Computational Infrastructure for Geodynamics (CIG)

NSF Award Number EAR-0426271

Five-year Strategic Plan: Sept. 1, 2004 to Aug. 31, 2009

Submitted Dec. 1, 2004

Table of Contents

1. An Overview of CIG

The Computational Infrastructure for Geodynamics (CIG) develops, supports, and disseminates community-accessible software for the greater geodynamics community from model developers to end-users. The software is being developed for problems ranging widely from mantle and core dynamics, crustal and earthquake dynamics, magma migration, seismology, and related topics. With a high level of community participation, CIG leverages the state-of-the-art in scientific computing into a suite of open-source tools and codes. The infrastructure now under development will consist of:

- a coordinated effort to develop reusable, well documented and open-source geodynamics software;
- the basic building blocks – an infrastructure layer – of software by which state-of-the-art modeling codes can be quickly assembled;
- extension of existing software frameworks to interlink multiple codes and data through a superstructure layer;
- strategic partnerships with the larger world of computational science and geoinformatics;
- specialized training and workshops for both the geodynamics and larger Earth science communities.

CIG is now putting into place a small core team of dedicated software architects and engineers whose work will be guided by scientific objectives formulated by the scientific community. This Software Development Team will provide software services to the community in terms of programming, documentation, training, and support. Guidance for the programmers will come from a Science Steering Committee (SSC) and an Executive Committee (EC) whose emphasis is to identify and balance common needs across disciplines. Remaining in place is our plan to house the software development team in the Center for Innovative Technologies (CIT^2), a new facility in Pasadena.

Since the official start date of the project, Sept. 1, 2004, the PIs and members of the community have been diligently putting the project into place. Under the Chairpersonship of Louise Kellogg, the Writing Committee's last act was to appoint a Nominating Committee (P. Olson, ch., G. Masters & M. Simons) who put together a slate of candidates for the first Executive Committee that were then voted into office by the Member Institutions. The outcome of that election process was as follows:

- Mark A. Richards, (Chairman, Oct., 2007), University of California
- Marc Spiegelman (Vice Chairman , Oct., 2006), Columbia University
- Bill Appelbe (At-Large, Oct., 2007), Victorian Partnership for Advanced Computing

- Carl Gable (At-Large, Oct., 2005), Los Alamos National Laboratory

They are joined by the Director (M. Gurnis) and Chief Software Architect (M. Aivazis) in an *ex-officio* capacity.

The Director and other members of the Executive Committee worked to put this first Strategic Plan together. The EC met on the Caltech campus Nov. 8-9 and dealt with pressing issues now facing CIG, many of which are described in this document.

2 An overview of the strategic plan

CIG strategy is centered on four interrelated goals: Development of a software repository and web site, development of a software framework, organization and facilitation of community participation, and the training of the software users.

2.1 Repository & Web site. The software repository and attendant web site are central to CIG's objectives of developing a comprehensive system of software verification and validation. Moreover, the repository is critical to bring modern software engineering practices to our community and CIG's software development team. Our goal is to have a single repository for developer use that allows multiple developers to work on large pieces of modular software. Careful attention will be paid to the means by which CIG software is downloaded and installed on user computers (the build procedure). This is a critical goal because nothing annoys and distracts users more than software that is difficult to download and install correctly. Software that is dependent on multiple levels of libraries, such as the software CIG will develop, is particularly susceptible to build problems. Thus, a critical goal for CIG is getting an effective build procedure working in year one. The build procedure will evolve as a key principle in all CIG software development.

CIG's method of development will be agile computing, which encourages the developers to construct test cases for each new piece of functionality before the actual code is written. The result is a suite of unit tests that exercise a large fraction of the code base. Further, as software defects are detected and repaired, our developers will construct tests that verify the absence of the flaw and make them part of our regression test suite. In order to provide a continuously updated assessment of the state of the code, the entire code base will be checked out of the source repository and will be automatically rebuilt on a regular basis on all supported platforms, perhaps even nightly. After a successful build, these suites will be automatically run and a report will be generated that contains the build and test failures and the platforms on which they occurred. This system of regression testing will be established by the end of year one and updated based on current practices in subsequent years.

The web site will be used for both software downloads as well as a means by which codes, including CIG software, can be benchmarked and compared to existing benchmark results (the later point is needed for the community to participate in our system of software verification and validation). This is essential for increasing the overall quality of our science especially as it moves into realm of complex, multi-physics and multi-scale simulations. Our intermediate term goal is for CIG to develop a standard procedure by which geophysical codes can be compared. This will mean the establishment of benchmarks that are sensible numerically (verification) and geophysically (validation). A common and permanent repository will be created for benchmark results. CIG will collaborate with individual communities sequentially in order to develop this procedure: CIG staff will adapt and refine this procedure as it gains experience. In year one we will work with a community interested in fault interactions (mainly scientists interested in the interpretation of GPS data). We will also start working with the mantle convection community. We hope to start working with the computational seismology community next and proceed on to magma-migration, the geodynamo, and long term crustal dynamics, not necessarily in this order.

2.2 The software framework. The development of a software framework for geophysical modeling and its upkeep are central to CIG's strategy. The framework will consist of software on two levels, as detailed in the NSF proposal:

- An *infrastructure layer* of software by which community and state-of-the-art modeling codes will be quickly assembled

- A *superstructure layer* of software by which multiple codes and data can be interlinked through the extension of existing software frameworks

CIG's resources are limited and entirely new software will only be generated when needed. CIG will rely on lower level components developed in more science-neutral arenas, including the many DoE supported projects, and higher-level frameworks to interlink software as much as possible. Exactly how this framework, or hierarchy of software, is developed and what exists at the different layers will be slowly developed through community involvement.

As a first step in moving forward, CIG will work with several communities starting in year one to make existing geophysical software more modular, using science-neutral libraries, and driven by a superstructure framework. Based on the high-degree of community participation we have identified the continued development of a finite element code used to model the interaction of faults on earthquake cycle and longer time-scales. This project will allow us to build a version of "Tecton" out of some common components and then to have the code scale from single workstations to large parallel computers. The code will also allow us to test the build procedure. From a managerial prospective, this will also allow us to test and refine the ability for members of the software development team to work closely with members of the community, in this case Charles Williams of RPI. We also hope to repeat the process in year one on a mantle convection code. This software will give the community some usable examples to work with before the end of year one.

The development of this software depends on our ability to start putting into place a software development team that has a balance of skills that will allow it to develop modular and useful software and to work effectively with the larger community. Currently projected CIG funds will allow for three software engineers by the start of year two in addition to the half-time support of Matt Knepley in the MSC Division at Argonne National Laboratory. There will also be moderate funds available for smaller subcontracts.

Establishing the subaward with ANL and initiating the collaboration with Matt Knepley are top priorities in year one. For all CIG technical staff, as well as those individuals working on contracts or subawards, we will establish clear tasks and timetables for the delivery of those tasks. Establishing the tasks and time-line for the ANL subaward are the first test of this procedure.

Our goal is to have at least two Software Engineers (SE) hired by the end of year one and possibly three. But quality will be critical so we will prudently make appointments after close scrutiny of each candidate. Two types of software engineers will be recruited: Numerical modelers or analysts and experts in software integration. An example of the first type of SE would be an individual with extensive experience with numerical models, especially an individual who has developed models or algorithms, in an area of applied physics (including geophysics) or engineering. An example of the second would be an individual with extensive experience with integrating low-level libraries, such as equation solvers, into high-level computer codes or integrating software components with scripting languages for scientific or engineering applications. Caltech will open these positions in late 2004 and we will advertise in venues such as Monster.com as well as trade journals and spread the word through the CIG community by means of bulk e-mails and the web site. Critical to the development of the software framework will be recruitment of outstanding software engineers who can participate in agile software development while also effectively interacting with the CIG community. One of the principles of agile software development is the need for close (almost daily) interaction between software developers and end users of the software (scientists) to minimize the risk of misunderstandings and ensure effective communication occurs.

2.3 Organizing community participation. Equally important for CIG, is the guidance from the scientific community on what this infrastructure should accomplish for their evolving research needs. This is accomplished through community oversight of CIG, committees, and workshops. Two key principles guide CIG's interaction with the community:

- *Openness*: All CIG reviews, meeting minutes, and other documents will be openly available to the community in a timely fashion, for review and public discussion, unless this conflicts with individual or institutional privacy rights (such as salary level or personal information). Openness minimizes the risk of actual or perceived bias or conflicts of interest within CIG.
- *Interaction*: All CIG committees and workshops should have an open and balanced representation of both the scientific and software communities. Interaction minimizes the risk of balkanizing the community CIG serves.

The priority of the first year is to convene the Science Steering Committee (SSC) and hold a series of workshops to gather input from the community. The Science Steering Committee will be elected by the member representatives through a voting mechanism described in the By-Laws; in the first year, we hope to have the SSC in place before the summer begins. These workshops and committee meetings will be the mechanism by we will layout a coherent software system that can deliver the core functionality of geodynamics (e.g. solve solid-mechanics, wave-propagation/seismology, coupled fluid/solid mechanics and dynamo calculations) in an open and extensible fashion. This exercise requires a systematic evaluation of current models to understand the techniques that are known to currently work, then to design and publish a set of abstract interfaces and data-structures that will allow them to interoperate. This initial design will focus on identifying a common subset of meshes, discretizations, and algorithms that are most commonly used and have promise for significant re-use in different disciplinary problems. The design will also include a coherent and consistent super-structure for defining model input and output, job submission, run archiving and interaction with external data-bases and products such as plate models. In year one, CIG will commit to co-supporting a workshop with SCEC for the finite element modeling of fault interactions as well as a renewed mantle convection workshop (the latter to be held in Boulder, CO). We will actively work with the computational seismology community to attempt to bring this community together; we will actively work with IRIS for co-funding of such a workshop.

2.4 User training. A key objective of CIG is the widespread adoption of CIG-developed software by the general earth science community and, in particular, by geophysicists. Both significant training and comprehensive documentation will be required. Potential users will want to know how to use CIG software as well as understand the underlying algorithms and implementation details. Based on current research directions, sophisticated users who require a deeper understanding of the workings of the software will likely have different backgrounds and interests. One class of user will consist of traditional computational scientists who want to use several different computational components (such as a solver and a mesher) to create working codes that incorporate algorithmic innovations. Another class of user will consist of those who attempt to string together several components of CIG software with non-CIG data analysis tools. In addition, there will be a class of less-technically-demanding users who will want to use CIG codes in a standard manner.

Several mechanisms will be used to train CIG users: (A) small, focused workshops; (B) visits to the CIG site to work with CIG staff and other users; and (C) the production of training manuals and a web site. CIG will use these mechanisms to provide training for both the expert and non-expert user.

No training workshops will be held in year one, but plan to have our first workshop in the later half of year two and then subsequently hold several training sessions per year. However, the technical writer will be recruited in year one and start working on the development of user manuals and training material immediately. The first task will likely be a manual for the use of the reengineered finite element code for the fault interaction community.

3. Annual Goals and Milestones:

3.a All goals and milestones for Year 1 (Sept. 1, 2004 – Aug. 31, 2005)

I. Software Repository
- a. Donated software available on web
- b. Adoption of version control, regression testing, and uniform build procedure
- c. Community accepted benchmarks
    - i. Mantle Convection Community
    - ii. Fault-interactions on earthquake time-scales.
    - iii. Seismic wave propagation

II. Computational Framework
- a. Specification of APIs for geophysical codes & adoption of standard formats for storing data (such as netCDF or HDF5)
- b. Frozen codes with common interfaces
    - i. eqsim
    - ii. "Tecton"
    - iii. CitcomS
    - iv. SpecFEM3D
- c. Pyre bindings for PETSc routines of geophysical interest
- d. Linking existing community codes to PETSc
    - i. "Tecton"
    - ii. CitComS (incompressible & compressible)
- e. Start interfacing codes with mesh generation tools

III. Organizing Community Participation
- a. Joint meeting of the Executive committee and newly established Science Steering Committee (SSC)
- b. Mantle convection workshop in Boulder, CO.
- c. Finite element modeling of fault interactions workshop (Co-sponsorship with SCEC)
- d. Computational seismology (will seek co-sponsorship with IRIS)

IV. User Training
- a. Start user manuals
    - i. CIG software repository
    - ii. Reengineered codes that use PETSc
- b. Start development of material for training sessions

3.b Software Repository Years 1-5

Year 1 Sept. 1, 2004 – Aug. 31, 2005	Year 2 Sept. 1, 2005 – Aug. 31, 2006	Year 3 Sept. 1, 2006 – Aug. 31, 2007	Year 4 Sept. 1, 2007 – Aug. 31, 2008	Year 5 Sept. 1, 2008 – Aug. 31, 2009
• Donated software available on web	• Maintain distribution of donated software available on web • Start distribution of new CIG generated Software	• Maintain distribution of existing and new software on web	• Maintain distribution of existing and new software on web	• Maintain distribution of existing and new software on web
• Adoption of version control, regression testing, and uniform build procedure	• Maintenance and update of version control, regression testing, and uniform build procedure	• Maintenance and update of version control, regression testing, and uniform build procedure	• Maintenance and update of version control, regression testing, and uniform build procedure	• Maintenance and update of version control, regression testing, and uniform build procedure
• Community accepted benchmarks: - Mantle Convection Community - Fault-interactions on earthquake time-scales. - Seismic wave propagation	• Community accepted benchmarks - Complete for mantle convection community - Complete for fault-interactions group - Continue to work on seismic wave propagation - Initiate for geodynamo problem - Initiate for long term crustal dynamics	• Community accepted benchmarks - Complete for geodynamo - Complete for long term crustal dynamics - Initiate for magma migration	• Community accepted benchmarks - Complete for magma migration	• Continued posting and revision of all community benchmarks

3.c Software Framework Years 1-5

Year 1 Sept. 1, 2004 – Aug. 31, 2005	Year 2 Sept. 1, 2005 – Aug. 31, 2006	Year 3 Sept. 1, 2006 – Aug. 31, 2007	Year 4 Sept. 1, 2007 – Aug. 31, 2008	Year 5 Sept. 1, 2008 – Aug. 31, 2009
• Reengineering: • Frozen codes with common interfaces - eqsim - “Tecton” - CitcomS - SpecFEM3D • Infrastructure Layer: • Specification of APIs for geophysical codes & adoption of standard formats for storing data (such as netCDF or HDF5) • Linking existing community codes to PETSc - “Tecton” - CitComS (incompressible & compressible) • Pyre bindings for PETSc routines of geophysical interest • Start interfacing codes with mesh generation	•Reengineering: • Frozen codes with common interfaces - geodynamo code to be donated - Two 1-D synthetic seismogram codes - Magma Migration code • Infrastructure Layer: • Final adoption of APIs for geophysical codes • Interfacing with mesh generation • Pyre bindings for SCRIP • Develop new solvers with Libraries as proposed by members, such as: •Advection-diffusion modules •Project Specific Layer •Develop post-processing of solver-linkage modules specific by members, such as •Geodynamic model-seismic wave linkage module •Gravity, geoid, topography, magnetic field post processing module • Superstructure Layer: • Link existing solvers or codes with superstructure layers as proposed by members • Continue to generalize Pyre coupers and exchangers for geophysical code	• Infrastructure Layer: • Develop new solvers with libraries as proposed by members •Project Specific Layer •Develop post-processing of solver-linkage modules specific by members • Superstructure Layer: • Link new solvers with superstructure layers as proposed by members	• Infrastructure Layer: • Develop new solvers with libraries as proposed by members •Project Specific Layer •Develop post-processing of solver-linkage modules specific by members • Superstructure Layer: • Link new solvers with superstructure layers as proposed by members	• Infrastructure Layer: • Develop new solvers with libraries as proposed by members •Project Specific Layer •Develop post-processing of solver-linkage modules specific by members • Superstructure Layer: • Link new solvers with superstructure layers as proposed by members

3.d Organizing Community Participation Years1-5

Year 1 Sept. 1, 2004 – Aug. 31, 2005	Year 2 Sept. 1, 2005 – Aug. 31, 2006	Year 3 Sept. 1, 2006 – Aug. 31, 2007	Year 4 Sept. 1, 2007 – Aug. 31, 2008	Year 5 Sept. 1, 2008 – Aug. 31, 2009
•2 Executive Committee (EC) meetings • Joint meeting of the Executive committee and newly established Science Steering Committee (SSC) • Mantle convection workshop in Boulder, CO. • Finite element modeling of fault interactions workshop (Co-sponsorship with SCEC) • Computational seismology (will seek co-sponsorship with IRIS)	• 1 EC meeting • 1 Science Steering Committee • Long-term crustal dynamics workshop (co-sponsor with “GeoMod” European group) • Geodynamo community workshop (find international partner) • Future Directions of coupling codes workshop (co-sponsor with GeoFramework)	• 1 EC meeting • 1 Science Steering Committee • Sponsor or co-sponsor about 3 workshops as proposed by members	• 1 EC meeting • 1 Science Steering Committee • Sponsor or co-sponsor about 3 workshops as proposed by members	• 1 EC meeting • 1 Science Steering Committee • Sponsor or co-sponsor about 3 workshops as proposed by members

3.e User Training Years 1-5

Year 1 Sept. 1, 2004 – Aug. 31, 2005	Year 2 Sept. 1, 2005 – Aug. 31, 2006	Year 3 Sept. 1, 2006 – Aug. 31, 2007	Year 4 Sept. 1, 2007 – Aug. 31, 2008	Year 5 Sept. 1, 2008 – Aug. 31, 2009
• Start user manuals - CIG software repository - Reengineered codes that use PETSc • Start development of material for training sessions	•Applying the Pyre Framework (co-sponsor with GeoFramework) • Training session on use of reengineered FE code for fault group (with SCEC) •Sponsor visitations of students, post-docs and faculty to CIG site • Training session on use of reengineered FE code for mantle convection (hold in conjunction with topical science conference) • Start user manuals - Finish for CIG software repository - Finish for reengineered codes that use PETSc • Material for training sessions - Finish for Reegineered codes with PETSs	•Develop training material and web site for new CIG software •Sponsor visitations of students, post-docs and faculty to CIG site •Hold specialized training session at CIG for newly developed software •Hold specialized training session at other site in conjunction with topical conference, preferable one that is sponsored by another group.	•Develop training material and web site for new CIG software •Sponsor visitations of students, post-docs and faculty to CIG site •Hold specialized training session at CIG for newly developed software •Hold specialized training session at other site in conjunction with topical conference, preferable one that is sponsored by another group.	•Develop training material and web site for new CIG software •Sponsor visitations of students, post-docs and faculty to CIG site •Hold specialized training session at CIG for newly developed software •Hold specialized training session at other site in conjunction with topical conference, preferable one that is sponsored by another group.

4. Allocation of resources by goal

Sept. 1, 2004 – Aug. 31, 2005

Software Engineer #1 (Numerical Analyst/Modeler) [Total 0.5 FTE]
- I.c (benchmarks): 0.1 FTE
- II.a (uniform data): 0.2 FTE
- II.b (common interfaces): 0.2 FTE

Software Engineer #2 (Software Integration) [Total 0.5 FTE]
- I.c (benchmarks): 0.2 FTE
- II.a (uniform data): 0.1 FTE
- II.b (common interfaces): 0.2 FTE

ANL Subcontract [Total 0.5 FTE]
- I.c (uniform build): 0.1 FTE
- II.d (Tecton & CitcomS): 0.2 FTE
- II.b (Pyre/PETSc): 0.15 FTE
- III.a (Convection workshop): 0.05 FTE

Director [Total 0.08 FTE paid by CIG]
- III (organize community): 0.08 FTE

Chief Software Architect [Total 0.08 FTE]
- II (framework): 0.08 FTE

Secretary [Total 0.25 FTE]
- III (community workshops): 0.2 FTE
- Paperwork for Management: 0.05 FTE

Technical Writer/Web master [Total 0.5 FTE]
- I.a (software on web): 0.2 FTE
- I.c. (benchmarks on web): 0.2 FTE
- III (community workshops): 0.1 FTE

Supplies and Expenses [$25K]
- I, II, III

Travel [$12K]
- I,II,III

Participant Costs [$88K]
- III

4. Membership

4.1 CIG Members and Member representatives:

Argonne National Laboratory (MSC)
California Institute of Technology
Colorado State University
Columbia University
Harvard University
Johns Hopkins University
Los Alamos National Laboratory (ES)
Massachusetts Institute of Technology
Oregon State University
Pennsylvania State University
Princeton University
Purdue University
Rensselaer Polytechnic Institute
State University of New York at Stony Brook
U.S. Geological Survey (Menlo Park)
University of Texas at Austin
University of California San Diego
University of California, Berkeley
University of California, Davis
University of California, Los Angeles
University of Colorado
University of Maine
University of Maryland
University of Michigan
University of Minnesota
University of Southern California
University of Washington
Washington University
Woods Hole Oceanographic Institution

4.2 CIG Foreign Affiliates and representatives:

Australian National University
Monash University
University of Sydney
Victorian Partnership for Advanced Computing

4.3 Strategy for keeping members informed

Member representatives and individuals within the larger CIG community (including those at member institutions) will be kept informed through several different means.

1. Through e-mail. Official member representatives are informed about significant CIG business and upcoming events (including business meetings, workshops, and recent software releases and posting on our web site) through an official e-mail list. CIG will also provide a moderated, list-server for interested general users, who can subscribe through the web-site.

2. Through the http://Geodynamics.org web site. The upcoming CIG calendar of events is posted and continuously revised. Nearly all CIG documents, including proposals submitted to CIG, the annual revision of the CIG Strategic Plan, By-Laws, etc., are posted of this site. The Web site is the principal means for standard software downloads, sharing of community benchmark, specifications of standards, and distribution of user & training manuals.

3. The annual CIG Business meeting. This meeting will be open to all and will be a forum for open discussions of the working of CIG, including past and upcoming activities & the Strategic Plan. In year one, this meeting will be held in conjunction with the Fall AGU meeting in San Francisco. Depending on the success of this meeting, this may be the venue for subsequent years.

4. CIG sponsored and co-sponsored workshops and training sessions. The current status of CIG will be presented at these workshops and we expect that CIG members will attend such workshops.

6. Annual science plan

CIG does not carry out science directly, but by providing the infrastructure to modeling specialists and other earth scientists who use modeling tools we facilitate the science funded elsewhere. In general, CIG will provide the infrastructure to allow several areas of computational geophysics to move forward and solve previously intractable problems. In this plan, we briefly describe the specific areas of science that we will likely facilitate in the coming year.

An expanding area of research is the study of crustal deformation on short-time scales in tectonically active areas opened by the wide (and growing) availability of dense, continuous GPS networks. The research is quickly expanding in anticipation of the PBO component of EarthScope, dense continuous geodetic networks in Southern California, Japan, and Tawain, and earth deformation imaged through the InSAR method. Forefront research is focused on the nature of the strain and stress field before and after earthquakes, silent "earthquakes", the interaction between active faults in a plate boundary zone, and magmatic inflation and deflation. Computational models play a fundamental role in all of these studies. CIG will be working with a group of researchers to coordinate development and validation of 3D quasi-static, finite-element codes for modeling crustal deformation; develop deformation models with observed topography, fault geometries, rheological properties, geologic slip rates, geodetic motions, and earthquake histories; and use these models to infer fault slip, rheologic structure, and fault interactions through stress transfer.

CIG will directly impact this research by providing assistance to this community in the inter comparison of different modeling codes, essential for the verification and validation of the methods used in modeling. We will work toward the adoption of standard file formats (for all areas of computational geophysics) and this will allow this group to easily compare results. We will work closely to reengineer one of their community codes with parallel equation solvers. The adoption of standard I/O and parallel solvers will allow the community to more easily use the finite element codes and run progressively larger problems on Beowulf clusters and other parallel machines. We would not expect to see substantial science in this first year, although the software would likely be in use. Next year, we would expect explicit scientific results, probably for problems tailored to Southern California earthquakes and deformation within Pacific-rim subduction zones, including Cascadia.

We will also be working closely with the mantle convection community. Current research in this area is motivated by global seismic tomographic inversions (facilitated by GNS), questions emerging from interdisciplinary research in the CSEDI program, questions emerging from interdisciplinary research in the MARGINS program (Especially the Subduction Factory initiative), and regional studies of the western U.S. now gearing up in anticipation of the USArray component of EarthScope. Some of the most challenging computational issues arise for convection in multicomponent convection, such as thermo-chemical convection. A related problem will be in the integration of melting and chemical transport into convection codes. There is substantial interest in coupling models of mantle and crustal dynamics. CIG will directly impact this field through the adoption of common file formats, establishing benchmarks, and then developing a common procedure to compare different methods and codes. In the last ten years, studies of mantle convection have dramatically expanded into more complex physical (rheological, multiphase, geometrical) systems and there is a need for the community to validate and verify results. The CIG staff will also be working with several mantle convection codes to use common components and this will allow more users to have access to validated mantle convection software by the end of the first year. We expect that the software will be used to study lower mantle dynamics, including the nature of the complex region at the base of the mantle, superplumes, the layering of mantle convection, the fixity and/or drift of hot-spot plumes, the interaction of the continental lithosphere with mantle convection, the origin of mantle convection, and subduction zones, including the dynamics of the mantle wedge.

An important aspect of our adoption of standard interfaces for codes and common file formats is that we are actively planning to coordinate this effort with observational initiatives. Obviously, much of the standardization will happen internal to CIG (through the upcoming workshops in crustal deformation and mantle convection and through CIG software engineers). However, we plan to develop with EarthScope and/or IRIS a standard format for the storage of both tomography models (3-D) and geodynamic models (3-D plus time). This simple standardization should vastly expand interdisciplinary research. Some obvious benefits would be in mantle convection that is being actively investigated observationally by seismologists and theoretically by modelers. For studies of superplumes (a subject of substantial interest) one could rapidly get answers to questions such as: What is the flow field predicted by a particular seismic inversion? What seismic waveforms are predicted for a specific realization of a thermo-chemical superplume? Seismic tomographic inversions from the USArray deployment could be fed directly into a flow model for mantle flow and crustal deformation. Much of the infrastructure development for this will likely be completed in year one with scientific results emerging in year two.

7. Five Year Management Plan

CIG will need the expertise, vision, and guidance of the community if it is to remain a nimble and evolving organization. Consequently, we are adopting a *community-centric* management structure that draws upon features of successful NSF-supported community infrastructure projects in the Earth sciences. The management plan, outlined here, has been codified in a set of by-laws available on our web site (http://Geodynamics.org).

7.a. Institutional Membership and Executive and Science Standing Committees.

CIG will be an institutionally-based organization governed by an Executive Committee. The structure of CIG will recognize member institutions, which are educational and not-for-profit organizations with a sustained commitment to CIG objectives, and a number of foreign affiliate members. The Member Institutions will change over time because CIG is an *open organization*, available to any institution seeking to collaborate on the development of open-source software for computational geodynamics and related disciplines.

The Executive Committee will be the primary decision-making body of CIG; it will meet at least twice per year to approve the annual science plan, management plan, and budget, and to deal with major business items, including the election of a Nominating Committee. With the Director, the Executive Committee will handle the day-to-day decision-making responsibilities through its regular meetings, teleconferences, AccessGrid sessions, and electronic mail. The Executive Committee will have seven members. It will have four voting members: the Chairman, the vice Chairman, and two members at-large. These members will be elected by representatives of member institutions for staggered three-year terms. The three nonvoting members are the Director, the Chief Software Architect, and the Chairman of the Science Steering Committee. The Executive Committee will have the authority to approve proposal submissions and contractual arrangements for CIG.

CIG will have a Science Steering Committee ultimately consisting of seven elected members including a chairperson. The committee will balance expertise in both geoscience and computational sciences and provide guidance within all of the sub-disciplines of computational geodynamics. Their principal duties will be to assess the competing objectives and needs of all the sub-disciplines covered by CIG, it will provide initial assessment of proposals submitted to CIG, and it will revise the Five Year Strategic Plan. Recommendations from the SSC will be passed on the EC.

7.b. Administration.

The Director will be the Chief Executive Officer of the organization and will bear ultimate responsibility for its programs and budget. The Director's responsibilities will include: (a) devising a fair and effective process for the development of the Strategic Plan, based on proposals or work plans such as those submitted to the Executive Committee by the Science Steering Committee, and overseeing the plan's implementation, (b) acting as P.I. 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, (c) ensuring that funds are properly allocated to various CIG activities, (d) overseeing the preparation of technical reports.

The Chief Software Architect (CSA) will serve as a non-voting member of the Executive Committee. His role will be to provide advice and perspective to the Executive Committee on the overall composition, integration, and balance between software development activities of the organization. He will provide frequent assessments of our software, identify new opportunities in both computational science and methods for software development, and provide evaluations of prospective members of the Software Development Team. The Executive Committee retains the authority to appoint the CSA.

7.c. Formulating CIG Priorities and Management of its Resources.

Concepts and plans for CIG activities will come directly from member institutions and their elected committees. Ideas and plans will move from members to the Science Steering Committee and finally to the Executive Committee. As part of the development of the Strategic Plan , the SSC will formulate a prioritized list of tasks for software development for the coming year, how these tasks are both inter-related and related to the broader needs of the community, and then transmit this as a recommendation to the Executive Committee. On at least a yearly basis, the Executive Committee will allocate resources to specific software development tasks. Following this allocation of resources, the EC will periodically appoint small committees to interface directly with the software development team (SDT).

It is expected that members of the SSC will be fully engaged in a dialog with the user community and active users of CIG software. Besides the constant dialog that such committee members would naturally have with the community, CIG will have a formal process for bringing new ideas up from the community. On a continual basis, users from Member Institutions will be able to submit one-page proposals for new CIG software development tasks. These proposals can be submitted at any time and are posted on the web for the community to read and evaluate. There will be a comments page where members of the user community can add scientific comments and evaluation. Periodically, but at least once per year, the SSC will evaluate these proposals in light of other information obtained from the community, formulate a prioritized list of tasks, and then submit it to the Executive Committee.

At its disposal, the Executive Committee will have resources to respond to the evolving community needs expressed through these task lists, including the Software Development Team and funds for contracts. However, the Executive Committee will also put into place two mechanisms for generating new resources and funds for CIG.

- *Augmented funding*. CIG will agree to develop additional software upon receipt of augmented funding. For example, a PI at a Member Institution may submit a science proposal to a federal agency in which the proposed work is either wholly or in part dependent upon software not yet available. This software would presumably be more specialized than the highest priority and core CIG tasks, but still encompassed within the mission of CIG and needs of the community. Following submission of a one-page proposal as described above, the Executive Committee will determine whether or not CIG can develop this software. If CIG can develop the software, the EC will detail the resources and funding required on a form for attachment to the PI's proposal. If the proposal successfully passes through peer review and the federal agency agrees to fund the project with augmentation to CIG funding, we will develop the software.

- *Collaborative proposals*. In time, CIG will have a specialized staff with skills in software development, numerical analysis, information technology, and related fields, skills not readily accessible within the geoscience community. We believe that members of the community will formulate collaborative research projects with SDT members. If such collaborative projects are judged to be of high merit for CIG by the EC, CIG will develop collaborative proposals. We expect one target of opportunity to be federal programs that require collaboration between scientists from both information technology and the domain sciences, such as the geosciences. It would be expected that such projects would provide funding for both external PIs and members of the SDT.

Software developed through either of these two mechanisms will be open source and made available to the community without restriction, like all CIG software. During its first two years of operation CIG must by necessity focus on a core set of objectives, and would most likely be unable to respond to proposals through these two mechanisms of expansion and funding. However, these two approaches would likely play an increasingly large role within CIG after its formative period.

8. Annual CIG allocations and expenditures

Category	Amount Allocated in NSF Award	Expected Expenditure
Senior Personnel	$24,000	$24,000
Other Professional (Technical)	$175,000	$194,645
Secretarial	$8,750	$8,750
Total Salaries	$207,750	$227,395
Fringe Benefits	$55,053	$60,260
Total Salaries & Fringe	$262,803	$287,655
Travel	$12,000	$12,000
Participant costs	$88,000	$88,000
Material & Supplies	$49,852	$25,000
Subaward	$120,558	$120,558
Total Other Direct	$170,410	$145,558
Indirect Costs	$216,787	$216,787
Total Direct & Indirect	$750,000	$750,000

Differences. $24,852 in direct costs are moved from the supplies and expenses category to the salaries and fringe benefit category. Since CIG will remain on campus for the first year and will not be hosting many visitors this frees up money that would have been spent on computers and office equipment and supplies. The money is moved to the technical salary category so that we will have enough salary to have the third software engineer on staff by the end of the year. When we move to year 2, we have support for 3 SE's and so we want to have all three in place by the time the second year starts.

9. Additional funding

none