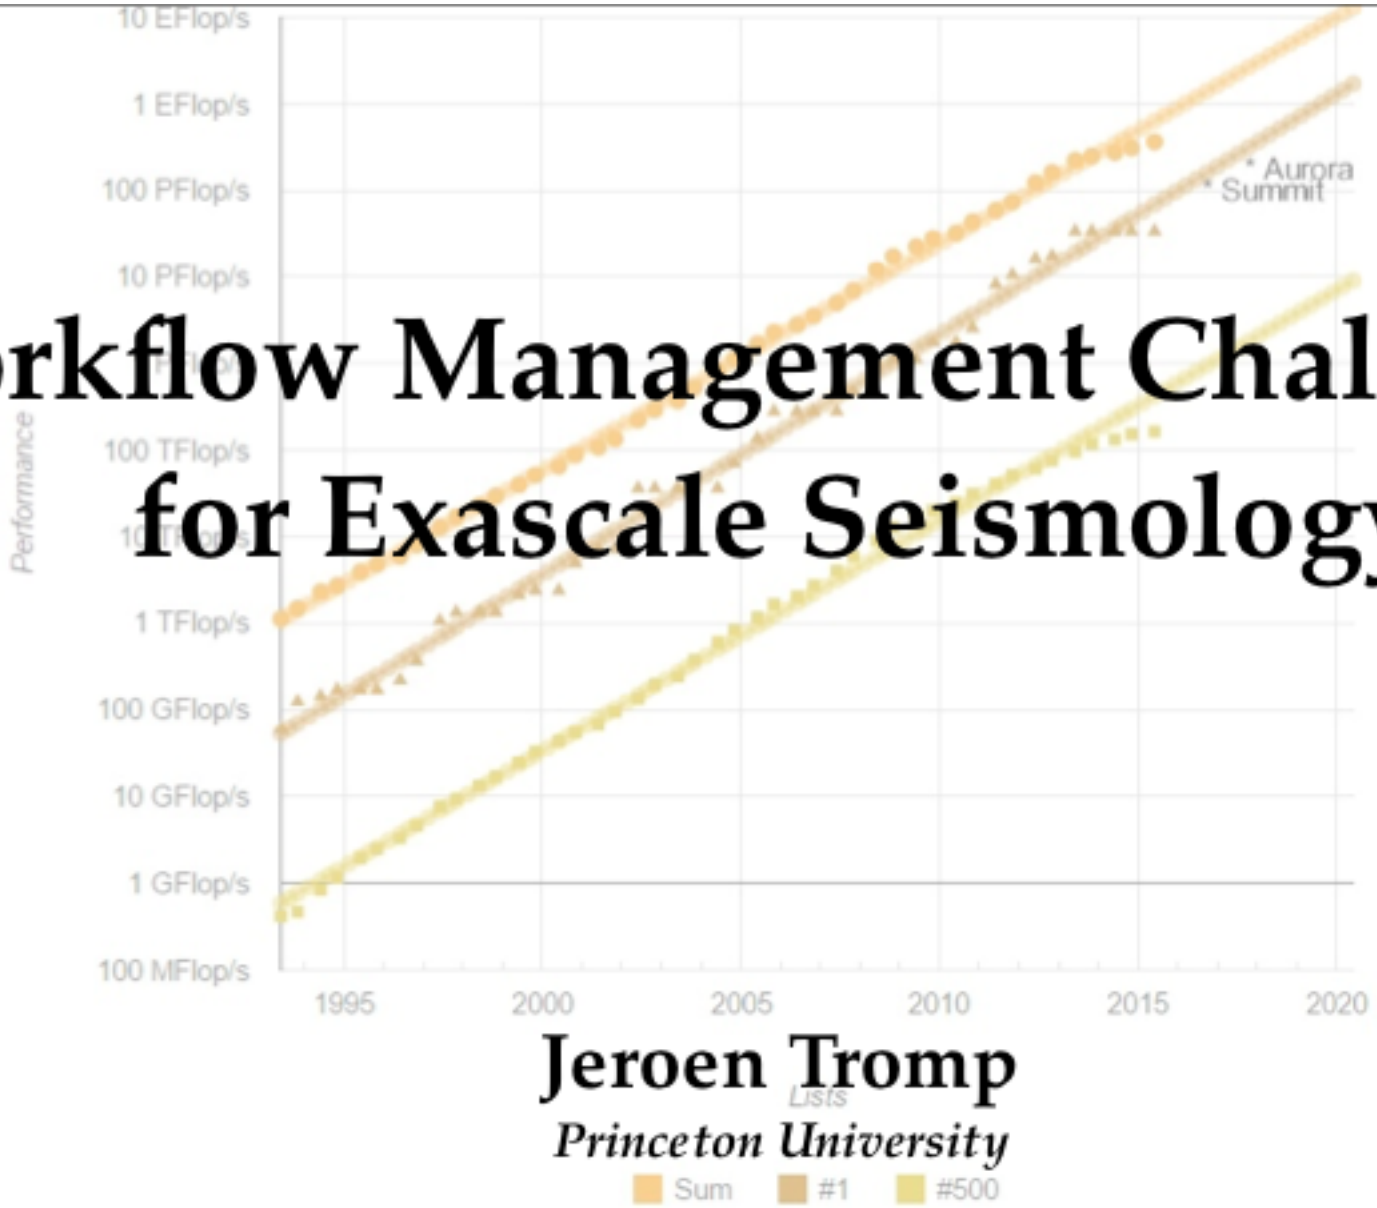

Workflow Management Challenges for Exascale Seismology

Jeroen Tromp

Princeton University

Ebru Bozdağ, Judith Hill, Shantenu Jha, Dimitri Komatitsch, Matthieu Lefebvre, Wenjie Lei, Daniel Peter, Norbert Podhorszki, David Pugmire, Youyi Ruan, James Smith & Matteo Turilli

Colorado School of Mines, University of Marseille, Princeton University, Rutgers University, KAUST & ORNL

Data Tsunami in Regional & Global Seismology

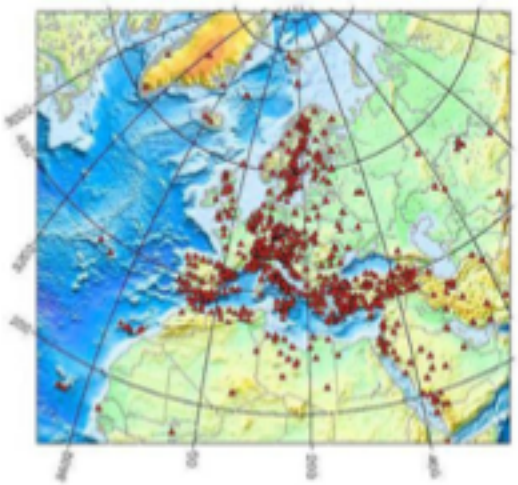

[www.geo.uib.no]

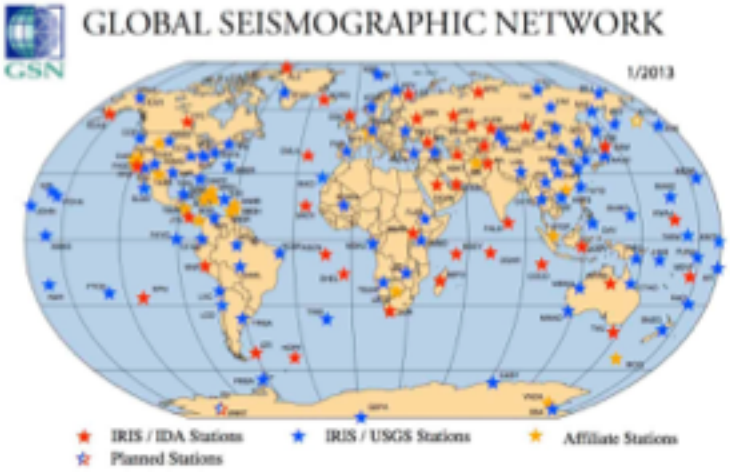

[www.iris.edu]

[web.mst.edu]

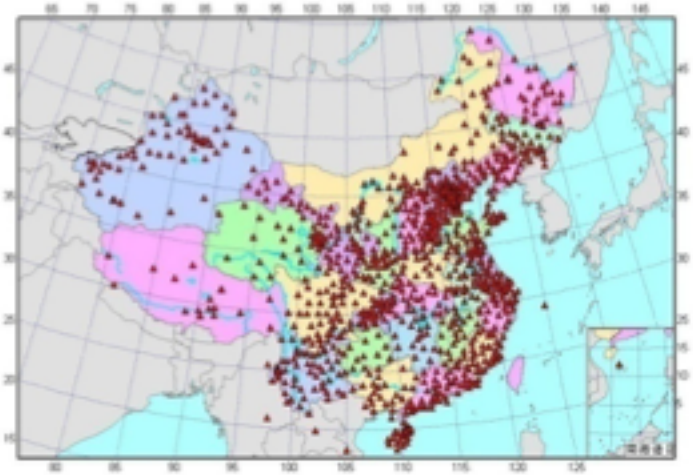

[data.earthquake.cn]

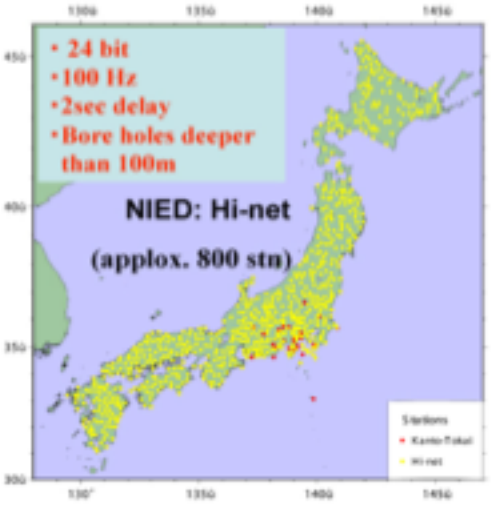

[drh.edm.bosai.go.jp]

MERMAID/MariScope

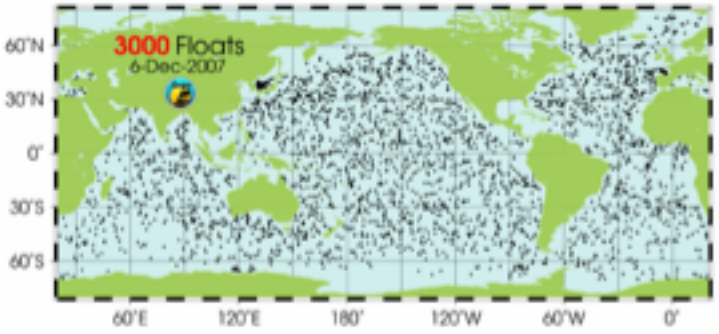

[Simons et al, 2006]

Open Source Forward & Inverse Modeling Software

Spectral-element solvers SPECFEM3D & SPECFEM3D_GLOBE

- 3D crust and mantle models
- Topography & Bathymetry
- Rotation
- Ellipticity
- Gravitation
- Anisotropy
- Attenuation
- Adjoint capabilities
- GPU accelerated

Computational Infrastructure for Geodynamics (CIG)
www.geodynamics.org

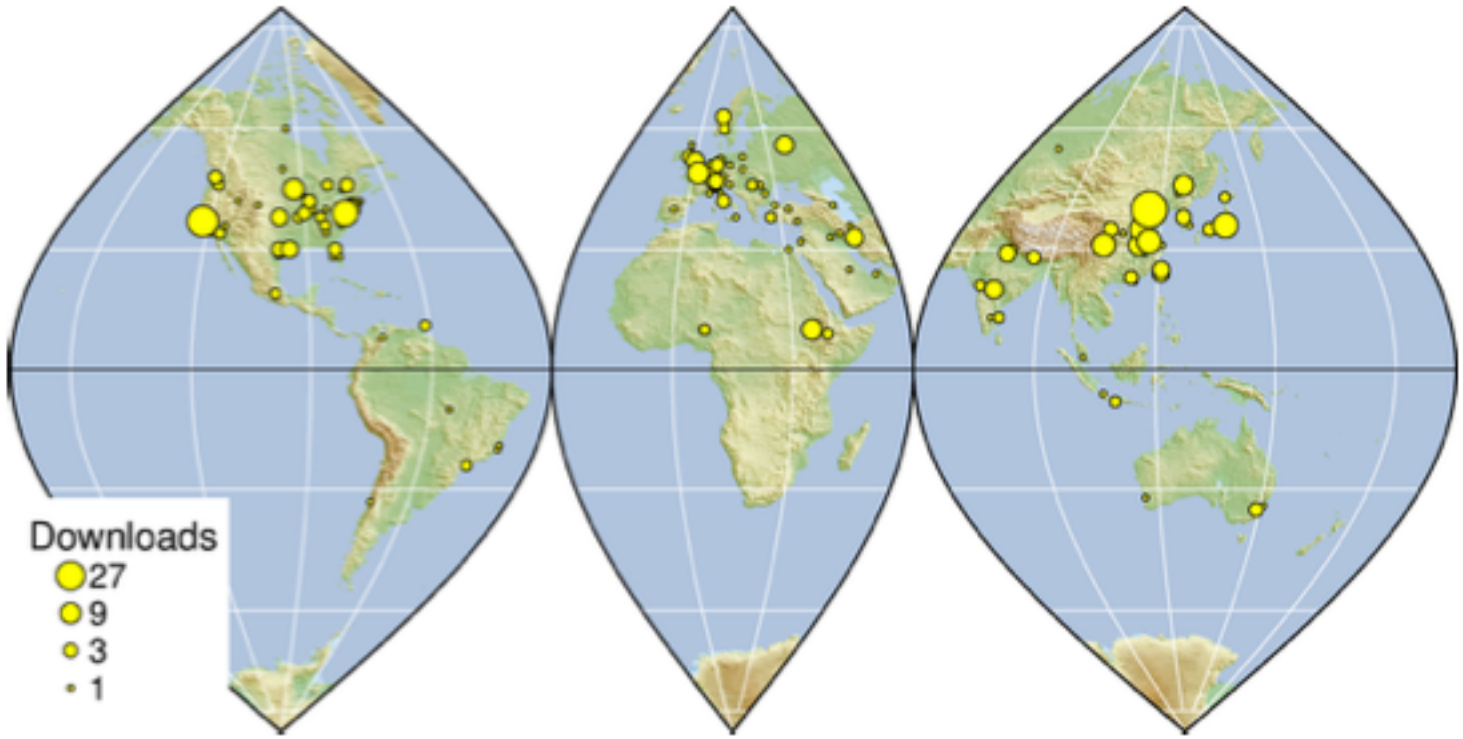

Exascale Goal: Use Complete Seismograms

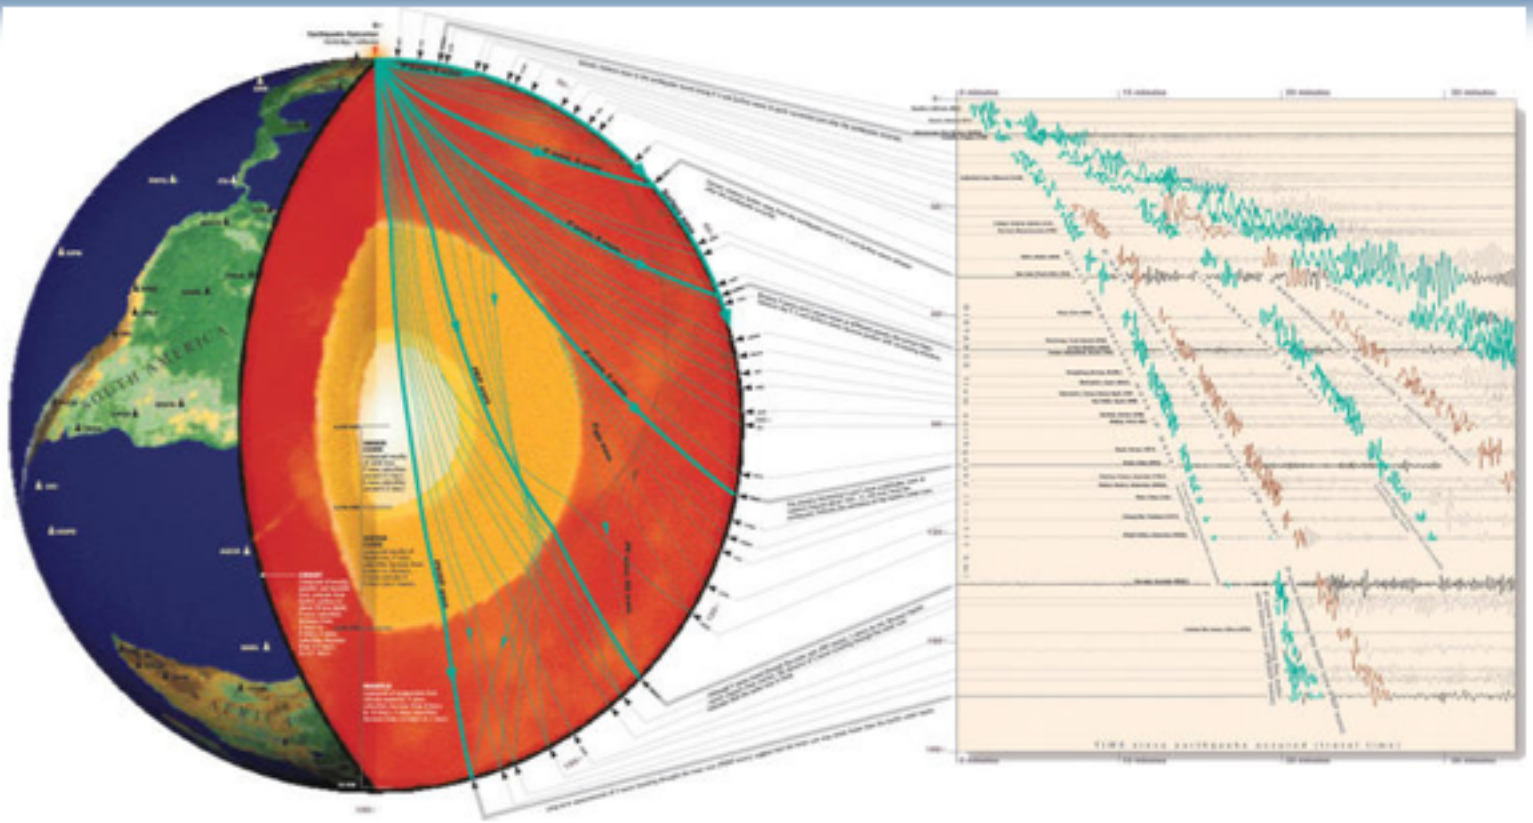

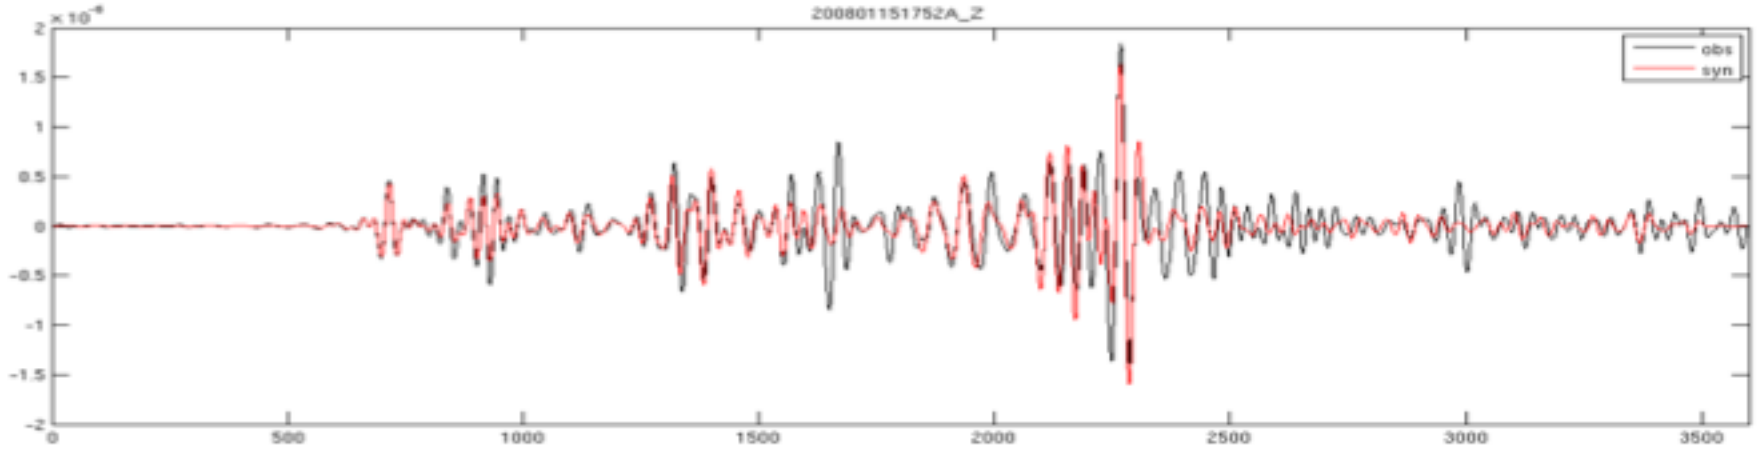

We Really Use "Full Waveforms"

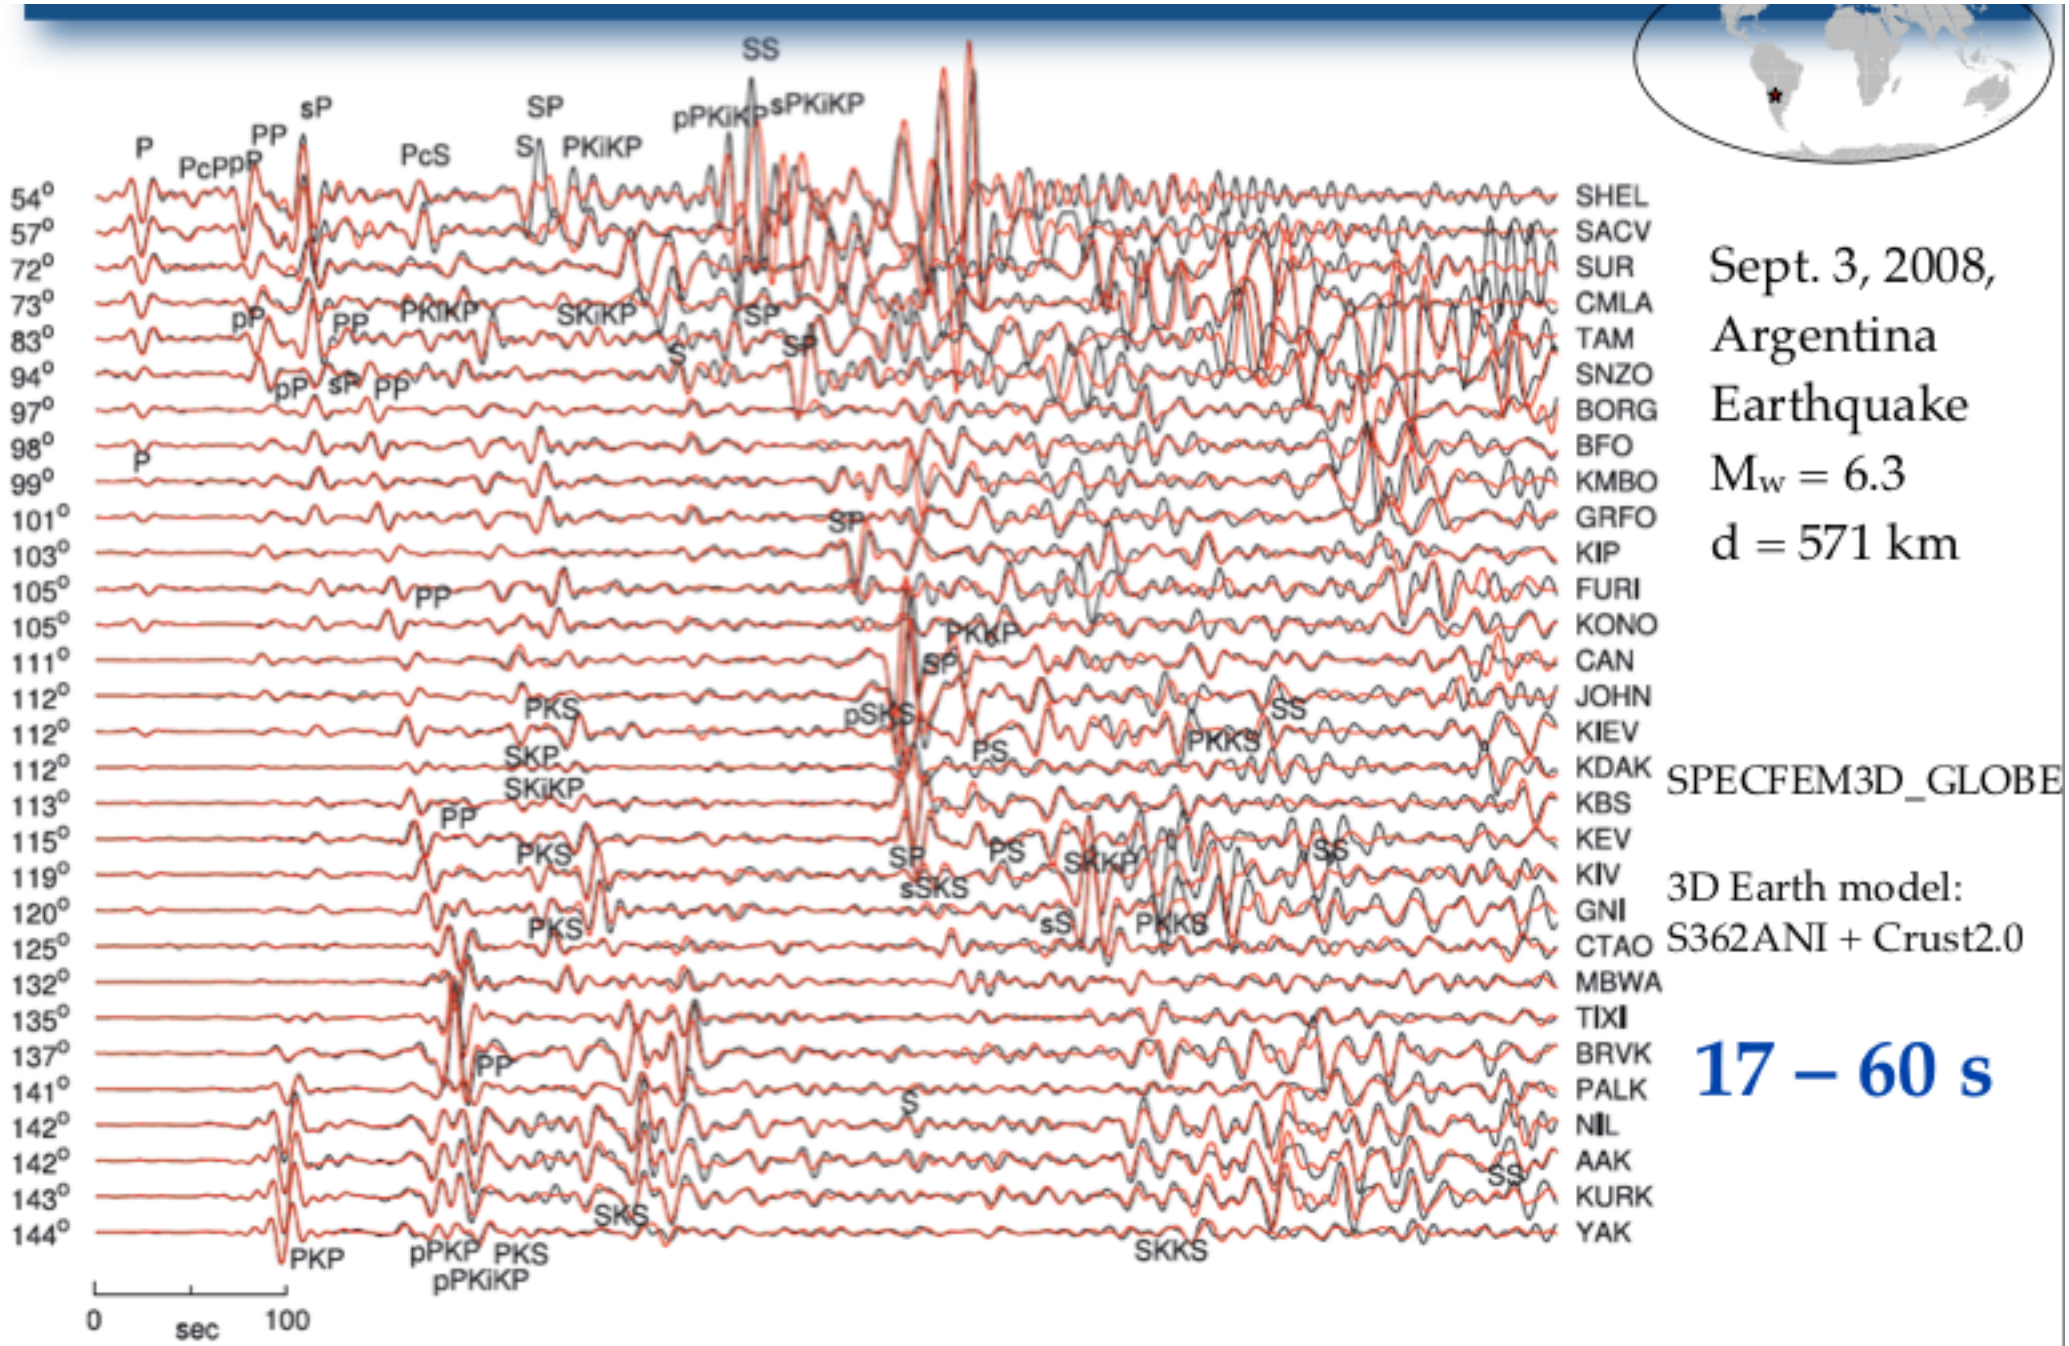

Exascale Goal: Use All Available Data

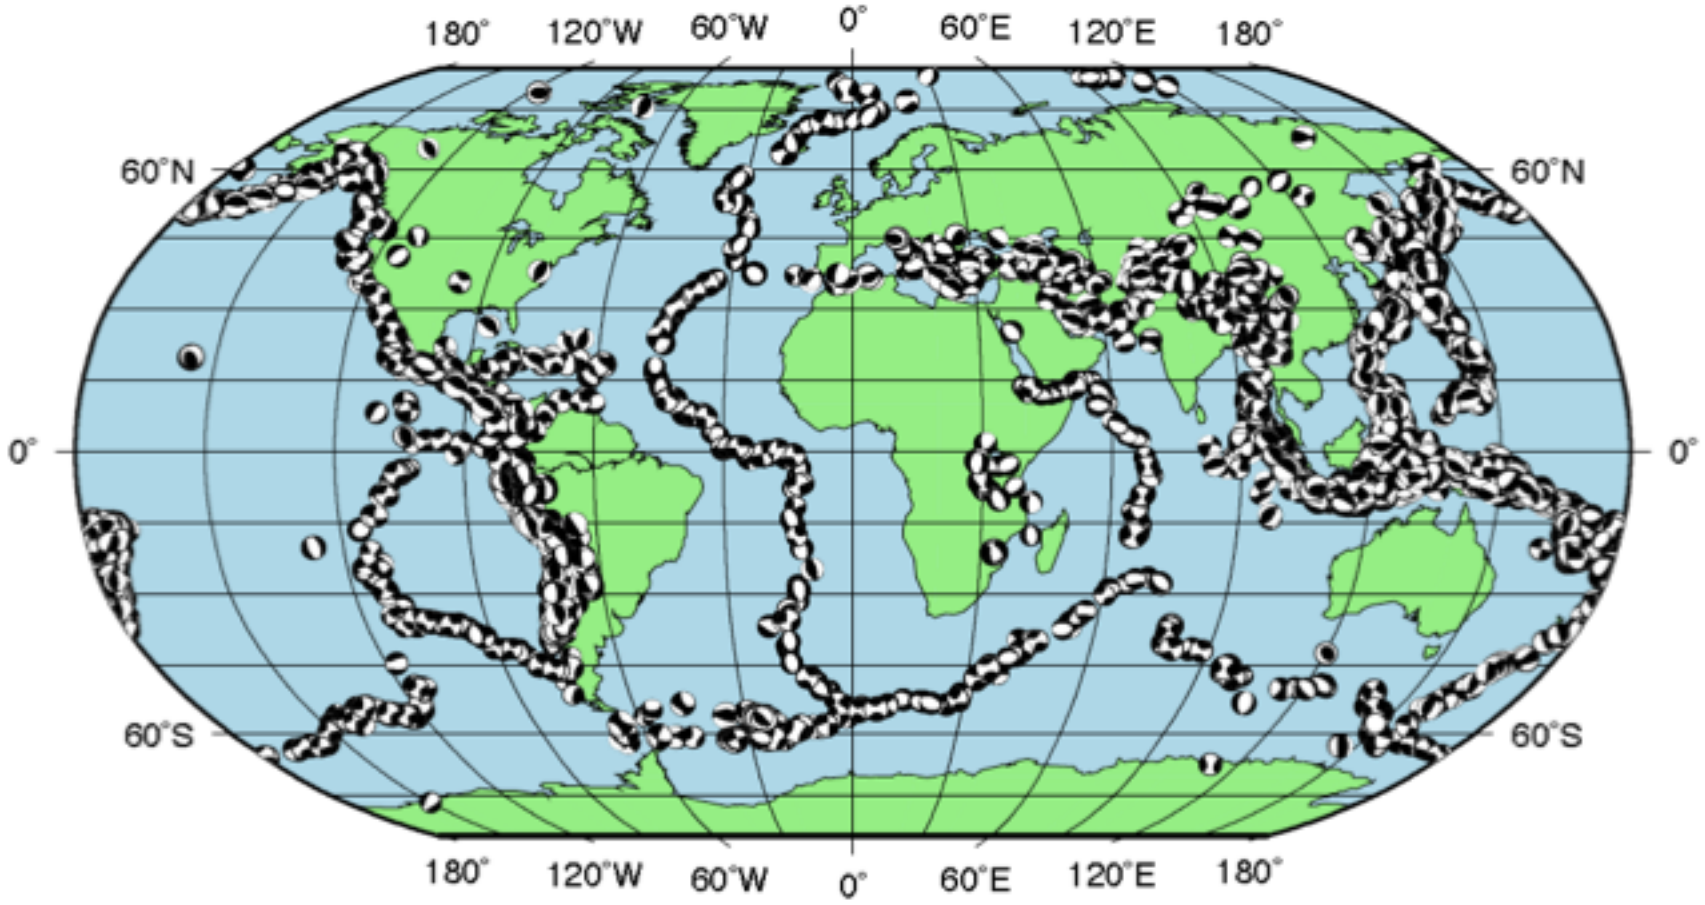

More than 6,000 earthquakes ($5.5 \leq Mw \leq 7.0$) since 1999

Exascale Goal: Higher-Frequency Body Waves

Short term goal (2018): 9 s ("Summit")
Long term goal (2021): 1 s (exascale)

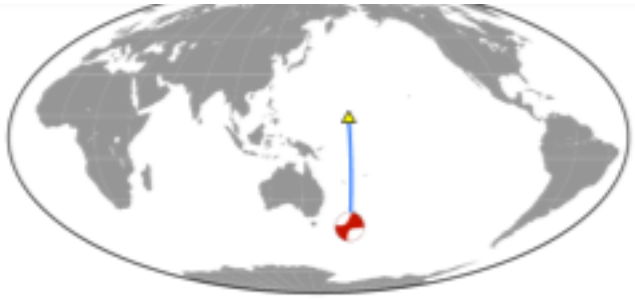

Station: KWAJ
Δ= 52°

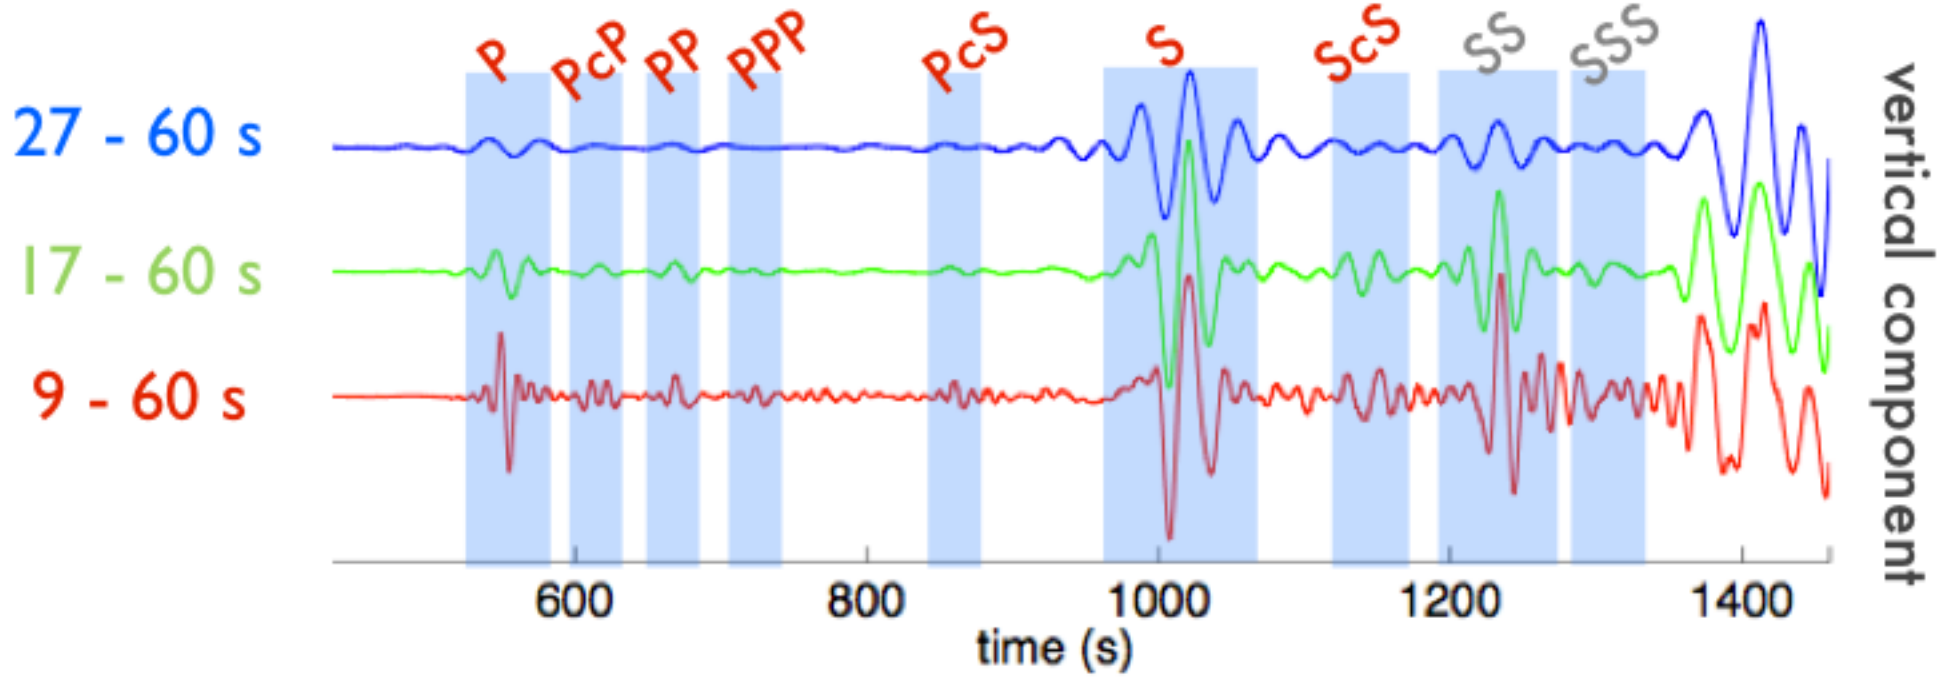

Ebru Bozdağ

Exascale Seismology Science Goals

Global Seismology:

- Use data with a shortest period of ~1 Hz
- Use all available events with magnitudes greater than ~ 5.5
- Use entire 200 minutes long, three-component seismograms
- Workflow stabilization & management
- Allow for transverse isotropy with a random symmetry axis
- Allow for variations in attenuation
- Facilitate uncertainty quantification
- Opportunities for ML/DL/AI in data selection & assimilation
- Data mining, feature extraction & visualization

Exascale Seismology Science Goals

Exploration Seismology:

- Anelastic full waveform inversion with 10-100 Hz three-component data
- Assimilation of Petabytes of data from tens of thousands (I/O!)
- Uncertainty quantification

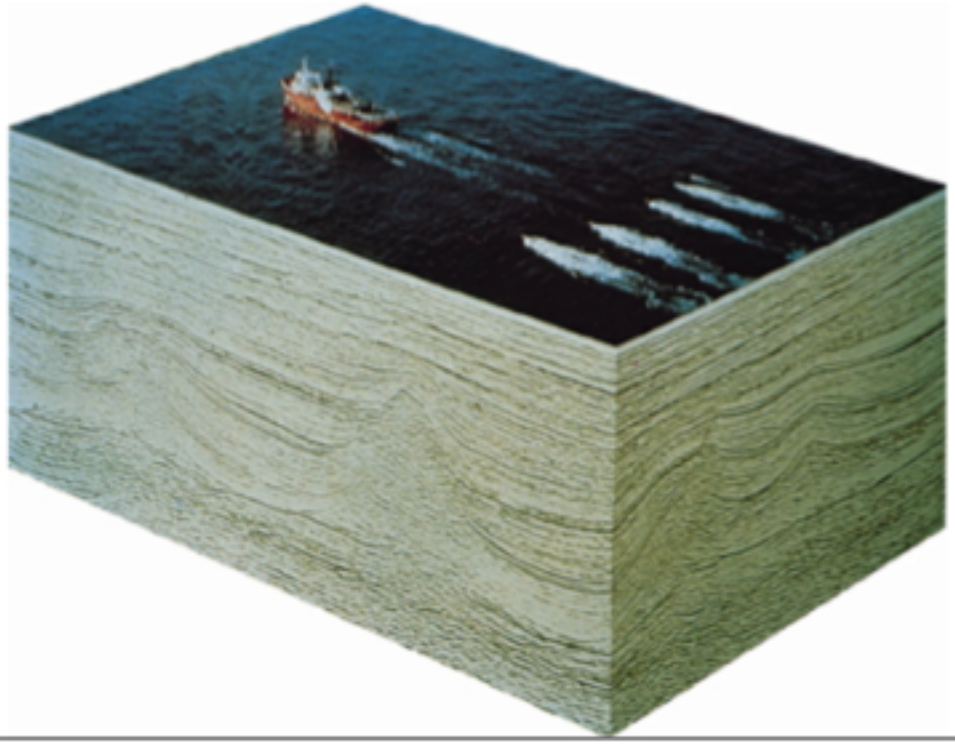

Global Adjoint Tomography: Earthquakes

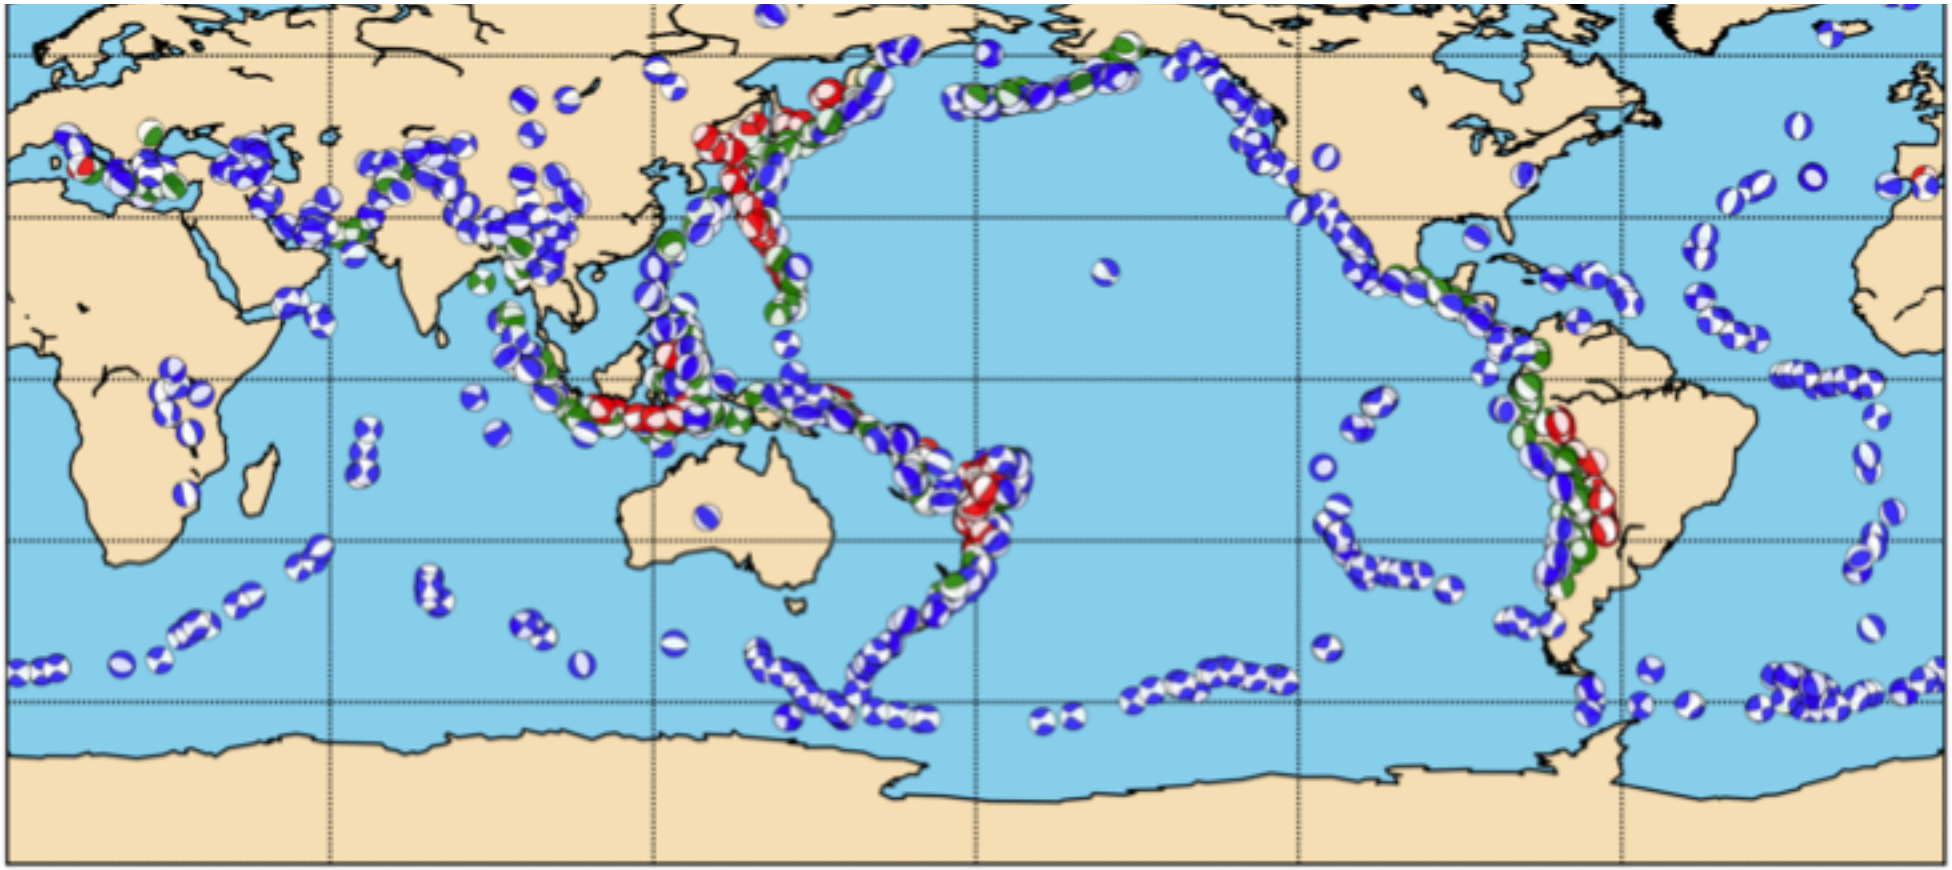

1,480 events

Global Adjoint Tomography: Stations

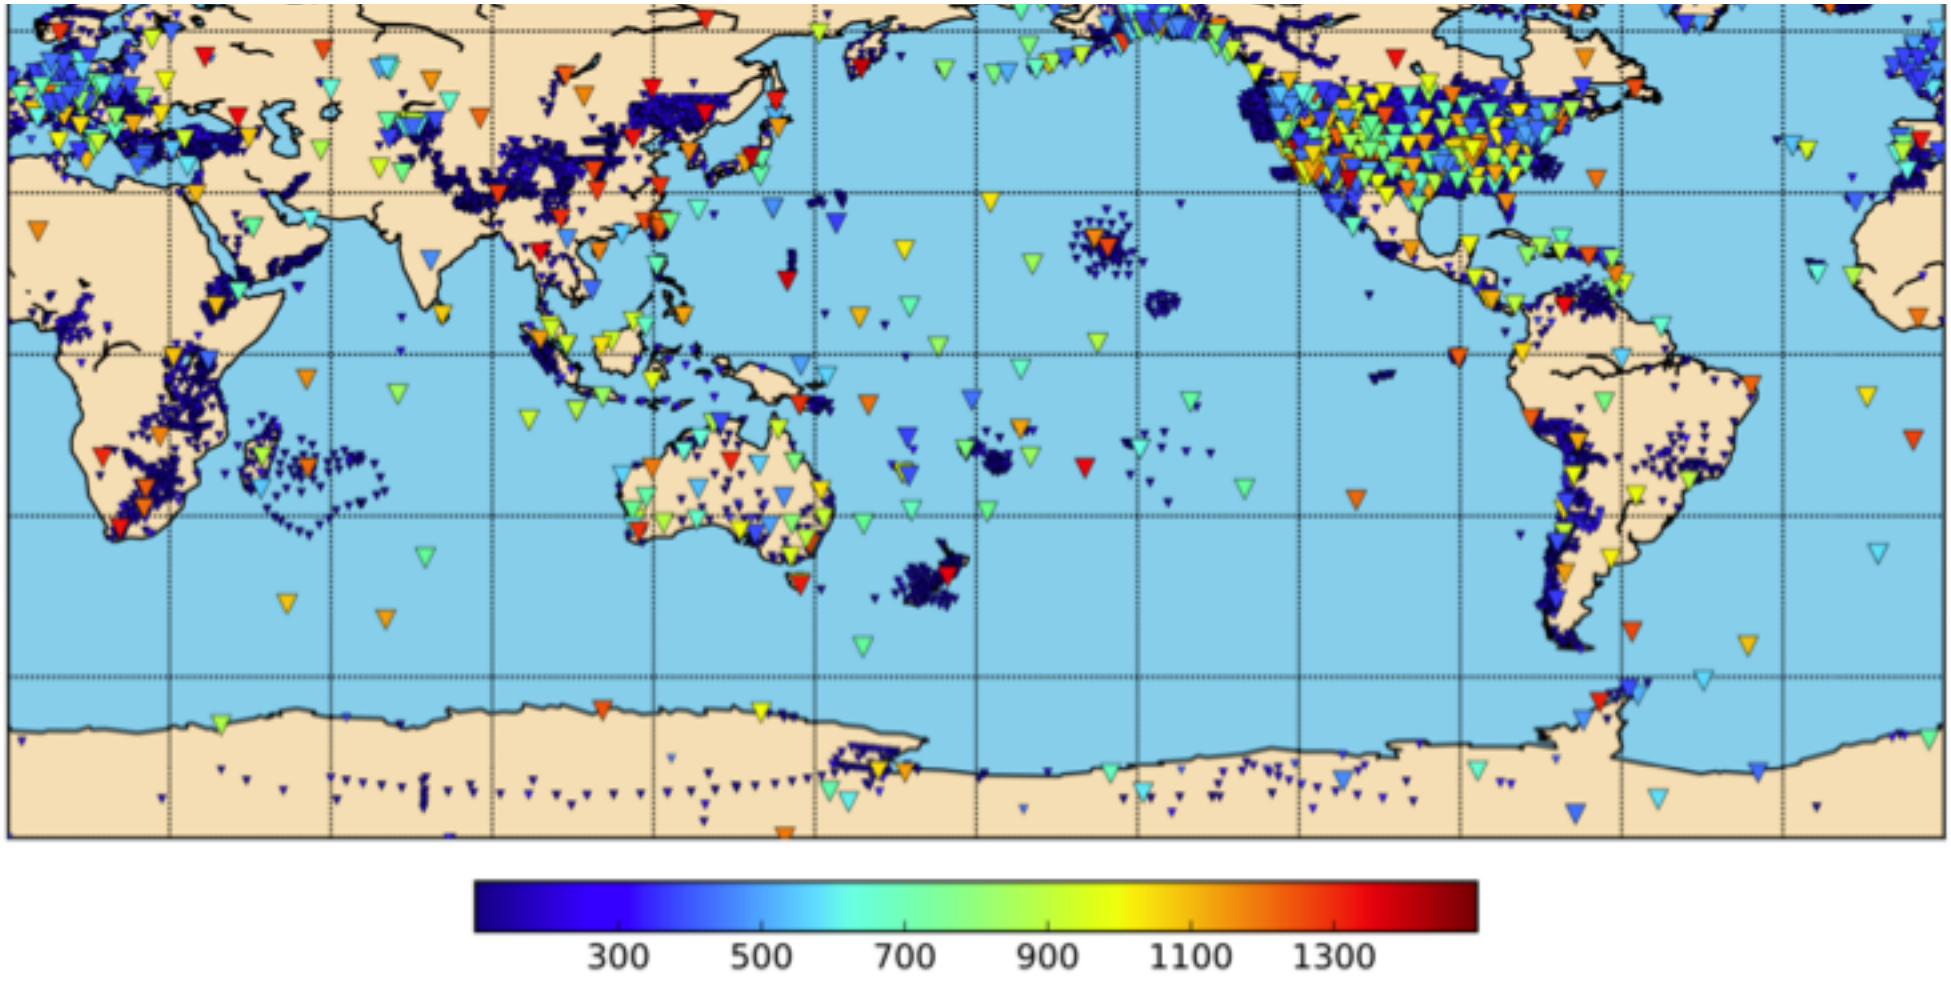

11,800 permanent and temporary seismographic stations

Hotspots & Mantle Plumes

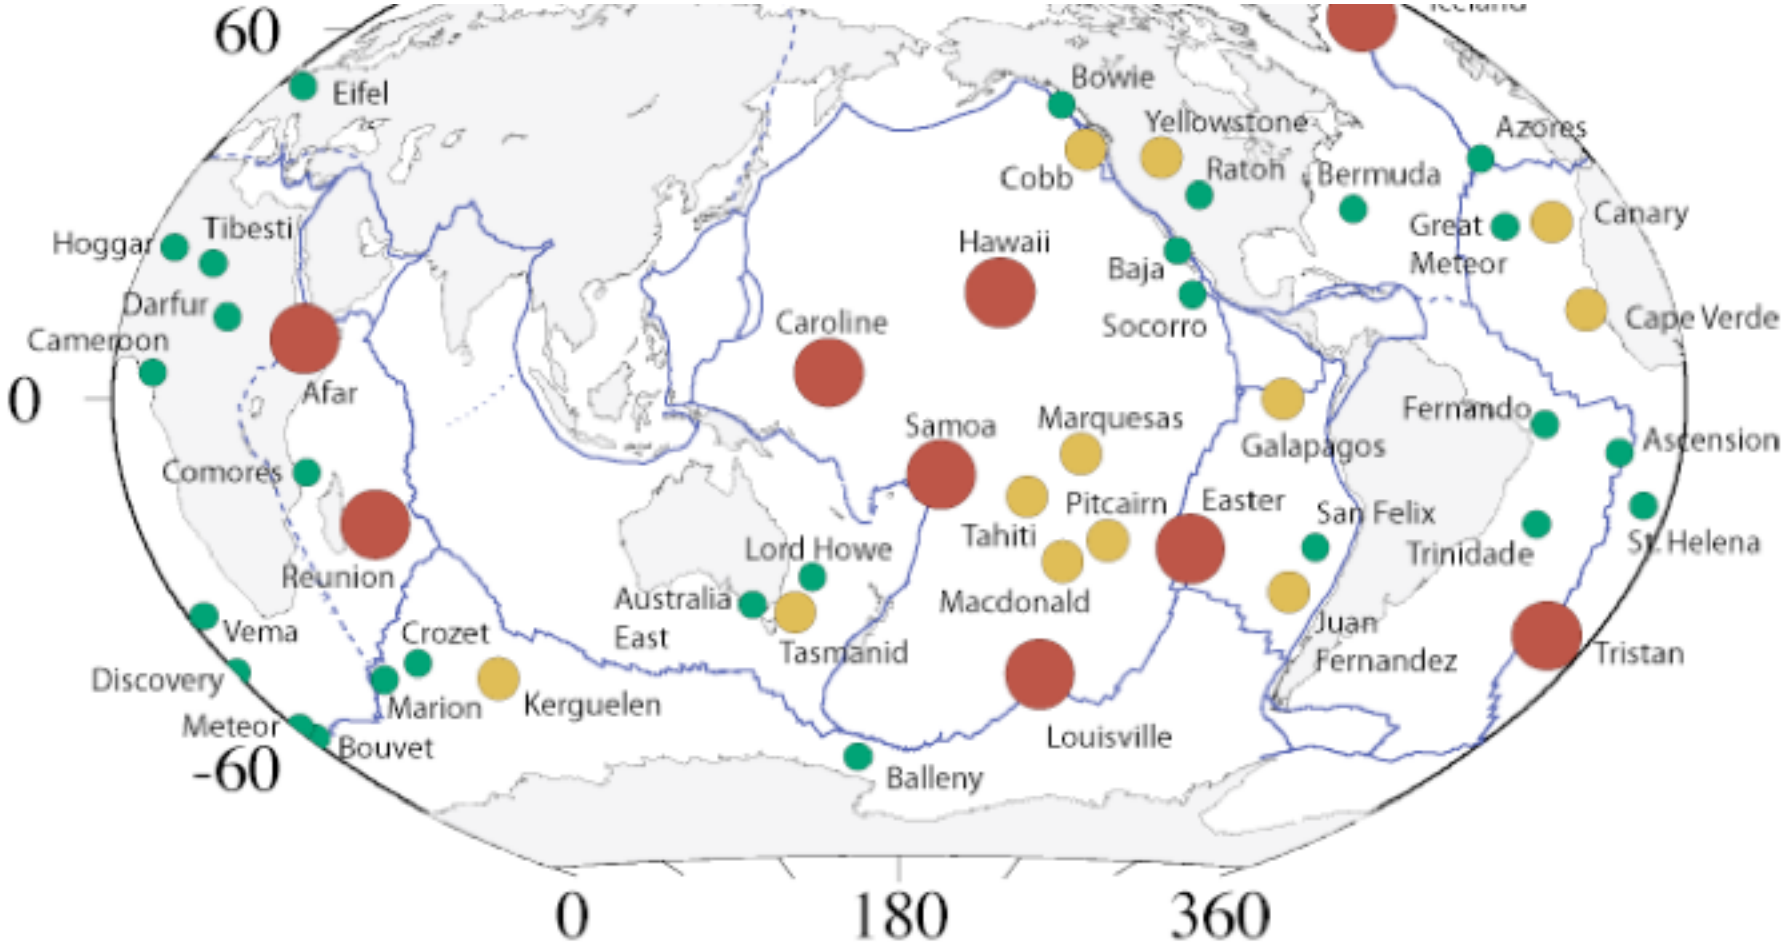

Model GLAD-M25

Indian Ocean at 250 Km

Seychelles

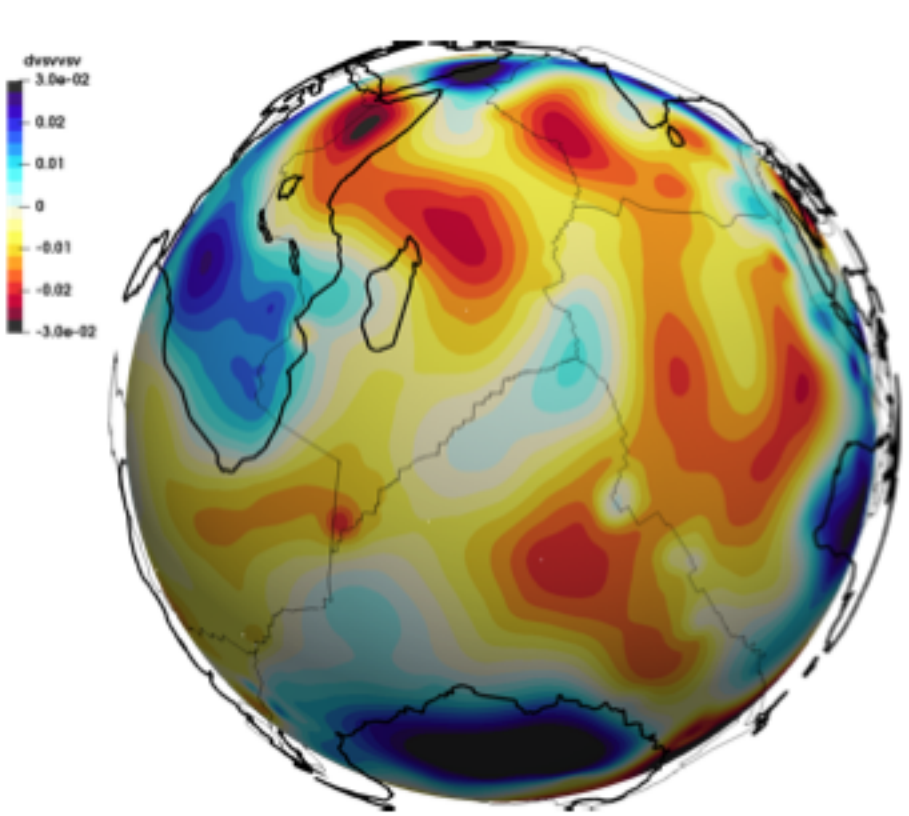

GLAD-M15

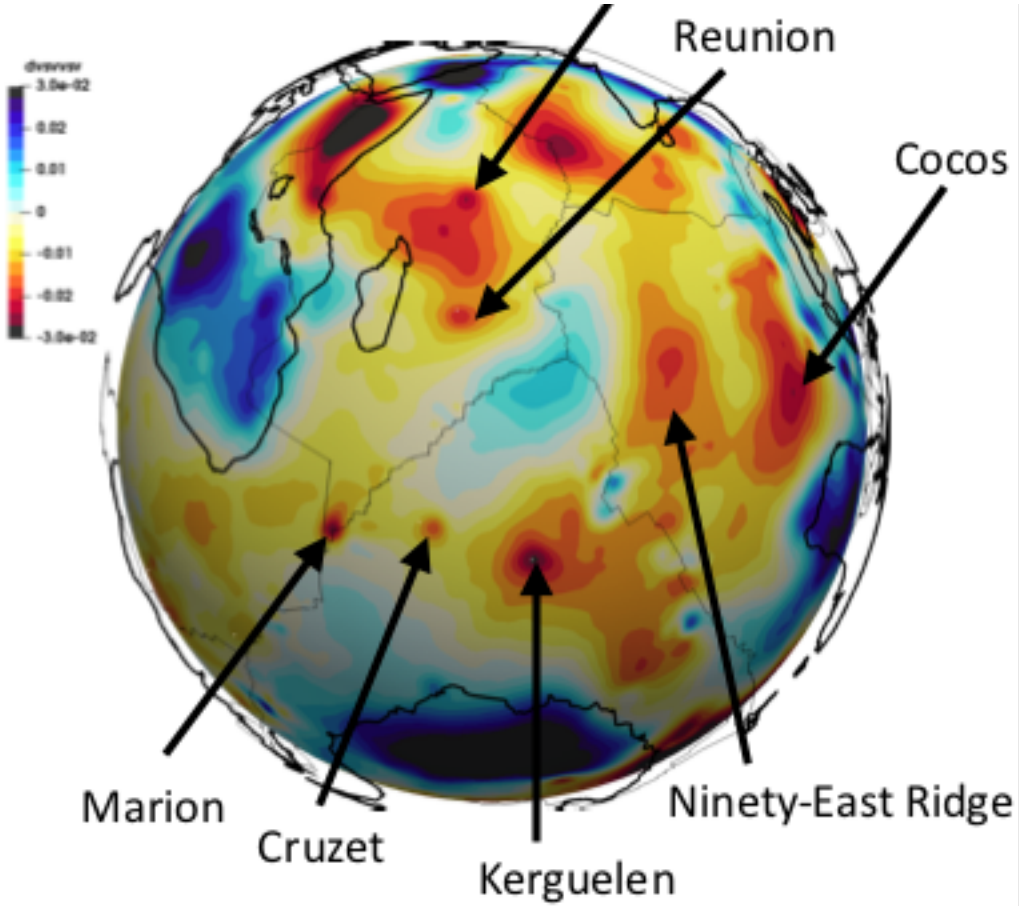

GLAD-M25

Multi-Scale Mantle Plumes

Stable lower-mantle plumes followed by small upper-mantle plumes: primary & secondary plumes

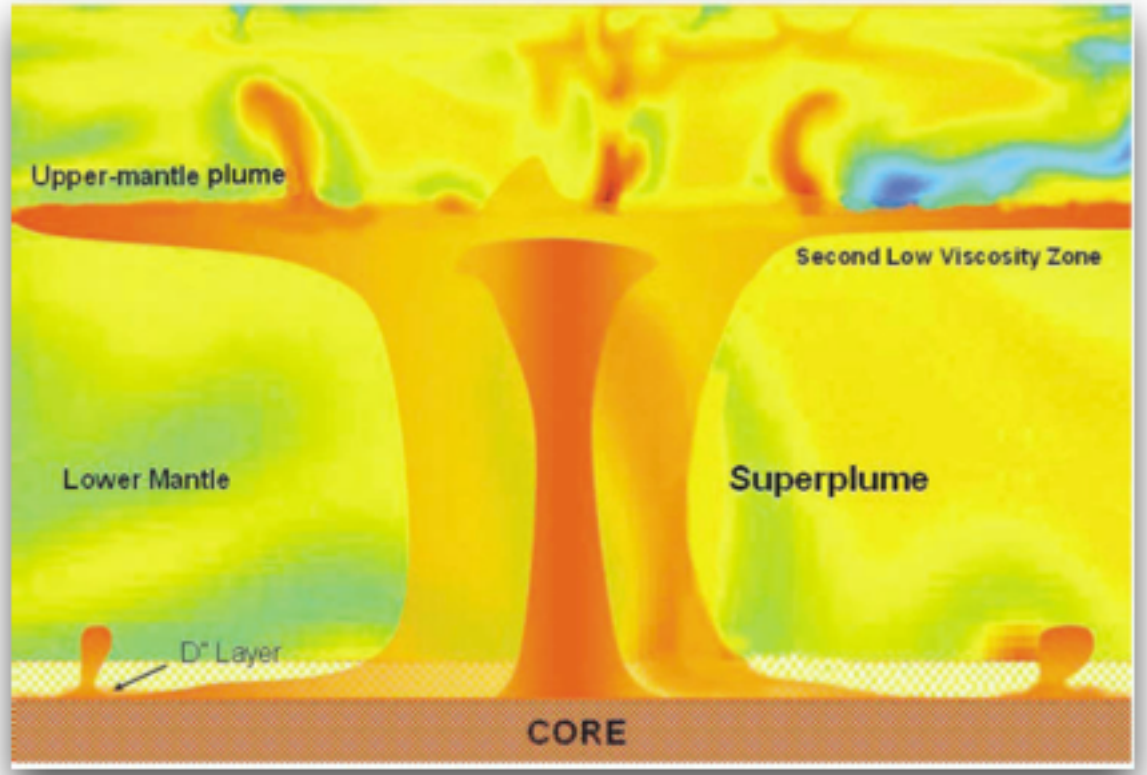

Controlled by:

- mantle viscosity
- thermal conductivity
- thermal expansivity
- phase transitions

Matyska & Yuen, 2007

Tonga Volcanoes in 15th Iteration Model M15

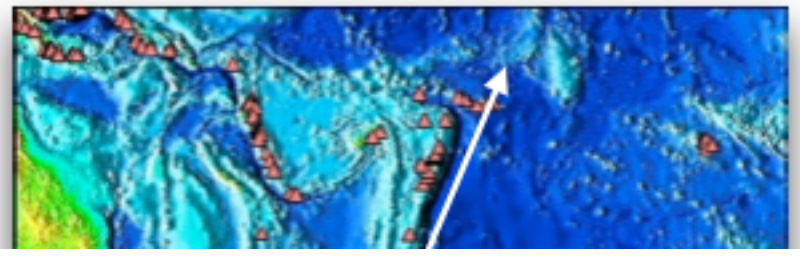

Raoul Island

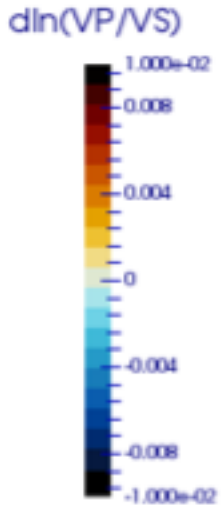

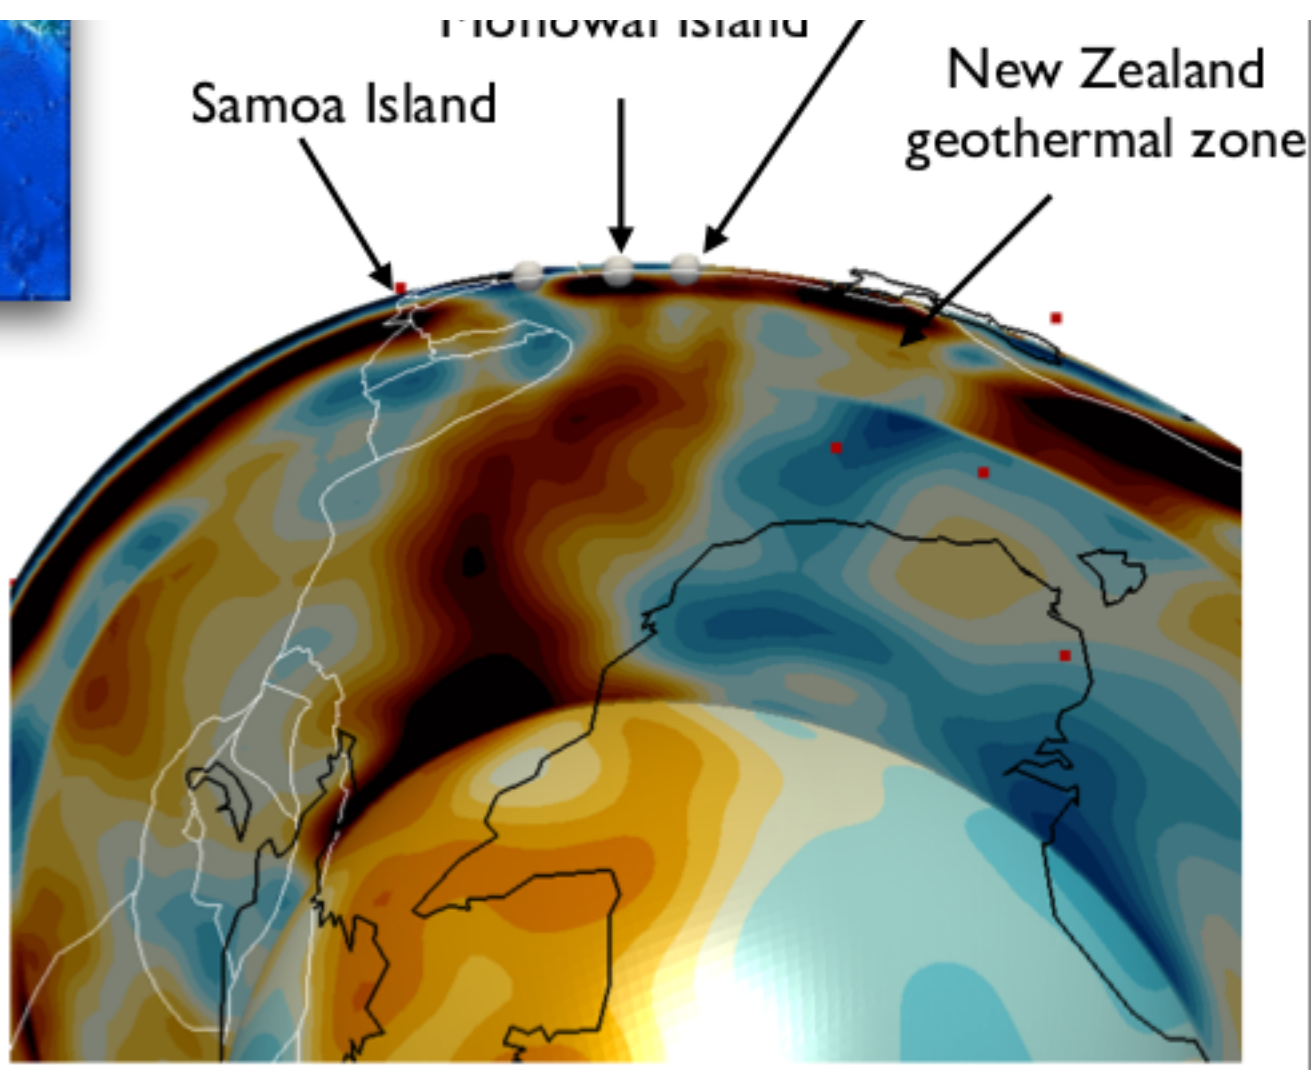

Pillars of the Mantle II

Imaging Earth's Interior with Adjoint Tomography

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY
COLORADOSCHOOLOFMINES EARTH • ENERGY • ENVIRONMENT
KAUST
Aix*Marseille université
PRINCETON UNIVERSITY

David Pugmire & Ebru Bozdağ

Global Adjoint Tomography Workflow

Challenge#1 Data Volume

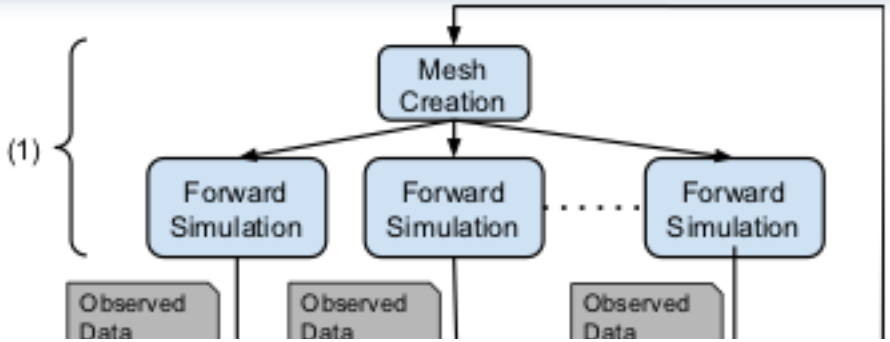

1 PB wavefield files

8 million 120-min 20 Hz seismograms (6 TB)

10 TB kernels

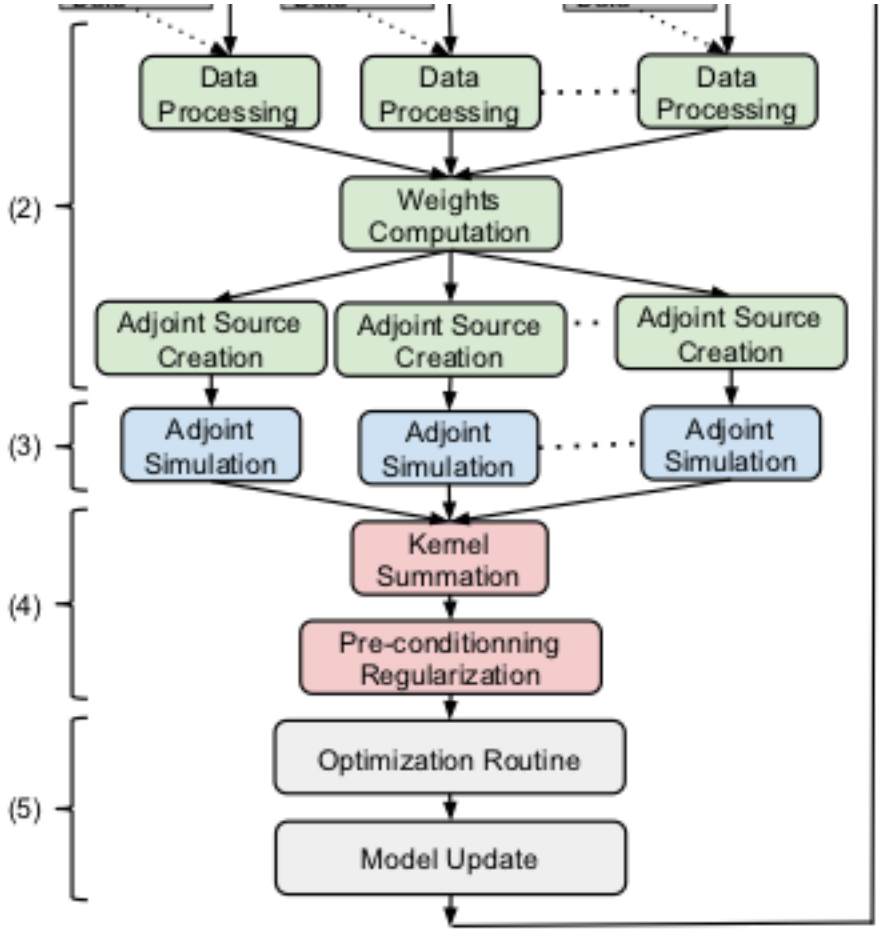

Global Adjoint Tomography Workflow

Challenge#1 Data Volume

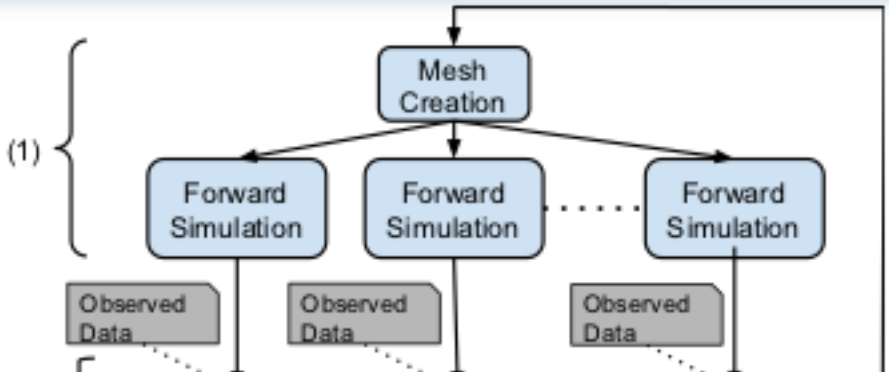

Adaptable Seismic Data Format
ASDF (Krischer et al. 2016)

Adaptable I/O System
ADIOS (Liu et al. 2014)

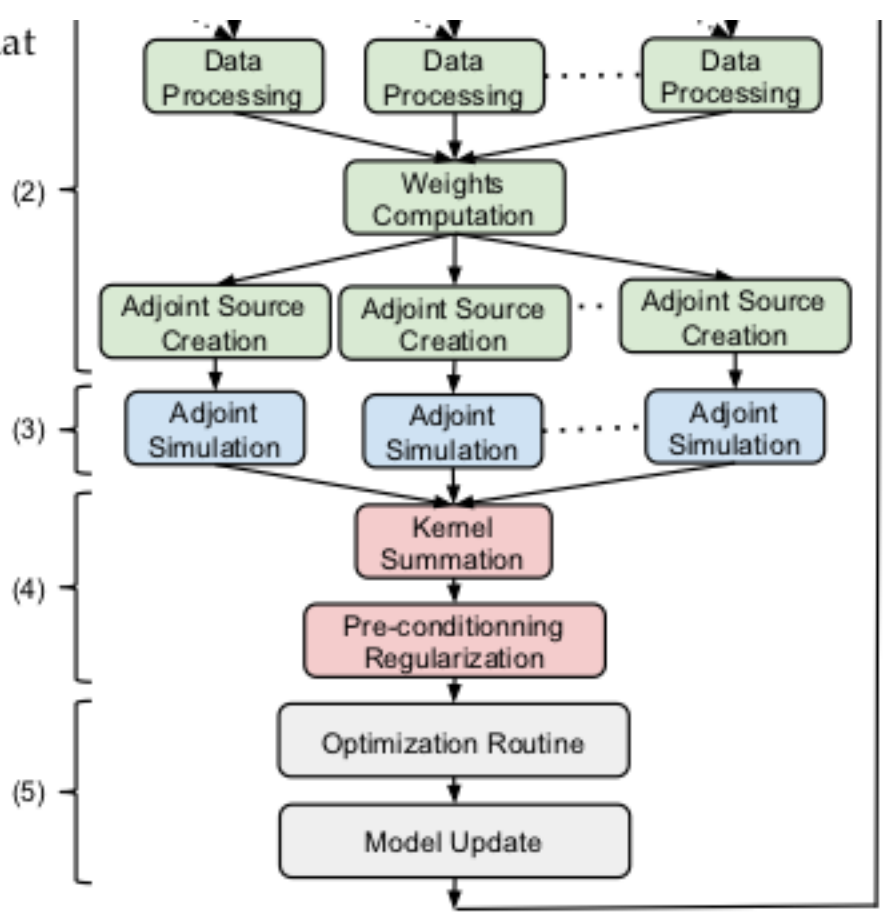

Adaptable Seismic Data Format

- Collaboration involving Princeton University, Münich 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

a single shot of earthquake into one file

- Take advantage of parallel processing
- Use modern file format as container (PHDF5)
- Store provenance inside the file for reproducibility
- Use existing standards when possible (e.g., StationXML, QuakeML)
- Open wiki for development (seismic-data.org)

Krischer et al. 2016

Global Adjoint Tomography Workflow

Challenge #2
Expensive Simulations
&
Complex Workflow

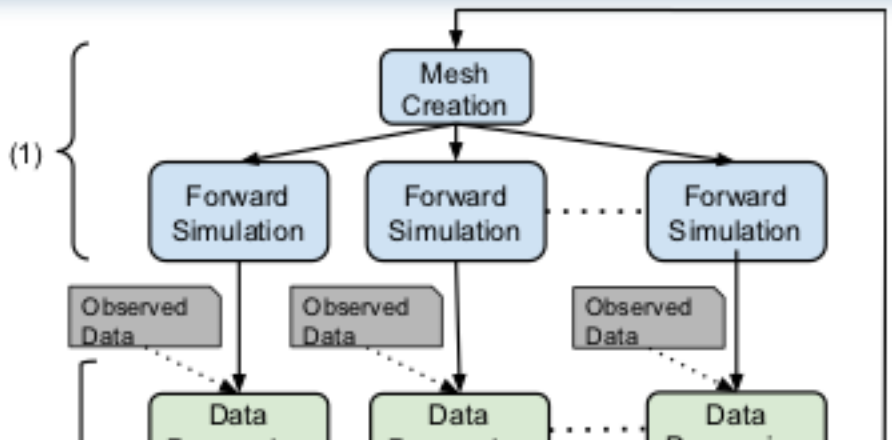

3 million core hours
for forward simulation

0.1 million core hours

for data processing

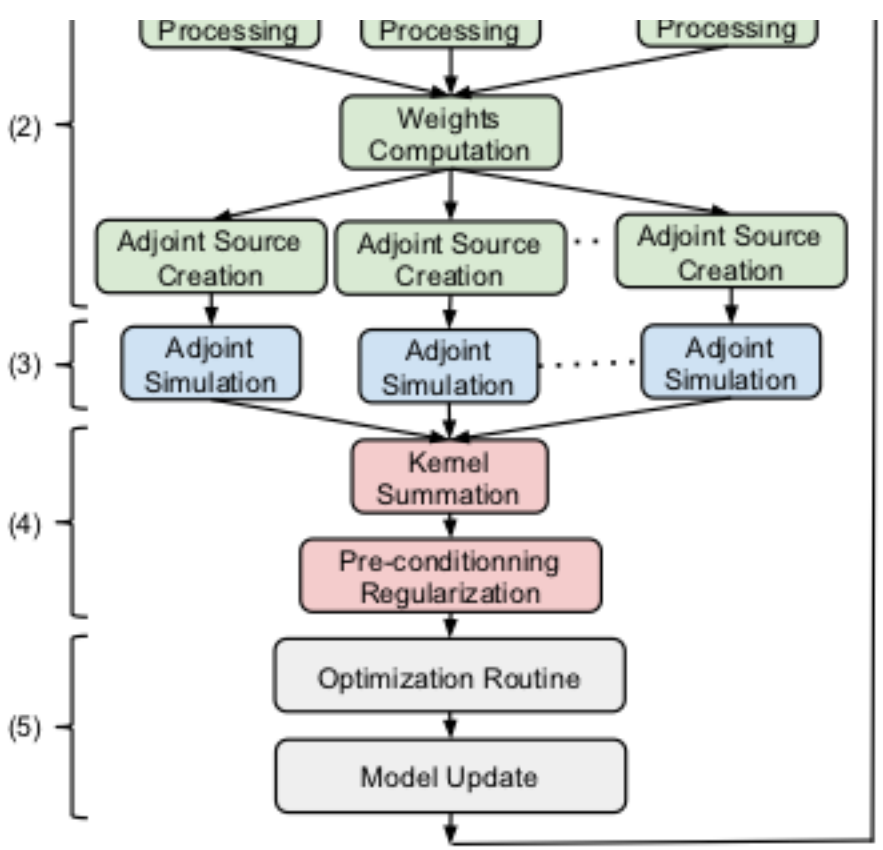

6 million core hours for adjoint simulation

1 million core hours for line search

Global Adjoint Tomography Workflow Management

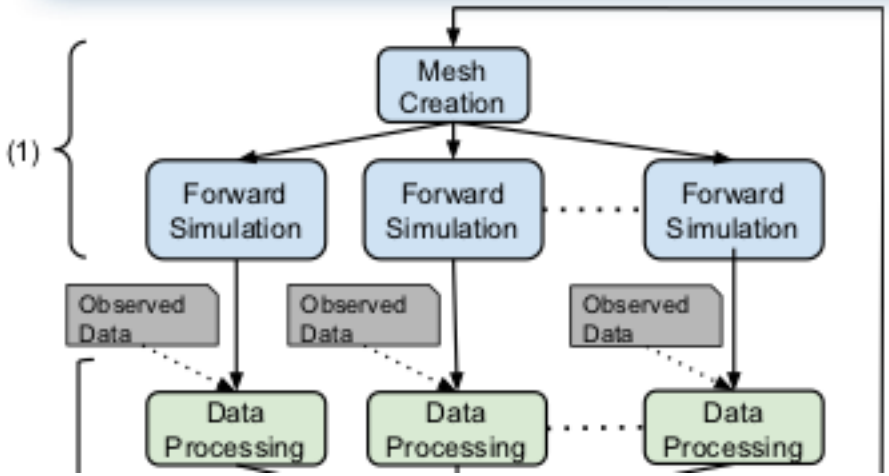

Main sources of trouble:

- Hardware failures
- Human errors

We are implementing the RADICAL EnTK workflow management toolkit:

- Automation, save time & human

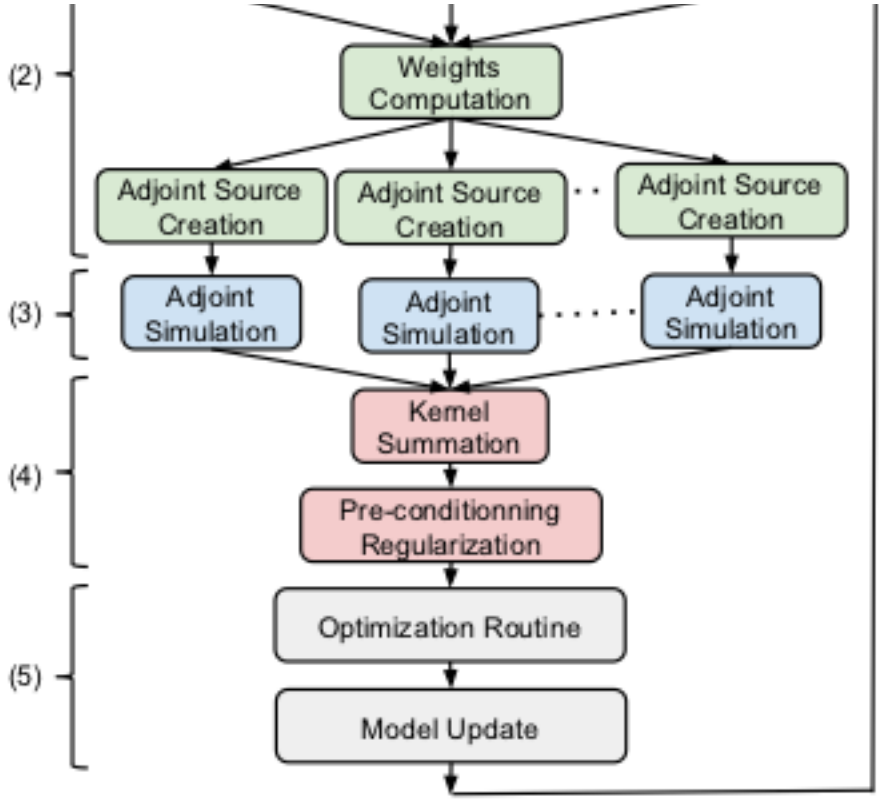

- *Automation*: save time & human effort for repeated tasks
- *Efficiency*: acceleration, taking full advantage of HPC systems
- *Fault tolerance*: automated job failure detection & recovery

Shantenu Jha (Rutgers)

Conclusions

Exascale readiness requires the following investments:

- Continual open source software development, e.g., GPUs & Intel Phi
- All software needs to be under source control (e.g., GitHub) and needs to be automatically and continually tested (e.g., BuildBot,

Travis, Jenkins)

- Workflow management tools (e.g., RADICAL EnTK)
- Modern seismological data format with full provenance (ASDF)
- Standard for the Exchange of Earth Models