Parallel Adaptive Mantle Convection Simulation

Carsten Burstedde*, Omar Ghattas*, Mike Gurnis†, Georg Stadler*, Eh Tan‡, Tiankai Tu*, Lucas C. Wilcox*, Shijie Zhong§

* Institute for Computational Engineering and Sciences (ICES), The University of Texas at Austin
† Seismological Laboratory, California Institute of Technology, Pasadena, California
‡ Computational Infrastructure for Geodynamics, California Institute of Technology, Pasadena, California
§ Department of Physics, University of Colorado, Boulder, Colorado

Adaptive mantle convection simulation on supercomputers

Our goal is to conduct global mantle convection simulations that can resolve faulted plate boundaries, down to 1 km scales. Uniform resolution leads to trillion element meshes, which are intractable even on petascale supercomputers. Thus parallel mesh adaptivity is essential.
We present Rhea, a new generation mantle convection code designed to scale to hundreds of thousands of cores. Rhea is built on ALPS, a parallel octree-based adaptive finite element library that supports new distributed data structures and parallel algorithms for dynamic coarsening, refinement, rebalancing, and repartitioning of the mesh.

1. Mantle convection

Model equations

The dynamics of mantle convection are governed by equations for the conservation of mass, momentum, and energy. A simplified version reads:

$$\frac{\partial T}{\partial t} + u \cdot \nabla T - \nabla^2 T - \gamma = 0 \qquad \text{(AD)}$$

$$\nabla \cdot \left[\eta(T)\left(\nabla u + \nabla^\top u\right)\right] - \nabla p + Ra\, T e_r = 0 \qquad \text{(S1)}$$

$$\nabla \cdot u = 0 \qquad \text{(S2)}$$

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

The need for mesh adaptation

Resolving phenomena at faulted plate boundaries requires ~ 1km resolution. On a uniform mesh of the Earth, this results in $\sim 10^{12}$ elements, well beyond the reach of even petascale supercomputers. Adaptive mesh refinement places resolution only where needed, resulting in 3 orders of magnitude reduction in number of elements. Below, we show snapshots of the thermal field T at three time instants (left) and corresponding adapted meshes (right). The mesh adaption, which is based on error indicators, resolves the rising plumes and the instabilities at the top layer.

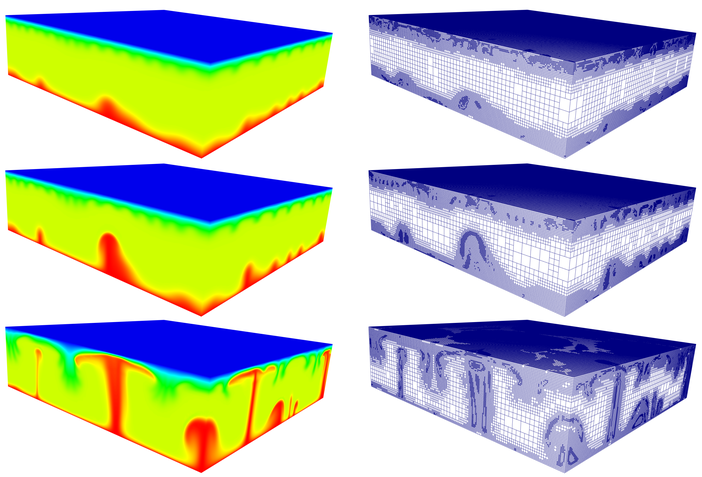

Operator-split time integration

The coupled system (AD), (S1), (S2) is split into an advection-diffusion time step, followed by a variable-viscosity Stokes solve to update the flow field. The mesh is adaptively refined, coarsened and redistributed every k_o timesteps.

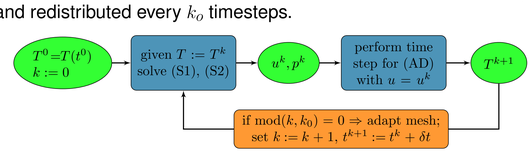

2. Parallel adaptive mesh refinement/coarsening (AMRC)

Parallel octree-based adaptive meshes

Octree-based adaptive finite element meshes are used to discretize the mantle convection equations. The mesh structure is dynamically partitioned to processors after each AMRC step according to a space filling curve (SFC). The locality-preserving property of SFCs leads to good parallel partitioning and good cache performance.

On the right we show the Z-ordered space filling curve leading to a one-to-one correspondence of octree leaves and finite elements and to a partitioning of the mesh among three processors.

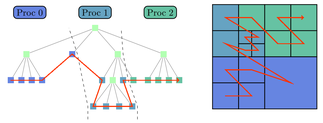

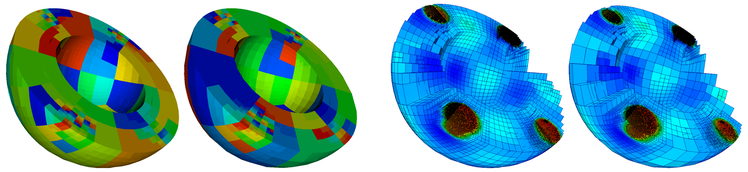

Above we show a partitioning of a spherical shell on 1024 cores for two time steps using the forest-of-octree library P4EST (left) and the corresponding adapted meshes and temperatures (right). Time evolution uses the high-order discontinuous Galerkin library MANGLL.

Parallel scalability of ALPS

Parallel scalability for adaptive solve of the advection-diffusion equation (no Stokes solve). Fixed-size scalability (right): Speedups based on total runtime plotted against the number of cores for four different problems. Weak scalability (below): On the left is a breakdown of total run time into components related to numerical PDE integration (blue) and AMRC functions (all other colors), with increasing number of cores from 1 to 62,464.

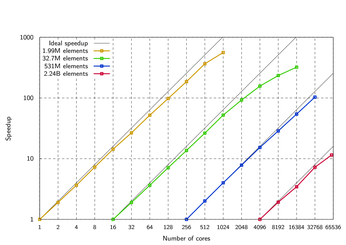

Problem size increases isogranularly at roughly 131,000 elements per core (largest problem has approximately 7.9 billion elements). Overall, AMRC consumes about 10% or less of the run time. On the right is the parallel efficiency measured in total processed elements per core per total run time. Despite a 62K-fold increase in problem size and number of cores, parallel efficiency remains above 50%.

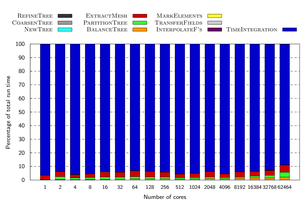

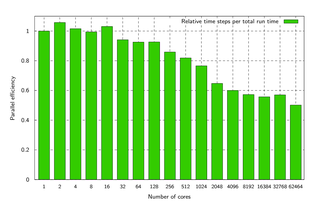

3. Scalable variable viscosity Stokes solver

The Stokes equations (S1–S2) are discretized by trilinear hexahedra for both velocity and pressure. Stabilization is provided by polynomial pressure projection. To address the heterogeneities stemming from variable viscosity and mesh size, we employ a preconditioned MINRES solver. The preconditioner is one algebraic multigrid (AMG) V-cycle on an approximately factored form of the Stokes system, with an F_P-type approximation of the pressure Schur complement. The AMG solver is *BoomerAMG* from the LLNL *hypre* package.

#cores	#elem	#elem/core	#dof	MINRES #iterations
1	67.2K	67.2K	271K	57
8	514K	64.2K	2.06M	47
64	4.20M	65.7K	16.8M	51
512	33.2M	64.9K	133M	60
4096	267M	65.3K	1.07B	67
8192	539M	65.9K	2.17B	68

The table demonstrates that the number of MINRES iterations and the cpu time for the Stokes solve is almost insensitive to a 8192$\times$ increase in problem size, $O(10^3)$ variation in viscosity, and 4 levels of non-uniform mesh refinement.

4. Adaptive mantle convection simulation

Parallel scalability

We present end-to-end simulations that resolve the range of length/time scales that we expect to be present in global models of mantle convection. Overall weak scalability shows that the time is dominated by the Stokes solve. The overhead due to mesh adaptation is negligible. The overall efficiency on 16,384 cores is about 25% compared to 1 core.

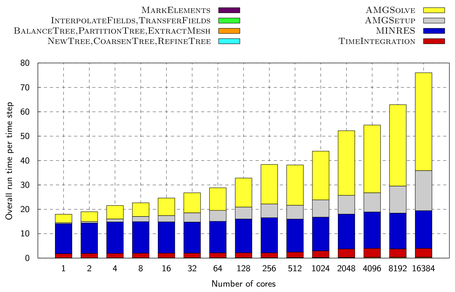

A yielding problem

Next we adopt a viscosity law that displays yielding under high stress. This simplified approach dynamically achieves large-scale motions with narrow zones of low viscosity such as those associated with plate tectonics and faulted plate boundaries. The model uses a temperature- and pressure-dependent viscosity that yields under high deviatoric stress according to

$$\eta = \begin{cases} \min\left\{10\exp(-6.9T), \frac{\sigma_y}{2\dot{\epsilon}}\right\}, & z > 0.9 \\ 0.8\exp(-6.9T), & 0.9 \geq z > 0.77 \\ 50\exp(-6.9T), & z \leq 0.77 \end{cases}$$

where σ_y is the yield stress and $\dot{\epsilon}$ is the second invariant of the deviatoric strain rate tensor. Here, the viscosities range over four orders of magnitude.

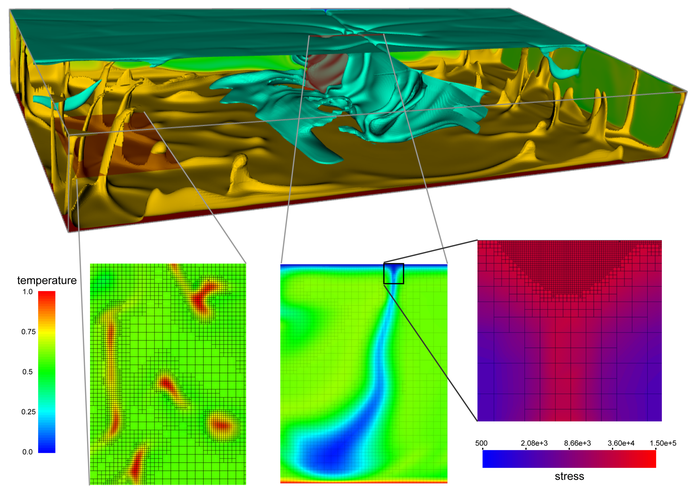

Mantle convection with yielding. Top: Temperature isosurfaces at $T = 0.3$ (cyan), $T = 0.8$ (orange). Bottom left: Horizontal slice of temperature showing that grid refinement follows the temperature gradient. Bottom center: Vertical slice of temperature showing the downwelling slab and the yielding zone. Bottom right: A zoom-in of the yielding zone, where the finest grid of ~ 1.5 km resolution covers the region of highest stress.

Acknowledgement

This work was partially supported by NSF's PetaApps program (grants OCI-0749334, OCI-0749045, and OCI-0748898), NSF Earth Sciences (EAR-0426271), DOE Office of Science's SciDAC program (grant DE-FC02-06ER25782) and NSF grants CNS-0619838, and DMS-0724746. We acknowledge many helpful discussions with the *hypre* developers Rob Falgout and Ulrike Yang, and with George Biros. We thank TACC for their outstanding support, in particular Bill Barth, Tommy Minyard, Karl Schulz and Romy Schneider.