Opportunities and Challenges in Computational Geophysics

Michael Gurnis

CIG & Seismological Laboratory, Caltech

Omar Ghattas

ICES, Jackson School of Geosciences, TACC,
Univ. of Texas, Austin

Computational Infrastructure for Geodynamics

National Science Foundation
WHERE DISCOVERIES BEGIN

Also From CIG

Marc Spiegelman

Applied Physics/Applied Math

Lamont-Doherty Earth Observatory, Columbia University

Chairman, Executive Committee

Brad Aagaard

U.S. Geological Survey, Menlo Park

Earthquake Hazards

Chairman, Science Steering Committee

Peter Olson

Earth & Planetary Sciences

Johns Hopkins University

Executive Committee (Former Chair SSC)

Outline

- What is CIG?
- Major solid-earth questions, challenges & opportunities
  - Mantle convection-plate tectonics-magma dynamics
  - Earthquake cycle
- Achievement Highlights
- A geophysics community advancing forward with computational infrastructure: An example
- General needs in applied math, computer science, and scientific computing
- Partnering computational science with specific disciplines, CIG-2

Computational Infrastructure for Geodynamics

National Science Foundation
WHERE DISCOVERIES BEGIN

What is CIG? (1)

- Computational Infrastructure in Geodynamics. An NSF Center that started in September, 2004; Beginning year 4 of 5.
- Core Philosophy of CIG Proposal: Computation has become an essential component of our science and we have successfully built our tools in isolation. In order to accelerate progress, CIG is a partnership between Earth Sciences and Computational Sciences and an attempt to deploy software engineering and the latest computational science tools in a new generation of infrastructure.

What is CIG? (2)

- Under community control with an Executive Committee and Science Steering Committee elected by 41 Member Institutions.
- Managed by Caltech
- Core resources (In house: 4.5 engineers, 1 technical writer; 1.5 admin/sys-adm. Outside: 0.5 FTE in support of PETSc at ANL; ~1 FTE outsourced)
- PetaApps Mantle Convection award (~3 FTE for 3 yrs: UT-Austin, CU, Penn, CIT)
- Support several workshops & training sessions each year
- Work in multiple sub-disciplines:
  - Mantle Convection
  - Magma dynamics
  - Short time-scale tectonics (EQ cycle)
  - Long time-scale tectonics (mountain building)
  - Computational seismology
  - Geodynamo

CIG is a Community Organization

- Argonne National Laboratory (MSC)
- Arizona University
- Brown University
- California Institute of Technology
- Colorado School of Mines
- Colorado State University
- Columbia University
- Cornell University
- Georgia Institute of Technology
- Harvard University
- Johns Hopkins University
- Lawrence Livermore National Laboratory
- Los Alamos National Laboratory (ES)
- Massachusetts Institute of Technology
- Oregon State University
- Pennsylvania State University
- Princeton University
- Purdue University
- Rensselaer Polytechnic Institute
- State University of New York at Buffalo
- State University of New York at Stony Brook
- U.S. Geological Survey (Menlo Park)
- University of California, Berkeley
- University of California, Davis
- University of California, Los Angeles
- University of California San Diego
- University of Colorado
- University of Hawaii
- University of Maine
- University of Maryland
- University of Michigan
- University of Minnesota
- University of Missouri-Columbia
- University of Nevada, Reno
- University of Oregon
- University of Southern California
- University of Texas at Austin
- University of Washington
- Virginia Polytechnic Institute and State University
- Washington University
- Woods Hole Oceanographic Institution

Foreign Affiliates

- Australian National University
- Geological Survey of Norway (NGU)
- GNS Science
- Monash University
- Munich University LMU
- University College London
- University of Science and Technology of China
- University of Sydney
- Victorian Partnership for Advanced Computing

Big Earth Science Questions

- Mantle convection-plate tectonics-magma dynamics
  - Integrates disparate components of how the planet works as a dynamic system
  - Multi-scale & multi-physics are key
  - Closely linked with the major observational initiatives
    - USArray component of EarthScope
    - Subduction Factory component of MARGINS
    - EarthChem Data Base
- Earthquake cycle
  - Major questions of significance for both basic science and hazard reduction
  - Multi-scale
  - Closely linked with Major observation initiatives:
    - PBO & SAFOD components of EarthScope
    - SEIZE component of MARGINS

Mantle Convection - Plate Tectonics

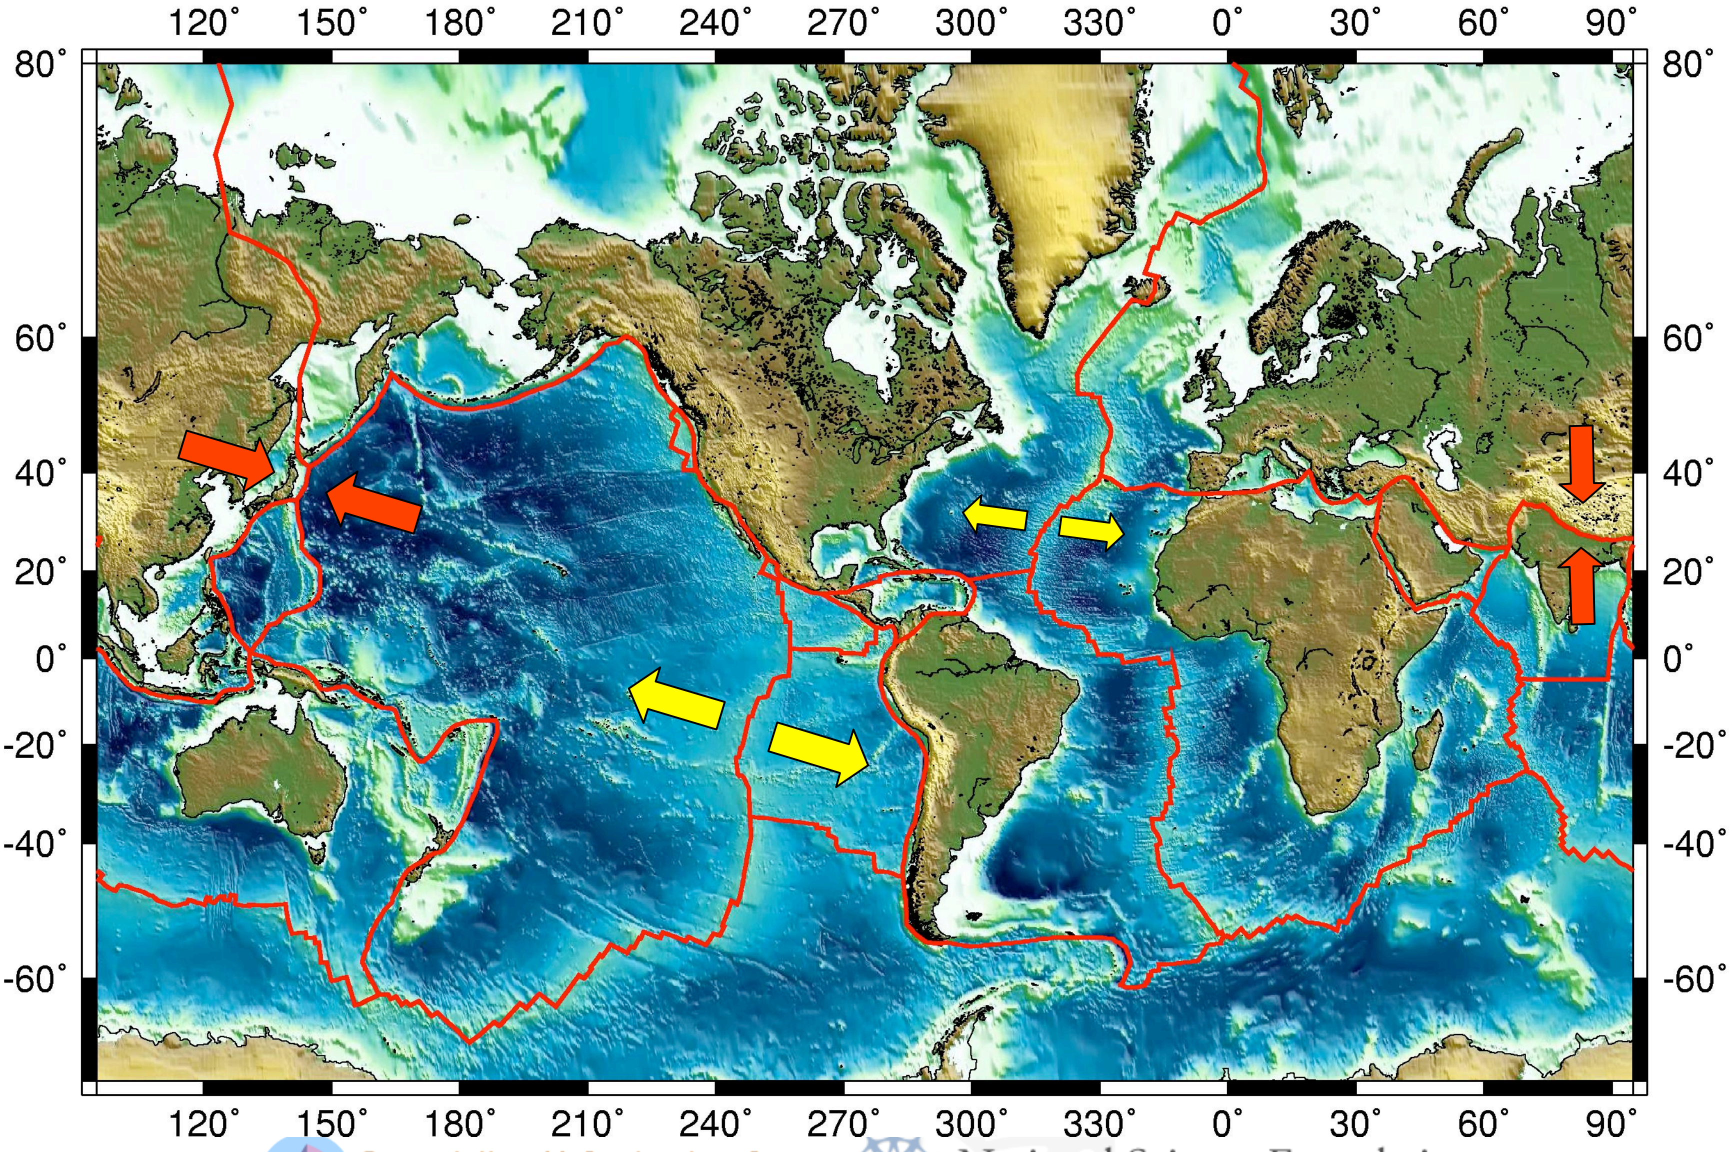

Computational Infrastructure for Geodynamics

National Science Foundation
WHERE DISCOVERIES BEGIN

Mantle Convection - Plate Tectonics

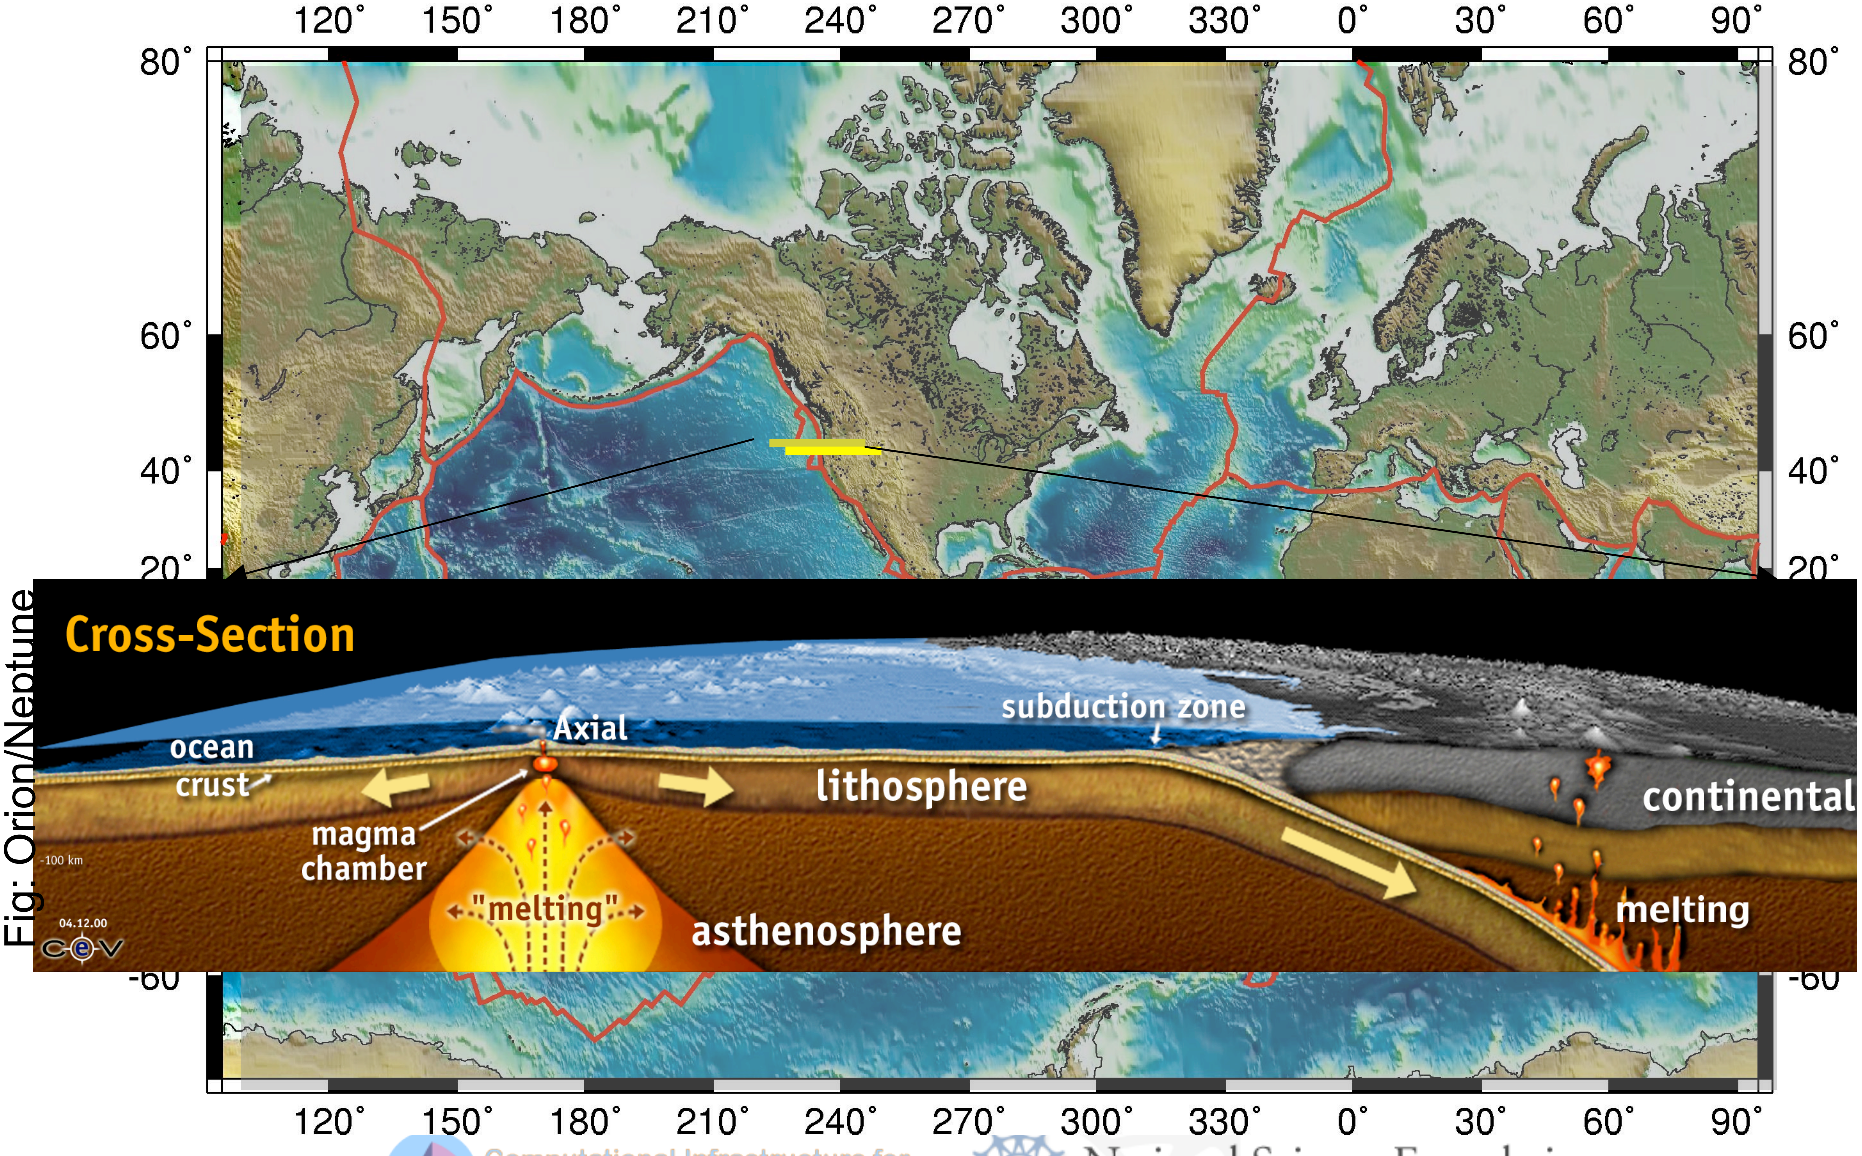

Fig: Orion/Neptune

Integrated models of plate boundaries

Tan et al.

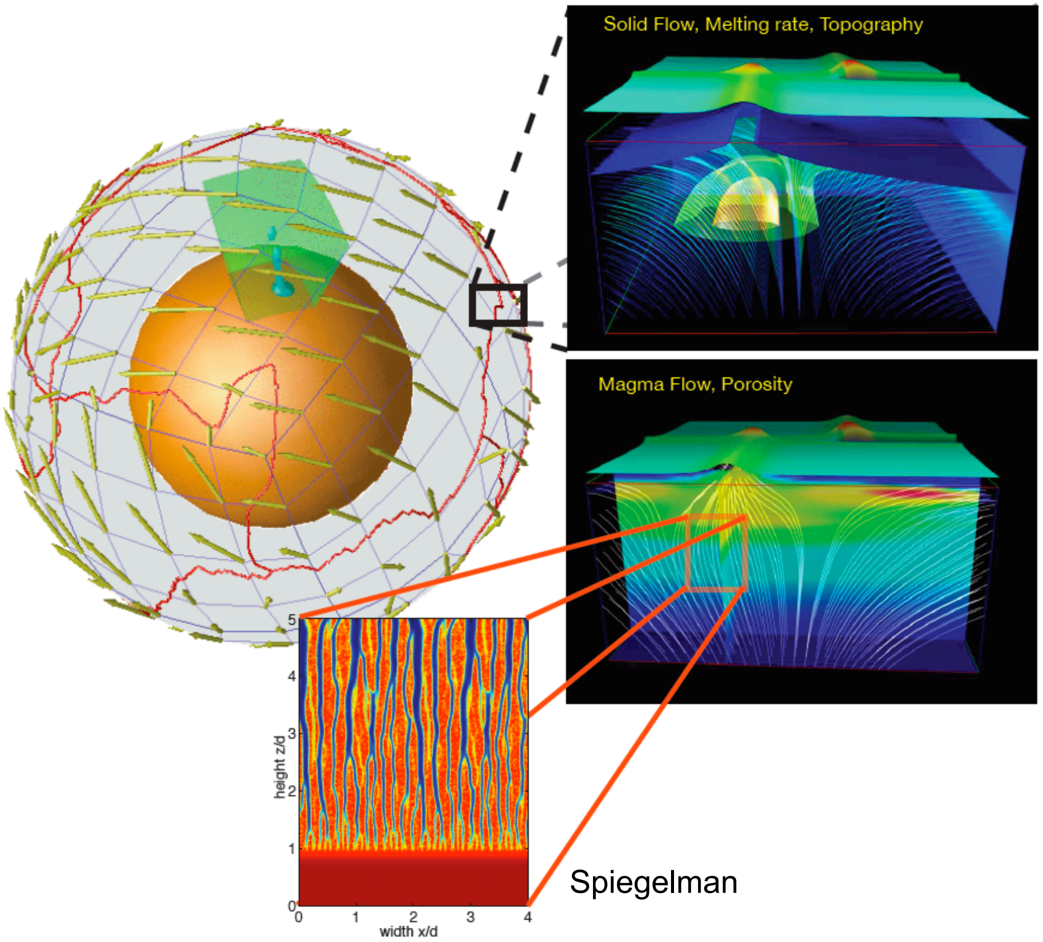

Spiegelman

Spiegelman

Mantle Convection-Plate tectonics: Earth Evolution

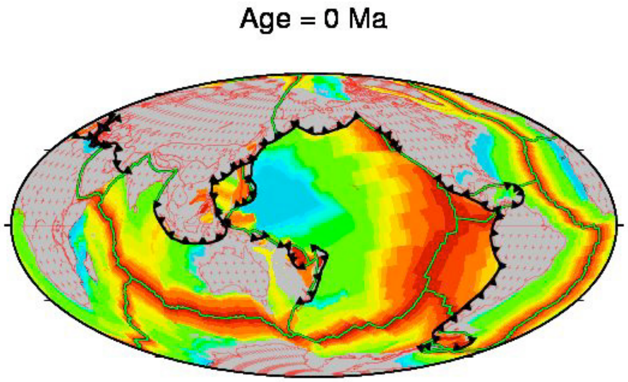

Gurnis et al.[2007]

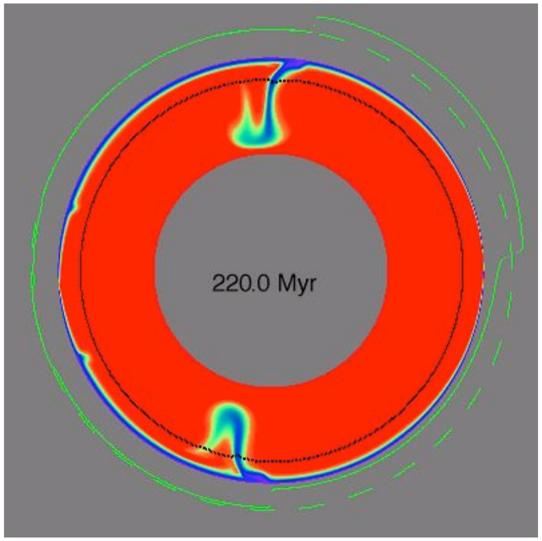

Forward

Inverse

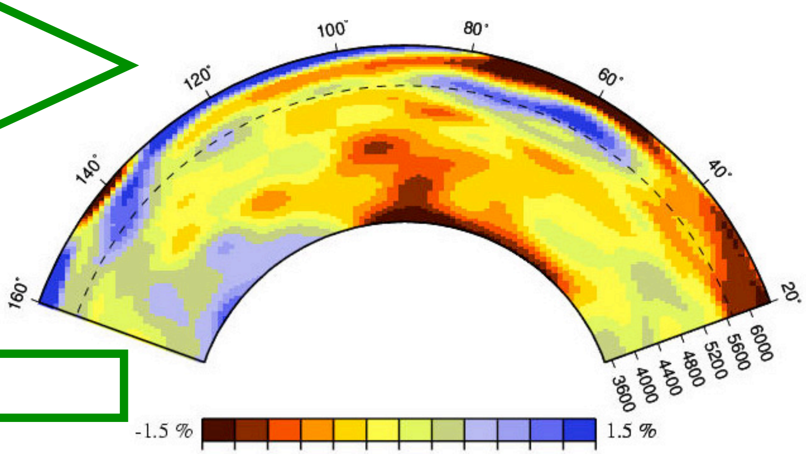

[Zhong & Gurnis]

S20RTS [Ritsema et al.]

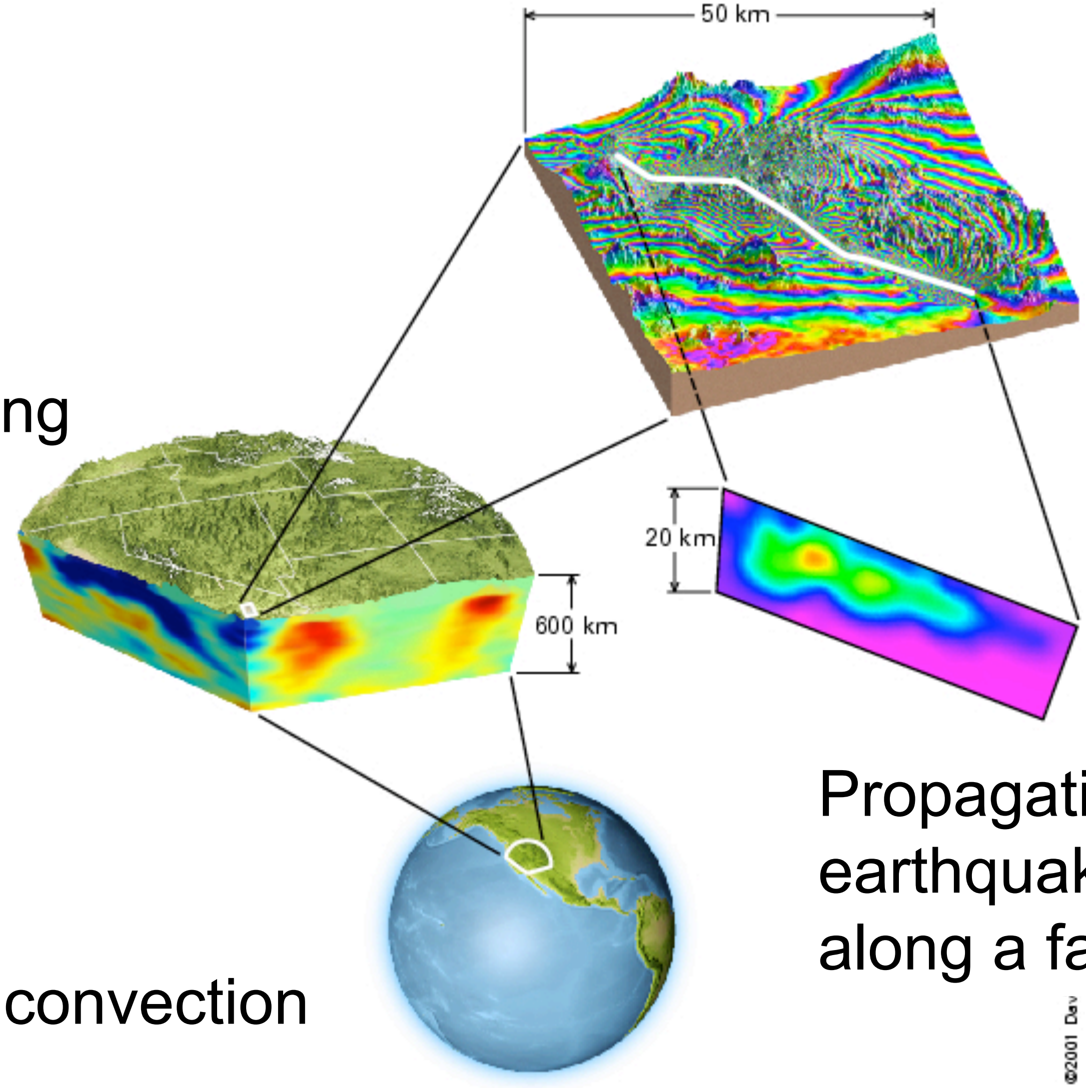

Faults

Ground shaking

Mountain building

Fault zones

Propagation of an earthquake rupture along a fault

Mantle convection

Long-term plate motions

Achievement Highlights

Mantle Convection
CitcomS3.0

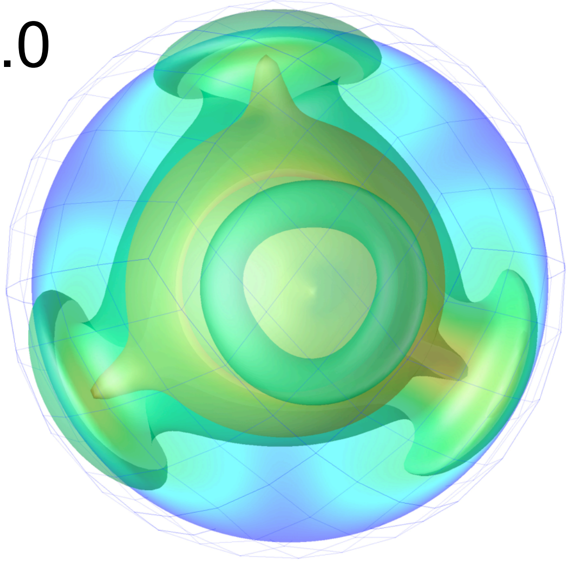

Earthquake cycle
PyLith

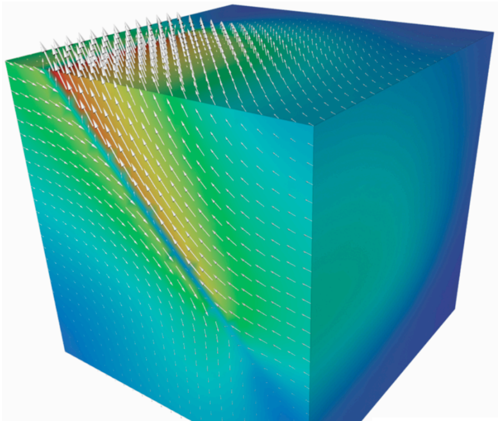

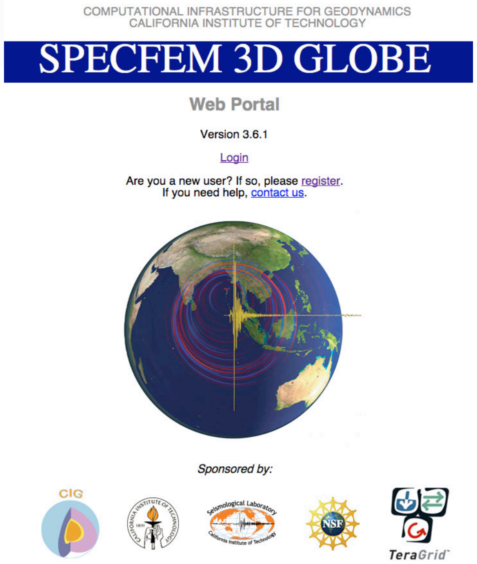

COMPUTATIONAL INFRASTRUCTURE FOR GEODYNAMICS
CALIFORNIA INSTITUTE OF TECHNOLOGY

SPECFEM 3D GLOBE

Web Portal

Version 3.6.1

Login

Are you a new user? If so, please register.
If you need help, contact us.

Sponsored by:

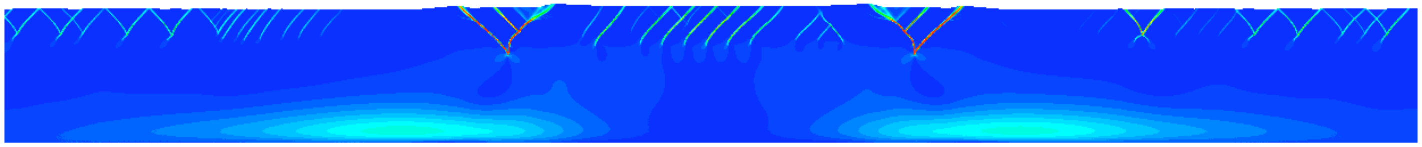

Long-time
scale tectonics: GALE

Seismology
Science Gateway

CitcomS 3.0: Compressible thermo-chemical convection with multiple equations of state

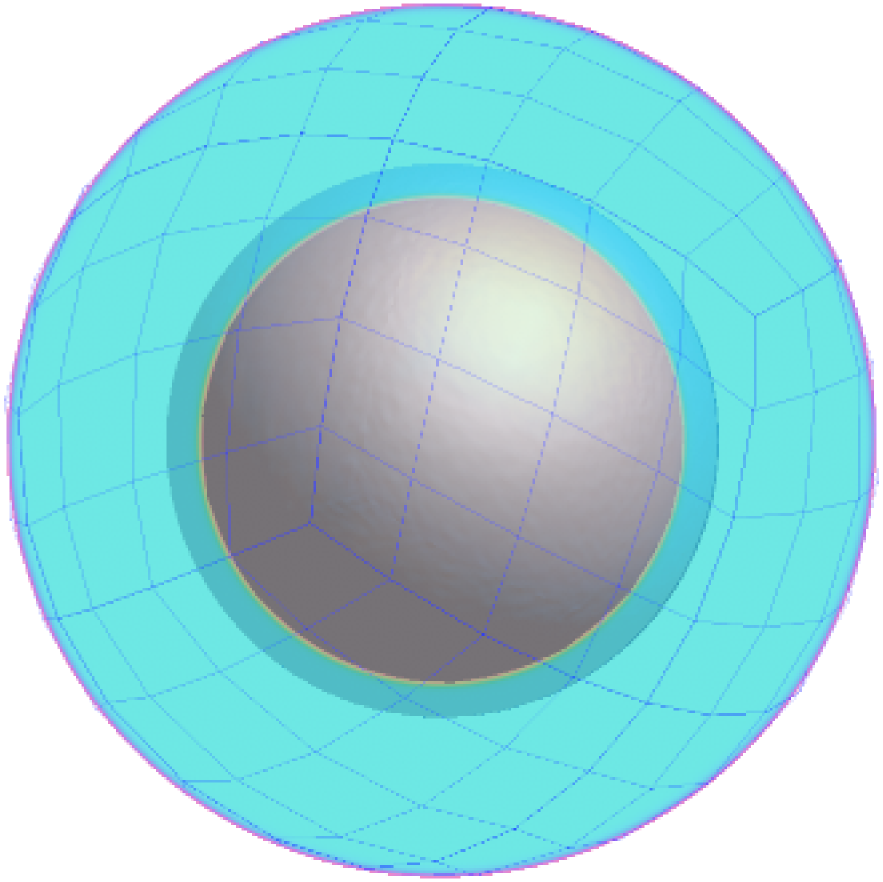

CIG Partners with ‘science neutral’ CS researchers

- Argonne National Laboratory and support of the PETSc team: Developed PyLith code with Sieve as a spin-off
- Center for Advanced Computer Research (CACR), Caltech: Applying & enhancing Python superstructure framework (Pyre)
- Victorian Partnership for Advanced Computing -- a not-for-profit software engineering company -- developed GALE
- University of Texas (ICES) Developing Peta-scale AMR-mantle convection code: possibility of a generic variable viscosity Stokes solver
- Texas A&M University Dept. of Mathematics. Development of geodynamics test suite for with generic AMR tool-kit: deal.II [in development]

Earthquake Cycle Community

A Vibrant community: many young scientists, a vision that capitalizes on EarthScope and with clear needs for modern computation tools

Create flexible, computationally efficient software for simulation of crustal deformation across spatial scales ranging from meters to hundreds of kilometers and temporal scales ranging from milliseconds to hundreds of years.

Supported "Community Finite Element Modeling" workshops for the last five summers (the last three with major CIG funding)

PyLith

CIG Developed a new package between members of the community, Brad Aagaard (USGS) and Charles Williams (RPI), with direct involvement of a generic computational science team (PETSc at ANL).

- Modular: Users can swap modules to run the problem of interest
- Scalable: Code runs on one to a thousand processors efficiently
- Extensible: Expert users can add functionality to solve their problem without polluting main code

Spinoff:
The team needed a Parallel FE Object for both parallel storage and manipulation.
In that PETSc is an abstraction for Linear Algebra, **Sieve** is an abstraction for FEM objects

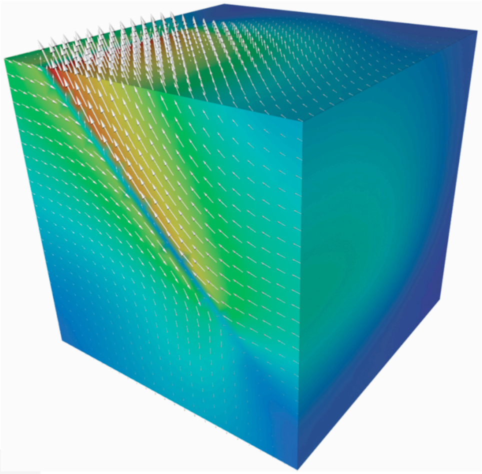

The coming golden age of computational sciences

- The 21st century will witness breakthroughs in addressing many of the scientific and societal grand challenge problems
- These grand challenge problems cannot be addressed without mathematical modeling and computer simulation
- The golden age is being enabled by:
    - Greater fidelity of mathematical models
    - Explosion of validation-quality observational data
    - Advances in computing systems, software, algorithms
- But fundamental challenges must be overcome:
    - Continued improvements in models of relevant phenomena
    - Advances in algorithms and software for complex models

Modeling challenges for geophysics

- Complex nonlinear phenomena
- Multiphysics couplings
- Wide range of length and time scales
- Highly anisotropic and heterogeneous media
- Complex geometries
- Phenomena only indirectly observable

Computational challenges for geophysics

- stable and accurate discretization for multiphysics problems
- systematic “upscaling” for multiscale problems
- scalable linear and nonlinear solvers capable of dealing with anisotropic, heterogeneous, coupled problems
- methods for 4D data assimilation and inverse problems
- uncertainty quantification and propagation methods for expensive simulations
- adaptive mesh refinement/coarsening methods
- geologically-aware mesh generation
- ultra-scale visualization
- numerical/geometric algorithms that scale to 10^5 cores
- algorithms for exploiting many-core processors
- fault-tolerant algorithms

Challenges for CIG

- CIG is leading the geophysics community into the modern age of modular, portable, extensible, robust, scalable open-source community software
    - Refactoring of existing research codes
    - Development of clean-sheet-of-paper codes
- Not just a software engineering effort: scientific computing research issues have been tackled along the way
- Much of the low-hanging fruit has been plucked
- The remaining challenges include many of the most difficult ones facing computational science
- These challenges are particularly pervasive and severe in geophysics problems, and require problem-driven solutions

Two examples of advances at the interfaces of computational geodynamics and scientific computing

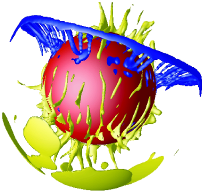

- New generation mantle convection code capable of scaling to petascale systems
- Aimed at first global mantle convection simulations that can resolve faulted plate boundaries
- Collaboration among computational geophysics, computational math, computer science (NSF PetaApps OCI-0749334; PIs G. Biros, O. Ghattas, M. Gurnis, S. Zhong)
- Two examples from project (C. Burstedde, G. Stadler, T. Tu, L. Wilcox)
  - Example 1: Parallel octree-based dynamic adaptive mesh refinement
  - Example 2: Scalable variable-viscosity Stokes solver for highly heterogeneous media

Parallel octree-based dynamic adaptive mesh refinement

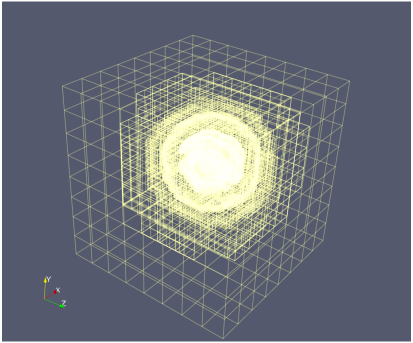

Adaptively refined mesh

- AMR essential for resolving physical phenomena that vary over a wide range of spatial scales
- $O(10^3)$ reduction in unknowns for mantle convection simulations that resolve faulted plate boundaries

Parallel octree-based dynamic adaptive mesh refinement

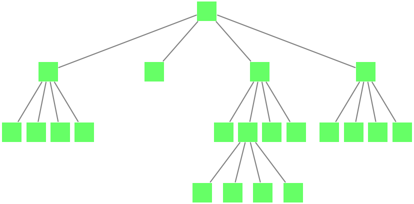

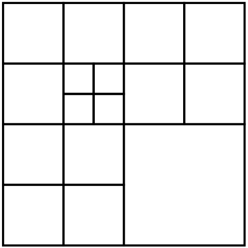

- Octree leaves map to hexahedral finite elements
- Octrees facilitate treatment of local adaptivity
- Tradeoff between numerical and geometric complexity

Parallel octree-based dynamic adaptive mesh refinement

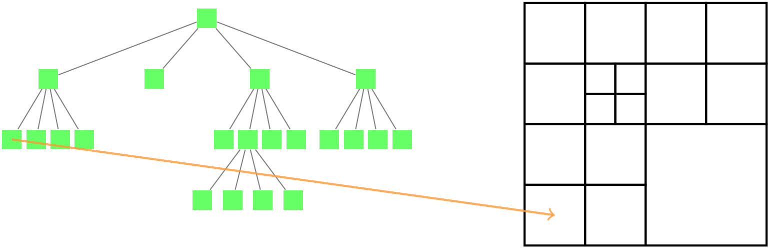

- Octree leaves map to hexahedral finite elements
- Octrees facilitate treatment of local adaptivity
- Tradeoff between numerical and geometric complexity

Parallel octree-based dynamic adaptive mesh refinement

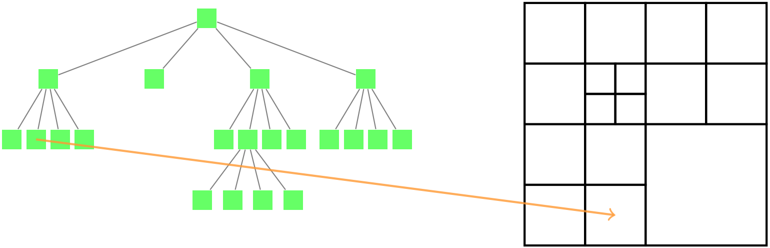

- Octree leaves map to hexahedral finite elements
- Octrees facilitate treatment of local adaptivity
- Tradeoff between numerical and geometric complexity

Parallel octree-based dynamic adaptive mesh refinement

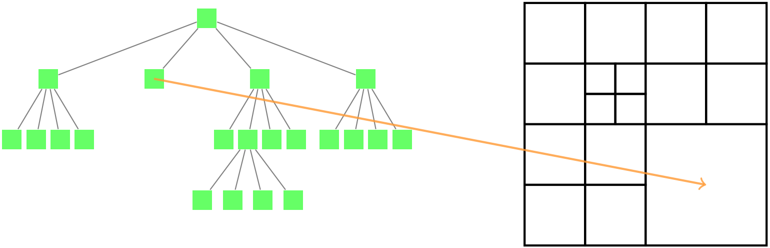

- Octree leaves map to hexahedral finite elements
- Octrees facilitate treatment of local adaptivity
- Tradeoff between numerical and geometric complexity

Parallel octree-based dynamic adaptive mesh refinement

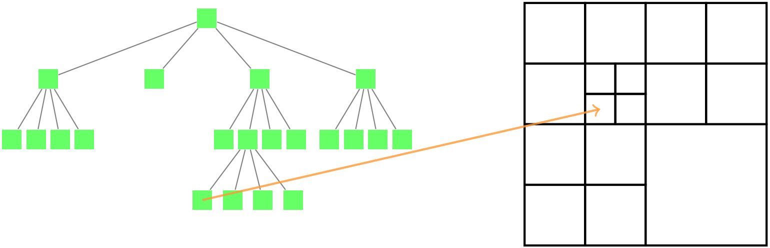

- Octree leaves map to hexahedral finite elements
- Octrees facilitate treatment of local adaptivity
- Tradeoff between numerical and geometric complexity

Parallel octree-based dynamic adaptive mesh refinement

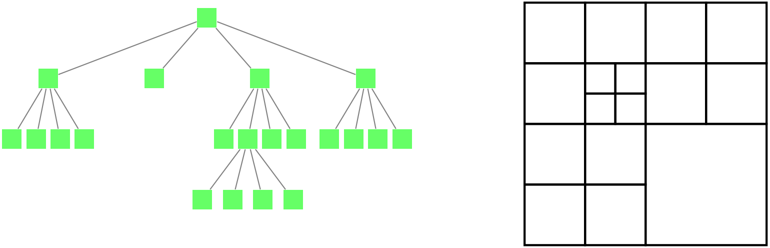

- Octree decomposed according to Z-ordered space-filling curve
- Underlying mesh operations parallelized according to the octree partitioning
- Locality-preserving property leads to good parallel partitioning, good cache performance

Parallel octree-based dynamic adaptive mesh refinement

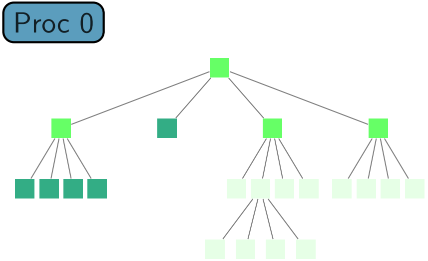

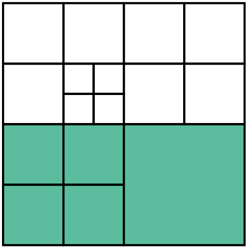

- Octree decomposed according to Z-ordered space-filling curve
- Underlying mesh operations parallelized according to the octree partitioning
- Locality-preserving property leads to good parallel partitioning, good cache performance

Parallel octree-based dynamic adaptive mesh refinement

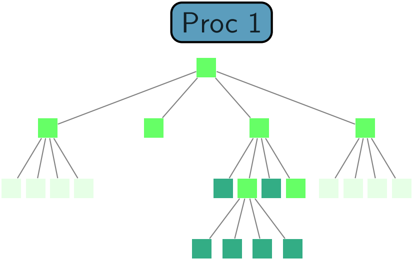

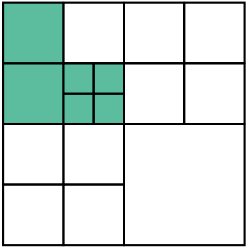

- Octree decomposed according to Z-ordered space-filling curve
- Underlying mesh operations parallelized according to the octree partitioning
- Locality-preserving property leads to good parallel partitioning, good cache performance

Parallel octree-based dynamic adaptive mesh refinement

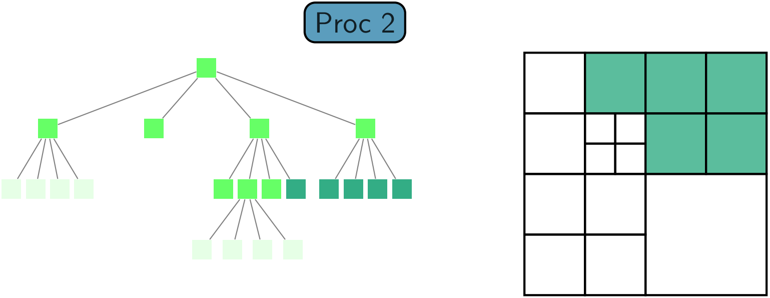

- Octree decomposed according to Z-ordered space-filling curve
- Underlying mesh operations parallelized according to the octree partitioning
- Locality-preserving property leads to good parallel partitioning, good cache performance

Parallel octree-based dynamic adaptive mesh refinement

- For time-dependent problems the mesh needs to be adapted dynamically
- Mesh adaptation components must be as scalable as the solver
- Dynamic load balancing problem

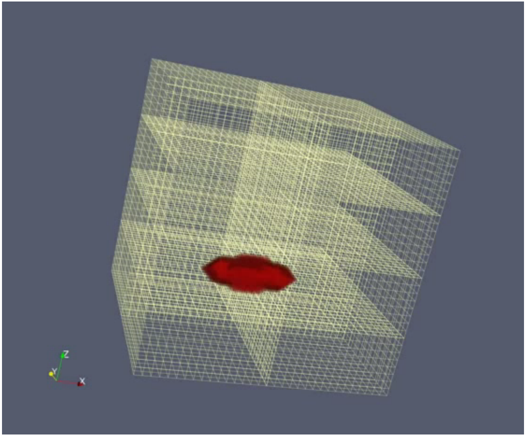

Parallel AMR performance for convection–diffusion

Wall clock time breakdown demonstrates excellent scalability of AMR components

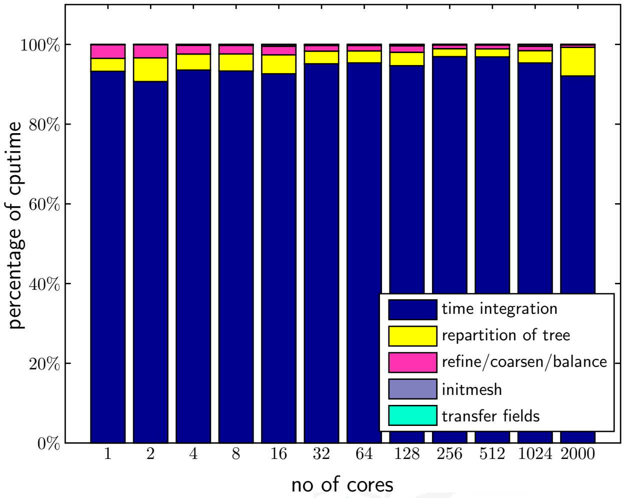

Parallel AMR performance for convection–diffusion

Excelent overall efficiency for weak scaling from 2 to 2000 cores

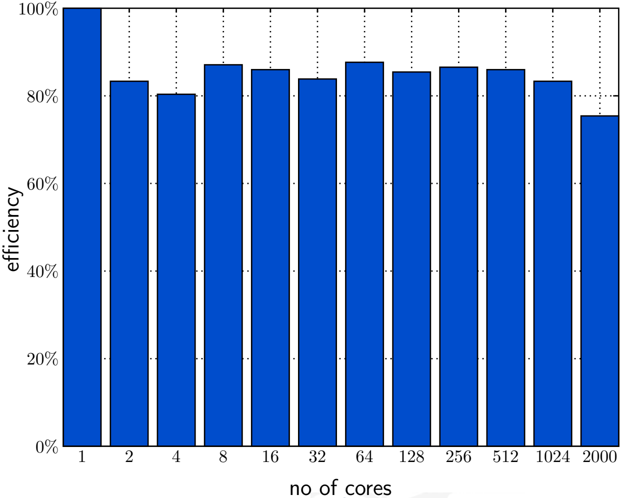

Variable-viscosity Stokes solver for mantle convection

Also a key component of several other solid earth geophysics problems

$$\frac{\partial T}{\partial t} + u \cdot \nabla T - \nabla^2 T - \gamma = 0 \tag{AD}$$

$$\nabla \cdot \Big[\big[\eta(T)(\nabla u + \nabla^\top u)\big) \Big] - \nabla p + \mathsf{Ra}\, T e_r = 0 \tag{S1}$$

$$\nabla \cdot u = 0 \tag{S2}$$

Variables:

- T . . . temperature
- u . . . velocity
- p . . . pressure

Parameters:

- Ra $\sim 10^6 - 10^9$ . . . Rayleigh number
- γ . . . heat production rate
- $\eta(T) \cong \eta_o \exp(1 - E_o T)$ . . . viscosity
- e_r . . . radial direction

Variable-viscosity Stokes solver for mantle convection

- Trilinear FEM for temperature, velocity, and pressure (the framework allows the use of higher-order FEM)
- Conforming approximation is enforced by algebraic elimination of hanging nodes
- Solver operates only on anchored nodes
- FEM stabilization:
    - Streamline Upwind/Petrov Galerkin (SUPG) for advection-diffusion system
    - Polynomial pressure projection for stabilization of Stokes equation
- *hypre* AMG/F_p-preconditioned MINRES for Stokes solve
- α-time-stepping for advection-diffusion equation

Variable-viscosity Stokes solver for mantle convection

Excellent algorithmic scalability for variable-coefficient Stokes solver

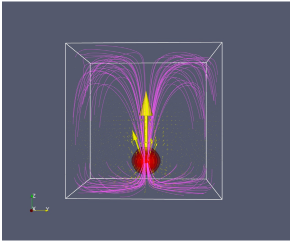

- Krylov solver: MINRES
- Preconditioner: One V-cycle of algebraic multigrid on approximately factored Stokes system (F_P approx of Schur complement) using *BoomerAMG* from *hypre*

no. of cores:	1	8	64	512
no. of elements:	61.4K	491K	3.9M	31.5M
no. of DOFs:	420K	2.5M	26.2M	204.5M
no. of iter:	26	33	39	45
solve time (sec):	73.8	128.4	170.5	327.6

Number of MINRES iterations for Stokes solve almost insensitive to 500X increase in problem size, $O(10^3)$ variation in viscosity, and 4 levels of mesh refinement

Computational Infrastructure for Geodynamics

National Science Foundation
WHERE DISCOVERIES BEGIN

The path forward for CIG

- CIG must broaden its mission and agenda, and engage a broader community of applied mathematicians, computer scientists, and scientific computing experts in the geophysics big science problems
- CIG has begun to partner with the wider computational science community to address modeling and computational challenges
- First steps: "computational geophysics meets scientific computing" workshop held at UT-Austin in Oct. 2006
    - Brought together two communities to identify and assess challenges and opportunities at the interfaces of applied math/CS and frontier problems in short and long term tectonics, seismology, mantle convection, magma migration, and geodynamo

Conclusions

- Has never been a more exciting time in computational geophysics
- Tremendous opportunities to advance our understanding of the earthquake cycle and of the dynamics of the Earth as a whole
- Enabled by rapidly-expanding cyberinfrastructure (including petascale computing) and explosion of observational data
- Manifold challenges lie ahead (identified by workshop):
    - geology-aware mesh generation
    - mesh adaptivity for strongly multiscale problems
    - solvers/preconditioners for strong heterogeneities/anisotropies
    - nonlinear solvers and time integrators for multiphysics problems
    - data management and sci vis for complex geometries
    - scalable parallel algorithms
    - data assimilation and inverse methods
- Deep collaborations *essential* between geophysics and applied math/CS; CIG already beginning to engender these