


CitcomS

User Manual
Version 3.3.0



Eh Tan
Michael Gurnis
Luis Armendariz
Leif Strand
Susan Kientz

CitcomS User Manual

October 9, 2014

Contents

I Preface	9
II Chapters	13
1 Introduction	15
1.1 About CitcomS	15
1.2 History	15
1.3 Governing Equations	18
1.4 Numerical Methods	20
1.5 Meshes and Geometry	21
2 Installation and Getting Help	25
2.1 Introduction	25
2.2 Getting Help	25
2.3 System Requirements	25
2.3.1 C Compiler	26
2.3.2 MPI Library	26
2.3.2.1 MPI C Compiler Command	26
2.4 Downloading and Unpacking Source	26
2.5 Installation Procedure	27
2.5.1 Installing to a Secondary Location (Optional)	27
2.6 Configuration	27
2.6.1 Configure Usage	28
2.6.2 Environment Variables	28
2.6.3 MPI Configuration	28
2.6.3.1 Manually Specifying MPI include and lib Directories	28
2.6.3.2 Manually Specifying MPI include and lib Directories and an Alternative Compiler	29
2.7 HDF5 Configuration (Optional)	29
2.7.1 Additional Tools	29
2.7.1.1 NumPy	29
2.7.1.2 PyTables	29
2.7.1.3 HDFView	30
2.7.1.4 OpenDXutils	30
2.8 GGRD Configuration (Optional)	30
2.9 Installing from the Software Repository	30
2.9.1 Tools You Will Need	30
2.9.2 Download Source from Subversion	31
2.9.3 Generating the GNU Build System	31

3	Running CitcomS	33
3.1	Using CitcomS	33
3.2	Changing Parameters	33
3.3	Uniprocessor Example	33
3.3.0.1	Example: Uniprocessor, <code>example0</code>	34
3.4	Multiprocessor Example	34
3.4.0.2	Example: Multiprocessor, <code>example1</code>	34
3.4.1	Output Formats	36
4	Working with CitcomS HDF5 Files	37
4.1	Introduction	37
4.2	About HDF5	37
4.3	Input Parameters	38
4.3.1	Optimizing Parallel I/O	38
4.4	Data Layout	39
4.5	Accessing Data	40
4.5.1	Inspecting Structures	40
4.5.2	Converting to ASCII Files	40
4.5.3	Accessing Data in Python	40
4.5.4	Accessing Data Using HDFView	40
5	Postprocessing and Graphics	43
5.1	Introduction	43
5.2	Postprocessing on a Beowulf Cluster	43
5.3	Postprocessing in a Non-Cluster Environment	44
5.4	Using OpenDX for Regional Sphere Visualization	44
5.5	Using OpenDX for Full Sphere Visualization	46
5.6	Using OpenDX for HDF5 Visualization	46
5.7	Using MayaVi for Visualization	47
5.8	Using GMT Commands for Visualization	49
6	Cookbooks	51
6.1	Introduction	51
6.2	Cookbook 1: Global Model	51
6.2.1	Problem	51
6.2.2	Solution	52
6.2.2.1	Example: Global Model, <code>cookbook1</code>	52
6.2.3	Discussion	53
6.3	Cookbook 2: Domain Size and Velocity Boundary Conditions	54
6.3.1	Problem	54
6.3.2	Solution	54
6.3.2.1	Example: Velocity Boundary Conditions, <code>cookbook2</code>	54
6.3.3	Discussion	55
6.4	Cookbook 3: Temperature-Dependent Viscosity	56
6.4.1	Problem	56
6.4.2	Solution	56
6.4.2.1	Example: Temperature-Dependent Viscosity, <code>cookbook3</code>	57
6.5	Cookbook 4: Regionally Refined Meshes	60
6.5.1	Problem	60
6.5.2	Solution	60
6.5.2.1	Example: Regionally Refined Meshes, <code>cookbook4</code>	60
6.5.3	Discussion	61
6.6	Cookbook 5: Subduction Models with Trench Rollback	62
6.6.1	Problem	62

6.6.2	Solution	62
6.6.2.1	Example: Subduction Models with Trench Rollback, <code>cookbook5</code>	63
6.6.3	Discussion	64
6.7	Cookbook 6: Pseudo-Free-Surface Formulation	65
6.7.1	Problem	65
6.7.2	Solution	65
6.7.3	Discussion	65
6.8	Cookbook 7: Thermo-Chemical Convection	67
6.8.1	Problem	67
6.8.2	Solution	67
6.8.2.1	Example: Thermo-Chemical Convection, <code>cookbook7</code>	68
6.8.3	Discussion	69
6.9	Cookbook 8: Compressible Steady-State Convection	71
6.9.1	Problem	71
6.9.2	Solution	71
6.9.2.1	Example: Compressible Steady-State Convection, <code>cookbook8</code>	73
6.9.3	Discussion	75

III Appendices 77

A Input Parameters for CitcomS 79

A.1	Input Parameters Grouped by Functionality	79
A.1.1	Parameters that Control Input Files	79
A.1.2	Parameters that Control Output Files	80
A.1.3	Mesh and Processors Setup	80
A.1.4	Domain Size	81
A.1.5	Restarting the Code	81
A.1.6	Run Length	81
A.1.7	Initial Conditions	81
A.1.8	Boundary Conditions	83
A.1.9	Non-Dimensional Numbers	83
A.1.10	Depth Information	83
A.1.11	Viscosity	83
A.1.12	Phase Change Information	85
A.1.13	Momentum Equation Solver Parameters	85
A.1.14	Energy Equation Solver Parameters	86
A.1.15	Age Information	86
A.1.16	Debugging Information	86
A.1.17	HDF5 Output Parameters	86
A.1.18	Tracer Parameters	87
A.1.19	Dimensional Information	87
A.1.20	Required Information	87

B CitcomS Input File Format 89

B.1	Introduction	89
B.2	Coordinate Files	89
B.3	Velocity Boundary Condition Files	90
B.4	Material Files	90
B.5	Lithosphere Age Files	91
B.6	Tracer Files	91
B.7	Reference State Files	91

C	CitcomS Output File Format	93
C.1	Introduction	93
C.2	Postprocessed Cap Output (test-case.cap00.10)	93
C.3	Postprocessed Opt Output (test-case.opt00.10)	93
C.4	Postprocessed Surf Output (test-case.surf0.10)	93
C.5	Time Output (test-case.time)	94
C.6	ASCII Output	94
C.6.1	Coordinate Output (test-case.coord.0)	94
C.6.2	Velocity and Temperature Output (test-case.velo.0.10)	94
C.6.3	Viscosity Output (test-case.visc.0.10)	94
C.6.4	Material Output (test-case.mat.0)	94
C.6.5	Surface Variables Output (test-case.surf.0.10 and test-case.botm.0.10)	94
C.6.6	Stress Output (test-case.stress.0.10)	94
C.6.7	Pressure Output (test-case.pressure.0.10)	95
C.6.8	Horizontal Average Output (test-case.horiz_avg.0.10)	95
C.6.9	Geoid Output (test-case.geoid.10)	95
C.6.10	Tracer Output (test-case.tracer.0.10)	95
C.6.11	Composition Output (test-case.comp_el.0.10 and test-case.comp_nd.0.10)	95
C.6.12	Heating Output (test-case.heating.0.10)	95
C.7	HDF5 Output (test-case.h5)	96
C.8	Misc. Binary Output	96
C.8.1	Checkpoint Output (test-case.checkpoint.0.10)	96
C.8.2	Domain Output (test-case.domain)	96
C.8.3	Coordinate Binary Output (test-case.coord_bin.0)	96
C.8.4	Seismic Output (test-case.seismic.0.10)	96
D	License	97

List of Figures

1.1	Global Node Numbering. Left: Global node numbering starts at the base of arrow A (θ_{\min} , ϕ_{\min} , r_{inner}), and advances from 1 at the base to n_{nodes} at the tip. Upon reaching the tip, numbering continues from the base of arrow B ($n_{\text{nodes}}+1$) to its tip ($2 n_{\text{nodes}}$), and so on for all nodes on the plane $\phi = \phi_{\min}$. Right: After completing each theta radius plane, the ϕ index is incremented and numbering commences from (θ_{\min} , r_{inner}) as on the left.	21
1.2	Topological connectivity of the 12 caps. N is the north pole and S is the south pole. The red line marks the 0 degree meridian.	22
1.3	Orthographic projection of processors from a full CitcomS mesh in which there are 16 processors in map view for each cap. The CitcomS cap is shown as distinct colors while the processor domains within the caps are indicated by the intensity of the color. This example was produced for a run with 2 processors in radius such that the total number of processors was $12 \times 16 \times 2 = 384$	23
3.1	Computational Domain. Map view on the configuration of the top layer of the computational nodes and the processors.	36
4.1	A screenshot of HDFView. The left panel shows the hierarchy of the groups and datasets. The right panel shows a 2D slice of a dataset. The bottom panel shows the metadata associated with the selected group or dataset.	41
5.1	Regional Model Visualized with OpenDX. A snapshot of an upwelling (blue isosurface) with a slice of the temperature field (bisecting plane).	45
5.2	How to import CitcomS HDF5 data. The <code>CitcomSImportHDF5</code> module is highlighted in the Tools panel.	47
5.3	Example MayaVi2 visualization of <code>cookbook1.100.h5</code> . Here we display the velocity field as vectors, as well a slice of the temperature field. Note the visualization pipeline in the tree view on the left panel.	48
5.4	Temperature field generated by <code>plot_layer.py</code> . This image is the mid-layer (layer 5) of the sample data <code>visual/samples/regtest.cap00.100</code> . The resolution of the image is poor because the data has only $9 \times 9 \times 9$ nodes.	49
5.5	Temperature fields generated by <code>plot_annulus.py</code> . The top panel is the horizontal cross section at the mid layer in Hammer-Aitoff projection. The thick line marks the surface track of the radial cross section. The bottom panel is the radial cross section. The cross section starts at 0° latitude and 0° longitude and has an azimuth of 45° . The image is from the sample data <code>visual/samples/fulltest.cap*.100</code> . There is always a gap on the left (starting point) of the cross section. Users familiar with GMT are welcome to contribute a fix to the script.	50
6.1	Global Model “Caps.” Left (A): Three-dimensional perspective image showing seven of the 12 spherical caps used in a full CitcomS run. Right (B): The temperature field at 1081 km depth from a Cookbook 1 run.	51
6.2	Cookbook 1: Global Model. This image, created by OpenDX, depicts a slice through a spherical model of convection, with warmer colors indicating upwelling and the cooler colors showing downwelling.	53

6.3	Cookbook 2: Velocity Boundary Conditions. This model, visualized with OpenDX, highlights a region of the sphere and the heated upwellings (warm colors), downwellings (cool colors), and the velocities (yellow arrows).	56
6.4	Cookbook 3: Temperature-Dependent Viscosity. This model, visualized with OpenDX, highlights a region that features both the heated upwelling (warm colors) and the even distribution of the velocities (arrows).	59
6.5	Cookbook 4: Regionally Refined Mesh. This model shows the temperature (upwelling – warm colors, downwelling – cool colors) and the uneven distribution of the velocities (yellow arrows). Note the refined mesh for the left and bottom regions of the model (inset).	61
6.6	Cookbook 5: Left (A): Initially, both plates have Euler poles (magenta and green dots) situated far apart from the plate, so the velocities are approximately constant (about 5 cm/yr for the left plate and about 2 cm/yr for the right plate). Right (B): After 30 Ma a pole jump occurred for the right slab only, so the Euler pole moved closer.	63
6.7	Cookbook 5: The pole jump after 30 Ma produced a flat slab on one side (fore side) and a steep slab on the other side (back side).	64
6.8	Cookbook 6: Graphs of topography profiles.	66
6.9	Cookbook 7: The composition and velocity field at the 15th step. The arrows are the velocity vectors. The composition field is shown in an isosurface of 0.7 and in a cross section.	69
6.10	Cookbook 8: The steady state temperature field at the 10,000th time step. A tetrahedral symmetric convection pattern is developed. Two temperature isosurfaces of 0.4 and 0.8 are shown.	76

Part I
Preface

Preface

About This Document

This document is organized into three parts. Part I consists of traditional book front matter, including this preface. Part II describes the capabilities of CitcomS and the details of implementation, including several “cookbooks” of short tutorials. Part III provides appendices and references.

The style of this publication is based on the Apple Publications Style Guide (developer.apple.com/documentation/UserExperience/Conceptual/APStyleGuide/AppleStyleGuide2003.pdf), as recommended by Python.org (www.python.org). The documentation was produced using L^AT_EX, specifically the pdf_lat_ex tool from the TeXLive distribution (www.tug.org/texlive).

Errors and bug fixes in this manual should be directed to the CIG Mantle Convection Mailing List (cig-mc@geodynamics.org).

Who Will Use This Document

This documentation is aimed at two categories of users: scientists who prefer to use prepackaged and specialized analysis tools, and experienced computational Earth scientists. Of the latter, there are likely to be two classes of users: those who just run models, and those who modify the source code. Users who modify the source are likely to have familiarity with scripting, software installation, and programming, but are not necessarily professional programmers.

The manual was written for the usage of CitcomS on a variety of different platforms. CitcomS has run on shared memory computers (Sun, Hewlett-Packard, SGI, and IBM), commercial distributed memory machines (Intel and Cray/SGI), clusters (including machines on the NSF TeraGrid), Linux PCs and Mac OS X desktops.

Citation

Computational Infrastructure for Geodynamics (CIG) is making this source code available to you in the hope that the software will enhance your research in geophysics. The underlying C code for the finite element package and the Python bindings for the framework were donated to CIG in July of 2005. A number of individuals have contributed a significant portion of their careers toward the development of CitcomS and Pyre. It is essential that you recognize these individuals in the normal scientific practice by citing the appropriate peer reviewed papers and making appropriate acknowledgements.

The CitcomS development team asks that you cite both of the following:

- Zhong, S., M.T. Zuber, L.N. Moresi, and M. Gurnis (2000), The role of temperature-dependent viscosity and surface plates in spherical shell models of mantle convection. *J. Geophys. Res.*, 105, 11,063-11,082
- Tan, E., E. Choi, P. Thoutireddy, M. Gurnis, and M. Aivazis (2006), GeoFramework: Coupling multiple models of mantle convection within a computational framework, *Geochem., Geophys., Geosyst.* 7, Q06001, doi:10.1029/2005GC001155.

Additionally, if you are using tracers in CitcomS, please cite the following:

- McNamara, A.K., and S. Zhong (2004), Thermochemical structures within a spherical mantle: Superplumes or Piles?, *J. Geophys. Res.*, 109, B07402, doi:10.1029/2003JB002847.

The developers also request that in your oral presentations and in your paper acknowledgements that you indicate your use of this code, the authors of the code, and CIG (geodynamics.org).

Support

Pyre development was funded by the U.S. Dept. of Energy's Advanced Simulation and Computing program (www.sandia.gov/NNSA/ASC) and the National Science Foundation's (www.nsf.gov) Information Technology Research (ITR) program (grant #0205653). Continued support of CitcomS is based upon work supported by the National Science Foundation under Grant No. EAR-0406751. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Conventions

Throughout this documentation, any mention of "username" is meant to indicate the user, meaning you should substitute your account name in its place.

Part II
Chapters

Chapter 1

Introduction

CitcomS is a finite element code designed to solve thermal convection problems relevant to earth's mantle released under the GNU General Public License (see Appendix D on page 97). Written in C, the code runs on a variety of parallel processing computers, including shared and distributed memory platforms.

1.1 About CitcomS

CitcomS is a finite element code written in C that solves for thermo-chemical convection within a spherical shell. It can solve for problems within either a full or a restricted spherical domain. Although the code is capable of solving many different kinds of convection problems using the flexibility of finite elements, there are aspects of CitcomS which make it well-suited for solving problems in which the plate tectonic history is incorporated. Variable viscosity, including temperature-, pressure-, position-, composition-, and stress-dependent viscosity are all possible.

The fundamental basis for the numerical solution of any time-dependent convection problem is the sequential solution of an equation of motion and an energy equation. Convection problems are initially valued with boundary conditions, including all of the problems which are solved with CitcomS. The normal sequence of steps for the solution of convection problems starts with an initial temperature field. First, the momentum equation is solved. The solution of this equation gives us the velocity from which we then solve the advection-diffusion equation, giving us a new temperature. CitcomS uses this interleaved strategy. It is possible to run convection backward in time so as to guess an initial condition for a normal forward running initial and boundary value problem. However, users should be aware that even specialists in mantle convection modeling are just now starting to explore methods in this area and, as such, this is an emerging area of research.

This code uses an iterative solution scheme to solve for the velocity and pressure and, as such, a converged solution cannot be guaranteed. Nested inside the iterative scheme, the code uses either a conjugate gradient solver or a full multi-grid solver to solve the discretized matrix equations.

1.2 History

Citcom (for **C**alifornia **I**nstitute of **T**echnology **C**onvection in the **M**antle) was originally written in the early 1990s by Louis Moresi. Although the code for three-dimensional problems was incorporated from its inception, early versions of the software only solved for time-dependent convection problems within two-dimensional Cartesian domains. Moresi's original code turned out to be incredibly modular and easily extensible. Consequently, the fundamental finite element infrastructure which Louis wrote is still in place and forms the basis for much of the code contained in the present release.

In the mid-1990s Moresi wrote versions of the code that solved the equations within three-dimensional Cartesian domains. Then Shijie Zhong successfully parallelized Citcom using message passing routines on a limited release Intel supercomputer. Zhong then created a spherical version of the code which he named CitcomS. Lijie Han then created a regional version of CitcomS as well as an alternate version of message passing for an arbitrarily large number of processors. Clint Conrad created the first Beowulf implementations of the code using the MPI library, then Conrad and Eh Tan re-coded the message passing of the fully spherical version so that problems run on arbitrarily large numbers of

processors could also be solved. A plethora of different versions of Citcom exist both on computers at the California Institute of Technology and around the world.

Consequently, by 2002, there were so many different versions of the code that some rationalization was in order. The software was migrated into a version control system and Eh Tan and Eun-seo Choi created a version of CitcomS that generates either a fully spherical or regional model. CitcomS was released to the community through the former GeoFramework project as version 1.0 and 1.1.

By 2004, in order to increase the functionality of CitcomS, the developers began to reengineer the code into an object-oriented environment specifically so it could work with a Python-based modeling framework called Pyre. This release of the software is essentially the product of those reengineering efforts. Eh Tan was the principal developer of CitcomS, with considerable help from Eun-seo Choi, Puru Thoutireddy, and Michael Aivazis.

CitcomS is one component of a larger collection of software encompassed by the former GeoFramework project, a collaboration between the Center for Advanced Computing Research (CACR) (www.cacr.caltech.edu) and the Seismological Laboratory (www.gps.caltech.edu/seismo), both at Caltech, and the Victorian Partnership for Advanced Computing (www.vpac.org) in Australia. The GeoFramework project developed a suite of tools to model multi-scale deformation for Earth science problems. This effort was motivated by the need to understand interactions between the long-term evolution of plate tectonics and shorter term processes such as the evolution of faults during and between earthquakes. During 2005 and 2006 much of the remaining software developed by GeoFramework was released under a GPL license and made available from CIG (geodynamics.org).

The second major release of CitcomS (2.0) incorporated the software framework Pyre, free surface modeling methods, and stress boundary conditions on the top and bottom surfaces. In the summer of 2005, as part of the 2.0.1 release, CIG replaced the old build procedure with the GNU Build System. A subsequent release, version 2.0.2, could compile and run on 64-bit systems.

The release of CitcomS 2.1 incorporated new features and functionality, the most important being the use of HDF5 (a parallel version of the Hierarchical Data Format). The HDF5 format allows you to deal with the massive data output created for production runs (see Chapter 4 on page 37). This version accepted `.cfg` files on input, which are easier to create and read. Other improvements included the incorporation of geoid calculations that had been left out of earlier releases, as well as new scripts to allow results to be visualized with MayaVi2 (svn.enthought.com/enthought/wiki/MayaVi) in addition to Generic Mapping Tools (GMT) (gmt.soest.hawaii.edu) and OpenDX (www.opendx.org). Instructions were provided on using this version as a preinstalled package on some of the NSF TeraGrid sites.

The release of CitcomS 2.2 incorporated the ability of tracing particles in the flow. The tracer code was developed by Allen McNamara and Shijie Zhong in 2004 and donated to CIG in early 2007. The tracer code has a wide range of applications in the mantle convection. It can be used in tracing the trajectory of passive particles, in delineating the top boundary of subducted slabs to define the low viscosity wedges, or in tracking the evolution of the chemical composition field.

The release of CitcomS 3.0 contains many new features, including two implementations of compressible convection, one by Wei Leng and Shijie Zhong and the other by Eh Tan; the ability to resume computation from previous checkpoints; multi-component chemical convection; a fixed non-Newtonian solver; an exchanger package for solver coupling; removing the rigid body rotation component from the velocity by Shijie Zhong; and an option to disable monitoring of maximum temperature. In addition, a rheology option for pseudo-plasticity, composition dependent viscosity and heat generation, compressed ASCII output, and an easier way for mesh refinement for the radial coordinate were added by Thorsten Becker.

Additional changes that are not backward compatible include (1) the viscosity field at element level is not smoothed (this might slow down the convergence but will represent the viscosity field more accurately); (2) the Lenardic filter on temperature is disabled by default; (3) the rigid body rotation component is removed from the velocity by default; (4) use of a better pseudo-random number generator to generate the initial tracer; (5) the type of input parameter `color` is changed from a boolean to an integer; (6) setting `restart=on` will resume the computation from the checkpoint files and will not need the tracer files (the old way of reading initial temperature from `velo` files can be achieved by `tic_method=-1`); and (7) the input parameter `reset_initial_composition` becomes obsolete. Among the seven backward-incompatible changes, the first four will affect the results.

The release of CitcomS 3.1 added the ability to convert temperature and composition to seismic velocities. The seismic velocities output can be uploaded to the CIG Seismology Web Portal for SPECFEM3D simulation to generate synthetic seismograms. Other enhancements include adding a self-gravitational effect on geoid and adding the Consistent-Boundary-Flux (CBF) method to compute dynamic topography, contributed by Shijie Zhong. Also,

GMT/NetCDF grd input was added for surface velocity boundary conditions, initial temperature, material dependence and local Rayleigh number in surface layers, contributed by Thorsten Becker. Mike Gurnis contributed added capability to read in time- and geographic-dependent, top surface temperature boundary conditions. Finally, this version allows multi-component chemical viscosity.

Some of the aforementioned features are backward-incompatible, including:

1. The global mesh is changed. The coordinates of all nodes are slightly shifted laterally. The change is not visually discernible, but does affect the solution compared to previous versions.
2. The convergence criterion of the Stokes solver is changed. A single parameter “accuracy” is used to control the accuracy of the continuity and momentum equations. The default value of “accuracy” was 10^{-6} and is now 10^{-4} .
3. The input parameter “mantle_temp” is moved from the [CitcomS.solver.param] to the [CitcomS.solver.ic] section.

The release of CitcomS 3.2 added numerous changes since 3.1.1 which was released in 2009. These include improved implementations of internal stress, radial layer dependent viscosity, and velocity boundary conditions. Improvements have also been made in solver convergence controls, VTK output and several bug fixes. Other enhancements include:

- Trial implementation of anisotropic viscosity.
- Added support of multigrid solver in Exchanger.
- Improved handling of grd-read velocity and traction boundary conditions.
- Changed default parameters for z_410, z_lmantle, z_lith to more Earth-like parameters.
- Added test implementation of steady state strain-rate weakening via PDEPV psrw=on parameter.
- Added Murnaghan’s integrated linear EoS.
- Added augmented Lagrangian component of stiffness matrix in cbf topo.
- Performance optimizations for CUDA code. Contributed by Leif Strand.
- Fixed several bugs in the code and manual, including:
 - Fixed a bug in pseudo free surface, reported by Robert Moucha.
 - Correctly handle orphan tracers.
 - Replaced fixed accuracy setting. Bug reported by Rob Moucha.
 - Fixed temperature perturbation bug for multiple processors.

This release (3.3) removes the Python dependencies from CitcomS and provides a tool to convert Python style parameter files to the original CitcomS parameter files. Python support was proving problematic because of incompatibility with versions of Python newer than 2.6 (final release 2008) and other internal issues. The removal of Python should simplify much of the code while retaining almost all functionality. A full list of the changes includes:

- Removed Python sections of CitcomS
- Wrote Python to standard parameter file conversion tool (Py2C)
- Converted Python cookbooks parameter files to original style parameter files
- Added output of parameters to a pidXXXXXX style file
- Updated the manual to correspond to the new changes

1.3 Governing Equations

With CitcomS, the mantle is treated as an anelastic, compressible, viscous spherical shell under Truncated Anelastic Liquid Approximation. With these assumptions, thermal convection is governed by the equations for conservation of mass, momentum, and energy:

$$(\rho u_i)_{,i} = 0 \quad (1.1)$$

$$-P_{,i} + \left(\eta(u_{i,j} + u_{j,i} - \frac{2}{3}u_{k,k}\delta_{ij}) \right)_{,i} - \delta\rho g\delta_{ir} = 0 \quad (1.2)$$

$$\rho c_P (T_{,t} + u_i T_{,i}) = \rho c_P \kappa T_{,ii} - \rho \alpha g u_r T + \Phi + \rho (Q_{L,t} + u_i Q_{L,i}) + \rho H \quad (1.3)$$

where ρ is the density, u is the velocity, P is the dynamic pressure, η is the viscosity, δ_{ij} is the Kronecker delta tensor, $\delta\rho$ is the density anomaly, g is the gravitational acceleration, T is the temperature, c_P is the heat capacity, κ is the thermal diffusivity, α is the thermal expansivity, Φ is the viscous dissipation, Q_L is the latent heat due to phase transitions, and H is the heat production rate. The expression $X_{,y}$ represents the derivative of X with respect to y , where i and j are spatial indices, r is the radial direction, and t is time. With phase transitions and temperature and composition variations, the density anomalies are:

$$\delta\rho = -\alpha\bar{\rho}(T - \bar{T}_a) + \delta\rho_{ph}\Gamma + \delta\rho_{ch}C \quad (1.4)$$

where $\bar{\rho}$ is the radial profile of density, \bar{T}_a is the radial profile of adiabatic temperature, $\delta\rho_{ph}$ is the density jump across a phase change, $\delta\rho_{ch}$ is the density difference between the compositions, Γ is the phase function, and C is the composition. The phase function is defined as:

$$\pi = \bar{\rho}g(1 - r - d_{ph}) - \gamma_{ph}(T - T_{ph}) \quad (1.5)$$

$$\Gamma = \frac{1}{2} \left(1 + \tanh \left(\frac{\pi}{\bar{\rho}g w_{ph}} \right) \right) \quad (1.6)$$

where π is the reduced pressure, d_{ph} and T_{ph} are the ambient depth and temperature of a phase change, γ_{ph} is the Clapeyron slope of a phase change, and w_{ph} is the width of a phase transition.

These equations lead to the following normalization in which primed quantities are nondimensional:

$$\rho = \rho_0 \rho' \quad (1.7)$$

$$\alpha = \alpha_0 \alpha' \quad (1.8)$$

$$g = g_0 g' \quad (1.9)$$

$$\kappa = \kappa_0 \kappa' \quad (1.10)$$

$$\eta = \eta_0 \eta' \quad (1.11)$$

$$c_P = c_{P0} c_P' \quad (1.12)$$

$$x_i = R_0 x_i' \quad (1.13)$$

$$u_i = \frac{\kappa_0}{R_0} u_i' \quad (1.14)$$

$$T_0 = \Delta T T_0' \quad (1.15)$$

$$T = \Delta T (T' + T_0') \quad (1.16)$$

$$t = \frac{R_0^2}{\kappa_0} t' \quad (1.17)$$

$$H = \frac{\kappa_0}{R_0^2} c_{P0} \Delta T H' \quad (1.18)$$

$$P = \frac{\eta_0 \kappa_0}{R_0^2} P' \quad (1.19)$$

$$d_{ph} = R_0 d_{ph}' \quad (1.20)$$

$$\gamma_{ph} = \frac{\rho_0 g_0 R_0}{\Delta T} \gamma_{ph}' \quad (1.21)$$

where ρ_0 is the reference density, α_0 is the reference thermal expansivity, g_0 is the reference gravity, κ_0 is the reference thermal diffusivity, η_0 is a reference viscosity, c_{P0} is the reference heat capacity, R_0 is the radius of the Earth, T_0 is the temperature at the surface, and ΔT is the temperature drop from the core-mantle boundary (CMB) to the surface. Dropping the primes, the equations become:

$$u_{i,i} + \frac{1}{\bar{\rho}} \frac{d\bar{\rho}}{dr} u_r = 0 \quad (1.22)$$

$$-P_{,i} + \left(\eta (u_{i,j} + u_{j,i} - \frac{2}{3} u_{k,k} \delta_{ij}) \right)_{,i} + (Ra \bar{\rho} \alpha T - Ra_b \Gamma - Ra_c C) g \delta_{ir} = 0 \quad (1.23)$$

$$\begin{aligned} \bar{\rho} c_P (T_{,t} + u_i T_{,i}) \left(1 + 2\Gamma (1 - \Gamma) \frac{\gamma_{ph}^2}{d_{ph}} \frac{Ra_b}{Ra} Di (T + T_0) \right) &= \bar{\rho} c_P \kappa T_{,ii} \\ -\bar{\rho} \alpha g u_r Di (T + T_0) \left(1 + 2\Gamma (1 - \Gamma) \frac{\gamma_{ph}}{d_{ph}} \frac{Ra_b}{Ra} \right) &+ \frac{Di}{Ra} \Phi + \bar{\rho} H \end{aligned} \quad (1.24)$$

where Ra , the thermal Rayleigh number, is defined as:

$$Ra = \frac{\rho_0 g_0 \alpha_0 \Delta T R_0^3}{\eta_0 \kappa_0} \quad (1.25)$$

This is not the usual definition of the Raleigh number that is based on layer thickness; it is based on the radius of the Earth R_0 . So for mantle convection problems where R_0 is slightly more than twice the layer thickness, our Ra is about a factor of 10 larger than by the usual definition. The phase-change Rayleigh number, Ra_b , the chemical Rayleigh number, Ra_c , the internal heating Rayleigh number, Ra_H , and the dissipation number Di , are defined as:

$$Ra_b = Ra \frac{\delta \rho_{ph}}{\rho_0 \alpha_0 \Delta T} \quad (1.26)$$

$$Ra_c = Ra \frac{\delta \rho_{ch}}{\rho_0 \alpha_0 \Delta T} \quad (1.27)$$

$$Ra_H = Ra_H \frac{R_0^3 - R_{CMB}^3}{3R_0^3} \quad (1.28)$$

$$Di = \frac{\alpha_0 g_0 R_0}{c_{P0}} \quad (1.29)$$

1.4 Numerical Methods

The governing equations are solved with the finite element method [9]. CitcomS employs an Uzawa algorithm to solve the momentum equation coupled with the incompressibility constraints [4, 10]. The energy equation is solved with a Steamline-Upwind Petrov-Galerkin method [11]. Brick elements are used, such as eight velocity nodes with trilinear shape functions and one constant pressure node for each element. The use of brick elements in 3D (or rectangular elements in 2D) is important for accurately determining the pressure, which controls the dynamic topography, in incompressible Stokes flow [9]. The discrete form of Equations 1.22 and 1.23 may be written in the following matrix form [6]:

$$(\mathbf{B}^T + \mathbf{C}) u = 0 \quad (1.30)$$

$$\mathbf{A}u + \mathbf{B}p = f \quad (1.31)$$

where \mathbf{A} is the ‘‘stiffness’’ matrix, u is a vector of unknown velocities, \mathbf{B} is the discrete gradient operator, \mathbf{C} is the second term in Equation 1.22, p is a vector of unknown pressures, and f is a vector composed of the body and boundary forces acting on the fluid. The individual entries of \mathbf{A} , \mathbf{B} , \mathbf{C} and f are obtained using a standard finite element formulation; see [6] for the explicit entries.

In the incompressible case, \mathbf{C} is zero. Equation 1.31 can be transformed by premultiplying by $\mathbf{B}^T \mathbf{A}^{-1}$ and using Equation 1.30 to eliminate the velocity unknowns:

$$\mathbf{B}^T \mathbf{A}^{-1} \mathbf{B}p = \mathbf{B}^T \mathbf{A}^{-1} f \quad (1.32)$$

This equation is solved using the Uzawa algorithm, an established method for solving the minimization of a dual function [12], which simultaneously yields the velocity field. A conjugate gradient scheme [10, 13] is used for this iteration and forms the basis for the technique used in CitcomS.

In the compressible case, there are two different strategies to solve Equations 1.30 and 1.31. The first strategy is to add another layer of iterations when solving Equation 1.32. The right-hand-side vector is updated by the velocity solution of the previous iteration. This equation can be solved using the same conjugate gradient scheme as the incompressible case.

$$\mathbf{B}^T \mathbf{A}^{-1} \mathbf{B} p^{(i)} = \mathbf{B}^T \mathbf{A}^{-1} f - \mathbf{C} u^{(i-1)} \quad (1.33)$$

The second strategy is to transform Equation 1.31 by premultiplying by $(\mathbf{B}^T + \mathbf{C}) \mathbf{A}^{-1}$ and using Equation 1.30 to eliminate the velocity unknowns:

$$(\mathbf{B}^T + \mathbf{C}) \mathbf{A}^{-1} \mathbf{B} p = (\mathbf{B}^T + \mathbf{C}) \mathbf{A}^{-1} f \quad (1.34)$$

This equation is solved using a bi-conjugate gradient stabilized scheme.

1.5 Meshes and Geometry

There are two forms of meshes and geometries for CitcomS. By default CitcomS will produce a mesh within a regional geometry that is bound by lines of constant latitude and longitude. There is an option to generate a global mesh of a spherical shell.

For a regional mesh, CitcomS uses meshes that are regular, although considerable flexibility exists for grid refinement in the regional models. There is an option for mesh refinement in which the mesh is refined as a function of latitude, longitude, or radius. Such refinement is suitable for higher resolutions near boundary layers or within the center of the map domain, but is incapable of increasing the resolution near a curvilinear feature, as a plate boundary, unless that plate boundary is orientated north-south or east-west.

In regional meshes, θ (or x) is the colatitude measured from the north pole, ϕ (or y) is the east longitude, and z is the radius. θ and ϕ are in units of radians. Figure 1.1 shows the organization of the mesh in a regional problem. The numbering of the nodes is z -direction first, then x -direction, then y -direction. This numbering convention is used for the input and output data.




Figure 1.1: Global Node Numbering. Left: Global node numbering starts at the base of arrow A (θ_{\min} , ϕ_{\min} , r_{inner}), and advances from 1 at the base to $node_z$ at the tip. Upon reaching the tip, numbering continues from the base of arrow B ($node_z + 1$) to its tip ($2 \cdot node_z$), and so on for all nodes on the plane $\phi = \phi_{\min}$. Right: After completing each θ radius plane, the ϕ index is incremented and numbering commences from $(\theta_{\min}, r_{\text{inner}})$ as on the left.

For a global mesh, CitcomS is also capable of generating a mesh for an entire spherical shell in which elements in map view are approximately equal in area. In the full spherical mode, CitcomS has 12 caps numbered 0 to 11 (Figure 1.2).




Figure 1.2: Topological connectivity of the 12 caps. N is the north pole and S is the south pole. The red line marks the 0 degree meridian.

The caps are approximately square in map view so that the edges of the square are oriented diagonally with respect to latitude and longitude. The four corners of the domain are connected by great circles (Figure 1.3). One would normally associate at least one processor with one cap. However, CitcomS can further decompose the domain such that additional processors are used to divide caps uniformly along the two edges of the caps (Figure 1.3) as well as in radius.




Figure 1.3: Orthographic projection of processors from a full CitcomS mesh in which there are 16 processors in map view for each cap. The CitcomS cap is shown as distinct colors while the processor domains within the caps are indicated by the intensity of the color. This example was produced for a run with 2 processors in radius such that the total number of processors was $12 \times 16 \times 2 = 384$.

Chapter 2

Installation and Getting Help

2.1 Introduction

CitcomS has been tested on Linux, Mac OS X, Sun Solaris and several NSF TeraGrid platforms. You should have no trouble installing CitcomS on most Unix-like systems.

Most users will install CitcomS from the released source package. The following sections will lead you through the installation process. Advanced users and software developers may be interested in downloading the latest CitcomS source code directly from the CIG source code repository, instead of using the prepared source package; see Section 2.9 later in this chapter.

2.2 Getting Help

For help, send e-mail to the CIG Mantle Convection Mailing List (cig-mc@geodynamics.org). You can subscribe to the Mailing List and view archived discussion at the Geodynamics Mail Lists web page (geodynamics.org/cig/lists).

2.3 System Requirements

Installation of CitcomS requires the following:

- A C compiler
- An MPI library

You can install all the dependencies in one command on most Linux machines. On Debian, Ubuntu or similar distributions, use this command (as `root`):

```
$ apt-get install build-essential python-dev \  
openmpi-dev openmpi-bin
```

On Red Hat, Fedora, CentOS, OpenSUSE or similar distributions, use this command (as `root`):

```
$ yum install make automake gcc gcc-c++ kernel-devel \  
python python-devel openmpi openmpi-libs openmpi-devel
```

On Mac OS X, Python is installed by default. You will need to install a C compiler and MPI library (see sections below).

MPI installations are typically configured for a particular compiler, and provide a special wrapper command to invoke the right compiler. Therefore, the choice of MPI implementation often determines which C compiler to use.

2.3.1 C Compiler

On Unix or Linux systems, there is a high likelihood that a usable C compiler is already installed. To check, type `cc` at the shell prompt:

```
$ cc
cc: no input files
$
```

On Linux, if the `cc` command is not found, install GCC using the package manager for your distribution.

The Mac OS X version of GCC is included in a software development suite called Xcode. Xcode is available as a free download at the Apple Developer Connection (developer.apple.com).

Warning: If you are using an Intel compiler on an Itanium CPU, do not use the `-O3` optimization flag as reports indicate that this optimization level will generate incorrect codes. For any compiler, you should always be careful about the correctness of the compiled codes when using an `-O3` or higher optimization level.

2.3.2 MPI Library

CitcomS requires a library which implements the MPI standard (either version 1 or 2). Several free, open-source implementations of MPI are available.

A popular choice is MPICH (www-unix.mcs.anl.gov/mpi/mpich). Other choices include LAM/MPI (www.lam-mpi.org) and Open MPI (www.open-mpi.org). Installing MPI from source involves walking through the standard GNU build procedure (`configure && make && make install`) while logged in as `root`.

Linux users may have a prebuilt MPI package available for their distribution. On Mac OS X, the Fink package manager offers a prepackaged version of LAM/MPI (www.lam-mpi.org); so if you have Fink (fink.sourceforge.net) installed, simply enter the following command from a Terminal window to install LAM/MPI:

```
$ fink install lammpi lammpi-dev
```

2.3.2.1 MPI C Compiler Command

Once you have an MPI library installed, make sure its C compiler command is on your `PATH`. Unfortunately, the name of this command varies from one MPI implementation to the next. The CitcomS configuration script searches for the following MPI C command names:

```
mpicc hcc mpcc mpcc_r mpxlc cmpicc
```

2.4 Downloading and Unpacking Source

To obtain CitcomS, go to the Geodynamics Software Packages web page (geodynamics.org/cig/software/citcoms), download the source archive and unpack it using the `tar` command:

```
$ tar xzf CitcomS-3.3.0.tar.gz
```

If you don't have GNU Tar, try the following command instead:

```
$ gunzip -c CitcomS-3.3.0.tar.gz | tar xf -
```

2.5 Installation Procedure

After unpacking the source, use the following procedure to install CitcomS:

1. Navigate (i.e., `cd`) to the directory containing the CitcomS source.

```
$ cd CitcomS-3.3.0
```

2. Type `./configure` to configure the package for your system.

```
$ ./configure
```

3. Type `make` to build the package.

```
$ make
```

If you are content to run CitcomS from the build directory, then you are done. Upon successful completion, the `make` command creates a script called `citcoms` in the `bin` subdirectory; this is the script you will use to run CitcomS. You may wish to add the `bin` directory to your `PATH`.

For more details about `configure`, see Section 2.6 below.

2.5.1 Installing to a Secondary Location (Optional)

Optionally, after building CitcomS, you can install it in a secondary location using the `make install` command. This is not necessary for using CitcomS and is not recommended for most situations. It is documented only for the sake of completeness.

By default, CitcomS is configured to install under `/usr/local`, which is not writable by normal users. To install as an ordinary user, give `make install` the `prefix` option, specifying a directory to which you have write access:

```
$ make install prefix=$HOME/cig
```

The above command will install CitcomS under `$HOME/cig`. Afterwards, you may wish to add `PREFIX/bin` (`$HOME/cig/bin`, in this example) to your `PATH`.

After running `make install`, you may (if desired) run `make clean` in the build directory to save disk space. You are also free to delete the source/build directory altogether. (Note that `make install` installs the examples under `PREFIX/share/CitcomS/examples`.)

2.6 Configuration

The `configure` script checks for various system features. As it runs, it prints messages informing you of which features it is checking for. Upon successful completion, it generates a `Makefile` in each source directory of the package. It also generates a `config.h` header file, which contains system-dependent definitions.

The `configure` script will attempt to guess the correct values of various installation parameters. In the event that the default values used by `configure` are incorrect for your system, or `configure` is unable to guess the value of a certain parameter, you may have to specify the correct value by hand.

Important: If the `configure` script fails, and you don't know what went wrong, examine the log file `config.log`. This file contains a detailed transcript of all the checks `configure` performed. More importantly, it includes the error output (if any) from your compiler. When seeking help for `configure` failures on the CIG Mantle Convection Mailing List (cig-mc@geodynamics.org), please send `config.log` as an attachment.

Upon successful completion, `configure` will print a brief configuration summary.

2.6.1 Configure Usage

For a detailed list of `configure` variables and options, give `configure` the `--help` option:

```
$ ./configure --help
```

The following is a summary of the variables and options that are important when installing CitcomS.

2.6.2 Environment Variables

Environment variables may be specified as arguments to `configure`, e.g.,

```
$ ./configure CC=icc # use the Intel compiler
```

Variable	Description
CC	C compiler command. This is usually set to the name of an MPI wrapper command, e.g., <pre>./configure \ CC=/opt/mpich-1.2.6/bin/mpicc</pre> See Section 2.6.3 for details and examples.
CFLAGS	C compiler flags. This can be set for optimization flags.
CPPFLAGS	C preprocessor flags; e.g., <code>-I<dir></code> if you have headers in a nonstandard directory.
LDFLAGS	Linker flags; e.g., <code>-L<dir></code> if you have libraries in a nonstandard directory.
LIBS	Library flags; e.g., <code>-l<lib></code> if additional library is needed

2.6.3 MPI Configuration

By default, `configure` will search for a C compiler using your `PATH` environment variable. It prefers MPI wrapper commands (such as `mpicc`) to ordinary compiler commands (such as `cc` or `gcc`). You may specify the compiler command name manually using the `CC` variable:

```
$ ./configure CC=/opt/mpich-1.2.6/bin/mpicc
```

The `configure` script will test for the presence of the MPI header (`mpi.h`) and an MPI library using the C compiler command. If `CC` is set to an MPI wrapper command such as `mpicc`, and/or the MPI header and library files are installed in a standard location (i.e., `/usr/include` and `/usr/lib`), these `configure` tests should succeed without difficulty.

But if `CC` is set to an ordinary compiler command name (e.g., `cc` or `gcc`) and MPI is installed in a non-standard location, you must manually specify `CPPFLAGS` and `LDFLAGS`, so that the compiler can find the MPI header files and libraries.

2.6.3.1 Manually Specifying MPI `include` and `lib` Directories

```
$ ./configure \  

CPPFLAGS="-I/opt/mpich-1.2.6/include" \  

LDFLAGS="-L/opt/mpich-1.2.6/lib -lmpich"
```

2.6.3.2 Manually Specifying MPI include and lib Directories and an Alternative Compiler

```
$ ./configure \
CC=icc \
CPPFLAGS="-I/opt/mpich-1.2.6/include" \
LDFLAGS="-L/opt/mpich-1.2.6/lib -lmpich"
```

Note that it may be necessary to specify the name of the MPI library itself in LDFLAGS using the `-l` compiler option. If a library name is not given – or if the given option doesn't work – `configure` will automatically try linking using `-lmpi` and, if that fails, `-lmpich`.

2.7 HDF5 Configuration (Optional)

For writing its output in binary format, CitcomS requires parallel HDF5 (PHDF5). In turn, PHDF5 requires an MPI library with MPI-IO support and a parallel file system. If an existing installation of the PHDF5 library is not available on your cluster, you can compile it from source by following the instructions in the file `release_docs/INSTALL_parallel` under the HDF5 source tree. Under Debian Linux, you may simply install the `libhdf5-mpich`, `libhdf5-mpich-dev` and `hdf5-tools` packages.

NOTE: Most post-processing and visualization scripts discussed in Chapter 5 only understand ASCII output. An additional step to convert the HDF5 output files to ASCII files is necessary before using those scripts.

By default, CitcomS will attempt to auto-detect your PHDF5 installation, and will disable HDF5 support if it is not found. You may specify the location of your PHDF5 installation by setting the `PHDF5_HOME` environment variable to the appropriate installation prefix.

```
$ export PHDF5_HOME=/opt/phdf5/1.6.5
$ ./configure --with-hdf5
```

2.7.1 Additional Tools

While the following software is not necessary for the normal operation of CitcomS, you may find it useful for accessing CitcomS data in HDF5 files.

2.7.1.1 NumPy

NumPy is an extension to Python which adds support for multi-dimensional arrays for use in scientific computing. You may download NumPy from the NumPy home page (numpy.scipy.org). To compile and install this extension, download it and issue the following commands after extracting it:

```
$ cd numpy-1.0
$ python setup.py install --prefix=$HOME/cig
```

Alternatively, under Debian Linux you can install the `python-numpy` package. On Gentoo Linux, NumPy is available in the `dev-python/numpy` ebuild.

2.7.1.2 PyTables

PyTables is an extension to Python and can expose HDF5 array datasets as Python NumPy arrays. It is available at PyTables (www.pytables.org).

To compile and install this extension, download the latest stable version and issue the following commands:

```
$ cd pytables-1.3.3
$ python setup.py install --prefix=$HOME/cig
```

To install on Debian Linux, you may use the `python-tables` package instead. On Gentoo Linux, it is available in the `dev-python/pytables` ebuild.

2.7.1.3 HDFView

HDFView is a visual tool written in Java for browsing and editing HDF5 files. You may download it from the HDFView home page (hdf.ncsa.uiuc.edu/hdf-java-html/hdfview).

2.7.1.4 OpenDXutils

In order to import HDF5 files into OpenDX, you need the OpenDXutils package from the Cactus project. Go to the OpenDXutils package website (www.cactuscode.org/Visualization/openDX) and follow the instructions to download and install the package. Note that you will need to set both `DXMODULES` and `DXMDF` environment variables before running OpenDX to load the package.

2.8 GGRD Configuration (Optional)

CitcomS can read data from files for initial temperature, boundary conditions, material control and others. By default, these files are in ASCII format, and the data in the files must be ordered in the same way of the node ordering. As a result, whenever the mesh is changed, these data files need to be regenerated to satisfy the ordering requirement. This restriction can be lifted if using GRD data files. GRD files are binary files in NetCDF format (Network Common Data Form). GRD files are widely used in GMT (Generic Mapping Tools) as well.

The GRD file support can be enabled at configuration time. It requires the installation of HC package (geodynamics.org/cig/software/packages/mc/hc), a 1D global mantle circulation solver which can compute velocities, tractions, and geoid for simple density distributions and plate velocities. To install HC, you will also need the GMT and NetCDF packages (library and header files). See the `README.TXT` file in HC for instructions on installation. After HC is installed, note the directory (for example: `$HOME/cig/HC-1_0`) of HC installation, and execute these commands:

```
$ export HC_HOME=$HOME/cig/HC-1_0
$ ./configure --with-ggrd
```

You may need to set `GMTHOME` and `NETCDFHOME` environment variables if these packages is not installed in the system directory.

2.9 Installing from the Software Repository

The CitcomS source code is available via Git at the Geodynamics website (geodynamics.org). This allows users to view the revision history of the code and check out the most recent development version of the software.

NOTE: If you are content with the prepared source package, you may skip this section.

2.9.1 Tools You Will Need

In addition to the usual system requirements, you will need a handful of additional development tools installed in order to work with the source from the CIG software repository.

First, you must have a Git client installed. To check, type `git`; it should return a usage message.

```
$ git
usage: git [-version] [-help] [-C <path>] [-c name=value]
[-exec-path[=<path>]] [-html-path] [-man-path] [-info-path]
[-p|--paginate|--no-pager] [--no-replace-objects] [--bare]
[-git-dir=<path>] [--work-tree=<path>] [--namespace=<name>] <command>
[<args>]
```

For more information on Git, we recommend one of the following websites:

<http://github.com/> <http://git-scm.com/> <https://www.atlassian.com/git/>

Second, you must have the GNU tools Autoconf, Automake, and Libtool installed. To check, enter the following commands:

```
$ autoconf --version
$ automake --version
$ libtoolize --version
```

For more information about these GNU tools, see the GNU website (www.gnu.org/software). The CitcomS v3.3.0 source package was created with Autoconf 2.68, Automake 1.11.2, and Libtool 2.4.2.

2.9.2 Download Source from Subversion

To check out the latest version of the software, use the `git clone` command:

```
$ git clone https://github.com/geodynamics/citcoms
```

This will create the local directory `citcoms` (if it doesn't already exist) and fill it with the latest CitcomS source from the CIG software repository.

The CitcomS directory thus created is called a *working copy*. To merge the latest changes into an existing working copy, use the `git pull` update command:

```
$ cd citcoms
$ git pull
```

This will preserve any local changes you have made to your working copy.

2.9.3 Generating the GNU Build System

Your working directory should now contain a fresh checkout of CitcomS:

```
$ ls
CitcomS
$
```

Before you can run `configure` or `make`, you must generate the necessary files using the GNU tools. The easiest way to do this is to run `autoreconf -i`:

```
$ cd CitcomS
$ autoreconf -i
$ ./configure
$ make
```

The `autoreconf` tool runs Autoconf to generate the `configure` script from `configure.ac`. It also runs Automake to generate `Makefile.in` from `Makefile.am` in each source directory.

The `configure` script stores dependency information for source files in hidden subdirectories called `.deps`. If source files are added, deleted, moved, or renamed – which may happen if you run `svn update` to merge the latest changes – the `make` command may report errors for source files which no longer exist. If this happens, clean your existing configuration using the `make distclean` command, and then re-run `configure` and `make`:

```
$ make distclean
$ ./configure
$ make
```

The `make distclean` command deletes all files generated by `configure`, including the Makefiles themselves! Therefore, after running `make distclean`, you will not be able to run `make` again until you re-run `configure`.

Chapter 3

Running CitcomS

3.1 Using CitcomS

Upon the completion of a successful build, two binary executables, `CitcomSRegional` and `CitcomSFull`, are always placed under the `bin` directory. These programs are compiled from pure C code and do not use Python or the Pyre framework. Each program has the same usage:

```
$ mpirun [mpi_options] CitcomSRegional inputfile
$ mpirun [mpi_options] CitcomSFull inputfile
```

Two input file examples, one for a regional spherical model and one for a full spherical model, are provided in the `examples/Regional` and `examples/Full` directories, respectively. The meaning of the input parameters is described in Appendix A on page 79.

3.2 Changing Parameters

A configuration file should be used to set the input parameters for CitcomS. A configuration file is a plain text file whose format is based on the Windows INI file. The format of a configuration file is as follows:

```
#CitcomS.component1.component2
# this is a comment
property1=value1
property2=value2 ; this is another comment
property3=val1,val2,val3 ; this property is specified by a list of values
```

Upon termination of each run, all of the parameters are logged in a `pidXXXXXX` file, where `XXXXXX` is the process id of the CitcomS application. Section names such as `#CitcomS.component1.component2` are optional, but are recommended for the purpose of organizing the contents of the configuration file

3.3 Uniprocessor Example

CitcomS runs similarly in full spherical or regional modes. For the purpose of this example, you will perform a test run of the regional version on a workstation. The CitcomS source package contains an `examples` directory (the `make install` command installs the examples under `PREFIX/share/CitcomS/examples`, where `PREFIX` is the CitcomS installation directory). In this directory, you will find the `example0` configuration file. Switch to the `examples` directory and execute the following on the command line:

```
$ CitcomSRegional example0
```

or, equivalently,

```
$ mpirun -np 1 CitcomSRegional example0
```

3.3.0.1 Example: Uniprocessor, example0

```

# CitcomS
cpu_limits_in_seconds=360000000
minstep=5

# CitcomS.controller
storage_spacing=1

# CitcomS.solver
datafile=example0
rayleigh=100000

# CitcomS.solver.mesher
fi_max=1
fi_min=0
nodex=17
nodey=17
nodez=9
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.ic
num_perturbations=1
perturbl=1
perturblayer=5
perturbm=1
perturbmag=0.05

# CitcomS.solver.visc
num_mat=4

```

3.4 Multiprocessor Example

In the `examples` directory, type the following at the command line:

```
$ mpirun -np 4 CitcomSRegional example1
```

3.4.0.2 Example: Multiprocessor, example1

```

# CitcomS
cpu_limits_in_seconds=360000000
minstep=70

# CitcomS.controller
storage_spacing=10

# CitcomS.solver
datafile=example1
rayleigh=100000

# CitcomS.solver.mesher
fi_max=1
fi_min=0

```

```

nodex=17
nodey=17
nodez=9
nprocx=2
nprocy=2
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.ic
num_perturbations=1
perturbl=1
perturblayer=5
perturbm=1
perturbmag=0.05

# CitcomS.solver.visc
num_mat=4

```

This example has:

- 2 processors in colatitude (x -coordinate), specified by the `nprocx=2` line in the configuration file
- 2 processors in longitude (y -direction), specified by the `nprocy=2` line
- 1 processor in the radial (z -direction). The default value for the `nprocx`, `nprocy`, `nprocz` parameters is 1. Hence the `nprocz=1` line is omitted from the configuration file.
- The total number of processors used is given by `nprocx` × `nprocy` × `nprocz`. Hence this model uses 4 processors in total.
- There are 17 nodes in x (theta) direction.
- There are 17 nodes in y (fi) direction.
- There are 9 nodes in z (radius_inner) direction.

The model will run for 70 time steps and the code will output the results every 10 time steps.

It is important to realize that within the example script (and in finite element method, FEM) the term “node” refers to the mesh points defining the corners of the elements. In `example1`, this is indicated with:

```

nodex=17
nodey=17
nodez=9

```

These quantities refer to the total number of FEM nodes in a given direction for the complete problem, and for the example it works out that within a given processor there will be $9 \times 9 \times 9$ nodes. Note that in the x -direction (or y) that for the entire problem there are 17 nodes and there is one node shared between two processors. This shared node is duplicated in two adjacent processors.

Unfortunately, “nodes” sometimes also refer to the individual computers which make up a cluster or supercomputer or the number of CPUs in a given computer. We run one “process” on each node. In the example script this is indicated with:

```

nprocx=2
nprocy=2

```




Figure 3.1: Computational Domain. Map view on the configuration of the top layer of the computational nodes and the processors.

3.4.1 Output Formats

The possible output formats in CitcomS are `ascii`, `ascii-gz`, `hdf5`, `vtk` with `ascii` being the default. The output format is specified by the `output_format` configuration parameter. The ASCII output can potentially take a lot of disk space. CitcomS can write `gzip` compressed output directly. You can run the example script with:

```
output_format=ascii-gz
```

Be warned that the post-process scripts do not understand this output format yet.

Instead of writing many ASCII files, CitcomS can write its results into a single HDF5 (Hierarchical Data Format) file per time step. These HDF5 files take less disk space than all the ASCII files combined and don't require additional post-processing to be visualized in OpenDX. In order to use this feature, you must compile CitcomS with the parallel HDF5 library if you haven't done so already (see Section 2.7 on page 29). You can run the example script with:

```
output_format=hdf5
```

The output files will be stored in `./` with a filename prefix `example1` and a filename suffix `h5`. See Chapter 4 on the next page for more information on how to work with the HDF5 output.

Chapter 4

Working with CitcomS HDF5 Files

4.1 Introduction

A typical run of CitcomS can create thousands if not millions of ASCII output files. This situation is inefficient since it requires an extra post-processing step for assembling the results from each processor (see Chapter 5 on page 43). Since the v2.1 release, CitcomS solves this problem when running the software on computers that have parallel file systems by assembling a binary HDF5 file in parallel I/O mode. You can skip this chapter if you are not using the HDF5 output format.

4.2 About HDF5

The Hierarchical Data Format (HDF) is a portable file format developed at the National Center for Supercomputing Applications (NCSA) (hdf.ncsa.uiuc.edu/HDF5). It is designed for storing, retrieving, analyzing, visualizing, and converting scientific data. The current and most popular version is HDF5, which stores multi-dimensional arrays together with ancillary data in a portable self-describing format. It uses a hierarchical structure that provides application programmers with a host of options for organizing how data is stored in HDF5 files.

HDF5 files are organized in a hierarchical structure, similar to a Unix file system. Two types of primary objects, *groups* and *datasets*, are stored in this structure. A group contains instances of zero or more groups or datasets, while a dataset stores a multi-dimensional array of data elements. Both kinds of objects are accompanied by supporting metadata.

A dataset is physically stored in two parts: a header and a data array. The header contains miscellaneous metadata describing the dataset as well as information that is needed to interpret the array portion of the dataset. Essentially, it includes the name, datatype, dataspace, and storage layout of the dataset. The name is a text string identifying the dataset. The datatype describes the type of the data array elements. The dataspace defines the dimensionality of the dataset, i.e., the size and shape of the multi-dimensional array. The dimensions of a dataset can be either fixed or unlimited (extensible). The storage layout specifies how the data arrays are arranged in the file.

The data array contains the values of the array elements and can be either stored together in a contiguous file space or split into smaller *chunks* stored at any allocated location. Chunks are defined as equally-sized multi-dimensional subarrays (blocks) of the whole data array and each chunk is stored in a separate contiguous file space. Extensible datasets whose dimensions can grow are required to be stored in chunks. One dimension is increased by allocating new chunks at the end of the file to cover the extension.

HDF5 also supports access to portions (or selections) of a dataset by *hyperslabs*, which consist of a subarray or strided subarray of the multi-dimensional dataset. The selection is performed in the file dataspace for the dataset. HDF5 also supports parallel I/O. Parallel access is supported through MPI-IO. The file and datasets are collectively created/opened by all participating processes. Each process accesses part of a dataset by defining its own file dataspace for that dataset. When accessing data, the data transfer property specifies whether each process will perform independent I/O or all processes will perform collective I/O.

4.3 Input Parameters

To enable HDF5 output in CitcomS, all you need to do is include the following section in your `.cfg` input file.

```
[CitcomS.solver.output]
output_format = hdf5
```

Alternatively, you can specify the option `--solver.output.output_format=hdf5` on the command line. The resulting filenames will start with the value of your specified `datafile` input parameter and end with `.h5`.

4.3.1 Optimizing Parallel I/O

There are several platform-dependent parameters used by the HDF5 library and the underlying MPI-IO routines to optimize the performance of parallel I/O. The optimal values for these parameters may vary from file system to file system. Ideally, before compiling CitcomS, the build procedure would configure these parameters based on your platform, but this is not implemented currently.

In order to facilitate the process of gathering I/O performance data from a variety of parallel file systems such as GPFS, PVFS, IBRIX FusionFS, etc., you can specify the following parameters in the CitcomS input file. You may use these parameters to tune the performance of the parallel I/O on your system.

All values are assumed to be specified in bytes, unless otherwise indicated.

1. MPI file hints for collective buffering file access.
 - (a) `cb_block_size`: Target nodes will access data in chunks of this size.
 - (b) `cb_buffer_size`: Specifies the total buffer space on each target node. Set this parameter to a multiple of `cb_block_size`.
2. HDF5 file access properties.
 - (a) `sieve_buf_size`: Maximum size of data sieve buffer.
 - (b) `output_alignment`: Alignment interval size in bytes.
 - (c) `output_alignment_threshold`: Allocation requests above this size will be aligned on a memory address that is a multiple of `output_alignment`.
 - (d) `cache_rdcc_nelmts`: Maximum number of chunks that can be stored in the raw data chunk cache.
 - (e) `cache_rdcc_nbytes`: Size of raw data chunk cache in bytes.

For more details, you can refer to the following references:

- **MPI-2: Extensions to the Message-Passing Interface, section 9.2.8** (www-unix.mcs.anl.gov/mpi/mpi-standard/mpi-report-2.0/node182.htm) provides a list of MPI-IO reserved file hints, also available in Section 7.2.8 of the MPI-2 reference book.
- **HDF5 documentation** (hdf.ncsa.uiuc.edu/HDF5/doc/UG/08_TheFile.html), Chapter 2, Section 7.3, contains a list of HDF5 file access properties.
- **HDF5 User's Guide: Data Caching** (hdf.ncsa.uiuc.edu/HDF5/doc/Caching.html) offers a short explanation of raw data chunk caching.

You should also refer to the documentation for your particular parallel file system and check if there are any other MPI or HDF5 hints that could improve your parallel I/O performance.

Finally, here is an example section that would appear in a typical CitcomS input file:

```
[CitcomS.solver.output]
cb_block_size = 1048576          # 1 MiB
cb_buffer_size = 4194304        # 4 MiB
sieve_buf_size = 1048576        # 1 MiB
output_alignment = 262144       # 256 KiB
output_alignment_threshold = 524288 # 512 KiB
cache_rdcc_nelmts = 521
cache_rdcc_nbytes = 1048576
```

4.4 Data Layout

Time independent data (e.g., node coordinates, grid connectivity) are saved in a file (for example, `test-case.h5`) at the first output stage. Subsequently, each output stage will save time dependent data (e.g., velocity, and temperature) in a separate file (for example, `test-case.100.h5` contains data of time step 100). Most of the output data from CitcomS is specified at the nodes of a logically Cartesian grid and is therefore well represented by multi-dimensional arrays. A cap dimension is defined for addressing each of the CitcomS caps, followed by three spatial indices (i, j, k) in the case of 3D data, two spatial indices (i, j) in the case of 2D surface data, or one spatial index (k) in the case of 1D radial data. An additional dimension is provided for storing the components of vector and tensor data. These data arrays are stored in different groups. Each group has a descriptive name. In addition, there is an `/input` group archiving the input parameters of the model. Two sample `.h5` files are provided in `visual/samples/` directory.

Dataset	Shape
<code>/input</code>	N/A
<code>/coord</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez}, 3)$
<code>/connectivity</code>	$(\text{cap_elements}, 8)$

Table 4.1: Layout of the time-independent data file

Dataset	Shape
<code>/velocity</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez}, 3)$
<code>/temperature</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez})$
<code>/viscosity</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez})$
<code>/pressure</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez})$
<code>/stress</code>	$(\text{caps}, \text{nodex}, \text{nodey}, \text{nodez}, 6)$
<code>/geoid</code>	(8) See Section C.6.9 on page 95 for details
<code>/surf/velocity</code>	$(\text{caps}, \text{nodex}, \text{nodey}, 2)$
<code>/surf/heatflux</code>	$(\text{caps}, \text{nodex}, \text{nodey})$
<code>/surf/topography</code>	$(\text{caps}, \text{nodex}, \text{nodey})$
<code>/botm/velocity</code>	$(\text{caps}, \text{nodex}, \text{nodey}, 2)$
<code>/botm/heatflux</code>	$(\text{caps}, \text{nodex}, \text{nodey})$
<code>/botm/topography</code>	$(\text{caps}, \text{nodex}, \text{nodey})$
<code>/horiz_avg/temperature</code>	$(\text{caps}, \text{nodez})$
<code>/horiz_avg/velocity_xy</code>	$(\text{caps}, \text{nodez})$
<code>/horiz_avg/velocity_z</code>	$(\text{caps}, \text{nodez})$

Table 4.2: Layout of the time-dependent data file

4.5 Accessing Data

As previously indicated, HDF5 is a self-describing binary file format. As such we can use a variety of tools to inspect the structure of an HDF5 file, as well as for retrieving data in any order.

4.5.1 Inspecting Structures

To quickly inspect the structure of an HDF5 file, you can use the command `h5ls` which is included with the HDF5 software:

```
$ h5ls -r file.h5
```

4.5.2 Converting to ASCII Files

You can convert the HDF5 files to the ASCII combined capfiles described in Appendix C.6 for specific time steps by using the command included in CitcomS:

```
$ h5tocap modelname step1 [step2 [...] ]
```

You can also convert the HDF5 files to the `velo` files described in Appendix C.6.2 on page 94 for restart purposes by using the command included in CitcomS:

```
$ h5tovelo modelname step
```

4.5.3 Accessing Data in Python

The small Python script `h5tocap.py` provides a good example of using the PyTables extension module to access the data contained in the CitcomS HDF5 files. Using PyTables, datasets can be retrieved from disk as NumPy arrays. The retrieval avoids unnecessary copying of data by using hyperslabs, which take advantage of Python's powerful array slice-indexing.

For example, obtaining the node coordinates, temperature, and topography values over the entire surface of the sphere for time step 100 can be done easily with the following code snippet:

```
import tables

h5file = tables.openFile('samples/cookbook1.h5', 'r')
surface_coords = h5file.root.coord[0:12, :, :, -1, :]

data100 = tables.openFile('samples/cookbook1.100.h5', 'r')
surface_temperature = data100.root.temperature[0:12, :, :, -1]
surface_topography = data100.root.surf.topography[0:12, :, :]
```

In this case, the slice `0:12` refers to all caps explicitly, while the empty slice `“:”` refers to the entire extent of the corresponding dimension. The values of `-1` above refer to the last `z`-index, which corresponds to the location of the surface nodes on each of the caps. Finally, note how both HDF5 datasets and groups are conveniently accessible as Python attributes on the PyTables file object.

For more details, refer to the documentation for PyTables (www.pytables.org) and NumPy (numpy.scipy.org).

4.5.4 Accessing Data Using HDFView

NCSA HDFView is a visual tool for accessing HDF files. You can use it for viewing the internal file hierarchy in a tree structure, creating new files, adding or deleting groups and datasets, and modifying existing datasets. HDFView is capable of displaying 2D slices of multi-dimensional datasets, with navigation arrow buttons that enable you to range over the entire extent of a third dimension. An example of using HDFView is found in Figure 4.1 which uses data from Cookbook 1 in Section 6.2 on page 51.

The screenshot shows the HDFView application window. The title bar reads "HDFView". The menu bar includes "File", "Window", "Tools", and "Help". The "File/URL" field shows the path: `/home/luis/CitcomS/visual/samples/cookbook1.h5`.

The left panel displays a tree view of the file structure:

- cookbook1.h5
 - connectivity
 - coord
 - input
- cookbook1.100.h5
 - botm
 - heatflux
 - topography
 - velocity
 - geoid
 - horiz_avg
 - temperature
 - velocity_xy
 - velocity_z
 - pressure
 - stress
 - surf
 - heatflux
 - topography
 - velocity
 - temperature
 - velocity
 - viscosity

The right panel, titled "TableView - /home/luis/CitcomS/visual/samples/cookbook1.h5 - /connecti...", displays a 2D data slice as a table with 22 rows and 8 columns. The columns are labeled 1 through 6, with an additional column on the far right. The data values range from 0 to 114.

	1	2	3	4	5	6	
1	0	9	90	81	1	10	91
2	1	10	91	82	2	11	92
3	2	11	92	83	3	12	93
4	3	12	93	84	4	13	94
5	4	13	94	85	5	14	95
6	5	14	95	86	6	15	96
7	6	15	96	87	7	16	97
8	7	16	97	88	8	17	98
9	9	18	99	90	10	19	100
10	10	19	100	91	11	20	101
11	11	20	101	92	12	21	102
12	12	21	102	93	13	22	103
13	13	22	103	94	14	23	104
14	14	23	104	95	15	24	105
15	15	24	105	96	16	25	106
16	16	25	106	97	17	26	107
17	18	27	108	99	19	28	109
18	19	28	109	100	20	29	110
19	20	29	110	101	21	30	111
20	21	30	111	102	22	31	112
21	22	31	112	103	23	32	113
22	23	32	113	104	24	33	114

The bottom panel shows the metadata for the selected group or dataset:

```

cp = 1200.0
density_above = 1030.0
density_below = 6600.0
surftemp = 273.0
z_lith = 0.014
z_410 = 0.06435
z_lmantle = 0.105
z_cmb = 0.439
radius_km = 6371.0
  
```

At the bottom of the window, there are tabs for "Log Info" and "Metadata".

Figure 4.1: A screenshot of HDFView. The left panel shows the hierarchy of the groups and datasets. The right panel shows a 2D slice of a dataset. The bottom panel shows the metadata associated with the selected group or dataset.

Chapter 5

Postprocessing and Graphics

5.1 Introduction

Once you have run CitcomS, you would have a series of output files (potentially spread throughout the file systems of your Beowulf cluster or in a set of directories on your parallel file system). You now have to retrieve and combine the data for the time step (or age) of interest. To visualize your results, it is recommended that you use the open-source Open Visualization Data Explorer, better known as OpenDX. The software is available from the OpenDX website (www.opendx.org). If you are using Linux, OpenDX is usually available as package `dx` or `opendx` in your distribution. If you are using Mac OS X, OpenDX is available via Fink (fink.sourceforge.net). We provide experimental scripts for a 3D visualization program called MayaVi2; see Section 5.7 on page 47. Scripts using GMT commands to plot 2D cross sections of temperature field are also provided; see Section 5.8 on page 49.

5.2 Postprocessing on a Beowulf Cluster

Generally, the results from your CitcomS run will be distributed on disks attached to individual nodes of your Beowulf cluster. The output files are written in each node under the directory that you specified as the `datadir` property in the input file. To examine those files, log in to a node and change directories to the one you specified with a prefix. For example, if you set `datadir=/scratch/username` and `datafile=test-case` in your input file, then your output files will be written to `/scratch/username` and will have the prefix `test-case`.

If you select HDF5 for the output format, the output files will have a `.h5` suffix. You won't need to postprocess the `.h5` files; you can visualize the results using OpenDX. See Section 5.6 on page 46 for more details.

If you select compressed ASCII (`ascii-gz`) for the output format, there is no post-processing script for it. You need to decompress and rename the output files before you can use the post-processing script described below.

If you select ASCII for the output format, you will have many output files. An example of a filename for the velocity output is `test-case.velo.2.10` where `test-case` is the model prefix, `velo` means that this is a velocity data file, `2` corresponds to the processor number (i.e., it is output by the 2nd processor), and `10` corresponds to the time step.

If your run used the time-dependent velocity boundary conditions (`solver.param.file_vbcs=on`), the log file will also have a `current age` line that lists the time-step numbers and their corresponding times. To choose an age to export for postprocessing, you have to determine which time step corresponds to the age that interests you by looking at the log file.

When you execute a CitcomS run, your input parameters will be saved in a file `pidxxxxx.cfg` where `xxxxx` is usually a five-digit number for the process ID. This pidfile contains most of the input parameters, which can be useful for archiving and postprocessing.

The ASCII output files of CitcomS need to be postprocessed before you can perform the visualization. The script `autocombine.py` can postprocess (retrieve and combine) CitcomS output; it will retrieve CitcomS data to the current directory and combine the output into a few files.

Using `autocombine.py`, retrieve and combine data of time-step 10:

```
$ autocombine.py mpirun.nodes pid12345.cfg 10
```

This reads the MPI machinefile (`mpirun.nodes`) and the CitcomS pidfile (`pid12345.cfg`), then calls other scripts to do the actual job. The general usage for `autocombine.py` is:

```
$ autocombine.py [machinefile] [pidfile] \
[step1] [step2 or more ...]
```

If your Beowulf cluster uses `ssh` (rather than `rsh`) to access the computation nodes, you must manually edit `visual/batchpaste.s` and replace ‘`rsh`’ with ‘`ssh`’ in the script.

Once `autocombine.py` has run, you will have 2 files (or 24 files for the full spherical version of CitcomS) formatted as follows:

```
test-case.cap00.10
test-case.cap00.10.general
```

The former file is the data file containing simulation results and is referred to as the “capfile”; its format can be found in Appendix C. The latter file is the OpenDX header for the data.

If you have enabled some `optional_output`, and the optional fields are node-based (e.g. `comp_nd`, `stress`, or `pressure`), these fields will be combined in separated files, with filenames:

```
test-case.opt00.10
test-case.opt00.10.general
```

The former file is the data file containing optional fields and is referred to as the “optfile.” Its format is column-based ASCII data. The latter file is the OpenDX header for the data. The header file contains information on the ordering of the optional fields.

5.3 Postprocessing in a Non-Cluster Environment

If you run CitcomS in a non-cluster environment or all of your data can be accessed from the local machine, you can still use `autocombine.py` to combine the data. In this case, the `machinefile` is not needed and can be replaced by `localhost`, such as:

```
$ autocombine.py localhost [pidfile] \
[step1] [step2 or more ...]
```

5.4 Using OpenDX for Regional Sphere Visualization

OpenDX modules designed for CitcomS can be found in the source directory called `visual`. The optional `make install` command installs the OpenDX modules under `PREFIX/share/CitcomS/visual` (where `PREFIX` defaults to `/usr/local`).

In this example, you will use `visRegional.net` to visualize the results of regional CitcomS.

1. Launch OpenDX by typing:

```
$ dx
```

You will see an OpenDX Data Explorer window.

2. Click **Edit Visual Programs** and select `visRegional.net` from the file selector.
3. You will see a **Visual Program Editor** window.
4. Select the **import** tab in the **Visual Program Editor**’s main window and double-click on the **FileSelector** block, which will open a **Control Panel** window.
5. In the **CitcomSImport** filename input box, select the header file of your postprocessed data, e.g.,

5.5 Using OpenDX for Full Sphere Visualization

1. After launching OpenDX, you will see an OpenDX Data Explorer window.
2. Click **Edit Visual Programs** and select `visFull.net` from the file selector.
3. You will see a **Visual Program Editor** window.
4. Select the **import** tab in the **Visual Program Editor**'s main window and double-click on the **FileSelector** block, which will open a **Control Panel** window.
5. In the **Format String** of **CitcomSFullImport** input box, select one of the 12 header files of your postprocessed data, e.g.,

```
samples/fulltest.cap00.100.general
```

6. The results of a full spherical CitcomS consist of 12 cap files. In order to import the 12 cap files at the same time, edit the filename with the cap number replaced by printf-styled format string `%02d`, e.g.,

```
samples/fulltest.cap%02d.100.general
```

7. In the pull-down menu, select **Execute** ▸ **Execute on Change**.
8. A new window will appear with the image of the model.
9. If you want to zoom, rotate, change projection, and otherwise manipulate the view of the model, experiment with the menu **Options** ▸ **View Control**.

Additional options in the control panel window for the spherical model include:

- **CitcomSFullImport reduced** can increase or reduce the resolution of the velocity vectors and grids. Reducing resolution is often necessary for visualizing large dataset interactively. Note that the resolution of other fields (e.g., temperature and viscosity) is not reduced.
- **Core radius** is the size of the orange sphere which represents the core-mantle boundary.
- **Isosurface value** is the temperature isosurfaces. You can change the values of the isosurfaces, including adding additional isosurfaces or deleting existing ones.
- **Latitude of normal axis** and **Longitude of normal axis** are the directions of the normal axis to the cross-section plane.

5.6 Using OpenDX for HDF5 Visualization

If you use the HDF5 output format (`solver.output.output_format=hdf5`), you can directly visualize the data without postprocessing. First, you need to install and set up the **OpenDXutils** package (see Section 2.7.1.4 on page 30). Then, open either `visRegional.net` or `visFull.net` in OpenDX.

1. In the pull-down menu, select **File** ▸ **Load Macro**, and select the file `CitcomSImportHDF5.net` to load it.
2. In the **Tools** panel of the main window, select the `CitcomSImportHDF5` module in the **Macros** category (highlighted in Figure 5.2).
3. Place the module in the work space and rewire the network as shown in Figure 5.2.
4. There are four input tabs in the `CitcomSImportHDF5` module.
 - (a) The first tab (connected to the left **FileSelector** module) specifies the HDF5 file containing time-independent information (e.g., `samples/cookbook1.h5`).




Figure 5.2: How to import CitcomS HDF5 data. The CitcomSImportHDF5 module is highlighted in the Tools panel.

- (b) The second tab (connected to the right FileSelector module) specifies the HDF5 file containing time-dependent information (e.g., `samples/cookbook1.100.h5`).
- (c) The third tab (connected to the Integer module) specifies the resolution reduction factor.
- (d) The fourth tab (unconnected, default to 0) specifies which cap(s) to import.
 - i. To visualize a regional model, leave this tab unchanged.
 - ii. To visualize a full spherical model, double-click this module and change the value of this tab to `{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11}`.
 - iii. You can also specify a subset of caps to visualize.

5. After the data are imported, the visualization can be manipulated as described in previous sections.

5.7 Using MayaVi for Visualization

This distribution also comes with scripts for visualizing CitcomS using MayaVi2. MayaVi2 is the successor of MayaVi for 2D/3D scientific data visualization. It is an interactive program allowing elaborate plots of scientific data. MayaVi2 relies on the Visualization Toolkit (VTK) and allows easy scripting in Python. To install MayaVi2, you should follow the instructions on the MayaVi2 website (<https://svn.enthought.com/enthought/wiki/>

GrabbingAndBuilding) and on the SciPy website (www.scipy.org/Cookbook/MayaVi/Installation). Besides MayaVi2, before using the script you will need to install PyTables (see Section 2.7.1.2 on page 29) and PyVTK; see PyVTK website (cens.ioc.ee/projects/pyvtk) for instructions.

The script acts as an extension module of MayaVi2. The installation of the script is still preliminary. You will need to include `PREFIX/share/CitcomS/visual/Mayavi2/citcoms_plugins/` in your `PYTHONPATH` environment variable. You also need to create a directory called `~/ .mayavi2` and copy into it the following:

```
PREFIX/share/CitcomS/visual/Mayavi2/mayavi_custom_ui.py
```

To run the script, type this command:

```
$ mayavi2_citcoms_display.py file.step.h5
```

This will launch MayaVi2 and load the data from `file.step.h5`. An example of using this script is found in Figure 5.3. You may adjust your visualization pipeline by adding any appropriate VTK filters and modules. For more details about this process, refer to *The Visualization Toolkit User's Guide* (ISBN 1-930934-13-0) or see the VTK website (www.vtk.org).




Figure 5.3: Example MayaVi2 visualization of `cookbook1.100.h5`. Here we display the velocity field as vectors, as well as a slice of the temperature field. Note the visualization pipeline in the tree view on the left panel.

5.8 Using GMT Commands for Visualization

The Generic Mapping Tools (GMT) is a collection of command-line tools for manipulating geographic and Cartesian data that produces PostScript (PS) or Encapsulated PostScript File (EPS) illustrations. GMT is widely used in the geophysics community. Two scripts `plot_layer.py` and `plot_annulus.py`, which can plot the temperature field in horizontal and radial cross sections, respectively, are provided. These scripts use GMT commands to generate EPS images. GMT is very customizable. Users might wish to customize the scripts for their need.

The usage of `plot_layer.py` is:

```
$ plot_layer.py modelname caps step layer
```

The script will look for the capfile(s) `modelname.cap00.step`, if `caps` is 1, or `modelname.cap00.step`, `modelname.cap01.step`, ... through `modelname.cap11.step`, if `caps` is 12. A slice of the capfile at the specified radial layer is then plotted in Mercator projection, if `caps` is 1, or in Hammer-Aitoff projection, if `caps` is 12.




Figure 5.4: Temperature field generated by `plot_layer.py`. This image is the mid-layer (layer 5) of the sample data `visual/samples/regtest.cap00.100`. The resolution of the image is poor because the data has only $9 \times 9 \times 9$ nodes.

The usage of `plot_annulus.py` is:

```
$ plot_annulus.py modelname step
```

The script will ask you a few questions on how to specify the great circle path. There are three options: (1) a starting point and an azimuth, (2) a starting point and another point on the great circle, and (3) a starting point and a rotation pole. The great circle path is used to construct the radial cross section.




Figure 5.5: Temperature fields generated by `plot_annulus.py`. The top panel is the horizontal cross section at the mid layer in Hammer-Aitoff projection. The thick line marks the surface track of the radial cross section. The bottom panel is the radial cross section. The cross section starts at 0° latitude and 0° longitude and has an azimuth of 45° . The image is from the sample data `visual/samples/fulltest.cap*.100`. There is always a gap on the left (starting point) of the cross section. Users familiar with GMT are welcome to contribute a fix to the script.

Chapter 6

Cookbooks

6.1 Introduction

These cookbook examples are meant to serve as a guide to some of the types of problems CitcomS can solve. Cookbook examples range from regional to full spherical shell problems that address traditional mantle convection problems. These cookbook examples are distributed with the package under the `examples` directory. However, you might need to edit these example scripts slightly to launch the job on your cluster (see Section 3.4 on page 34 for more information).

Cookbooks 1 to 4 introduce the basic parameters and are suitable for all users. Cookbook 5 introduces time-dependent velocity boundary conditions, reading in initial temperature field from files and tuning of the advection solver. Cookbook 6 introduces the pseudo-free-surface formulation for studying short wavelength dynamics topography. Cookbook 7 introduces thermo-chemical convection problem and tuning of the Conjugate Gradient velocity solver. Cookbook 8 introduces compressible convection problem, checkpointing/restarting, geoid and tuning of the Multigrid and Uzawa velocity solver.

6.2 Cookbook 1: Global Model

6.2.1 Problem

This example solves for thermal convection within a full spherical shell domain. The full spherical version of CitcomS is designed to run on a cluster that decomposes the spherical shell into 12 equal “caps” and then distributes the calculation for caps onto separate processors. To run CitcomS with the full solver parameter set, it is recommended that you have a minimum of 12 processors available on your cluster.




Figure 6.1: Global Model “Caps.” Left (A): Three-dimensional perspective image showing seven of the 12 spherical caps used in a full CitcomS run. Right (B): The temperature field at 1081 km depth from a Cookbook 1 run.

6.2.2 Solution

You will use `cookbook1`. The first set of parameters specifies that you are going to use the full spherical version of the solver. The default is to use the regional spherical version, so this must be set.

```
solver=full
```

The second set of parameters specifies the number of time steps (100), how often full outputs of the computation are created (25), and the prefix of output filenames (`cookbook1`).

```
steps=100
storage_spacing=25
datafile=cookbook1
```

The following parameters control the temperature field for the initial conditions. The last set of parameters includes the number of perturbations to the initial temperature (1), the number of nodal lines of the perturbation in the longitudinal direction, e.g., the spherical harmonic order (3), the number of nodal lines in the latitudinal direction, e.g., the spherical harmonic degree (2), which layer the perturbation is on (5), and the amplitude of the perturbation (0.05). Note that although the number of perturbations is assigned here as `num_perturbations=1`, that is actually the default value.

```
num_perturbations=1
perturb1=3
perturbm=2
perturb1ayer=5
perturbmag=0.05
```

This example is executed by typing

```
$mpirun -np 12 CitcomSFull cookbook1
```

6.2.2.1 Example: Global Model, `cookbook1`

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=100
solver=full

# CitcomS.controller
storage_spacing=25

# CitcomS.solver
datafile=cookbook1
rayleigh=100000

# CitcomS.solver.mesher
fi_max=1
fi_min=0
nodex=9
nodey=9
nodez=9
nproc_surf=12
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.ic
num_perturbations=1
```

```
perturb1=3  
perturblayer=5  
perturbm=2  
perturbmag=0.05
```

```
# CitcomS.solver.visc  
num_mat=4
```

Once you have run the model, you can visualize the results using OpenDX, described in the previous chapter. When you invoke “Edit Visual Program,” select `visFull.net`.




Figure 6.2: Cookbook 1: Global Model. This image, created by OpenDX, depicts a slice through a spherical model of convection, with warmer colors indicating upwelling and the cooler colors showing downwelling.

6.2.3 Discussion

You have generated a simple example of the full CitcomS model, with minimal parameter alterations. With a default Rayleigh number of 10^5 and perturbation on the initial temperature field which has a degree-3 and an order-2 pattern, after 100 time steps, the convection pattern remains relatively steady.

6.3 Cookbook 2: Domain Size and Velocity Boundary Conditions

6.3.1 Problem

This example solves for thermal convection with velocity boundary conditions imposed on the top surface within a given region of a sphere. This requires using the regional version of CitcomS.

6.3.2 Solution

You will use `cookbook2`. The parameters specify the number of time steps (60), the prefix of the output filenames (`cookbook2`), and how often outputs will be created (30).

```
steps=60
storage_spacing=30
datafile=cookbook2
```

The `solver.mesher` facility has several properties involved in the generation of the computational mesh. Continue to use the default values for the physical portion of the domain in which you are interested. However, try modifying the layout of the mesh as shown. The parameters specify two processors per cap in the x - and y -direction, one processor per cap in the z -direction (by default), 17 nodes per cap in the x - and y -direction, 9 nodes in the z -direction, and the minimum/maximum extents in the colatitudal, longitudinal, and radial (`theta`, `fi`, `radius`) directions. The unit for `theta` and `fi` is radian.

```
nprocx=2
nprocy=2
nodex=17
nodey=17
nodez=9

theta_min=0.7854
theta_max=1.5708
fi_min=0.0
fi_max=0.7854
radius_inner=0.55
radius_outer=1.0
```

The following parameters allow you to impose a uniform velocity across the top surface. You have a velocity which is purely in the colatitude direction with a non-dimensional velocity of 100 (see Equation 1.14 on page 19 for how to dimensionalize it).

```
topvbc=1
topvbxval=100
topvbyval=0
```

In addition, the initial temperature field has a linear conductive profile. The amplitude of initial temperature perturbation is set to zero using the following parameters.

```
num_perturbations=1
perturbmag=0.0
```

6.3.2.1 Example: Velocity Boundary Conditions, `cookbook2`

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=60
```

```
# CitcomS.controller
storage_spacing=30      ; how often outputs are created

# CitcomS.solver
datafile=cookbook2     ; prefix of output filenames
rayleigh=100000

# Modify the layout of the mesh
# CitcomS.solver.mesher
nodex=17
nodey=17
nodez=9
nprocx=2
nprocy=2

theta_max=1.5708
theta_min=0.7854
fi_max=0.7854
fi_min=0.0
radius_inner=0.55
radius_outer=1.0

# Import a uniform velocity across the top surface.
# CitcomS.solver.bc
topvbc=1
topvbxval=100
topvbyval=0

# In addition, set the initial temperature perturbation to zero.
# CitcomS.solver.ic
num_perturbations=1
perturbl=1
perturblayer=5
perturbm=1
perturbmag=0.0

# CitcomS.solver.visc
num_mat=4
```

6.3.3 Discussion

Using OpenDX to visualize the results (Figure 6.3), this model allows you to create a plate-driven convection in which there is a thermal upwelling on one wall, a thermal downwelling on another, and uniform horizontal velocity across the top. The downwelling is not exactly subduction because the sidewalls have reflecting boundary conditions. This means that there is a symmetrical downwelling in a vertical domain on the other side.




Figure 6.3: Cookbook 2: Velocity Boundary Conditions. This model, visualized with OpenDX, highlights a region of the sphere and the heated upwellings (warm colors), downwellings (cool colors), and the velocities (yellow arrows).

6.4 Cookbook 3: Temperature-Dependent Viscosity

6.4.1 Problem

A common problem in geophysics is the exploration of natural convection in the presence of variable viscosity, including temperature-dependent or stress-dependent viscosity.

6.4.2 Solution

You will use `cookbook3`. The parameters specify the number of time steps (200), how often outputs will be created (25), the prefix of the output filenames (`cookbook3`), and the Rayleigh number (10^6).

```
steps=200
monitoringFrequency=25
datafile=cookbook3
rayleigh=1e6
```

The parameters specify whether the viscosity should be updated at every time step (`on`), the number of viscous layers (4), the viscosity of each layer (1,1,1,1), whether the viscosity is temperature dependent (`on`), the type of viscosity law to use (4), the activation energy of each layer (0.2,0.2,0.2,0.2), the temperature offset of each layer (0,0,0,0), the activation volume of each layer (0,0,0,0), whether to apply the minimum cutoff (`on`), the value of the minimum cutoff (1.0), whether to apply the maximum cutoff (`on`), and the value of the maximum cutoff (100.0).

```
VISC_UPDATE = on
num_mat = 4
visc0 = 1,1,1,1
TDEPV = on
rheol = 4
viscE = 0.2,0.2,0.2,0.2
viscT = 0,0,0,0
viscZ = 0,0,0,0
```

```

VMIN = on
visc_min = 1.0
VMAX = on
visc_max = 100.0 ===== VISC_UPDATE=on
num_mat=4
visc0=1,1,1,1
TDEPV=on
rheol=4
viscE=0.2,0.2,0.2,0.2
viscT=0,0,0,0
viscZ=0,0,0,0
VMIN=on
visc_min=1.0
VMAX=on
visc_max=100.0 >>>> python-removal

```

The range of the layers are determined by the following parameters: layer 1 extends from the top surface to a depth of `solver.const.z_lith` (default to 0.014, or 90 km dimensionally); layer 2 is below layer 1 and extends to a depth of `solver.const.z_410` (default to 0.06435, or 410 km dimensionally); layer 3 is below layer 2 and extends to a depth of `solver.const.z_lmantle` (default to 0.105, or 670 km dimensionally); layer 4 is below layer 3 and extends to the bottom. These depth parameters also control the depth of various phase transitions.

There are several viscosity laws coded in the program, which can be selected by the `rheol` parameter. All available viscosity laws can be found in Appendix A.1.11. For `rheol=4`, the temperature dependence of the viscosity of each layer is determined by:

$$visc = visc0 \times \exp\left(\frac{viscE + viscZ \times (1 - r)}{T^* + viscT} - \frac{viscE + viscZ \times (1 - r_{inner})}{1 + viscT}\right) \quad (6.1)$$

where $T^* = \min(\max(T, 0), 1)$. This equation is non-dimensionalized. The last term on the right-hand side normalizes the viscosity such that $visc = visc0$ when $T = 1$ and $r = r_{inner}$. Comparing Equation 6.1 with the dimensional Arrhenius activation equation:

$$\begin{aligned} \eta &= \eta_0 \exp\left(\frac{E_a + PV_a}{RT}\right) = \eta_0 \exp\left(\frac{E_a + PV_a}{R\Delta T (T' + T'_0)}\right) \\ &= \eta_0 \exp\left(\frac{E_a + \rho g (1 - r) V_a}{R\Delta T (T' + T'_0)}\right) = \eta_0 \exp\left(\frac{E_a/R\Delta T + (1 - r) \rho g V_a/R\Delta T}{T' + T'_0}\right) \end{aligned} \quad (6.2)$$

where E_a and V_a are the activation energy and activation volume, P is the hydrostatic pressure, R is the universal gas constant, and other symbols are as defined in Section 1.3, it is clear that

$$viscE = \frac{E_a}{R\Delta T} \quad (6.3)$$

$$viscZ = \frac{\rho g V_a}{R\Delta T} \quad (6.4)$$

$$viscT = T_0 \quad (6.5)$$

6.4.2.1 Example: Temperature-Dependent Viscosity, `cookbook3`

```

# CitcomS
cpu_limits_in_seconds=360000000
minstep=200

```

```

# CitcomS.controller
storage_spacing=25

```

```
# CitcomS.solver
datafile=cookbook3
rayleigh=1e6

# CitcomS.solver.mesher
fi_max=1
fi_min=0
nodex=17
nodey=17
nodez=9
nprocx=2
nprocy=2
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.ic
num_perturbations=1
perturbl=1
perturblayer=5
perturbm=1
perturbmag=0.05

# CitcomS.solver.visc
TDEPV=on
VISC_UPDATE=on
VMAX=on
VMIN=on
num_mat=4
rheol=4
visc0=1,1,1,1
viscE=0.2,0.2,0.2,0.2
viscT=0,0,0,0
viscZ=0,0,0,0
visc_max=100.0
visc_min=1.0
```




Figure 6.4: Cookbook 3: Temperature-Dependent Viscosity. This model, visualized with OpenDX, highlights a region that features both the heated upwelling (warm colors) and the even distribution of the velocities (arrows).

6.5 Cookbook 4: Regionally Refined Meshes

6.5.1 Problem

Frequently for convection problems, particularly those with variable viscosity, there are features with strong gradients in temperature or viscosity that can be better resolved with refined meshes. For example, for the problem just studied in Cookbook 3, temperature-dependent viscosity, a higher resolution is required for a narrow hot upwelling while a lower resolution is sufficient for the wider cold downwelling.

6.5.2 Solution

The parameter that controls whether a mesh is to be uniform or refined is `coor`. Set it to 1 (`coor=1`) in order to read the uneven mesh point coordinates from an input file (specified by `coor_file`). The format of the coordinate input file is described in Appendix B.2 on page 89.

```
coor=1
coor_file=coor.dat
```

The computational domain is bounded in colatitude between radian 1 and 2, in latitude between radian 0 and 1, and in radius between 0.55 and 1, as determined by the content of `coor_file`.

6.5.2.1 Example: Regionally Refined Meshes, `cookbook4`

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=250

# CitcomS.controller
storage_spacing=50

# CitcomS.solver
datafile=cookbook4
rayleigh=1e6

# CitcomS.solver.mesher
coor=1
coor_file=coor.dat
fi_max=1
fi_min=0
nodex=33
nodey=17
nodez=17
nprocx=1
nprocy=1
nprocz=1
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.ic
num_perturbations=1
perturb1=1
perturb1ayer=10
perturbm=0
perturbmag=0.05
```

```

# CitcomS.solver.visc
TDEPV=on
VISC_UPDATE=on
VMAX=on
VMIN=on
num_mat=4
rheol=4
visc0=1,1,1,1
viscE=0.2,0.2,0.2,0.2
viscT=0,0,0,0
viscZ=0,0,0,0
visc_max=100.0
visc_min=1.0

```




Figure 6.5: Cookbook 4: Regionally Refined Meshes. This model shows the temperature (upwelling – warm colors, downwelling – cool colors) and the uneven distribution of the velocities (yellow arrows). Note the refined mesh for the left and bottom regions of the model (inset).

6.5.3 Discussion

The resulting model is like a 2D model, but extends in the longitudinal direction. To set up this type of model, the initial temperature gradient was perturbed only in the longitudinal direction, by setting the parameters `perturb1=1` and `perturbm=0`. The model results show a thin thermal upwelling on one wall and a wide thermal downwelling on the opposite wall. Also, at the bottom of the model a thin layer of hot material can be shown. Note the higher resolution in the narrow regions with hot material and the lower resolution for the wide thermal downwelling region.

6.6 Cookbook 5: Subduction Models with Trench Rollback

6.6.1 Problem

A common issue to address is the problem arising when the position of the oceanic trench is not constant in time (trench rollback) for a subduction zone. In addition, the trench rollback speed may vary in time, which can be caused, for example, by a rotation pole jump.

6.6.2 Solution

In order to introduce in a convection model a trench rollback you have to prepare the velocity files. The following scenario is proposed for this cookbook: there are two plates centered along the equator, with two different Euler poles (see Figure 6.6). Initially, both plates have Euler poles situated far apart, so the velocities are approximately constant (about 5 cm/yr for the left plate and about 2 cm/yr for the right plate). After 30 Ma, a pole jump occurred for the right slab only, so the Euler pole moved closer.

Since we impose a top velocity boundary in our model, the following parameters should be turned on:

```
# CitcomS.solver.bc
topvbc=1

# CitcomS.solver.param
file_vbcs=on
```

The following two lines specify the location of the velocity files and the starting age for the model:

```
start_age=55
vel_bound_file=./velocity/bvel.dat
```

Since the starting age is set to 55 Ma, there will be 57 velocity files, one for each Ma (`vel.dat0`, `vel.dat1`, ..., `vel.dat55`, `vel.dat56`). The format of the velocity file is described in Appendix B.3 on page 90. Note that the velocity file uses a unit of cm/yr.

When `tic_method=-1` is set, the initial temperature field is read from setup files. The files have the same naming scheme and format as the CitcomS `velo` output, as described in Appendix C.6.2 on page 94. The location and prefix of the file is specified in `datadir_old` and `datafile_old`, and the time step of the `velo` file in `solution_cycles_init`. In this example, a processor of MPI rank `n` will read file `./ic/cookbook5.velo.n.0`

```
# CitcomS.solver
datadir_old=./ic
datafile_old=cookbook5

# CitcomS.solver.ic
tic_method=-1
solution_cycles_init=0
```

The descending slab will cause some numerical instability in the temperature solver. To reduce the numerical instability, you will use a smaller time-step size (0.75 of the maximum time-step size for a stable solution). The change in maximum temperature will also be monitored. If the maximum temperature increases too much (> 5%) between time steps, the temperature solver will rerun with half time-step size.

```
# CitcomS.solver.tsolver
finetunedt=0.75
monitor_max_T=on
```




Figure 6.6: Cookbook 5: Left (A): Initially, both plates have Euler poles (magenta and green dots) situated far apart from the plate, so the velocities are approximately constant (about 5 cm/yr for the left plate and about 2 cm/yr for the right plate). Right (B): After 30 Ma a pole jump occurred for the right slab only, so the Euler pole moved closer.

6.6.2.1 Example: Subduction Models with Trench Rollback, cookbook5

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=1100
```

```
# CitcomS.controller
storage_spacing=100
```

```
# CitcomS.solver
datadir_old=./ic
datafile=cookbook5
datafile_old=cookbook5
rayleigh=4.07e+08
```

```
# CitcomS.solver.mesher
coor=1
coor_file=./coor.dat
fi_max=1
fi_min=0
nodex=17
nodey=65
nodez=33
nprocx=1
nprocy=2
nprocz=1
theta_max=2.0708
theta_min=1.0708
```

```
# CitcomS.solver.tsolver
finetunedt=0.75
monitor_max_T=on
```

```
# CitcomS.solver.bc
topvbc=1
```

```
# CitcomS.solver.ic
num_perturbations=1
perturbl=1
```

```

perturblayer=5
perturbm=1
perturbmag=0.05
solution_cycles_init=0
tic_method=-1

# CitcomS.solver.param
file_vbcs=on
start_age=55
vel_bound_file=./velocity/bvel.dat

# CitcomS.solver.visc
TDEPV=on
VMAX=on
VMIN=on
num_mat=4
visc0=100,0.003,1,2
viscE=24,24,24,24
viscT=0.182,0.182,0.182,0.182
visc_max=100.0
visc_min=0.01

```




Figure 6.7: Cookbook 5: The pole jump after 30 Ma produced a flat slab on one side (fore side) and a steep slab on the other side (back side).

6.6.3 Discussion

The results for this problem are presented in Figure 6.7. Since the two Euler poles are kept fixed for the first 30 Ma, the shape of the subducting slab will be the same along the trench. At 25 Ma, the Euler pole for the right plate jumps toward the plate. Therefore, the velocity along the trench varies from about 1 cm/yr to about 2.5 cm/yr. This will produce in time a flat slab at one side of the model, and a steep slab at the other side.

6.7 Cookbook 6: Pseudo-Free-Surface Formulation

6.7.1 Problem

Free-slip boundary conditions are typically applied on the top surface of mantle convection models, and the dynamic topography is obtained by assuming that the normal stress on the top surface is instantaneously compensated by the deformed surface. This type of boundary condition works well for long-wavelength topography, but a free-surface formulation becomes necessary in cases where intermediate to short wavelength topography is of interest or the lithosphere has a very high effective viscosity.

The basic algorithm [16, 17] consists of four steps.

1. On the nodes of the top surface, compute the topography increment by integrating normal velocity over time.
2. Calculate the normal traction on the top surface based on the accumulated topography up to the current time step, and add it to the forcing term in Equation 1.31.
3. Update the velocity field with the changed forcing term.
4. If the velocity field has not converged yet, repeat steps 1 to 3.

6.7.2 Solution

To verify that the above algorithm works, you will run two different CitcomS models with each boundary condition (BC) (free-slip and pseudo-free-surface) and compare the topography computed accordingly. The scripts in `cookbook6.cfg` include the following parameters.

- Domain size: $45^\circ \times 45^\circ \times 1200$ km
- Mesh size: $61 \times 61 \times 25$ nodes with mesh refinement (`coor=1`)
- Boundary Conditions (BC): free-slip/free-surface for top; free-slip for all the other surfaces
- Initial Conditions (IC): spherical hot blob with diameter of $\frac{1}{4}$ of the radial dimension placed at the center of the domain

The pseudo-free-surface boundary condition is enabled by:

```
# CitcomS.solver.bc
topvbc=2
pseudo_free_surf=on
```

The time step size is fixed to 7.77×10^{-10} (about 1000 yrs), instead of determined dynamically by the velocity solution.

```
# CitcomS.solver.tsolver
fixed_timestep=7.77e-10
```

6.7.3 Discussion

The plots of the topography profile are consistent with results described in [16]. In the graphs in Figure 6.8, the solid lines indicate free-slip and the dashed lines indicate free-surface.




Figure 6.8: Cookbook 6: Graphs of topography profiles.

6.8 Cookbook 7: Thermo-Chemical Convection

6.8.1 Problem

This example solves for thermo-chemical convection within a full spherical shell domain. Composition heterogeneity exists in the Earth mantle. The density anomalies due to the composition heterogeneity, as well as due to the thermal heterogeneity, drive the convection flow.

6.8.2 Solution

You will use `cookbook7`. Most of the parameters you have encountered in previous cookbooks. The initial condition is the same as in Cookbook 1, and the viscosity law is the same as in Cookbook 3.

A large number of ASCII files will be output. In order to keep the current directory tidy, the output files can be put in a subdirectory `output/` by specifying

```
datadir=output
```

Since you are interested in the composition field, the composition output should be enabled and some output files disabled:

```
output_optional=tracer, comp_nd
```

The most important parameters are in the `CitcomS.solver.tracer` section. Turn on the tracer module:

```
tracer=1
```

The number of tracers and their initial location must be set. Here, you specify `tracer_ic_method=0`. The tracers will be generated pseudo-randomly, with a total number equal to

```
tracers_per_element*(total number of finite elements).
```

If `tracer_ic_method=1`, all processors will read the location (and flavors; see next paragraph) of the tracers from the same file specified by `tracer_file`. If `tracer_ic_method=2`, each processor will read the location (and flavors; see next paragraph) of the tracers from the file specified by `datafile_old` (in the `CitcomS.solver` section) and `solution_cycles_init` (in the `CitcomS.solver.ic` section).

```
tracer_ic_method=0
tracers_per_element=20
tracer_file=tracer.dat
```

Each tracer can have a “flavor” attached to it. The meaning of the flavor depends on the application. Here, the flavor indicates the chemical species of the tracer. There are two flavors of tracers (`tracer_flavors=2`). A tracer of flavor 0 is of the “normal,” or ambient, chemical composition, while a tracer of flavor 1 is of the “anomalous” chemical composition. Because the tracers are automatically generated, you need to specify how to assign the flavor to each tracer. If `ic_method_for_flavors=0`, tracers above `z_interface` are assigned to flavor 0, and tracers below `z_interface` to flavor 1.

```
tracer_flavors=2
ic_method_for_flavors=0
z_interface=0.7
```

The thermo-chemical convection module is turned on by

```
chemical_buoyancy=1
```

The composition field is determined by the ratio method (`buoy_type=1`) [19]. The density anomaly of the anomalous chemical composition is defined by $B = \Delta\rho_{ch}/(\rho_0\alpha_0\Delta T)$, the buoyancy ratio.

```
buoy_type=1
buoyancy_ratio=0.5
```

The code keeps track of the sum of the bulk composition. In a perfect world, the sum of the bulk composition would not change with time. But due to numerical issues, the sum of the bulk composition tends to decrease with time. If `tracer_ic_method=2`, the code will read in the location of the tracers, as well as the initial sum of the bulk composition, from the previous run.

The regular grid is an auxiliary grid that helps locate the tracers [18]. The optimal grid spacing of the regular grid depends on the size of the CitcomS mesh. A general rule of thumb is that the grid spacing of the regular grid should be less than $18/(nodex \times nprocx)$. These parameters are used only in the full spherical version.

```
regular_grid_deltheta=1.0
regular_grid_delphi=1.0
```

If you encounter any error, look at the end of the tracer log files (`cookbook7.tracer_log.*`) for error messages.

You will use the conjugate gradient solver to solve the matrix equation $\mathbf{A}x = b$ for x . The conjugate gradient solver is more efficient than the multigrid solver (`Solver=multigrid`) for small problems (e.g., less than $17 \times 17 \times 17$ nodes per processor). The desired accuracy and maximum iterations are set by `accuracy` and `vlowstep`. For the Uzawa algorithm, the maximum iterations are set by `piterations`. Both `vlowstep` and `piterations` should be large integers, but you can experiment with a larger value for `accuracy` to speed up the computation.

```
Solver=cgrad
accuracy=1e-03
vlowstep=1000
piterations=1000
```

6.8.2.1 Example: Thermo-Chemical Convection, `cookbook7`

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=15
solver=full

# CitcomS.controller
storage_spacing=5

# CitcomS.solver
datadir=output
datafile=cookbook7
rayleigh=1e7

# CitcomS.solver.mesher
fi_max=1
fi_min=0
nodex=9
nodey=9
nodez=9
nproc_surf=12
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.vsolver
Solver=cgrad
accuracy=1e-04
piterations=1000
```

```

vlowstep=1000

# CitcomS.solver.ic
num_perturbations=1
perturbl=3
perturblayer=5
perturbm=2
perturbmag=0.05

# CitcomS.solver.output
output_optional=tracer,comp_nd

# CitcomS.solver.tracer

# CitcomS.solver.visc
TDEPV=on
VISC_UPDATE=on
VMAX=on
VMIN=on
num_mat=4
rheol=4
visc0=1,1,1,1
viscE=0.2,0.2,0.2,0.2
viscT=0,0,0,0
viscZ=0,0,0,0
visc_max=100.0
visc_min=1.0

```




Figure 6.9: Cookbook 7: The composition and velocity field at the 15th step. The arrows are the velocity vectors. The composition field is shown in an isosurface of 0.7 and in a cross section.

6.8.3 Discussion

When the model is running, it will output the progress of the run to the screen. This line has information about the grid:

Problem has 9 x 9 x 9 nodes per cap, 6930 nodes and 6144 elements in total

The following lines give the radial coordinate and reference density of each nodes and their viscosity layer:

nz	radius	depth	rho	layer
1	0.550000	0.450000	1.000000e+00	4
2	0.606250	0.393750	1.000000e+00	4
3	0.662500	0.337500	1.000000e+00	4
4	0.718750	0.281250	1.000000e+00	4
5	0.775000	0.225000	1.000000e+00	4
6	0.831250	0.168750	1.000000e+00	4
7	0.887500	0.112500	1.000000e+00	4
8	0.943750	0.056250	1.000000e+00	2
9	1.000000	0.000000	1.000000e+00	1

This line gives the perturbation parameters used to construct the initial temperature:

Initial temperature perturbation: layer=5 mag=0.05 l=3 m=2

This line gives the magnitude the right-hand side vector in Equation 1.31:

Momentum equation force 2.207120371e+03

The following lines give the convergence progress of the Stokes solver, where v is the volume averaged norm of velocity, p is the volume averaged norm of pressure, div/v is the volume averaged norm of $\nabla \cdot (\bar{\rho}v)$ divided by v , dv/v is the volume averaged norm of velocity change divided by v , and dp/p is the volume averaged norm of pressure change divided by p . The convergence of the Stokes solver is achieved if either (1) div/v is less than accuracy or (2) both dv/v and dp/p are less than accuracy for two consecutive iterations. In this example, the second condition is met.

```
(000) 7.0 s v=1.470010e+02 p=0.000000e+00 div/v=5.37e+00 dv/v=1.00e+00 dp/p=1.00e+00
⋮
(090) 151.3 s v=1.012752e+02 p=1.304499e+04 div/v=5.54e-04 dv/v=1.84e-05 dp/p=9.20e-05
(091) 151.8 s v=1.012752e+02 p=1.304512e+04 div/v=5.66e-04 dv/v=2.07e-05 dp/p=7.45e-05
```

This line gives the rotation angle and the angular coordinates of rotation pole of the removed net angular momentum:

Angular momentum: rot=9.917258e-02 tr=8.924735e+01 fr=1.083630e+02

These lines give the heat flux across the top and bottom surfaces:

surface heat flux= 2.227892
bottom heat flux= 2.216542

The surface heat flux H_{surf} can be converted to the Nusselt number Nu_{surf} by:

$$Nu_{surf} = H_{surf} \times (r_{outer} - r_{inner}) \frac{r_{outer}}{r_{inner}} \quad (6.6)$$

The results for this problem are presented in Figure 6.9. The buoyancy ratio in this model is too low to stabilize the chemical layer. A few thermo-chemical plumes are rising from the lower mantle, especially the ones at the 4, 6, 9, 10, and 12 o'clock directions. The resolution of this model is fairly low. The composition isosurface is slightly discontinuous across the cap boundary. A model of higher resolution will not have this kind of artifact.

6.9 Cookbook 8: Compressible Steady-State Convection

6.9.1 Problem

This example is a benchmark problem for compressible thermal convection. The Stokes solver in CitcomS has been benchmarked and validated against a semi-analytical solution. However, no analytical solution exists for the benchmark on the energy equation solver, which is nonlinear. The steady-state solution is usually used for the comparison with other numerical solutions.

6.9.2 Solution

This cookbook example will run for 10,000 time steps to reach steady state. It will use 12 processors and take 1 to 2 days to finish on a modern computer. At every 1,000th time-step interval, a checkpoint for the internal state of the solver is saved.

```
checkpointFrequency=1000
```

If the solver is interrupted before finishing the computation, one can resume the computation from the checkpointed state. To shorten the computation time, a checkpoint at the 9,000th time step is provided in the `examples/Cookbook8/` directory. To restart the computation from this checkpoint, set these parameters:

```
restart=on
solution_cycles_init=9000
```

If you restart from the checkpoint, the computation will resume from the 9,000th time-step. Note that the `coord` files are only output at the 0th time-step, so you will need to run the model without restarting for one time-step to get the `coord` files.

A compressible convection model has four dimensionless numbers: the Rayleigh number, the dissipation number, the Gruneisen parameter $\gamma = \alpha_0 K_{S0} / \rho_0 c_{P0}$, and the non-dimensionalized absolute surface temperature, which are defined in Section 1.3. The Rayleigh number of CitcomS is scaled by the radius of the Earth. If scaled to the thickness of the mantle, the effective Rayleigh number is 7×10^3 ($= 7.68175583 \times 10^4 \times (r_{outer} - r_{inner})^3$). Similarly, the effective dissipation number is 0.275 ($= 0.5 \times (r_{outer} - r_{inner})$). Under these non-dimensional numbers, the convection is of low vigor and low compressibility.

```
rayleigh=7.68175583e4
dissipation_number=0.5
gruneisen=1.25
surfaceT=0.1
```

Since we are going to use the multigrid solver, the grid size is specified by:

```
nodex=33
nodey=33
nodez=33
levels=5
```

The additional parameter `levels` specifies the nested levels of multigrid units and is subjected to the following constraint:

$$\text{nodex} = 1 + \text{nprocx} \times \text{mgunitx} \times 2^{\text{levels}-1} \quad (6.7)$$

where `mgunitx`, which **must be an integer**, is an input parameter.

If `reference_state=1`, then constant gravity, heat capacity, thermal expansivity, and $\rho_r = \exp\left(\frac{D_i}{\gamma}(1-r)\right)$ are used as the reference state. If `reference_state=0`, the reference state is read from a file `refstate_file`. See Appendix A for the file format.

```
reference_state=1
refstate_file=ref.dat
```

We are interested in the geoid, dynamic topography, and heat flux. The dynamic topography and heat flux are computed on the surface grids. The geoid is computed in the spherical harmonics, with a maximal degree of 20.

```
output_optional=geoid, surf, botm
output_ll_max=20
```

The initial temperature is a conductive profile with a single spherical harmonic perturbation. The perturbation is located at mid-depth and is defined as:

$$mag \times \sin\left(\frac{(r - r_{in})\pi}{r_{out} - r_{in}}\right) (\sin(m\phi) + \cos(m\phi)) P_{lm}(\cos\theta)$$

```
tic_method=3
num_perturbations=1
perturb1=3
perturbm=2
perturb1ayer=17
perturbmag=0.01
```

You will need the output of dynamic topography and the geoid. The dynamics topography will be computed by the Consistent Boundary Flux (CBF) method, and the effect of self-gravitation is included in the geoid. The maximum spherical harmonics degree for the geoid is 20.

```
output_optional=geoid, surf, botm
use_cbf_topo=on
self_gravitation=on
output_ll_max=20
```

Various parameters tune the performance of the solver. The maximum size of each time step is determined dynamically by the Courant criterion. To enhance the stability of the energy equation solver, you will only use three quarters of the maximum Courant time-step size.

```
finetunedt=0.75
```

You will use the multigrid solver to solve the matrix equation $\mathbf{A}x = b$ for x . The multigrid solver is more efficient than the conjugate gradient solver (`Solver=cgrad`) for larger problems (e.g., more than $17 \times 17 \times 17$ nodes per processor). Several parameters control the behavior of the multigrid solver: `mg_cycle=1` for the V cycle and 2 for the W cycle; `down_heavy` and `up_heavy` are the number of smoothing cycles for downward/upward smoothing; `vlowstep` and `vhighstep` are the number of smoothing passes at lowest/highest levels; and `max_mg_cycles` is the maximum number of multigrid cycles per solve. All these parameters should be small integers.

```
Solver=multigrid
mg_cycle=1
down_heavy=2
up_heavy=2
vlowstep=20
vhighstep=2
max_mg_cycles=50
```

The following parameter turn on the pre-conditioner for the matrix equation solver (either multigrid or conjugate gradient).

```
precond=on
```

The stiffness matrix uses augmented Lagrangian formulation to improve the convergence for large viscosity variations [3]. These parameters specify whether to enable the formulation and how much weight to use for the formulation.

```
aug_lagr=on
aug_number=2.0e3
```

The discrete Stokes equations 1.30 and 1.31 are solved using the Uzawa algorithm, which iteratively updates the pressure and velocity solutions. Three variations of the Uzawa algorithm are used in CitcomS, one for the incompressible case, and the other two for the compressible case. One parameter, `piterations`, common to the three variations, specifies the maximum number of iterations and the desired residual level for the continuity equation 1.30. Sometimes, larger value of `piterations` is required for convergence if complicated velocity boundary conditions are used.

```
piterations=375
```

For the compressible case, two choices of the Uzawa algorithm are available. If `uzawa=cg`, the algorithm described in Equation 1.33 is used. In this case, an additional parameter controls the maximum number of outer iterations.

```
uzawa=cg
compress_iter_maxstep=100
```

If `uzawa=bicg`, the algorithm described in Equation 1.34 is used, and no additional parameter is needed.

The overall accuracy of the velocity solver is controlled by a single parameter. The solver is converged if the residuals of Equation 1.30 and 1.31 both are smaller than `accuracy`, or if the changes in the velocity and pressure both are smaller than `accuracy` for the last two iterations.

```
accuracy=0.001
```

Finally, the net angular momentum of the velocity solution is removed. The net angular momentum and rigid body rotation are unconstrained by the Stokes equations, if free-slip boundary conditions are used for the top and bottom boundaries in a full spherical model. That is, you can add an arbitrary amount of rotation to the velocity solution, and the resultant velocity is still a valid solution of the Stokes equations. Since no external torque is applied to the mantle in free-slip boundary conditions, the angular momentum must be constant in time. Enforcing the angular momentum to be zero by removing it from the velocity solution is often desirable.

```
remove_rigid_rotation=off
remove_angular_momentum=on
```

However, for models with imposed plate velocity, it is advisable to turn off both `remove_rigid_rotation` and `remove_angular_momentum`.

6.9.2.1 Example: Compressible Steady-State Convection, cookbook8

```
# CitcomS
cpu_limits_in_seconds=360000000
minstep=10000
solver=full

# CitcomS.controller
checkpointFrequency=1000
storage_spacing=1000

# CitcomS.solver
datadir=output
datadir_old=restart
datafile=cookbook8
datafile_old=cookbook8
dissipation_number=0.5
gruneisen=1.25
rayleigh=7.68175583e4
```

```
surfaceT=0.1

# CitcomS.solver.mesher
coor=1
coor_file=coord.dat
fi_max=1
fi_min=0
levels=5
mgunitx=2
mgunity=2
mgunitz=2
nodex=33
nodey=33
nodez=33
nproc_surf=12
theta_max=2.0708
theta_min=1.0708

# CitcomS.solver.tsolver
finetunedt=0.75

# CitcomS.solver.vsolver
Solver=multigrid
accuracy=0.001
aug_lagr=on
aug_number=2.0e3
compress_iter_maxstep=100
down_heavy=2
max_mg_cycles=50
mg_cycle=1
piterations=375
precond=on
remove_angular_momentum=on
remove_rigid_rotation=off
up_heavy=2
uzawa=cg
vhighstep=2
vlowstep=20

# CitcomS.solver.ic
num_perturbations=1
perturbl=3
perturblayer=17
perturbm=2
perturbmag=0.01
restart=off
solution_cycles_init=9000
tic_method=3

# CitcomS.solver.output
output_ll_max=20
output_optional=geoid,surf,botm
self_gravitation=on
use_cbf_topo=on
```

```

# CitcomS.solver.param
reference_state=1
refstate_file=ref.dat

# CitcomS.solver.visc
TDEPV=on
VISC_UPDATE=on
VMAX=on
VMIN=on
num_mat=4
rheol=1
visc0=1,1,1,1
viscE=2.99573,2.99573,2.99573,2.99573
viscT=0.5,0.5,0.5,0.5
visc_max=1e+06
visc_min=0.001
visc_smooth_method=1

```

6.9.3 Discussion

The results for this problem are presented in Figure 6.10. A tetrahedral symmetric pattern is developed for the convection. The surface heat flux H_{surf} at the steady state is 3.892, and the bottom heat flux H_{botm} is 12.817. The heat flux imbalance ($H_{botm}r_{inner}^2/H_{surf}r_{outer}^2 - 1$) is -0.38%. The Nussult number, converted from the surface heat flux by Equation 6.6, is 3.182. The total viscous heating is 7.68991, and the total adiabatic cooling is 7.71719. Under a steady state, these two terms should be exactly balanced.

The geoid is output as spherical harmonic coefficients. A post-processing program is provided to project the geoid coefficients to a regular (longitude, latitude) grid. To use a regular grid of 1° spacing, run this command:

```
$ visual/project_geoid cookbook8.geoid.0.10000 geoid.xyz 361 181
```

It will generate a file `geoid.xyz`. There are 3 columns in the file, which are (longitude, latitude, geoid). The geoid is in a unit of meters. Several dimensional constants are required for the geoid computation. These constants have sensible default values (in SI units) for the Earth. Note that the temperature drop from the core-mantle boundary to the surface (ΔT in Equation 1.16) is derived from these constants and the Rayleigh number.

```

radius=6.371e+06
density=3340
thermdiff=1e-06
gravacc=9.81
thermexp=3e-05
refvisc=1e+21
cp=1200
density_above=1030
density_below=6600

```




Figure 6.10: Cookbook 8: The steady state temperature field at the 10,000th time step. A tetrahedral symmetric convection pattern is developed. Two temperature isosurfaces of 0.4 and 0.8 are shown.

Part III
Appendices

Appendix A

Input Parameters for CitcomS

A.1 Input Parameters Grouped by Functionality

This section explains the meaning of the input parameters for CitcomS. These parameters are grouped by their functionality. Parameters are given with their default values.

A.1.1 Parameters that Control Input Files

reference_state=1 refstate_file="refstate.dat"	If reference_state is set to 1, a simple reference state of $p_r = \exp((1-r)D_i/\gamma)$ is used, with constant gravity, thermal expansivity, and heat capacity. If reference_state is set to 0, the reference state is read from a file refstate_file.
mineral_physics_model=3	Using which mineral physics model to convert the seismic velocities.
file_vbcs=off vel_bound_file="bvel.dat"	If file_vbcs is set to on, the top surface velocity boundary conditions are read in from files which have location and name specified by vel_bound_file. Requires setting topvbc=1 to take effect. If you wish to have a uniform top surface velocity boundary condition or some simple geometric pattern, then file_vbcs should be set to zero.
mat_control=off mat_file="mat.dat"	If mat_control is set to on, then the time- and positional-dependent viscosity factor is defined from the files specified by mat_file. These parameters allow you to define the material group of each element (such as a moving weak zone). Not working in this version.
lith_age=off lith_age_file="age.dat" lith_age_time=off	If lith_age is set to on, then the age of each surface nodes is read from the files specified by lith_age_file. These parameters control the thermal age of the top thermal boundary condition. If lith_age_time is on, the files are time-dependent.

A.1.2 Parameters that Control Output Files

<code>output_format=ascii</code>	Choose the format and layout of the output files. Can be either <code>ascii</code> , <code>ascii-gz</code> or <code>hdf5</code> . If <code>ascii-gz</code> is chosen, the code places gzipped files into <code>data_dir</code> , and will put all time-step output into subdirectories of <code>data_dir</code> . The same naming logic holds for reading old velo files.
<code>output_optional="surf, botm, tracer"</code>	Choose additional output, including <code>surf</code> , <code>botm</code> , <code>geoid</code> , <code>seismic</code> , <code>stress</code> , <code>pressure</code> , <code>connectivity</code> , <code>horiz_avg</code> , <code>tracer</code> , <code>heating</code> , <code>comp_el</code> and <code>comp_nd</code> .
<code>datadir="."</code>	Controls the location of output files.
<code>datafile="regtest"</code>	Controls the prefix of output file names such as <code>regtest.xxx</code> . Cannot contain the "/" character if <code>output_format=ascii</code> .
<code>storage_spacing=10</code> (in non-Pyre version) or <code>monitoringFrequency=100</code> (in Pyre version)	Controls the interval between output files. CitcomS dynamically determines the size of the time step; this means that you might not get an output at the exact time required, but you can always get close depending on how small this number is. Do not make this number too small since outputs slow the code down and you may end up with an unmanageable number of output files.
<code>checkpointFrequency=100</code>	The time-step interval between checkpoint output, which can be used later to resume the computation.
<code>output_ll_max=20</code>	This parameter controls the maximum degree of spherical harmonics coefficients for geoid output.
<code>self_gravitation=off</code>	Considering the effect the self gravitation on the geoid or not.
<code>use_cbf_topo=off</code>	Using the Consistent Boundary Flux (CBF) method to compute the dynamic topography or not.

A.1.3 Mesh and Processors Setup

<code>nproc_surf=1</code>	This specifies the number of spherical caps of the mesh; must be 1 for regional spherical model and 12 for full spherical model.
<code>nprocx=1</code> <code>nprocy=1</code> <code>nprocz=1</code>	These specify the number of processors in each spherical cap. For a full spherical model, <code>nprocx</code> must be equal to <code>nprocy</code>
<code>nodex=9</code> <code>nodey=9</code> <code>nodez=9</code>	These specify the number of FEM nodes in each spherical cap. These parameters are not used in the C version if multigrid solver is used. For a full spherical model, <code>nodex</code> must be equal to <code>nodey</code>
<code>mgunitx=8</code> <code>mgunity=8</code> <code>mgunitz=8</code> <code>levels=1</code>	These specify the nested level of multigrid units. Used by multigrid solver only. These parameters are not completely independent to each other. The constraint of Equation 6.7 must be satisfied. These parameters are not used in the C version if conjgrad solver is used. For a full spherical model, <code>mgunitx</code> must be equal to <code>mgunity</code>

A.1.4 Domain Size

<pre>theta_min=1.0708 theta_max=2.0708 fi_min=0 fi_max=1</pre>	<p>These parameters specify the horizontal extent of the computational domain. <code>theta_min</code> and <code>theta_max</code> are the colatitude measured in radians from the north pole. <code>fi_min</code> and <code>fi_max</code> are the longitudes measured from the prime meridian eastward in radians. Only in regional CitcomS.</p>
<pre>radius_inner=0.55 radius_outer=1.0</pre>	<p>These parameters specify the radial extent of the computational domain. <code>radius_inner</code> and <code>radius_outer</code> are the inner and outer radii in non-dimensional units. It is probably more convenient to normalize lengths by the radius of the Earth. If you do this, then the Rayleigh number must be calculated with the radius of the Earth, not the thickness of the mantle. The core mantle boundary is close to having a non-dimensional radius of 0.55.</p>
<pre>coor=0 coor_file="coor.dat" coor_refine=0.1,0.15,0.1,0.2</pre>	<p>If <code>coor=0</code>, there will be uniform mesh in the latitudinal, longitudinal, and radial directions.</p> <p>If <code>coor=1</code>, the coordinate is reading from the file specified by <code>coor_file</code>. This is used to have a regular but uneven spacing between elements.</p> <p>If <code>coor=2</code>, there will be uniform mesh in the latitudinal and longitudinal directions. The mesh in the radial direction is generated according to <code>coor_refine</code>, which is a vector of 4 numbers. The 1st value of <code>coor_refine</code> specifies the radius fraction of the bottom layer, the 2nd value specifies the fraction of the nodes in the bottom layer, the 3rd value specifies the top layer fraction, and the last value specifies the top layer node fraction.</p>

A.1.5 Restarting the Code

<pre>restart=off</pre>	<p>If <code>restart</code> is on, each processor will resume the computation from the checkpoint files.</p>
<pre>post_p=off</pre>	<p>Similar to <code>restart</code>, except that the model will then only run for 1 time step, which can be useful to regenerate the flow field and calculate the associated observables.</p>
<pre>datadir_old="." datafile_old="regtest" solution_cycles_init=0</pre>	<p>If <code>restart</code> is on, for example, processor #5 will read its initial conditions from checkpoint file <code>regtest.chkpt.5.0</code> in this case.</p>

A.1.6 Run Length

<pre>minstep=1 maxtotstep=1000000 (only in pure C version) or steps=1 (only in Pyre version)</pre>	<p>The maximum and minimum number of time steps for the model, including the 0th time step.</p>
<pre>cpu_limits_in_seconds= 360000000</pre>	<p>Controls the termination of the code based on total wall clock time used. Available only in pure C version.</p>

A.1.7 Initial Conditions

<pre>tic_method=0</pre>	<p>Which method to use to generate the initial temperature field.</p>
-------------------------	-----------------------------------------------------------------------

<pre>datadir_old="." datafile_old="regtest" solution_cycles_init=0 zero_elapsed_time=on</pre>	<p>If <code>tic_method=-1</code>, the initial temperature is read from files. For example, processor #5 will read its initial temperature from old velo file <code>regtest.velo.5.0</code> in this case. If <code>zero_elapsed_time</code> is on, the initial time is set to zero. If it is off and <code>tic_method=-1</code>, the initial time is read in from the old velo files. Note that this option has no effect when <code>restart=on</code>.</p>
<pre>num_perturbations=1 perturbmag=0.05 perturbl=1 perturbm=1 perturblayer=5 half_space_age=40 mantle_temp=1.0 blob_center=[-999, -999, -999] blob_radius=0.063 blob_dT=0.18</pre>	<p>If <code>tic_method=0</code>. The initial temperature is a linear temperature gradient with perturbations at specific layers, where <code>num_perturbations</code> specifies the number of perturbations, and <code>perturblayer</code> specifies the layers to be perturbed, representing the number of the mesh node in radial direction. There must be as many entries as <code>num_perturbations</code> in a comma-separated list. The perturbation added to each layer is given by:</p> $mag \times \cos(m\phi) \times P_{lm}(\cos\theta)$ <p>for the full sphere, and by:</p> $mag \times \cos\left(\frac{(\theta - \theta_{min})l\pi}{\theta_{max} - \theta_{min}}\right) \times \cos\left(\frac{(\phi - \phi_{min})m\pi}{\phi_{max} - \phi_{min}}\right)$ <p>for the regional sphere.</p> <p>If <code>tic_method=1</code>, <code>T</code> is 1 everywhere, except a cold thermal boundary layer at the top, whose temperature is determined by the half-space cooling model and <code>half_space_age</code> (in million of years, Myrs).</p> <p>If <code>tic_method=2</code>, <code>T</code> is <code>mantle_temp</code> everywhere, except for a warm spherical blob and a cold thermal boundary layer at the top, whose temperature is determined by the half-space cooling model and <code>half_space_age</code> (in Myrs). The location of the blob is default to the center of the computational domain.</p> <p>If <code>tic_method=3</code>, the initial temperature is a conductive profile with perturbations to all layers. The perturbation is given by:</p> $mag \times \sin\left(\frac{(r - r_{in})\pi}{r_{out} - r_{in}}\right) (\sin(m\phi) + \cos(m\phi)) P_{lm}(\cos\theta)$ <p>for the full sphere, and by:</p> $mag \times \sin\left(\frac{(r - r_{in})\pi}{r_{out} - r_{in}}\right) \times \cos\left(\frac{(\theta - \theta_{min})l\pi}{\theta_{max} - \theta_{min}}\right) \times \cos\left(\frac{(\phi - \phi_{min})m\pi}{\phi_{max} - \phi_{min}}\right)$ <p>for the regional sphere.</p> <p>If <code>tic_method=4</code>, the initial temperature is read from <code>grd</code> files.</p> <p>If <code>tic_method=10</code>, <code>T</code> is <code>mantle_temp</code> everywhere, except for a cold thermal boundary layer at the top and perturbations at all layers, similar to <code>tic_method=3</code>.</p> <p>If <code>tic_method=11</code>, <code>T</code> is <code>mantle_temp</code> everywhere, except for a hot thermal boundary layer at the bottom and perturbations at all layers, similar to <code>tic_method=3</code>.</p> <p>If <code>tic_method=12</code>, <code>T</code> is <code>mantle_temp</code> everywhere, except for a cold thermal boundary layer, a hot thermal boundary layer at the bottom and perturbations at all layers, similar to <code>tic_method=3</code>.</p> <p>If <code>tic_method=90</code>, <code>T</code> is 0 everywhere, except for a single perturbation at the middle layer. This initial temperature is good for comparison with analytical solutions.</p>

A.1.8 Boundary Conditions

topvbc=0 topvbxval=0.0 topvbyval=0.0	Surface velocity boundary condition parameters. If <code>topvbc</code> is 0, <code>topvbxval</code> and <code>topvbyval</code> specify the tangential surface stress (stress BC). If <code>topvbc</code> is 1, <code>topvbxval</code> and <code>topvbyval</code> specify the tangential surface velocity (velocity BC). The surface normal velocity is zero in these two cases (impermeable BC).
botvbc=0 botvbxval=0.0 botvbyval=0.0	As above, but for bottom velocity boundary conditions.
side_sbcs=off	Enable traction boundary condition for the sidewalls or not. Must be <code>on</code> for coupled model.
pseudo_free_surf=off	Enable pseudo free surface or not.
toptbc=1 toptbcval=0.0	Surface temperature boundary conditions. If <code>toptbc</code> is 0, <code>toptbcval</code> specifies the surface heat flux (not working in this version). If <code>toptbc</code> is 1, the <code>toptbcval</code> specifies the surface temperature.
bottbc=1 bottbcval=1.0	As above, but for bottom temperature boundary conditions.
temperature_bound_adj=off depth_bound_adj=0.157 width_bound_adj=0.08727 lith_age_depth=0.0314	Additional parameters for temperature boundary conditions when <code>lith_age</code> is on.

A.1.9 Non-Dimensional Numbers

rayleigh=1.0e+5	This specifies the Rayleigh number, which is one of the most important parameters you may want to change.
dissipation_number=0.0	The dissipation number.
gruneisen=0.0	The Gruneisen parameter. When this parameter is 0, the code treats it as infinity (i.e., incompressible case).
surfaceT=0.1	The non-dimensional value of surface temperature.
Q0=0.0	This specifies the internal heating number.

A.1.10 Depth Information

z_lith=0.014 z_410=0.06435 z_lmantle=0.105 z_cmb=0.439	These specify the non-dimensional depth of the Moho, 410km discontinuity, 660km discontinuity and D". These parameters are used to determine the depth of viscosity layers and phase changes (see next two sections). The names are only suggestive. You are free to refer <code>z_lith</code> to an arbitrary depth, for example.
-----------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

A.1.11 Viscosity

Viscosity=system	This parameter must be set as indicated.
visc_smooth_method=3	This specifies which method to use to smooth viscosity projection for the multigrid solver.
VISC_UPDATE=on	If <code>VISC_UPDATE</code> is on, viscosity will be updated every time step. Otherwise, viscosity will be time-independent.

num_mat=4	This specifies the number of material layers. Material 1 is at depth between 0 and z_lith, material 2 between z_lith and z_410, material 3 between z_410 and z_lmantle, and material 4 between z_lmantle and the bottom. If mat_control is on, then a multiplicative factor is applied to the viscosity, as defined below.
visc0=1,1,1,1	The pre-exponent factor of layered viscosity structure (η_0 in the equations below). There must be as many entries as num_mat in a comma-separated list.
TDEPV=off	Enable temperature dependence or not.
rheol=3	When rheol=1, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(E_\eta(T_\eta - T))$ When rheol=2, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(-T/T_\eta)$ When rheol=3, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(\frac{E_\eta}{T^*+T_\eta} - \frac{E_\eta}{1+T_\eta})$, where $T^* = \min(\max(T, 0), 1)$. When rheol=4, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(\frac{E_\eta+Z_\eta(1-r)}{T^*+T_\eta})$ When rheol=5, same as rheol=3, except the viscosity cut-off is applied before mat_control. When rheol=6, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(E_\eta(T_\eta - T^*) + Z_\eta(1-r))$ When rheol=7, temperature-dependent viscosity is computed by $\eta = \eta_0 \times \exp(\frac{E_\eta+Z_\eta(1-r)}{T+T_\eta} - \frac{E_\eta+Z_\eta(1-r_{inner})}{1+T_\eta})$, where r_{inner} is the inner radius of the mesh. When rheol=8, same as rheol=3, except viscosity reduction is applied when the temperature exceeds the solidus. When rheol=9, same as rheol=3, except using T , not T^* . When rheol=10, same as rheol=8, except using T , not T^* .
viscE=1,1,1,1 viscT=1,1,1,1 viscZ=1,1,1,1	Parameters defining viscosity law (E_η , T_η , and Z_η in the equations above, respectively). There must be as many entries as num_mat in a comma-separated list.
SDEPV=off	Enable stress dependence (non-Newtonian) or not.
sdepv_expt=1,1,1,1 sdepv_misfit=0.001	If SDEPV is on, these specify the exponent in the viscosity law and the criterion of convergence test. There must be as many entries as num_mat in a comma-separated list.
PDEPV=off	Pseudo-plastic rheology, implemented by adding a plastic viscosity $\eta_p = \frac{\sigma_y}{(2\epsilon_{II}+10^{-7})+\eta_p^0}$ where the yield stress is defined as $\sigma_y = \min(a + b(1-z), y)$ and ϵ_{II} is the second shear strain rate invariant (all non-dimensional). The effective viscosity η' is computed from η and η_p as $\eta' = \frac{\eta\eta_p}{(\eta+\eta_p)}$ (pdepv_eff=on) or $\eta' = \min(\eta, \eta_p)$ (pdepv_eff=off).
pdepv_a=1e20,1e20,1e20,1e20 pdepv_b=0,0,0,0 pdepv_y=1e20,1e20,1e20,1e20	Parameters for σ_y .
pdepv_eff=on	Effective or minimum viscosity “plasticity” (see above).
pdepv_offset=0	Offset for plastic viscosity η_p^0 .
CDEPV=off	Compositionally dependent viscosity pre-factor (only for tracer=on). Will assign two pre-multipliers for $C = 0$ and $C = 1$ and compute a geometric mean for all elements depending on the nodal composition. Only works for two flavors and assuming C varies between $[0; 1]$.
cdepv_ff=1,1	Flavor-wise viscosity pre-factors for $C = 0$ and $C = 1$, respectively.
low_visc_channel=off low_visc_wedge=off	Used in conjunction with tracers. The tracers define the upper boundary of the subducted slab.

lv_min_radius=0.9764 lv_max_radius=0.9921 lv_channel_thickness=0.0047 lv_reduction=0.5	If <code>low_visc_channel</code> or <code>low_visc_wedge</code> is on, this specifies the radial extents of the low viscosity zones and the viscosity reduction factor, respectively.
VMIN=off visc_min=0.001	If VMIN is on, minimum viscosity is cut off at <code>visc_min</code> .
VMAX=off visc_max=1000	If VMAX is on, maximum viscosity is cut off at <code>visc_max</code> .

A.1.12 Phase Change Information

Ra_410=0.0 clapeyron410=0.0235 transT410=0.78 width410=0.0058	These specify the phase change parameters (phase change Rayleigh number, Clapeyron slope, ambient temperature, and phase change width, respectively). The depth of this phase change is specified by <code>z_410</code> .
Ra_670=0.0 clapeyron670=-0.0235 transT670=0.875 width670=0.0058	As above. The depth of this phase change is specified by <code>z_lmantle</code> .
Ra_cmb=0.0 clapeyroncmb=-0.0235 transTcmb=0.875 widthcmb=0.0058	As above. The depth of this phase change is specified by <code>z_cmb</code> .

A.1.13 Momentum Equation Solver Parameters

stokes_flow_only=off	If you wish only to solve for the velocity once (e.g., Stokes flow) then change this parameter to <code>on</code> . However, if you want to do a convection problem, this should be <code>off</code> .
remove_rigid_rotation=on remove_angular_momentum=on	Whether to remove rigid body rotation component or net angular momentum from the velocity solution. Only effective for global models.
Solver=cgrad	Can be either <code>cgrad</code> for conjugate gradient solver or <code>multigrid</code> for multigrid solver for the outer loop of the momentum solver.
node_assemble=on	Whether to assemble stiffness matrix at the node level or not.
mg_cycle=1 down_heavy=3 up_heavy=3 vlowstep=1000 vhighstep=3 max_mg_cycles=50	Parameters for <code>multigrid</code> and <code>cgrad</code> solvers. For <code>multigrid</code> solver: <code>mg_cycle=1</code> for V cycle and 2 for W cycle; <code>down_heavy</code> and <code>up_heavy</code> are the number of smoothing passes for downward/upward smoothing; <code>vlowstep</code> and <code>vhighstep</code> are the number of smoothing passes at lowest/highest levels; <code>max_mg_cycles</code> is the maximum number of multigrid cycles used per solve; all these parameters should be small integers. For <code>cgrad</code> solver, <code>vlowstep</code> is the maximum iterations of the conjugate gradient solver and should be a large integer.
piterations=1000	Maximum iterations of the outer loop for the momentum solver.
accuracy=1.0e-4	Convergence criterion for the momentum solver.
uzawa=cg compress_iter_maxstep=100	Can be either <code>cg</code> or <code>bicg</code> . If set to <code>cg</code> , additional parameters control the maximum iterations.
precond=on	Whether to use the preconditioner.
aug_lagr=on aug_number=2.0e3	Whether to use augmented stiff matrix and the weight of the augmented stiff matrix.

A.1.14 Energy Equation Solver Parameters

ADV=on	If on, solves the energy equation.
fixed_timestep=0.0	If it is equal to 0, the size of the time step is variable and is determined dynamically. Otherwise, the size of the time step is fixed at the specified value.
finetunedt=0.9	Set the size of the time step to the specified fraction of a maximum stable advection time step. Must be between 0 and 1.
adv_sub_iterations=2 adv_gamma=0.5	The number of iterations and the weight of each iteration for the energy predictor-corrector solver.
filter_temp=off	Using a Lenardic filter to remove the overshoots and undershoots of the temperature field or not. The filter conserves the total energy.
monitor_max_T=on	If on, the maximum value of the current and previous temperature fields are compared. If the maximum temperature increases more than 5%, the energy equation solver is rerun with a smaller time-step size. Keep this parameter on to prevent numerical instability.
inputdiffusivity=1	Currently, don't change this parameter. It is used only in problems which are integrated backward in time.

A.1.15 Age Information

start_age=40.0	Set initial age (in Myrs). This age determines which files of various time-dependent input to read in.
reset_startage=off	If on, the initial age is set to <code>start_age</code> . If it is <code>off</code> and <code>tic_method=-1</code> , the initial age is read in from previous output.

A.1.16 Debugging Information

DESCRIBE=off BEGINNER=off VERBOSE=off	These parameters affect the echo behavior of the CitcomS parser. Only in non-Pyre version.
verbose=off	This is used for debugging. If <code>verbose</code> is <code>on</code> , additional information is output to a <code>.info</code> file.
see_convergence=on	If <code>on</code> , the velocity residual will be output on the screen for every iteration of the momentum solver.

A.1.17 HDF5 Output Parameters

cb_block_size=1048576 cb_buffer_size=4194304	Size for collective buffer in MPI-IO.
sieve_buf_size=1048576	Size of data sieve buffer.
output_alignment=262144 output_alignment_threshold= 524288	Memory alignment.
cache_mdc_nelmts=10330 cache_rdcc_nelmts=521 cache_rdcc_nbytes=1048576	Cache size for chunked dataset.

A.1.18 Tracer Parameters

<code>tracer=off</code>	If on, enables the tracers.
<code>tracer_ic_method=0</code> <code>tracers_per_element=10</code> <code>tracer_file="tracer.dat"</code>	Specify the initial location of the tracers. If <code>tracer_ic_method=0</code> , the tracers are generated randomly, with the number of tracers per element specified by <code>tracers_per_element</code> . If <code>tracer_ic_method=1</code> , the location of the tracers is read from a file specified in <code>tracer_file</code> . If <code>tracer_ic_method=2</code> , the location of the tracers is read from the old tracer output, similar to <code>tic_method=-1</code> .
<code>tracer_flavors=0</code>	The number of flavors among the tracers. If set to 0, flavor counting is turned off.
<code>ic_method_for_flavors=0</code> <code>z_interface=0.7</code>	Specify the initial flavors of the tracers. Used only when <code>tracer_ic_method=0</code> . Currently only <code>ic_method_for_flavors=0</code> is supported. A layered structure is generated. <code>z_interface</code> must be a comma-separated list with $(\text{tracer_flavors}-1)$ entries. Tracers above the 1st radius in <code>z_interface</code> are of flavor 0, between the 1st and 2nd radii are of flavor 1, ... etc.
<code>regular_grid_deltheta=1.0</code> <code>regular_grid_delphi=1.0</code>	The grid spacing of the regular grid. The regular grid is an auxiliary grid to help locate the tracers.
<code>chemical_buoyancy=on</code>	If on, enables thermo-chemical convection.
<code>buoy_type=1</code>	If <code>buoy_type=1</code> , the composition field is determined by the ratio method [19]. (No other methods are implemented yet.)
<code>buoyancy_ratio=1.0</code>	The ratio of chemical density anomaly to the reference density, must be a comma-separated list with $(\text{tracer_flavors}-1)$ entries.
<code>itracer_warnings=on</code>	The warning level of the tracer module.
<code>tracer_enriched=off</code> <code>Q0_enriched=0.0</code>	Whether the composition anomaly is associated with radioactive heating anomaly. If on, specifies the internal heating number inside the composition anomaly.

A.1.19 Dimensional Information

<code>radius=6371e3</code> <code>density=3340.0</code> <code>thermdiff=1.0e-6</code> <code>gravacc=9.81</code> <code>thermexp=3.0e-5</code> <code>refvisc=1.0e+21</code> <code>cp=1200</code> <code>density_above=1030.0</code> <code>density_below=6600.0</code>	Various dimensional information in SI units.
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------

A.1.20 Required Information

<code>Problem=convection</code> <code>Geometry=sphere</code>	For this version of CitcomS, all of these parameters must be set as indicated.
-----------------------------------------------------------------	--------------------------------------------------------------------------------

Appendix B

CitcomS Input File Format

B.1 Introduction

CitcomS expects Unix-styled ASCII files (i.e., no carriage character following new line character) for all input files. This can be a nuisance in DOS/Windows systems. You may want to find a text editor that can write Unix-style ASCII files. In the following, words in normal `courier` must be input exactly as shown, while *italicized* words should be substituted by your values. All parameters are in non-dimensional units unless specified.

B.2 Coordinate Files

For regional version of CitcomS, the mesh must be regular, but the mesh spacing may be unequal. The `coor_file` has the format:

```
nsd= 1
1      theta1
2      theta2
...    ...
nodex theta_nodex
nsd= 2
1      phi1
2      phi2
...    ...
nodey phi_nodey
nsd= 3
1      r1
2      r2
...    ...
nodez r_nodez
```

For the full spherical version of CitcomS, the mesh of each cap must be regular and equidistant in the horizontal dimension. Only the vertical dimension is specified by `coor_file`. The `coor_file` has the format:

```
nsd= 3
1      r1
2      r2
...    ...
nodez r_nodez
```

B.3 Velocity Boundary Condition Files

If `vel_bound_file` is set to `bvel.dat`, then it is necessary to have one `bvel.dat` file per cap for each million-year interval. For example, if `start_age=83`, then the following files are needed for the regional mesh:

```
bvel.dat84
bvel.dat83
bvel.dat82
...
bvel.dat0
```

TIP: Even though the model starts at 83 million years before the present, by default it will attempt to read in both a file for 84 and a file for 83 million years ago. To deal with this, create a duplicate copy of `bvel.dat83` and call it `bvel.dat84`

In this example, the model age will start at 83 Ma and decrease toward 0 Ma. At any particular age, the velocity boundary conditions are interpolated in time according to the nearest `bvel.dat` files. Once the age becomes negative, the velocity boundary conditions are specified by the content of `bvel.dat0`.

TIP: By setting `start_age=0` and providing identical `bvel.dat0` and `bvel.dat1`, you can effectively get time-invariant velocity boundary conditions.

For the global mesh, each of the 12 caps requires one file for each million-year interval. For example, these files are needed for an 82-million-year age:

```
bvel.dat82.0
bvel.dat82.1
bvel.dat82.2
...
bvel.dat82.11
```

Each velocity boundary condition file has the format:

V_x V_y

The units of V_x and V_y are cm/yr.

B.4 Material Files

In this version of CitcomS, the implementation of material support is not working. The material output has been disabled and is documented here for completion. If `mat_file` is set to `mat.dat`, then it is necessary to have one `mat.dat` file for each million-year interval. For example, if `start_age=83`, then the following files are needed:

```
mat.dat84
mat.dat83
mat.dat82
...
mat.dat0
```

The same note about `vel_bound_file` (see Section B.3 above) applies. The format of `mat_file` is:

n viscosity_factor

B.5 Lithosphere Age Files

If `lith_age_file` is set to `lith.dat`, then it is necessary to have one `bvel.dat` file for each million-year interval. For example, if `start_age=83`, then the following files are needed for the regional mesh:

```
lith.dat84
lith.dat83
lith.dat82
...
lith.dat0
```

The same note about `vel_bound_file` (see Section B.3 above) still applies. The input age is in millions of years. For the global mesh, each of the 12 caps requires one file for each million-year interval. For example, these files are needed for an 82-million-year age:

```
lith.dat82.0
lith.dat82.1
lith.dat82.2
...
lith.dat82.11
```

The format of `lith_age_file` is:

```
n age
```

B.6 Tracer Files

This file contains the initial location of all tracers. The first line is the number of tracers and the number of columns in the file. The first three columns are the coordinates of the tracers. The fourth column, if present, indicates the flavor of the tracers.

```
num_tracers num_columns
theta phi radius [flavor]
```

B.7 Reference State Files

This file contains the profiles of the reference state. This file must contain at least `nodesz` lines, and each line must contain 7 columns, where the meaning of each column is:

```
 $\rho_r$   $g$   $\alpha$   $c_P$  reserved reserved reserved
```


Appendix C

CitcomS Output File Format

C.1 Introduction

The format of the output files of CitcomS is described here. In the following sections, the model prefix is assumed as `test-case`, the processor number as 0, and the time step as 10. All outputs are in non-dimensional units unless specified.

C.2 Postprocessed Cap Output (`test-case.cap00.10`)

The command `autocombine.py` produces 1 cap file for regional CitcomS and 12 cap files for the full CitcomS. The first line of the cap file is a comment describing the node geometry (`nodex × nodey × nodez`). The rest of the file is column-based, where the meaning of each column is:

```
colatitude longitude radius vel_colat vel_lon vel_r temperature viscosity
```

C.3 Postprocessed Opt Output (`test-case.opt00.10`)

If the optional output contains node-based fields (e.g., `comp_nd`, `pressure`, and `stress`), the command `autocombine.py` produces 1 opt file for regional CitcomS and 12 opt files for the full CitcomS. The first line of the opt file is a comment describing the node geometry (`nodex × nodey × nodez`). The rest of the file is column-based, where the meaning of each column is defined in the corresponding `.general` file. For example, if the `.general` file contains these two lines, then the first column in the opt file is the composition, the 2nd column is the pressure, and the last six columns are the stress.

```
field = composition, pressure, stress  
structure = scalar, scalar, 6-vector
```

C.4 Postprocessed Surf Output (`test-case.surf0.10`)

An undocumented (and unmaintained) post-processing script `batchsurf.py` can combine the `coord`, `surf`, and `botm` output. The format of the combined surf file is listed for completion. The first line of the surf file is a comment describing the node geometry (`nodex × nodey`). The rest of the file is column-based, where the meaning of each column is:

```
colatitude longitude topography heat_flux vel_colat vel_lon
```

C.5 Time Output (`test-case.time`)

This file reports non-dimensional elapsed model time and spent CPU time (in seconds), and it is only outputted on computer node 0. The meaning of each column is:

```
step total_t delta_t total_cpu_time step_cpu_time
```

C.6 ASCII Output

C.6.1 Coordinate Output (`test-case.coord.0`)

This file is only outputted at the 0th time step. The first line is a header. The rest of the file is column-based, where the meaning of each column is:

```
colatitude longitude radius
```

C.6.2 Velocity and Temperature Output (`test-case.velo.0.10`)

The first two lines of this file are headers. The rest of the file is column-based, where the meaning of each column is:

```
vel_colat vel_lon vel_r temperature
```

C.6.3 Viscosity Output (`test-case.visc.0.10`)

The first line of this file is a header. The rest of the file is column-based, where the meaning of the only column is:

```
viscosity
```

C.6.4 Material Output (`test-case.mat.0`)

In this version of CitcomS, the implementation of material support is not working. The material output has been disabled and is documented here for completion. This file is only outputted at the 0th time step. There is no header. The meaning of each column is:

```
element_number layer material
```

C.6.5 Surface Variables Output (`test-case.surf.0.10` and `test-case.botm.0.10`)

The first line of each file is a header. The rest of each file is column-based, where the meaning of each column is:

```
topography heat_flux vel_colat vel_lon
```

C.6.6 Stress Output (`test-case.stress.0.10`)

The first two lines of the file are headers. The rest of the file is column-based, where the meaning of each column is:

```
 $S_{\theta\theta}$   $S_{\phi\phi}$   $S_{rr}$   $S_{\theta\phi}$   $S_{\theta r}$   $S_{\phi r}$ 
```

You can use the above values to form the following symmetric stress tensor:

$$\begin{bmatrix} S_{\theta\theta} & S_{\theta\phi} & S_{\theta r} \\ S_{\phi\theta} & S_{\phi\phi} & S_{\phi r} \\ S_{r\theta} & S_{r\phi} & S_{rr} \end{bmatrix} \quad (\text{C.1})$$

C.6.7 Pressure Output (`test-case.pressure.0.10`)

The first line of the file is a header. The rest of the file is column-based, where the meaning of the only column is:

```
pressure
```

C.6.8 Horizontal Average Output (`test-case.horiz_avg.0.10`)

The first line of the file is a header. The rest of the file is column-based, where the meaning of each column is:

```
radius temperature RMS(V_horizontal) RMS(V_vertical)
```

C.6.9 Geoid Output (`test-case.geoid.10`)

The first line of the file is a header. The rest of the file is column-based, where the meaning of each column is:

col 1	spherical harmonic degrees (l)
col 2	spherical harmonic orders (m)
col 3	cosine coefficients of total geoid (S_{lm})
col 4	sine coefficients of total geoid (C_{lm})
col 5	cosine coefficients of the geoid due to surface topography
col 6	sine coefficients of the geoid due to surface topography
col 7	cosine coefficients of the geoid due to internal density variation
col 8	sine coefficients of the geoid due to internal density variation

The units of the geoid coefficients are meters. The geoid field can be reconstructed by

$$\sum (S_{lm} \sin(m\phi) + C_{lm} \cos(m\phi)) P_{lm}(\cos\theta)$$

where P_{lm} is the associated Legendre polynomials.

C.6.10 Tracer Output (`test-case.tracer.0.10`)

The first line of the file is a header. The second field in the header is the number of tracers in the file. The third field in the header is the number of columns in the file. The rest of the file is column-based, where the meaning of the columns is:

```
theta phi radius [flavor]
```

C.6.11 Composition Output (`test-case.comp_e1.0.10` and `test-case.comp_nd.0.10`)

These files contain the composition field either on the element (`comp_e1`) or on the node (`comp_nd`). The format of the files is the same. The first two lines of the file are headers. The fourth field in the first line is the number of composition components. In the second line, the first field is the current bulk sum of composition 1. The second field is the initial bulk sum of composition 1. The third and fourth fields (if they exist) are the current and initial bulk sum of composition 2, etc. The rest of the file is column-based, where the meaning of the columns is:

```
composition1 [composition2 ... compositionN]
```

C.6.12 Heating Output (`test-case.heating.0.10`)

These files contain various heating terms on the element. The first two lines of the file are the headers. The rest of the file is column-based, where the meaning of the columns is:

```
adiabatic_heating viscous_heating latent_heating
```

C.7 HDF5 Output (`test-case.h5`)

The format and layout of HDF5 output is described in Section 4.4 on page 39.

C.8 Misc. Binary Output

C.8.1 Checkpoint Output (`test-case.checkpoint.0.10`)

These files are used for restarting the run. These files are not portable. You cannot run a model on machine A, copy the checkpoint files to machine B, and expect you can always restart the simulation on machine B. Their format is undocumented on purpose and will remain so.

C.8.2 Domain Output (`test-case.domain`)

This file is only outputted at the 0th time step. This file is output by rank-0 processor only. It contains the domain bounds of all processors in binary format. The file begins with four integers: `nproc`, `ncolumns`, `const1`, `const2`, where `nproc` is the number of processors, `ncolumns` is 10, `const1` and `const2` are used to detect binary incompatibility. The rest of the file has `nproc` \times `ncolumns` doubles, the columns contains the min/max of radius and the θ, ϕ coordinates of the four corners:

```
r_min r_max theta0 phi0 theta1 phi1 theta2 phi2 theta3 phi3
```

C.8.3 Coordinate Binary Output (`test-case.coord_bin.0`)

These files are only outputted at the 0th time step. These files contain the coordinates of nodes in binary format. The file begins with four integers: `nodex`, `nodey`, `nodez`, `const`, then followed by the θ, ϕ, r coordinates of all nodes.

C.8.4 Seismic Output (`test-case.seismic.0.10`)

These files contain the density, P wave velocity and S wave velocity in binary format. The density of all nodes is stored in double, then followed by P wave velocity of all nodes, and finally S wave velocity of all nodes.

Appendix D

License

GNU GENERAL PUBLIC LICENSE Version 2, June 1991. Copyright (C) 1989, 1991 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software – to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. Copyright the software, and
2. Offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program” below refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification.”) Each licensee is addressed as “you.”

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program’s source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- (a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- (b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- (c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- (a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- (b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- (c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version,” you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found. For example:

One line to give the program’s name and a brief idea of what it does. Copyright © (year) (name of author)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright © year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type 'show w'. This is free software, and you are welcome to redistribute it under certain conditions; type 'show c' for details.

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items – whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program 'Gnomovision' (which makes passes at compilers) written by James Hacker.

(signature of Ty Coon)

1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Bibliography

- [1] Moresi, L., M. Gurnis, and S. Zhong (2000), Plate tectonics and convection in the Earth's mantle: Toward a numerical simulation, *Comp. Sci. Engin.*, 2, 22-33.
- [2] Moresi, L.N., and V.S. Solomatov (1995), Numerical investigation of 2D convection with extremely large viscosity variations, *Phys. Fluid*, 7, 2,154-2,162.
- [3] Moresi, L., S. Zhong, and M.Gurnis (1996), The accuracy of finite element solutions of Stokes' flow with strongly varying viscosity, *Phys. Earth Planet. Inter.*, 97, 83-94.
- [4] Moresi, L.N., and M. Gurnis (1996), Constraints on the lateral strength of slabs from three-dimensional dynamic flow models, *Earth Planet. Sci. Lett.*, 138, 15-28.
- [5] Zhong, S., M. Gurnis, and L. Moresi (1998), The role of faults, nonlinear rheology, and viscosity structure in generating plates from instantaneous mantle flow models, *J. Geophys. Res.*, 103, 15,255-15,268.
- [6] Zhong, S., M.T. Zuber, L.N. Moresi, and M. Gurnis (2000), The role of temperature-dependent viscosity and surface plates in spherical shell models of mantle convection, *J. Geophys. Res.*, 105, 11,063-11,082.
- [7] Tan, E., M. Gurnis, and L. Han (2002), Slabs in the lower mantle and their modulation of plume formation, *Geochem. Geophys. Geosys.*, 3, 1067.
- [8] Conrad, C.P., and M. Gurnis (2003), Seismic tomography, surface uplift, and the breakup of Gondwanaland: Integrating mantle convection backwards in time, *Geochem. Geophys. Geosys.*, 4(3), 1031, doi:10.1029/2001GC000299.
- [9] Hughes, T.J.R. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.; 1987. 672 p.
- [10] Ramage, A., and A.J. Wathen (1994), Iterative solution techniques for the Stokes and Navier-Stokes equations, *Int. J. Numer. Methods. Fluids*, 19, 67-83.
- [11] Brooks, A.N. *A Petrov-Galerkin Finite Element Formulation for Convection Dominated Flows*. Unpublished doctoral thesis, California Institute of Technology, Pasadena, CA, 1981.
- [12] Cahouet, J., and J.-P. Chabard (1988), Some fast 3D finite element solvers for the generalized Stokes problem, *Int. J. Numer. Methods. Fluids*, 8, 869-895.
- [13] Atanga, J., and D. Silvester (1992), Iterative methods for stabilized mixed velocity-pressure finite elements, *Int. J. Numer. Methods. Fluids*, 14, 71-81.
- [14] Hager, B.H., and R.J. O'Connell (1981), A simple global model of plate dynamics and mantle convection, *J. Geophys. Res.*, 86, 4,843-4,867.
- [15] Tan, E., E. Choi, P. Thoutireddy, M. Gurnis, and M. Aivazis (2006), GeoFramework: Coupling multiple models of mantle convection within a computational framework, *Geochem., Geophys., Geosyst.* 7, Q06001, doi:10.1029/2005GC001155.
- [16] Zhong, S., M. Gurnis, and L. Moresi (1996), Free-surface formulation of mantle convection—I. Basic theory and application to plumes, *Geophys. J. Int.*, 127, 708-718.

- [17] Zhong, S., A. Paulson, and J. Wahr (2003), Three-dimensional finite-element modeling of Earth's viscoelastic deformation: effects of lateral variations in lithospheric thickness, *Geophys. J. Int.*, *155*, 679-695.
- [18] McNamara, A.K., and S. Zhong (2004), Thermochemical structures within a spherical mantle: Superplumes or Piles? *J. Geophys. Res.*, *109*, B07402, doi:10.1029/2003JB002847.
- [19] Tackley, P.J., and S.D. King (2003), Testing the tracer ratio method for modeling active compositional fields in mantle convection simulations, *Geochem. Geophys. Geosyst.*, *4*, 8302, doi:10.1029/2001GC000214.
- [20] Trampert, J., P. Vacher, and N. Vlaar (2001), Sensitivities of seismic velocities to temperature, pressure and composition in the lower mantle, *Phys. Earth Planet. Inter.*, *124*, 255-267.