coupling of plate tectonics to mantle convection is also a major focus of this project.

CORE CONVECTION: THE GEODYNAMO

The generation of Earth's magnetic field by magnetohydrodynamic convection of liquid iron in the outer core is a challenging problem both scientifically and computationally. Indeed, until recently, all models were kinematic, i.e. some component of either the magnetic field or the velocity field had to be prescibed. A major breakthrough came with the achievement of a completely self-consistent 3-D spherical dynamo model, using the DYNAMO code written by Gary Glatzmaier (Glatzmaier and Roberts 1995a; Glatzmaier and Roberts 1995b; Glatzmaier and Roberts 1996b). It is this model which is being further developed in this project.

DYNAMO Numerical Method and Parallelization

The DYNAMO code uses a spectral transform method, explanding variables in spherical harmonics azimuthally and Chebychev polynomials in radius, giving better radial resolution in the boundary layers where it is most needed. The equations are solved and timestepped in spectral space, while nonlinear products are calculated in grid space, requiring the transformation of fields and derivatives from spectral to grid, and nonlinear products from grid to spectral, every time step. Since a particular grid point is influenced by all spectral modes, the transforms have a global data dependency and efficient parallelization of such a method presents a significant problem. However, after considerable work by Tom Clune, Gary Glatzmaier, and Dan Katz of JPL, an efficient parallelization has been achieved which allows the code to be run on hundreds of nodes of a parallel supercomputer such as the Cray T3E.

The key to this efficient parallelization is the transpose method used in the transforms. In spectral space and grid space, an azimuthal decomposition is used, in which spectral modes or azimuthal regions are distributed across the nodes, with nodes containing all the radial information for the modes or azimuthal regions they contain. This decomposition was recently changed from 1-D to 2-D. A transform from spectral to grid space involves a Legendre transform in latitude, followed by a Fast Fourier transform in longitude. Prior to each transform, data is transposed across the nodes such that no communication is required during the transform, allowing them to proceed at a speed which is limited only by CPU performance. This method appears to be more efficient than using cross-node Legendre and Fast Fourier transforms, an approach which was taken previously for a related spectral transform code used to model mantle convection (Tackley et al. 1993).

Another major modification was to use a rhomboidal, instead of triangular, truncation of spherical harmonics in the horizontal expansions of the variables. Since the implicit treatment of the Coriolis force requires all spherical harmonic degrees for a given spherical harmonic order, to be solved simultaneously and therefore (most easily) on a single processor, the rhomboidal truncation, which uses the same number of degrees per order, provides much better load balance on the T3E. It also appears to be better for the (rotationally dominated) dynamo problem because it provides the same latitudinal resolution for all longitudinal modes (i.e., orders).

Considerable work has gone in to optimizing the reuse of cache on the DEC alpha CPUs. Modifications include reordering the data storage to be cache-friendly (done simultaneously with the parallel transpose operations), using level-3 BLAS for Legendre transforms and vendor-supplied libraries for the FFT, redistributing the data for the implicit solves so that several vectors can be computed on at once, and 'cache-aligning' the primary arrays. Several performance enhancements have also been implemented to take advantage of Cray T3E specific features, such as using stride-1 SHMEM calls wherever possible, and using E-registers (special hardware features on the T3E that allow the CPU to bypass cache for regularly strided data.) to do local transposes of data.

Performance

DYNAMO now demonstrates excellent scalability, and performance of up to 150 Gflops with 114 million grid points on a 512-node Cray T3E-600 (i.e., with 300 MHz CPUs), as summarized Table 1. Based on preliminary tests it is estimated that a performance of 650 Gflops should be acheived on a 1488 node Cray T3E-1200.

GRID			nproc	Gflops
nr	nlat	nlong		
129	384	384	128	25.4
			256	49.1
			512	94.6
129	768	768	256	62.8
			512	122.0
129	1150	768	512	150.0

Table 1: Performance of DYNAMO on a Cray T3E-600 on various numbers of processors nproc. 'nr', 'nlat' and 'nlong' are the number of grid points in the radial, latitudinal and longitudinal directions respectively.

Scientific Results

Many scientific breakthroughs have been produced using DYNAMO, including the first self-consistent geodynamo simulation (Glatzmaier and Roberts 1995a), the first simulation of a geomagnetic field reversal (Glatzmaier and Roberts 1995b), and the discovery of differential rotation of the solid inner core, i.e., the inner core rotates at a slighly different rate from the Earth's surface (Glatzmaier and Roberts 1996b). Presently, several studies are being persued by project team members Gary Glatzmaier, Paul Roberts and Peter Olson, and in collaboration with Ulrich Christensen at the University of Goettingen, Germany (e.g., (Christensen, Olson, and Glatzmaier 1998; Glatzmaier and Roberts 1996a; Glatzmaier and Roberts 1997; Olson and Glatzmaier 1996). Here we briefly discuss three examples of present studies.

Core/Mantle Coupling- Earth. Perhaps the most important mode of mantle:core coupling is thermal, in which the thermal structure of the lower mantle determines the heterogeneous pattern of heat flux out of the core at the core-mantle boundary (CMB). Simulations are being performed to determine how the pattern of heat flux affects the frequency of magnetic dipole reversals. So far, seven different simulations, all identical except for the imposed CMB heat flux pattern, have been run at least 100 kyr (five magnetic diffusion times), and some for well over 200 kyr. All but one reversed at least once. Figure 1 shows an example. The only case not to reverse has a heat flow pattern that is axisymmetric and symmetric with respect to the equator with peak CMB heat flux occurring at the poles and at the equator and minimum heat flux at mid latitude. We speculate that a heat flow pattern with this type of symmetry may have controlled the geomagnetic field during the so-called superchrons, constant polarity epochs that lasted tens of millions of years with apparently no reversals.

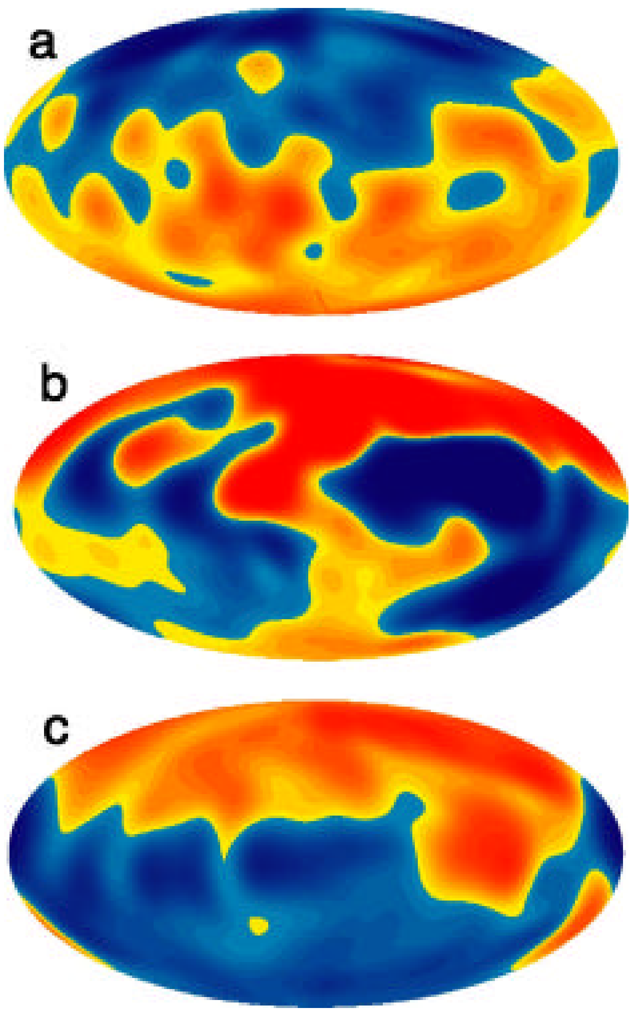

Figure 1. Three snapshots, separated by about 3 kyr, of the radial component of the simulated geomagnetic field over the CMB during one of the dipole reversals. Red (blue) represents outward (inward) directed field.

Core-Mantle Coupling- Io. Jupiter's moon Io is a prime example of a body for which mantle-core coupling may play an important role. Recent measurements by the Galileo spacecraft indicate that Io has an active component to its magnetic field. Circulation in the core is probably driven by differential heating and cooling from the mantle. Io's mantle is particularly interesting, being unique among planetary mantles so far considered because convection is driven by non-uniform tidal dissipation rather than uniform radiogenic heating or an isothermal hot lower boundary. Tackley, Glatzmaier, and Gerald Schubert at UCLA have been simulating mantle convection in Io. Results to date indicate a mean flow driven by the distribution of tidal heating, with small-scale instabilities superimposed on the mean flow. As Rayleigh number is increased these small-scale instabilities increasingly distribute the heat. Patterns of heat flux across the core-mantle boundary are being used as a boundary condition for core dynamo simulation. The resulting spherically-symmetric part of the temperature gradient is subadiabatic, i.e., convectively stable in the traditional sense. However, the strongly variable flux at the CMB locally cools or heats the core fluid just below the CMB enough to drive thermal convection in the core. The interaction of this motion with the basic planetary rotation mechanisms maintains zonal shear and helical flows that induce a complicated, time-dependent magnetic field from the simple, constant, Jovian background field.

Ancient Earth. Gary Glatzmaier and Paul Roberts have applied numerical dynamo models to the ancient Earth, at a time when the crystallizing inner core was only one quarter of its present size. The results are compared to the corresponding model with the present-day inner core size. The models dynamo models are driven by heat coming partially from the latent heat released as core fluid freezes onto the inner core boundary, partially from the internal energy of the core itself, plus the buoyancy produced as the light components of core fluid segregate during that freezing. Because the inner core plays a much smaller role in the ancient Earth model, dynamo action is less efficient, and requires greater heat to be extracted from the core, compared to the present-day configuration. This has significant implications for Earth history, for it is known that the Earth has possessed a magnetic field over most of its existence. Our findings make it harder than before to explain this.

MANTLE CONVECTION

Although solid rock, Earth’s 2900 km deep mantle convects by viscous creep at effectively infinite Prandtl number, at velocities of order centimeters per year. This convection is intimately linked with plate tectonics; indeed, plate tectonics and mantle convection are best viewed as different aspects of the same, coupled system. Despite the low velocities, the Rayleigh number is of order 10^7, and thus convective features are small compared to the size of the domain, and the convection is quite time-dependent. One modeling challenge is thus resolution: resolving features less than 100 km thick in a spherical domain of depth 2900 km and circumference 40,000 km. However, the biggest modeling challenge is the rapid spatial variation of physical properties, and in particular, viscosity, which is strongly dependent on temperature, pressure and stress, and thus varies by many orders of magnitude over the smallest lengthscales in the system (the thermal boundary layers). These strong variations cause convergence problems for iterative solution schemes and accuracy issues for all schemes.

The largest unresolved scientific issue is how plate tectonics arises from mantle convection, since plates do not arise from purely viscous creep and are not observed on any other terrestrial planet or moon undergoing mantle convection. For our present spherical modeling, this fundamental issue is sidestepped by using a parameterization of plates, but the attempt to solve the mantle:plate system self-consistently is discussed in the last section of this paper.

The TERRA code

The TERRA code, originally developed by Baumgardner (Baumgardner 1985), uses a finite-element multigrid method to model mantle convection in a three-dimensional spherical shell. In the azimuthal directions a triangular mesh is generated by projecting a regular isocahedron onto the sphere, then repeatedly subdividing the resulting spherical triangles by factors of two (Baumgardner and Frederickson 1985). This approach gives nearly uniform cell size over the sphere. In the radial directions, mesh refinement is used to give greater resolution in the boundary layers. An elliptical multigrid solver is used to solve for the velocity and pressure fields at each timestep, while timestepping is performed using a version of the MPDATA advection technique (Smolarkiewicz 1984). The multigrid approach is very efficient, giving a linear scaling of execution time with number of points, efficient storage (presently about 100 words per grid cell), and straightforward parallelization.

Strongly Varying Material Properties

As mentioned above, the spatially rapid variations of material properties, particularly viscosity, present one of the greatest challenges for the solution scheme. Thus, one of the focuses of the recent efforts on TERRA has been algorithm development, to enhance our ability to treat some of the important physical effects in mantle convection with greater realism. For example, the multigrid algorithm used to solve the elliptic momentum equation for the vector velocity field at each grid point on each time step has been improved. The tensor operator for this system now has coefficients that vary strongly with position and time due to the temperature, depth, and stress dependence of the deformational properties of the mantle. We can now run models in which the material strength varies by as a factor of 25 from one grid point to the next. With this enhanced capability, TERRA's ability to treat lithospheric plates in a realistic manner is increased substantially. It allows us to use a yield stress criterion to model strain localization and weakening at plate margins. This enhanced robustness has been accomplished mainly by the implementation of matrix-dependent transfers (Reusken 1993) and the Galerkin coarse-grid approximation into the multigrid scheme, by project postdoc Woo-Sun Yang working in conjunction with Baumgardner.

Plate Treatment

In TERRA, plates are parameterized as a boundary condition. Plates are modeled as perfectly rigid patches in the topmost layer of nodes in the spherical grid. The motion of each plate is described by a single Euler rotation vector. In a dynamical calculation the set of plate rotation vectors is obtained via a Newton iteration procedure that minimizes the net torque on each plate. The resulting piecewise uniform surface velocity field is then applied as a velocity boundary condition in solving the momentum equation for the velocity field in the interior of the spherical shell domain. The spatial identity of the plates is carried by sets of particles, one set for each plate, that are moved in a Lagrangian manner across the surface. A set of rules is applied at plate boundaries to treat the removal of particles at zones of convergence and addition of particles and creation of new plate in zones of divergence.

Parallelization and Performance

The TERRA code is parallelized using a two-dimensional domain decomposition in the azimuthal directions, implemented by Bunge (Bunge and Baumgardner 1994). In this scheme, each of the 10 diamond-shaped blocks corresponding to pairs of the original isocahedral triangles is subdivided by the number of nodes, with each node holding a matching piece of each diamond, and all the associated radial information. The multigrid method is highly suitable for parallel processing and high parallel efficiencies of 85% or more have been accomplished without a large amount of tuning. Under the HPCC project, two parallel tuning modifications have been implemented: eliminating a piece of serial coding at the lowest levels of the multigrid solver, and removing an even power-of-two restriction of the number of processors. The first task involved mapping the calculation at the lowest grid levels, instead of to a single processor as had been done previously, rather to one fourth of the remaining active processors. This strategy accomplishes the work at these lowest levels of the multigrid algorithm in parallel instead of serially. It required writing new communication routines to handle the new communication patterns as well as modifying the data structures at these lower levels.

The largest amount of tuning work has been to optimize the reuse of cache on the DEC Alpha processor and the use of the hardware stream buffers (which allow more efficient transfer of data from main memory to secondary cache) on the Cray T3E. This work was done by Baumgardner and Spencer Swift of SGI/Cray, respectively. In optimizing cache usage, the primary strategy was to rewrite the routines that apply the various 21-point finite element stencils to achieve higher cache reuse. This involved unrolling loops to load 1-D lines of data instead of 2-D slices into the T3D's small 1024-word primary data cache.

The TERRA code achieved the HPCC 100 Gflops milestone October 29, 1998, on 1024 nodes of a Cray T3E-1200 for a calculation that included lithospheric plates intimately coupled to the underlying convecting mantle for the present day Earth similar to model results reported by (Bunge et al. 1998). The computational mesh for this problem contained over 42 million grid points for a total of 170 million velocity/pressure degrees of freedom in the momentum system of equations. A single time step in this calculation required on the order of only 20 wall clock seconds as a result of the high efficiency of the multigrid algorithm used in solving this system. The spatial resolution of this calculation was approximately 30 km at the Earth's surface. One output of this run was a map of the heat flux into the mantle at the core/mantle boundary which was used as an input boundary condition for a core dynamo calculation with the DYNAMO code, also as part of this milestone effort. The TERRA calculation gave a performance on the T3E-1200 of 115.8 Gflops, or 113 Mflops per processor. This represents 9.4 percent ot the theoretical peak processor performance of 1200 Mflops. TERRA displays almost perfectly linear speedup to the number of processors (1024) used in this milestone calculation.

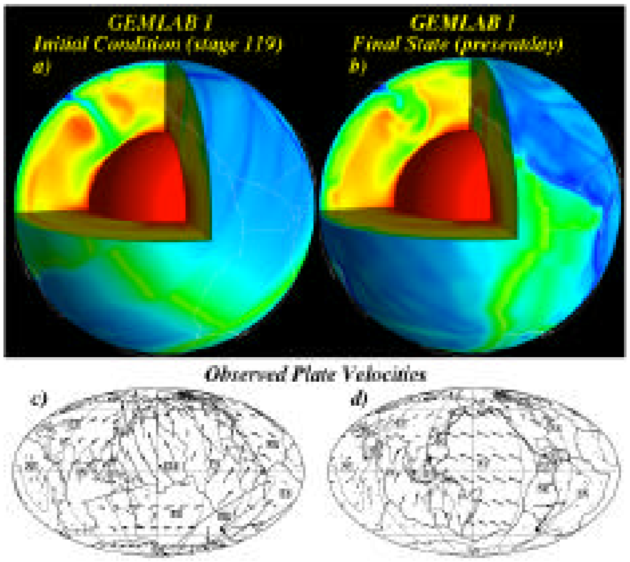

Figure 2. a) Cut-away of the 3D temperature field (from 50 to 2900 km depth) for a starting model seen from the Pacific hemisphere. Blue is cold, and red is hot. Present-day plate boundaries are drawn. b) same as a) but for present day after 119 Myrs evolution. c) and d) maps of plate boundaries for the two times. After (Bunge et al., 1998).

Scientific Results

The TERRA code is now a powerful, efficient tool which can be run on platforms from PCs to the largest parallel supercomputers and is being used for a range of studies by project team members and other researchers around the country. Studies fall into two categories: (i) Simulations designed to investigate the basic fluid dynamics of mantle convection. These usually have long integration times, of order billions of years, e.g., (Bunge, Richards, and Baumgardner 1996; Bunge, Richards, and Baumgardner 1997). (ii) Simulations designed to reproduce in detail the most recent period of Earth history- that for which reliable reconstructions of plate motions exist, which is about 150 million years, less than one overturn time of the mantle convection system (Bunge et al. 1998). These simulations produce output which is directly comparable with various observations, such as 3-D maps of Earth's interior derived from seismic observations (Megnin et al. 1997), or variations in surface topography and plate stress. Such models can also produce estimates of the heat flux distribution through the core-mantle boundary at various times in the past, which can be used as boundary conditions for the DYNAMO code.

An example of such a simulation from a recent study (Bunge et al. 1998) is shown in Figure 2. In this simulation the geologically-reconstructed plate motions for the last 119 million years are imposed as a surface boundary condition. As an initial condition the system is run for a long time with the 119 million year plate configuration, before switching on the evolving configuration. The resulting 3-D heterogeneity structure resembles that observed in the actual Earth using seismic imaging techniques (tomography).

The Quest for a Self-Consistent Plate Tectonics

Lithospheric plate tectonics and mantle convection are different aspects of the same, coupled system, yet mantle convection simulations do not exhibit plate tectonic behavior unless it is imposed by the modeler. This is because present convection models assume that the mantle is a viscous fluid, and 'realistic' temperature-dependent viscosity results in an stiff, immobile, rigid lid. Thus, for example, the present plate treatment in TERRA requires the direct specification of plates by the modeler, with the plates then being tracked using particles, and rules governing the interaction of plate particles at the plate boundaries. While such approaches are facilitating some important research, it is ultimately necessary to identify the correct self-consistent description of plate tectonics and mantle convection, in which both components arise naturally out of a unified material description of rock deformation as a function of temperature, pressure, and stress. Team members Tackley and David Bercivici at the University of Hawaii have been investigating such approaches, with some success (Bercovici 1998; Tackley 1998).

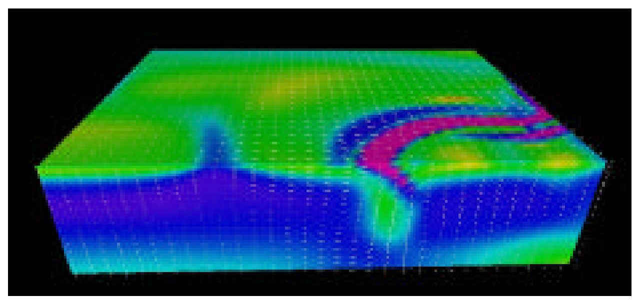

Figure 3. Cross-section through a mantle convection simulation in which strong plates and weak plate boundaries form self-consistently from the material properties. Shown is the log of the viscosity field, varying by 5 orders of magnitude, with red to green being high viscosity and blue to purple being low viscosity. Single-sided subduction (to the right of center) and a passive rift (left) are both distinguishing characteristics of plate tectonics.

The cold upper boundary layer of the convection is so stiff (because the viscosity is temperature-dependent) that viscous creep is no longer the important deformation mechanism, and other processes such as elastic deformation, brittle failure (causing faults) and plastic failure become important. Through poorly-understood interactions, these processes cause localization which results in the formation of weak plate boundaries in the strong, cold boundary layer. These modes of deformation also introduce history-dependence into the system.

Incorporating such complexities into numerical mantle convection models is extremely difficult, because (i) of the greatly increased numerical complexity, particularly with finite-strain elasticity, (ii) they cause timescales and lengthscales which are much smaller than those associated with regular convection, for example, when faults form and earthquakes occur, (iii) the physical properties and material behaviors are not well known.

However, encouraging progress is being made, as shown in Figure 3. This shows a calculation made with the STAG code, a finite-volume multigrid code which can model large viscosity variations and complex material properties and is being used as a testbed for different material descriptions (Tackley 1993). The calculation shown uses a viscous rheology but with the key additions of plastic yielding (i.e., the material has a maximum strength and gives way once that strength is reached), and strain weakening, i.e., stiff material becomes weaker when deformed, allowing it to deform faster and leading to localization. This calculation illustrates the self-consistent formation of plate boundaries, with a single-sided subduction zone and passive rift clearly visible. Future work will include elasticity and brittle failure, the incorporation of which is being studied by project postdoc Moritz Heimel at UCLA. When our understanding of the appropriate deformation mechanisms has been refined, these will be included in the TERRA model.

Acknowledgements. This work was supported by NASA under Cooperative Agreement No. NCCS5-147.

REFERENCES

Baumgardner, J. R. 1985. "3-Dimensional Treatment Of Convective Flow In the Earths Mantle." *Journal Of Statistical Physics* 39 (5-6):501-520.

Baumgardner, J. R., and P. O. Frederickson. 1985. "Icosahedral Discretization Of the 2-Sphere." *Siam Journal On Numerical Analysis* 22 (6):1107-1115.

Bercovici, D. 1998. "Generation of Plate Tectonics from lithosphere-mantle flow and void-volatile self-lubrication." *Earth Planet. Sci. Lett.* 154:139-151.

Bunge, H.-P., and J. R. Baumgardner. 1994. "Mantle convection modeling on parallel virtual machines." *Computers in Physics* 9 (2):207-215.

Bunge, H. P., M. A. Richards, and J. R. Baumgardner. 1996. "Effect of depth-dependent viscosity on the planform of mantle convection." *Nature* 379 (6564):436-438.

Bunge, H. P., M. A. Richards, and J. R. Baumgardner. 1997. "A sensitivity study of 3-dimensional spherical mantle convection at 10^8 Rayleigh number - Effects of depth-dependent viscosity, heating mode, and an endothermic phase change." *J. Geophys. Res.* 102 (B6):11991-12007.

Bunge, H. P., M. A. Richards, C. Lithgowbertelloni, J. R. Baumgardner, S. P. Grand, and B. A. Romanowicz. 1998. "Time Scales and Heterogeneous Structure in Geodynamic Earth Models." *Science* 280 (5360):91-95.

Christensen, U., P. Olson, and G. A. Glatzmaier. 1998. "A Dynamo Model Interpretation Of Geomagnetic-Field Structures." *Geophysical Research Letters* 25 (10):1565-1568.

Glatzmaier, G. A., and P. H. Roberts. 1995a. "A 3-Dimensional Convective Dynamo Solution With Rotating and Finitely Conducting Inner-Core and Mantle." *Physics Of the Earth and Planetary Interiors* 91 (1-3):63-75.

Glatzmaier, G. A., and P. H. Roberts. 1995b. "A 3-Dimensional Self-Consistent Computer-Simulation Of a Geomagnetic-Field Reversal." *Nature* 377 (6546):203-209.

Glatzmaier, G. A., and P. H. Roberts. 1996a. "An Anelastic Evolutionary Geodynamo Simulation Driven By Compositional and Thermal-Convection." *Physica D* 97 (1-3):81-94.

Glatzmaier, G. A., and P. H. Roberts. 1996b. "Rotation and Magnetism Of Earths Inner-Core." *Science* 274 (5294):1887-1891.

Glatzmaier, G. A., and P. H. Roberts. 1997. "Simulating the Geodynamo." *Contemporary Physics* 38 (4):269-288.

Megnin, C., H. P. Bunge, B. Romanowicz, and M. A. Richards. 1997. "Imaging 3-D Spherical Convection Models - What Can Seismic Tomography Tell Us About Mantle Dynamics." *Geophysical Research Letters* 24 (11):1299-1302.

Olson, P., and G. A. Glatzmaier. 1996. "Magnetoconvection and Thermal Coupling Of the Earths Core and Mantle." *Philosophical Transactions Of the Royal Society Of London Series a Mathematicalphysical and Engineering Sciences* 354 (1711):1413-1424.

Reusken, A. 1993. "Multigrid With Matrix-Dependent Transfer Operators For a Singular Perturbation Problem." *Computing* 50 (3):199-211.

Smolarkiewicz, P. K. 1984. "A fully multidimensional positive definite advection transport algorithm with small implicit diffusion." *J. Comp. Phys.* 54 (2):325-362.

Tackley, P. J. 1993. "Effects of strongly temperature-dependent viscosity on time-dependent, 3-dimensional models of mantle convection." *Geophys. Res. Lett.* 20 (20):2187-2190.

Tackley, P.J. 1998. "Self-consistent generation of tectonic plates in three-dimensional mantle convection." *Earth Planet. Sci Lett.* 157:9-22.

Tackley, P. J., D. J. Stevenson, G. A. Glatzmaier, and G. Schubert. 1993. "Effects of an endothermic phase transition at 670 km depth in a spherical model of convection in the Earth's mantle." *Nature* 361 (6414):699-704.

BIOGRAPHY. Prof. Paul J. Tackley received a B.A. in Natural Sciences from Cambridge University, England, in 1987 and a Ph.D. in Geophysics from Caltech in 1994, immediately thereafter taking up a faculty position in the Department of Earth and Space Sciences at UCLA. He was the first person to use massively-parallel computers to model mantle convection, starting with Caltech’s Intel Delta system in 1992, and the results generated have appeared on the cover of *Nature* and *Physics World*, and have been featured in many popular magazines including *National Geographic*. He was also the first to develop a 3-D multigrid code capable of modeling mantle convection with strongly varying viscosity. Recently he received a Faculty Fellowship from the David and Lucile Packard Foundation.