Geoinformatics: Facility Support:

Computational Infrastructure for Geodynamics

Principal Investigator:
Louise Kellogg, *UC Davis*

Proposal Writing Committee:

Brad Aagaard, *USGS*
Wolfgang Bangerth, *TAMU*
Thorsten Becker, *USC*
Bruce Buffett, *UC Berkeley*
Michael Gurnis, *Caltech*
Peter van Keken, *Michigan*
Alan Levander, *Rice*
Marc Spiegelman, *Columbia*

Submitted to NSF's Geoinformatics Program
July 8, 2009

Project Summary

We request funding for the Computational Infrastructure for Geodynamics (CIG), whose mission is to provide the software cyberinfrastructure that facilitates transformative research in geodynamics by a broad community of scientists ranging from model developers to end-users. The computational infrastructure provided by CIG is driven by community needs and defined by community participation, and includes: (a) reusable, well-documented open-source geodynamics codes and tools; (b) support for a virtual community that facilitates intellectual exchange and partnerships between computational scientists and earth scientists; and (c) specialized training and workshops for the earth science community. Software will be developed and supported for geodynamics, seismology, and associated disciplines to address challenging scientific problems in the deep earth, lithosphere, and crust, ranging in scale from global mantle structure and circulation to magma- and fault-scale processes.

CIG directly employs a core software development team of dedicated software engineers whose work is guided by scientific objectives formulated by the geodynamics community. Specialized development is also done through subcontracts to institutions with specific expertise. Development directions are determined through community input via Working Groups, a Science Steering Committee, and an Executive Committee whose collective charge is to identify and balance common needs across disciplines. CIG will be housed at the University of California, Davis, for the benefit of the national and international earth science community. The move of CIG from its current location at Caltech to a new home at UC Davis provides an opportunity to add additional expertise in areas such as scientific visualization, tectonics, and remote collaboration technologies.

Intellectual Merit: The first five years of CIG have facilitated new scientific insights by providing well-engineered codes and techniques that have garnered wide use in the geophysical community. CIG has also fostered an interdisciplinary community of geophysicists and computational scientists who are working together towards new methods for computational geophysics. As a result of these efforts, new scientific questions have opened up and new avenues for advancing computational geophysics are emerging. We propose to continue support and development of state-of-the-art codes and tools, aiming especially at methods for addressing multi-scale and multi-physics problems for vastly improved modeling in deep earth dynamics, lithospheric dynamics, magma migration, earthquake dynamics and crustal dynamics, seismology, and related fields.

Broader Impacts: CIG has had and will have a broad impact upon our science through the creation of programming support and software infrastructure that will outlast the research efforts of any individual researcher and increase the rate of scientific discovery by providing the means to routinely and rapidly solve formerly intractable problems. CIG will not only provide the means of linking geodynamics with computational science and geoinformatics, but it will foster communication between different subdisciplines. CIG will impact geoscience through multidisciplinary and interdisciplinary conferences, workshops, and collaborations, and by educating a new generation of Earth scientists in state-of-the-art methods of computational geodynamics. As a consequence, it will also diffuse knowledge about modern software development methodologies in the community and thereby put the development of geodynamics codes, within but also outside of CIG, on a more modern and sustainable path. Equally importantly, the need to provide scalable and robust solvers and algorithms will continue to drive maturation of such methods in the computational sciences and engineering.

1 Introduction

Scientific computation has played an essential role in solid Earth geosciences since the plate-tectonics revolution. Quantitative numerical models provide a critical link between physical and chemical observations at Earth's surface and dynamic processes in the interior, including convection in Earth's mantle and core, plate tectonics, and the dynamics of plate boundaries with their attendant mountain building, volcanism and earthquakes. With increases in computational power and sophistication of algorithms, solid Earth geoscience has seen a proliferation of powerful and predictive models for all of these problems. However, as in many fields, codes developed for research reflect the specific interests of the individual developer/investigator and are often not intended for general use, with associated documentation or technical support. More critically, the research codes are usually not designed to be scalable, interoperable, or have reusable components, making it extremely difficult to extract or combine the best features of these models to leverage algorithmic advances or explore the fundamental interactions between the different Earth systems. Such coupled, interdisciplinary problems are clearly becoming the focus of research in solid Earth geodynamics. Examples include the influence of mantle convection on the geodynamo, the interaction of the lithosphere, plate boundaries and the deeper mantle (e.g., Figure 1), the role of fluids in lithospheric deformation, and evolution of the planet. Examples of such problems and their computational requirements are discussed in detail in Section 3. To improve our understanding of these problems, and provide better integration with observations, requires a new level of coordination in software development and design.

In response to the existing computational environment, the Computational Infrastructure for Geodynamics (CIG) was established in 2004 with funding from NSF as a community-driven partnership between Solid-Earth Science and Computational Science. The role of CIG is to facilitate scientific discovery by providing computational infrastructure to the geoscience community, with development directions guided by community-defined transformative research questions and a software engineering philosophy that maintains the flexibility needed for scientists to pursue individual avenues of scientific enquiry. Specifically, CIG provides advanced computational tools, along with the associated training and education, to a broad community of users and developers to enable them to more effectively explore and understand the dynamics of our planet. From its inception, the long-term vision of CIG has been to develop interoperable, free, open-source computational tools that support the development of models of coupled Earth systems by individual research groups. In 2010, CIG will move to the University of California, Davis, under the leadership of Louise Kellogg. Director-Designate Kellogg has been involved with CIG since its beginning. The move provides an opportunity to add additional expertise in both the computational sciences and the geosciences. We intend to continue CIG's effort to help engineer a new generation of numerical modeling software and associated computational tools, training, community building, and technical support that enable researchers to meet the most pressing scientific challenges in geodynamics.

The long-range vision of interoperable, open-source tools is driven by the specific computational and scientific needs of a wide range of geodynamics subdisciplines, including computational seismology, mantle convection, magma dynamics, short-term and long-term lithospheric deformation and the geodynamo. CIG has supported these disciplines with a suite of codes and a broad range of scientific, training and educational workshops. Specific accomplishments are detailed in Section 2 and include nine principal codes that are documented, supported, and developed using modern software engineering methods (e.g., version control, regression testing, and support for multiple platforms). CIG has also played a role in the development of several experimental codes for multiphysics and multiscale problems.

The suite of CIG codes ranges from established geodynamics codes (e.g. [1, 2]) to completely new developments built on advanced computational libraries (e.g., [3, 4]). Both are necessary to facilitate benchmarking of each capability of the codes, aand to provide users the flexibility to define their own research directions. An initial emphasis was given to specific sub-disciplines, such as Mantle Convection or Seismic Wave Propagation, because these projects were well-defined and afforded critical experience in developing and supporting complex codes for a diverse community. This work, however, emphasized that to make further progress and enable exploration of problems at the intersections of the disciplines, requires a more flexible code infrastructure that exploits common components and overlapping interests of the various communities. CIG uses a community-based approach to guide development of models and and to promote integration of observations. Forward models, including geodynamics codes developed and supported by CIG, are increasingly used as a tool to retrieve information and gain insight from data. CIG codes and activities have been developed to aid and complement a wide range of other solid Earth community activities, including IRIS [5], EarthScope [6], MARGINS [7], SCEC [8], COMPRES [9], and CIDER [10].

In its first five years, CIG has facilitated new scientific insights in the geosciences by providing well-engineered codes and techniques for use by researchers, and has fostered an interdisciplinary community of geophysicists and computational scientists who are working together towards new methods for computational geophysics. At the same time, major initiatives in the solid Earth sciences such as EarthScope, MARGINS, and COMPRES, are providing new data and insights into Earth's interior. These wide ranging activities are opening up new scientific questions, at the same time that new opportunities are emerging for advancing computational geophysics. These new scientific frontiers require well-engineered, tested codes that are readily usable by the scientific community, and require a geoscience community that is prepared to use state-of-the-art computational methods. The objective of CIG is to facilitate scientific progress through development and dissemination of such codes and techniques beyond the level that is usually possible through single investigator grants or small collaborations.

In the remainder of this proposal, we review the current state and accomplishments of CIG (Section 2) and emphasize the outstanding scientific problems in geodynamics and their attendant computational challenges (Section 3). We then review critical advances in computational science that enable progress on attacking these problems (Section 4). The final sections describe a concrete set of plans and a revised CIG management structure that builds on our current experience to deliver the next generation of computational tools for Solid Earth Dynamics.

2 CIG accomplishments

This section first summarizes the principal activities of CIG to date before we discuss outstanding scientific and computational problems that CIG hopes to help facilitate in the future. Section 5 will emphasize the important lessons learned from our initial experience that we feel are critical for future progress.

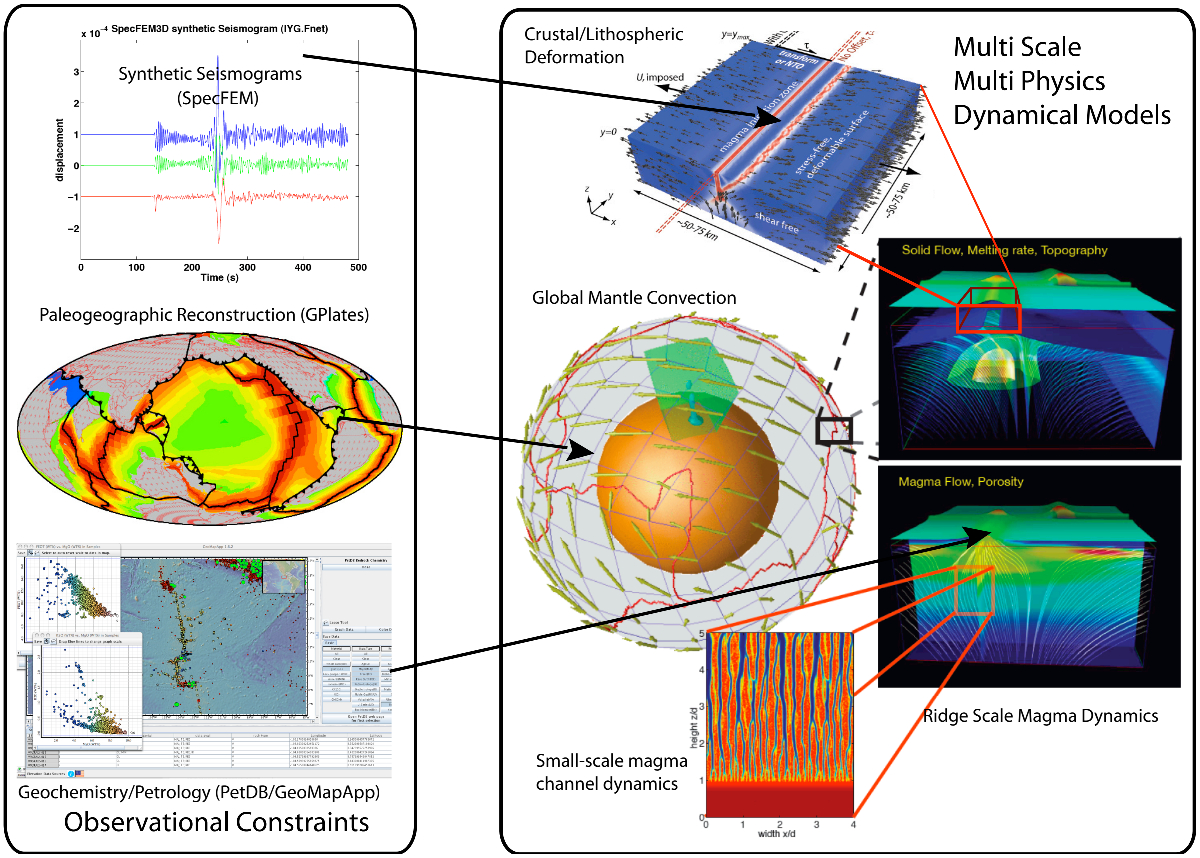

Figure 1: Examples of computations and data available to help us understand solid earth dynamics at a range of scales. In this case, the dynamics of plate boundary processes and their interaction with global mantle flow. RIGHT PANEL: Example output of computational codes for global mantle convection (*CitcomS*), mid-ocean ridge flow with melting and melt transport [11], crustal scale magma injection and faulting (*Gale*), and small scale reactive melt channel formation [12]. Each model was designed to consider a particular scale or set of processes. The challenge is how to permit users to combine these models, as needed, to explore the dynamics of the coupled interactions and use them to make inferences from geophysical and geochemical data. LEFT PANEL: Example data used to test and drive models, including seismic waveforms (*SPECFEM*), plate reconstructions (*GPlates*), and geochemistry (*PetDB*).

CIG can claim many accomplishments over the last four and a half years. We have accelerated development of new, well-tested and documented software, enabled innovative science, and trained members of the community including early-career scientists, graduate students, and postdocs. In turn, these changes have strengthened the Earth sciences as a whole. Through the application of modern software engineering techniques, concentrated development, and much closer links to Computational Science, CIG has helped transform the way computational software is developed and maintained by the U.S. geodynamics community. We have completed a detailed self-evaluation of CIG activities since its inception in September of 2004. The documentation collected for this self-evaluation are contained in a report EXPANDING COMPUTATIONAL INFRASTRUCTURE: THE FIRST FIVE YEARS OF CIG, which we include as a Supplementary Document. Here we briefly

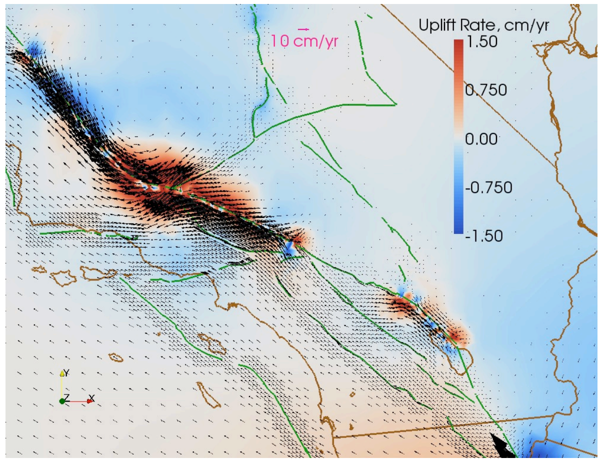

Figure 2: Three-dimensional viscoelastic model of southern California deformation over 15k years from earthquakes on 58 faults, solved with *PyLith* from work by C. Williams (formerly at RPI). Vectors and shading show horizontal and vertical motion relative to the central United States for 10 years spanning the magnitude 7.9 1857 earthquake on the San Andreas fault. Fault geometry from the Southern California Earthquake Center (SCEC) Community Fault Model and elastic physical properties from the SCEC Community Velocity Model.

summarize the principal results.

CIG has developed, released, enhanced, and maintained nine principle software packages, including four major codes described in some detail below, for the solution of mantle convection, seismic wave propagation, and tectonic problems. *CitcomS* (for mantle convection) and *SPECFEM* (for wave propagation) provide a powerful open-source community infrastructure for the solution of both regional and global problems with the most recent releases allowing the seamless transfer of data between these two codes. All of these capabilities are available through a Science Gateway on the NSF TeraGrid. Although both codes were inherited from prior efforts, they have become community codes through the efforts of CIG. Prior to the establishment of CIG, *CitcomS* had many different development lines in the community with little coordination and no oversight. We now maintain a single documented version through a centralized repository. We also promoted the use of *SPECFEM*, which was not available under an open source agreement before CIG. We developed and released two new finite-element packages for the solution of long-term and seismic-cycle tectonic problems (Fig. 2), *Gale* and *PyLith*, respectively, that bring new, open-source functionality to the U.S. community. Software development activities are now underway that promise to dramatically increase the range of resolvable scales in geodynamic calculations through the use of Adaptive Mesh Refinement (AMR).

Through these software releases, CIG has facilitated a wide range of scientific activities, as detailed in the 66 brief abstracts in the Supplementary Documents. Research using CIG software has, for example, focused on enabling direct links between mantle convection and seismic observations, mineral physics, Earth tides and gravity, and observations of surface vertical motions and sea level change as well as elucidating the dynamics of plate tectonics, subduction, plumes, and the interior of other planets. Our software has been used to improve images of Earth's interior through adjoint tomography (Fig. 3), directly test tomographic models with seismic waveforms, and improve and automate moment tensor inversions of earthquakes. It has also been used in some of the first three-dimensional models of the initiation and growth of faults in extensional tectonic envi-

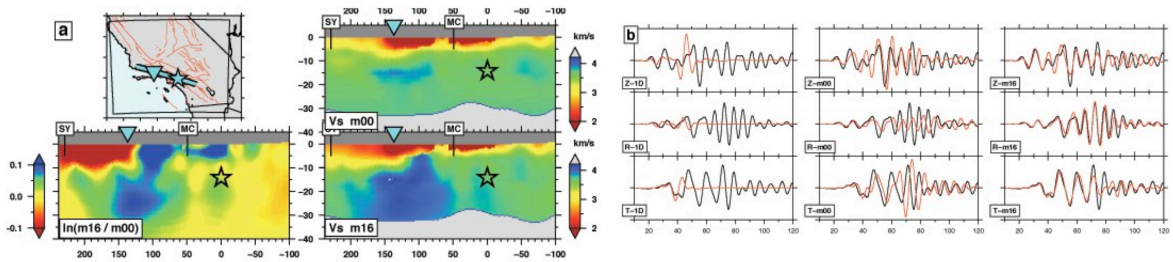

Figure 3: Improvement in tomographic images of the crust afforded by full waveform seismic tomography through forward and adjoint *SPECFEM* simulations. (a) Geometry and initial, intermediate and final shear wave anomalies after 16 full waveform inversion iterations. (b) Initial, intermediate and final predicted and actual waveforms. From *Tape et al.* [13].

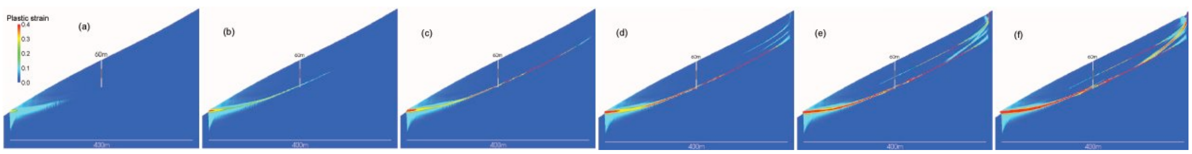

Figure 4: Dynamic model of a landslide undergoing failure with deformation localized on a dominant slip plane as a function of time. This was an unanticipated application of *SNAC* for surface-process studies (from C. Stark, *LDEO* [14]). *SNAC* was originally designed for larger-scale geodynamic problems involving elasto-viscoplastic deformation of the whole crust and lithosphere using the same software framework and components used in *Gale*.

ronments. CIG software has facilitated studies of the interplay of crustal extension and melting, hill slope failure (Fig. 4) and surface inflation associated with volcanic intrusion.

CIG applies modern software engineering to the codes it develops and supports by using good software design practices, including automated build and test procedures, development and use of benchmarks and test cases, and by providing documentation of codes. The software repository and attendant web site (`http://geodynamics.org`) are central to facilitating collaboration and sharing validated open-source software. We primarily use a single, open repository for developer and community use, a bug-tracking database to allow developers and external participants to register and comment on bugs or request new functionality, and support list-servs for each sub-discipline, and editable Web pages for community input.

CIG has facilitated much stronger, mutually beneficial links between earth sciences and computational sciences. The earth sciences have benefited computational scientists by providing the context to push the envelope of methods that had previously been developed on idealized problems. Real-world applications in geodynamics, with all their practical intricacies such as strongly inhomogenous or anisotropic coefficients and geometries, provide an important testbed for new methods. Conversely, the earth sciences have benefited from computational sciences by adopting and using well-tested and innovative software and computational practices. As examples of this interaction, we mention the development of the *PETSc* package *Sieve*, which allows much greater flexibility in the storage and manipulation of finite-element objects on a parallel computer. This package is now a core component of *PyLith* for simulations of the complete earthquake cycle. Likewise, our research with the Institute of Computational Sciences and Engineering (ICES) at the University of Texas at Austin has allowed the largest, most efficient scaling of Adaptive Mesh Refinement on a massively parallel computer [15]; global simulations of mantle convection with individual faulted plate boundaries that resolve features at 1-km scales will be within reach this

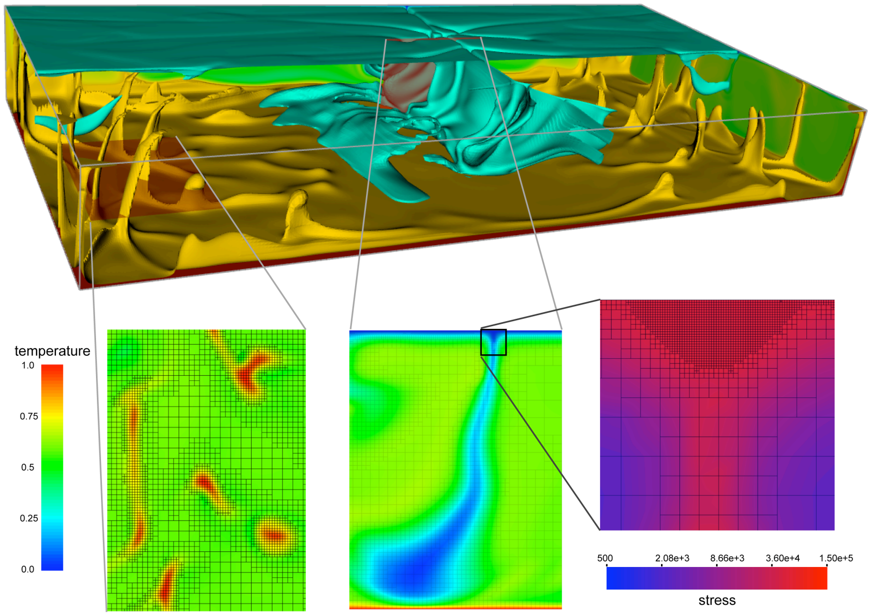

Figure 5: Mantle convection with yielding using Adaptive Mesh Refinement using the *Rhea* code developed in collaboration between the University of Texas at Austin and CIG. Top: Temperature isosurfaces. Bottom: Horizontal and vertical slices of temperature showing that grid refinement follows the temperature gradient. The finest grid has a resolution of approximately 1.5 km. From Burstedde *et al.* [15].

year (Fig. 5). Our collaboration with the developers of *deal.II* at Texas A&M University are giving us much greater flexibility as we push the envelope on multiscale models of plate boundaries with AMR. Furthermore, being part of the Center for Advanced Computer Research (CACR) at Caltech has allowed CIG to deploy and maintain our software with the most advanced tools in software engineering. The upcoming move of CIG headquarters to the University of California, Davis, will provide new opportunities for collaboration with computational scientists, applied mathematicians, and computer scientists at UC Davis, Lawrence Berkeley National Lab, and Lawrence Livermore National Lab, bringing additional expertise on numerical methods, scientific visualization, and high-performance computing.

The broader impacts of CIG have been enhanced by nearly two dozen workshops and training sessions we have organized since mid-2005 (see Supplementary Material). Some of these workshops have focused on specific technical and scientific questions, such as the 2005 workshop on Compressible Mantle Convection or the 2007 workshop on Adaptive Mesh Refinement, which are essential for accelerated progress in geodynamics. Other workshops have been decidedly interdisciplinary, actively building new bridges with other areas of computational science (such as the 2006 Austin and 2008 Santa Fe workshops). Training sessions, such as those at the EarthScope and IRIS national meetings and at the CIDER 2008 summer tutorial, have educated new users, including early career scientists, in the use of the community software. Of the more than 900 participants in CIG workshops and training, more than 400 have been been graduate students or postdoctoral scholars. We believe that these workshops have strengthened the technical sophistication of the U.S. Earth science community and helped prepare the next generation of geodynamics researchers to use state-of-the-art numerical methods. Workshops, training sessions and other community building mechanisms are essential to the success of CIG.

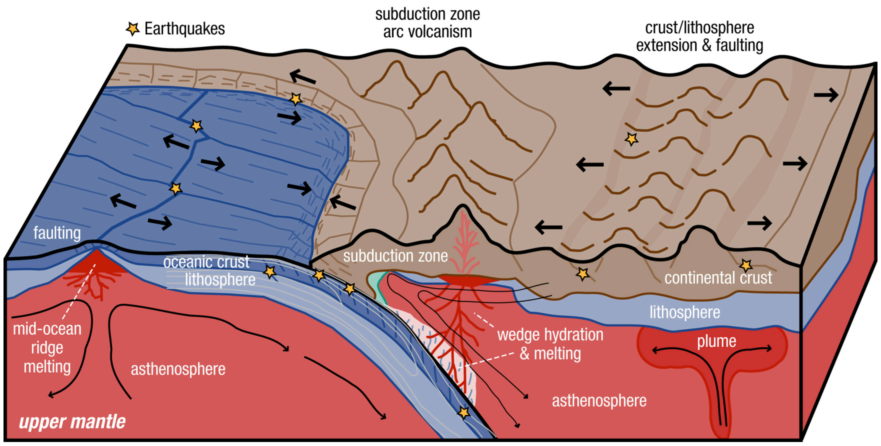

Figure 6: Sketch of the lithosphere-mantle system with key components labeled.

3 Outstanding Scientific and Computational Goals

During its initial phase, CIG has gained significant expertise in developing and supporting advanced computational tools for a diverse community of users and developers. Nevertheless, the number of important and compelling scientific problems that would clearly benefit from a continuing, coordinated effort of software design, implementation and support on behalf of the geodynamics community has not diminished. CIG's effort will balance providing the functionality available in the current codes, while developing and supporting the next generation of computational tools and codes that will enable individual researchers to more readily explore more advanced problems. This section highlights examples from across the sub-disciplines that have emerged as focus areas from a multi-year planning process. This process included input from both Earth Science and Computational Science communities through a series of business meetings, on-line discussions at `http://geodynamics.org`, and workshops, the most recent being the 2009 planning workshop, *Opportunities and Challenges in Computational Geophysics* in Pasadena, CA.

A common theme among these scientific problems is the need to understand the couplings and interactions between the fundamental subsystems. Examples include coupling between mantle convection and plates, plate boundary dynamics with magmatism and brittle deformation, fault interaction, coupling between core and mantle convection, and better integration of observations and geodynamic models. Figures 1 and 6 show some schematic examples, with Figure 6 labeling some of the major tectonic features and processes. Leveraging the common mathematical and computational components shared by many of these problems into reusable and interoperable modules will require continued collaboration among Earth Scientists and Computational Scientists. CIG is uniquely positioned to foster this dialog, and it is unlikely that such a unified design will arise through the efforts of individual investigators.

The remainder of this section provides specific examples and outstanding questions that such advanced codes could help address. Software developed by CIG will be guided by these and similar questions. We also highlight the specific computational requirements and challenges inherent in such codes.

3.1 Global Scale Dynamics

The two largest scale dynamic systems in Earth's interior are the convecting solid mantle (with plate tectonics as an important surface manifestation) and the convecting liquid core responsible for the geodynamo. These two systems operate on significantly different time scales, although they are coupled by the exchange of heat, mass and momentum across the core-mantle boundary. Fundamental questions remain about each system when taken in isolation; new challenges and opportunities emerge when these two systems are permitted to interact.

3.1.1 Deep mantle dynamics and plate tectonics

The theory of plate tectonics is a kinematic description of plate movements at Earth's surface and provides a framework to understand seismicity, volcanism, and mountain building at the boundaries between plates. However, the dynamic processes controlling plate tectonics are not understood: We do not fully grasp the distribution of forces on plates, how this distribution changes with time, and how it is linked to the gross thermal and chemical evolution of our planet. Computational models of mantle convection provide a self-consistent perspective on the dynamical system that drives plate tectonics and permit interpolation and extrapolation into time periods or regions where observational coverage is poor or unavailable. Models also provide a framework for interpreting a rich and diverse geological record. Models of mantle convection have considerable overlap with studies of lithospheric dynamics, long-term tectonics, and the dynamics of faulting and seismogenesis. Key questions identified by the scientific community that motivate the development of new computational tools include the following:

Subduction: How is subduction initiated and what controls the dynamics of subduction zones? Although computational models characterize how faulting in the seismogenic zone decouples the subducting plate from the overriding plate, they do not adequately explain initiation of a subduction zone, controls on the size of the seismogenic zone, and the transition to localized, aseismic deformation.

Origins of plate tectonics: When did plate tectonics start and what controls the secular changes in plate tectonics? It remains unclear how mantle convection is linked to the formation and breakup of super-continents, or how periods of enhanced mantle plume activity may influence plate tectonics and core dynamics.

Hot spots: What controls the formation and stability of hot spots such as Hawaii and Iceland? Given the poor seismological coverage of putative mantle plumes, theoretical tests of plume models and their alternatives are central to resolving a variety of questions, such as the viability of the "plume head and tail" model.

Stirring, melting and geochemical evolution: How does the stirring of Earth's mantle, coupled with melting, fractionation and transport at volcanic centers, control the geochemical evolution of the mantle and crust and interactions with the core? Observations of geochemical variability provide one of the principal constraints on the long-term evolution of the mantle and core. However, to infer dynamics from chemistry requires consistent computation of both large scale stirring and small scale magmatic processes.

Seismic anisotropy and rheology: What controls the formation of seismic anisotropy and how can we use observations of anisotropy to infer the flow fields in the mantle? The progressive deformation of rocks creates seismically detectable crystal textures. Therefore, if seismic observations can be linked to experimental work through models of mantle convection, anisotropy can provide fundamental constraints on the rheology and flow of Earth's interior.

3.1.2 Core dynamics

Earth is unique among the terrestrial planets in that it has an internally generated magnetic field and an active system of plate tectonics. Recent advances in our understanding of the core reinforce the view that the mantle plays an important role in sustaining the magnetic field, primarily by regulating heat flow from the core. Evidence suggests that fluctuations in mantle convection are imprinted on the dynamics of the core and the resulting magnetic field. As a consequence, geological estimates of magnetic reversal rates, magnetic intensity and the large-scale structure of the field reveal the past dynamics of both the core and the mantle. However, our ability to interpret this record is limited by shortcomings in the models we use to understand the dynamics.

Origin of fields: An overarching goal of research on core dynamics is to understand the origin of Earth's magnetic field. Persistent decay of the magnetic field requires some mechanism for continual regeneration. Simulations of convection in the core have succeeded in generating large-scale magnetic fields without constraints on the form of the flow or the field. Many of these models predict dipole fields that are comparable in strength to Earth's field, yet many basic questions remain unanswered. How much energy is required to run a dynamo? What is the typical strength of the magnetic field inside the core? How do magnetic reversals occur and what does the reversal frequency tell us about interactions with the mantle? A resolution of these questions will rely on improvements in the numerical models.

Core-mantle interactions: Long-term variations in the magnetic field over timescales of 10^7 to 10^8 years point to an influence from the mantle. Models of mantle convection that assimilate geological information could be used to drive core dynamics models with a spatially and temporally varying heat flow at the core-mantle boundary. It is also possible for the core to exert an influence on the mantle, possibly through the strong latitudinal dependence in the convective heat flow in the core. Developing the capability to couple models of mantle and core dynamics opens new opportunities to synthesize diverse sets of observations. For example, predictions of field variations from coupled models could be tested against the paleo-magnetic record to refine the dynamics of both the mantle and the core.

3.2 Plate Boundary Scale Dynamics

Mantle convection provides the large-scale context for understanding solid-earth dynamics. In particular, the broad scale of convection is controlled by the presence of strong plates with extremely weak, narrow boundaries. Understanding the dynamics at plate boundaries and its interaction with large-scale mantle flow, requires computational tools that can introduce fluids (magma and/or volatiles) into solids while allowing a wide range of mechanical behaviors (ductile or brittle).

3.2.1 Lithospheric-scale tectonics

Faults along plate boundaries form a main component of the broader dynamics of lithospheric systems. Plates are properly viewed as a component of mantle convection. However, in many cases, the long-term (millions of years) evolution of the lithosphere can usefully be described with a shallow mantle and crust focus. Key outstanding science questions include:

Fluids: What role do magmatic and hydrous fluids play in controlling the evolution of deforming regions of the lithosphere? Recent laboratory studies demonstrate that the presence of melt affects the strength of the lithosphere and hence the style of deformation. Consequently, models should allow consistent coupling of solid deformation with melt transport.

Rheology: What is the rheological structure of the lithosphere? Tectonic deformation, as observed at Earth's surface, reflects plate motions and motion of the less viscous asthenosphere "filtered" through the rheological response of the lithosphere. The strength of the upper and lower crust relative to the upper mantle is thought to control the style of deformation and seismicity at both inter- and intra-plate settings, but this expectation needs to be tested more thoroughly within the context of large-scale models.

Memory: How does structural "memory" associated with previous faulting affect the initiation of new plate boundaries and the fracture of continental and oceanic plates? This is relevant for modeling subduction initiation and the thermal evolution of Earth throughout supercontinental (Wilson) cycles.

Interconnections: To what degree are climate, tectonics and surface processes coupled? Such interactions can influence the evolution of compressional orogens such as the Andes or the Himalayas, where topography feeds into regional climate dynamics. Other examples include sedimentary basins, where the interaction of faulting and surface processes can affect drainage patterns. Simulations can quantify the conditions under which strong coupling between the surface and subsurface exists, and the extent to which regional climate change feeds into the style of deformation in an orogen.

3.2.2 Fault system dynamics

Understanding fault dynamics and the time-dependent inelastic response of the crust and mantle throughout the seismic cycle has immediate societal relevance through seismic hazard assessment. Models are an essential component for forecasting long-term, and perhaps eventually, short-term earthquake probabilities. Key outstanding science questions include:

Rheology: What physical processes control fault strength throughout the seismic cycle? Numerous mechanisms have been proposed to explain certain features, but none of the current numerical models provide satisfactory explanations for the apparent disparity between static and dynamic strength, as well as extreme localization of dynamic slip.

Fluids: What role do fluids play in faulting? Fluids are thought to influence triggering of earthquakes by dynamic stresses during the passage of seismic waves, episodes of creep transients in subduction zones, and rapid weakening in dynamic slip. Simulations are needed to determine the precise mechanisms and explain when they become important relative to other processes, such as fracture on secondary faults and heat generation.

Memory: Is the spatio-temporal variation in earthquakes and crustal stress due to existing geometrical heterogeneity of faults systems, or are they controlled by dynamic processes during rupture? Models can be used to quantify the relative importance of the complex geometry of active fault systems versus the local state of stress and dynamic processes on individual faults in assessing seismic hazard.

Interconnections: In what ways are earthquakes, the short time-scale manifestations of strain localization and damage within a fault system, related to the larger scale formation and evolution of plate boundaries? Resolving these issues will require simulations with a more comprehensive description of plate boundary mechanics involving elasticity, fluid flow, magma dynamics, and in some regions ice dynamics. Brittle and viscoelastic crustal deformation are also influenced by hydrological processes (e.g., increased surface loads associated with filling man-made reservoirs) and glacial isostatic adjustment (e.g., reduced surface loads associated with glacial retreat).

3.3 Integration with Observations

The ultimate goal of simulation is to constrain models by comparing computations with observations. Progress on the key science questions in geodynamics can therefore only be achieved by combining geological, geochemical and geophysical observations with quantitative model predictions. Recently, several efforts at the national and international level have constructed new geophysical observatories that provide a dramatic increase in the number and quality of observations. For example, EarthScope now provides data streams that reveal complex transient and episodic deformation of the crust via continuous recordings from seismometers, borehole strain meters and GPS receivers throughout the western United States. The continuous recordings from seismometers also capture wave propagation through Earth's upper mantle in a variety of tectonic regimes. A key role of CIG is to help develop the computational tools that can use these, as well as other geological, geochemical and geophysical observations, to infer information about Earth's structure.

3.3.1 Seismic imaging of Earth's interior

Inversion of seismic waveforms has become an increasingly routine technique that provides an order of magnitude increase in resolution of global seismic velocity, anisotropy, and attenuation structure over more conventional approaches based on travel times. The range of applications includes imaging Earth's structure at global, regional and local scales, earthquake source region characterization, hydrocarbon exploration, monitoring CO_2 sequestration, subsurface waste and groundwater management. In addition, seismic imaging is important because its results can be cross-correlated with predictions of other codes to validate models or provide initial conditions. Key questions that can be answered using advanced imaging techniques include:

Regional seismic investigations: What is the detailed structure of the lithosphere-asthenosphere boundary, the upper mantle beneath subduction zones and rifts, secondary convection cells in the upper mantle, and the zones surrounding descending slabs and ascending plumes? Combining new imaging techniques with recently available seismic data streams will provide a substantial improvement in the resolution of these structures.

Crust and mantle: How does the structure of rift and arc magmatic system evolve from their points of origin in the mantle to their points of exit at the surface? Imaging these structures requires inversion of generally non-overlapping seismic frequency bands via coupling of passive seismic observations with higher resolution, regional crustal controlled-source observations.

Joint inversion: An ultimate goal of geodynamics is an internally consistent Earth model. Constructing such a model will necessarily involve joint inversion of seismic data with other geophysical (and geochemical) data through geodynamics models that capture both global and plate-boundary scale dynamics.

3.3.2 Inverse problems

One of the central goals of geodynamic modeling and CIG is to enable the use of available data to infer Earth's properties through computational models. Seismic imaging is a special case of these inverse problems in that it is possible to analytically differentiate outputs (synthetic seismograms) with respect to inputs (seismic velocities and source locations). This is not commonly the case for more indirect observations from GPS, geochemistry, or paleo-magnetism, which provide information about plate velocities, magma composition, or magnetic field reversal frequencies. In such cases Bayesian inversion offers an alternative approach at least for cases with a modest number of parameters to be inferred. In this case codes only need to compute predictions of observable quantities for arbitrary input models. However, this capability requires that programs (e.g., forward models) are written in a way so that they can be run as subroutines of optimization algorithms.

3.4 Common Computational Requirements and Challenges

While the scientific problems addressed by CIG have a very broad scope, many of them share common computational requirements and challenges. With good software engineering and a concerted collaboration between Earth Scientists and Computational Scientists, it will be possible to design geodynamics codes comprised of reusable and interoperable components that can be applied to a wide range of problems. It should be stressed that we are *not* proposing to develop a monolithic code. Rather, CIG recognizes that advances in computational methods and libraries facilitate developing a suite of more flexible codes that share common components, providing the advantage of addressing a much wider range of problems with the same computational tools. However, some problems are not amenable to this approach, and CIG will continue to explore the most fruitful code design that is driven by the scientific needs of the geodynamics community through close and ongoing collaboration between earth scientists and computational scientists and engineers.

Mantle convection and lithospheric deformation provide example problems where coordinated design and development could pay off. Both problems require robust, scalable Stokes solvers that can handle an overlapping range of spatially variable and often nonlinear solid rheologies. Additionally, they need multi-scale methods to resolve boundary layers, plate boundaries and emergent shear zones resulting from plasticity. The mantle convection and lithospheric deformation communities are also both interested in the role of fluids and magma in modifying solid-state deformation. Because most commonly used formulations for magma migration are extensions of Stokes flow, it makes sense to extend the next generation of mantle and lithospheric codes to include magma dynamics. Thus, we envision building both mantle convection and lithospheric deformations models on a common computational toolkit.

Fault and earthquake dynamics also have considerable need for multi-scale, multi-physics methods. However, in contrast to mantle convection and lithospheric deformation, these problems are dominated by brittle and anelastic deformation with slip across complex, intersecting, nonplanar fault surfaces. Unstructured discretization schemes (e.g., unstructured tetrahedral meshes) are often the most practical approach for these problems, and CIG has developed efficient methods supporting parallel, unstructured finite-element meshes (*Sieve*). In developing the next generation of codes for earthquake dynamics, it is unclear whether alternative implementations of slip on nonplanar fault surfaces will enable application of the multi-scale and multi-physics methods developed for mantle convection and lithospheric deformation, or whether these problems will require approaches tailored to unstructured discretizations.

Likewise, while seismic wave propagation is an explicit and linear forward problem, it also benefits from efficient, multi-scale mesh generation. Analysis of wave propagation for Earth models produced by geodynamic models requires stable projection schemes for mapping physical and seismic properties to different finite element spaces. Furthermore, seismic imaging has significantly advanced in recent years with the use of adjoint methods in waveform tomography. Such methods may be applicable to more general geodynamic inverse problems.

Geodynamo models present particularly difficult multi-scale computational challenges arising from the vast range of spatial and temporal scales that must be resolved for realistic values of viscosity and other physical properties. Improved discretization techniques are needed to alleviate the obstacles associated with spectral techniques, which impose tight constraints on the structure of the grid and are not well suited to massively parallel computations. Geodynamo models also require stable-time stepping schemes for integrating over time scales spanning several orders of magnitude. Experience gained in developing multi-scale methods for other geodynamic codes may open new avenues for applying multi-scale techniques to geodynamo models.

4 Emerging Computational Methodologies

The scientific and computational goals outlined in the previous section are ambitious; however, we believe important and continuing developments in computation put these goals within reach. Through interaction between geophysicists and computational scientists, CIG has already successfully applied current computational science methodologies and software practices to geodynamics, with codes using state-of-the-art numerical algorithms as well as more established methods. In particular, all newly developed software is built on widely-used, well-tested and extensively documented software packages. All of CIG's codes undergo extensive testing to guarantee and maintain software quality.

CIG will continue to build applications based on well-established methodologies and software packages, although some of these methods may not have previously been applied to geodynamics problems. We will ensure that our codes are developed using or permitting future integration with the following modern and emerging methodologies:

Adaptive mesh refinement (AMR). By refining meshes only where necessary to resolve features of the solution, AMR has been shown in a variety of disciplines to enable solution of problems that are intractable with uniformly refined or *a priori* generated meshes. Today, AMR support is available through a number of packages such as *deal.II* (developed at Texas A&M and the recipient of the prestigious J. H. Wilkinson Prize for Numerical Software) for which a suite of geodynamics applications is already available; *Sieve* (the foundation of CIG's *PyLith* code); or the massively parallel *Rhea* mantle convection code developed in association with CIG. CIG has consistently identified AMR as an important technology for advancing geodynamics codes, and it will play a significant role of future software development in CIG.

Adjoints and inversion. The computation of adjoint solutions plays a crucial role in most inversion and data assimilation techniques, and can also be used in error control and mesh refinement methods. Since all relevant geodynamics problems are time dependent, the ability to solve adjoint problems requires storage and retrieval of forward solutions as well as solving linearized backward-in-time equations. For nonlinear problems, this capability is difficult to retrofit and will, therefore, be part of initial designs of future codes. Since the ultimate aim of most of geodynamics computing is understanding (inverting) or matching (assimilating) observed behavior, we consider adjoint solvers a crucial technology.

Where differentiation of outputs with respect to inputs is not possible (see Section 3.3.2), we will ensure that codes can be used in a Bayesian inversion framework [16, 17] by allowing codes to run as subroutines of Monte Carlo model parameter generators, and to postprocess solutions for arbitrary observables.

Multiphysics coupling. Many of the applications we ultimately want to consider in geodynamics require the interaction of different physical processes across a broad range of temporal and spatial scales. Examples include bringing thermodynamics into mantle convection solvers and considering the flow of melt and other fluids in long-term deformation solvers. Furthermore, the grand challenge problems of coupling codes from entire areas such as mantle convection and plate motion; core dynamics and mantle convection; or seismics and mantle convection involve multiphysics coupling.

To simultaneously solve these disparate problems we will use techniques from multiphysics approaches that have matured over the past decade. This includes, among others, physics-based preconditioners, multiscale time integrators, and stable projection schemes at interfaces between different physical regions. We will also use libraries like *FEniCS* [18] or *deal.II* [19, 20] that make defining the combined semilinear forms of coupled problems much simpler, as well as offer significant support to develop physics-based and other blocked preconditioners.

Scientific visualization. To date CIG has focussed on providing software for geodynamics modeling; however, as models become larger and more complex there is an emerging need for tools that enable more rapid construction and interpretation of the resulting data. By using open-source, platform-independent methods for interactive visualization of 3D data such as *3DVisualizer* [21], *Visit* [22] or *Paraview* [23] that are built on flexible, extendable interfaces such as *VRUI* [24] or *VTK* [25], CIG will enhance the value to researchers and students of geodynamics codes.

Reliance on existing packages. CIG has fared well in its reliance on existing packages for common components of its software. It allows us to share software development, testing, documentation, and training resources with other communities. Equally importantly, we share in the experience of what numerical methods work well. As a consequence, we will continue to base our software on packages like *PETSc*, *Trilinos*, *hypre* (for parallel linear algebra, solvers, and preconditioners), *StGermain*, *deal.II* or *FEniCS* (for discretizations and numerical methods) for most things not particular to geodynamics. CIG however, always remains open to new methods and approaches that can be demonstrated to meet community needs.

Trends in parallel, distributed and heterogeneous computing. It is clear that large-scale computing is currently moving into new directions with the integration of many compute cores on single chips and the availability of graphic processing units (GPUs) to augment the compute power of main processors. As a large-scale user of parallel computing, CIG will need to address the resulting challenges for its software.

In keeping with our philosophy of using proven technology, we do not intend to rewrite our software for individual architectures. However, we will closely follow the various standards that are being defined around these new programming paradigms (for example OpenMP [26], OpenCL [27], or CUDA [28]) as well as platform-independent support libraries such as the Threading Building Blocks [29, 30]. We will use these technologies on a case by case basis, depending on the suitability of our applications for a certain kind of hardware and when we deem the software layer that supports this hardware to be sufficiently mature, interoperable between different systems, and stable with regard to future hardware developments. We will also be able to share in the expertise we receive through libraries like *PETSc*, *Trilinos*, or *deal.II* that face the same challenges. For example, *deal.II* already has support for multicore systems through the *Threading Building Blocks* library, and *Trilinos* is considering a similar technology. Furthermore, work is underway for GPU support in *PETSc* with initial work focusing on fast methods with dense operations, such as Fourier transforms, multipole methods, and wavelet transforms, and accelerated multigrid methods.

5 The Lessons of CIG

As we go forward, it is essential that we critically evaluate our progress during the first four and half years of CIG. As discussed above, CIG has developed new software that has contributed to a growing list of scientific contributions and established stronger links between the fields of geoscience and computational science. Perhaps not surprisingly for a new and complex effort, some CIG activities have been more successful than others. In this section, we describe the character of the software development that unfolded through CIG, elaborating on selected examples. In summary, we learned that (1) users and developers must work closely together as the best codes are produced from efforts that have both scientific and computational leadership, and (2) education and training of users at all levels are essential, especially for advanced codes and methods.

At the start of CIG, we quickly realized that each of the subdisciplines were distinct in terms of their technical expertise, their ability to develop codes, their reception to software engineering, and how close they were to achieving their goals with existing software. Some disciplines had enough experience in developing software in their research groups and never developed strong links with CIG. At the other end of the spectrum, some groups had little experience with code development and were more willing to collaborate with computational scientists. Most subcommunities fell between these extremes.

The mantle convection community has a long history of code development in which the products of software development have been passed on by several generations of faculty and graduate students, and the leadership of CIG has drawn heavily on the experience of this community with code development, benchmarking efforts, and the challenges of supporting codes developed by individual researchers. Much of CIG's development for the mantle convection community focussed on *CitcomS*, a finite-element code that solves for thermal convection within a spherical shell. When CIG took over maintenance of *CitcomS*, a wide range of versions were in use, as the community had not adopted a single repository, and different research groups had developed significant enhancements to the code that were not in the main repository. By advocating the use of software repositories and giving technically proficient users access, we reduced the number of alternative versions of the software and shortened the time required for features developed by others to be merged back into the main software. Furthermore, CIG collaborated with the community in the development of new packages (like *Rhea* [15]) that transcend the current capabilities of *CitcomS*.

CIG also contributed to the mantle convection community and software development through workshops. We saw an important change in outlook between the two mantle convection workshops in 2005 and 2008 (Table S1). The first workshop focused primarily on scientific questions, while computational methods were only reviewed. Three years later the discussion expanded to address both scientific and computational challenges and provided detailed discussion of new methods, equation solvers and code verification all of which are essential for progress.

Software development in tectonics unfolded differently and resulted in two new codes to address the needs of two distinct communities. These developments followed different paths, reflecting differences in interests and technical abilities within the two subcommunities. In short-time scale tectonics (i.e., the earthquake cycle), an existing working group (evolving out of SCEC) that included a few individuals with computational experience teamed with CIG engineers to develop *PyLith*. *PyLith* integrates the functionality of *EqSim* [31, 32], for dynamic rupture, and a version of *Tecton* [33], for quasi-static problems. *PyLith* uses a hierarchy of software functionality as it was initially envisioned by CIG. As part of the general toolkit needed for the solution of many of the problems that CIG encounters, Matthew Knepley at Argonne National Laboratory (ANL) developed infrastructure for parallel storage and manipulation of general finite-element meshes (*Sieve*).

Developments in the long-time scale community followed a different path. At a Tectonic Modeling workshop in 2005 (Tables S1 and S2), the community requested development of new open-source software that could handle large deformations with viscoplastic rheologies, such that fault evolution would unfold in three dimensions as a function of time using the Arbitrary Lagrangian Eulerian method. Although such methods had long been used for 2-D problems, the expertise did not exist within the U.S. community to develop a scalable 3-D code. CIG realized that we could leverage technology developed by our partners, the Victorian Partnership for Advanced Computing and Monash University in Australia, and develop such a code with common components. The end result was *Gale*, a parallel, two- and three-dimensional, implicit finite-element code.

With working versions of *CitcomS*, *PyLith*, and *Gale*, our self-evaluation found wide differences in the ability to use the software. *CitcomS* has a large and expanding base of users who investigate a range of challenging problems in deep mantle dynamics, plate tectonics, and planetary science. *PyLith* is used by dozens of researchers within the short-term tectonics community to model crustal deformation associated with volcanoes and earthquake faulting, often in problems involving complex geologic structure. Many users cite the extensive documentation along with rapid response from the developers as one of the main advantages of using *PyLith*. *Gale* has been used to calculate some of the first 3-D models with nucleating and evolving faults, although the user base for this code has been slow to build, in part likely due to the complexity of the code. These cases illustrate that the success of software development is strongly affected by how closely the developers work with the users and the technical sophistication of the users. It is clear that in order to serve beginning and expert users, CIG must focus more effort on integrating users in the development process and intensify training efforts for our community. In response to this need, we plan an increased emphasis on workshops, education, and other activities supporting community engagements.

6 Proposed future development for CIG

6.1 Infrastructure

CIG will continue to use modern software engineering technology to minimize the risks of software development in light of our evolving requirements. All of our software will be open source as described in our policy on SOFTWARE OWNERSHIP AND LICENSING (Supplementary Material). The software repository and attendant web site will remain central to CIG's objectives of facilitating collaboration and sharing of validated open-source software. We will maintain unit and regression testing for those software packages that currently have such functionality while adding it to all new software packages. We will continue to build the software in the repositories nightly or automatically in response to software commits. The nightly regression testing generates an electronic report that contains the build and test failures (including the platforms on which they occurred) and allows our community of developers to rapidly identify when a change in a repository component or platform has caused an error or inconsistency. Regression testing gives users of the repository confidence in the robustness of our software. A wide range of open software packages will be used, such as *Subversion* and *Mercurial* for repositories, *Roundup* for bug tracking, and *build-bot* for building and regression testing, while continuously maintaining an open eye to new features and packages that we can deploy in support of CIG.

CIG expects to receive software donations and we plan to follow the guidelines laid out in our SOFTWARE SUPPORT POLICY (Supplementary Material). Essentially, we have a policy of support in which we only commit resources to those codes that have been validated, benchmarked, and qualify as state-of-the-art in the opinion of our Steering Committee.

6.2 Ongoing development of existing CIG software

Plate tectonics and deep mantle dynamics Many of the important questions in plate tectonics and mantle dynamics will rely on better integration with mineral physics and seismic wave propagation. The development of adjoint and multiphysics capabilities is also crucial. CIG is clearly headed down this path with our existing code *CitcomS*, an experimental code *Rhea* (soon to be released), and a new system proposed below.

CitcomS is a versatile tool which we propose to maintain. The latest version now has several compressible formulations for thermo-chemical convection in a full sphere [1] giving us the ability to incorporate multiple equations of state simultaneously. Users have begun to exploit this new capability, and we plan to include additional equations of state, such as thermodynamic formulations of realistic mantle compositions [34]. Given the computational complexity of realistic thermodynamic formulations we intend to implement this new feature using look-up tables. *CitcomS* scales well up to several hundreds of processors, bringing 30 km resolution globally easily within reach of present-day department clusters and TERAGRID facilities. We plan to extend this capability by making existing solvers more robust for variable viscosity problems using recently developed preconditioners [35]. We also expect to recode low-level routines to take advantage of GPUs, as recently demonstrated for conjugate gradient and multigrid solvers. We recently completed an integration of *CitcomS* with *SPECFEM*, allowing the direct calculation of seismograms for fully heterogeneous convection models. This capability will need to be improved as users begin to test the seismic predictions with data. Finally, we expect to introduce an adjoint of the *CitcomS* energy equation [36] back into the main line of the code.

Core dynamics CIG currently maintains a serial code, *MAG*, for numerical simulations of the geodynamo. This code is well tested, fully documented, and broadly used in the geodynamo community. Individual components of the geodynamo model can be independently controlled, allowing users a great deal of flexibility in the choice of problems that can be solved. Over the initial phase of CIG a number of additional features were added to *MAG* to improve the graphical output and facilitate comparisons between predictions and observations. CIG also developed a Science Gateway for *MAG*, which greatly expands the potential user base by providing an intuitive graphical interface and eliminating the need for novice users to install the code on a local machine. Ongoing developments will focus on adding to the repository a parallel version of *MAG*, or equivalent spectral model, by donation of an existing code. Several suitable codes have already been developed and CIG will work with the authors who express interest in collaborating with CIG to make one or more of these codes available to the broader community.

Lithosphere-scale tectonics and surface processes *Gale* currently serves as an important tool for studying the evolution of fault systems, mountain building, and rifting. However, it has limited potential for development much beyond its current features as a result of design complexity. Therefore, we will take a two-pronged approach to code development. We will continue to maintain *Gale*, add new features, and address portability issues. The most ambitious planned addition to *Gale* is the introduction of magma migration using an existing suite of modules that are based on the same software framework. We also plan to increase the flexibility of meshing in *Gale*, implement more general boundary conditions, and introduce an interface that allows users to couple *Gale* to surface process codes. These new features will provide important capabilities for the study of lithospheric deformation, allowing users to investigate interactions between magma migration and crustal-scaling faulting and interactions between climate and tectonics.

In addition, we plan to develop a new code that will offer all of the features in *Gale* (and more) through software that utilizes common components in other CIG software. As discussed in the next section, the new code will greatly enhance the range of spatial scales for lithospheric deformation and facilitate integration with mantle dynamics. We also plan to couple lithospheric deformation with regional-scale seismic imaging, analogous to the current integration of *CitcomS* and *SPECFEM*.

Fault system dynamics Ongoing development in fault system dynamics will focus on *PyLith* to support research on earthquake physics and the behavior of plate tectonics on the seismic time scale. One of the primary tasks will be coupling simulations of earthquake rupture propagation with interseismic deformation in order to resolve the entire seismic cycle, from dynamic deformation on time scales of milliseconds to minutes to quasi-static deformation on time scales of days to thousands of years. This will empower researchers to confront many basic questions regarding the behavior of fault systems, including the interactions between dynamic processes during earthquake rupture and spatio-temporal patterns of seismic activity, aseismic transients, episodic tremor, and crustal stress. Other goals for *PyLith* development include improving the efficiency of the solvers through the use of multi-grid solvers and preconditioners well-suited for dislocation implementations (e.g., constraints on relative displacements).

Magma dynamics Magma dynamics is a natural multiphysics extension of mantle convection and lithospheric deformation that links the dynamics of a low-viscosity fluid to a strong, but deformable solid matrix. Because fluid and solid flow can be highly coupled, the dynamics of both materials need to be solved within the same code. Moreover, these codes should use the same libraries as those used for mantle convection and lithospheric deformation to ensure that the software is transferable between applications. To test the efficiency and accuracy of the different code bases, we have developed the Magma Dynamics Demonstration Suite (MADDs) as a set of benchmark problems of increasing complexity. MADDs has currently been implemented using the same code base as *Gale* and as a combination of *FEniCS* and *PETSc* libraries, which permit automatic code-generation and physics-based block preconditioners. Initial work has also begun on a *deal.II* implementation. A version of MADDs will be included as a component of *Gale* to allow modeling of fluid flow.

Seismic imaging CIG currently maintains *SPECFEM* as a tool for both global and regional analysis of seismic observations. Ongoing development of the regional version *SPECFEM Sesame* will include more flexible meshing for regional models using existing meshing software. Current successes in 2D mesh generation will be extended to 3D problems. CIG will expand the seismology portal interface to *SPECFEM*, include regional simulations with local, regional and teleseismic events, and streamline the mesh generation for user-defined regional models. CIG also will encourage donations from developers of other maturing seismic imaging codes that are used for making high resolution images at regional scales from scattered waves, such as those using the inverse generalized Radon transform and Kirchhoff integral inversion methods. Our goal is to add and enhance existing codes for regional seismic analysis. CIG will continue to facilitate comparison of numerical methods and benchmarking of codes, in part through cosponsoring of workshops on synthetic seismogram calculation and imaging science with the European seismic imaging initiative, QUEST (a follow-on to its computational forward seismology initiative, SPICE).

6.3 New development pathways for geodynamics software

In addition to ensuring the continued availability of our existing codes as outlined in the previous section, CIG intends to develop a series of new codes. Detailed specifications and the exact set of numerical methods to be used will be refined through discussions within the CIG working groups and with the broader geodynamics and computational science communities.

Plate tectonics and deep mantle dynamics New development efforts will focus on codes with adaptive mesh refinement so as to better resolve plate boundaries and other sharp rheological features in the solid Earth. As an intermediate step, we will soon take advantage of *Rhea*, which uses the *Alps* library for adaptive mesh refinement [15]. Because the functionality of *Alps* will be integrated into *deal.II*, we propose to build a new code based on the *deal.II* library which already has a well-documented thermal convection solver that is both adaptive and parallel. This will allow us to have modules that will solve for melting within restricted parts of the overall global domain that itself will have adaptive mesh refinement. Multiphysics at plate boundaries, for example, is an unresolved problem and it is essential that we not build an entire code that locks the community into one formulation. Overall, the code will allow the user to select between a spherical or Cartesian geometry and will have thermo-chemical advection, compressible and incompressible solvers, AMR, adjoints and a standard model of melting as available features.

Lithosphere-scale tectonics and surface processes Given the similar computational requirements with mantle convection, we propose building a new lithospheric-scale tectonics code from many of the same components as the new AMR mantle convection code. The new code will also allow flexibility in the choice of model components (including multiple models of fluid transport). AMR is essential to resolve the localization that occurs within faults and shear zones as a result of plasticity, damage and melting. Tracking plastic failure or damage through a field approach and using a semi-Lagragian method offers a potential path forward that eliminates many of the short-comings associated with the Lagrangian techniques currently employed by *Gale* and *SNAC*. Semi-Lagrangian methods advect fields along characteristics and the number of unknowns scales only with the mesh size (and quadrature order) of the underlying finite element. Most importantly, semi-Lagrangian methods do not have the usual (CFL) limit on time steps, allowing advection of fields across multiple elements in a single step. This ability is critical for AMR techniques, which may have elements sizes that vary by many orders of magnitude. Semi-Lagrangian methods are already implemented in *Rhea*.

Fault system dynamics As studies of fault systems migrate from simplistic, idealized geometries and rheologies towards realistic geologic structure with nonplanar, intersecting fault surfaces and nonlinear rheologies with localized deformation, software must provide resolution across wider ranges of spatial and temporal scales, integration of elastodynamics with heat and fluid flow, and greater scalability. While *PyLith* is a modular, state-of-the-art code and provides great flexibility, our future needs will likely demand a more refined approach in order to achieve the desired resolution and scalability. One potential development pathway involves integrating the structured adaptive mesh refinement used in *Rhea* with finite-element basis functions that can generate displacement jumps (dislocations) within elements. An alternative approach combines improved methods for refining and coarsening unstructured meshes with a more traditional cohesive element dislocation formulation, such as the one currently used in *PyLith*. Finally, we will investigate alternative strategies such as meshless methods that may alleviate many of the issues associated with multi-scale techniques and complex geometries. Additional capabilities, such as coupling of heat and fluid flow with elastodynamics, may favor the use of *FEniCS* or *deal.II*, which would require the migration of techniques for fault slip into those tools.

Magma dynamics Developments in magma dynamics will be carefully coordinated with those for mantle convection and lithospheric deformation. Using *deal.II* in the development of new codes for mantle convection and lithospheric deformation facilitates inclusion of magma dynamics in those codes, because *deal.II* provides support for AMR in addition to a rich library of higher order elements necessary for magma dynamics. We expect that there will also be a parallel development of codes using *FEniCS* and *PETSc* libraries for fully unstructured simplicial meshes (triangles and tetrahedra). *FEniCS* allows the introduction of multiple formulations for magma dynamics (and solid deformation), through high-level automatic code generation. *FEniCS* and *deal.II* are both high-level C++ libraries for generalized finite elements and complement each other. We expect developments in one library to be readily implemented in the other (e.g., semi-Lagrangian methods for fully unstructured grids and physics-based block pre-conditioners for magma have already been implemented in a hybrid *FEniCS/PETSc* code). Ultimately, we foresee magma dynamics and fluid transport being an optional extension of all the solids codes.

Seismic imaging Regional seismic investigations require advances in imaging due to the recent acquisition and deployment of large numbers of high quality broadband seismometers by EarthScope and other international initiatives. We propose to support this rapid expansion of data collection by adding the ability to use P and S teleseismic body waves as seismic sources in *SPECFEM Sesame*. The synthetic waveforms can be compared to observations, such as travel-time anomalies identified in tomography, amplitude anomalies, receiver functions, and different possible anisotropic structures. A second important development is to extend the adjoint capabilities to *SPECFEM Sesame* for waveform tomography at regional and local scales. CIG will also support extensions of the adjoint capability in both the regional and global versions of *SPECFEM* to include anisotropy and attenuation, allowing for a more complete description of Earth's mechanical properties.

Further work is also required to relate the outputs of geodynamic models to inputs for calculation of synthetic waveforms. A link already exists to map *CitcomS* outputs to *SPECFEM*. We also need to develop the opposite capability by using the output of seismic inversion as a constraint or initial condition on geodynamic models. Ultimately, we seek to establish interfaces for inputs/outputs of Earth models for integration of convection, tectonic and seismic modeling at both global and regional scales.

Core dynamics New developments in core dynamics stand at a crossroads. Most researchers agree that the current models are inadequate to solve the most pressing science questions. However, there is less consensus on the path forward. Current models are based primarily on spectral methods, although examples of models based on finite differences, finite elements or spectral elements can be found in recent literature. CIG can play an influential role in future developments by facilitating workshops that bring together researchers in Earth Science and Computational Science. Experience and expertise gained by CIG in other areas of geodynamics will undoubtedly be a great help. Contingency funds in later years of this proposal ould be used to support the development of a core dynamics model once a consensus emerges from the community.

Scientific visualization As geodynamics models grow more complex, researchers find it increasingly challenging to interpret the results. We plan to take advantage of the partnership with the UC Davis KeckCAVES [37] which is developing new methods for interactive exploration of complex data using visualization on a wide range of hardware platforms including desktop computers, GeoWalls [38], and immersive systems [24]. Figure 7 illustrates the use of visualization for a mantle convection model run using CIG codes. The model of the Alaskan subduction zone was constructed using *CitComCU*. Interactive visualization [21] was used to validate the geometry of the

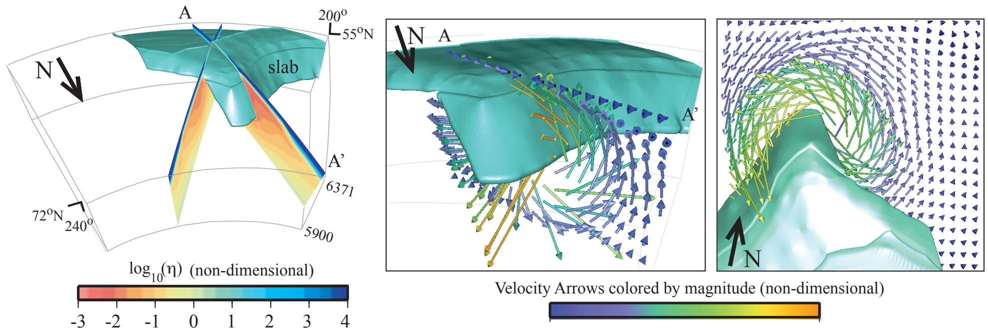

Figure 7: Visualization of a model of the Alaskan subduction zone, showing an isosurface of slices of viscosity with arrows representing velocity [40].

slab, which was constructed using available geophysical data including seismicity and has more than 100 million finite-element nodes (for a resolution of 2.35 km to 20 km). *Visualizer* was also used to interpret the results of the *CitComCU* model output. The software developed at the Keck-CAVES follows the same standards as CIG and is provided as open source, supported, documented, and validated software that runs on a wide range of platforms from a desktop computer to a large immersive system [21]. We plan to continue development with visualization of complex, time-varying data using techniques for visualizing volumetric, vector, and tensor data. Visualization of AMR grids provides a particular challenge to which we bring expertise [39].

New problems/New methods The new development pathways discussed in this section represent the current state of an ongoing discussion between the Earth Science and Computational Science communities fostered by CIG. Some of these developments are direct results of CIG activities (e.g. *Sieve* or *Rhea*), others reflect recent developments in computational science (e.g. FEniCS), but few were actually envisioned when CIG began. Moreover, we expect that new methods, models and problems will arise as this project progresses and CIG will remain responsive to these new developments. CIG maintains both formal and informal procedures for incorporating new ideas from the community (see management section) and will seek out and foster new collaborations and projects if desired by the geoscience community.

6.4 Verification and Validation

CIG's software development methodology adheres to verification and validation best practices to the extent possible. Here, *verification* means that codes are tested to ensure that they correctly and accurately solve the problems they are intended to solve [41]. While this testing should be an obvious step, many research codes are not verified by comparing with complex, known solutions; even fewer development efforts have automated methods to ensure that continued development of the code does not break existing features. To ensure a structured development process that addresses these issues, CIG uses a large suite of unit tests for our newly developed software; for example, *PyLith* alone has more than 1,100 unit tests run nightly. We also run automated comparisons of generated solutions against at least one analytical result for each of our solvers, and have a buildbot-based [42] system that compiles and builds our software packages at regular intervals on multiple platforms. CIG has developed an open-source tool (*cigma*) that compares different solutions to a given problem – computed using different codes, meshes, and shape functions – against each other and against analytical solutions (see Supplementary Material). Solutions and continuously generated test results are published in our manuals and web pages. These efforts ensure the high quality that makes CIG's software so useful to the community. The capacity to provide this level of software development sophistication is one of the major advantages of a consortium effort like CIG; it is otherwise difficult for individual researchers to incorporate verification testing in the context of individual research grants.

On the other hand, *validation* refers to the process of ensuring that the equations in our software represent an accurate approximation of the real world system they are supposed to model [41]. This lies at the core of computational geodynamics research: to determine which models accurately represent our planet. CIG has, to date, largely built its software on widely accepted models that have already been validated; the mantle convection community in particular has a long history of carrying out benchmarks for validation of codes. However, CIG actively validates codes and facilitates validation efforts by the geodynamics community. A core design criterion for our software is that it is extensible and allows for the use of different models that can then be compared against experimental and observational data. For example, *SPECFEM* can use different Earth or regional seismic models and can be coupled with *CitcomS*; *CitcomS* can use different viscosity models and has incorporated the capability to deal with compressibility; and both *Gale* and *PyLith* can use different rheology models. Consequently, at the minimum CIG's software allows its users to validate it against field data from other sources. CIG has also directly participated in community activities in which a range of numerical codes attempt to match the results of laboratory experiments.

CIG intends to continue to emphasize sound software development practices. We will continue to insist on automated testsuites for all codes we actively support and will continue to build extensive unit tests for software developed as part of CIG. We put similar emphasis on the quality of the packages upon which our software is built; for example *PETSc*, *StGermain*, and *deal.II* run some 400, 500 and 1,650 tests every night, respectively, and have automated builds on a number of different platforms with dozens of different configurations. This level of testing ensures that the foundation upon which we build our software is solid. Similarly, seismology will develop benchmarks for testing imaging codes, in all probability jointly with the Quest group.

Through our emphasis on testing, we are also increasing the level of sophistication of software written by the geodynamics community at large. Several community-driven new codes, such as *PyLith*, have substantial automated testsuites. This is a significant step forward from previous practices. We therefore believe that CIG, by setting an example, is able to influence the production of software even outside its immediate reach and lead to more professionally written codes and more robust scientific results.

6.5 Education and Outreach

Experience from the initial phase of CIG has helped to identify the broad needs of a diverse community. On the one hand, we seek to expand the base of users and increase the level of computational expertise across the Earth Sciences as a whole. We also strive to increase the number of expert users who can expand the capabilities of the existing software suite. By making their innovations available to the broader community, we hope to accelerate the pace of scientific discovery. In order to meet the diverse needs of users, we propose a flexible strategy for delivering education and training. Our strategy is twofold: first, to provide education and training for new and potential users that prepare them to use CIG computational methods and, second, build the computational geoscience community through partnerships with established programs such as the NSF-funded "On the Cutting Edge" program for professional development of geoscience faculty [43], and the Center for Deep Earth Research (CIDER) [10] which provides summer programs for graduate students and undergraduates.

Many new users of CIG software have little prior training in physical modeling and computational methods, so there is a clear need to provide introductory instruction in the basic use of the codes and the underlying methods. This instruction includes hands-on training at smaller scientific meetings and workshops, following the successful model established at previous meetings of IRIS, EarthScope and CIDER. We also propose to continue the longer and more focused CIG-sponsored workshops that are organized around specific scientific disciplines and specific software in the repository. These longer workshops would be augmented by numerical modeling workshops that cover a series of computational themes. Most, if not all, of these themes would be common across all CIG software and would involve participants with diverse interests. We plan to hold two computational workshops per year on topics that progressively increase in sophistication over the duration of the grant. The goal is to provide a comprehensive introduction to the principal numerical methods used in geodynamics. Education modules and self-directed tutorials from previous sessions would be available through the CIG website. We expect the education modules to be a valuable resource for researchers who are unable to attend the workshops or graduate students who start their research part at different times during the duration of the grant. We also expect the education materials to be useful for instructors teaching upper-level undergraduate and graduate classes in geodynamics and numerical modeling. Individual instructors would be encouraged to contribute simulations and animations to the CIG Education Webpage for a broader audience. We will partner with David Mogk of Montana State University to make these materials available through the Cutting Edge geoscience education programs (a letter of commitment is attached.) Our goals are to showcase the role of computing in Earth Sciences and help educators share materials.

A smaller set of users require intermediate or advanced training to help them extend the functionality of CIG software. This type of training entails extended, focused discussion between developers and users. Intermediate training would begin via discussions at the longer and more focused CIG workshops. Additional and more advanced training will require extensive discussions through virtual meetings as well as occasional site visits (developers traveling to users and users traveling to developers). Several projects within CIG that involve developers at multiple locations already use these approaches for effective collaboration.

7 Management Plan

In 2010, CIG will move to new headquarters at the University of California, Davis, under the directorship of the PI, Louise Kellogg. The new director was selected by the Executive Committee after a national search. Caltech and UC Davis will work closely to ensure a seamless transition of the central facility and staff, and uninterrupted support of CIG codes and CIG activites. In particular we will keep the core repository at Caltech until a mirror is operational. A subcontract to Caltech will support software engineering efforts and ensure ongoing support for CIG codes during the transition. Code development for projects such as *PyLith* or *Rhea*, which are principally distributed will not be affected by the transition. UC Davis will house CIG in offices adjacent to the

Computer Science, Mathematics, and Geology departments, providing ready access to expertise in numerical methods, gridding, high-performance computing, networking, scientific visualization, remote collaboration, geophysics, and tectonics. Through a recent faculty hiring initiative in computational science and engineering, the university has built clusters of expertise in computational geophysics, computational materials (including earth materials), and complex systems, and established a computer cluster which will be available to CIG developers. CIG will also have access to high-speed networking and state-of-the art scientific visualization facilities through the KeckCAVES (visualization dedicated to the geosciences) and the Institute for Data Analysis and Visualization (IDAV) which supports interdisciplinary computer science research.

During and after the transition to the new location, CIG will continue to draw upon the expertise, vision, and guidance of the community to remain a flexible and evolving organization that readily adapts to changes in priorities and computational opportunities. We will maintain our community-centric management structure, which was initially modeled on the experience of other successful NSF-supported community infrastructure projects and has evolved over the last five years to meet the needs of the CIG community. The management plan, outlined here, is codified in a set of bylaws (see Supplementary Material).

Institutional membership and committees. CIG is an institutionally-based organization governed by its membership organizations, via an elected Executive Committee (EC) and an elected Science Steering Committee (SSC). CIG's member institutions include educational, government or non-profit organizations in the US with a sustained commitment to CIG objectives, as well as non-voting foreign affiliate members. The current membership of CIG includes 42 member institutions and 9 foreign affiliates (a complete list is included in the Supplementary Material). CIG welcomes new active partners; requests for membership can be made by institutions at any time, and are approved by a simple vote of the electorate. The latter consists of the existing institutional representatives. CIG is governed by two committees, the Executive Committee and the Science Steering Committee, whose members are elected by the membership. In addition, a number of informal or *ad hoc* committees have been established as needed. These include topical working groups, described below, and an organizing committee for each workshop. CIG makes an effort to represent a diverse range of scientific disciplines, institutions, and points of view on these committees, and especially makes an effort to involve both early- and mid-career scientists in planning and decision-making.

The **Executive Committee (EC)** is the primary decision-making body of CIG; as is current practice, it will meet at least once per year to approve the annual Strategic Plan (SP), management plan, and budget, and to deal with major business items. With the Director, the EC will handle additional business responsibilities through its regular meetings, teleconferences, and electronic mail. The Executive Committee has five voting members (a Chair, Vice-Chair, and three members at-large) and two non-voting members (the Director and the SSC Chair). Exclusive of the Director, these members will be elected by representatives of member institutions for staggered three-year terms. The Executive Committee will have the authority to approve proposal submissions and contractual arrangements for CIG. The EC also appoints a **Nominating Committee** charged with proposing candidates for membership on the elected committees.

The **Science Steering Committee (SSC)** provides guidance across all of the sub-disciplines of computational geodynamics, while maintaining an overall balance of Geoscience and Computational Science. Its principal duty will be to assess the objectives and needs of the sub-disciplines covered by CIG, identify common themes, and engage in an ongoing, open discussion on software design to maximize the range of potential applications and enhance the reusability of CIG codes.

As CIG has matured, we have identified a need for advice and a broader perspective on CIG activities from the greater geoscience and computational science communities. An **External Advisory Council** will be established, consisting of 6 scientists from educational, non-profit, industry, and governmental institutions with a major commitment to research in fields related to computational geodynamics. Members of the council will be elected by the Executive Committee for three year terms. The Advisory Council will meet annually to review all aspects of CIG, provide advice as requested by the Executive Committee, and report to the Director and the Chairman of the EC. The CIG bylaws will be amended accordingly to establish the Advisory Council.

Working Groups are an important part of the CIG management structure. These *ad hoc* groups are organized around scientific disciplines and are usually tied to development or support of one or more of the CIG codes. Each working group has a chair who acts as a contact person for both the community and the SSC. The working groups provide critical input to the SSC in establishing science priorities, and help to ensure that CIG is responsive to community needs. Currently, the most active working groups are those in Computational Seismology, Long-Term Tectonics, Magma Migration, Mantle Convection, and Short-Term Crustal Dynamics.

Administration and project management. The Director is the Chief Executive Officer of the organization and bears ultimate responsibility for its programs and budget. The Director's responsibilities (in coordination with the EC and SSC) will include: (a) overall management of the core CIG activities, including supervision of CIG staff, (b) devising a fair and effective process for the development of the science plan, based on proposals or work plans such as those submitted to the SSC, (c) ensuring that appropriate project milestones and usage metrics are established, documented, and reported, (d) acting as PI on proposals submitted by the core CIG facility, retaining final authority to make and implement decisions on grants awarded to the core facility and contracts, (e) ensuring that funds are properly allocated to various CIG activities, and (f) representing CIG to the broader community and to the National Science Foundation.

In the current CIG management structure, many day-to-day organizational and administrative responsibilities are currently being handled by the Director, who also maintains his professional academic commitments of teaching and research. For continued CIG development, which is occurring with more external partners, project management has become critical and requires full-time commitment. Therefore, we propose to add a Chief Software Engineer/Project Manager (CSE) who will be a full-time employee of CIG with a strong background in computational science and will work on behalf of the community to assist the Director and the EC and SSC. The CSE will act as direct liaison between the core Software Development Team (SDT), external developers and the greater user community. The CSE's responsibilities will include (a) day-to-day management of the core SDT, (b) tracking and coordinating with external development projects, (c) helping establish and track project milestones and usage metrics, and (d) overseeing the preparation of technical reports. In addition, the CSE will provide critical input on software design and architecture, and will help identify new opportunities in both computational science and methods for software development. The CSE will enable the Director to interact more with the scientific community and ensure that CIG is responsive to community needs.

Formulating CIG priorities and management of its resources. The management and governing structure is designed to balance CIG's activities between ongoing support and development of the established codes and infrastructure, and response to community needs for new codes and infrastructure. Thus, for example, CIG is committed to providing ongoing support for a widely used family of mantle convection codes (*CitcomS*) while at the same time introducing new methods such as AMR.

Concepts and plans for CIG activities come directly from member institutions, working groups and individual scientists. The Working Groups, Science Steering Committee, and Director are conduits for formal and informal dialog among the CIG community. In addition to the informal dialog that such committee members naturally have as members of the scientific community, CIG has, and will continue to use, a formal process for bringing new ideas in from the community. At any time, users from Member Institutions may submit one-page proposals for new CIG software development. These proposals will continue to be posted on the web for the community to read and evaluate, with a form where members of the user community can add comments and evaluation. At least once per year (and more often as needed), the SSC will evaluate these proposed CIG activities, will formulate a prioritized list of tasks, and develop a yearly strategic plan for CIG. The SSC will work in consultation with the software development team and the Director to assess how tasks are inter-related and related to the broader needs of the community. To make this process as productive as possible, the Director and SSC will be on the lookout for opportunities for new activities and will work with those who are in the process of proposing a new effort to ensure that it is within the scope of CIG's mission.

On at least a yearly basis, the EC will evaluate the task list developed by the SSC and decide how to allocate resources to specific software development tasks. At its disposal, the EC will have resources to respond to the evolving community needs expressed through these task lists, including the SDT and funds for contracts.

Metrics. CIG currently uses three primary metrics to assess the usage and impact of CIG: software downloads, participation in workshops, and contributed citations documenting the scientific advances facilitated by CIG software. Software downloads have steadily increased as CIG codes expand and users are trained in their use, with more than 3500 downloads in the last year (see Supplementary Material for a breakdown of the history of downloads of different software packages). CIG workshops have involved more than 900 participants in the last five years; more than 40% of the participants at CIG workshops were students or postdocs. These workshops provide both training for new users and forums for planning development of CIG codes. Members of the community have supplied citations for publications that relied on CIG infrastructure (see references in Supplemental Material). This list is likely incomplete, and should be seen as a sample of scientific outcomes supported by CIG. Many of the publications have appeared in the last 18 months (since Jan. 2008), indicating that CIG's scientific influence is emerging. CIG will continue to monitor these basic metrics, and will develop new metrics as appropriate. The strongest indication of the impact of CIG may be the growing base of active researchers using CIG software to address important scientific questions and their attendance at the variety of CIG workshops.

8 Results of Prior NSF Support

"Computational Infrastructure for Geodynamics, CIG", PIs Michael C. Gurnis and Michael Aivazis, California Institute of Technology, EAR-0426271, $6,750,000; 9/1/2004–8/31/2009.

This grant established the Computational Infrastructure for Geodynamics and the software development, community building, and training programs that form the basis for the current proposal. The primary accomplishments of CIG are detailed in Section 2, "CIG Accomplishments", and Section 5, "The Lessons of CIG". Supplementary material accompanying this proposal describes the CIG software, licensing, distribution policies, and documentation that were developed with this grant. The Supplementary material also includes documentation of the intellectual and broader impacts of CIG, including downloads of software, participation in CIG programs, abstracts of scientific results submitted by the CIG, and a partial bibliography of papers using CIG codes.

"CI-TEAM Implementation Project: Enabling Interactive Visual Exploration and Remote Collaboration for the Geosciences and Physical Sciences", PI: Louise Kellogg, co-PIs: James Crutchfield, Bernd Hamann, Dawn Sumner, Magali Billen, OCI-0753407, $920,672, 04/01/2008–03/31/2011

This continuing cyber-infrastructure grant supports a team of geoscientists, computer scientists, and physicists to develop immersive virtual reality (VR) methods for visual data analysis in an educational context. In the first 1.5 years of the project, we provided interdisciplinary training in scientific visualization to undergraduate and graduate students from geology, physics, and computer science, developed techniques for hierarchical representation and visualization of digital elevation data to be used by students in geology [44], developed tools for virtual geologic mapping [45], released a volume visualization code for use with geodynamics models (including CIG codes), seismic tomography, neutron tomography, serial thin-sections, and other 3D data [21, 24], and developed an accessible, easy-to-use Python programming interface for VR environments aimed at students. Using a LIDAR visualization tool [46], students developed techniques for identification and representation of tree canopies in scanned point data (for research in environmental sciences and geomorphology), tools for interactively exploring low-dimensional nonlinear dynamical systems, constructed Virtual Earth Science Education Modules to help students improve spatial reasoning skills, and conducted experiments on remote collaboration [47]. We prototyped and demonstrated an inexpensive portable immersive visualization environment for use in classroom, outreach, or research environments. We used the above tools and associated data in undergraduate courses in Geology Physics, Mechanical and Aerospace Engineering, Computer Science, in Freshman Seminars in Geology, and in a high-school science enrichment program (involving more than 100 students), and in an outreach program at the California State Fair (reaching more than 16,000 people). Mentored research, education, and training are integrated; projects are carried out by teams of faculty, postdocs, grad students, and undergraduates, providing opportunities for more advanced students (grad students and postdocs) to mentor undergraduates. Our products include software released under an open-source license, two Ph.D. dissertations, several undergraduate theses, presentations and the publications cited above.

References

[1] E. Tan, W. Leng, S. Zhong, and W. Gurnis, “Citcoms v3.0 – a compressible thermo–chemical mantle convection code,” in *EOS Trans. AGU*, ser. Fall Meet. Suppl., Abstract, vol. 88, 2007, pp. DI14A–01.

[2] D. Komatitsch and J. Tromp, “Spectral-element simulations of global seismic wave propagation -i. validation,” *Geophys. J. Int.*, vol. 149, 2002.

[3] “PETSc: Portable, Extensible Toolkit for Scientific Computation,” http://www-unix.mcs.anl.gov/petsc/petsc-2.

[4] “The StGermain framework: enabling communities of scientists to iteratively develop computational codes,” http://www.stgermainproject.org/.

[5] “IRIS: Incorporated Research Institutions for Seismology,” http://www.iris.edu.

[6] “EarthScope,” http://earthscope.org.

[7] “MARGINS,” http://www.nsf-margins.org.

[8] “SCEC: Southern California Earthquake Center,” http://www.scec.org.

[9] “Compres,” http://www.compres.stonybrook.edu/.

[10] “CIDER: Cooperative Institute for Dynamic Earth Research,” http://www.deep-earth.org.

[11] M. Spiegelman and D. W. Sparks, “Everybody wants a super model: Towards an integrated theory of mid-ocean ridge dynamic,” *EOS Trans. AGU*, vol. 79, no. 45, p. F39, 1998, fall Meeting, Invited.

[12] M. Spiegelman, P. B. Kelemen, and E. Aharonov, “Causes and consequences of flow organization during melt transport: The reaction infiltration instability in compactible media,” *J. Geophysical Research*, vol. 106, no. B2, pp. 2061–2077, 2001, www.ldeo.columbia.edu/∼mspieg/SolFlow/.

[13] C. Tape, Q. Liu, A. Maggi, and J. Tromp, “Adjoint tomography of the southern california crust,” *Nature*, vol. submitted, 2009.

[14] C. P. Stark, E. Choi, and M. Convertino, “Landslide rupture and length–depth sclaing,” *Earth and Planetary Science Letters*, 2009, submitted.

[15] C. Burstedde, O. Ghattas, M. Gurnis, G. Stadler, E. Tan, T. Tu, L. C. Wilcox, and S. Zhong, “Scalable adaptive mantle convection simulation on petascale supercomputers,” *International Conference for High Performance Computing, Networking, Storage, and Analysis (ACM/IEEE Supercomputing 2008)*, pp. 1–15, 2008.

[16] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2004.

[17] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*. Springer, 2006.

[18] "FEniCS: Software for automated solution of differential equations," http://fenics.org/.

[19] W. Bangerth, R. Hartmann, and G. Kanschat, "deal.II – a general purpose object oriented finite element library," *ACM Trans. Math. Softw.*, vol. 33, no. 4, pp. 24/1–24/27, 2007.

[20] ——, `deal.II` *Differential Equations Analysis Library, Technical Reference*, 2008, `http://www.dealii.org/`.

[21] M. I. Billen, O. Kreylos, B. Hamann, M. A. Jadamec, L. H. Kellogg, O. Staadt, and D. Y. Sumner, "A geoscience perspective on immersive 3d gridded data visualization," *Computers and Geosciences*, vol. 34, p. 1056, 2008.

[22] "Visit: Software that delivers parallel interactive visualization," http://www.llnl.gov/visit/.

[23] "Paraview: An open-source, multi-platform data analysis and visualization application," http://paraview.org/.

[24] O. Kreylos, "Environment-independent vr development," in *ISVC 2008, Part I, LNCS 5358*, e. a. Bebis, G., Ed., vol. LNCS 53. Berlin Heidelberg: Springer-Verlag, 2008, p. 901.

[25] "VTK: Visualization toolkit," http://vtk.org/.

[26] "OpenMP: The OpenMP API specification for parallel programming," http://www.openmp.org.

[27] "OpenCL: The open standard for parallel programming of heterogeneous systems," http://www.khronos.org/opencl/.

[28] "CUDA: Compute Unified Device Architecture," http://www.nvidia.com/object/cuda_home.html.

[29] "Threading Building Blocks," http://www.threadingbuildingblocks.org.

[30] J. Reinders, *Intel Threading Building Blocks*. O'Reilly, 2007.

[31] B. T. Aagaard, T. H. Heaton, and J. F. Hall, "Dynamic earthquake ruptures in the presence of lithostatic normal stresses: Implications for friction models and heat production," *Bulletin of the Seismological Society of America*, vol. 91, no. 6, pp. 1765–1796, 2001.

[32] B. T. Aagaard, J. F. Hall, and T. H. Heaton, "Characterization of near-source ground motions with earthquake simulations," *Earthquake Spectra*, vol. 17, no. 2, pp. 177–207, 2001.

[33] C. A. Williams and G. Wadge, "An accurate and efficient technique for including the effects of topography in three-dimensional elastic deformation models with applications to radar interferometry," *J. Geophys. Res.*, vol. 105, pp. 8103–8120, 2000.

[34] L. Stixrude and C. Lithgow-Bertelloni, "Thermodynamics of mantle minerals: 1. Physical properties," *Geophysical Journal International*, vol. 162, pp. 610–632, 2005, doi: 10.1111/j.1365-246X.2005.02642.x.

[35] D. A. May and L. Moresi, “Preconditioned iterative methods for Stokes flow problems arising in computational geodynamics,” *Physics of the Earth and Planetary Interiors*, vol. 171, pp. 33–47, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/B6V6S-4T5JHY6-2/2/2d30f943c948d97fb5797a1f89b73698

[36] L. Liu and M. Gurnis, “Simultaneous inversion of mantle properties and initial conditions using an adjoint of mantle convection,” *Journal of Geophysical Research*, vol. 113, p. B08405, 2008.

[37] “Keckcaves,” http:keckcaves.org.

[38] “Geowall Consortium,” http:www.geowall.org.

[39] O. Kreylos, G. Weber, E. Bethel, J. Shalf, B. Hamann, and K. Joy, “Remote interactive direct volume rendering of AMR data,” Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-49954, 2002.

[40] M. Jadamec, “Three-dimensional lithosphere and mantle dynamics: Models of the subduction-transform plate boundary system in southern alaska,” Ph.D. dissertation, University of California, Davis, 2009.

[41] D. E. Post and L. G. Votta, “Computational science demands a new paradigm,” *Physics Today*, vol. 58, no. 1, pp. 35–41, 2005.

[42] “Buildbot: A system to automate the compile/test cycle,” http://buildbot.net/.

[43] “On the Cutting Edge,” http://serc.carleton.edu/NAGTWorkshops/index.html.

[44] T. Bernardin, E. Cowgill, R. Gold, B. Hamann, O. Kreylos, and A. Schmitt, “Real-time terrain mapping,” H. Hagen, Ed., vol. Scientific Visualization: Challenges for the Future. Los Alamitos, California.: IEEE Computer Society Press, 2009, p. in press.

[45] L. Kellogg, G. Bawden, T. Bernardin, M. Billen, E. Cowgill, B. Hamann, M. Jadamec, O. Kreylos, O. Staadt, and D. Y. Sumner, “Interactive visualization to advance earthquake simulation,” *Pure Appl. Geoph.*, vol. 165, p. 621, 2008.

[46] O. Kreylos, G. Bawden, and L. Kellogg, “Immersive visualization and analysis of LiDAR data,” in *ISVC 2008, Part I, LNCS 5358*, e. a. Bebis, G., Ed., vol. LNCS 53,. Berlin Heidelberg: Springer-Verlag, 2008, p. 846.

[47] G. Kurillo, R. Bajcsy, K. Nahrsted, and O. Kreylos, “Immersive 3d environment for remote collaboration and training of physical activities,” *Proceedings of the IEEE Virtual Reality 2008 conference*, p. 269, 2008.