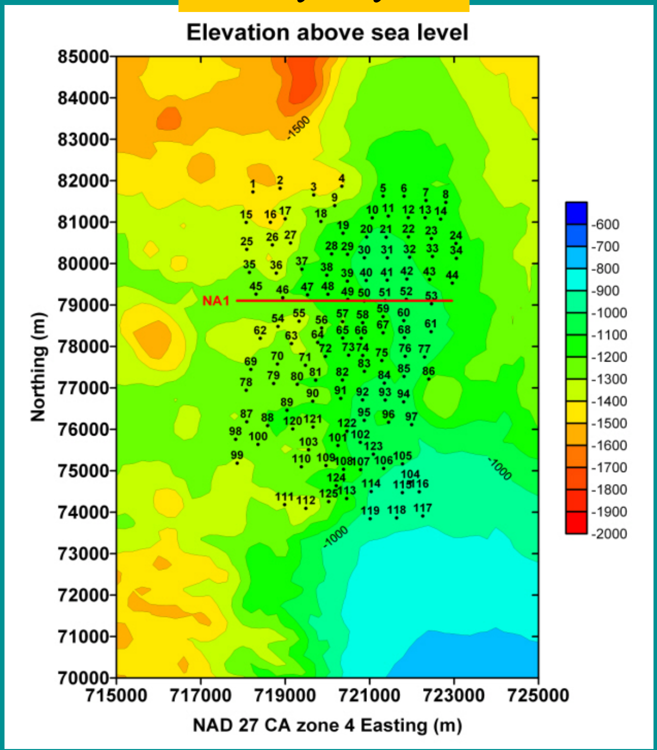

3D Conductivity Image

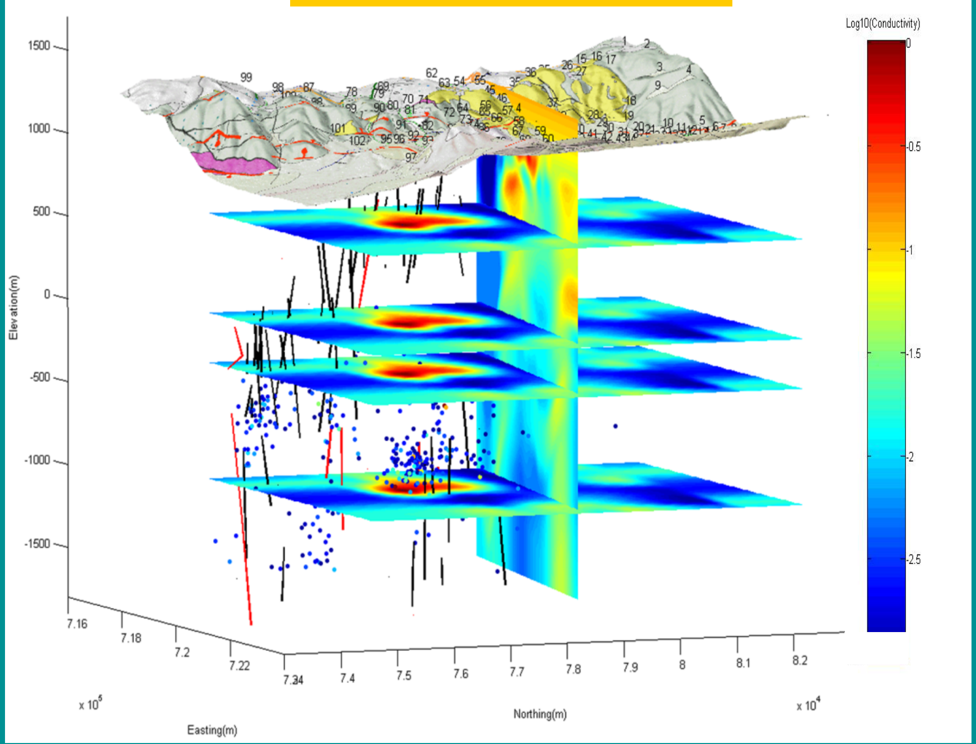

Seismic Imaging: Laplace-Fourier Domain

10 to 100X Computational Demands

337 shot gathers
151 detectors/shot
maximum offsets 15km

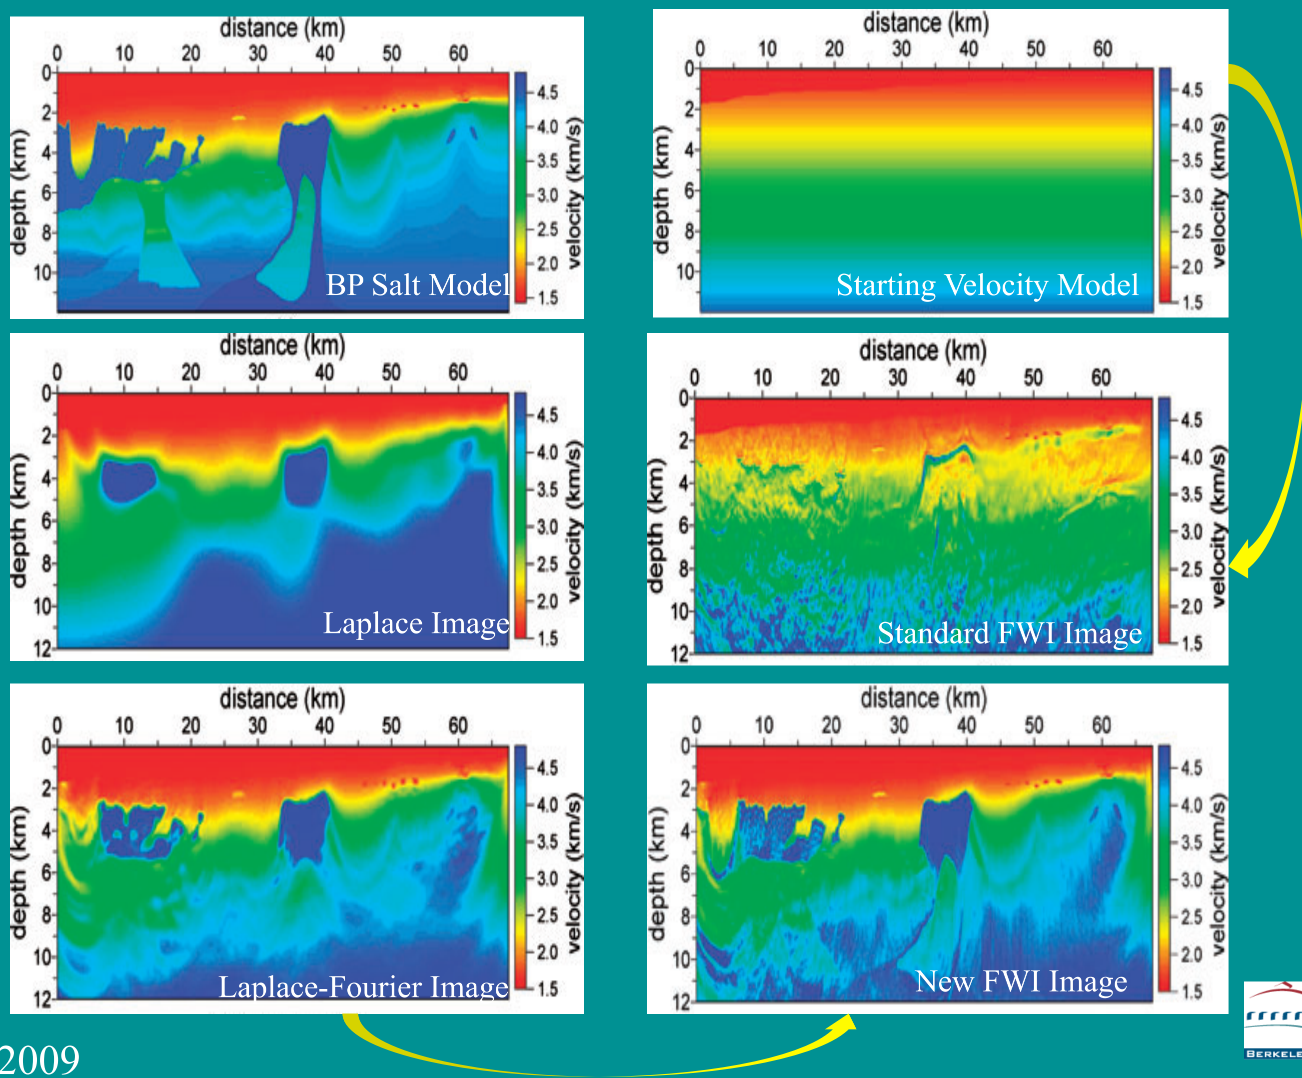

s = 10.5 to 0.5
Δ=0.5

s = 10.5 to 0.5
f = 6 to 0.5
Δ=0.5

Taken from Shin & Cha, 2009

Computing Alternatives

- Multi-Core Geophysical Imaging
  - Multiple Imaging Experiments
    - » Necessary to reduce model uncertainty
    - » Never exhaust the modeling possibilities & scenarios
  - Costly !!!
    - » Millions of Dollars Computing Expenses Incurred Yearly
- Cheaper and Faster Way to Compute ??
  - New Hardware & Computing Architectures
    - » GPU's
    - » FPGA's
  - Painful Process to Migrate to New Platforms

Geophysical Inverse Modeling GPU Platforms

- Main computational bottleneck:

  Sparse Matrix-Vector Multiplication (SpMV) in iterative Krylov solvers
- Krylov solvers used for solving the forward modeling problem
- Non-contiguous memory access limits performance of SpMV
- Proper memory alignment is key to achieving high SpMV performance

GPU Iterative Krylov Solvers Implemented Thus Far

- QMR (Quasi-minimum-residual) and BiCG (BiConjugate Gradient) methods for complex-symmetric matrices:
- Needed for methods:
  - CSEM: Exploration and environmental studies
  - MT: Crustal studies, geothermal
  - SIP: Environmental studies
  - Seismic problem in Laplace-Fourier domain

Environmental Imaging Problems

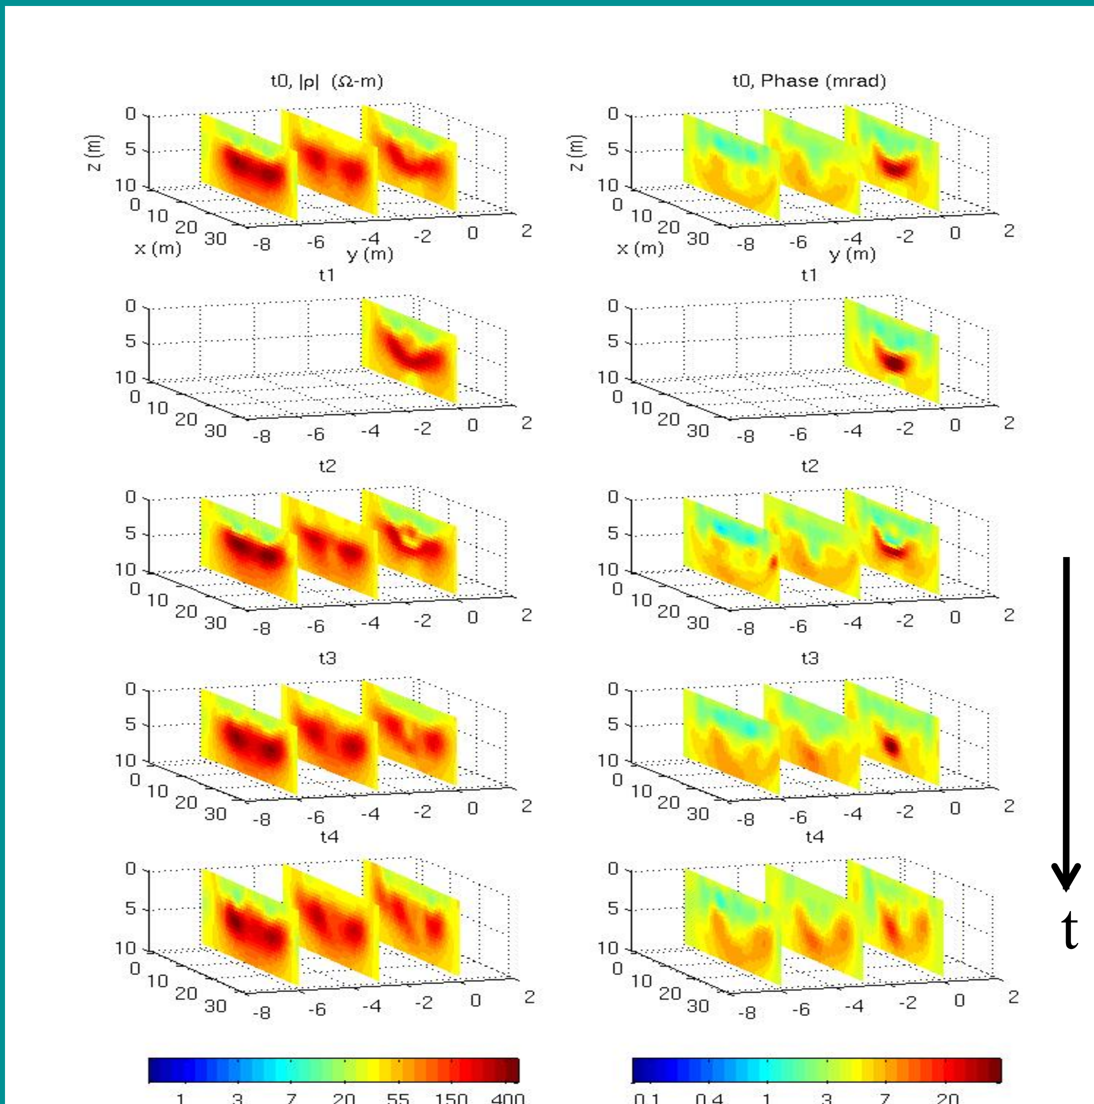

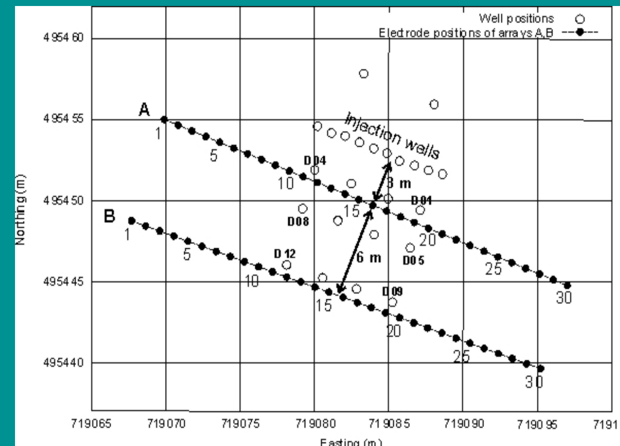

SIP method (spectral induced polarization) provides indirect information about hydrological subsurface properties

CSEM Imaging

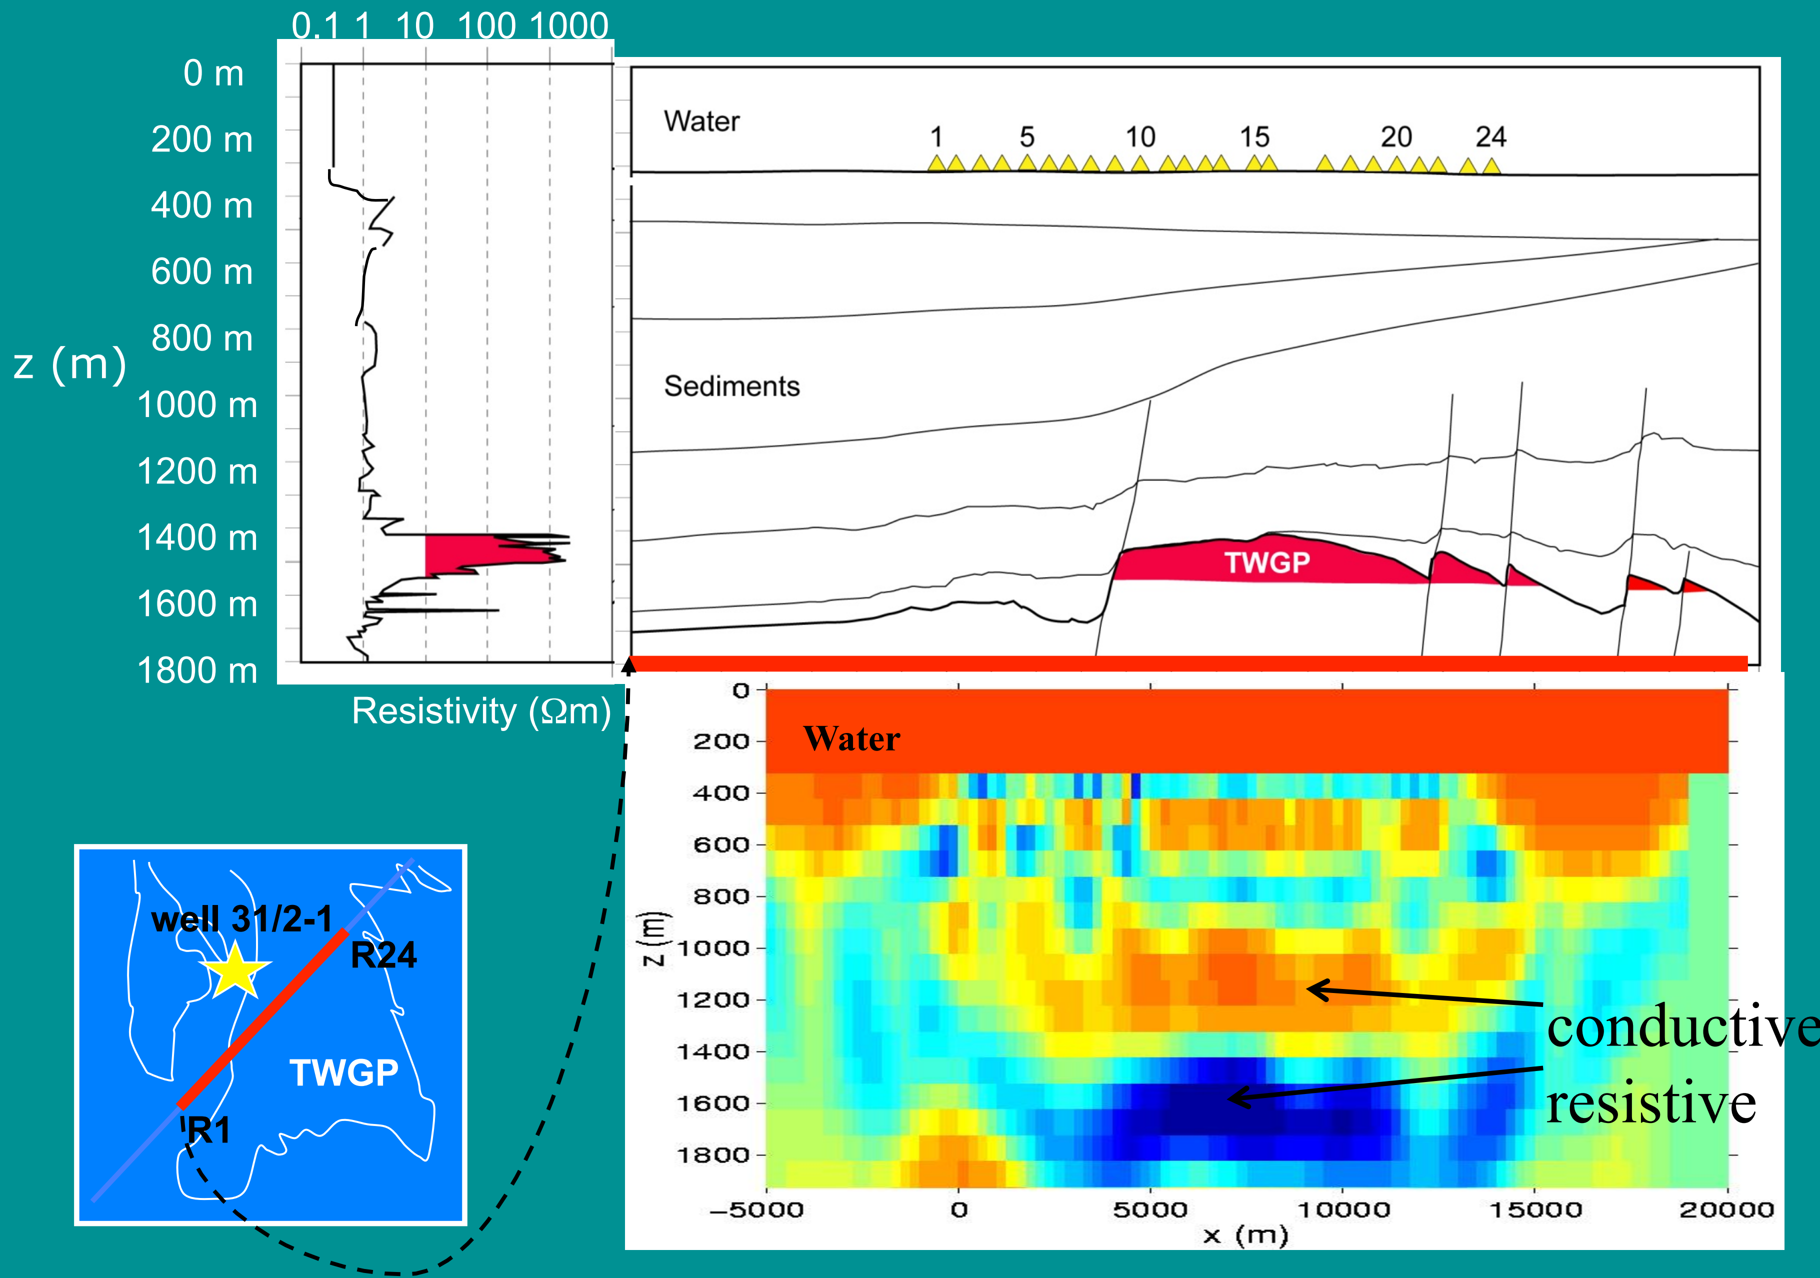

QMR Solver Performance

- GPU computing speed-up achieved on Dirac (NVIDIA Tesla C2050) compared to CPU-performance (Intel Nehalem 2.4 GHz, 8MB cache, Quad core with 8 cores per node):

  1 - 40

Other Krylov Solvers Implemented:

- CG-solvers for modeling problems involving real arithmetic: For example needed for electrical resistivity tomography
- 3D-Poisson problem on 158 x 110 x 165 grid
  - CG-solver 1 (our own implementation): 9 sec
  - CUSP-BiCG solver: 21 sec
  - CUSP-CG solver: 15 sec
  - AZTEC CPU solution: 42 sec

Challenges

- Limited memory on GPU. In CPU-world we can just increase number of CPUs for solving big problems
- In GPU-world: “Parallelize” GPUs ?
- Preconditioning methods needed for some geophysical modeling problems
- Efficient co-processing: Available CPUs should not stay idle

Conclusions

- 3D Imaging on Multi-Core Machines
  - demonstrated need for energy exploration
  - 1000's to 10000's compute cores
  - expensive

- Computing Alternatives Being Investigated
  - GPU's & FPGA's
  - 40 CPU's ≈ 1 GPU
  - many technical issues to be resolved for large scale imaging

BERKELEY LAB

ACKNOWLEDGEMENTS

Funding & Field Data Sets Were Provided By These Institutions:

US Department of Energy
Office of Science
Geothermal Program Office
ExxonMobil Corporation
Chevron Corporation

Joint Imaging of EM and Seismic Data

- Issues
  - Rock Physics Model
    - » links attributes to underlying hydrological parameters
    - » too simplistic
    - » difficult or impossible to define robust/realistic model
  - Differing Resolution in the Data
    - » EM data 10x lower resolution compared to seismic
  - RTM & FWI of Seismic Data
    - » requires very good starting velocity model
    - » velocity can be difficult or impossible to define
    - » huge modeling cost due to very large data volumes

      (10,000's of shots; 100,000's traces per shot)

Berkeley Lab

Joint Imaging of EM and Seismic Data

- A way forward
  - Abandon Rock Physics Model
    - » assume conductivity and velocity structurally correlated
    - » employ cross gradients: $t = \nabla\sigma \times \nabla\upsilon$
    - » $t = 0 \Rightarrow \nabla\sigma \,||\, \nabla\upsilon$; $\nabla\sigma = 0$ and/or $\nabla\upsilon = 0$
  - Equalize Resolution in the Data
    - » treating seismic and EM data on equal terms
    - » Laplace-Fourier transform seismic data – Shin & Cha 2009

$$\hat{g}(s) = \int_0^\infty g(t) e^{-st} dt$$

$g(\hat{s})$ *and* s *are complex*

Acoustic Wave Equation

Propagating Wave

Time Domain

$$\left[\frac{1}{v^2}\frac{\partial^2}{\partial t^2}+\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)\right]p(x,y,z,t)=-s(t).$$

Fourier/Frequency Domain

$$\left[-\frac{\omega^2}{v^2}+\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)\right]p(x,y,z,\omega)=-s(\omega).$$

Damped Diffusive Wave

Laplace/Fourier Domain

$$\left[\frac{s^2}{v^2}+\left(\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2}+\frac{\partial^2}{\partial z^2}\right)\right]\hat{p}(x,y,z,s)=-\hat{s}(s).$$

similar to physics & similar resolution with EM fields

skin depth: $\delta = \upsilon / s_r$

JOINT IMAGING FORMULATION

$$\varphi = \alpha \sum_{j=1}^{N} \left\{ \left(d_{em_j}^{obs} - d_{em_j}^{p} \right) \varepsilon_j \right\}^2 + \beta \sum_{l=1}^{M} \left\{ \left(\hat{d}_{s_l}^{obs} - \hat{d}_{l_l}^{p} \right) \chi_j \right\}^2$$

$$+ \lambda_{em} \boldsymbol{\sigma}^{\mathrm{T}} \mathbf{W}^{\mathrm{T}} \mathbf{W} \boldsymbol{\sigma} + \lambda_s \boldsymbol{\nu}^{\mathrm{T}} \mathbf{W}^{\mathrm{T}} \mathbf{W} \boldsymbol{\nu} + \tau \sum_{i=1}^{m_c} \boldsymbol{t}_i \cdot \boldsymbol{t}_i$$

d_{em}^{obs} and d_{em}^{p} are N observed and predicted EM data

$\hat{d}_s^{obs}$ and $\hat{d}_s^{obs}$ are M observed and predicted Laplace-Fourier seismic data

ε and χ = EM and seismic data weights

$\boldsymbol{\sigma}$ = m conductivity parameters

$\boldsymbol{\nu}$ = m acoustic velocity parameters

$\mathbf{W} = \nabla^2$ operator; constructs a smooth model

λ_{em} and λ_s = conductivity & velocity tradeoff parameters

α and β = scaling factors for EM and seismic data types

$\boldsymbol{t}$ are m_c cross gradient structural constraints; τ is a penalty parameter

Initial Results

marine example

f = 0.25

s =5,4,3,2,1
f = 0

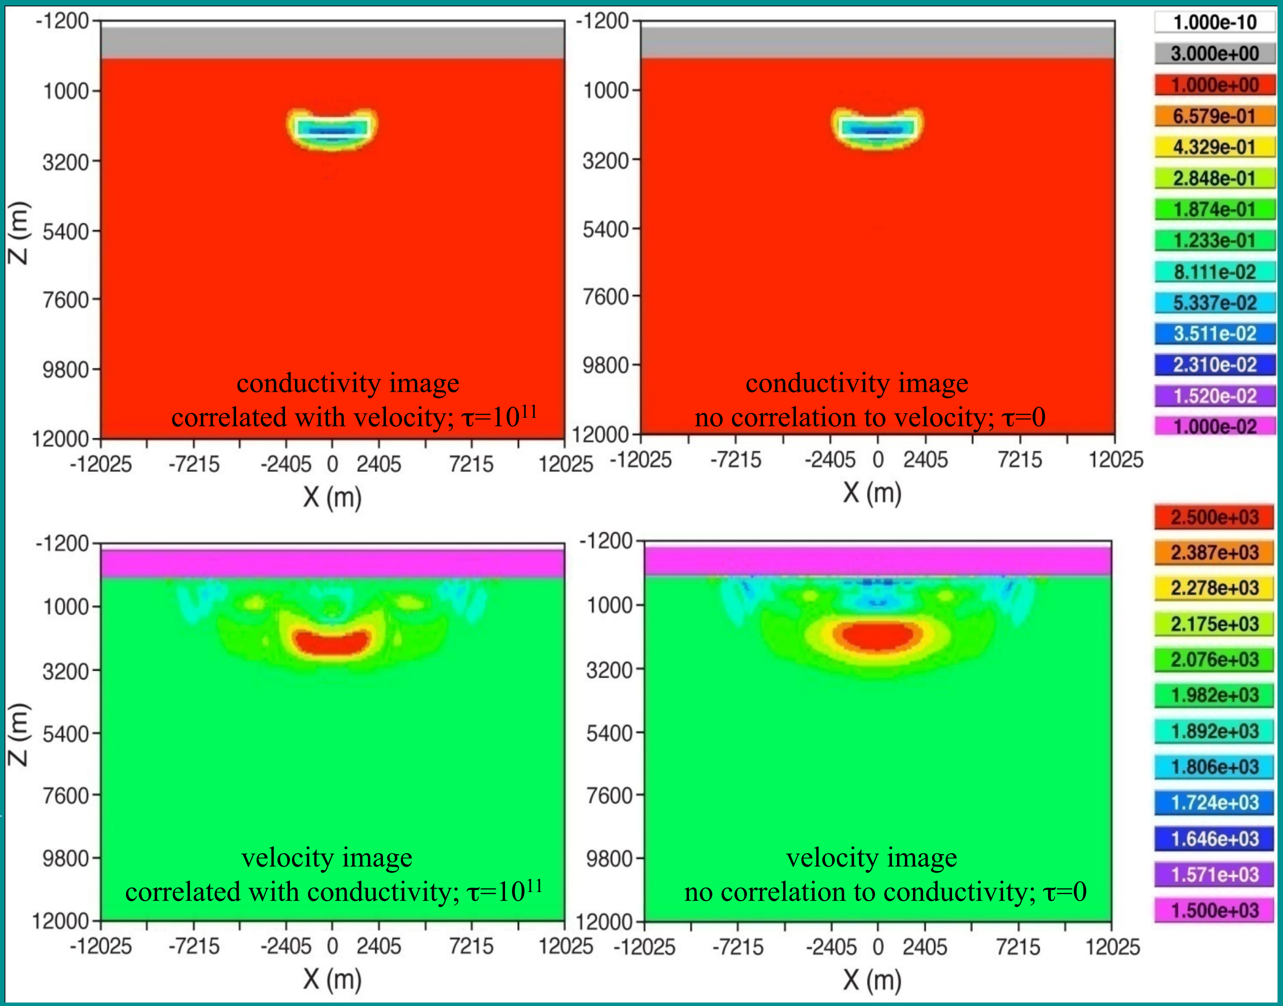

CSEM
16 km max. offsets
17 shots
161 detectors/shot

seismic
12-16 km offsets
85 shots
121-161
detectors/shot

Correlations with MEQ Data

Vp/Vs Ratio Map
700 m below the surface

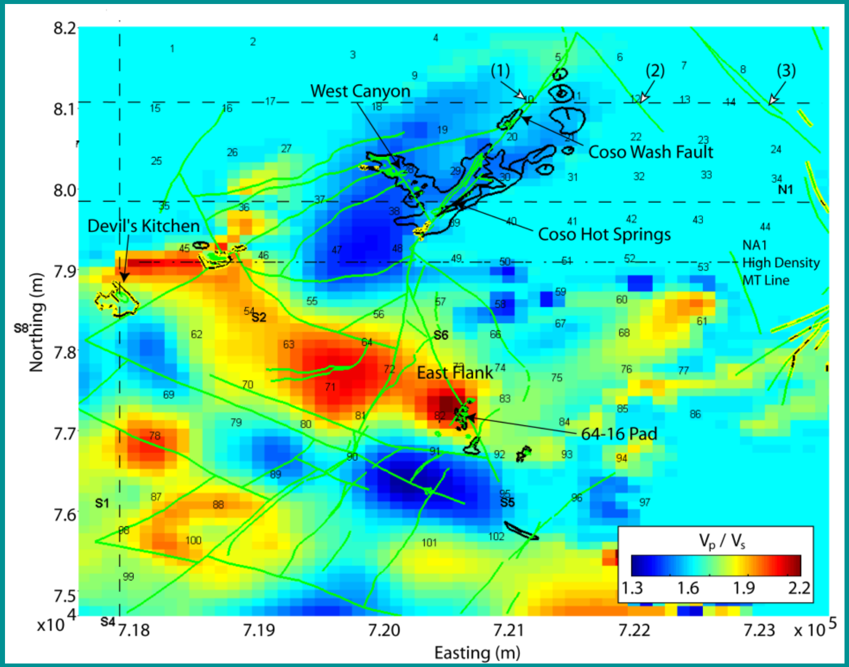

Conductivity Map
700 m below the surface