

ADVANCING SOLID EARTH SYSTEM SCIENCE THROUGH HIGH- PERFORMANCE COMPUTING



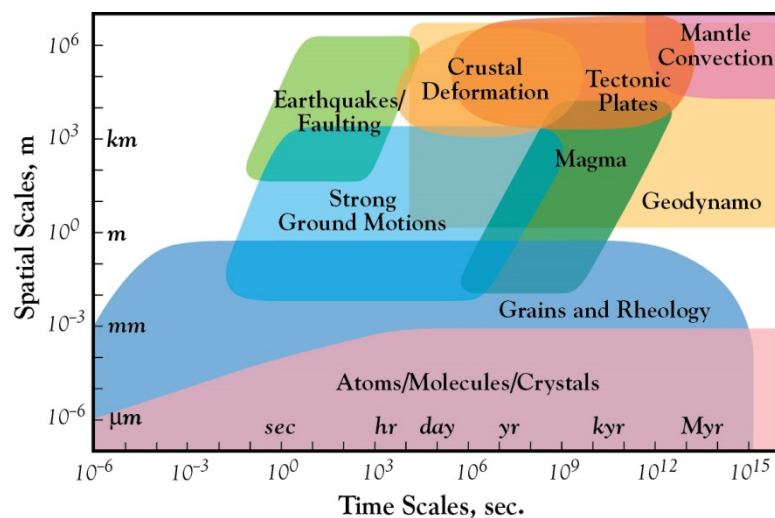
Prepared by (in alphabetical order)

Lorraine Hwang, Tom Jordan, Louise Kellogg, Jeroen Tromp, and Ray Willemann

June 2014

Executive Summary

Solid Earth System Science brings together geophysics, geology, physics, engineering, chemistry, applied mathematics, and computational science to study interconnected systems operating on extreme time and length scales. From the subatomic interactions that control the properties of earth materials at extreme conditions; to global scale coupling of ice, oceans, and the solid earth, from near instantaneous events such as earthquakes and landslides, to processes spanning hundreds of millions of years such as evolution of the Earth's magnetic field, these processes are fundamental to the evolution and dynamics of the Earth and critical to human life, safety, and wellbeing. The scientific research accommodating these scales pushes and challenges the frontiers of numerical and computational methods.



The nation has made a substantial investment to acquire high quality, high resolution, and continuous data related to Earth's structure, composition, geologic events, and geodynamic processes. Partnering with researchers in computational sciences and applied mathematics, the solid earth science community has been working at, and advancing, the cutting edge of numerical and computational technology. Geodynamics and seismology are leading drivers of computational applications and high-performance computing, and these advances are enabling new scientific discoveries. As the fields are poised on new frontiers of computation and science, it has become apparent that progress, and especially transformative science, requires dedicated resources to fully leverage new computational capabilities and new methods simulation and data processing. Dedicated resources provide an opportunity to advance the frontiers of both the solid earth system science and computational science.

The need and opportunity arise because solid earth science is increasingly a system science, in which coupled dynamical systems operating at different scales must be understood together. Modeling this coupling is essential, but is limited by both the readiness of the communities to interface codes and the computational power required to handle the increased amount of data and simulate geological processes over the needed time and spatial scales.

The opportunity arises because seismology and geodynamics have already created many of the computational building blocks needed to transform the way that solid earth system science is done. The field has historically both driven and benefited from advances in technology, and investigators also have a long history of banding together to meet technological and data needs of entire communities (forming consortia and organizations such as IRIS, CIG, SCEC, and UNAVCO). Paused on the cusp of this transformation, we face high hurdles that include: lack of sufficient and sustained access to leadership-class computing for high resolution, long duration simulations and inversions; difficulty sustaining intensive collaborations between solid earth scientists and computational scientists for continuous development of next generation tools; limited access and support services for community-wide use of simulations and inversions requiring mesoscale computational resources.

Understanding these fundamental processes enhances our ability to protect the planet and improves the well-being and economic competitiveness of the nation. High-resolution imaging of Earth's subsurface structure combined with predictions of ground shaking provide information essential to preventing and mitigating damage from earthquakes such the 1994 magnitude 6.7 Northridge, CA earthquake, which resulted in approximately 60 fatalities and caused \$20 billion in damage, and the 2011 magnitude 5.8 Mineral, VA earthquake, which caused more than \$200 million in damage and was felt by more people than any other US earthquake in history. Understanding the dynamics of the crust and upper mantle at plate boundaries is essential to understanding the risks of great earthquakes and associated tsunami in the Pacific Northwest, Alaska, and other critically important locations. Similarly, computational modeling of the Earth's magnetic field and its interaction with space weather is essential to assessing risks of magnetic storms to the power and communication infrastructure; widespread blackouts from grid failures have historically resulted in upwards of \$6 billion in economic loss. Currently, the capacity for the high performance computation required to address societal needs is lacking and is holding back scientific advances in the U.S., limiting global competitiveness.

To achieve this vision requires the dedication of resources both human and physical capital. The computational infrastructure largely exists within key federal facilities but needs to be leveraged, tuned, and devoted to geodynamics and seismology, to rise to the numerous distinctive domain-specific challenges, which include: designing an overall computational environment, selecting appropriate hardware, developing software and algorithms for specialized computing architectures, integrating software into well-designed or well-established workflows, and scheduling resources, and drawing on the scientific capabilities of a large and diverse community of users. Hence, the high performance computing systems that have been created at many supercomputer centers and national laboratories may need to be reconfigured with the rhythm of geodynamics research as a primary design consideration. This challenge will in turn promote innovation in the computational sciences and computer engineering domains.

The computational hardware cannot be exploited without the creation of the human infrastructure. Dedicated expert staff is required, including system administrators, developers of algorithms, software engineers, programmers, and visualization, data, networking and outreach experts, to enable cross-pollination between teams of experts from other domains.

Specifically, the solid Earth system science research community requires:

1. NSF programs to sustain deep collaborations among solid earth scientists and computational scientists focused on extreme-scale computing in solid earth system science.
2. Coherent processes for the allocation of HPC cycles and storage that are coordinated with NSF science funding.
3. Support for software engineering teams dedicated to solving extreme-scale problems in solid-earth system science.
4. Dedicated allocations of computers with sufficient memory and data storage, accommodating both the need for large amounts of mesoscale computing and pathways and access to leadership-class computing.
5. Strategic alliances between NSF, DOE, and other agencies, in extreme-scale computing.

Acknowledgements

This white paper originated from a series of workshops and discussions. Foundational discussions were hosted by Lawrence Livermore National Laboratory on November 8-9, 2012. Further discussion was held at a breakout group on High-Performance Computing at the EarthScope National Meeting, May 13-15, 2013, in Raleigh North Carolina. The role of and need for high-performance computing also emerged from the EarthCube End User Workshop on Modeling, held April 22-23, 2013, in Boulder, Colorado. We would like to thank all who participated in the lively discussions. We benefited from numerous contributions by multiple authors and workshop participants; this document does not necessarily reflect their views.

CIG is supported by the National Science Foundation under Grant No. EAR-0949446. IRIS activities are partially supported by several Cooperative Agreements with the National Science Foundation, including award 1261681, Seismological Facilities for the Advancement of Geoscience and EarthScope. SCEC is supported by Cooperative Agreements with the National Science Foundation and the USGS. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We especially thank the following for their contributions: Wolfgang Bangerth, Greg Beroza, Steve Bohlen, Bruce Buffett, Eric Heien, Alan Levander, Philip Maechling, Sarah Minson, Barbara Romanowicz, Carl Tape.

Table of Contents

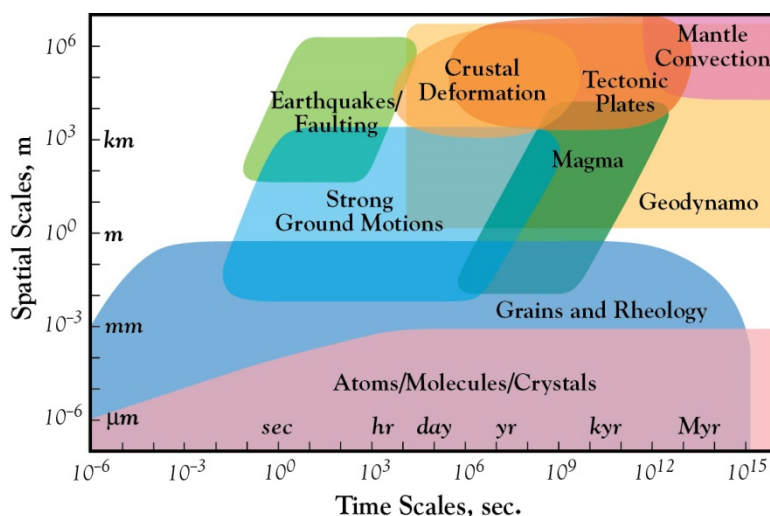
Executive Summary	i
Acknowledgements	iv
1. Introduction	1
2. Scientific Drivers	2
2.1 Seismology	2
2.3 Geodynamics	3
3. Societal Impacts	5
Mesoscale Computing: Advancing Practice among a Broad Research Community	5
Energy and Natural Resources: Near-Subsurface Imaging in the 21st Century.....	7
Seismic Vulnerability: Safety of Nuclear Power Plants and Other Critical Infrastructure.....	7
Real-Time Alerts: Tsunamis, Volcanic Eruptions, Earthquake Sequences, and Strong Shaking	8
Rapid Response: Facilitating Emergency Response and Early Warning During a Crisis.....	9
4. Requirements	9
5. Implementation	11
6. Partnerships	13
7. Conclusions	13
8. References	14
Appendix A. Partnership Opportunities.....	16

1. Introduction

Geodynamics and seismology employ quantitative physical and mathematical methods to study the solid Earth, both to gain a fundamental understanding of the dynamics of the planet on which we live, and to address critical societal problems related to natural resources, natural hazards, and human impact on the landscape. The past decade has seen an explosion in the quantity and quality of observational data available about the Earth, as well as increasing sophistication in our understanding of coupled Earth processes. Geophysicists require increasingly complex numerical methods to understand the solid Earth at

various spatial and temporal scales. During every decade since the 1950s, seismology and geodynamics have been among the scientific fields driving development of computing capability and other information technology. Innovations driven partly by demands from geophysics include digital filtering, fast Fourier transforms and other data processing techniques, numerical simulation methods to test concepts of earthquake rupture, seismic wave propagation, convection in the mantle and the outer core, and other dynamic processes in the Earth's interior, and innovations in data collection transfer, processing and management. In the last decade, at least three computational models in seismology or geodynamics have won or been finalists for the Gordon Bell Prize for outstanding achievement in high-performance computing (HPC) applications.

As sensors become more readily deployed in large numbers in national networks (such as USArray and the Plate Boundary Observatory) and in dense regional networks, and as data telemetry and storage capacities increase, solid earth scientists are faced with a wealth of data. At the same time, advances in both numerical methods and computational technology now make it possible to carry out forward models of geodynamics and seismic wave propagation at unprecedented resolution. The infrastructure available today for collecting, exchanging, and comparing multiple types of geophysical data is advancing rapidly, partly because continued development is the goal of large important initiatives in the National Science Foundation (NSF) such as EarthCube, in the Office of Science and the National Nuclear Security Administration (NNSA) at the Department of Energy (DOE), and in the U.S. Geological Survey (USGS), the Nuclear Regulatory Commission (NRC), National Aeronautics and Space Administration (NASA), and other government agencies.



The study of geophysics spans disciplines to help us understand Earth processes on a subatomic to global scale and from microseconds to hundreds of millions of years. The wide span of scales in space and time are challenging to accommodate using today's computational resources.

But while these initiatives support a wide range of rapid improvements to our understanding of the solid Earth, the available base of cutting-edge computational facilities that are well configured for geophysical research has lagged behind and is deficient, both in comparison to other information technology resources and in comparison to computational resources for Earth science in other advanced countries. The inadequacy of computational facilities accessible for geophysical research in the U.S. is such that only a few individuals who have gained *ad hoc* access to appropriate national computational resources can match research results by foreign communities who are using data collected and managed by the U.S.

2. Scientific Drivers

Seismology

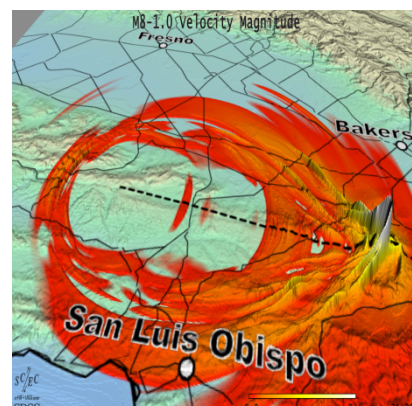
Seismology has a long tradition of using and advancing the development of high-performance computing to study the physics of earthquake processes and their impacts, and to image the Earth's interior. HPC will continue to enable the field to move in essential new directions through dynamic rupture modeling, simulating long-time earthquake catalogs, and modeling complex crustal and deep earth structure and its wave propagation effects – taking full advantage of the wealth of information contained in the waveforms of continuous seismic data streams.

Dynamic Rupture Modeling

Improved computational capabilities have allowed progressively more realistic simulations of spontaneous failure with realistic failure criteria, heterogeneous stress fields, and perturbations from planar fault geometries. Geometric complexity has a strong effect on the evolution of slip and rupture time, which results in the generation high frequency waves that manifest as strong ground motion that causes damage. Progress is being driven partly by evolving techniques, such as adaptive mesh refinement, and a growing community of experienced researchers. With sufficient HPC resources, the next several years could see computational modeling that includes higher-resolution, more detailed dynamic simulations of earthquake rupture on geometrically realistic faults.

Earthquake Simulators

Earthquake simulator algorithms approximate our understanding of earthquake physics and the geometry of fault systems to develop simulated long-term seismicity catalogs to inform efforts in earthquake forecasting. The more closely that the physics and geometry of faulting is included in earthquake simulators, the better the true behavior of the Earth can be modeled and the resulting catalogs be applied to earthquake forecasting. Accounting for full fault physics and fault system geometry remains computationally intractable with current computing capability, so the demand is clear, and significant progress is being made towards identifying critical parameters.



<http://visservices.sdsc.edu/projects/sc/m8/1.0/>

Velocity snapshot from a M8 earthquake illustrating super-shear wave propagation entering the Big Bend section of the San Andreas Fault. After Cui et.al (2010).

Wave Propagation in Complex Structures

For a thorough assessment of damage potential, engineering considerations require a predictive understanding of strong ground motion to high frequencies. This requires a stochastic approach to the representation of Earth structure, and perhaps to wave propagation modeling as well. Plastic yielding and other nonlinearities are important in the extreme near field to bound shaking intensity and to accommodate the large strains that result from slip on geometrically complex faults. Progress is being made using deterministic simulations carried out in randomly varying media to predict the effects of scattered waves on earthquake ground motions.

Mining Large Waveform Data Streams

Cross-correlation based earthquake monitoring has refined our understanding of earthquake sequences and tremors, offering a more complete view of earthquake activity. Near-real-time waveform matching could be particularly valuable during earthquake swarms, volcanic crises, and early in aftershock sequences when information is pouring out of the Earth at a rapid pace overwhelming traditional methods. This method is computationally challenging, scaling with time squared and requires close proximity of continuous waveform data to HPC.

Seismic Imaging

The abundance of high quality data along with advances in numerical methods and growth of HPC capabilities is transforming our view of the Earth. Both controlled source and earthquake source imaging leverage large data sets and full waveform methods to investigate the complex three-dimensional structure of the Earth from local to global scales. Advancements in regional and global seismology and exploration seismology often closely follow one another. Methods such as full waveform or adjoint tomography provide new opportunities for improving images of the Earth's interior and the earthquake rupture process by increasing the amount of usable data. The use of adjoint methods has only recently become feasible with the availability of accurate and fast three-dimensional wave propagation solvers and high-performance computing resources. Tomographic imaging has a broad range of applications from economic, e.g. natural resource development and natural hazard assessments, to basic sciences, e.g. understanding mantle dynamics and related surface processes and national security, e.g. manmade seismic events.

Geodynamics

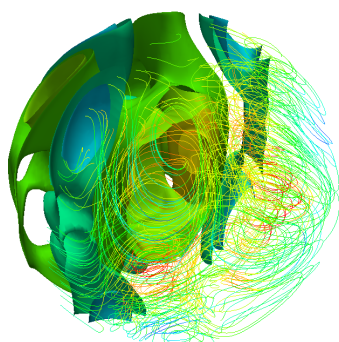
Geodynamics is an interdisciplinary field that brings together a large segment of the geosciences with applied mathematics and computational science. The field has a long history of utilizing and stimulating advances in high-performance computing to image the Earth's interior, to comprehend the role of lithospheric and mantle dynamics in tectonics, and to understand how core dynamics generate the Earth's magnetic field. Quantitative numerical models provide a critical link between physical and chemical observations at Earth's surface and dynamic processes in the interior. Powerful and predictive models have been developed to understand the dynamics at plate boundaries with their attendant mountain building, volcanism and earthquakes; and convection in the Earth's mantle and core. The field is well positioned to take advantage of computational resources necessary to improve resolution of these models giving new insights to an increasing number of important and compelling scientific problems.

Global-Scale and Tectonic-Scale Geodynamics

The convecting solid mantle drives plate tectonics providing the large-scale context for understanding solid-earth dynamics such as long-term tectonics, the dynamics of faulting and seismogenesis and the geodynamo. While fundamental questions remain about each geodynamical system considered in isolation to make problems tractable, new opportunities arise with the capability to simulate interaction amongst them. Mantle convection is coupled to the geodynamo by the exchange of heat, mass and momentum across the core-mantle boundary. The broad scale of convection itself is controlled by the presence of strong, cold plates with extremely weak, narrow boundaries. Crustal dynamics at plate boundaries interacts with large-scale mantle flow and introduces the presence of magmas and volatiles into solids into a system with a wide range of ductile and brittle behaviors.

Recent progress on the key science questions in geodynamics is being achieved by combining geological, geochemical and geophysical observations with quantitative model predictions. National and international efforts have constructed new geophysical observatories that are providing a dramatic increase in the number and quality of data streams that reveal complex transient and episodic deformation of the crust via continuous recordings. Using these data in joint inversions is integral in achieving the ultimate goal of an internally consistent Earth model, but requires even better integration of observations and geodynamic models.

The next generation codes based on adaptive mesh refinement methods are creating opportunities for realistic, observationally constrained modeling of temperature-dependent, three-dimensional mantle flow and its interaction with other earth systems. With appropriate computing resources, simulations could address important questions raised by advances in seismic imaging about lithospheric subduction and delamination and the stability of chemically distinct large deep reservoirs and the structure at the core-mantle boundary.



Visualization of the outer core shell dynamo simulation results from Calypso. Half of the shell is shown as pressure isosurfaces and half as velocity streamlines colored with the vertical magnitude of the magnetic field. In both sections, the columnar convection cells are visible. *Hiroaki Matsui*

Geodynamo

Numerical geodynamo models have become an increasingly important tool in the study of the dynamics and evolution of the Earth's core. The problem is challenging because of strong interactions between the convective flow and the magnetic fields. While the convective flow is ultimately responsible for generating the magnetic field, the magnetic field exerts a substantial influence on the structure of the convective flow.

Dramatic improvements in computational capabilities have driven rapid progress in the scientific questions that can be addressed using numerical simulations. Numerical geodynamo models have succeeded in produced Earth-like fields with a strong dipole component. Many of these models also exhibit spontaneous polarity reversals. A more recent trend has been the integration of geodynamo and mantle convection models to investigate how changes in the magnitude and spatial pattern of heat flow across the core-mantle boundary can change the frequency

of magnetic reversals. Because this heat flow is mainly controlled by mantle convection, paleomagnetic observations of changes in the reversal rate may offer unexpected insights into the dynamics of the mantle.

The low viscosity of liquid iron permits flow across a vast range of length scales representing a challenge to current computing capabilities. The combined effect of planetary rotation and a strong magnetic field can substantially alter the structure of small-scale turbulence. Reducing the magnitude of the viscous force in the core to 10^{-8} maybe sufficiently close to the expected magnitude of 10^{-15} such that viscosity becomes unimportant to the dynamics of convection in the core. However, with reduction in viscosity the numerical burden quickly becomes prohibitive. Each simulation requires one million time steps or more. A few million time steps at a force ratio of 10^{-8} can be completed approximately over several weeks with good scaling on 10^5 computer cores. The development of pseudo-spectral techniques may demonstrate both the scaling and resolution needed to run on HPC and make transformative advances in understanding the geodynamo.

3. Societal Impacts

In addition to enabling progress on key scientific questions such as those outlined above, advancing access to high-performance computing is also critically important to the broader goals of the scientific community, including technology and workforce development. This section outlines these broader impacts in more detail.

Mesoscale Computing: Advancing Practice among a Broad Research Community

Earth science programs often begin with a project to describe a given process or resolve a specific structure and evolve into an attempt to understand how that process or structure interacts with a system of other phenomena. As the complexity of a system and the number of non-linearly interacting processes grow, so do the requirements for computational resources. Simulations of the complete system grow more challenging, of course, but at the same time there is an exponential increase in the number of subsystem simulations required to establish a sufficient understanding of individual interactions.

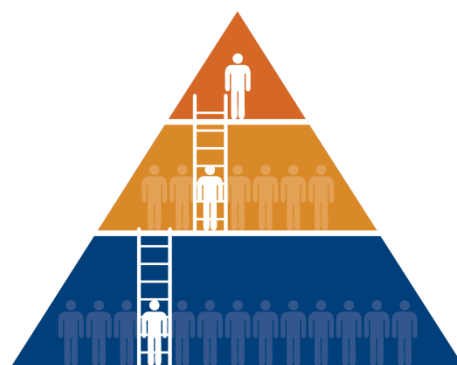
Despite this, only a very small proportion of researchers in seismology and geodynamics are making use of even moderately large national facilities for HPC. Reaching the top level of leadership class computing for a researcher requires a highly specialized skill set that is not widespread among solid earth scientists, and the gap between potential and actual achievements is growing wider. That is, the rate of data acquisition continues to increase rapidly so that sequential computing is progressively more inadequate, while the scale of cutting edge computing continues to enlarge so that “entry requirements” to ensure best use of HPC resources is becoming greater. HPC resources committed to seismological and geodynamics research can bridge this widening gap, with a strategic impact that extends far beyond the users currently ready to exploit leadership class capabilities.

Seismologists have long shared source codes, and today seismology and geodesy are ahead of other scientific fields in curating and unrestricted sharing of their “crown jewels”, data and data products.

Geodynamicists today are ahead of other fields in developing community software based on open source principles and optimized for HPC. Nevertheless, failure to bridge the gap between potential and achieved results from use of HPC is increasingly problematic as new techniques are developed to extract more information from the relentlessly growing volume and diversity of data. Indeed, a sense of mutual frustration has arisen, with domain scientists facing high hurdles to begin using massively parallel systems and new hardware such as GPUs, while HPC professionals are reluctant to allocate computing time and devote personnel to research groups that are not prepared to make optimal use of these systems.

With NSF support, the geophysical community already undertakes a wide variety of loosely-linked activities to bring about a strategic change. These activities include implementation of community software, development of standardized formats and exchange protocols for data and products, creation of curated open-access data systems, and ongoing professional education comprising workshops, short courses, and advanced studies institutes. The organizations that carry out this work – usually emphasizing selected activities but without neglecting the others include the EarthScope National Office (ENO), IRIS, the Computational Infrastructure for Geodynamics (CIG), Cooperative Institute for Deep Earth Research (CIDER), the Southern California Earthquake Center (SCEC), and UNAVCO. During each of the past four years, for example, CIDER summer programs featured presentations and tutorials on CIG resources aimed at preparing a diverse community of early-career scientists to access and make use of HPC resources. Tutorials on how to use and access HPC facilities are currently offered by national facilities such as the Extreme Science and Engineering Discovery Environment (XSEDE), but partnerships with SCEC, CIDER, CIG, IRIS, UNAVCO, and the EarthScope National Office (ENO) could adapt these tutorials for Earth scientists and bring them to venues that would facilitate broader participation.

Collectively, SCEC, CIDER, CIG, IRIS, UNAVCO, and ENO have laid much of the groundwork to build a metaphorical ladder for climbing the HPC pyramid that will improve the scientific return from the broadly-defined research community. The professional programming staff of CIG, an NSF geoinformatics facility, could expand to adapt and optimize selected existing codes to run on identified HPC resources. With new resources, IRIS's recently established mirror of seismographic data could be expanded to include InSAR, LIDAR, and UNAVCO GPS data, and be reorganized for optimal use in HPC workflows. Better sharing results from simulations of the geodynamo and mantle convection, and of deformational and structural models of the Earth's surface and interior from data inversions could facilitate their use as tools in further research and assimilation of further data types and more recently acquired data.



The ladder to leadership class computing begins with everyday research computing on simple desktop computers, to middle tier clusters, and ultimately to the largest computers in the world.

The envisioned infrastructure creates the possibility of computing environments on remote workstations with readily usable interfaces for configuring workflows for efficient execution, for post processing and visualization and for disseminating important model results to researchers in different disciplines of the geosciences. The transformative potential is analogous – but far broader – to portable

seismology 30 years ago, when IRIS programs opened access broadband digital instrumentation, moving the technology from the purview of a few researchers at select institutions to the norm across the field.

Energy and Natural Resources: Near-Subsurface Imaging in the 21st Century

Energy production increasingly involves active reservoir management to extract hydrocarbons and geothermal energy from non-traditional formations and mature fields. Injection programs both for energy production and waste disposal have been associated with induced seismicity. The phenomenon has increased the overall rate of seismicity in the central and eastern U.S. during the 1990's around identified injection locales, but remains poorly understood [Ellsworth *et al.*, 2012]. In an energy-constrained world and one in which carbon sequestration may become an important strategy to combat climate change, understanding the linkages between crustal structure, stress, fluid movement, and extraction and injection methodologies will contribute to hazard mitigation allowing these resources to be developed fully and safely.

In recent years, the petroleum industry has implemented new techniques for reservoir imaging and resource exploration that use orders of magnitude more seismic data than were collected previously. The industry's new approaches to data processing are based partly on fundamental research at universities and, with access to sufficient computing resources, academic seismologists could build on industry practices and adapt them for subsurface imaging to achieve benefits for broader society, such as:

- Detailed mapping across the nation of structural units which might hold tight gas, partly to make provisional plans for infrastructure development and to begin developing environmental and other regulations before leasing, intensive lateral drilling, and hydrofracturing.
- Mapping porous units and impermeable cap rocks with the potential to be used for carbon sequestration, monitoring for migration of fluids injected in the course of sequestration projects, and detection of risks to reservoir integrity and associated hazards;
- Imaging hydrothermal systems that might be exploited for geothermal power;
- Aquifer imaging for both potable water and energy production;
- Large-scale subsurface mapping to discover formations with the potential to be mapped in detail by industry to develop reserves of other natural resources, such as rare earth elements.

Seismic Vulnerability: Safety of Nuclear Power Plants and Other Critical Infrastructure

Data collection efforts that contribute to the National Earthquake Hazard Reduction Program (NEHRP) are advancing: the USGS's Advanced National Seismic System (ANSS) is growing and the NSF's EarthScope Project has created an opportunity to collect complementary seismic and geodetic data across the U.S. Additional computing resources would enable researchers to take better advantage of the disaster mitigation potential of these data. Massive cross-correlation of multiple data streams with templates from past earthquakes can discover the extent to which faults display "repeating earthquakes" and document evolution of those events, elucidating the subtle changes in the states of fault systems [Schaff and Richards, 2011]. HPC-enabled adjoint tomography of ANSS, EarthScope, and other data could extract additional information about how attenuation varies with frequency, location,

and propagation direction. Numerical simulations of rupture processes and of focusing of nonlinear strong motion by basins and other subsurface structures could guide prioritization of mitigation efforts and siting of critical infrastructure [Aagaard *et al.*, 2008]. Computationally-demanding multi-disciplinary analyses such as hybrid simulation techniques can link seismological results with civil engineering results from NSF's Network for Earthquake Engineering Simulations (NEES).

NEHRP has made progress in quantifying seismic hazard and facilitating mitigation, but reports from the National Research Council (NRC) outline both the benefits that could accrue from more complete data collection and analysis [NRC, 2006] and the decades of work that are necessary to achieve true resilience [NRC, 2011]. Since the magnitude 9.0 Tohoku-Oki earthquake during March 2011, prompt impacts and ongoing challenges from disablement of nuclear power plants in Japan have been recurring reminders of the direct and cascading effects that can result from damage to critical infrastructure. The magnitude 5.8 earthquake near Mineral, Virginia, during August 2011 – and the temporary shutdown of the nearby North Anna nuclear power plant – demonstrated anew that it would be perilous to neglect earthquake hazard in the central and eastern U.S. as the nation considers renewed growth in the use of nuclear power, in many cases from plants proposed in regions of mapped seismic hazard [Andrews and Folger, 2012].

Real-Time Alerts: Tsunamis, Volcanic Eruptions, Earthquake Sequences, and Strong Shaking

A geodynamics-dedicated HPC facility would have a clear impact on tsunami early warning. The improvements to tsunami early warning would come not only from having additional computational resources to quickly produce tsunami forecasts but also from a capability to rapidly produce high-quality slip models, a key input into issuing accurate tsunami warnings.

Earth scientists often deploy dense arrays of seismometers and GPS receivers during a swarm of earthquakes or diffuse tremor activity, when a volcanic eruption may be imminent, or to record aftershocks after a large earthquake – sometimes utilizing IRIS's Rapid Array Mobilization Pool and/or supported by a RAPID award from NSF. Especially when rapid deployments are executed as community projects, use of shared HPC facilities to compute and distribute advanced products quickly would offer opportunities for benefits to broader society, extending to early warning of an eruption or strong shaking from an aftershock [Brown *et al.*, 2011].

The sequence of earthquakes near L'Aquila, Italy, during 2009 culminated in a magnitude 6.3 event that killed more than 300 people and displaced more than 65,000 residents. Retrospective analysis of the Italian government's actions as the sequence proceeded, before the main shock, demonstrates that subjective expert judgment no longer represents the state of the science for making decisions during a seismic crisis [Jordan *et al.*, 2011]. Instead, operational earthquake forecasting requires advanced configuration of computing facilities and prompt reallocation of resources for systems analysis [Barbot *et al.*, 2012], monitoring for changes in fault structure [Zaccarelli *et al.*, 2011], and scenario modeling to facilitate the best informed response.

HPC facilities could potentially also be useful for earthquake early warning, but this would require such substantial infrastructure development that it is difficult at this time to assess the potential impact of additional HPC resources on earthquake early warning.

Rapid Response: Facilitating Emergency Response and Early Warning During a Crisis

Even after a large damaging earthquake, prompt access to HPC facilities can benefit broader society, now by improving emergency response. Mapping specific rupture features – Where on the fault did rupture propagate at the highest speed? Where did the two sides slip furthest past each other? Where are the asperities that remained locked or slipped in an unexpected direction? – continues to be executed as a research effort, with even first results typically not available for several weeks.

Such rapid results could be useful for decision-making during emergency response and the earliest stages of recovery. For example, rupture propagation can direct and concentrate vibrational energy in a series of strong pulses that locally cause much greater damage than other locations. Using this knowledge in combination with other data can, as seen in recent earthquakes in Japan, produce detailed strong shaking maps to successfully direct responders where they are most needed. During recovery, knowing where patches of extreme slip are concentrating stress on adjacent fault segments identifies where the risk of more frequent or larger aftershocks is greater aiding in the decisions whether to reoccupy or evacuated a building.

Seismological results from earthquakes are often created as research products some time after the initial emergency response phase. To enable faster results, a researcher must not only have relevant experience but must also gain access to an HPC facility with time available, then resources must be requested, codes installed, and data - seismological, GPS, InSAR, LIDAR, and geological, organized. An HPC facility committed even partially to seismology and geodynamics could have the flexibility and nimbleness to aid emergency response, by pre-installing codes, enabling fast access to real-time data streams and by proactively partnering with researchers to re-prioritize compute cycles promptly after a significant earthquake.

4. Requirements

Specifically, the seismology and geodynamics research communities needs at present include:

1. High performance computing hardware dedicated to the geodynamics community.

The solid earth scientific community is currently effectively able to utilize computer clusters with thousands of nodes (tens of thousands of cores) with fast network, mixes of fat and thin nodes, for hundreds of core hours per year, and multiple petabytes of storage for data and simulations. If everyone in the scientific community had access to these capabilities, they would be fully utilized now. And these needs are expected to rapidly expand:

On a regional scale [e.g. Zhu and Tromp, 2013], a typical nonlinear iteration in an inverse problem requires thousands of distinct simulations, each harnessing several hundred compute cores, and all of which must be completed to complete a single iteration. Moreover, a typical inverse problem requires tens of iterations to achieve a scientifically valid result. The community

currently uses hundreds of millions of core compute hours annually and that within five years demand is likely to approach one billion core hours.

Full mantle convection simulations accurate to the scale of kilometers will require 10^7 to 10^9 degrees of freedom to accurately resolve temperature, pressure, and velocity – and must run for 10^3 to 10^5 steps to understand long-term dynamics. Simulations would require 10^4 cores for lengths of several days.

Accurately simulating conditions that advance our understanding of the geodynamo – at an Ekman number of 10^{-8} or smaller – require several weeks of computing time on 10^5 cores.

Thus, in order to support the variety and scale of computation required for advanced research, computing resources must include on the order of one million cores at the scale of existing NSF “Track 2” machines, each with several gigabytes of fast memory, and fast interconnect to support efficient communications between parallel processes. Moreover, an effective system would require several hundred terabytes of fast scratch disk and non-purged user storage, and backup storage on the order of a petabyte.

2. *Experts whose time is dedicated to supporting the hardware and to support access and usage by the scientific community.*

Dedicated expert staff includes system administrators, programmers, visualization expertise, data and networking expertise, and a staff assistant to coordinate interactions with existing community organizations, such as CIG, CIDER, IRIS, and SCEC.

3. *Floor & office space co-located in an existing HPC facility.*

A physical presence among HPC scientists in the scientific domains of Earth science is required to enable collaboration and consultation.

4. *Support for education, training, and cross-disciplinary collaboration, including online and in-person workshops, tutorials, and code manuals.*

A common theme in recent reports on use of computing resources in Earth science [OITI 2002, CyDras 2004, Cohen 2005, Petascale 2006; EarthCube End User Modeling Workshop Report 2013] is that the human component must be developed along with the deployment of technology. Particular areas that must be addressed include development of Earth science computation as a scientific field; the need for collaborative interactions between Earth scientists and specialists in mathematics, statistics, and computer science; and the development of a support structure for use of high-performance computing by the geodynamics community.

To support the envisioned research effort, the system would require high bandwidth access from external academic and computing centers, well-integrated visualization and graphics capabilities, and support of community software and scientific workflows. Such a system could not only accommodate production simulations focused on solving a wide breadth of geophysical forward and inverse problems, but also support code development and testing. If made widely available, the system could serve the computational needs of faculty, students and postdoctoral researchers in geodynamics.

5. Implementation

Many current directions in development of cyberinfrastructure for “big data” applications are not relevant or are insufficient for the requirements of solid earth system science research. Instead of using statistical methods to discover features inherent in the data, as is typical of “big data” mining, the goal of solid earth science is to illuminate structure and dynamics that exist independently of any single dataset. To accomplish this, a solid earth system science application can involve both forward numerical simulations of physical processes and inversion of data, as well as comparison of predictions from the simulations with features of the data to further refine models. Thus, the workflow for solid earth system science applications often involves intervals of intensive data access or interpretation interspersed with periods of intensive computation. As a result, there are numerous distinctive domain-specific challenges in designing an overall computational environment, selecting appropriate hardware, writing programs for specialized computing architectures, integrating programs into well-designed workflows, and scheduling resources to make efficient use of both the computational resources and the scientific capabilities of a large community of users.

With a national commitment to develop exascale computing, the lack of access to advanced computing for Earth science does not stem from an absence of facilities altogether. Rather, there is a need to establish and sustain strong connections between solid earth scientists and the community that develops and operates computing facilities. The hardware with vastly greater capabilities in numerical processing, data storage, and data transfer have been created at many supercomputer centers and national laboratories may need to be reconfigured with the rhythm of Earth science research as a design consideration. Some members of the teams of software engineers and computational scientists who work with scientists from different fields of study to facilitate successful use of high-performance computing may need to commit sustained effort to learn the distinct requirements of Earth science research.

Efficient use of shared computing would require adoption of community data and software, a step that the Earth science community is already pursuing and unquestionably ready to practice more comprehensively. Over many years, Earth scientists have embraced progressively better technical standards and community norms for open data exchange. From the creation of the International Seismological Centre in 1970 to the operations of today’s organizations such as IRIS, UNAVCO and Observatories and Research Facilities for European Seismology (ORFEUS) Data Management Centers which use international standards for formats and protocols to exchange data, the distributed nature of Earth science has impelled sharing of data, data products, and software. Even before invention of the World Wide Web, Earth scientists used journals and books, computer disks and tapes, and ftp sites to freely distribute source code and data. Better access to the internet and the prevalence of open source communities, has promoted a culture of sharing. The concept of trusted, well-documented and tested code is the next step in this evolution and is actively practice at organizations such as CIG. Wider adoption is limited by the financial resources available to support and maintain expert technical staff.

Specific implementation steps that the community is ready to pursue or is already pursuing include:

Creating software that is open, extensible, portable and scalable: CIG has made significant strides towards providing software that builds on widely used libraries for most components which has significantly affected the development of these libraries. As a result, the geodynamics and seismology community are advanced compared to many other scientific domains in sharing and co-developing openly available codes. Modifying CIG codes to scale to tens or hundreds of thousands of cores is achievable in collaboration with computational scientists intimately familiar with the details of large machines and experience in writing and analyzing codes at these scales. This expertise is available at several NSF and DOE facilities, and collaboration with the Earth science community could be facilitated by co-location of a facility at one of these centers.

Data and model formats: Seismologists, geodesists, and geodynamicists have been progressively embracing use of well-documented formats for exchange, archiving, and analysis. Together with other Earth scientists, these communities recognize the advantages of new approaches, such as self-describing formats, that are more robust and better enable interdisciplinary work. A community-spanning initiative would facilitate new standards within and between disciplines and would coordinate with practitioners worldwide, and across different levels of data products extending to earth models based on data from several disciplines.

Data flows: Many workflows in geodynamics Science are complex, requiring coordinated scheduling and execution of the components and management of the flow of very large data sets. Moving data, both observational data and simulation results, has become an urgent challenge to scalability, requiring explicit management. Earth scientists are beginning to learn to address both challenges by more tightly integrating the stages and learning from astrophysical surveys and other communities about modeling the performance of alternative data flows and more use of specialized middleware.

Graphics Processing Units: Visualization will benefit from use of GPUs, which should be included among the resources available for Earth science. In addition, ongoing projects are adapting common software libraries to better support GPUs for some types of Earth science applications with highly structured calculations. GPUs may not be suitable for all geoscience problems, but promising approaches are being considered for several applications and availability of a modest system would motivate and facilitate re-design of some existing scientific software.

Next generation software and algorithms: The demanding requirements of computational models in the geosciences have been a major driver in the development of all aspects of scientific computation. Computational research in geodynamics require combining stochastic and deterministic models, improving methods of model validation and verification, quantification of uncertainty, developing inverse methods and techniques for the identification of extreme events and critical transitions, and formulating novel numerical algorithms and implementations, along with the greatly enhanced use of data from the rapidly evolving observing systems. These have become important research problems in mathematics, statistics, and computer science. Optimal algorithms become more critical as larger problems are being solved on larger computers.

Continued advances require support for developing applicable mathematical and numerical methodologies across fields of geoscience. New methods require research by applied mathematicians, computational scientists, and statisticians, among others, that are motivated by geoscience problems.

Coherent processes for the allocation of HPC resources: Currently, processes for allocation of HPC resources occur through a complex and uncoordinated series of proposals and negotiations. Research funding is allocated through NSF programs in GEO and CISE; computing cycles are allocated through NSF XSEDE & PRAC, and DOE INCITE programs; run-time storage either goes with the computing cycle allocation or, for extended storage, is allocated by negotiation with HPC centers. Allocation of the time of software engineers is done through formal proposals (e.g. XSEDE) and informal negotiations with HPC centers. Better coordination of these allocation processes will support more effective use of extreme computing.

6. Partnerships

Advances in fundamental research and its application for the benefit of society in the Earth sciences extend across the missions of several agencies of the Federal government (see Appendix A). Partly because of this diversity of goals, the Earth sciences could leverage a variety of existing HPC resources as well as build on current trends in programs under development found in different Federal agencies. Partnerships can be an effective way to reach Earth science HPC goals by utilizing established facilities and expertise without substantial commitment of new resources.

7. Conclusions

A truism held among many Earth scientists –seismologists, geodesists, and geophysicists – is that ours is a “data driven” field. Today, we are limited not only by our ability to collect and exchange rapidly growing sets of data that are shared globally, but also by our ability to discern all of the information about the Earth that is latent in these data through our ability to model its dynamical processes. We are cyclically driven to better understand our data by creating better models, which in turn drives the collection of more and often higher quality data. The computing technology exists today to enable discovery of vastly more knowledge about the Earth, and to use that knowledge to benefit society. However, advancements have been slowed by the lack of access to the necessary computational resources. Promoting such a resource requires training and maintaining the human capital and a management structure to insure robust and balanced operational practices. In return, geodynamics will continue to help society discover and utilize new resources, mitigate the effects of our changing natural world, and advance basic science.

8. References

- Aagaard, B., T. Brocher, D. Dreger, A. Frankel, R. Graves, S. Harmsen, S. Larsen, K. McCandless, S. Nilsson, N. A. Petersson, A. Rodgers, B. Sjogreen and M. L. Zoback, (2008). Ground motion modeling of the 1906 San Francisco earthquake II: Ground motion estimates for the 1906 earthquake and scenario events, *Bull. Seismol. Soc. Amer.*, *98*, 1012-1046.
- Andrews, A., and P. Folger, 2012. Nuclear power plant design and seismic safety considerations, *Congressional Research Service*, R41805.
- Barbot, S., N. Lapusta, and J.-P. Avouac (2012). Toward predictive modeling of the seismic cycle, *Science*, *336*, 707–710.
- Brown, H.M., R.M. Allen, M. Hellweg, O. Khainovski, D. Neuhauser, and A. Souf (2011) Development of the ElarmS methodology for earthquake early warning: Realtime application in California and offline testing in Japan, *Soil Dyn. Earthquake Eng.*, *31*, 188-200, doi:10.1016/j.soildyn.2010.03.008
- Cohen (2005). High-Performance Computing Requirements for the Computational Earth Sciences Computational Infrastructure for Geodynamics (2013). A membership-governed organization that supports and promotes Earth science by developing and maintaining software for computational geophysics and related fields, supported by the National Science Foundation under award EAR-0949446, <http://www.geodynamics.org/>.
- Cui, Y., Olsen, K. B., Jordan, T. H., Lee, K., Zhou, J., Small, P., Roten, D., Ely, G., Panda, D. K., Chourasia, A., Levesque, J., Day, S. M., and Maechling, P. (2010) Scalable Earthquake Simulation on Petascale Supercomputers. In Proceedings of the 2010 ACM/IEEE International Conference for High-Performance Computing, Networking, Storage and Analysis, doi=10.1109/SC.2010.45 (SC10 Gordon Bell Finalist).
- CyRDAS (2004). Research and Development in the Atmospheric Sciences in the 21st Century. A Report from the Ad Hoc Committee for Cyberinfrastructure, Research, Development and Education in the Atmospheric Sciences. (CyRDAS). <http://www.cyrdas.org/>
- EarthCube Modeling Workshop for the Geosciences (2013) Executive Summary, <http://geodynamics.org/cig/community/workshops/Earthcube13/ExecSummary>, accessed September 30, 2013.
- Ellsworth, W.L., S.H. Hickman, A.L. Lleons, A. McGarr, A.J. Michael, and J.L. Rubinstein (2012). Are seismicity rate changes in the midcontinent natural or manmade?, *Seismol. Res. Lett.*, *83*, 416 (abstract).
- Jordan, T.H., Y.-T. Chen, P. Gasparini, R. Madariaga, I. Main, W. Marzocchi, G. Papadopoulos, G. Sobolev, K. Yamaoka, and J. Zschau (2011). Operational Earthquake Forecasting: State of Knowledge and Guidelines for Utilization, *Annals Geophys.*, *54*, 315–391, doi:10.4401/ag-5350
- National Research Council (2006). *Improved Seismic Monitoring – Improved Decision Making*, ISBN 978-0-309-09695-9, National Academies Press, 196 pp.

- National Research Council (2011). *National Research Resilience: Research, Implementation, and Outreach*, ISBN 978-0-309-18677-3, National Academies Press, 244 pp.
- OITI (2002). Information Technology Infrastructure Plan to Advance Ocean Sciences. Ocean Information Technology Infrastructure Steering Committee (OITI). <http://www.geo-prose.com/oiti/index.html>
- Petascale (2005 Technical Working Group and Ad Hoc Committee for a Petascale Collaboratory for the Geosciences. 2005. *Establishing a Petascale Collaboratory for the Geosciences: Technical and Budgetary Prospectus*. A Report to the Geosciences Community, UCAR/JOSS. 56 pp.
- Schaff, D.P., and P.G. Richards (2011). On finding and using repeating seismic events in and near China, *J. Geophys. Res.*, *116*, B03309, doi:10.1029/2010JB007895.
- Simons, M., S. E. Minson, A. Sladen, F. Ortega, J. Jiang, S. E. Owen, L. Meng, J-P. Ampuero, S. Wei, R. Chu, D. V. Helmberger, H. Kanamori, E. Hetland, A. W. Moore, and F. H. Webb (2011), The 2011 Magnitude 9.0 Tohoku-Oki Earthquake: Mosaicking the Megathrust from Seconds to Centuries, *Science*, *332*, 6036, 1421-1425, doi:10.1126/science.1206731.
- Zaccarelli, L., N. M. Shapiro, L. Faenza, G. Soldati, and A. Michelini (2011). Variations of crustal elastic properties during the 2009 L'Aquila earthquake inferred from cross-correlations of ambient seismic noise, *Geophys. Res. Lett.*, *38*, L24304.

Appendix A. Partnership Opportunities

National Science Foundation (NSF)

The NSF's Division of Earth Sciences (NSF/EAR) has a mission congruent with many goals of better use of HPC in the Earth sciences. Nevertheless, HPC needs are probably best addressed in the NSF at the level of the Geosciences Directorate (NSF/GEO) or agency-wide. The GEO Directorate already leads or participates in several current initiatives, including the NCAR Yellowstone Wyoming Supercomputer Center, EarthCube (with OCI), and Cyberinfrastructure for the 21st Century (CIF21) (agency-wide).

EarthCube, in particular, illustrates how tying HPC for Earth science with existing initiatives can be beneficial. Among other things, EarthCube aims to “lower adoption thresholds” which is aligned with the need to use meso-scale computing to broaden the EAR investigator community familiar with HPC.

Opportunities to partner with EarthCube-funded projects include:

- Ensuring fast access to data from several Geoscience disciplines at an HPC facility, and
- Creating broader impacts by enlarging the investigator community that uses HPC effectively through, for example, the development of user interfaces that enable configuration of customized workflows using pre-installed codes and datasets.

Note that none of the aforementioned initiatives successfully addresses the issue of HPC in Earth sciences. The Wyoming Supercomputer Center has proved in practice to be fully committed to the atmospheric sciences. The hardware configuration for the Yellowstone computer in Wyoming was driven by requirements of NCAR software – no Earth science data sets are mirrored at the site; computer science staff is not attuned to the needs of Earth science researcher; redirection of resources to support rapid response to geophysical disasters is not considered. The EarthCube Program has a complementary overall objective to create transformative infrastructure for integrated data management that is distinctly different from the goal of developing HPC resources.

Directorate-level leadership is important at NSF due to the linkages between GEO and other Directorates. Linkage with the Engineering Directorate in the use of HPC, for example, could create partnerships in the use of model results in strong motion research with the NEES program. A past partnership with the MPS directorate supported research between geosciences and mathematical sciences including work aimed at developing better algorithms. Broader and deeper opportunities could be explored partnering with the Computer and Information Science and Engineering Directorate (NSF/CISE). The relevance of HPC for Earth science to CISE strategic goals is demonstrated, for example, by the 2011 report on HPC from a task force organized by the Advisory Committee to the Office of Cyberinfrastructure, which includes a recommendation:

“By 2015–2016, NSF should ensure that academic researchers have access to a rich mix of HPC systems ... to include NSF investment in a few systems delivering sustained performance in the 10–50 petaflops range ... and at least one system capable of exceeding 250 petaflops ... [B]ecause the scale, resolution, and fidelity of scientific simulations are limited by the configuration and balance of HPC systems, it is crucial for NSF to engage the scientific and engineering community in shaping the solicitations for these next generation systems.”

In support of this strategic recommendation, CISE's Division of Advanced Cyberinfrastructure (ACI) has several programs that would contribute to development of HPC resources for the Earth Sciences. For example, as of November 2012 UIUC press report that the ACI petaflop computer Blue Waters is in production mode. SPECfem3D, a CIG code, is one of the benchmark codes used to measure its performance.

CISE could participate in a variety of ways in addition to creating HPC systems better configured and operated for Earth sciences research. For example, libraries for new algorithms that run more efficiently on massively parallel systems might be developed in partnership with the Exploiting Parallelism and Scalability program (XPS), which is administered jointly by the divisions of CISE. In recent solicitations the focus areas include "Domain-Specific Design", where the topics in which work is sought include "tools, frameworks, and libraries that support the development of domain-specific solutions to computational problems and are integrated with domain science."

National Nuclear Security Administration (NNSA)

In managing seismological and other research to improve US capability in monitoring nuclear tests in its Nonproliferation Research and Development programs, the NNSA can both contribute to and benefit from better use of HPC in the Earth sciences. Driven by concern about smaller yields and evasion in testing programs, NNSA has shifted its research emphasis from teleseismic methods to regional and local monitoring. Since monitoring seismic wave at higher frequencies closer to the source produces a more accurate picture of the source properties, cutting edge research in these bandwidths is a central part of the NNSA's research. As exemplified by the Source Physics Experiment (SPE), NNSA shares a need with earthquake science to go beyond elastic wave propagation and understand how nonlinear processes within an explosion or shear failure generate seismic waves.

SPE, like NEHRP, is a multi-year project that entails development of new infrastructure and expertise to reach its goals. SPE and other NNSA test monitoring research comprises a wide ranges of activities, including numerical simulations enabled by the use of cutting-edge HPC resources at Lawrence Livermore National Laboratory and Los Alamos National Laboratory. Partnerships between the NNSA, the national laboratories that it manages and with the broader Earth sciences effort could be mutually beneficial. Policy makers and government officials should be informed by a wide range of knowledgeable and invested professionals about synergistic benefits to national security and national resiliency from a commitment of HPC resources to Earth science research that such partnerships create. In addition,:

- With applications that offer wider range of benefits to broader society, HPC systems at NNSA-managed national laboratories will be more generally supported, more reliably funded, and more regularly improved to take advantage of technological advances.
- Computer science professionals at the NNSA-managed National Laboratories with expertise in developing computer architectures, algorithms, and work flows for massively parallel systems, will further hone their skills through work in different but related area in Earth science research,
- Seismological and other investigators at NNSA-managed National Laboratories and in the academic research community will work more efficiently and benefit more quickly from knowledge transfer through shared resources.

Department of Energy Office of Science (DOE/SC)

Basic Energy Sciences (BES) supports fundamental research to provide the foundations for new energy technologies and to support DOE missions in energy, environment, and national security. The BES program also plans, constructs, and operates major scientific user facilities to serve researchers from universities, national laboratories, and private institutions. The Geosciences research area supports basic experimental and theoretical research in geochemistry and geophysics. Geophysical research focuses on new approaches to understand the subsurface physical properties of fluids, rocks, and minerals and develops techniques for determining such properties at a distance; it seeks fundamental understanding of wave propagation physics in complex media and the fluid dynamics of complex fluids through porous and fractured subsurface rock units. The research area also emphasizes incorporating physical and chemical understanding of geological processes into multiscale computational modeling. BES capital funding is provided for equipment that includes augmenting computational capabilities.

Advanced Science Computing Research Program (ASCR) includes the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, which promotes transformational advances through large allocations of computer time and supporting resources of run time on HPC systems at Oak Ridge National Lab and Argonne National Lab. For projects with core funding source from DOE, NSF, and other agencies, INCITE supports projects that would not be possible or productive without petascale computing. Awards have supported research in earthquakes, carbon capture and storage, and climate by US- and non-US-based researchers at universities, national laboratories, and industry.

The Scientific Discovery Through Advanced Computing (SciDAC) program spans several parts of the Office of Science, including ASCR. Within the category Earth Systems Science, one of the two main research thrusts is climate and environmental sciences, including a goal to provide improved models for better understanding the movement of subsurface contamination. In addition, the original SciDAC program plan recognized that “Analysis of impacts of earthquakes [is] a complex problem of basic and strategic importance”. Aligning with SciDAC could be through partnering with selected current SciDAC Institutes, such as Frameworks, Algorithms, and Scalable Technologies for Mathematics (FASTMath), Quantification of Uncertainty in Extreme Scale Computations (QUEST), Institute for Sustained Performance, Energy and Resilience (SUPER), and Scalable Data Management, Analysis and Visualization (SDAV)

Department of Energy Divisions of Renewable Energy and Fossil Energy

Even outside of NNSA and DOE/SC, several Department of Energy activities intersect in obvious ways in which HPC can be better used to map structure and monitor temporal change in the Earth’s crust. Within the Energy Efficiency and Renewable Energy Division (DOE/EERE), the Geothermal Technologies Program pursues research and development to facilitate technology validation and deployment, reduce cost, and improve performance. Within the Fossil Energy Division (DOE/FE), the Carbon Capture and Storage program is using Regional Carbon Sequestration Partnerships to implement large-scale geologic storage and utilization projects that are intended to demonstrate the long-term, effective, and safe storage and utilization of carbon dioxide in geologic formations throughout the United States. Also

within DOE/FE, the Office of Oil and Natural Gas supports research and policy options to ensure environmentally sustainable domestic and global supplies of oil and natural gas.

Air Force Research Laboratory (AFRL)

Air Force Research Laboratory's Space Vehicles Directorate sponsors research at universities and private contractors in research related to seismic detection of underground nuclear tests. During FY 2013 solicitations were issued in several areas that can and does benefit from HPC resources:

Seismic Source Physics, understanding the properties of small seismic events, including the effect of depth, near source topography, and media properties.

Velocity Models, new techniques of determining three-dimensional, spatially variable velocity, fitting multiple data types, uncertainties and trade-offs.

Signal Analysis, includes detection and characterization of sources by waveform matching.

Synthetic Seismograms, methods of computing to regional and near teleseismic distances using three-dimensional velocity models, at frequencies of 1 Hz and higher.

To accommodate HPC requirements for these aims, AFRL has indirectly relied on seismologists at LANL and LLNL working on programs funded by NNSA to make use of HPC facilities at each of their own laboratories. If these HPC facilities were available to university-based investigators, managers at AFRL could be expected to fund a broader spectrum of projects.

U. S. Geological Survey (USGS)

As part of a successful pilot project for its multi-hazards initiative, the USGS has participated in several "Shakeout" activities over the past several years. Each activity is built partly on cutting edge seismology: simulating rupture and the associated strong ground shaking at critical facilities and across areas with high population density to develop realistic and meaningful scenarios. Progress towards developing a seismology HPC facility could justify a larger appropriation to the earthquake hazards program to be spent partly on its support for earthquake scenarios and related hazard assessments.

NASA

The increased collection of space-based data for geodetic and climate-change investigations yields a significant need for computational technologies and resources to solve problems including inverting geodetic and gravitational data to extract deformation on faults and in active volcanoes, and distinguishing the effects of deep earth structure and crust and mantle rheology versus ice sheet changes on uplift. Looking outward to other planets, the methods used to model Earth's interior dynamics are and have been directly applied to Mars, Venus, the Moon, Mercury, and planetary bodies of the outer solar system. Numerous past, ongoing, and future missions, combined with modeling, are yielding unprecedented insight into the dynamical behavior of these planets. NASA scientists currently use a number of geodynamics codes (including CIG codes) to model Earth deformation. We see opportunities for increased partnerships with NASA on advanced computing.