Toward Exascale Seismic Imaging & Inversion

Jeroen Tromp
Department of Geosciences
Program in Applied & Computational Mathematics
Princeton Institute for Computational Science & Engineering

Ebru Bozdağ, Dimitri Komatitsch, Lion Krischer, Matthieu Lefebvre, Wenjie Lei, Daniel Peter & James Smith
ORNL: Judy Hill, Norbert Podhorszki & David Pugmire

Software

SPECFEM3D_Cartesian

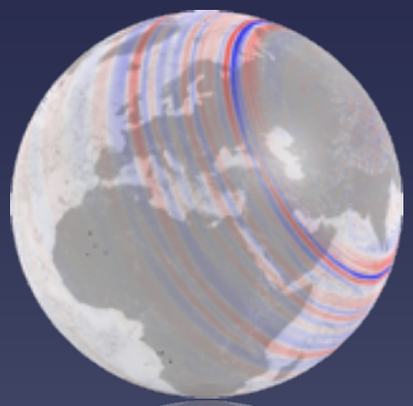

SPECFEM3D_GLOBE

Spectral-Element M

- Hexahedral finite-elements
- Gauss-Lobatto-Legrendre quadrature
- Diagonal mass matrix
- Explicit time-marching scheme

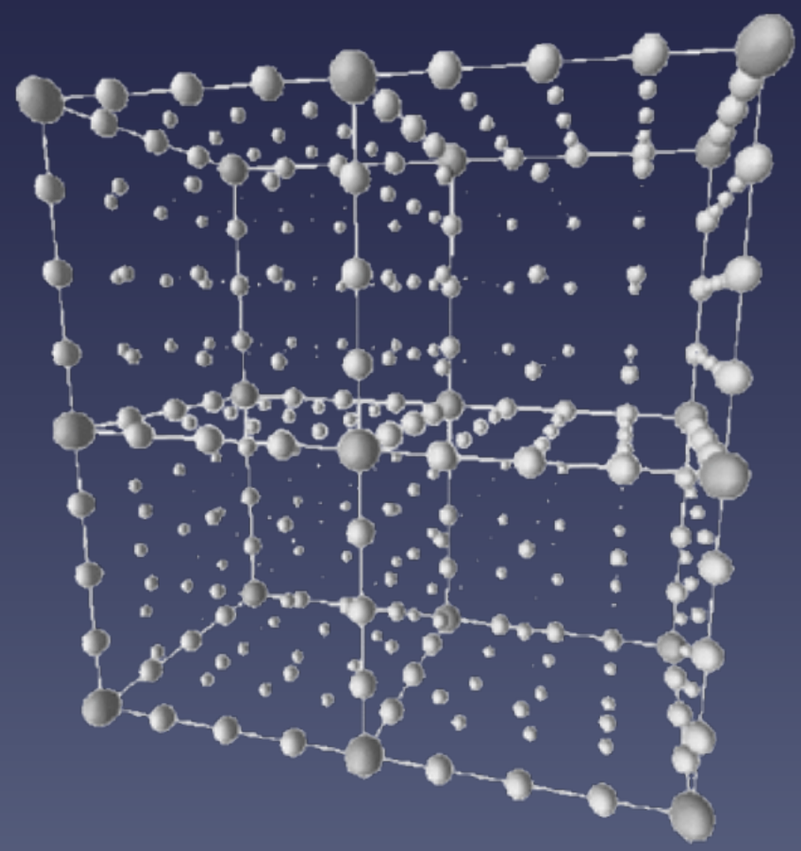

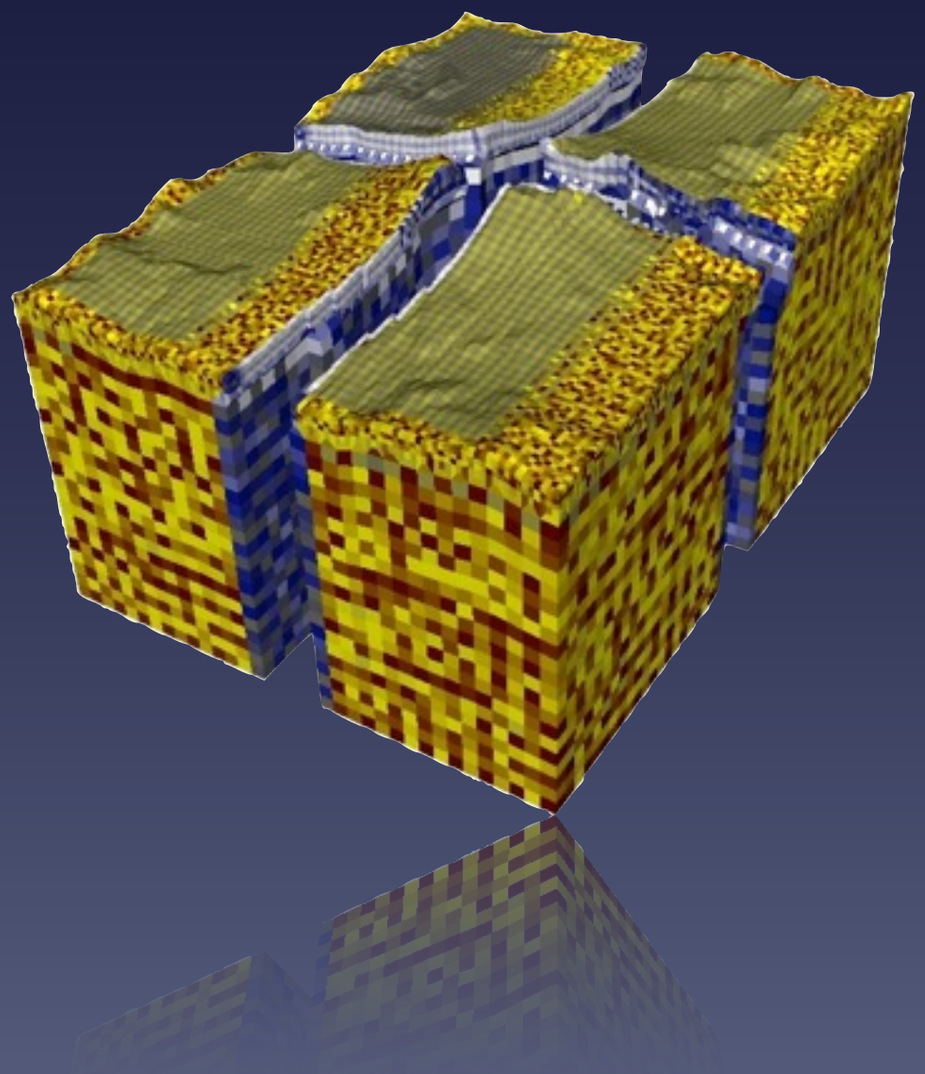

SPECFEM3D_Cartesian

SPECFEM3D

Open-Source

- Unstructured meshes
- Load-balanced mesh partitioning
- Fluid-solid coupling
- Anisotropy
- Attenuation
- Adjoint capabilities

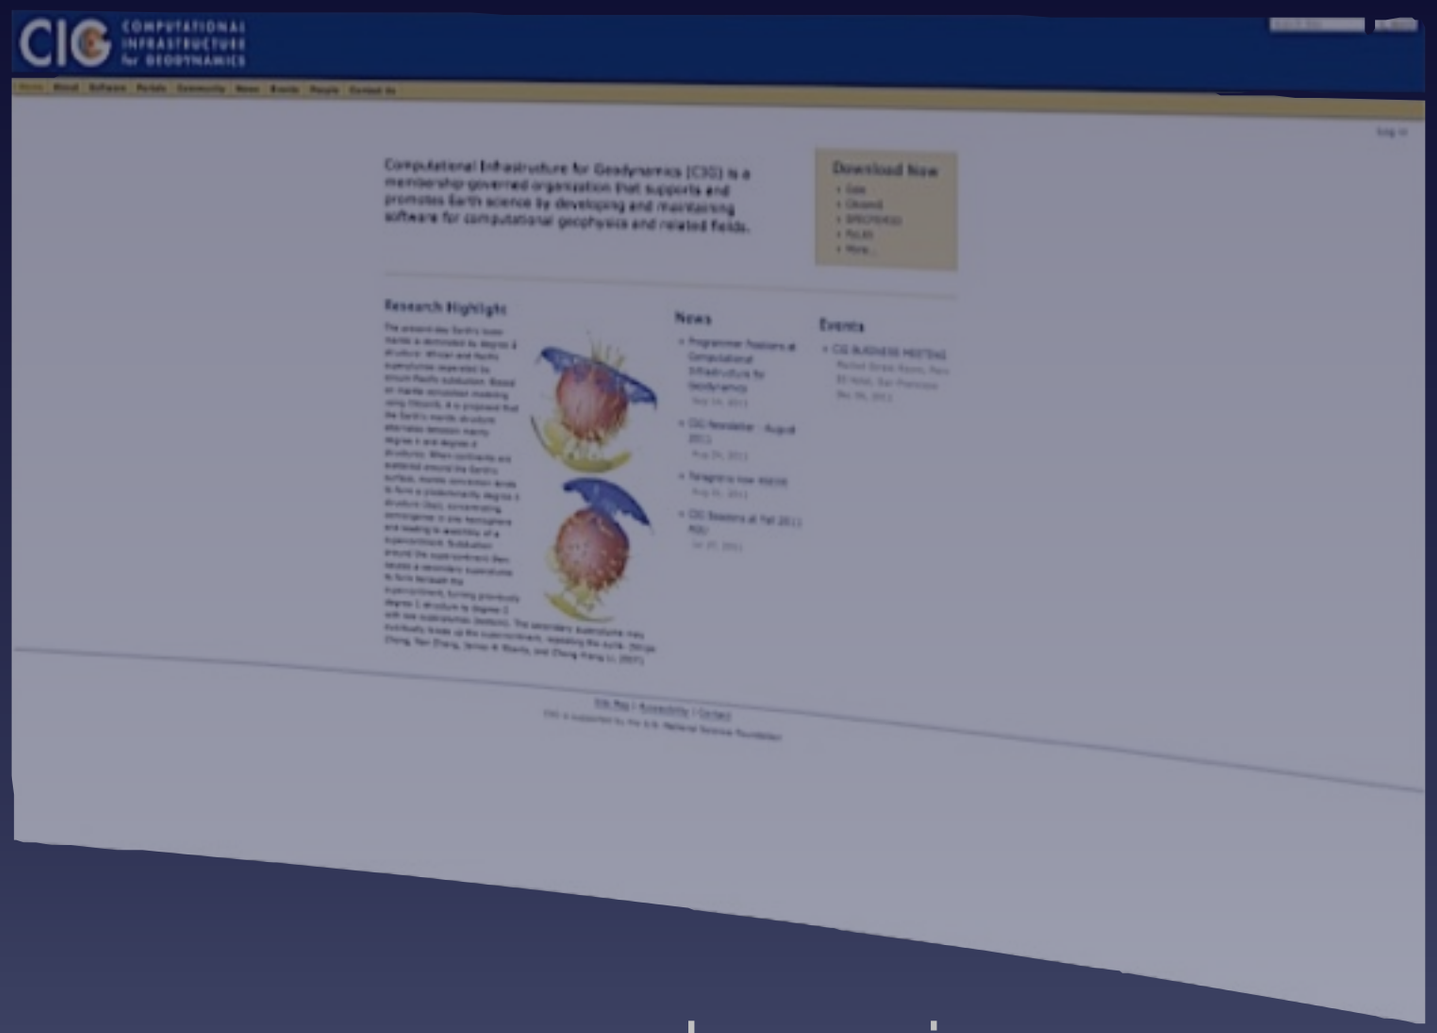

www.geodynamics.org

SPECFEM3D

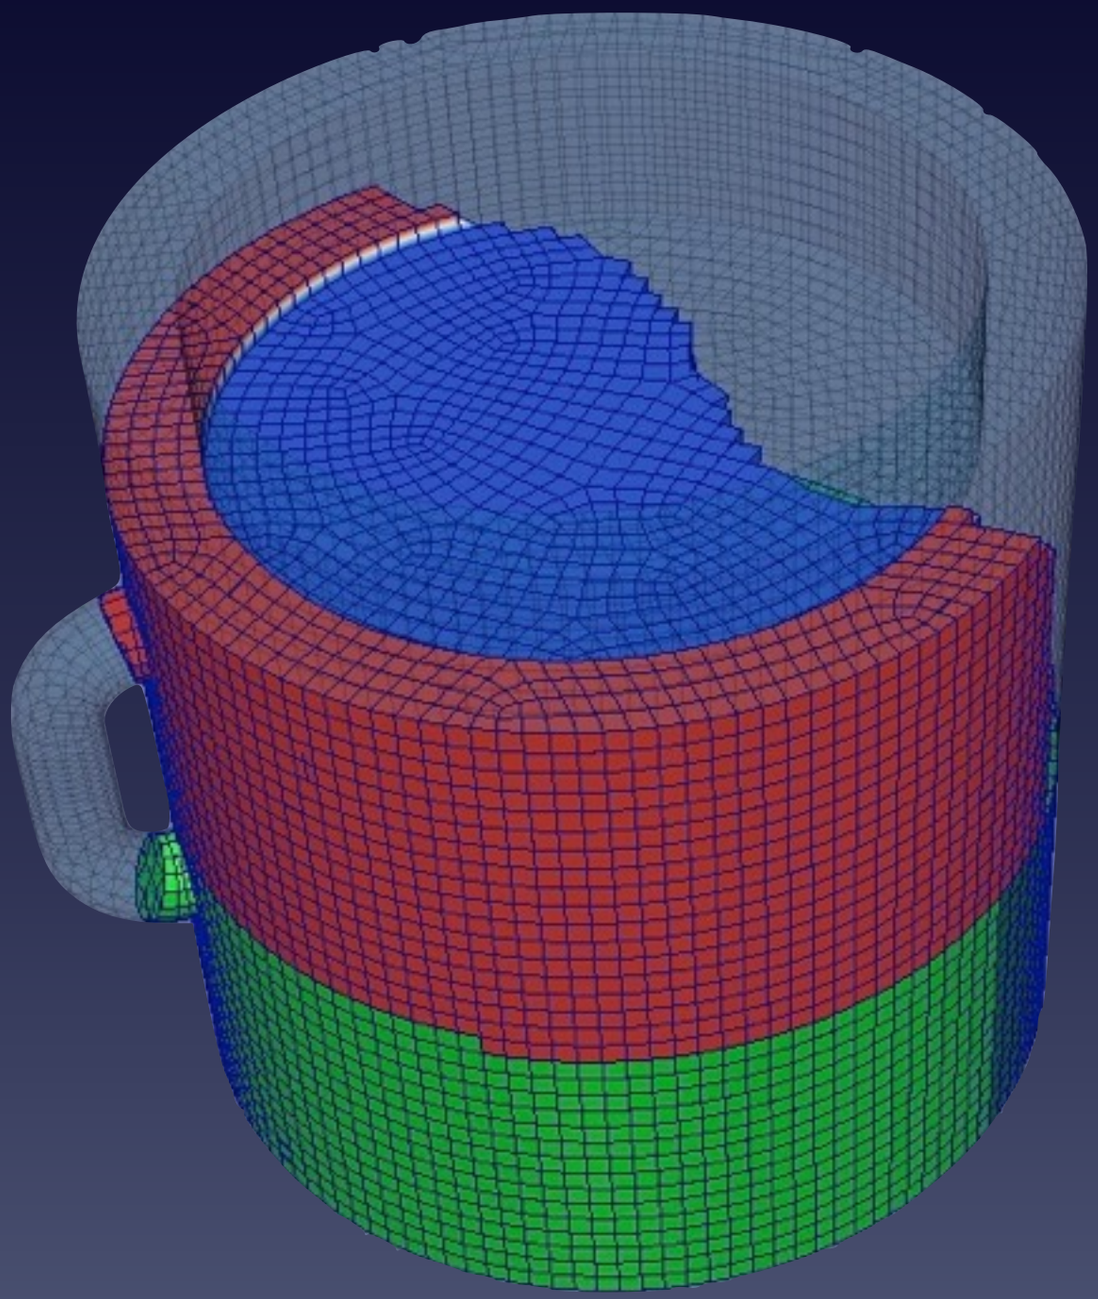

SPECFEM3D

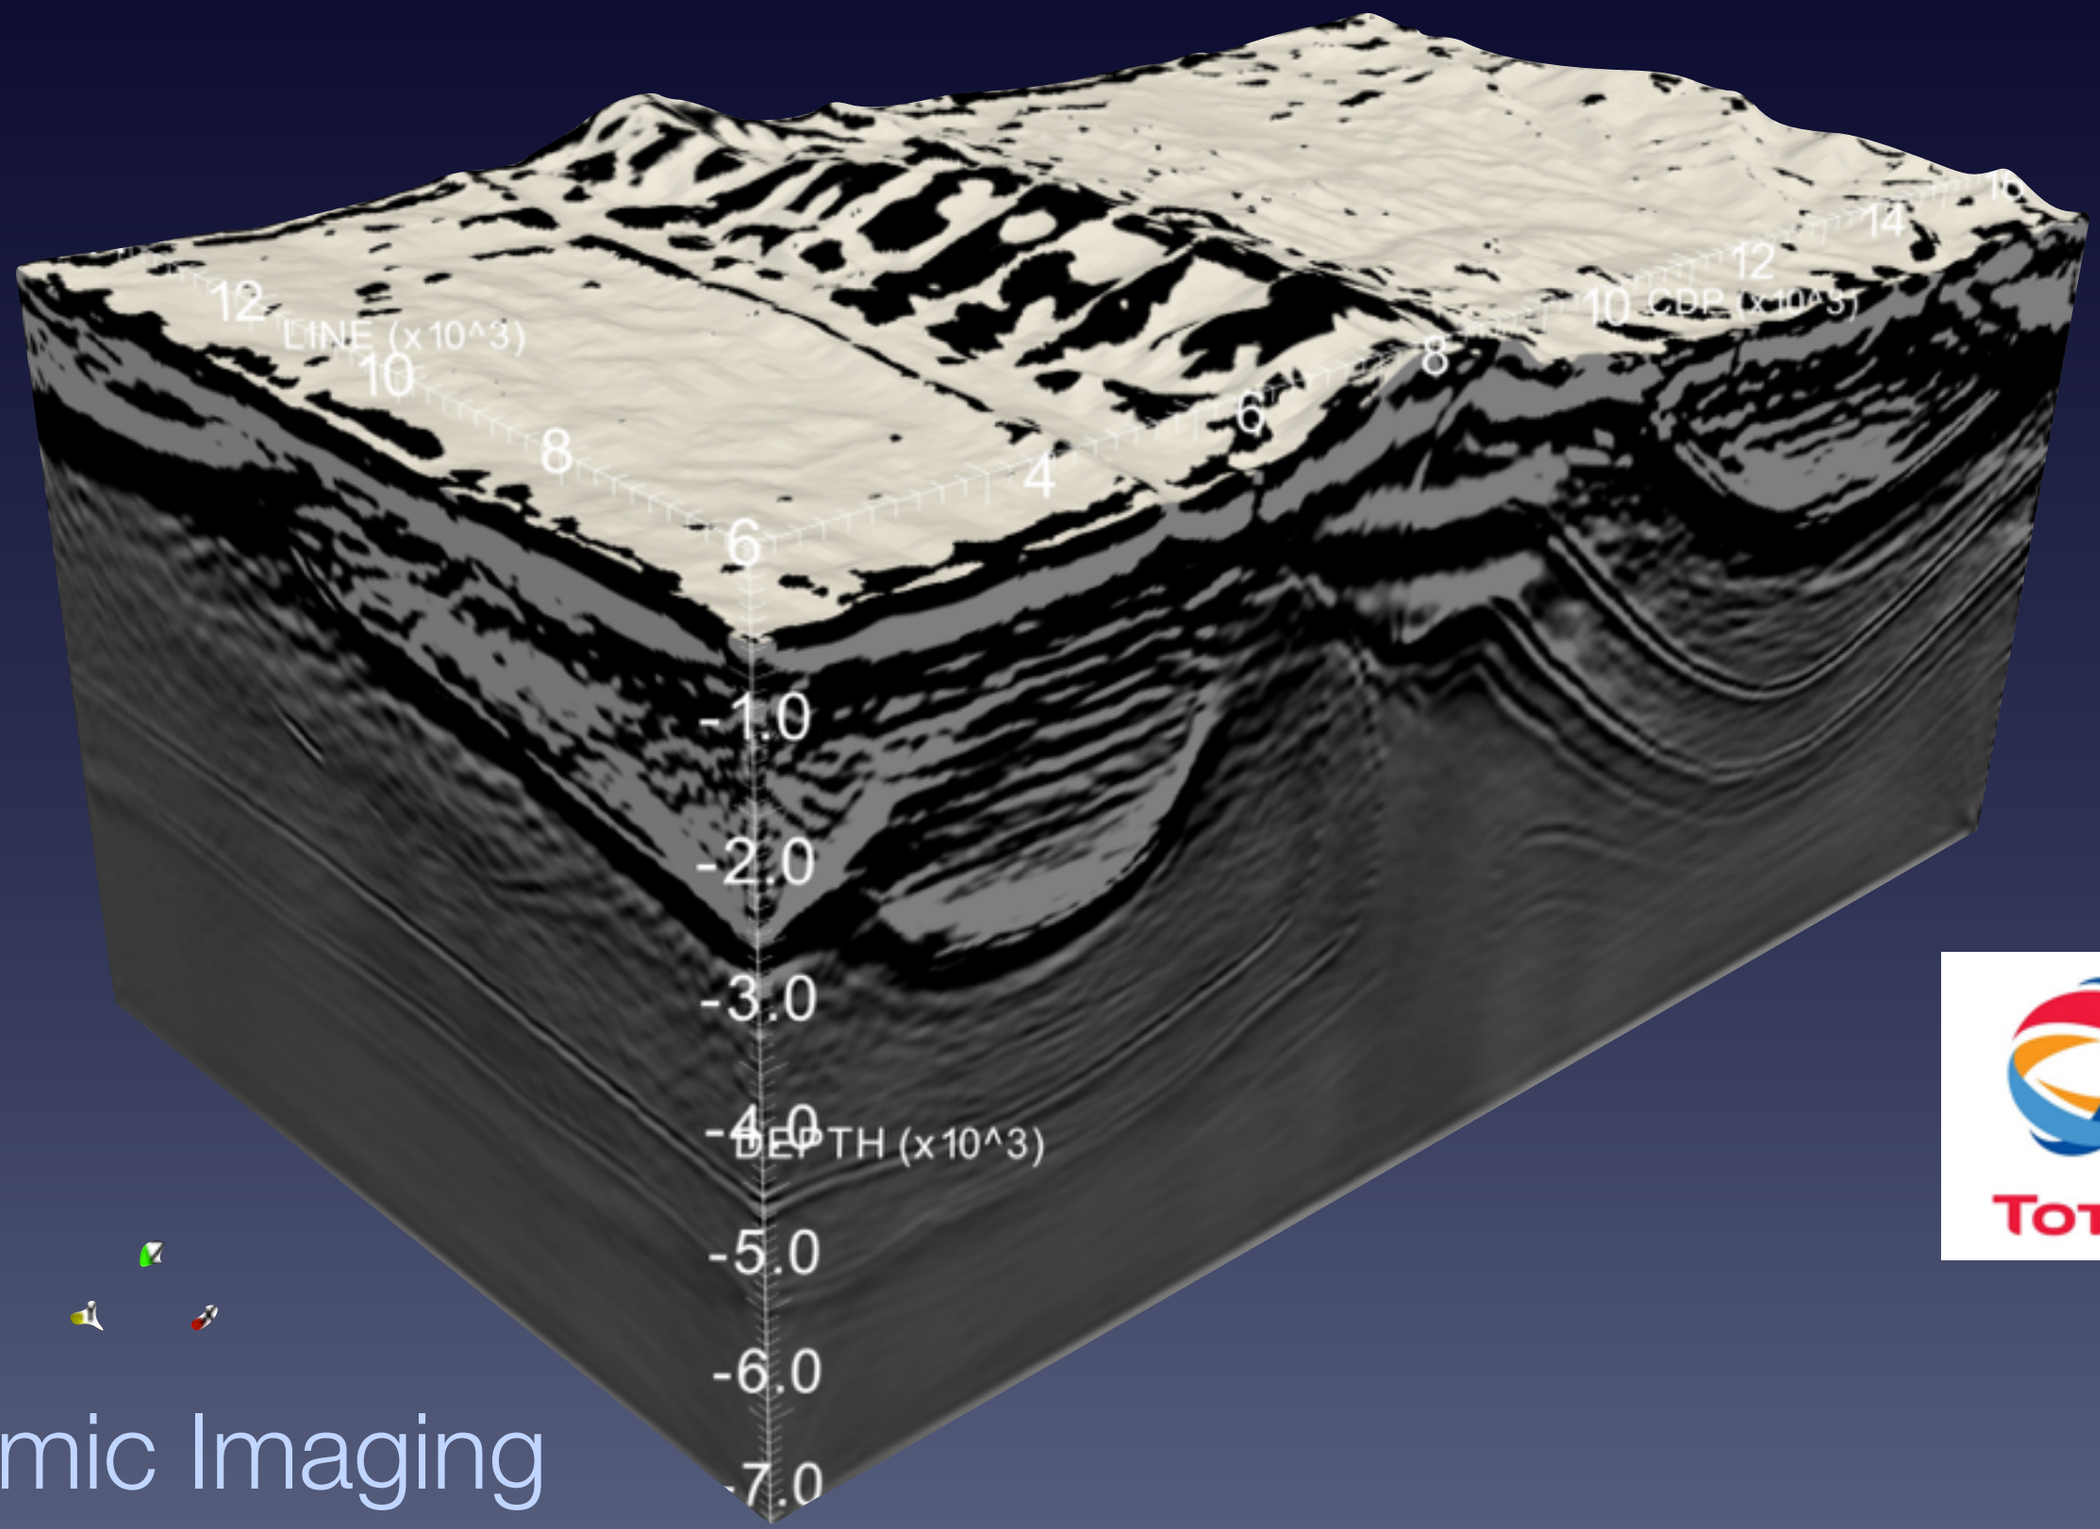

Seismic Imaging

Yang Luo

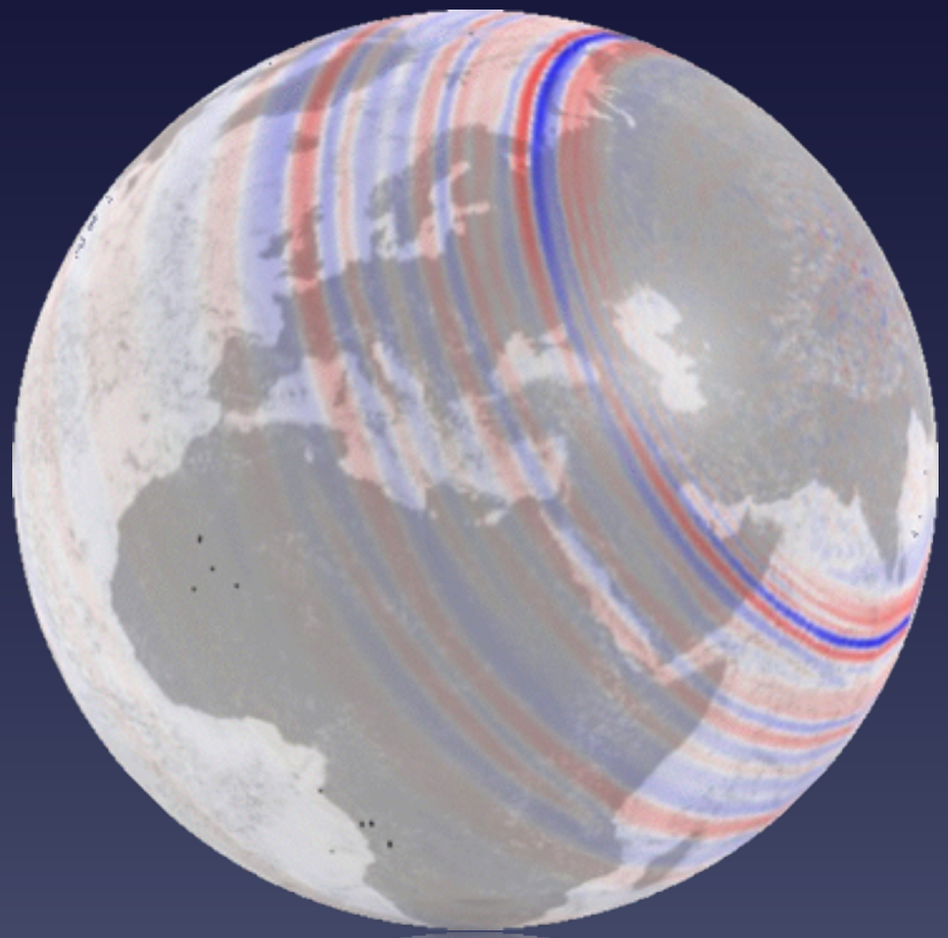

SPECFEM3D_GLOBE

SPECFEM3D_GLOBE

Open-Source

- 3D crust & mantle models
- Topography & bathymetry
- Rotation
- Ellipticity
- Gravitation
- Anisotropy
- Attenuation
- Adjoint capabilities

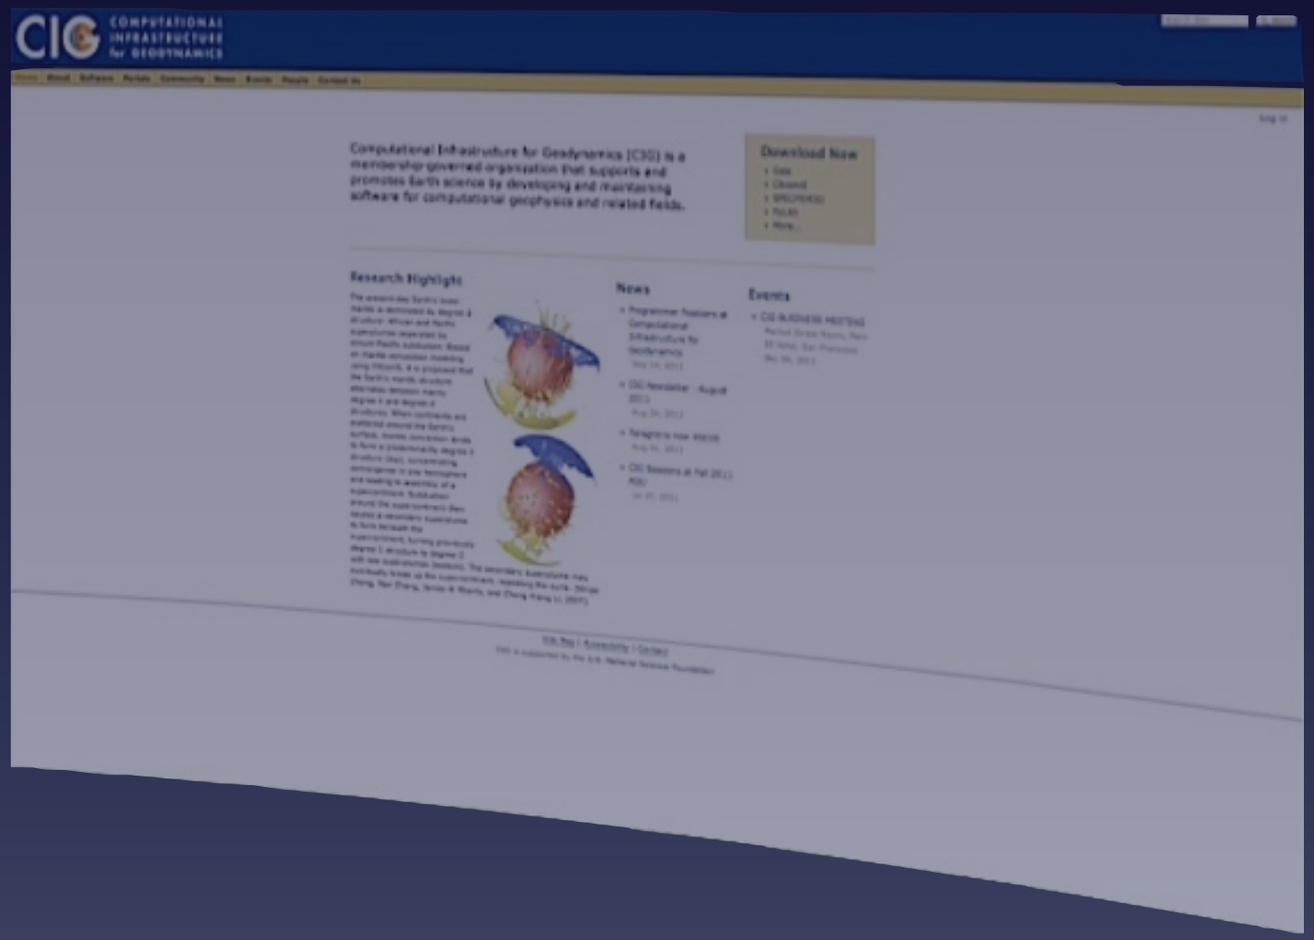

www.geodynamics.org

ShakeMovie

http://global.shakemovie.princeton.edu

- Automatically triggered by global earthquakes
- 1D & 3D synthetics
- Public outreach movies
- Provided by Incorporated Research Institutions for Seismology (IRIS)

Global Seismic Tomography

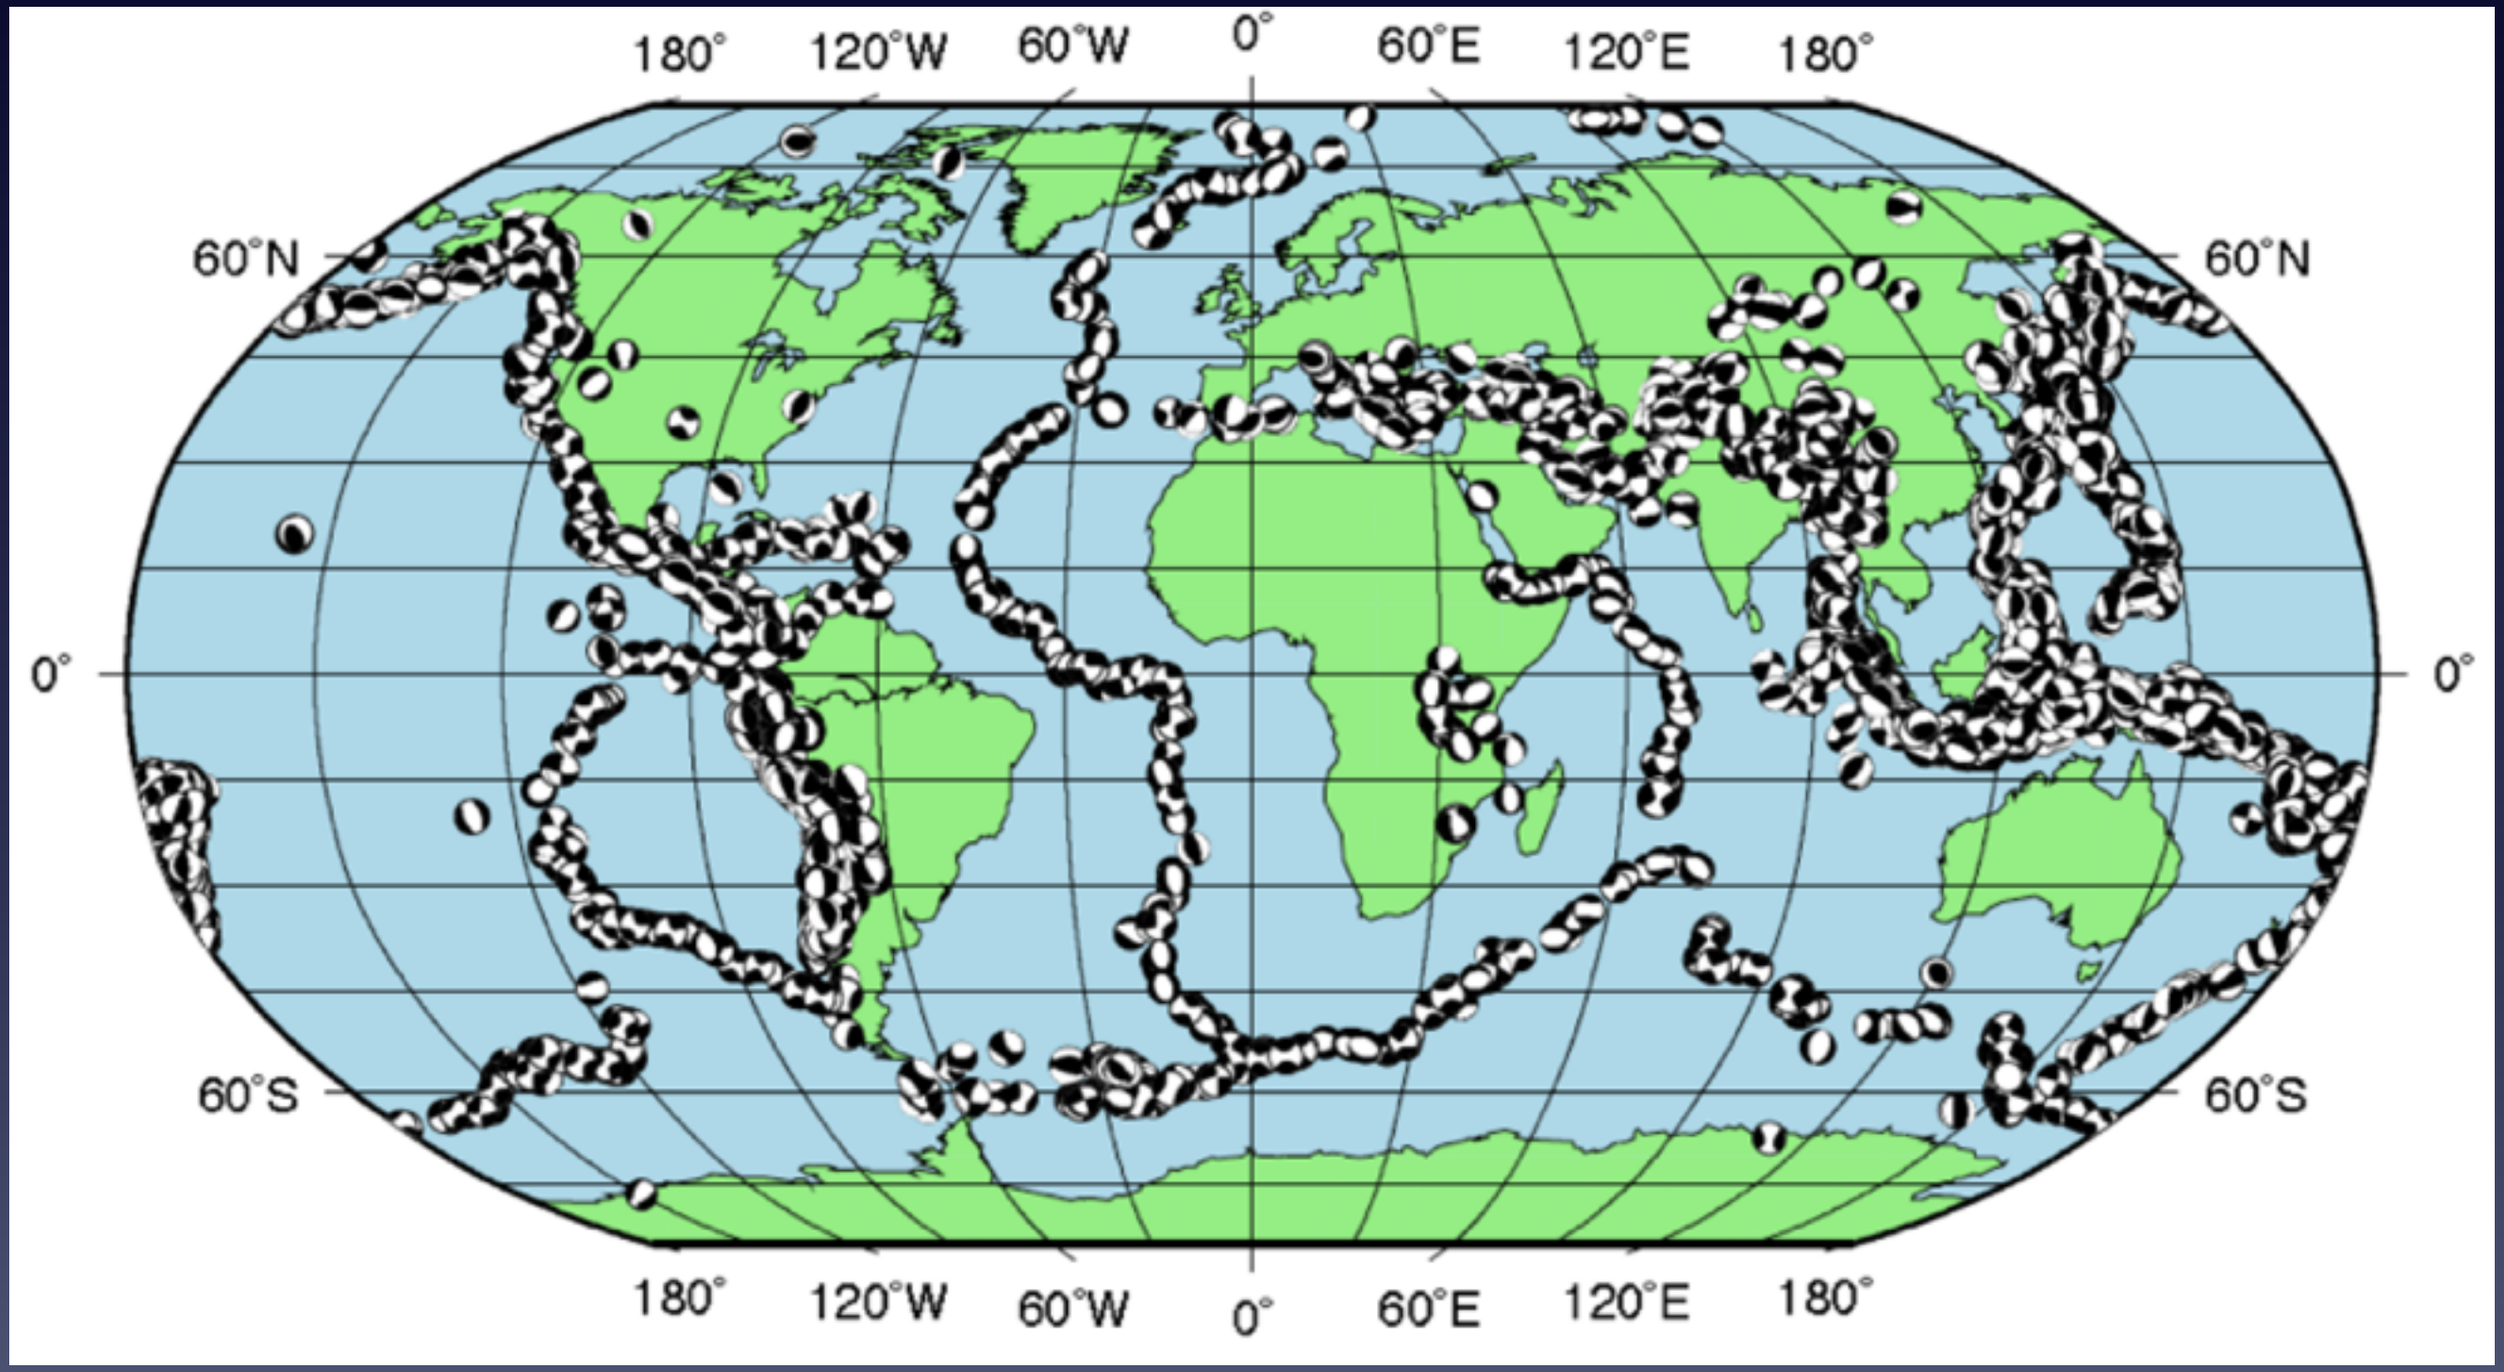

Assimilation of ~100 million data

High-performance computing

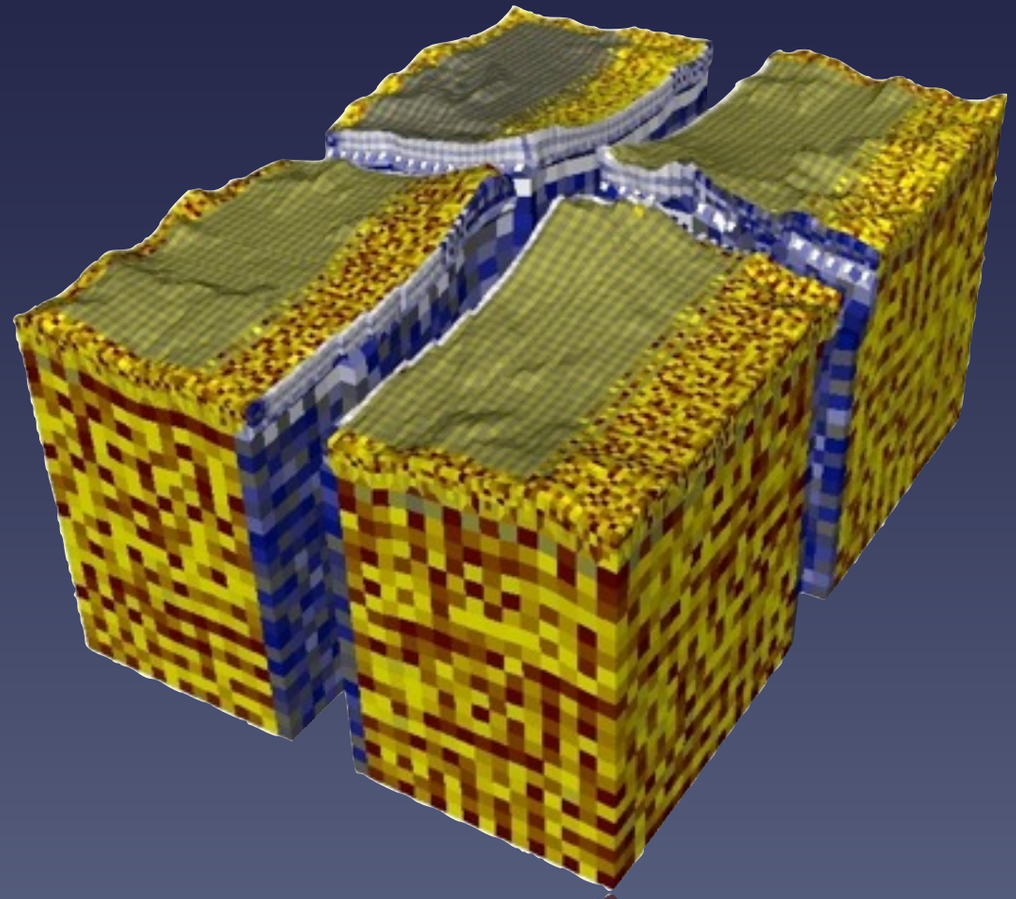

GPU C

Titan

Piz Daint

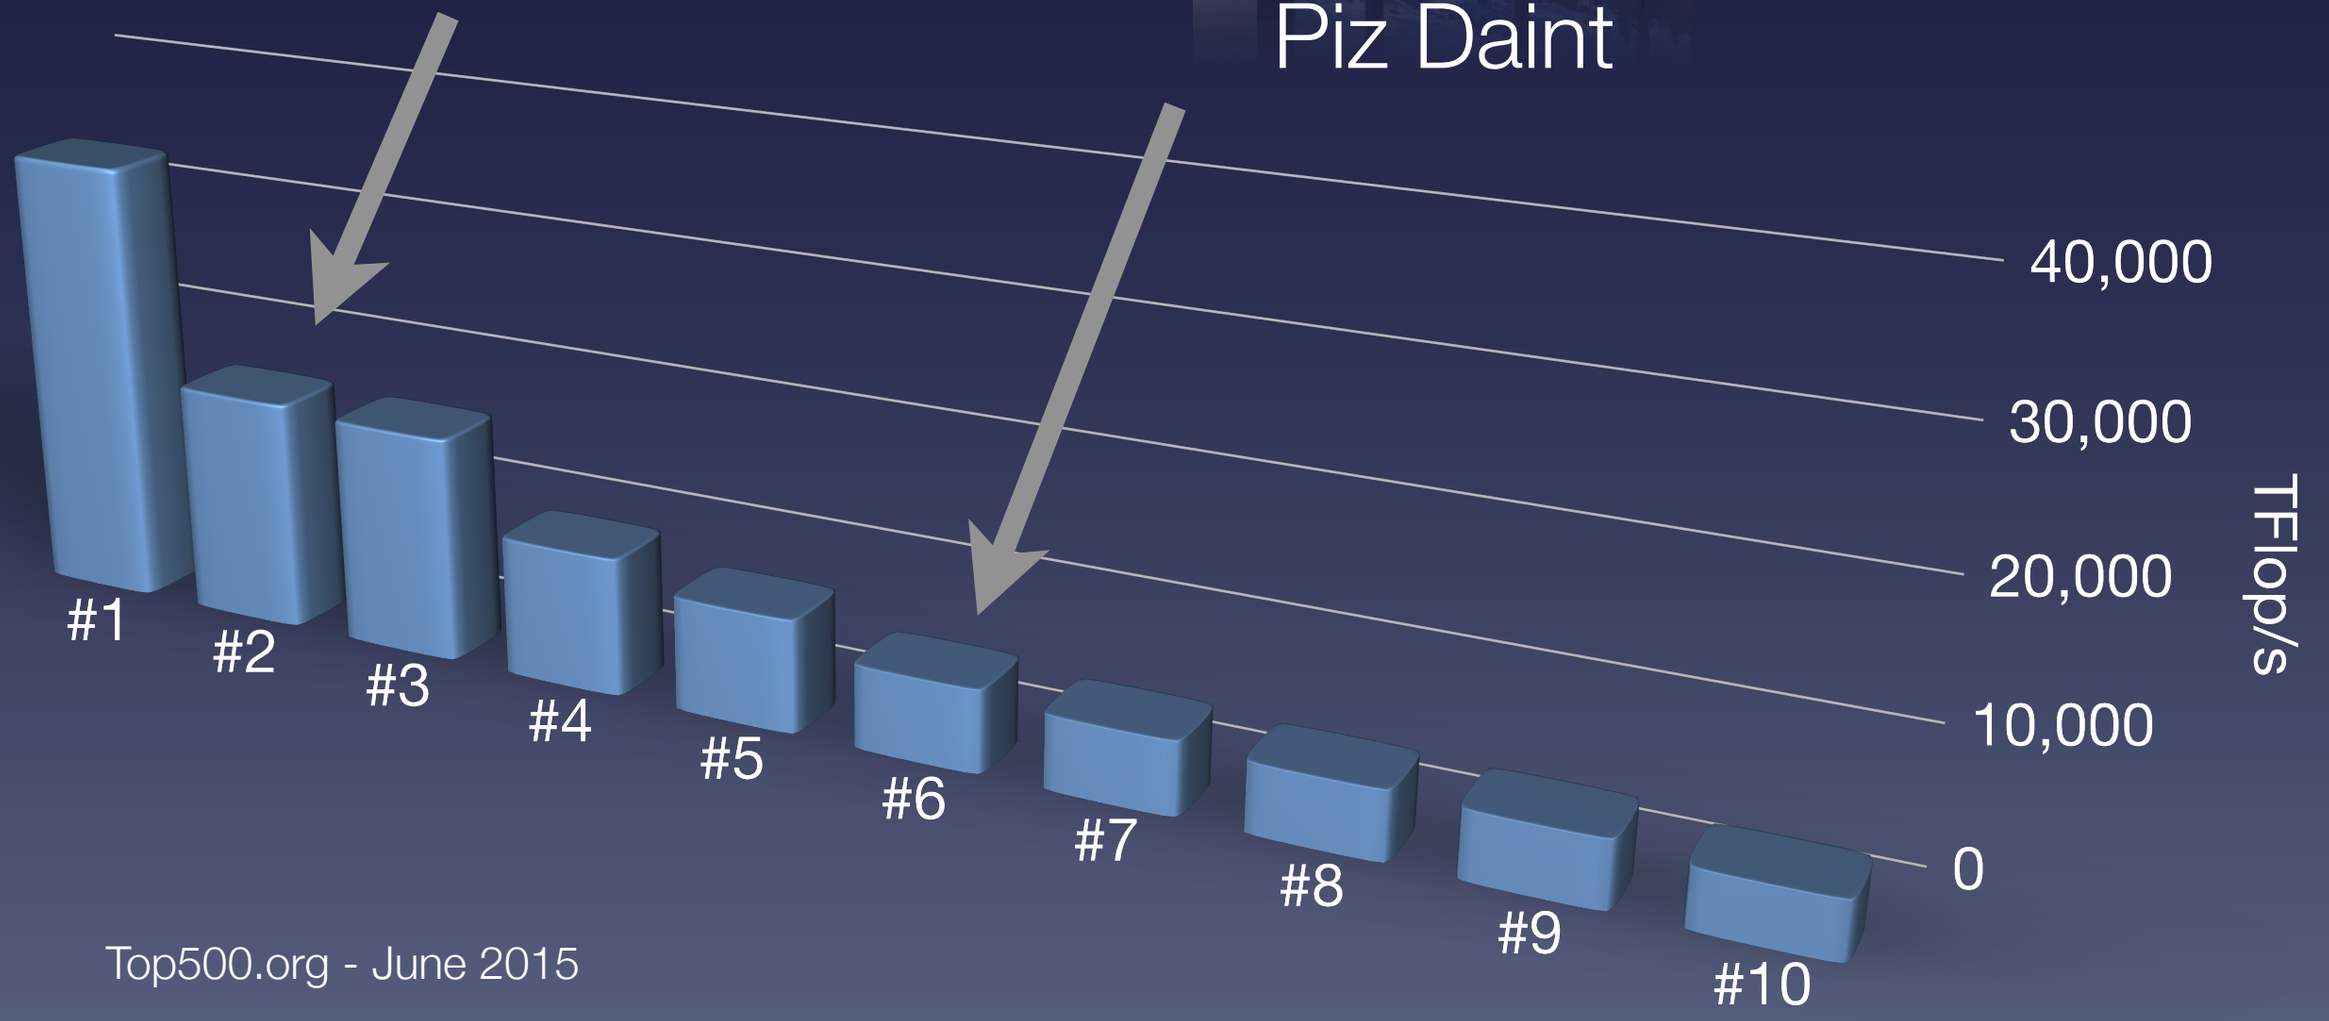

Top500.org - June 2015

GPU C

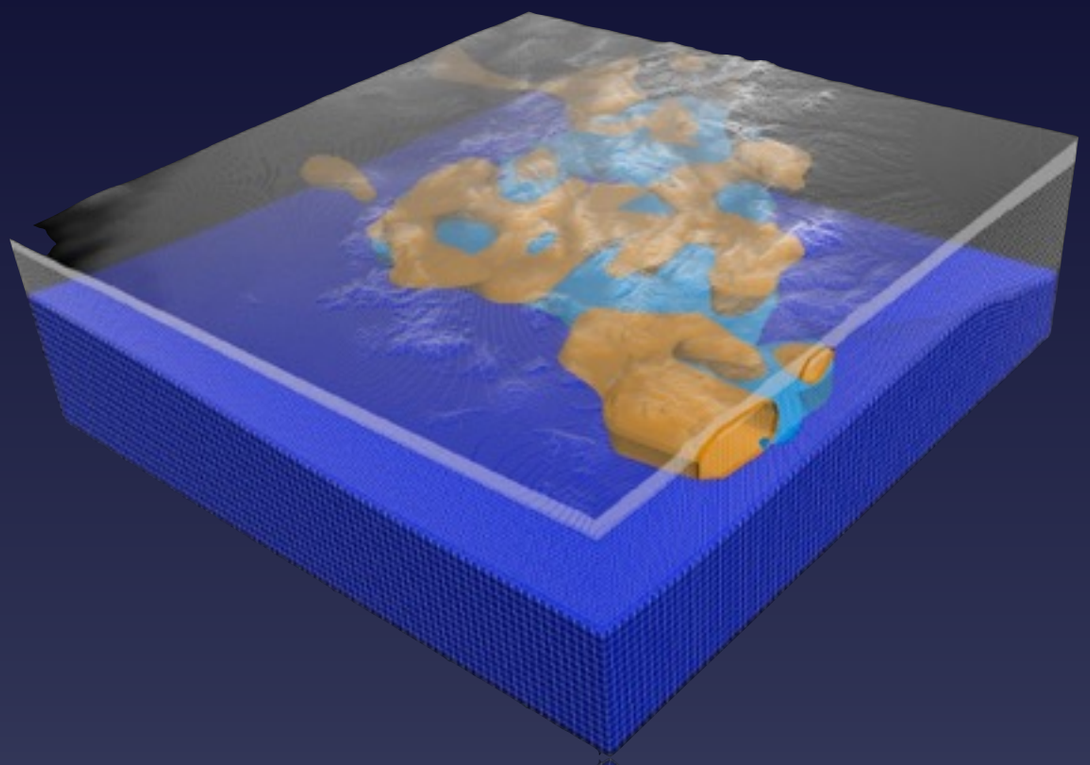

SPECFEM3D
~70,000 lines of code
Fortran90 / C / Cuda

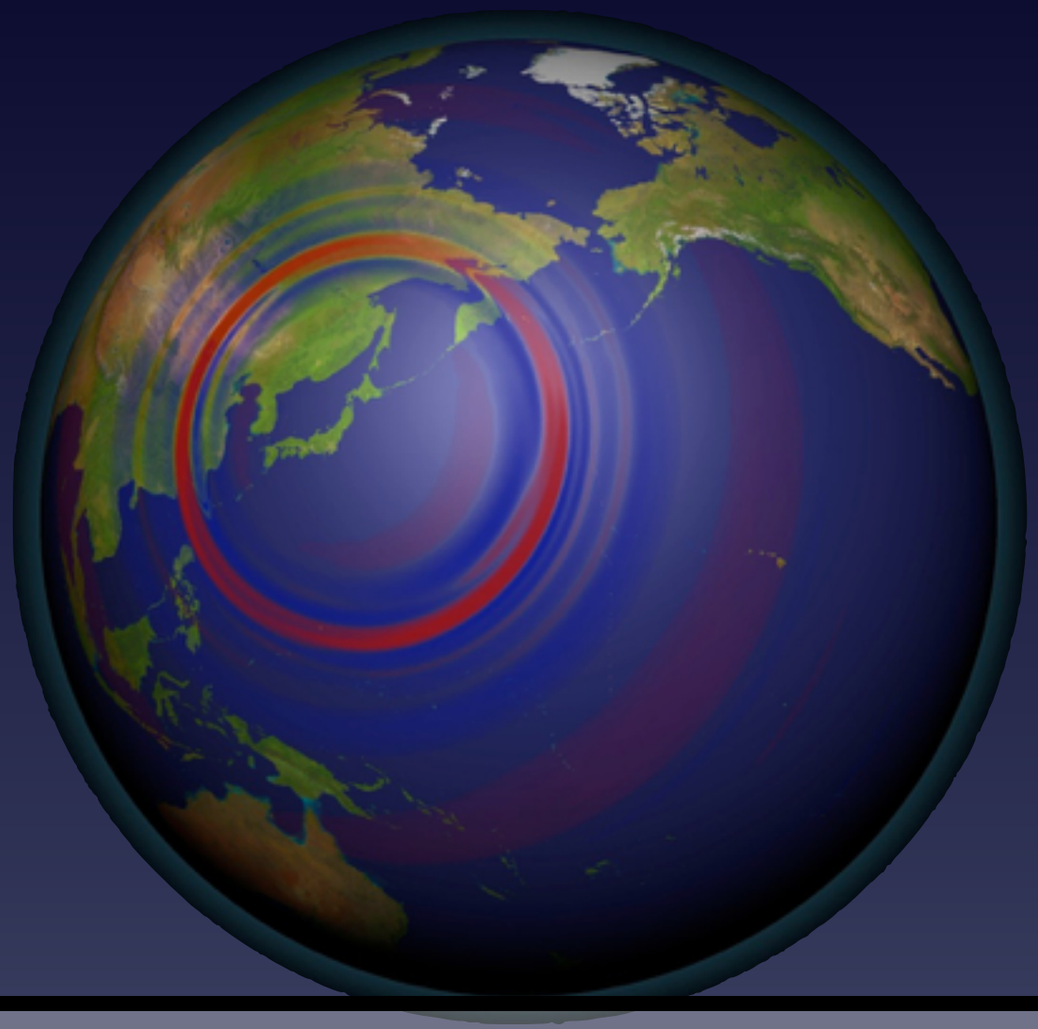

SPECFEM3D_GLOBE
~100,000 lines of code
Fortran90 / C / Cuda / OpenCL

GPU Portability

- Initial implementation: CUDA
  - In collaboration with NVIDIA (Peter Messmer & Cyril Zeller)
- Current implementation:
  - *BOAST : Bringing Optimization through Automatic Source-to-Source Transformations*
  - Kernels written in Ruby
  - Generates CUDA and OpenCL
  - Calls to kernels in C
  - Tuned for Fermi and Kepler architectures

RELATED LINKS

What is GPU Computing
GPU-Accelerated Applications
GPU Applications Catalog
Tesla GPU Accelerators
Where to Buy Tesla Accelerators

BENCHMARK REPORTS

Tesla K20 Benchmark Report
Computational Chemistry Benchmarks

Test Drive the World's Fastest GPU Accelerators

Your GPU-Accelerated Code Just Got Even Faster

Accelerate your code with NVIDIA® Kepler™ GPUs, the world's fastest accelerators. The Tesla® K40 GPU is our latest and fastest accelerator. With powerful features like GPU Boost and 12 GB of memory, Tesla K40 delivers up to 40% more performance compared to the Tesla K20X.

ACCELERATION ACROSS DOMAINS WITH GPUs

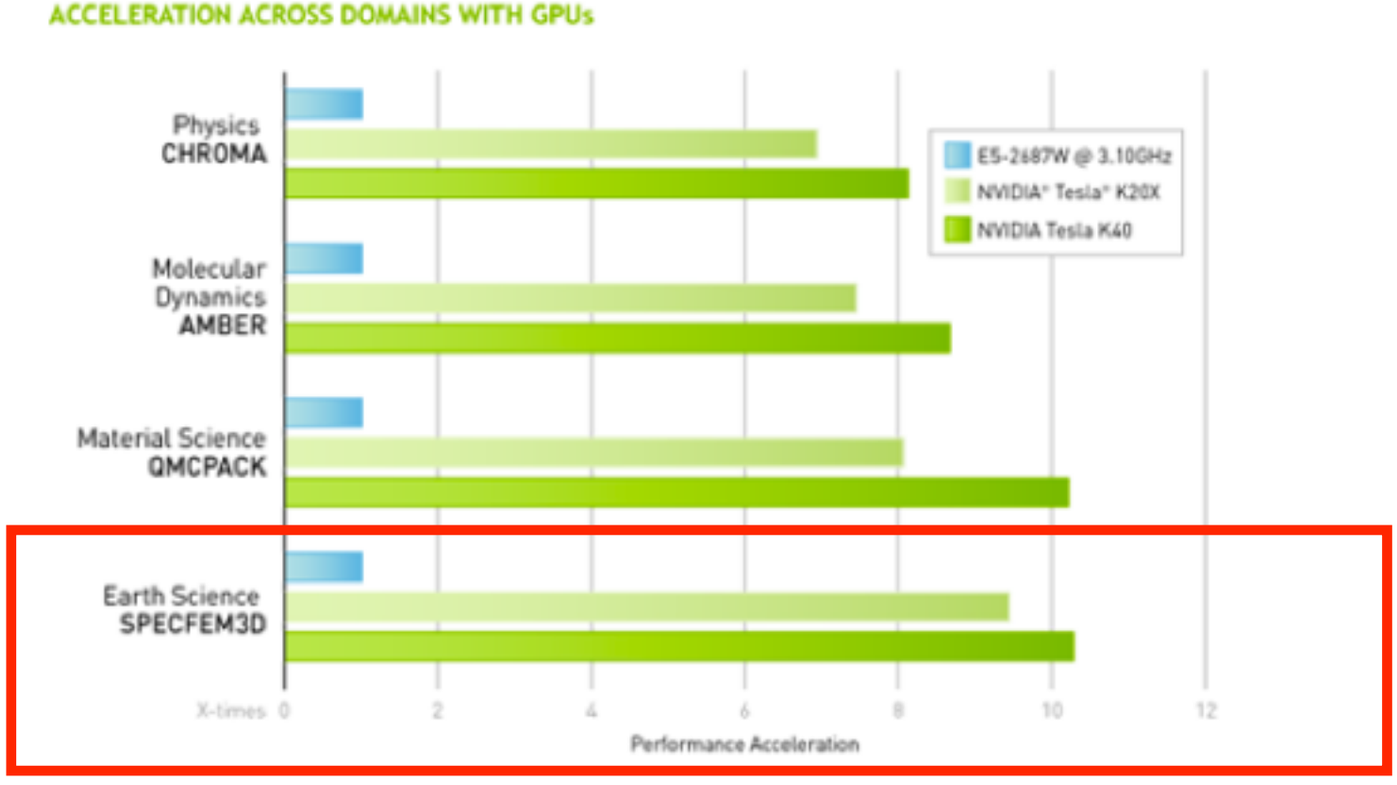

Daniel Peter

Performance

Strong Scaling

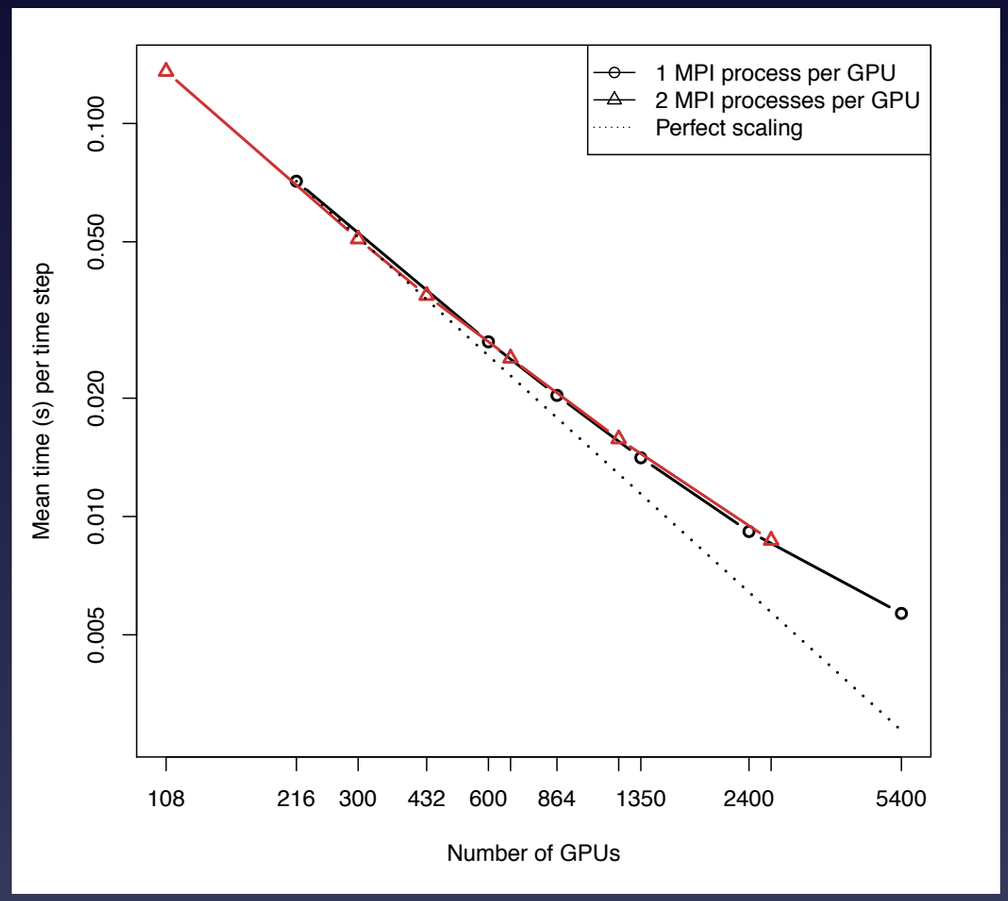

Weak Scaling

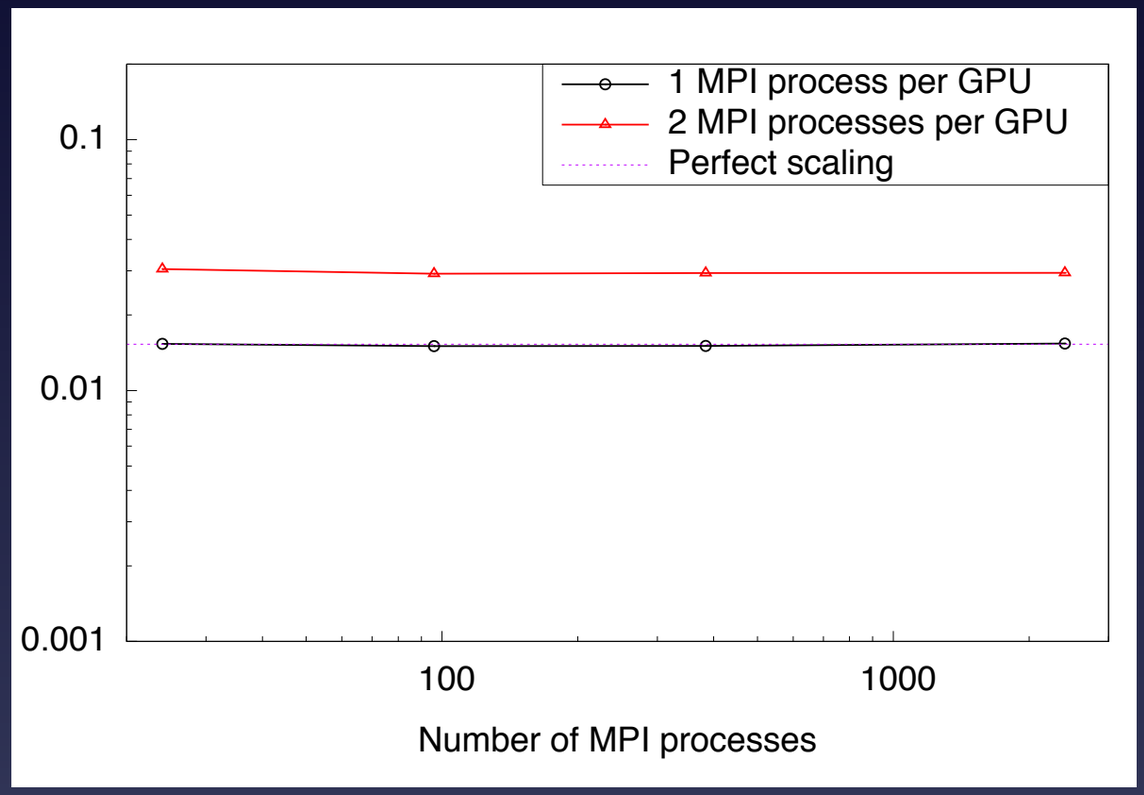

February 05, 2013

Four Applications Sustain One Petaflop on Blue Waters

HPC wire

July 18, 2012

Researchers Squeeze GPU Performance from 11 Big Science Apps

SPECFEM3D_GLOBE - Milestones

2015 **K computer simulation > 1.24 PFlops** (shortest period ~ 1.2 s)
on 82,134 nodes, 82,134 MPI ranks w/ 8 OpenMP threads, 657,072 cores

2013 **Blue Waters XE6 simulation > 1 PFlops** (shortest period < 2 s)
on 21,675 XE nodes, 693,600 MPI ranks, 693,600 cores

2008 **Kraken XT5 simulation > 160 TFlops** (shortest period ~1.72 s)
on 149,784 cores

2003 **Earth Simulator simulation > 5 TFlops** (shortest period ~5 s)
on 243 nodes, 1,944 MPI ranks, 1,944 cores

Titan

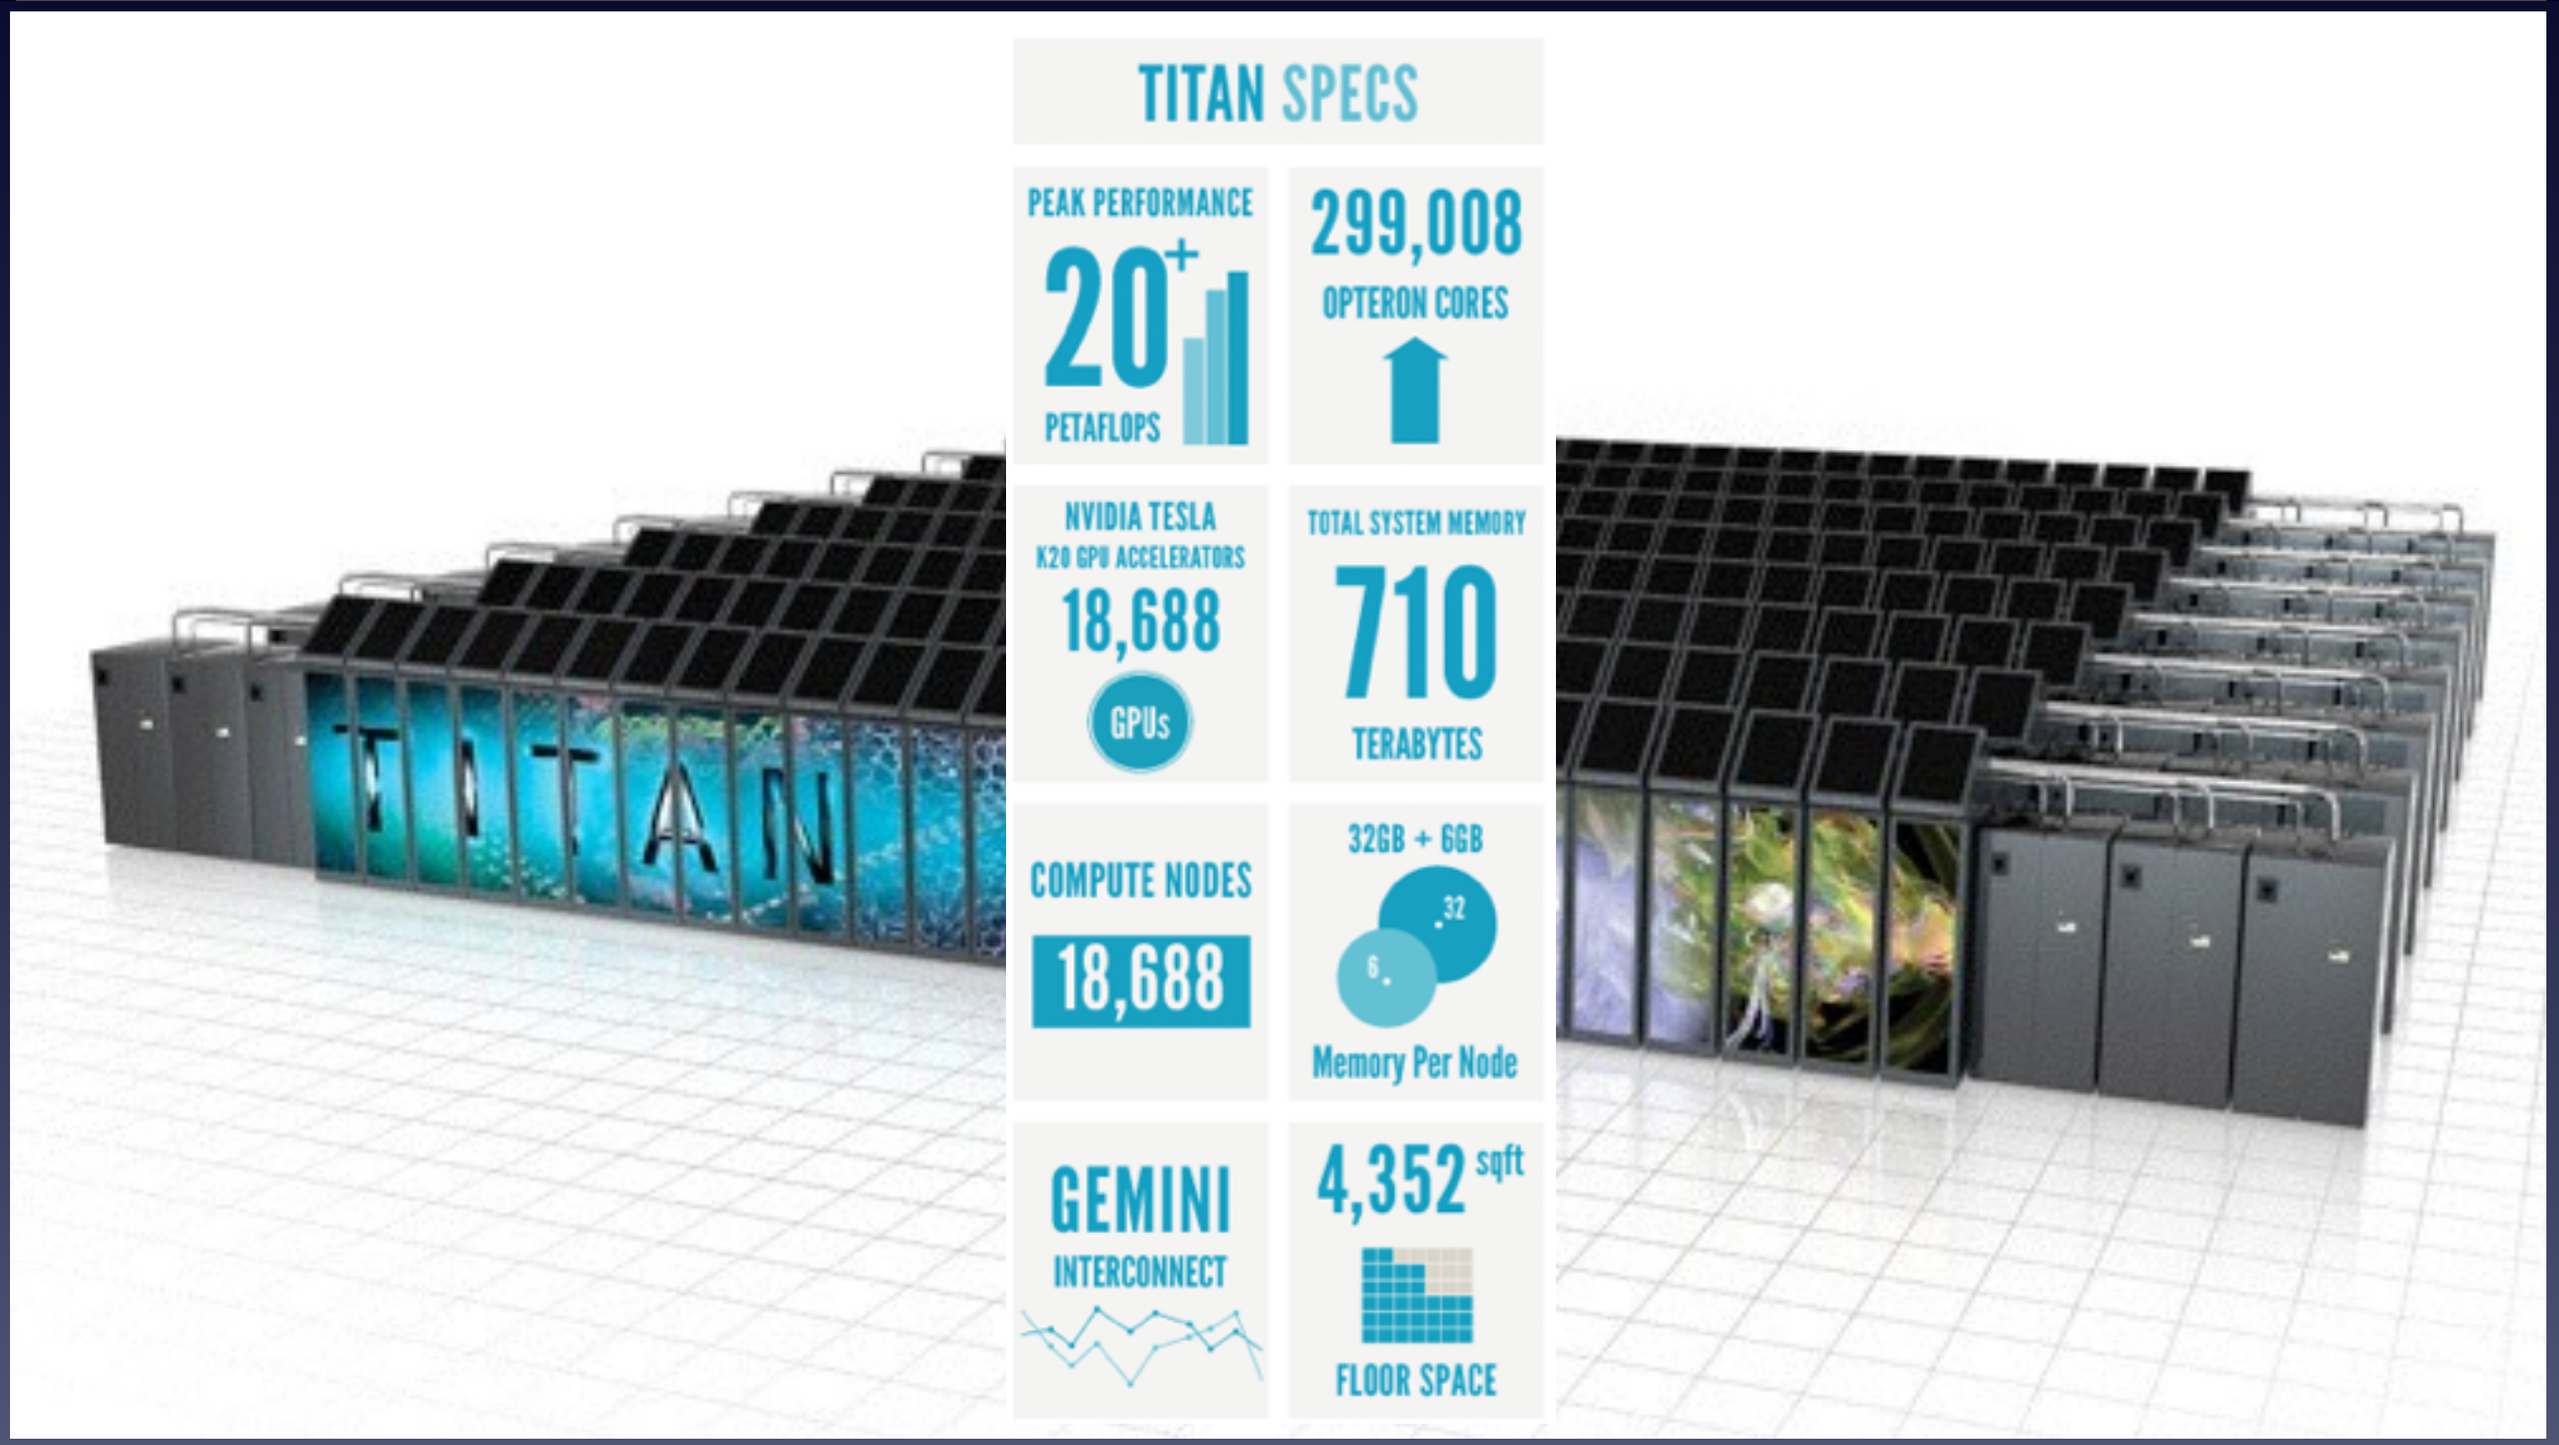

2013 - 2015 SPECFEM3D_GLOBE allocation: 250M core hours

Next Generation Machine: Summit

ORNL's Center for Accelerated Application Readiness (CAAR)

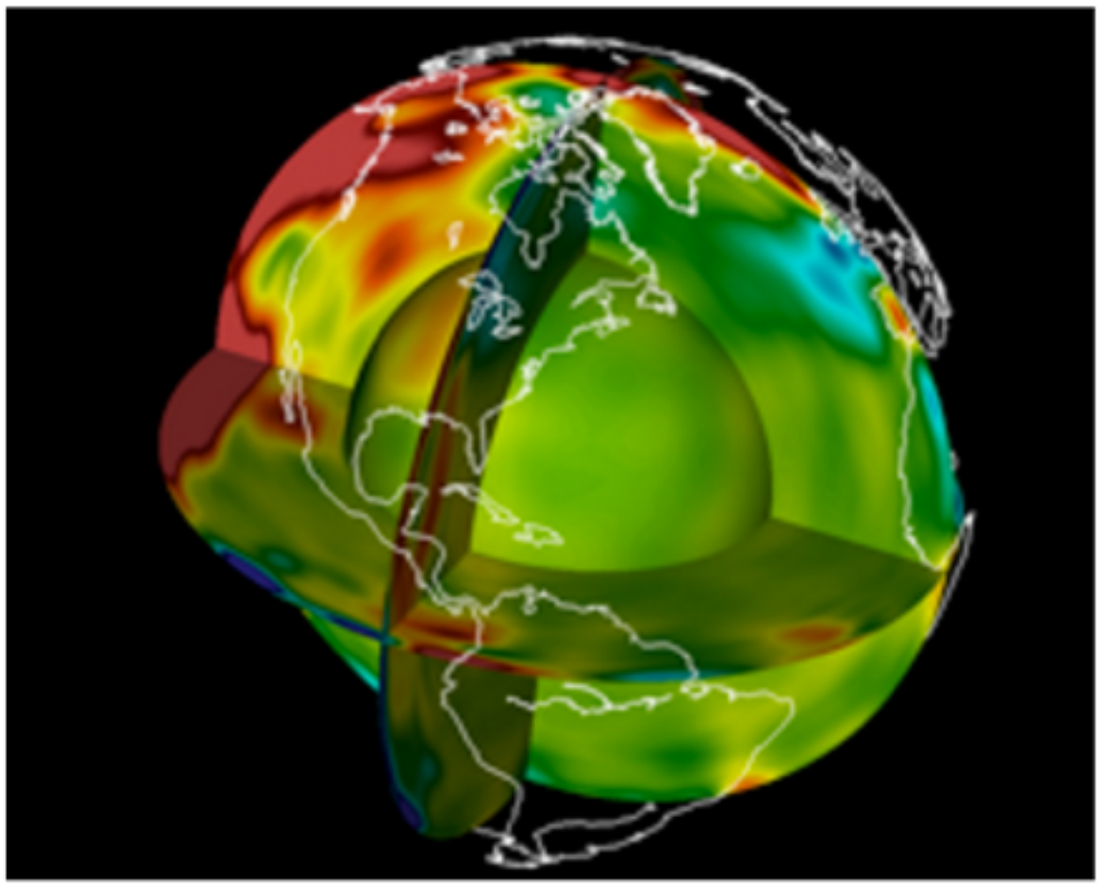

Code: SPECFEM

Science Domain: Seismology

Title: Mapping the Earth's Interior Using Big Data

PI: Jeroen Tromp, Princeton University

OLCF partnership with IBM, NVIDIA & Mellanox

Exascale Challenges

Taming Workflow Issues

Seismic Tomography Workflow

1. Current data formats are inadequate for fast, parallel I/O; ASDF: an Adaptable Seismic Data Format

N selected earthquakes

Request observed data

Extract SEED files

Convert to ASDF

N observed ASDF data

Run mesher: I ADIOS mesh file

Run N forward simulations

N synthetic ASDF data

Process data, select windows, make measurements & compute adjoint sources

Pre-processing (embarrassingly parallel)

N ASDF adjoint source files

Run N adjoint simulations

N ADIOS kernel files

Sum kernels: I ADIOS gradient file

Pre-condition & smooth the gradient

Determine step length

Update model: I ADIOS model file

Convergence?

Finish

No

Iterate

L-BFGS
Quasi-Newton

3. Workflow stabilization: Kepler, Swift & Pegasus

Post-processing (parallel)

2. Storage & visualization of Earth models: ADIOS with VisIt

Seismic Imaging & Inversion Challenges

- Cheap, abundant sensors
- Massive amounts of data
  - Industry data sets
  - Regional & global seismology data sets
  - Cross-correlation data sets for seismic interferometry
- On HPC systems, I/O is the bottleneck
- Adopt new data formats for fast parallel I/O (e.g., NetCDF, HDF5 & ADIOS)
- Data culling tools to reduce preprocessing time
- A standard for the exchange of Earth models
- Adopt workflow management tools (e.g., Kepler, Pegasus & Swift)
- Tools for data mining, feature extraction, visualization & virtualization (e.g., ParaView, VisIt)

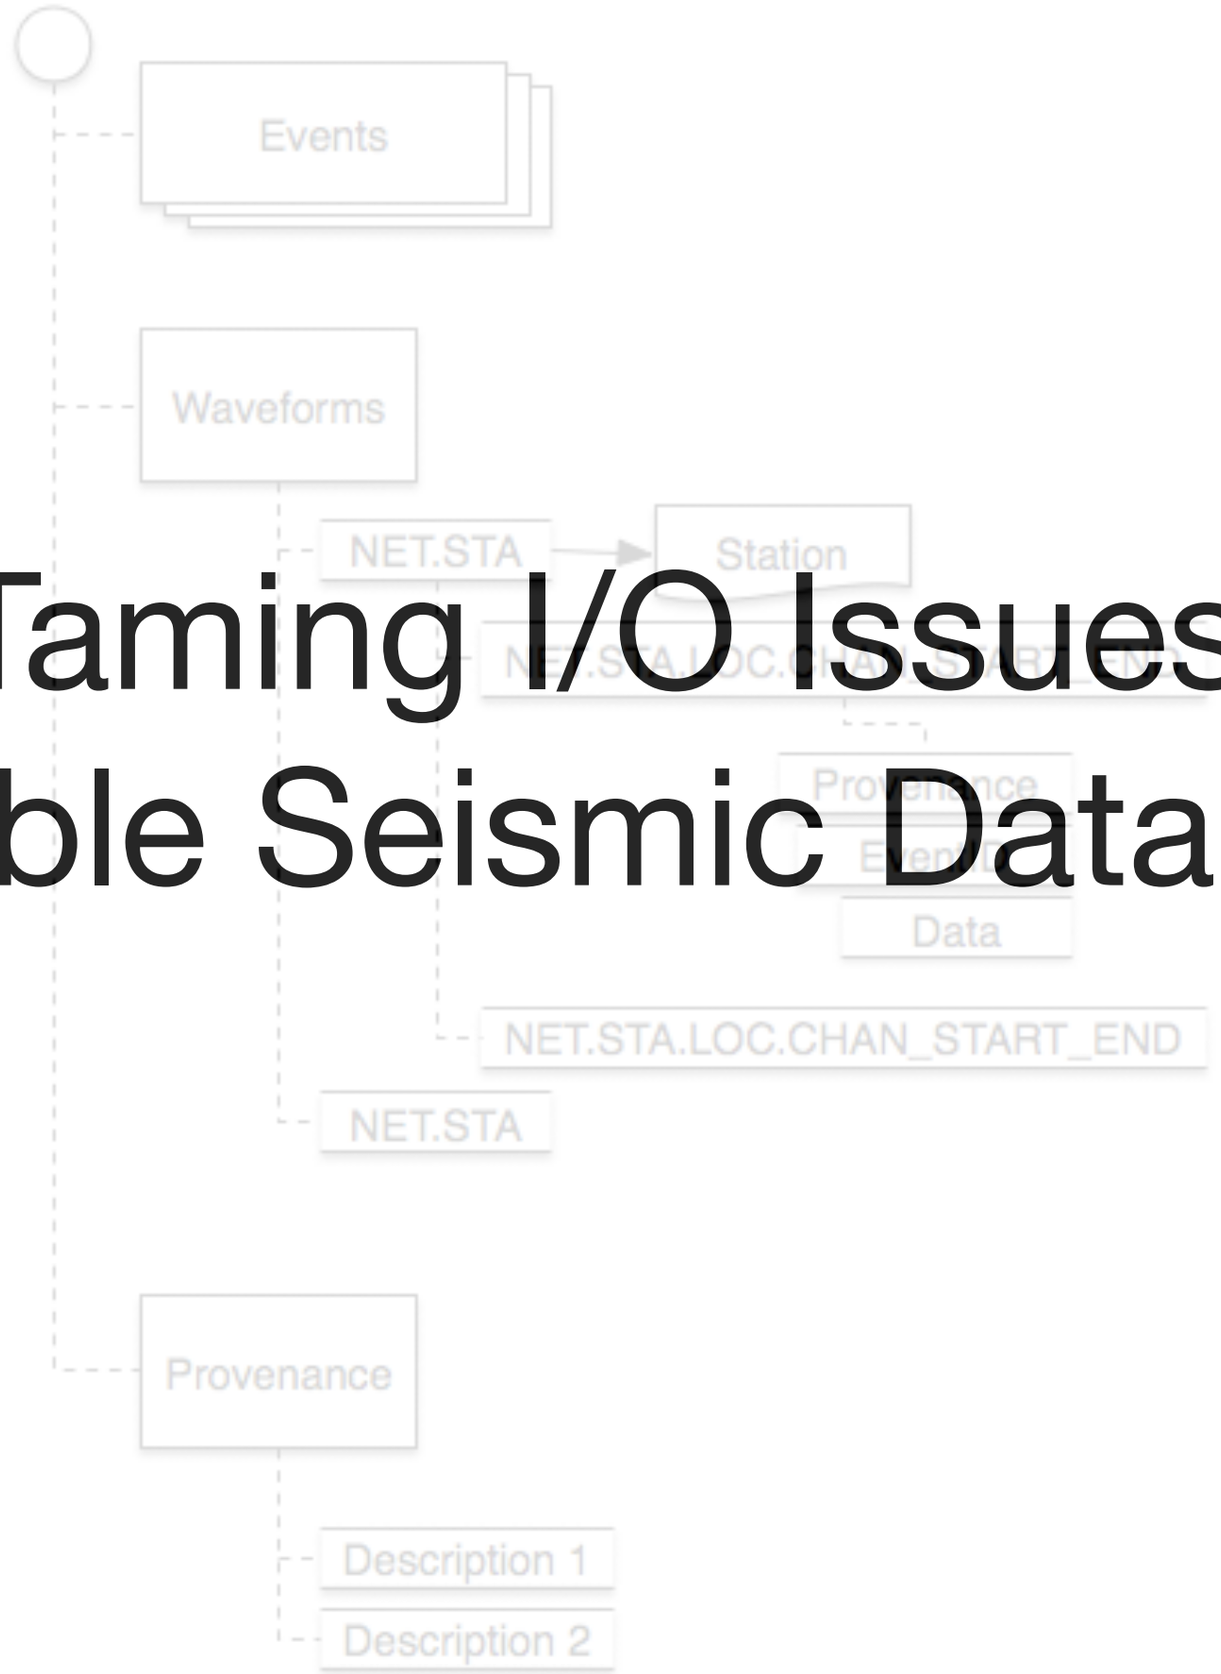

Taming I/O Issues: Adaptable Seismic Data Format

Data in Regional & Global Seismology

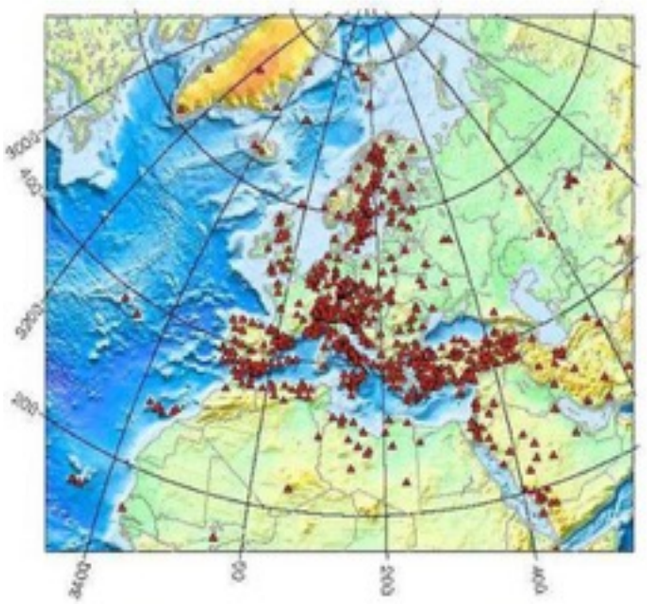

[www.geo.uib.no]

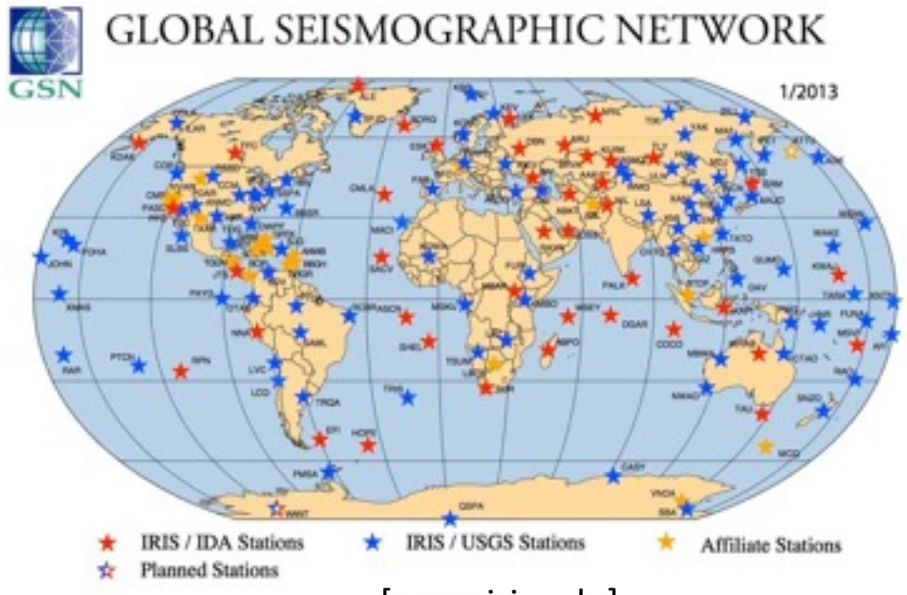

[www.iris.edu]

[web.mst.edu]

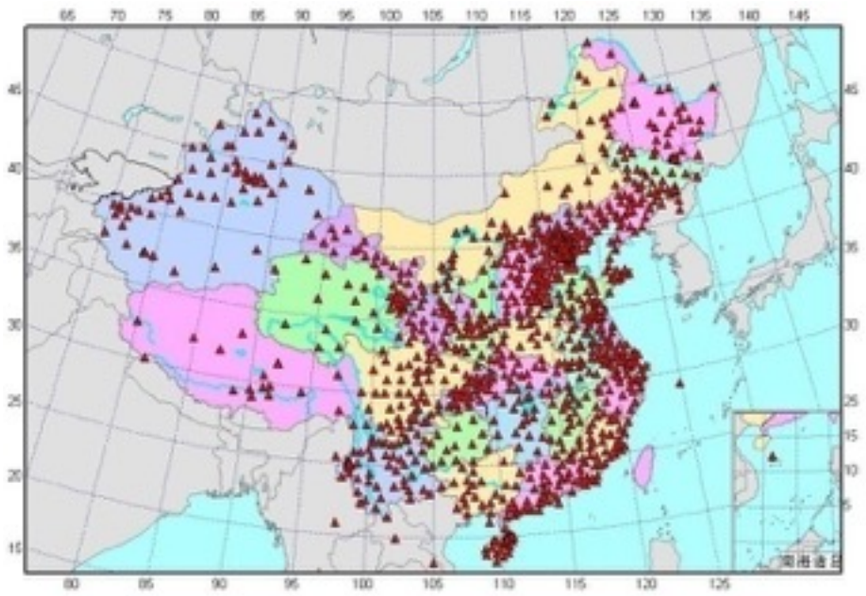

[data.earthquake.cn]

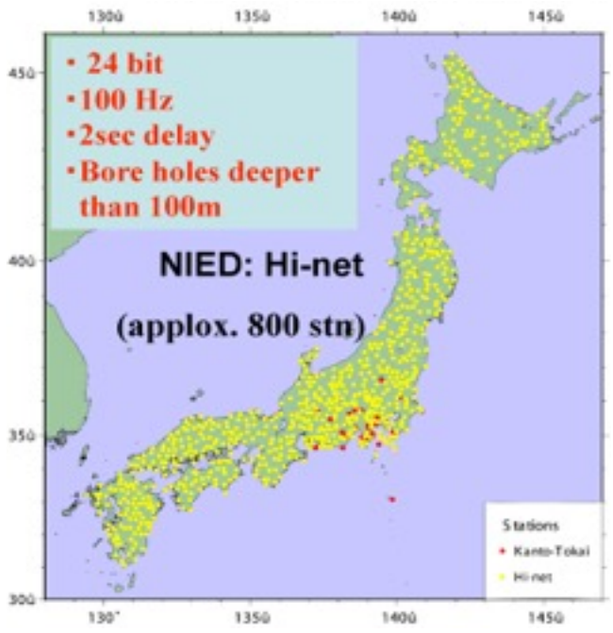

[drh.edm.bosai.go.jp]

MERMAID/MariScope

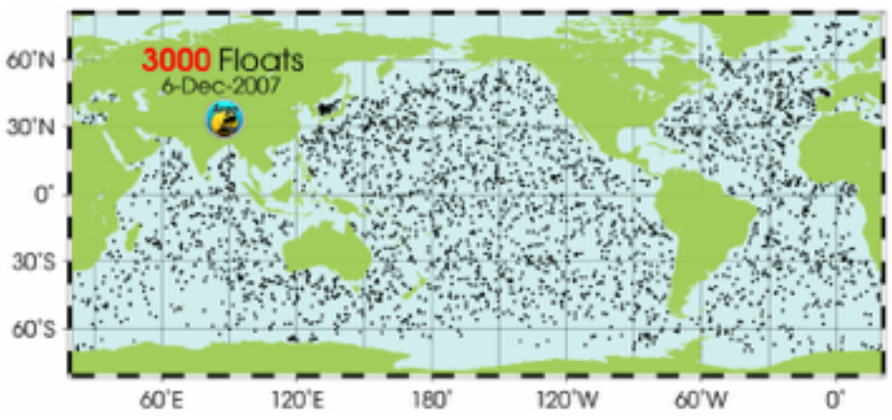

[Simons et al, 2006]

Data in Exploration Seismology

3D marine survey can involve 5,000 shots and 50,000 recorders

- Petabytes of data
- SEG-Y is the current standard
- Variable SEG-Y file structure
- SEG-Y programs do not always follow specifications

ASDF: Adaptable Seismic Data Format

- Collaboration involving Princeton University, Munich University (ObsPy) and Oak Ridge National Laboratory
- Increase I/O performance by combining all the time series for a single shot or earthquake into one file
- Take advantage of parallel processing
- Use modern file format as container (HDF5)
- Store provenance inside the file for reproducibility
- Use existing standards when possible (e.g., XML)
- Open wiki for development

Development Team

Michael Afanasiev, Jean-Paul (Pablo) Ampuero, Kangchen Bai, Piero Basini, Céline Blitz, Alexis Bottero, Ebru Bozdag, Emanuele Casarotti, Joseph Charles, Min Chen, Paul Cristini, Clément Durochat, Percy Galvez, Dominik Göddeke, Vala Hjörleifsdóttir, Sue Kientz, Dimitri Komatitsch, Jesús Labarta, Nicolas Le Goff, Pieyre Le Loher, Matthieu Lefebvre, Qinya Liu, David Luet, Yang Luo, Alessia Maggi, Federica Magnoni, Roland Martin, René Matzen, Dennis McRitchie, Matthias Meschede, Peter Messmer, David Michéa, Vadim Monteiller, Surendra Nadh Somala, Tarje Nissen-Meyer, Daniel Peter, Kevin Pouget, Max Rietmann, Elliott Sales de Andrade, Brian Savage, Bernhard Schuberth, Anne Sieminski, James Smith, Leif Strand, Carl Tape, Jeroen Tromp, Brice Videau, Jean-Pierre Vilotte, Zhinan Xie, Chang- Hua Zhang, Hejun Zhu

SPECFEM - Awards

2010 **BULL Joseph Fourier Prize winner**
for the partial GPU port of SPECFEM3D_GLOBE

2008 **ACM Gordon Bell Finalist**
for SPECFEM3D_GLOBE simulation reaching resolution of 1.72s shortest period

2003 **ACM Gordon Bell Award for Best Performance**
for SPECFEM3D_GLOBE simulation on the Earth Simulator