CIG Science Gateway and Community Codes for the Geodynamics Community

July 14, 2022

Project Overview

The Computational Infrastructure for Geodynamics (CIG), an NSF cyberinfrastructure 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: dynamo 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 Dynamics and Dynamos. It is widely accepted that the Earth's outer core consists of liquid iron alloy and that the geomagnetic field is sustained by its convection (so called geodynamo). Numerical simulations have played a large role in elucidating the fluid motion in the Earth's outer core and geodynamo processes. Although previous efforts (after Glatzmaier and Roberts, 1995) have 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 due to the small viscosity. The low viscosity results in a vast range of length scales of the flow required for a comprehensive simulation, ranging from the geometry of the outer core ($L \sim 1000$km) to the thickness of the boundary layer ($L \sim 0.1$m). Computational resources are still insufficient to achieve this level of resolution but the community is working to target a middle range ($L \sim 100$m) that can be achieved using the cutting edge numerical methods and high-end supercomputers available today.

Mantle and Lithospheric Dynamics. Mantle convection and lithospheric dynamics is at the heart of understanding how plate tectonics on Earth works. However, the processes governing plate tectonics remain poorly understood due the lack of direct observations from the deep part of the Earth and the highly non-linear nature of solid Earth deformation. Progress on fundamental questions, such as the dynamic origin of the tectonic plates that cover the surface, coupling between lithospheric deformation and convection deep inside the mantle, feedbacks between erosion and tectonic uplift, the drivers and dynamics of observed magmatism and volatile transport, and the evolution of plate boundary systems all require an interdisciplinary approach. Numerical models of mantle convection and lithospheric dynamics must therefore assimilate information from a wide range of disciplines, including seismology, geochemistry, mineral and rock physics, geodesy, and diverse geologic data sets.

The technical challenges associated with modeling mantle convection and lithospheric dynamics are substantial. They are characterized by strongly variable (i.e., stress-, temperature-, and pressure-dependent) viscosities. The lithosphere exhibits processes such as elastic flexure and brittle 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 and lithosphere are chemically heterogeneous, replete with silicate melts and volatiles, and have numerous pressure- and temperature-induced structural changes that affect its dynamics.

While substantial progress has been made in recent years in large part to due to improved numerical methods, scaling for massive 3D simulations, and new techniques for assimilating geophysical data sets, significant challenges remain across a wide range of topics. Here, we propose to develop and test new numerical methods in computationally expensive 2D and 3D simulations towards the following scientific goals:

1. Stabilization and accurate decomposition of non-linear viscoelastic-plastic deformation, which governs deformation within the lithosphere

2. Improvement of new nonlinear solvers for models using matrix-free preconditioners (Geometric Multigrid) and two-phase flow (see below).

3. Improving large scale IO (loading of large datasets, MPI combined with large amount of input data, faster graphical output, generic intermediate exchange format).

4. Assessing the accuracy and efficiency of particles versus compositional fields for tracking viscoelastic stresses.

5. Testing of various new features for simulating lithospheric dynamics
6. Improved methods for calculating crystal preferred orientations (CPOs) during dynamic simulations
7. Further integration with external software packages that provide a platform for geologic data assimilation and assembly of complex initial conditions.
8. Improving state of the art high-resolution global mantle flow models incorporating plate boundary databases, databases of subducted plates, and a grain-size dependent rheology.

The details of these topics and associated work plans are outlined within the Resource Requirements section.

Coupling Core and Mantle Dynamics. The timing of critical events in Earth's evolution, such as the evolution of Earth's inner core, are highly debated, as large uncertainties of core conductivity impede the definition of a well-defined cooling history. Long-term ($> 10^7$ years) temporal variations in the geomagnetic field, as seen in paleomagnetic data, can be tied to deep Earth processes predicted from numerical geodynamo simulations, and are therefore widely studied as a proxy to study the evolution of Earth's deep interior. With this approach, multiple possible timings for the onset of inner core nucleation (ICN) have been identified from the observation of extreme variations in the ancient geomagnetic field at 550 Ma and between 1.5-1.0 Ga. With the currently available geodynamo simulations, it is difficult to ascertain which – if any – of these signals corresponds to the onset of ICN. This is because the influence of mantle convection, and the resulting changes in amplitude and pattern of heat flux at the core-mantle boundary (CMB), on the geodynamo are still poorly understood. New estimates for the expected range of CMB heat flux variability from mantle convection models run in ASPECT can now be systematically coupled to geodynamo simulations that have demonstrated the ability to produce Earth-like magnetic fields.

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 inbetween and 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 enable quantitative predictions of how stress changes from fault interactions and fluid injection influence seismic cycles. 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 evidence of 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 source, and wave propagation effects. CIG aids the community by supporting 3D codes that provide a more-accurate representation of Earth's structure and properties including anisotropy, attenuation, and gravitational affects. 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, the nature of ULVZ's, lithospheric structure, and plate boundary zone complexity.

Infrastructure. An important role of CIG for the geoscience community lies in the maintenance of geodynamics software packages and in training and supporting the computational needs of their users. The XSEDE allocations enabled part of this work that pertains to large scale computations in two ways: First, running on the various leading edge systems allowed us to improve the software packages and their underlying software stack to run on XSEDE resources. Second, we developed installation and usage guides on XSEDE resources that allows users to get started quickly and without requiring extensive knowledge and research into the machine in question. See https://github.com/geodynamics/aspect/wiki for some examples.

Our work has enabled many scientists to be able to evaluate and use XSEDE resources who otherwise would not have been able to do so.

To that end we plan to:

- Port and tune the software to new XSEDE machines/compilers.
- Continuously maintain the whole software stack and report and fix bugs that appear.
- Perform performance tests and tuning of the codes.
- Create installation and usage guides.

Computational Experiments and Resource Requirements

Numerical Approaches

Calypso. Calypso is a code for magnetohydrodynamics (MHD) simulations in a rotating spherical shell to solve the geodynamo processes. It uses a spherical harmonic transform method and a finite difference method in the horizontal and vertical discretizations, respectively. Linear terms (e.g. diffusion, buoyancy, and 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. We were awarded 110,000 node*hours for TACC Frontera Pathway allocation from April 2021 to August 2022 for both ASPECT and Calypso. We optimized the spherical harmonic transform and visualization modules for the Frontera system.

Rayleigh. Rayleigh is an open-source community dynamo code developed through CIG (NSF support). It solves the fully nonlinear MHD equations of motion for a compressible fluid in a rotating spherical shell under the anelastic approximation and 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. This code has been performance tested extensively on XSEDE's Stampede2, NASA's SGI Pleiades system and Argonne's Blue Gene/Q system, Mira.

ASPECT. ASPECT is a CIG-developed code designed to solve the equations that describe thermally driven convection in the Earth's mantle and tectonic deformation in the Earth's lithosphere. It allows for both 2D and 3D models of arbitrary geometry (generally focused on segments or whole mantle models), adaptive mesh refinement in locations of scientific interest, easy modification of material, gravity, rheology and temperature models, and tracers to model geochemistry and material transport. Recent work has started investigating the effectiveness of GPU or other coprocessors in ASPECT simulations. Further details are available in [Kronbichler et al., 2012; Heister et al., 2017]. Based on improvements we made on XSEDE resources we were awarded 110,000 node*hours for TACC Frontera Pathway allocation from April 2021 to August 2022 for both ASPECT and Calypso. For ASPECT, we optimized the ASPECT's GMG solver and Calypso's visualization modules for the new Frontera system.

Leeds Spherical Dynamo (LSD) LSD solves the Boussinesq dynamo equations by representing velocity, and magnetic field, as poloidal and toroidal scalars [Willis et al., 2007]. It is pseudo-spectral; the variations in a sphere are expanded in spherical harmonics, and radial variations are discretized by finite differences with a nonequidistant grid using Chebyshev zeros as grid points. The nonlinear terms are evaluated by the transform method. Time stepping is using a predictor-corrector method and the time step is controlled using a CFL condition and error information from the corrector step. The LSD code is parallelized using MPI in both radial and in θ (meridional) directions [Matsui et al., 2016]. LSD will be used in a one-way coupled model together with ASPECT. ASPECT output will be used as (time-independent) boundary condition for LSD.

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 complex nonplanar fault geometry), implementation of a variety of finite-element types, 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, University of CNRS (France), and KAUST, 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]. The SEM 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] as well as on clusters of GPU accelerating graphics cards [Komatitsch, 2010].

Table 1: List of Websites

Code	Website
Calypso	`https://geodynamics.org/resources/calypso`
ASPECT	`https://geodynamics.org/resources/aspect`
PyLith	`https://geodynamics.org/resources/pylith`
SPECFEM3D_GLOBE	`https://geodynamics.org/resources/specfem3d`

Resource Requirements

CIG researchers used a significant portion of the past period's allocation for studies of the geodynamo, lithospheric deformation, and mantle convection.

CIG plans the following use of its proposed XSEDE resources during the period of October 1st, 2022 to September 30, 2023 in support of (1) scalability testing and code validation, (2) development of new numerical methods for better code performance, (3) science production models as described below, and (4) workshops training sessions and nurturing new geophysics users on XSEDE resources using Calypso, ASPECT, PyLith and

SPECFEM3D_GLOBE. New users will require additional SUs to conduct their research. CIG expects to support their feasibility testing and spin-up, which will enable researchers to apply for their own allocations. More details are provided below.

PyLith, Calypso, and SPECFEM3D_GLOBE development. The short-term tectonic finite element code PyLith is continuing development and scaling work lead by Brad Aagard (USGS), Matt Knepley (University of Buffalo), and Charles Williams (GNS). The geodynamo code Calypso is also continuing development and scaling work by Dr. Hiroaki Matsui (UC Davis). SPECFEM3D_GLOBE is principally developed by Prof. Jeroen Tromp (Princeton) and Dr. Daniel Peter(KAUST). The allocation will be used to establish the scaling performance and efficiency of each code, add functionality, and improve the support for many-core architectures. To perform simulations to ensure the validity of the codes and check their scalability and performance, we anticipate requiring up to 256 nodes for brief periods (10-20 hours) and estimate a total requirement of 5,000 SUs for each code (15,000 SUs in total) on Stampede2 for this development.

Calypso development: Data analysis and visualization of large scale dynamo simulation by Rayleigh. We performed large scale dynamo simulation using the Rayleigh code for the Earth and Jupiter under the support of DOE INCITE project with approximately 10 billion grids, $((N_r, N_\theta, N_\phi) = (512, 1536, 3072))$ for each snapshot. Parallel computing is also required for its visualization and data analysis. Dr. Hiroaki Matsui will develop a migration module to connect Rayleigh code and Calypso's visualization modules for visualization and data analysis of these data on Stampede2. We will request 200(times) x 0.2(hours) x 512(nodes) = 20,480 SUs for this study.

Calypso: Investigation of requirements of past geodynamo processes with small solid inner core Dr. Hiroaki Matsui is investigating required parameter ranges to sustain the intense magnetic field in the past Earth when the solid inner core is smaller than that in the present. For this study, we need to perform approximately 10 cases with changing 3 different aspect ratio of the inner core. Each case requires 32 nodes of Stampede2's Skylake node for 24 hours. Consequently, We request 23,040 SUs for the present research.

SPECFEM3D: Full waveform inversion workflows. We will demonstrate containerizing 2D and 3D full-waveform inversion (FWI) workflows for training workshops and classroom teaching, as well as production multi-scale FWI runs on personal laptops, clusters, and supercomputers. FWI workflows are well-defined and typically consist of three stages: 1) pre-processing stage (i.e., filtering time series, selecting measurement windows, measurements, and computation of adjoint sources, etc.), 2) numerical forward and adjoint simulations, 3) post-processing stage (i.e., summation of event kernels to compute gradients, smoothing and pre-conditioning gradients, line search, model updates, etc.).

The initial setup and maintenance of the entire FWI workflow, specifically the python environments for the pre-processing stage, require a non-trivial amount of work on different computer systems. Our goal is to use containers to reduce the complexity and increase the reproducibility of results. We have containerized a 2D FWI workflow and also demonstrated a sample case in 3D that can be run in both on workstations and in HPC environments

Containers are initially created and run on the workstations using Docker and on the TACC's Frontera system, where Singularity is used as the container platform. As part of this allocation, we will also demonstrate the containerized FWI workflows on XSEDE Stampede2 to further tune the containers to muliple HPC platforms. We request 5,000 SUs on Stampede2 for this effort.

Mantle Convection and Lithosphere Dynamics Modeling with ASPECT. Below, we outline distinct projects using ASPECT that will require XSEDE resources.

New matrix-free solver development. Dr. Marc Fehling and Prof. Wolfgang Bangerth developed general algorithms for parallel hp-adaptive finite element methods that work for continuous and discontinuous elements. To demonstrate the scalability of our implementation, we ran weak and strong scaling experiments on the expanse supercomputer which involved the solution of a Laplace problem in 2D with billions of unknowns and a Stokes problem in 3D. We will continue to develop new solvers and preconditioners for parallel hp-FEM to improve the computational efficiency with matrix-free and geometric multigrid methods. With an extension of XSEDE computing time, we will re-run our previous experiments including the new optimizations on the expanse supercomputer to measure the practical improvement. For this, we request 200,000 core-hours on Expanse.

Prof. Timo Heister will continue his work on improving I/O performance for checkpointing and graphical output in ASPECT to better support the large model outputs generated by the simulations mentioned below. We request 5,000 SUs on Stampede2 for benchmarking.

Tracking of material properties and stresses. Profs. John Naliboff, Bob Myhill, Julianne Dannberg, Rene Gassmoeller and Dr. Anne Glerum will continue their collaborative work on the implementation of new a viscoelastic-plastic (VEP) formulation in ASPECT. The majority of this work and testing the relative efficiency and accuracy of using particles versus compositional fields to track lithologies, viscoelastic stress tensor components, and additional properties in high-resolution 2D and 3D simulations. While related development work for a range of analytical benchmarks has occurred in 2D, the trade-offs between efficiency and accuracy are likely to be different in high-resolution and more complex 2D/3D simulations. Over the past reporting period, the majority of work on this topic has been dedicated to the underlying numerical VEP formulation and assessing the tradeoffs between particles and compositional fields in complex (e.g., realistic) 2D simulations of continental extension and inversion. The work proposed here will focus on extending these tests to complex 3D simulations. Here, we conservatively request 5,000 SUs on Stampede2 for this proposed work.

Global Models with Prescribed Plate Boundaries Dr. Arushi Saxena, Prof. Juliane Dannberg, and Prof. Rene Gassmoeller, will continue their work on global mantle convection models with prescribed plate boundaries. Examples of this type of model have been widely requested by the ASPECT user community, and this work will provide the basis for multiple groups to do related research projects. Here, the goal is to work out various techniques and issues for diverse data assimilation, solver stability, adaptive mesh refinement techniques, and highly non-linear material properties. In our models we have reached mid-size instantaneous global mantle convection models on Expanse that have about 200 Million DoFs. Each model took around 30,000 CPU hours to finish. For the next phase of the project we aim to improve our models in two ways. First, we plan to improve the representation of subducted slabs by importating data from the published global database SLAB2.0 [Hayes et al., 2018]. Second, we plan to improve our rheology by computing a variable, location dependent grain-size, which will influence our rheological model parameters. While we will use the same resolution as before, based on initial tests we expect the additional complexity to double our cost, about 60,000 core hours per model. We plan to run 25 models covering our previously explored parameter space, totaling 1.5 million core hours (SUs) on Expanse, or an equivalent amount on Stampede2.

New Rheologic Formulations for Tectonics. Profs. John Naliboff, Bob Myhill, Cedric Thieulot, Timo Heister, and additional collaborators will continue their work on developing and testing new rheological formulations that will be of significant use to the ASPECT mantle convection and lithospheric dynamics community. This work will include further testing of currently implemented plasticity stabilization in combination with adaptive mesh refinement, composite rheological formulations, and compressibility. While the majority of development work for these topics will be completed using small simulations on local resources, the resources requested here will be used for limited proof-of-concept production models that will be highlighted in the ASPECT repository. Although the size and required run times of lithospheric dynamics simulations vary dramatically, we have found a reasonable average for high-resolution simulations on Stampede2 is 15 nodes used over 48 hours (720 SUs). Here, we request resources to conduct the equivalent of 10 simulations (7,200 SUs).

New methods for two-phase fluid transport Profs. John Naliboff, Prof. Juliane Dannberg, and their PhD students will test recently implemented methods for comparing coupled versus uncoupled two-phase simulations in ASPECT, with a focus on applications to volatile transport in the convecting mantle and lithosphere. These comparisons will involve benchmarking in 2D and 3D, in combination with extensive testing of solver parameters. Following the completion of benchmarking against analytical or previously published solutions, a limited number of high-resolution 2D and 3D complex simulations will be conducted as proof-of-concept models. Development and testing of these features have been widely requested by the community, and the results will be made available and documented within the ASPECT repository. At this stage, the required run times for the range of planned simulations are difficulty to assess, and as such we request an amount (5,000 SUs) that should allow extensive testing in both 2D and 3D.

Integration and extension of CPO calculations. Dr. Menno Fraters will extend his work on integrating crystal/lattice preferred orientation (CPO/LPO) in ASPECT. This work will include more testing of the current implementation of the modified D-Rex CPO model in ASPECT, test better integration with the Geodynamic World builder which can create CPO initial conditions, and explore the addition of water models and CPO mobility models into the CPO implementation in ASPECT. We expect to use 6 nodes for 48 hours per model (288 SUs). Here we request resources on Stampede2 to conduct the equivalent of 40 simulations (12,000 SUs).

Heterogeneous boundary conditions for geodynamo simulations. Dr. Daniele Thallner will build on previously presented homogeneous numerical geodynamo models using LSD and run additional models that apply heterogeneous boundary conditions from mantle convection models in ASPECT. On average, these LSD models were using 420 SUs per model to be run for a representative amount of time (˜10 magnetic diffusion times) on a different supercomputer. Based on our scaling tests we expect an equivalent usage of computing time on Stampede2. For a systematic exploration of the effects of varying heat flux boundary conditions, we request resources to conduct the equivalent of 40 simulations, or 10 simulations with 4 different boundary conditions each (16,800 SUs).

Table 2: Summary of requested SUs for Stampede2

Software	Purpose	Requested SUs
Calypso	Development and Optimization	5,000
	Large data analysis and visualization	20,480
	Past Earth's dynamo model	23,040
SPECFEM3D_GLOBE	Development and Optimization	5,000
	FWI workflow	5,000
PyLith	Development and Optimization	5,000
ASPECT	Solver development	5,000
	Tracking Stress	5,000
	Lithospheric Dynamics	7,200
	Two-Phase Fluid transport	5,000
	CPO Calculations	12,000
LSD	Heterogeneous BCs for Geodynamo	16,800
	Total for Stampede2	114,520

Table 3: Summary of requested SUs for Expanse

Software	Purpose	Requested SUs
ASPECT	Solver development	200,000
	Global Plate Boundaries	1,500,000
	Total for Expanse	1,700,000

Summary In total, a yearly allocation of 109,520 SUs on Stampede2, 1,7000,000 SUs on Expanse will enable CIG to continue offering feature development, support, and training to users of these commonly used geophysics codes. This will also allow extensive studies of code accuracy, performance and validation using high-resolution simulations. We also request 10,000 GB on the data storage system Ranch.