Present and Future Computing Requirements

Large-Scale Geophysical Imaging and Simulation

Gregory A Newman

Lawrence Berkeley National Laboratory

Geophysics Department

NERSC BES Requirements for 2017
October 8-9, 2013
Gaithersburg, MD

Why HPC => Time to Solution

- Why is this important?
  - More Science Can Get Done
  - More Breakthroughs
  - More Publications
  - More Realistic Models
  - More Understanding
  - End Member Simulations
  - Time Sensitive Decisions

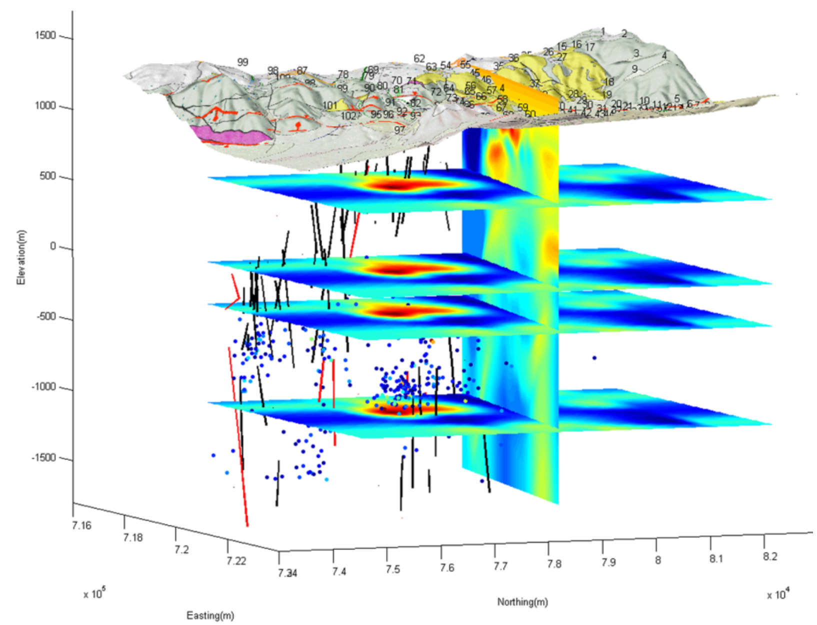

BERKELEY LAB

GEOPHYSICIAL IMAGING

- Seismic
  - 3D Reverse Time Migration
    - Large Scale Computations: 1,000s Cores, Weeks of Processing
  - 3D Elastic and Acousitc Full Waveform Inversion
    - Iterative reverse time migration
    - Promises Much Greater Image Fidelity
    - Formidable Numerical Issues – Local Minima, Very Good Starting Models Required
    - Frontier Research Area
    - Enormous Computation: 10,000's Cores, Months of Processing
- Electromagnetic (CSEM & MT)
  - 3D Full Waveform Inversion
    - Provides information on non-seismic attributes
    - Complements seismic imaging – through lower resolution
    - Constrained by seismic imaging
    - Computational demands also big: 1,000s to 10,000s cores
- Joint Seismic-Electromagnetic Imaging
  - The Holy Grail ?
    - Frontier Research Area
    - Grand Challenge Problem

Wave Equations for Geophysical Simulation and Imaging

Acoustic Waves

Time Domain

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

Frequency Domain

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

Electromagnetic Waves

$$\nabla \times \nabla \times \mathbf{E}_s + i\omega\mu\underline{\sigma}\,\mathbf{E}_s = \mathbf{S}.$$

Discretization Methods: Finite Differences, Finite Elements

Elastic Wave Field Simulation

First- order system for velocity –stress components

Laplace-Fourier Domain

$v_{x,y,z}$ - velocity components,

$\tau_{xx,xy,xz,yz,yy,zz}$ - stress components,

ρ - density,

λ and μ - Lame coefficients.

$$s\rho v_x = div\left(\vec{\tau}_x\right) + f_x,\ \vec{\tau}_x = \left(\tau_{xx},\tau_{xy},\tau_{xz}\right);$$
$$s\rho v_y = div\left(\vec{\tau}_y\right) + f_y,\ \vec{\tau}_y = \left(\tau_{xy},\tau_{yy},\tau_{yz}\right);$$
$$s\rho v_z = div\left(\vec{\tau}_z\right) + f_z,\ \vec{\tau}_z = \left(\tau_{xz},\tau_{yz},\tau_{zz}\right),$$
$$s\tau_{xy} = \mu\left(\partial_y v_x + \partial_x v_y\right);$$
$$s\tau_{xz} = \mu\left(\partial_z v_x + \partial_x v_z\right);$$
$$s\tau_{yz} = \mu\left(\partial_z v_y + \partial_y v_z\right);$$
$$s\tau_{xx} = \lambda div\left(\vec{v}\right) + 2\mu\partial_x v_x;$$
$$s\tau_{yy} = \lambda div\left(\vec{v}\right) + 2\mu\partial_y v_y;$$
$$s\tau_{zz} = \lambda div\left(\vec{v}\right) + 2\mu\partial_z v_z.$$

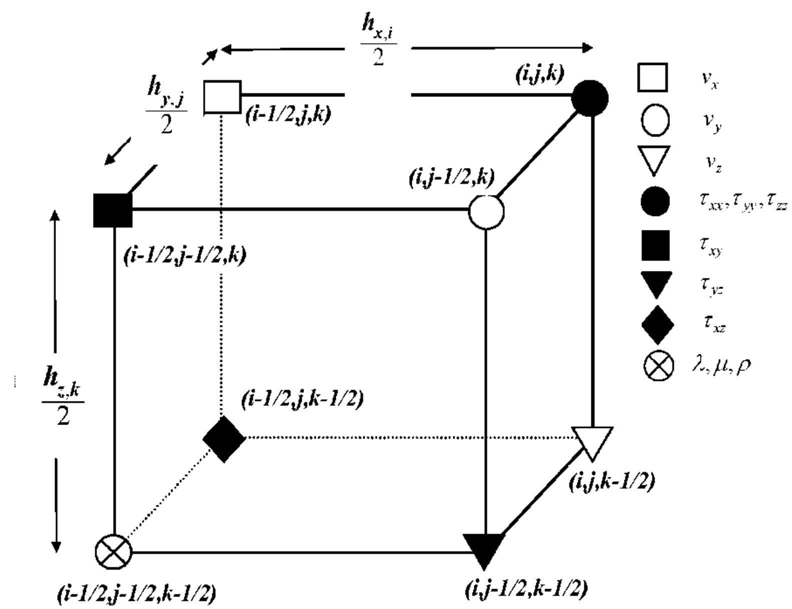

$$\hat{\mathbf{M}} = \begin{pmatrix} M_{xx} & M_{xy} & M_{xz} \\ M_{xy} & M_{yy} & M_{yz} \\ M_{xz} & M_{yz} & M_{zz} \end{pmatrix}$$

Forces $f_{x,y,z}$ are defined via $\nabla \bullet \hat{\mathbf{M}}$

Moment-Tensor components (R. Graves 1996)

LARGE-SCALE MODELING & IMAGING CONSIDERATIONS

- Require Large-Scale Complex Modeling and Imaging Solutions
  - 10's of million's field unknowns (fwd problem; Maxwell's & Poisson's, acoustic and elastic field wave equations)
    - Solved with finite difference approximations & iterative Krylov solvers
  - Imaging grids 400 nodes on a side
    - Exploit gradient optimization & implicit Gauss-Newton schemes, adjoint state methods

- Parallel Implementation
  - Domain Decomposition Techniques, MPI Interconnect fabric
  - Two levels of parallelization
    - Model Space (simulation and inversion mesh)
    - Data Space (each transmitter/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

HPC MODELING & INVERSE MODELING

- FURTHER CONSIDERATIONS
  - Sometimes Smaller Model Parameterizations Encountered
    - Induction logging, but still 1000's of fwd solves needed for imaging
    - Stochastic imaging

- Parallel Implementation Considerations
  - Will the application scale with processors employed => 10's to 10,000's ?
    - Reduction in time to solution
    - Efficient exploitation of resources
    - Shows capability to attack large scale problems that cannot be solved otherwise

Solver Selection

- Choice depends on problem:
  - Direct Solvers
    - Multiple Right Hand Side Solutions
    - Robust with Respect to Mesh Design
    - Requires Matrix Factorization – expensive and time consuming for large meshes
    - Parts of the solver solution inherently non-parallel (triangular forward and back solves)
    - Parallel Solvers: MUMPS, SUPER LU, PARDISO
  - Iterative Solvers
    - Single Right Hand Side Solution
    - Sensitive to Mesh Design – Preconditioning Required
    - Highly Efficient Solution Process for Large Meshes
    - Parallel Krylov Solvers: Your Own, PETSE and TRILINOS Libraries
    - Algebraic Muligrid and Preconditions

Some HPC Applications

- Resistivity Mapping of Hydrocarbons
- End Member Solutions (SEAMS Resistivity Model)
- Geothermal Resource Evaluation
- Joint EM & Seismic Imaging
- 3D Elastic Wave-Field Simulation & Imaging

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 curren
- Depth of interrogation ~ 3 to 4 km
- Sensitive to thin resistive beds

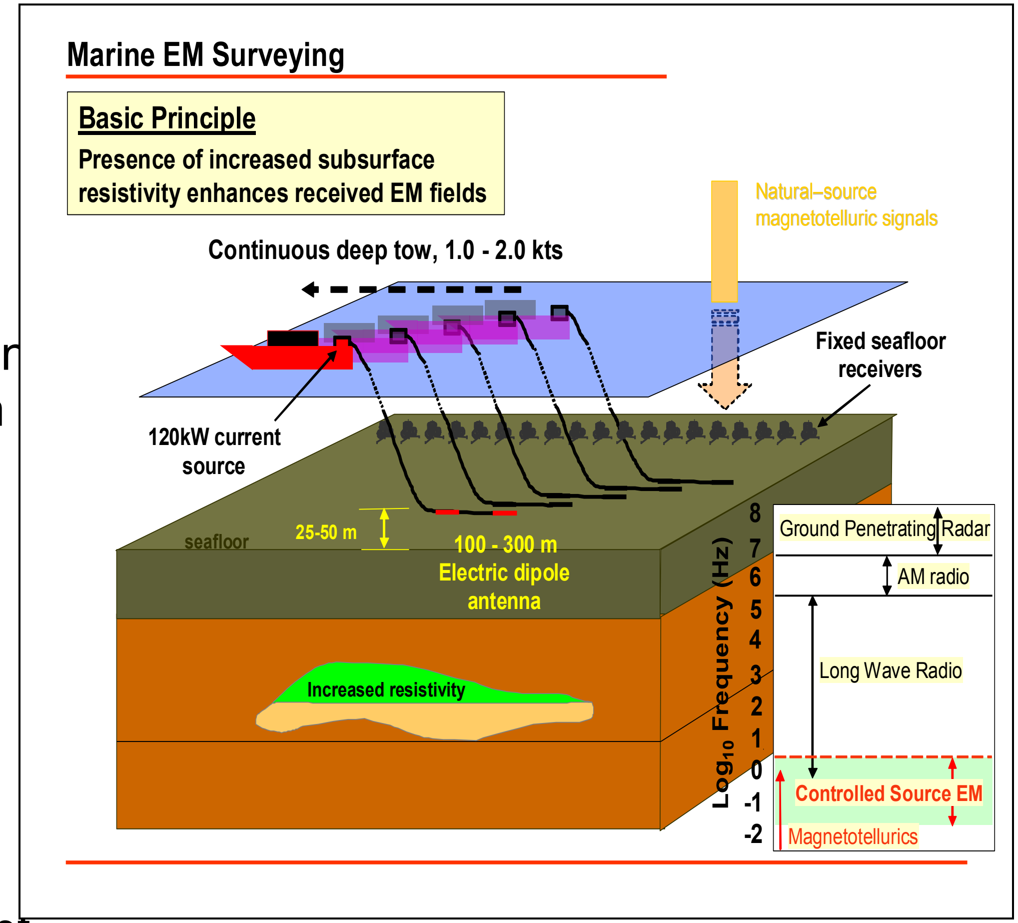

MT

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

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

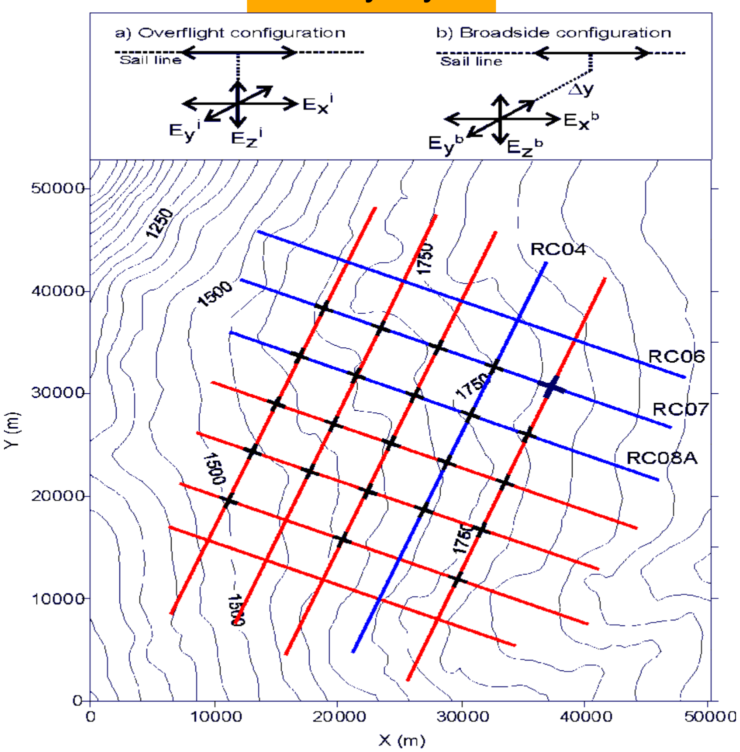

Overflight electric field

Broadside electric field

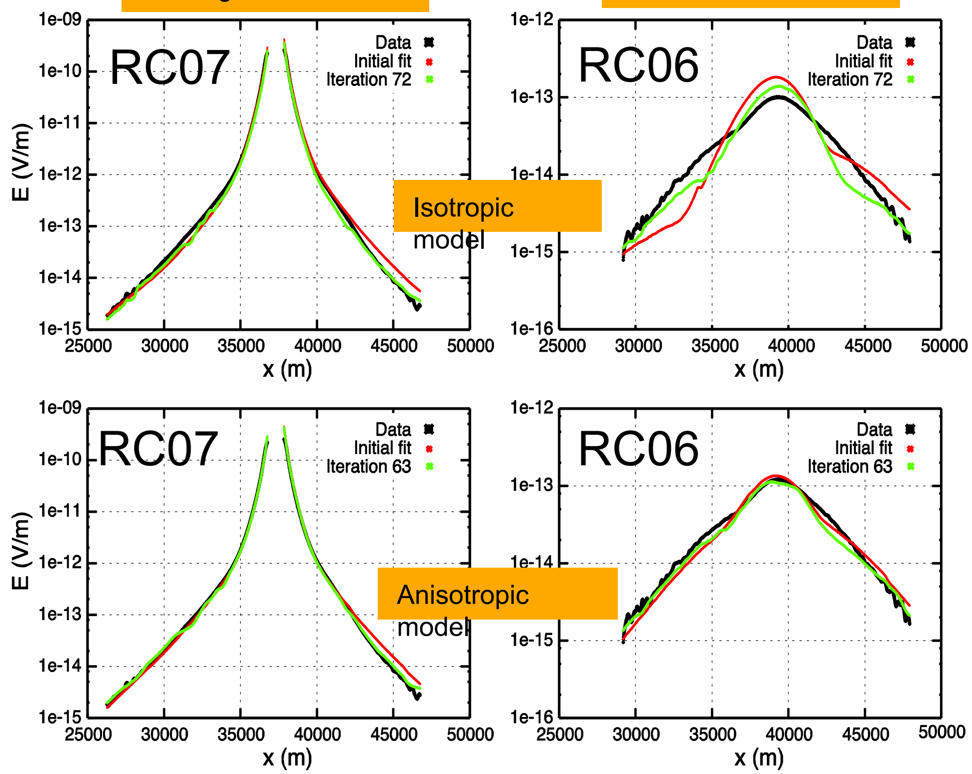

3D CSEM Resistivity Imaging

Offshore Brazil

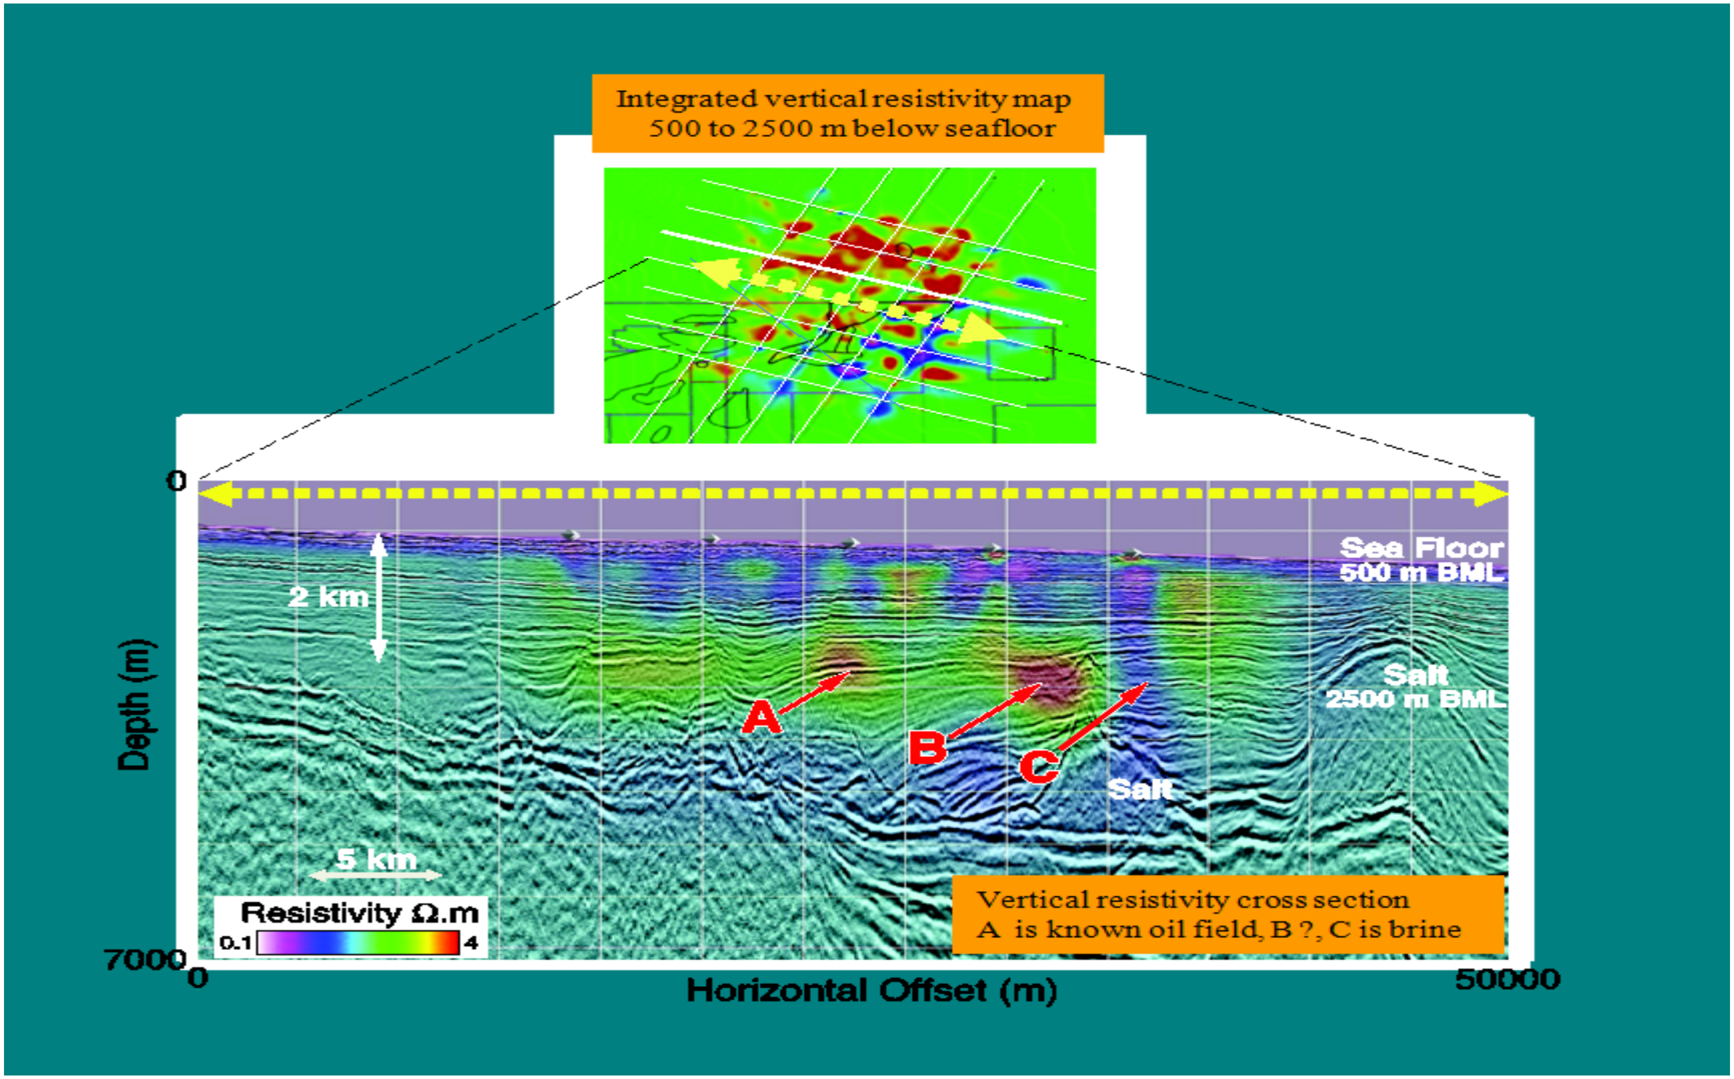

The Marmousi Model

Imaging results for shear velocity

- Exact model
- Inverted model (3,5,6 Hz)

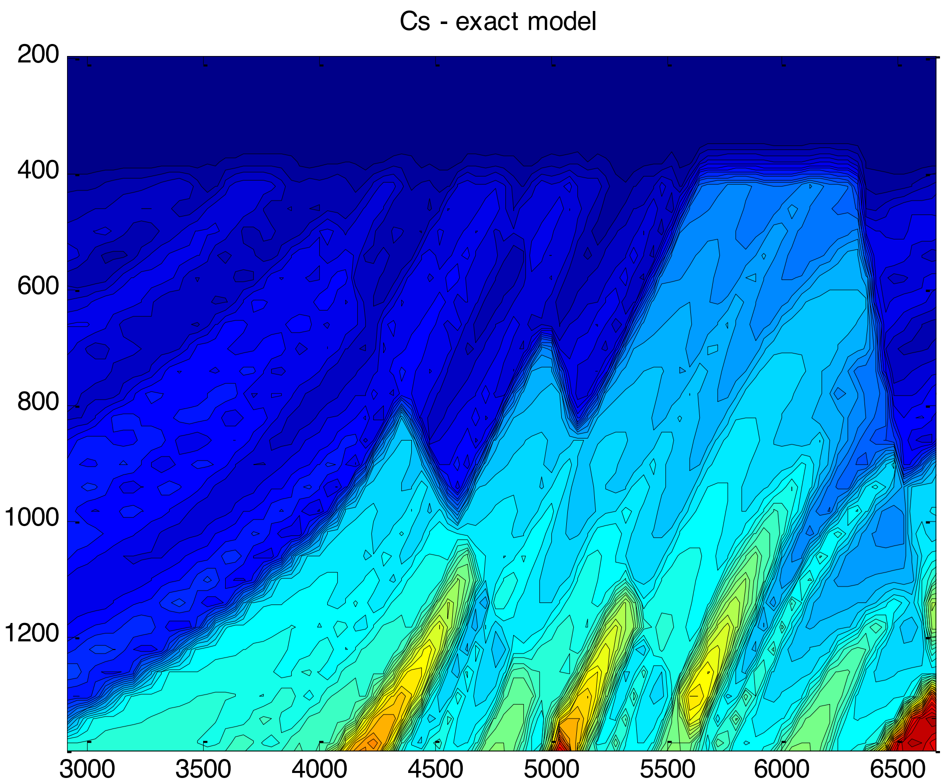

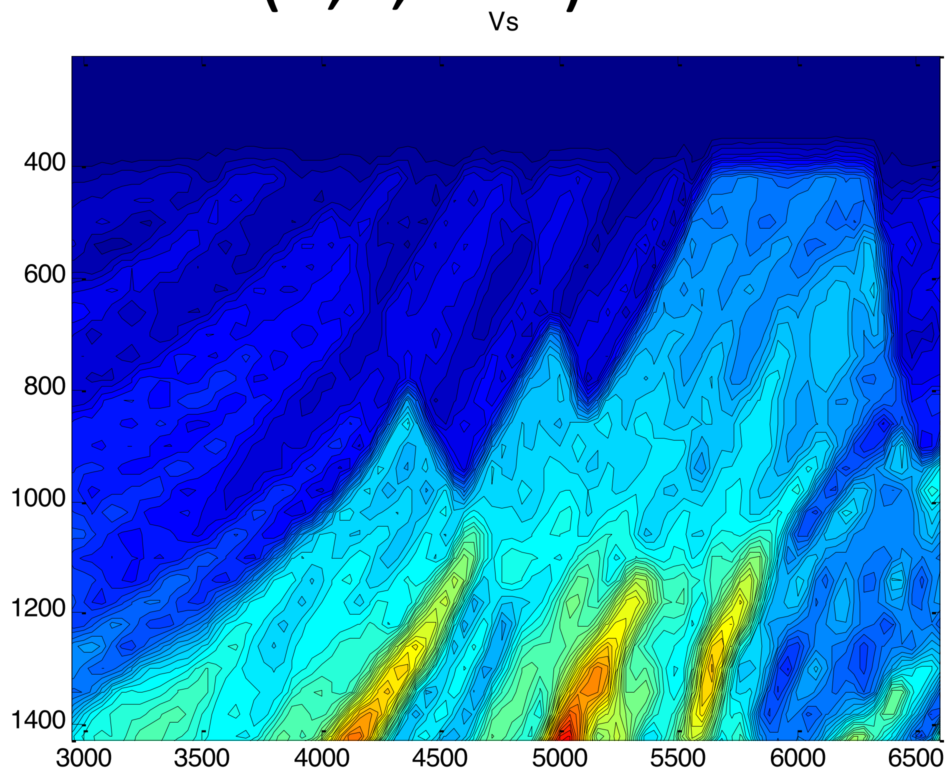

Measuring HPC Performance

Parallel scaling of the Elastic simulator

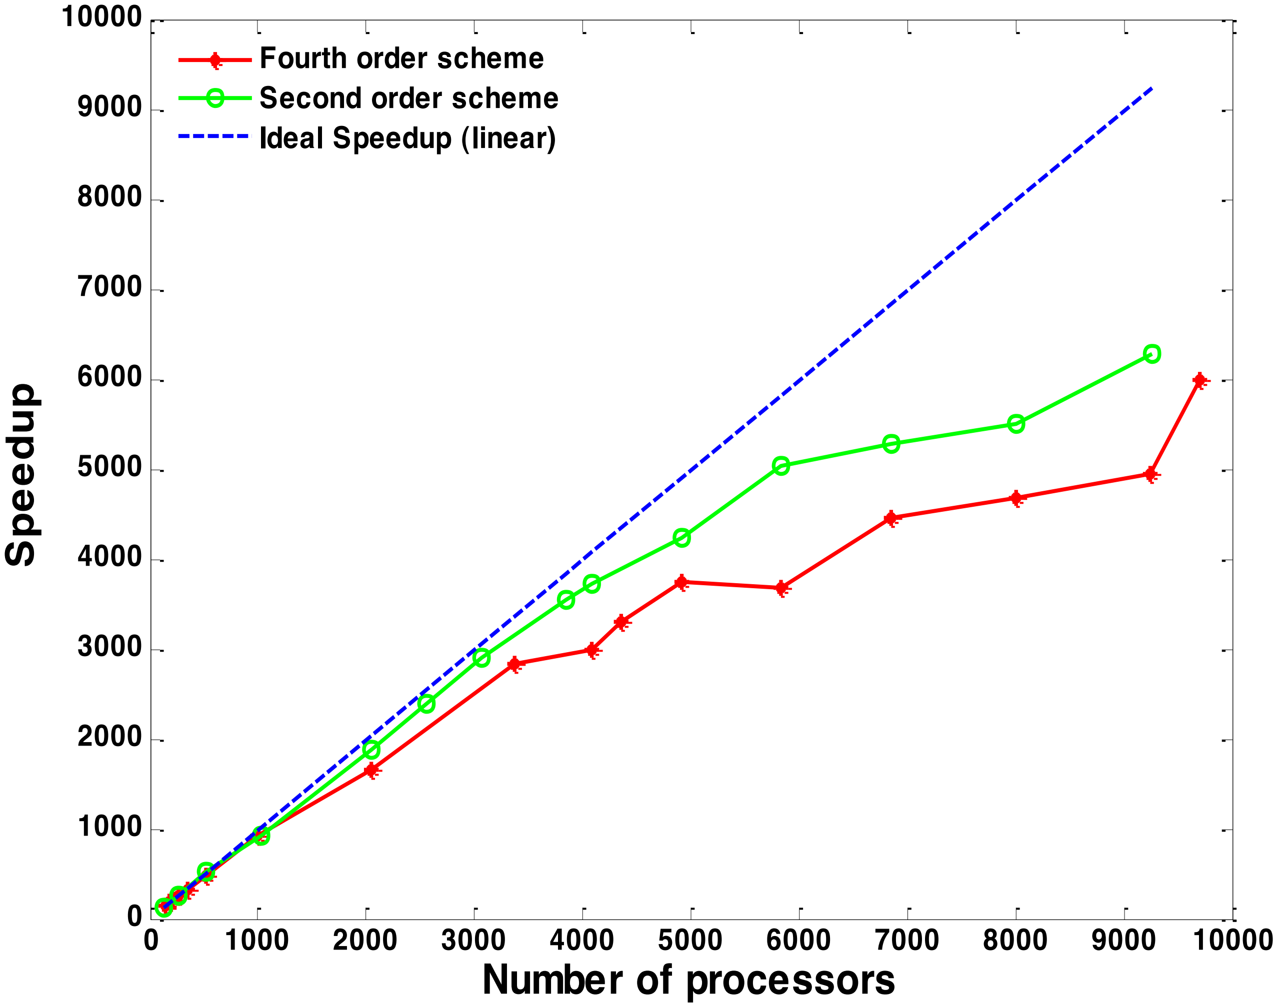

Speedup

$$S = \frac{T_1(n)}{T_p(n)} \leq p,$$

where p="ideal speedup",

$$T_1(n), T_p(n)$$

are the times for running a problem of size n on 1 and p processors

Scaling curves for a fixed-size (588x588x261) problem run on Cray XT4 – NERSC Franklin System

Geophysical Simulations on GPUs

Main challenge:

Manage memory access in most efficient way

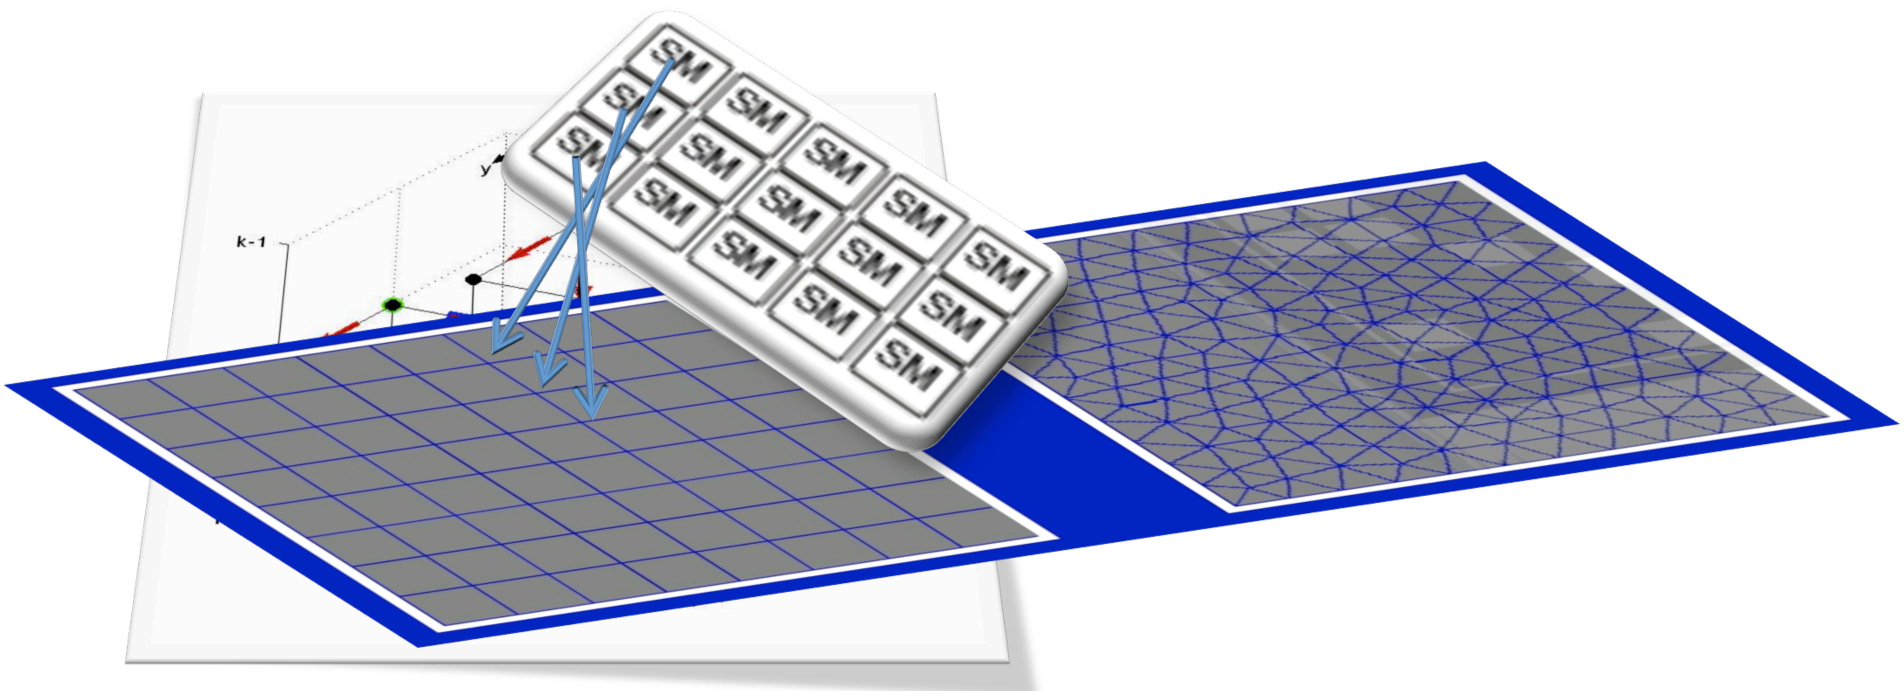

Iterative Krylov Solver Performance Tests

Typically used for EM problems:

CG, BiCG, QMR

Computing times
for 1000 Krylov solver iterations

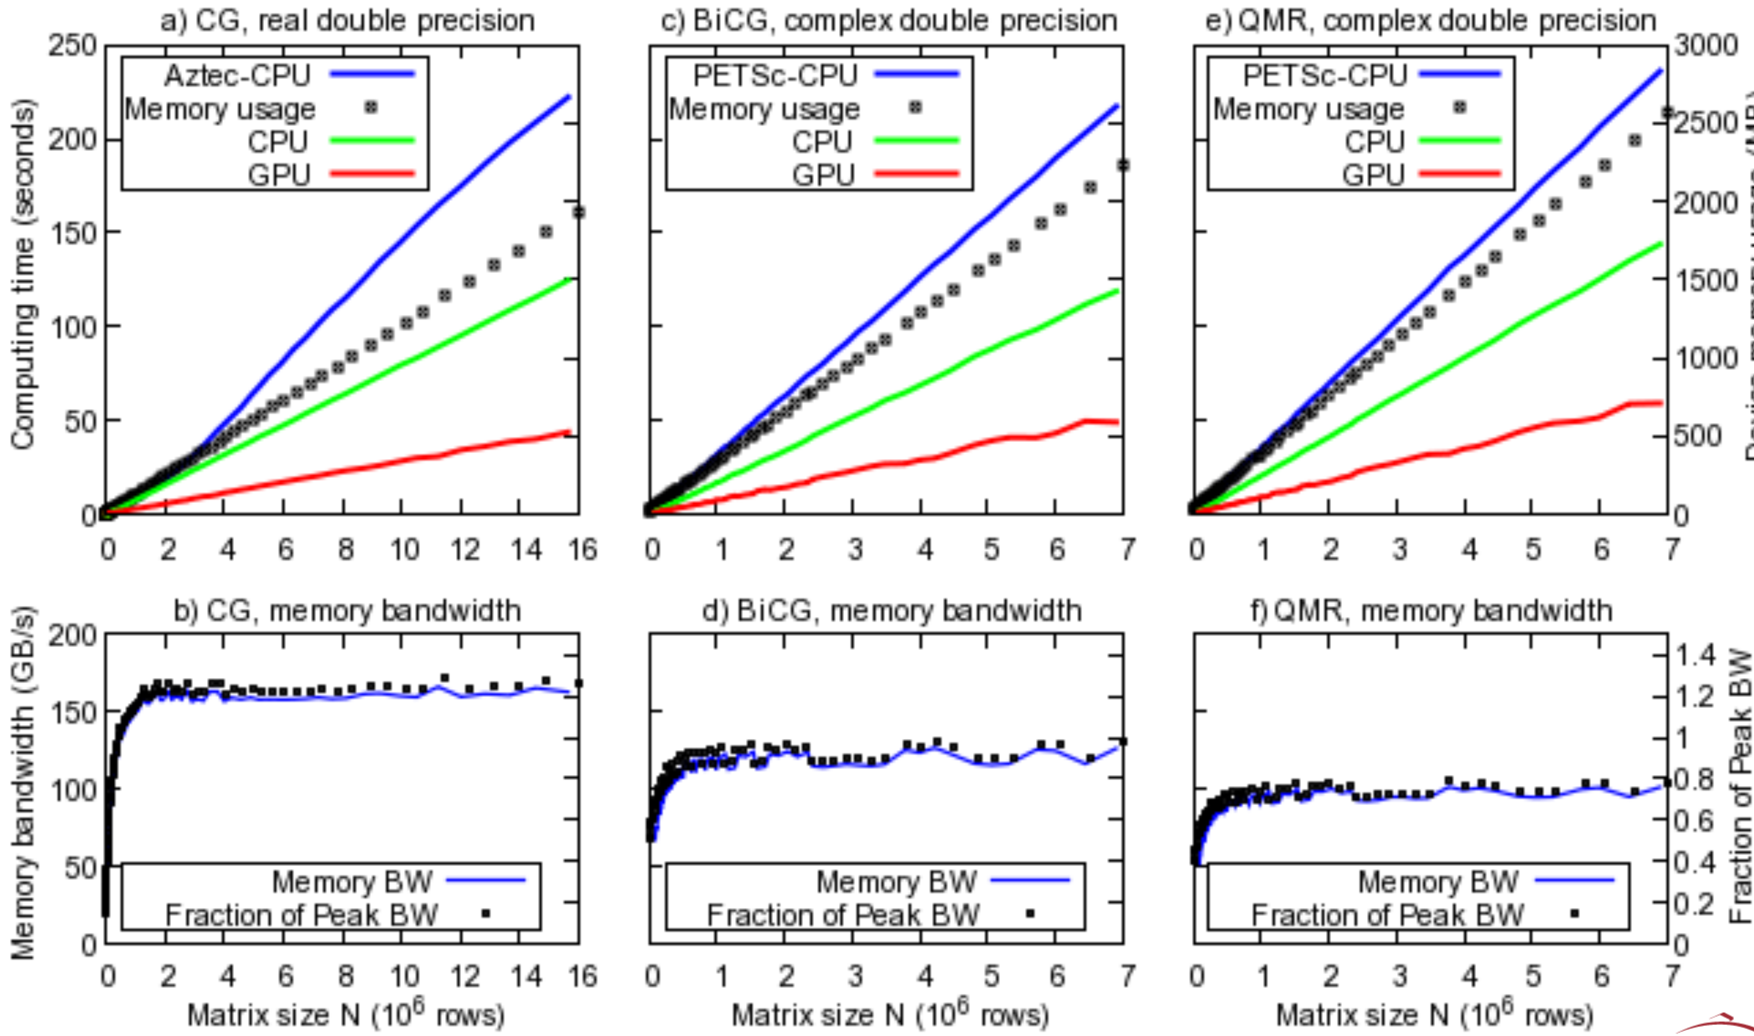

BERKELEY LAB

GPU/CPU-MPI Comparisons

- For Largest Problems Tested:
  - 1 GPU (448 processor cores)
    - Equivalent to 23 CPU's for CG iteration (DC & IP Problems)
    - Equivalent to 19 QMR and BiCG iterations (EM and MT)
- GPU's Impressive, but not good enough for now.
  - Marine CSEM and MT Imaging (Production)
    - Use routinely 64 to 512 CPU cores per fwd solve
    - What about multiple GPU's with MPI (too slow for now)
  - Elastic Wave field modeling and Imaging
    - Similar performance comparisons are expected

Current HPC Usage

- Machines currently using **Hopper and Edison at NERSC**
- Hours used in 2012-2013 is now approaching **16M**
- Biggest Jobs > **5000 to 20,000 compute cores**
- Run times per job **24 to 36 hours**
- Data read/written per run: **approximately 16 to 160 Gigabytes of data written to scratch mostly for for check pointing.**
- Maximum Memory used per (node **16 Gbytes**| core **0.5 Gbytes**| globally **2.5 to 5 Terrabytes**)
- Necessary software, services or infrastructure: **Fortan 90, 95, C, C++, MPI, Cuda**

HPC Requirements for 2017

- Science goals: **solve problems approaching 10^9 grid nodes; elastic wave field simulation and imaging, joint imaging experiments, treat larger data volumes**
- Compute hours needed (in units of Hopper hours) **> 30M**
- Faster Solvers: **massively parallel algebraic Mulitgird, designed Specifically from complex and complex-symmetric linear Systems.**
- Changes to memory needed per (**2-4x** core | **2-4x** node | **2-4x** globally)
- Changes to necessary software, services or infrastructure: **hybrid computing systems are coming (multi-core GPU-CPU-MPI interconnects)**
- Legacy Software: **porting to such hybrid machines will be an issue**