

# CIG Science Gateway and Community Codes for the Geodynamics Community

January 19, 2016

## Project Overview

The Computational Infrastructure for Geodynamics (CIG), an NSF cyber-infrastructure facility, aims to enhance the capabilities of the geodynamics community through developing scientific software that addresses many important unsolved problems in geophysics. CIG's strategy is to:

1. support the benchmarking and validation of its codes,
2. develop new codes and ensure they achieve good performance and scalability, and
3. assist new users by providing technical support, training, and small allocations of computation time.

These efforts have met with success, and the current CIG compute allocations on the XSEDE infrastructure have been used at a substantial rate to achieve these goals.

CIG supports the aforementioned efforts in the following areas of activity: geodynamo simulation, mantle dynamics, seismic wave propagation, and crustal and lithospheric dynamics on both million-year and earthquake time-scales.

In this proposal, we request support to continue these activities and to test next-generation, large-scale computational codes for use in geophysics. In the next section, we describe the major scientific questions and computing challenges that CIG focuses on. We then describe the codes and methodologies used and offer a justification of the requested resources.

## Science Objectives

**Core Dynamo and Dynamics.** Numerical simulations play a large role in elucidating fluid motion in the Earth’s outer core and the resulting geomagnetic field (geodynamo). Although previous work [after Glatzmaier and Roberts, 1995] has successfully reproduced some spatial and temporal characteristics of the geomagnetic field, a large discrepancy still exists between the parameters used in geodynamo simulations and actual values associated with the outer core. This discrepancy reflects the extremely low viscosity of the liquid outer core relative to surrounding layers. The low viscosity results in a vast range of length scales required for a comprehensive simulation, ranging from the geometry of the outer core ( $L \sim 1000\text{km}$ ) to the thickness of the boundary layer ( $L \sim 0.1\text{m}$ ). Computational resources are still insufficient to achieve this level of resolution, but the community is working to target a middle range ( $L \sim 100\text{m}$ ) that can be achieved using the cutting edge numerical methods and high-end supercomputers available today.

**Mantle Dynamics.** Mantle convection is at the heart of understanding how the Earth works, but the process remains at best poorly understood because the mantle is not accessible to direct observation. Progress on fundamental questions, such as the dynamic origin of the tectonic plates that cover the surface, layering and stratification within the mantle, evolution of the thermal history of the Earth and its geochemical cycles, the interpretation of seismic tomographic models of Earth’s interior structure, and the source of volcanic hotspots, all require an interdisciplinary approach. Numerical models of mantle convection must therefore assimilate information from a wide range of disciplines, including seismology, geochemistry, mineral and rock physics, geodesy, and tectonics.

The technical challenges associated with modeling mantle convection are substantial. Mantle convection is characterized by strongly variable (i.e., stress-, temperature-, and pressure-dependent) viscosities. The lithosphere exhibits processes such as fracture and shear zone deformation (strain localization) that are physically distinct from the viscous flow deeper in the mantle, and occur on fundamentally different (smaller) length scales. In addition, the mantle is chemically heterogeneous, is replete with silicate melts and volatiles, and has numerous pressure- and temperature-induced structural changes that affect its dynamics.

**Crustal and Lithospheric Dynamics: Million Year Timescales.** The lithosphere, with the embedded crust, represents the main thermal boundary layer of the Earth's heat engine and, as such, encompasses a wide range of pressure and temperature conditions with diverse deformational mechanisms. Recently, deep seismic profiling, receiver function analysis, and magnetotelluric sounding have greatly increased our understanding of crustal and lithospheric structure. Numerical modeling has become an essential step in the integration of these data into process-orientated models of mountain building, lithospheric stretching, sedimentary basin genesis, and plate boundary deformation.

Deformation of the lithosphere presents a number of challenges to numerical simulation. The deep lithospheric mantle encompasses a differential temperature of up to 1000°C and an effective viscosity contrast of many orders of magnitude. The complex physics of frictional materials is particularly challenging because it involves strain-localization, time- and rate-dependent yield strength and strain softening. Crustal deformation is a free-surface problem and sensitive to the complexities of the Earth's surface, including physical and chemical erosion, mass transport by rivers and ocean currents, and deposition of sediment. There are also broad implications for the feedback between erosion and tectonic uplift. Climate change during the late Cenozoic has influenced sediment (and thus geochemical) fluxes to the ocean and atmosphere, and the way in which crustal dynamics modulates the erosional response of the Earth to climate change remains an open question.

**Crustal Dynamics: Earthquake time-scales.** A rapidly advancing area of crustal geodynamics, one of great societal importance, is the problem of the physics of the earthquake cycle. Because of the recent development of the capability for high-accuracy measurement of deformation of the Earth's surface in real time, this field, long starved for data, is now a burgeoning observational science. Recent observations made with high precision space geodesy indicate that displacements caused by slow aseismic motions following earthquakes can be comparable to coseismic displacements, demonstrating substantial post-seismic evolution of strain and stress in addition to coseismic changes.

It has recently been recognized that relatively modest changes in stress can trigger earthquakes. Theoretical advances in rock mechanics have led to algorithms relating temporal variations in stress to changes in earthquake

activity, and are beginning to make possible quantitative predictions of how stress changes from fault interactions influence seismicity levels. For example, a 3D finite element model of the Coulomb stress has addressed whether the 1999 Hector Mine earthquake was triggered by the 1992 Landers earthquake. Although results from models such as these have been impressive, more definitive tests require an order of magnitude finer nodal spacing, meshes incorporating the actual elastic structure of the region, the interaction of many faults, and more realistic rheologies.

**Seismic Wave Propagation.** Seismology provides the means to image the three-dimensional structures within the Earth’s interior that are responsible for geodynamic processes. The foundation of computational seismology is the generation of synthetic seismograms and adjoint methods, used in the modeling and inversion for Earth structure, earthquake rupture, and wave propagation effects. CIG aids the community by making 3D codes that provide a more-accurate representation of Earth properties such as anisotropy, attenuation and gravitational affects available to the community. Such 3D codes are now revolutionizing seismology, by allowing a direct investigation of countless geodynamic topics such as the fate of subducted lithosphere, existence of mantle plumes, lithospheric structure, and plate boundary zone complexity.

**Infrastructure.** Investigation into these vital Earth science issues has generally been hampered by lack of computational power to model or simulate Earth structures. Most geophysical processes are complex, coupled, and impossible to solve analytically or simulate in a laboratory, hence, a long-term, sustained effort in model building and large-scale simulation is needed. Geophysical models and codes have reached a level of maturity that allows and requires large-scale 3D coupled simulations, but need substantial infrastructure in order to be run. CIG facilitates solutions by developing open-source geodynamics software that addresses such problems, and by supporting workshops, training sessions, and conferences in the above sub-disciplines. However, even with investments by universities and institutions in small- to medium-sized clusters, a large number of problems in geodynamics still require more powerful capabilities. CIG intends to continue developing and benchmarking its codes, conducting training sessions on its applications, and encouraging new users to try XSEDE resources to see if they can be applied effectively to

their research problem. This proposal details five major geodynamics software packages that CIG believes to be most important to the community performing research in mantle convection, planetary dynamos, seismology, and short/long time-scale tectonics.

## Computational Experiments and Resource Requirements

### Numerical Approaches

**Calypso.** Calypso is a newly developed code to perform magnetohydrodynamics (MHD) simulations in a rotating spherical shell modeled on the Earth’s outer core. Convection in this core is driven by the temperature difference between the outer and inner boundaries of the fluid shell. It uses a pseudospectral method in the horizontal discretization and a finite difference method in the radial discretization. Linear terms (e.g. diffusion, buoyancy, Coriolis force) are evaluated in spherical space, while non-linear terms (advection, Lorentz force, magnetic induction) are evaluated in the physical space. For time integration, Calypso uses a Crank-Nicolson scheme for the diffusion terms and second-order Adams-Bashforth scheme for the other terms.

**Rayleigh.** Rayleigh is an open-source community dynamo code developed by Nick Featherstone (CU Boulder) with sponsorship by CIG. This code solves the three-dimensional, nonlinear, MHD equations of motion for a compressible fluid in a rotating spherical shell under the anelastic approximation. Rayleigh employs a pseudo-spectral algorithm with spherical harmonic basis functions and mixed explicit/implicit time-stepping (Adams-Bashforth/Crank-Nicolson). A poloidal/toroidal representation ensures that the mass flux and magnetic field remain solenoidal. Rayleigh has been designed with petascale computing in mind; it has been performance tested extensively on NASA’s SGI Pleiades system and Argonne’s Blue Gene/Q system, Mira. Rayleigh demonstrates strong scaling with 80% of ideal efficiency up to 131,072 cores for  $2048^3$  problem sizes using pure MPI.

**ASPECT.** ASPECT is a CIG developed code intended to solve the equations that describe thermally driven convection with a focus on doing so in

the context of convection in the earth mantle. It allows for both 2D and 3D models of arbitrary shapes (generally focused on segments or whole mantle models), adaptive mesh refinement in locations of scientific interest, easy modification of material, gravity, viscosity and temperature models, and tracers to model geochemistry and material transport. Recent work has started investigating the effectiveness of GPU or MIC coprocessors in ASPECT simulation. Further details are available in [Kronbichler, *et al.* 2011].

**CitcomS.** CitcomS is a finite element code to solve thermo-chemical convection problems relevant to the planetary mantle in a 3D spherical geometry [Moresi and Solomatov, 1995; Moresi and Gurnis, 1996, Zhong *et al.*, 2000]. There are two forms of meshes and geometries for CitcomS, regional and spherical. CitcomS employs an Uzawa algorithm to solve the momentum equation coupled with the compressibility constraints [Moresi and Gurnis, 1996; Ramage and Wathen, 1994]. Nested inside the Uzawa algorithm, the code uses either a conjugate gradient solver or a multi-grid solver to solve the discretized matrix equations. The energy equation is discretized in the Streamline Upwind Petrov-Galerkin method [Brooks, 1981] and integrated with an explicit second-order predictor-corrector method.

**PyLith.** PyLith is a 2D and 3D finite-element code for modeling interseismic and seismic processes related to capturing the physics of earthquakes, including slow strain accumulation, sudden dynamic stress changes during earthquake rupture, and slow postseismic relaxation. Implicit time-stepping provides efficient time integration for quasi-static (interseismic deformation) problems, and explicit time-stepping provides efficient time integration for dynamic (rupture and wave propagation) problems. Key features of PyLith are its ability to accommodate unstructured meshes (which allows larger variations in discretization size and complex nonplanar fault geometry), implementation of a variety of finite-element types (e.g., higher order elements as well as conventional linear and parabolic tetrahedral and hexahedral elements), and implementation of a variety of fault and bulk constitutive models appropriate for the Earth’s lithosphere. The bulk constitutive models include linear and nonlinear viscoelastic models in addition to linear elastic models. PyLith uses PETSc [Balay *et al.*, 1997, 2001, 2004] to achieve fast, efficient, parallel solution of the partial differential equation.

**SPECFEM3D\_GLOBE.** In collaboration with Princeton, Caltech and the University of Pau (France), CIG offers this software, which simulates global and regional (continental-scale) seismic wave propagation using the spectral-element method (SEM). The SEM is a continuous Galerkin technique, which can easily be made discontinuous; it is then close to a particular case of the discontinuous Galerkin technique, with optimized efficiency because of its tensorized basis functions. In particular, it can accurately handle very distorted mesh elements [Oliveira and Seriani, 2011].

SPECFEM3D\_GLOBE has very good accuracy and convergence properties [De Basabe and Sen, 2007; Seriani and Oliveira, 2008]. The spectral element approach admits spectral rates of convergence and allows exploiting hp-convergence schemes. It is also very well suited to parallel implementation on very large supercomputers [Carrington *et al.*, 2008; Komatitsch *et al.*, 2010a] as well as on clusters of GPU accelerating graphics cards [Komatitsch, 2010b].

Further details regarding each code and downloads of the source are available at the following URL in Table 1 (Rayleigh is not publicly available yet).

Table 1: List of Websites

Code	Website
Calypso	<a href="https://geodynamics.org/cig/software/calypso/">https://geodynamics.org/cig/software/calypso/</a>
Rayleigh	<a href="https://www.youtube.com/watch?v=km0Bv6p2U08">https://www.youtube.com/watch?v=km0Bv6p2U08</a> <a href="https://www.youtube.com/watch?v=6u0P-pyJsXo">https://www.youtube.com/watch?v=6u0P-pyJsXo</a>
ASPECT	<a href="https://geodynamics.org/cig/software/aspect/">https://geodynamics.org/cig/software/aspect/</a>
CitcomS	<a href="https://geodynamics.org/cig/software/citcoms/">https://geodynamics.org/cig/software/citcoms/</a>
PyLith	<a href="https://geodynamics.org/cig/software/pylith/">https://geodynamics.org/cig/software/pylith/</a>
SPECFEM3D_GLOBE	<a href="https://geodynamics.org/cig/software/specfem3d_globe/">https://geodynamics.org/cig/software/specfem3d_globe/</a>

## Resource Requirements

CIG researchers used a significant portion of the past period’s allocation to perform code benchmarking in 2014-2015. In July 2015, CIG extended its allocation to March 2016 after work temporarily slowed due to the departure of a member of the code development team. In the upcoming period, we anticipate using SUs at higher rate than the previous period, due to the ongoing development, testing, and the ramp up of benchmarking and research use of CIG codes.

CIG plans the following use of its proposed XSEDE resources during the period of April 1st, 2016 to March 31, 2017 in support of (1) scalability testing and code validation, (2) development of new numerical methods for better code performance, (3) workshop training sessions, and (4) nurturing new geophysics users on XSEDE resources using Calypso, Rayleigh, ASPECT, CitcomS, PyLith and SPECFEM3D\_GLOBE. New users anticipate several millions of SUs will be required to conduct their research in which CIG expects to support the feasibility testing and spin up which will enable researchers to apply for their own allocations. More details are provided below.

**ASPECT, Calypso and Rayleigh development.** ASPECT mantle convection code and the Calypso and Rayleigh geodynamo codes are continuing development and scaling work. This will be primarily done by CIG researchers at TAMU, Dr. Hiroaki Matsui, and Dr. Nick Featherstone, respectively. The allocation will be used to establish the scaling performance and efficiency of each code, add functionality, and improve the support for Xeon Phi and/or GPGPU based computation. To perform simulations to ensure the validity of the codes and check their scalability and performance, we anticipate requiring up to 4096 cores for brief periods (1-4 hours) and estimate a total requirement of 50,000 SUs for each code (150,000 in total) on Stampede for this development. We also request 15,000 SUs on Maverick for development of GPGPU based computation.

**Geodynamo multi-scale convection modeling.** Dr. Hiroaki Matsui plans to investigate the dynamics of turbulence in a planetary dynamo using Calypso. Because the current model based on an accurate Earth core dynamo requires extremely high resolution (on the order of 10 million cores for



1000 hours), Dr. Matsui is currently working on improving model scale, code performance and sophisticated turbulence modeling. To establish the validity of this approach, it is necessary to do medium scale runs on Stampede. The code will also be developed further to improve support for accelerators. We request to run with (384,288,576) grids for 5,000,000 time steps using 1024 cores. It will require  $(1024 \times 480) = 490,520$  SUs.

**Present-Day Convection and Lithospheric Deformation Models.**

The observed velocity and deformation of Earth’s tectonics plate reflects forces arising from density variations in both the convecting mantle and lithosphere. Using recent models of the lithospheric and deep mantle density field, Dr. John Naliboff will use ASPECT to simultaneously model global, present-day mantle convection and lithospheric deformation. These models will build extensively on existing functionality in ASPECT, which includes assigning initial conditions based on prescribed surface topography and depth-dependent density structure. As resolving the Earth’s free surface and shallow density structure requires locally high-resolutions (5-10 km), we estimate that the models will likely contain 200-300 million DOFs based on previous global simulations in ASPECT. As the models solve for present-day convection and near surface deformation, they require only 1-5 total time steps. Using previous ASPECT models with a similar size (200 million DOFs) and short runtime (200,000 years) as a reference, we estimate each model will likely require approximately 50,000 SUs using 1500 processors. In order to rigorously test the effects of near-surface mesh resolution and examine a range of initial density and rheological structures, we estimate a total requirement of 2-3 million SUs to complete this study. Here, we request 500,000 SUs on Stampede in order to perform initial scaling and parameter tests.

**High-Resolution Continental Extension Models.** Dr. John Naliboff will use ASPECT to begin a project examining 3-D continental extension processes. At present, high-resolution 2-D numerical models of long-term continental extension provide significant insights into the tectonic processes controlling specific structural and geophysical observations in both active and prior regions of lithospheric extension (e.g., Brune *et al.* 2014; Naliboff and Buiter, 2015). Here, we propose to adapt and use ASPECT to conduct a series of high-resolution 3-D lithospheric extension simulations that

build on such 2-D simulations. To provide estimates of numerical robustness, the results of these models will be directly compared with equivalent simulations run with different long-term tectonics codes that utilize distinct numerical techniques. Broadly, this project aims to gain further understanding of continental extension processes and provide a framework for future 3-D lithospheric modeling comparisons.

Based on recent high-resolution ( $5 \times 10^6$  elements) lithospheric extension experiments run on a massively scalable long-term tectonics code, we estimate that a single model run with ASPECT will require 24 hours across approximately 5,000 CPUs (7,500 SUs). Using these estimates, 2 million SUs would provide resources to run both 100-200 high-resolution simulations and a suite of scaling tests on Stampede. Here, we request 200,000 SUs to perform initial scaling tests and examine a series of preliminary model designs.

**High-Resolution Subduction Models.** Prof. Margarete Jadamec will be using CitcomCU (a variant of CitcomS) to continue the research into high-resolution subduction models with multiple plates. Observational and experimental constraints indicate plate boundaries are inherently three-dimensional and are characterized by lateral strength variations. For example, a power law rheology (one that includes the effects of the dislocation creep) can explain both observations of seismic anisotropy and the decoupling of mantle flow from lithospheric plate motion, due to the dynamic reduction of mantle wedge viscosity (Jadamec and Billen, 2010, 2012). However, large viscosity variations occurring over short distances pose a challenge for computational codes, and models with complex 3D geometries require substantially greater numbers of elements, increasing the computational demands (Spera *et al.*, 1982; Moresi and Solomatov, 1995; Moresi *et al.*, 1996; Tackley, 1996; Zhong *et al.*, 2000; May and Moresi, 2008; Geenen *et al.*, 2009; Burstedde *et al.*, 2009; Jadamec and Billen, 2010; Furuichi *et al.*, 2011). High-resolution subduction models with multiple plates are particularly computationally challenging, because the model domain contains numerous regions with large viscosity contrasts and complex non-linear flow. In addition, the models incorporate a high level of geometric complexity due to the geophysical data incorporated to construct accurate geometries representative of actual plate boundaries. Such models require 3000-8000 processors for 1-3 hours and 2-3 high-resolution models with the optimized parameters run for 24 hours with a usable scientific result. (Jadamec and Fischer, 2013).

**Melt Impact on Mantle Flow and Plume Generation** Drs. Julianne Dannberg and Rene Gassmoeller will investigate how grain size evolution and grain size dependent rheology in the Earth’s mantle influence the mantle flow field and seismic properties of the mantle, in particular with regard to the evolution of mantle plumes. Feedback between grain size reduction due to high stresses and the implied reduction in viscosity leads to shear localisation and a viscosity minimum at the edges of mantle plumes, indicating a potential to ascend and erode the lithosphere faster than previously expected.

Related to the requirements of these earth-like high-resolution, highly non-linear models, they strive to improve ASPECTS’s scaling to a higher number of processors, also studying the effect of spatially varying viscosity on the performance, and including cases with and without melt migration.

Finally, they will investigate the self-consistent plume generation and ascent in interaction with the surrounding mantle flow, also including the influence of melt generation and migration on the distribution and composition of the generated magmas of Large Igneous Provinces and ocean island basalts compared with observations.

Previous research runs on similar resources showed that:

- 3D global mantle convection models with a resolution of 25 km (100 million DOFs) and a model runtime of 250 million years, took 90k SUs, running on 768 processors.
- two-phase flow models, a 3D model of a single mantle plume (600 x 600 x 300 km, 200 million DOFs, highest resolution of 0.5 km) and a model runtime of 200 thousand years running took 90k SUs with 1536 MPI processes
- grain-size evolution models, a global mantle convection models in 2D with a finest resolution of 5 km (6.5 million DOFs) and a model runtime of 300 million years, took 280k SUs, running on 768 processors.

Hence, we request 460k SUs for these simulations.

**Von Karmann sodium experiment.** Prof. Jean-Luc Guermond plans to perform initial tests in duplicating the Von Karmann sodium experiment on a computational platform. This initial work will confirm that the experiment can be replicated at lower resolutions, with the aim of using the results to apply for a larger allocation on Stampede at a later date.

**Earthquake Modeling** Prof. Surendra Somala plans to use SPECFEM3D to further seismic hazard studies in strong ground motions, critical ground motion simulations and seismic source inversions in 3D heterogenous media. Recent work indicates that the spectral element method (SPECFEM3D) is a promising tool to model dynamic rupture processes that resolve cohesive zones, which are typically observed from rotary shear experiment on crustal rocks. For such models to be successfully, the grid spacing must be finer than the cohesive zone width, yet the model domain must still be large enough to contain an entire fault plane that can accommodate large earthquake. A typical forward simulation of SPECFEM3D requires  $N_t \times N_e \times 9(2 \times N_{gll} + 9) \times N_{gll}^3$  floating point operations, where:

- $N_t$  is the number of time-steps per simulation
- $N_e$  is the number of spectral elements contained in the mesh
- $N_{gll}$  is the number of Gauss-Labatto-Legendre (GLL) nodes per element

Great earthquakes typically last for minutes. The timesteps required to resolve cohesive zones at meter length-scales are consequently on the order of milliseconds. Seismic velocities in the crust are approximately 3 km/s, but to include shallow low velocity layers the time step has to be further decreased by an order of magnitude. A few million spectral elements are required to model strong motion shaking so as to encompass regions of perceivable shaking. Typically 5 GLL nodes are found to give best performance in spectral element methods. On typical XSEDE facilities using 256 processors, one simulation requires several hours to a day. Investigations of hazard scenerios require 1000's of such runs totaling approximately 3 million SUs. We will request 15% of SUs (450k SUs) for initial testings.

In total, a yearly allocation of 2,491,520 SUs will enable CIG to continue offering support and training to users of these common geophysics codes. This will also allow extensive studies of code accuracy, performance and validation using high-resolution simulations. We also request 20,000 SUs on a visualization oriented system such as Maverick to assist in analyzing and visualizing simulation results and to develop GPGPU based computations.

Table 2: Summary of requested SUs

Software	Purpose	Requested SUs
Calypso/Rayleigh	Development	100,000
	Geodynamo Turbulence studies	491,520
	Von Karmann sodium experiment	200,000
ASPECT	Development/Overshoot	50,000
	Present-Day Global Deformation Models	500,000
	Continental Extension Models	200,000
CitcomS	High-resolution subduction models	500,000
SPECFEM3D	Earthquake Modeling	450,000
Total for Stampede		2,491,520
ParaView/VisIt	Visualization	5,000
CUDA Calypso	Development	15,000
Total for Maverick		20,000

## References

Balay, S., W.D. Gropp, L. Curfman McInnes, and B.F. Smith, Efficient Management of Parallelism in Object Oriented Numerical Software Libraries, *Modern Software Tools in Scientific Computing*, Edited by A.M. Bruaset and H. P. Langtangen, 163–202, Birkhauser Press, 1997.

Balay, S., K. Buschelman, W.D. Gropp, D. Kaushik, M.G. Knepley, L. Curfman McInnes, B.F. Smith, and H. Zhang, PETSc Web page, <http://www.mcs.anl.gov/petsc/>, 2001.

Balay, S., K. Buschelman, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L. Curfman McInnes, B.F. Smith, and H. Zhang, *PETSc Users Manual*, ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.

Brooks, A.N., A Petrov-Galerkin Finite Element Formulation for Convection Dominated Flows. Unpublished doctoral thesis, California Institute of Technology, Pasadena, CA, 1981.

Brune, S., C. Heine, M. Prez-Gussiny, and S. Sobolev, *Rift migration explains continental margin asymetry and crustal hyper-extension*, **377**, doi: 0.1038/ncomms5014, 2014.

Burstedde, C., O. Ghattas, G. Stadler, T. Tu, and L. Wilcox, Parallel

scalable adjoint-based adaptive solution of variable-viscosity stokes flow problems, *Computer Methods in Applied Mechanics and Engineering*, **198**(21-26), 1691–1700, 2009.

Carrington, L., D. Komatitsch, M. Laurenzano, M. Tikir, D. Micha, N. Le Goff, A. Snavely, and J. Tromp. High-frequency simulations of global seismic wave propagation using SPECFEM3D\_GLOBE on 62 thousand processor cores. *Proceedings of the ACM/IEEE Supercomputing SC2008 conference*, 1–11, doi: 10.1145/1413370.1413432. Article #60, Gordon Bell Prize finalist article, 2008.

De Basabe, J. D. and M. K. Sen. Grid dispersion and stability criteria of some common finite-element methods for acoustic and elastic wave equations. *Geophysics*, **72**(6), T81–95, doi: 10.1190/1.2785046, 2007.

Furuichi, M., D. May, and P. Tackley, Development of a stokes flow solver robust to large viscosity jumps using a schur complement approach with mixed precision arithmetic, *Journal of Computational Physics*, **230**, 8835–8851, 2011.

Geenen, T., M. ur Rehman, S. MacLachlan, G. Segal, C. Vuik, A. van den Berg, and W. Spakman, Scalable robust solvers for unstructured fe geodynamic modeling applications: Solving the stokes equation for models with large localized viscosity contrasts, *Geoch. Geoph. Geosyst*, **10**, Q09,002, 2009.

Glatzmaier, G.A., and P. H. Roberts. A three-dimensional self-consistent computer simulation of a geomagnetic field reversal. *Nature*, **377**, 203–209, 1995.

Jadamec, M. A., and M. I. Billen, Reconciling surface plate motions and rapid three-dimensional flow around a slab edge, *Nature*, **465**, 338–342, 2010.

Jadamec, M. A., and M. I. Billen, The role of rheology and slab shape on rapid mantle flow: Three- dimensional numerical models of the Alaska slab edge, *Journal of Geophysical Research*, **117** (B02304), doi:10.1029/2011JB008563, 2012.

Jadamec, M.A., and Fischer, K.M., Mantle Response to a Slab Gap and Three-dimensional Slab Interaction in South America, Abstract DI24A-05 presented at 2013, Fall Meeting, AGU San Francisco, CA, 9-13 December, 2013.

Komatitsch, D., L. P. Vinnik, and S. Chevrot. SHdiff/SVdiff splitting in an isotropic Earth, *J. Geophys. Res.*, **115**(B7):B07312, doi: 10.1029/2009JB006795, 2010a.

Komatitsch, D., D. Gddecke, G. Erlebacher, and D. Micha. Modeling

the propagation of elastic waves using spectral elements on a cluster of 192 GPUs. *Computer Science Research and Development*, **25(1-2)**, 75–82, doi: 10.1007/s00450-010-0109-1, 2010b.

Kronbichler, M., T. Heister, W. Bangerth, High Accuracy Mantle Convection Simulation through Modern Numerical Methods, Submitted to *Geophysics Journal International*, 2011.

Moresi, L.N., and V.S. Solomatov, Numerical investigation of 2D convection with extremely large viscosity variations, *Phys. Fluid*, **7**, 2154–2162, 1995.

Moresi, L.N., and M. Gurnis, Constraints on the lateral strength of slabs from three-dimensional dynamic flow models, *Earth Planet. Sci. Lett.*, **138**, 15–28, 1996.

Moresi, L., S. Zhong, and M. Gurnis, The accuracy of finite element solutions of Stokes flow with strongly varying viscosity, *Physics of the Earth and Planetary Interiors*, **97**, 83–94, 1996.

Naliboff, J. and S.H. Buiter, Rift reactivation and migration during multiphase extension, *Earth and Planetary Science Letters*, **421**, 58–67, doi: 10.1016/j.epsl.2015.03.050, 2015.

Oliveira, S.P., and G. Seriani. Effect of element distortion on the numerical dispersion of spectral element methods. *Communications in Computational Physics*, **9(4)**, 937–958, 2011.

Ramage, A., and A.J. Wathen, Iterative solution techniques for the Stokes and Navier-Stokes equations, *Int. J. Numer. Methods. Fluids*, **19**, 67–83, 1994.

Serian, G., and S.P. Oliveira. Dispersion analysis of spectral-element methods for elastic wave propagation, *Wave Motion*, **45**, 729–744. doi: 10.1016/j.wavemoti.2007.11.007, 2008

Tackley, P., Effects of strongly variable viscosity on three-dimensional compressible convection in planetary mantles, *Journal of geophysical research*, **101(B2)**, 3311–3332, 1996.

Zhong, S., M.T. Zuber, L. Moresi, and M. Gurnis, Role of temperature-dependent viscosity and surface plates in spherical shell models of mantle convection, *Journal of Geophysical Research*, **105(B5)**, 11,063–11,082, 2000.