Imaging the earth's interior with seismic waves, supercomputers, and PGAS

Scott French[1*] and Barbara Romanowicz[1,2]

[1] *Berkeley Seismological Laboratory, UC Berkeley, Berkeley, CA, USA*
[2] *Institut de Physique du Globe de Paris, Paris, France*
[*] *Now at: Google Inc.*

Motivation: Why seismic imaging?

- **Short answer**: Because we can’t just dig a big hole

- **Surface and space-based observations provide our only window into the evolution and interior dynamics of Earth**
  - Geophysical observations (bathymetry, geoid, heat flow, etc.)
  - Geochemical analysis (meteorites, lavas, xenoliths, etc.)
  - Geodesy (GPS, interferometry, etc.)

- **But how can we actually *see inside* of Earth’s mantle?**
  - **Seismic waves are unique among surface observables:** They carry the signature of the structures through which they have propagated
  - **Seismic imaging** techniques (e.g. tomography) leverage this and enable us to look within

Whole-mantle waveform tomography

- **Objective**: 3D model of material properties (elastic wave speed) throughout the earth's entire mantle (the outer 2890 km)
- **Observations**: Seismograms of natural earthquakes (hundreds)
- **Predictions**: Numerical simulations of seismic wave propagation

Reconciling observation and prediction

- Defines a nonlinear inverse problem
- Prediction (numerical simulation) is expensive: **500K – 1M** CPU hours
- Too costly for stochastic methods
- Must be solved iteratively

Left: Earlier waveform tomographic model SEMum2 covered only the upper ~ 800 km of the mantle (*French et al., 2013, Science*)

Waveform tomography in practice

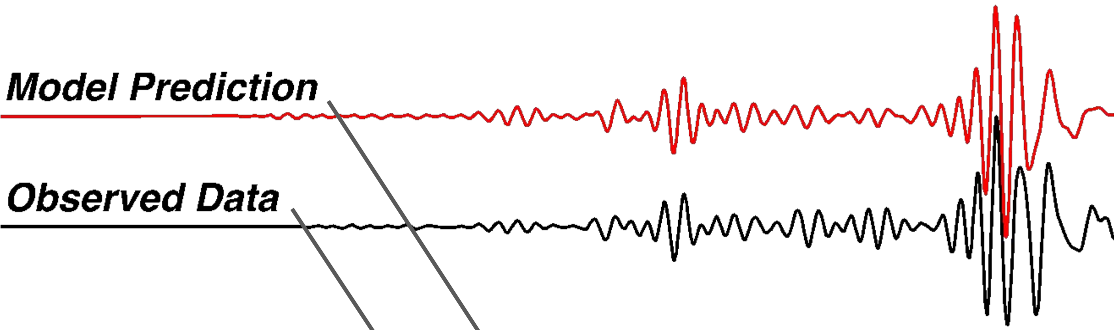

$$2\chi(\mathbf{m}) = \|\mathbf{d} - \mathbf{g}(\mathbf{m})\|^2_{\mathbf{C}_\mathbf{d}^{-1}} + \|\mathbf{m} - \mathbf{m}^p\|^2_{\mathbf{C}_\mathbf{m}^{-1}}$$

$$\mathbf{m}^{i+1} = \mathbf{m}^i + \left(\mathbf{C}_\mathbf{m}\mathbf{G}^T\mathbf{C}_\mathbf{d}^{-1}\mathbf{G} + \mathbb{I}\right)^{-1} \left(\mathbf{C}_\mathbf{m}\mathbf{G}^T\mathbf{C}_\mathbf{d}^{-1}\left[\mathbf{d} - \mathbf{g}(\mathbf{m}^i)\right] - \mathbf{m}^i + \mathbf{m}^p\right)$$

Key components:

- Data representation
- Misfit function
- Theoretical treatment of wave propagation
- Optimization scheme
- Starting model

Computational steps (until convergence):

- Wavefield simulation (SEM)
- Hessian and gradient computation (NACT)
- Assembly and solution of the update system

$$\mathbf{G}_{ij} = \partial \mathbf{g}_i(\mathbf{m}) / \partial \mathbf{m}_j$$

Step I: Wavefield simulation

Method of choice: Spectral Element Method

- Cheap time integration ($\mathbf{M}$ diagonal)
- Very low numerical dispersion
- Natural b.c. treatment (free surface)
- Straight-forward meshing and parallel decomposition

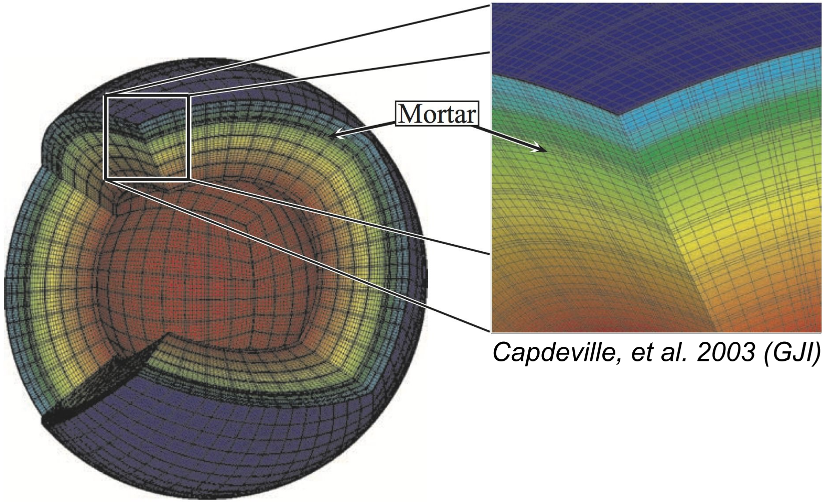

Capdeville, et al. 2003 (GJI)

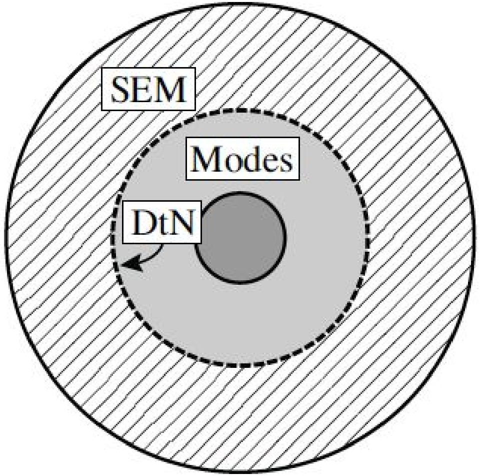

Coupled-SEM implementation

- Fortran 90 + MPI (+ OpenMP work-in-progress)
- Coupled to an analytical solution in the core (DtN operator)
- Anisotropic homogenization of thin layers: improved time stability, fewer integration steps (shorter simulations)
- Mortar method for non-conforming mesh refinement

Step II: Parallel Hessian assembly

Nonlinear Asymptotic Coupling Theory (NACT)

- Calculates $\mathbf{G}_{ij} = \partial\, \mathbf{g}_i(\mathbf{m}) \,/\, \partial\, \mathbf{m}_j$
  - Used in computing Hessian estimate $\mathbf{G}^T\mathbf{G}$ and gradient $-\mathbf{G}^T[\mathbf{d} - \mathbf{g}(\mathbf{m})]$
- **Data parallelism**: Each seismogram yields an independent strided row-panel of $\mathbf{G}$

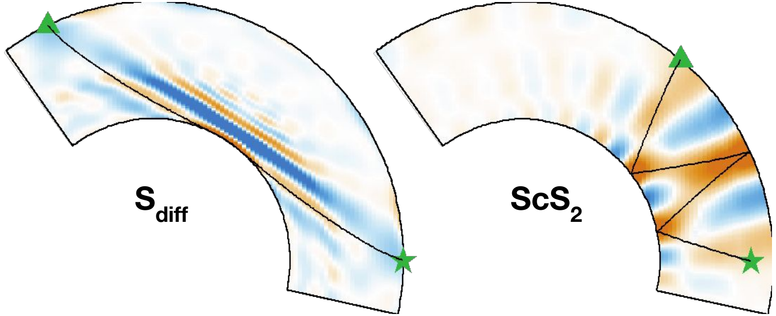

Above: Example NACT partial derivatives at arbitrary times during S_{diff} and ScS_2 arrivals

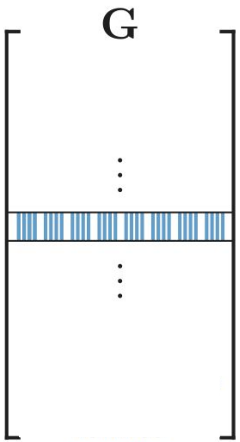

- $\mathbf{G}$ is non-sparse (10% nz) and unwieldy (~13 TiB in our recent work)
- Instead, we directly form $\mathbf{G}^T\mathbf{G}$ (~180 GiB) using a custom PGAS distributed matrix abstraction
- Implemented mainly in C, along with C++ (UPC++), OpenMP, and MPI-IO

Step II: Parallel Hessian assembly (continued)

Partitioned Global Address Space (PGAS) model

- Logically partitioned globally “shared” address space supporting one-sided access
- Excellent fit for distributed data structures + irregular access patterns

Distributed matrix abstraction (*French et al., IPDPS’15*)

- Based on UPC++: A set of PGAS extensions to C++ (*Zheng, et al. IPDPS’14*)
    - Modeled primarily on UPC, but adds: **Dynamic remote memory management** and **asynchronous remote tasks**
    - Key to implementing one-sided updates optimized for our use case (`+=` only, assume associative / commutative, can progress asynchronously)
- Distributed matrices use block-cyclic PBLAS-compatible format
    - ScaLAPACK used to solve the Gauss-Newton model update equation
- Performs significantly better than solutions based on `MPI_Accumulate`

Putting it all together at NERSC

- **SEM simulations**: Hopper
  - 500 - 1000 runs per iteration
  - 12 - 24 nodes, aggregated
  - ~ 90% of our allocation
- **Hessian estimation and Gauss-Newton updates**: Edison
  - 10 - 20 runs per iteration
  - 128 - 512 nodes, standalone
- **Three iterations, plus an additional round of simulations (training and validation)**
  - **~ 3.1M raw core hours**

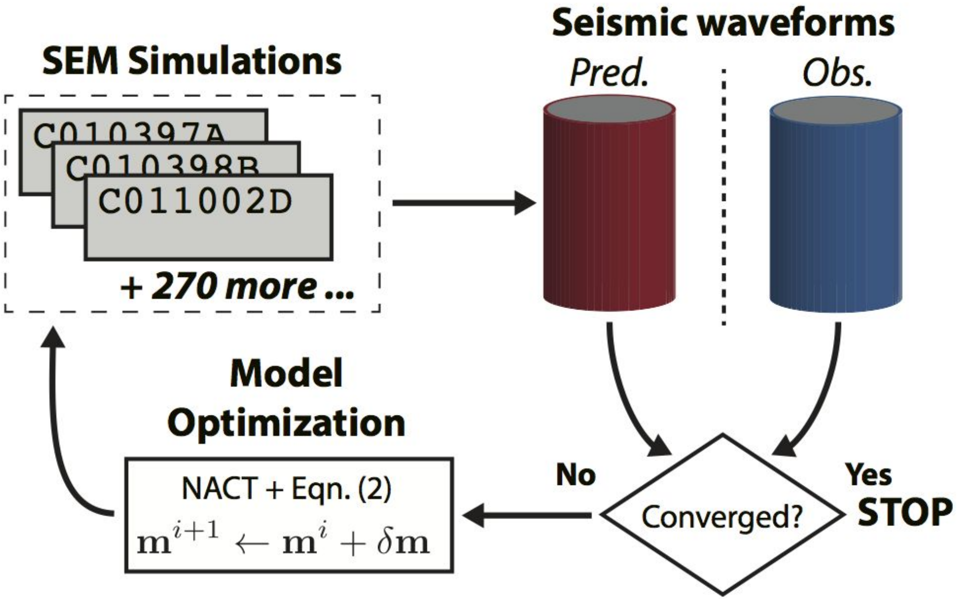

Above: An overview of the iterative waveform inversion procedure deployed at NERSC.

Scientific results: A whole-mantle model

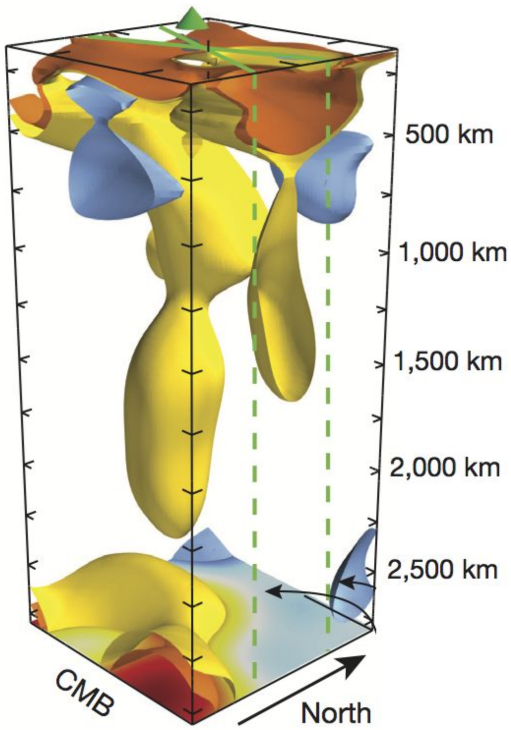

Above: 3D rendering of shear-velocity structure beneath the Hawaii hotspot.

Geophysical Journal International

Geophys. J. Int. (2014) **199,** 1303–1327 doi: 10.1093/gji/ggu334

GJI Seismology

Whole-mantle radially anisotropic shear velocity structure from spectral-element waveform tomography

S. W. French[1,*] and B. A. Romanowicz[1,2,3]

LETTER

doi:10.1038/nature14876

Broad plumes rooted at the base of the Earth's mantle beneath major hotspots

Scott W. French[1]† & Barbara Romanowicz[1,2,3]

- The **first** whole-mantle seismic model based on waveform tomography using numerical wavefield simulations
- Reveals **new details** of earth structure not seen in previous models based on approximate forward modeling techniques (especially low shear-velocity structures)

Scientific results: A whole-mantle model

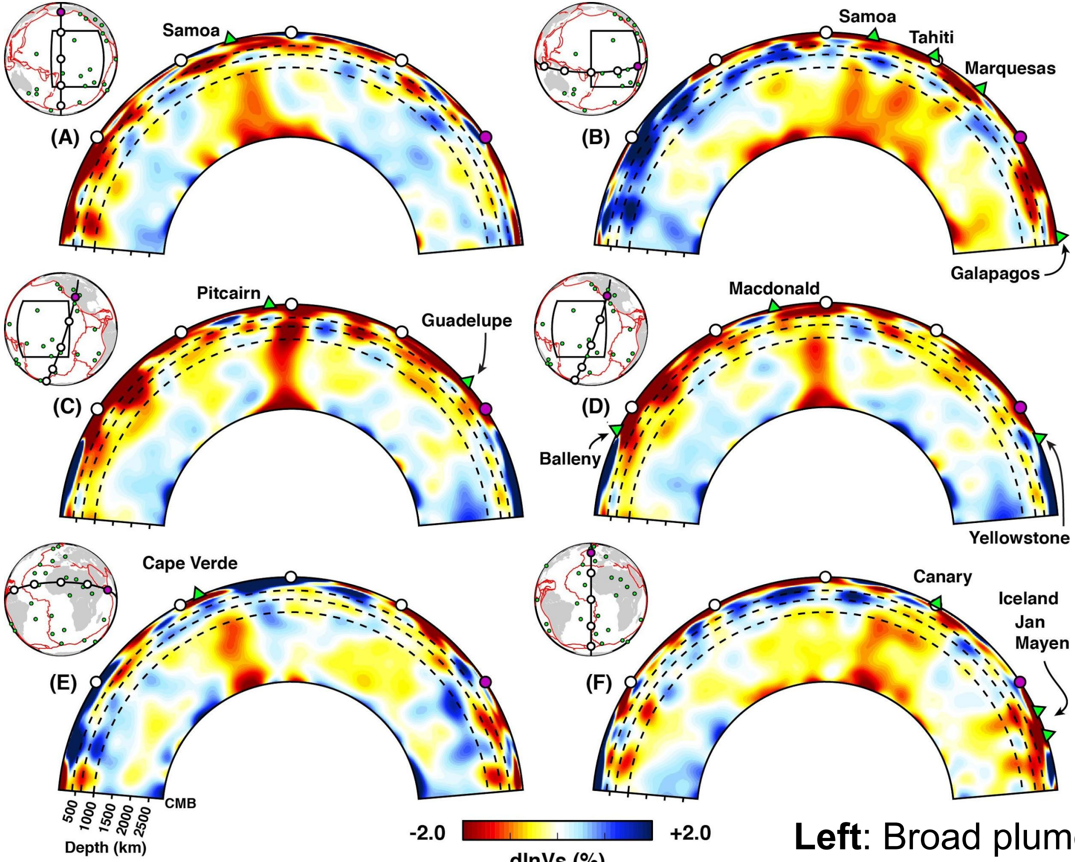

- **Unambiguous detection** of columnar low-velocity anomalies beneath major hotspots (plumes)
- Plumes are **unusually broad** in the lower mantle (deeper than 1000 km) and **clearly deflect** at that depth
- **Independently corroborated** by isotope signatures, localized seismic observations, regional high-resolution models, geodynamic modeling efforts

Left: Broad plumes in the earth's lower mantle, including those beneath Pitcairn, Samoa, Cape Verde, and other hotspots.

Conclusion

Scientific contributions

- **First-ever** whole-mantle seismic model based on numerical wavefield simulations
  - **Unambiguous detection** of “plumes” beneath major hotspots
- **Impact**: New constraints on future geodynamic models, present-day mantle circulation, Earth’s heat budget
- **Future directions (ongoing)**: Starting condition for high-resolution regional imaging, inversion for global anelastic structure

Made possible thanks to

- **NERSC resources**: Without access to NERSC resources, and the ease of scientific productivity thereon, this study would not have been possible
- **Powerful PGAS programming systems**: Access to UPC++ and discussions with the DEGAS group enabled us to extend our imaging to whole-mantle scale

Thank you

This work was supported by an NSF Graduate Research Fellowship to SWF, NSF grant EAR-1417229, and ERC Advanced Grant WAVETOMO. Computations were performed at the National Energy Research Scientific Computing Center, supported by the US Department of Energy Office of Science (contract DE-AC02-05CH11231).

Extra slides

Distributed matrix abstraction

Typical configuration: One UPC++ process per NUMA domain, many OpenMP threads

```
ConvergentMatrix<float,...> GtG( M, M );
...
// for each locally computed update
GtG.update( GtG_i, slice_idx_i );
```

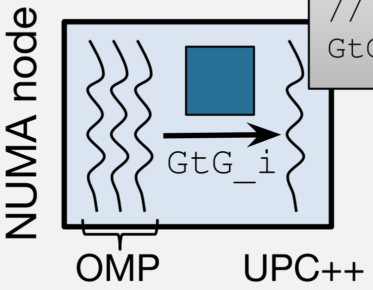

1. Bin `GtG_i` elements by target process
2. For each target:
   data movement: `upcxx::allocate`, `upcxx::copy`
   update task: `upcxx::async`

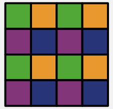

Eventually on all UPC++ processes ...

```
GtG.commit(); // barrier
// fetch local pointer
float *mat = GtG.get_local_data();
// ScaLAPACK
// MPI-IO collective write
```

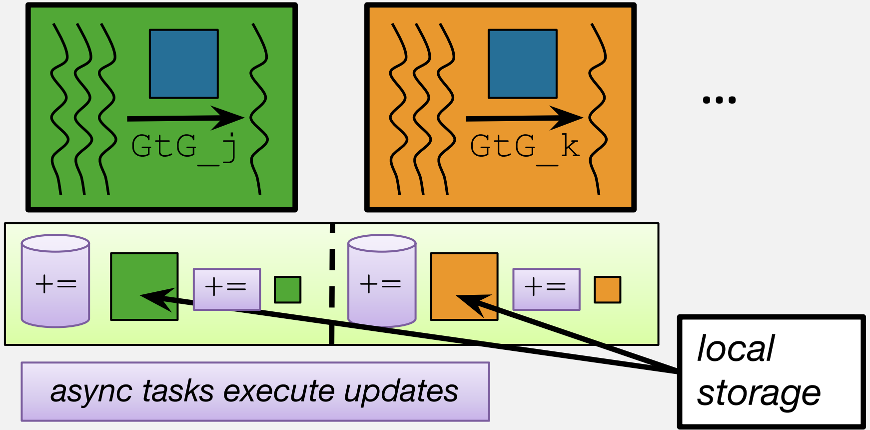

Strong scaling (Hessian estimation)

- Near complete overlap of computation and communication
  - Largest overhead growth at higher concurrency is binning
  - Readily scales to next-generation problem size

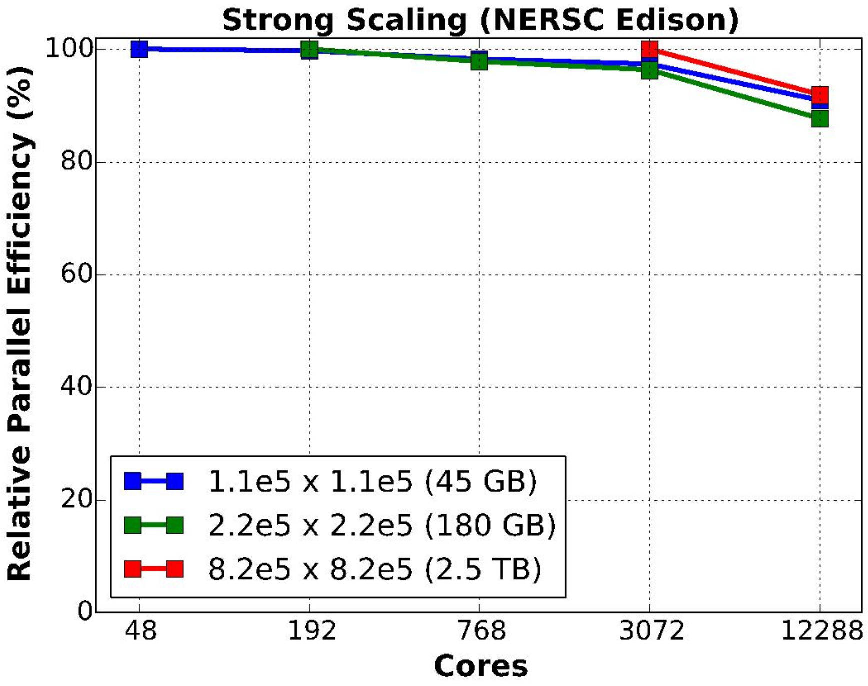

Setup

- NERSC Edison (Cray XC30)
  - 5,576 2x 12-core Intel IVB
  - 64 GB DDR3 per node
  - Cray Aries interconnect
- GNU Compilers 4.8.2 (-O3)
- GASNet-1.22 / UPC++ master
- Up to 12,288 cores
- Matrix size: 50GB – 2.5TB

Weak scaling vs. MPI (Hessian estimation)

- Distributed matrix size fixed (180 GB)
- Dataset size scaled w/ concurrency
  - 64 updates per MPI or UPC++ task + thread team (NUMA domain)

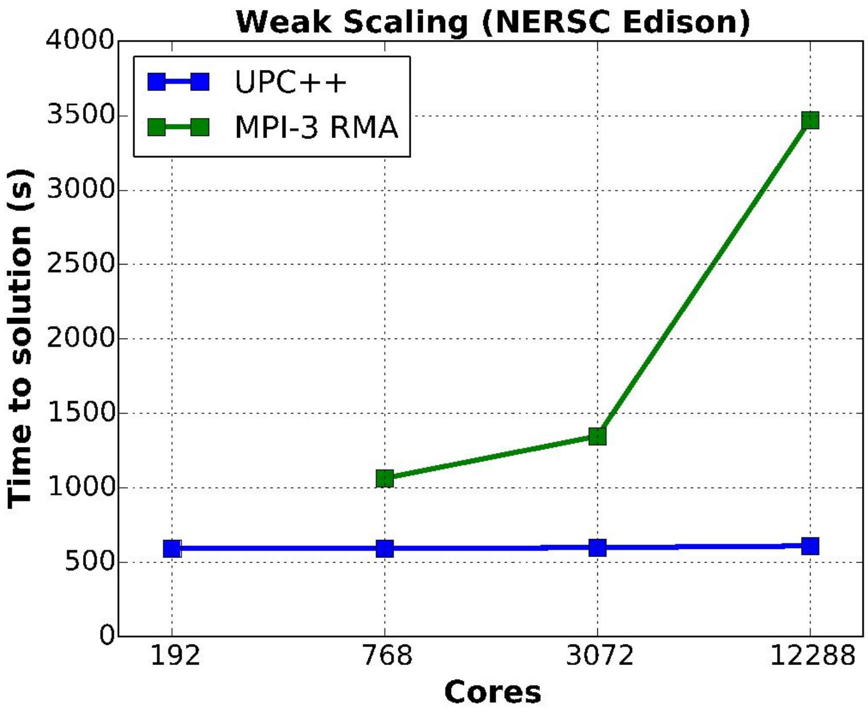

Setup

- NERSC Edison (Cray XC30)
- GNU Compilers 4.8.2 (-O3)
- Cray MPICH 7.0.3
- Up to 12,288 cores
- Matrix size: 180GB

Scientific results: Independent studies

GEOPHYSICS

Viscosity jump in Earth's mid-mantle

Maxwell L. Rudolph,[1]* Vedran Lekić,[2] Carolina Lithgow-Bertelloni[3]

Probabilistic inversion of Earth's non-hydrostatic geoid (gravitational equipotential surface), combined with geodynamic modeling

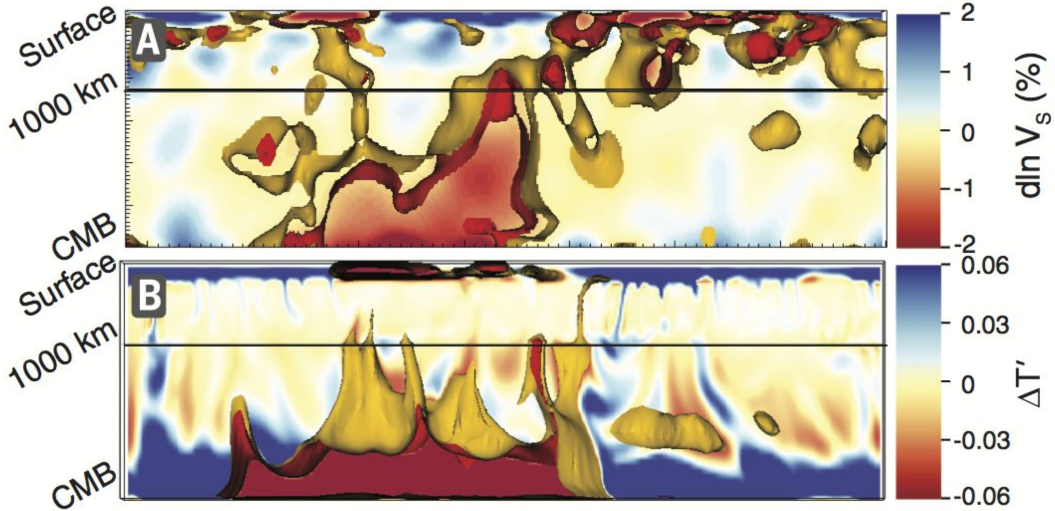

Left: Inferred viscosity step superimposed on shear-velocity variation in our seismic model (**top**); Geodynamic model of mantle convection with the implied viscosity contrast (**bottom**).

Rudolph, M., V. Lekic, and C. Lithgow-Bertelloni (2015), Viscosity jump in the Earth's mid mantle, Science, 360 (6266), 1349-1352

Scientific results: A whole-mantle model