Report on the

High Performance Computing in the Geosciences Workshop

Hosted by the
National Center for Atmospheric Research (NCAR)
25-27 September 2006

Prepared by the Workshop Program Committee:

Krista Laursen
Peter Backlund
Lawrence Buja
Peter Fox
Richard Loft
Michael Wiltberger

On behalf of the workshop participants

Preface

Editorial assistance during the preparation of this report was provided by the following workshop speakers, session leaders, and attendees:

Tom Bettge, NCAR
Frank Bryan, NCAR
Mick Follows, Massachusetts Institute of Technology (MIT)
Dale Haidvogel, Rutgers University
Chuck Hakkarinen, Belmont, California
Thomas Jordan, University of Southern California and Southern California Earthquake Center (SCEC)
Benjamin Kadlec, NCAR
James Kinter, Center for Ocean-Land-Atmosphere Studies (COLA)
L. Ruby Leung, Pacific Northwest National Laboratory (PNNL)
Dimitri Mavriplis, University of Wyoming
Chuck Meertens, UNAVCO
Dag Nummedal, Colorado School of Mines
Tim Palmer, European Center for Medium-Range Weather Forecasts (ECMWF)
Thomas “Zack” Powell, University of California at Berkeley
Jimmy Raeder, University of New Hampshire
David Randall, Colorado State University
Mark Rast, University of Colorado Laboratory for Atmospheric and Space Physics (LASP)
Lisa Cirbus Sloan, University of California at Santa Cruz
James Syvitski, University of Colorado at Boulder and INSTAAR
Robert Wilhelmson, University of Illinois at Urbana-Champaign and National Center for Supercomputing Applications (NCSA)
Kraig Winters, Scripps Institution of Oceanography
David Yuen, University of Minnesota

This material is based upon work supported by the National Science Foundation under Grant No. 0631272. 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.

Introduction

Organized and hosted by the National Center for Atmospheric Research (NCAR), the High Performance Computing in the Geosciences workshop was held at NCAR from 25-27 September 2006. The workshop was attended by 125 participants, including representatives of the university community, other national laboratories, supercomputing facilities, NCAR, and the National Science Foundation (NSF). The strong presence of NSF at the workshop was particularly noteworthy, with specific NSF attendees as follows:

Margaret Leinen – *Assistant Director for Geosciences*
Steve Meacham – *Office of Cyberinfrastructure (OCI) Program Director*
Art Goldstein – *Division of Earth Sciences (EAR) Acting Director*
Leonard Johnson – *EAR Continental Dynamics Program Director*
Rodey Batiza – *Division of Ocean Sciences (OCE) Ocean Drilling Program Director*
Eric Itsweire – *OCE Physical Oceanography Program Director*
Jarvis Moyers – *Division of Atmospheric Sciences (ATM) Director*
Cliff Jacobs – *ATM/UCAR and Lower Atmospheric Observing Facilities (ULAFOS) Head*
Sarah Ruth – *ATM/ULAFOS Program Coordinator*
Peter Milne – *ATM/Lower Atmospheric Research Section (LARS) Cross Disciplinary Activities Program Director*
Therese Moretto Jorgensen – *ATM/Upper Atmosphere Research Section (UARS) Science Assistant*

The goal of the workshop was to develop a compelling plan – supported by the geosciences research community – for implementation of the vision for a unified research cyberinfrastructure system called for in the “Petascale Collaboratory for the Geosciences” reports (see *http://www.joss.ucar.edu/joss_psg/meetings/petascale/*).

In the sections to follow, an overview of the workshop discussions is provided and a summary of key workshop outcomes and proposed next steps is presented. Appendix A includes a detailed summary of the various workshop sessions. All of the invited speaker presentations and breakout session reports referred to in this workshop report are available on-line at the following URL:

http://www.ncar.ucar.edu/Director/dcworkshop/agenda.jsp

Major Workshop Findings and Community Recommendations

As outlined in Appendix A, the Day Two breakout sessions provided workshop participants with the opportunity to consider and discuss a number of questions pertaining to collaboratory structuring and composition both from a disciplinary and a cross-disciplinary perspective. While the reports given by the various session leaders revealed some differences in opinion regarding how such a geosciences high performance computing (HPC) enterprise might be most optimally organized and utilized, more striking were the similarities that became apparent in terms of the desired components and working relationships within the collaboratory.

Readers of this report are encouraged to visit the workshop Web site at http://www.ncar.ucar.edu/Director/dcworkshop/agenda.jsp and review the reports assembled for each of the morning and afternoon breakout sessions. Here, key insights and recommendations to come out of all of the small group discussions are summarized.

- In general, there was consensus across the various disciplines that the geosciences community as a whole has problems suitable for petascale computing resources and that a need exists for access to large-scale computing resources. However, it should be noted that the timeframe within which access to petascale computing platforms is needed varies by discipline. For example, while atmospheric and earth science investigators are currently positioned to exploit petascale computing resources, hydrology and sedimentology researchers feel that more leadership and cohesion within their community will be needed in order for scientists in those areas to be able to make use of petascale computing resources within a five-year timeframe. Similarly, paleoclimate investigators participating in the workshop indicated that researchers in their field will need to learn how to make use of petascale resources.

- Many of the discipline areas cited a need for the collaboratory to provide access to a spectrum of small- to large-scale computing facilities to handle a broad range of computations. For example, solid earth scientists participating in the workshop outlined that while some of the computational problems studied in their community are adaptable to grid-based solutions (i.e., capacity computing performed on distributed cluster systems), other problems require access to capability computing resources.

- A number of the disciplinary breakout groups emphasized the need for both end-to-end community models and community data sets to be made available within the geosciences collaboratory. Additionally, access to forecast systems within the computing enterprise was also noted as desirable.

- The need for the collaboratory to provide end-to-end HPC research support was emphasized by the majority of workshop participants. Specifically, in addition to providing a gateway for access to needed computing resources, the collaboratory must also be structured to incorporate data visualization, archival, and analysis resources and software and application support services (e.g., code porting and development). Observational data should also be accessible via the collaboratory.

- Efficient build-up of the collaboratory can, and should, involve the leveraging of and affiliation with existing information technology (IT) development efforts and discipline centers (e.g., SCEC, CIG, GEON, CUAHSI, TeraGrid, etc.).

- The importance of creating the collaboratory using a distributed, grid-based framework was emphasized by several of the breakout groups. Rather than all HPC functions (e.g., computations, archival, visualization, etc.) taking place at one or a small number of institutions, the current capabilities and expertise available at a variety of locations should be leveraged and made accessible to collaboratory users.

- In terms of how computing resources provided within the collaboratory should be allocated, there was a general sense across the discipline areas that some form of grant process involving peer reviews of computing support requests should be implemented. Such a process should include mechanisms for reviewing not only the science that is proposed but also the computational feasibility of the proposal. A number of the participants also emphasized the importance of flexibility within whatever allocation model is put in place. This will be essential for ensuring that important event studies (e.g., tsunami, hurricane, and earthquake modeling) can be undertaken on a priority basis.

- A common theme across all of the breakout sessions (both disciplinary and cross-disciplinary) was the belief that the human dimensions of the collaboratory structure will be as important to the endeavor as the provision of needed hardware and software resources. Specifically, it was viewed as essential by the majority of attendees to involve not just domain and computer scientists in the collaboratory construction effort but also users, stakeholders, and social scientists. The provision of suitable community education and outreach, user education, and workforce training opportunities, as well as the establishment of appropriate governance and communication pathways, will also be important. To quote the Hydrology and Sedimentology breakout group, “cohesion among people is as important as hardware.”

- It was generally agreed that the collaboratory will foster new, stronger cross-disciplinary relationships, as research performed in one area will likely also be of benefit to investigators working in other disciplines. For example, studies of future water availability will be of benefit to climate, hydrology, and ground water researchers. In order to encourage the development of cross-disciplinary partnerships, the collaboratory should be equipped and configured to lower both the technical and sociological barriers to collaboration. Technical barriers to collaboration can be addressed by improving the interoperability of software, the availability of data, and the use of grid technology. While some of the approaches for promoting collaboration from a sociological perspective may be conventional (e.g., workshops, working groups, etc.), novel approaches to encourage the dissemination of innovations across disciplines should be implemented within the collaboratory.

- It was the opinion of the participants that if organized appropriately, a geosciences collaboratory will minimize replication and maximize productivity within the geosciences. For example, existing experts (e.g., code porting specialists) can be leveraged across the enterprise and on-demand computing can be made more readily available to greater numbers of investigators across the geosciences. Additionally, a well-conceived and structured collaboratory could promote greater data interoperability and more effective utilization by encouraging data format standardization and more consistent management and archival practices.

- A majority of the workshop participants expressed the opinion that an important mission of the collaboratory should be to focus resources on compelling science problems that require the collaboration of the entire geosciences community. Some examples of "interdisciplinary grand challenge" problems suggested during the workshop included:

    - Studies of the earth's past – Paleoclimate research is integrative of not just ocean, atmosphere, and ice studies but also of studies of the geological record. By adding explorations of plate tectonics and the Sun's variability (stellar evolution) as additional components to the study, the geosciences community would be able to assemble a project that is truly major in scope. Explorations of events like snowball earth, the Paleocene-Eocene Thermal Maximum (PETM), the Cretaceous-Tertiary (KT) event, and the ice ages will inform society's understanding of the Earth/Sun system's potential for variability.

    - Water cycle investigations – Understanding the Earth's water cycle is critical for human survival and would bring hydrologists and climate experts together in pursuit of a truly grand challenge mission.

  The performance of regional climate simulations and studies of how carbon dioxide is sequestered are two more examples of complex, large-scale research problems that could be studied using geoscience collaboratory resources. This list is not intended to be comprehensive or exhaustive. It is expected that an entire portfolio of such interdisciplinary grand challenges will emerge from the engagement of the geoscience community in the collaboratory planning process.

Next Steps: Moving Toward a GeoCollaboratory

As outlined in Appendix A, Dr. Tim Killeen concluded the workshop discussions on Day Three by outlining some key near-term steps to be taken to continue to move the geosciences collaboratory formation process forward. These are as follows:

- NCAR has been asked by NSF to serve as a catalyst for a community-led effort to define and begin to implement the desired geosciences enterprise. In this capacity, NCAR personnel will be working with both NSF and community representatives to develop and advance plans for building up the geosciences collaboratory.

- A new steering committee that is representative of scientific and technical expertise across the geosciences will be formed and will be charged to work with NCAR, NSF, and the broader community to develop an implementation plan for the collaboratory.

- A Town Hall meeting will be held during the December 2006 AGU meeting to provide an opportunity for further discussion of the geosciences collaboratory concept with members of the scientific community. Additional opportunities for such community outreach and dialogue will also be identified in the coming months in order to ensure that the greatest possible segment of the geosciences community is brought into discussions about the collaboratory development effort.

- A series of pilot projects should be identified and initiated in order to demonstrate the utility of forming a geosciences collaboratory.

- Annual workshops should be proposed and conducted in order to continue to build geosciences community consensus regarding science drivers for the collaboratory.

Conclusions

The High Performance Computing in the Geosciences Workshop brought together investigators from across the geosciences and provided them with a forum within which to discuss and debate the cyberinfrastructure needs of geoscience researchers. A very clear message coming out of the workshop is that the geosciences community is ready for – and needs – better cyberinfrastructure that will provide access not only to HPC and data handling (e.g., archival, analysis, visualization) resources but that also enables true inter- and cross-disciplinary collaboration. Additionally, workshop participants clearly articulated a need for end-to-end modeling and forecast systems. The provision of such capabilities within a geosciences collaboratory framework will be welcomed by the community. As an initial consensus and community building step, the workshop was highly successful and has set the stage for the collaboratory planning and implementation process to follow in the coming months.

Appendix A: Session Summaries

Day One – NSF and Community Perspectives on HPC Challenges

After welcoming the participants to the workshop, NCAR Director Tim Killeen discussed the workshop focus. As outlined by Dr. Killeen, the workshop was organized to provide representatives of the geosciences community with an opportunity to engage in dialogue about how best to develop the petascale scientific cyberinfrastructure needed to address unmet interdisciplinary and grand challenge research problems. The workshop was also intended to serve as a vehicle for the participants to discuss strategies for positioning the geosciences community within the larger high performance computing (HPC) landscape and to integrate the vision and needs of the community into a focused scientific research enterprise.

The remainder of the Day One agenda following Dr. Killeen's opening remarks was committed to invited talks given by representatives of NSF and the geosciences community. Dr. Steve Meacham of OCI gave the first of two NSF presentations and presented a summary of OCI's HPC investment foci, the TeraGrid development effort, and an OCI perspective on petascale computing challenges. Highlights from Dr. Meacham's talk included his discussion of the following:

- The *overall OCI cyberinfrastructure support portfolio*, which includes the provision of support for data, data visualization, and analysis efforts, high performance computing, the enhancement of virtual organizations, and learning and workforce development;
- The *OCI strategy for HPC resource provision* via the Track 1 (petaflop supercomputing systems), Track 2 (several hundred teraflops systems), and Track 3 (university HPC systems) investment plan;
- *The TeraGrid as an extensible, integrating architecture for HPC research*;
- *Unique challenges posed by petascale computing*, including how to effectively make use of large numbers of cores and threads, how best to fully exploit parallelism, the need for effective input/output (I/O) management, fault tolerance requirements, and how to effectively work with and exploit hybrid systems.

Following Dr. Meacham's talk, Dr. Margaret Leinen of NSF spoke about the unique cyberinfrastructure requirements of the geosciences. In her comments to the workshop participants, Dr. Leinen made the following key points:

- The NSF Geosciences Directorate (GEO) recognizes the maturity of the geosciences community in developing and using cyberinfrastructure.
- Observation and simulation are currently – and will increasingly become – an integral part of the pathway to new knowledge in the geosciences.
- The addition of significant new observational capabilities to the GEO portfolio is making the need for new cyberinfrastructure capabilities greater than ever before.

- All disciplines within the geosciences (e.g., oceanography, earth sciences, atmospheric sciences, space physics, etc.) have emphasized the need for access to HPC resources.
- The suitable provision of support services must be an essential component of the strategy to provide HPC resources to the geosciences community.
- The ability of researchers to manipulate and assimilate large data sets will be critical to the conduct of successful research.
- GEO understands that an investment will need to be made in the preparation of codes to run on petascale platforms.

At the conclusion of her talk, Dr. Leinen posed the following questions to the workshop attendees:

- From the perspective of the GEO community, what are the advantages and disadvantages of having more specific discipline/problem oriented supercomputing centers versus an all-purpose center? Are both valuable and do they serve different purposes?
- What could the geosciences community do with a discipline-oriented half-petaflop computing platform if it is made available by 2010?
- What balance should NSF strike between making larger supercomputing systems available to researchers versus investing in needed software development?

After the two high-level NSF presentations, the focus shifted to presentations given by representatives of the ocean, earth, and atmospheric sciences communities. In general terms, each of the community members invited to speak on the opening day of the workshop was asked to provide their perspective on important scientific research areas that could benefit from the presence of a geosciences collaboratory and on HPC challenges that currently exist within their geosciences discipline area.

The first of the community speakers – Dr. Dale Haidvogel of Rutgers University – gave a presentation outlining the supercomputing needs and challenges currently being confronted by the coastal oceanography community. As detailed by Dr. Haidvogel, the last five years have seen significant progress made in addressing the needs called out in the 2002 Ocean Information Technology Infrastructure (OITI) report entitled *An Information Technology Infrastructure Plan to Advance Ocean Sciences*. Specifically, the circulation models, observing networks, and data assimilation methods identified as requirements in the report are (or will soon be) in place. Specific challenges facing coastal oceanographers in the next five years are, as Dr. Haidvogel outlined, as follows:

- Performing multi-variate data assimilation;
- Addressing ecosystem complexity;
- Developing sub-gridscale parameterizations;
- Improving coastal observing networks;
- Applying robust models to unstructured grids; and
- Solving a variety of technical issues, including grid nesting.

Speaking as a representative of the earth sciences community, Dr. Tom Jordan of the University of Southern California and the Southern California Earthquake Center (SCEC) next gave a presentation describing the application of petascale computing to earthquake system science. As detailed by Dr. Jordan, the modeling of seismic shaking is, fundamentally, a system-level problem. Such efforts – which involve the study of nonlinear, multi-scale interactions among many geologic components – are both large in scope and large in scale. Additionally, simulations require the application of many types of high performance computing: capability computing to undertake big runs, capacity computing to perform many runs, data-intensive computing to handle large data volumes, and on-demand computing to allow for real-time tracking of event cascades. As discussed by Dr. Jordan, SCEC is working to advance the field of earthquake science through system-specific studies (e.g., the Puente Hills blind thrust system, the San Andreas fault system) in the southern California natural laboratory. Observations offered by Dr. Jordan at the end of his talk included the following:

- Terascale computing efforts have made better earthquake forecasts and ground motion predictions possible by demonstrating the feasibility of physics-based probabilistic seismic hazard analysis (PSHA). However, in order for physics-based PSHA to become operational, SCEC will require petascale capability, capacity, and data-intensive computing.
- National petascale initiatives will rely on the vertical integration of hardware, software, and personnel resources into a cyberinfrastructure that can support the application of petascale computing in domain problem areas.

The final invited speaker of the afternoon, Dr. Robert Wilhelmson of the National Center for Supercomputing Applications (NCSA), gave a presentation that focused on how lessons from the past can be used by the community to prepare for a future involving petascale computing. In his talk, Dr. Wilhelmson emphasized the importance of cyberinfrastructure in enabling research, forecasting, decision making, and education in the geosciences. In discussing the value of cyberenvironments, Dr. Wilhelmson outlined that many students spend 80% of their time doing low-level software and technology work and only 20% of their time on actual science. He advocated the need to flip these values by actively working to build needed cyberinfrastructure systems. Using the development of the George E. Brown, Jr. Network for Earthquake Engineering Simulation (NEES) as an example, Dr. Wilhelmson discussed the lessons learned from the design, development, and deployment of a community grid infrastructure. Namely:

- Communication is key;
- The project should be viewed as a partnership;
- The project must be kept moving forward;
- Ways should be identified for “users” to take responsibility;
- The community should be listened to but the project team must be prepared to lead; and
- The development effort should build on previous work.

Several common themes were apparent in the presentations given by Drs. Haidvogel, Jordan, and Wilhelmson. Specifically, geoscience researchers in various disciplines are voicing many of the same desires, including:

- *To obtain more complete descriptions of the systems being studied.* This is driving more complex phenomenology, a move from disciplinary to interdisciplinary models, and the utilization of coupled models.
- *To obtain end-to-end prediction systems.* This is producing a need for efficient data assimilation/inversion methods, workflow management systems, on-demand computing (to study natural disasters), and tools to perform geospatial data manipulation.
- *To capture multiscale phenomena.* This is driving a push for higher model resolutions, the inclusion of adaptive mesh refinement techniques and a move to unstructured grids, and the development of subgridscale parameterizations.
- *To achieve more model coverage.* This is leading to the inclusion of bigger regions in models, the study of more parameter values, and the incorporation of larger ensembles.

During the evening of the first day, a welcoming reception and dinner was held for workshop participants at the NCAR Mesa Laboratory. At the completion of the evening's social events, Dr. Tim Palmer of the European Center for Medium-Range Weather Forecasts (ECMWF) gave the last of the invited talks. During a presentation entitled "Petaflop Computing and Reliable Climate Prediction – a European Perspective," Dr. Palmer presented the argument that if there is currently no bigger problem facing society than climate change, then quantifying the threat of climate change using the best available resources (including HPC resources) must surely be of the highest priority to society.

Day Two – Community Dialogue on Building a Geosciences Enterprise

The second day of the workshop was structured to address the primary purpose of the workshop: namely, to provide a forum for geoscience community members to exchange ideas regarding how best to develop the cyberinfrastructure framework needed to address interdisciplinary and grand challenge research questions.

To set the stage for the participant discussions scheduled for the majority of Day Two, Drs. Frank Bryan and Richard Loft of NCAR began the day by giving presentations providing background and context for the concept of developing a geosciences HPC enterprise. In his talk, Dr. Bryan gave an overview of the community input process used to create the *Establishing a Petascale Collaboratory for the Geosciences* (PCG) reports and outlined the major recommendations made in these documents. As was emphasized by Dr. Bryan, the overarching recommendation made in the PCG reports is to "establish a Petascale Collaboratory for the Geosciences with the mission to provide leadership-class computational resources that will make it possible to address, and minimize the time to solution of, the most challenging problems facing the geosciences."

Dr. Loft followed this first Day Two talk with a presentation outlining a blueprint for a distributed domain-specific enterprise for the geosciences. In addition to re-emphasizing the findings outlined in the PCG reports, Dr. Loft discussed how the community views summarized in the PCG documents have further evolved to now incorporate key elements of the NSF cyberinfrastructure strategic plan. In particular, the development of a geosciences collaboratory should be based on a distributed model that leverages emerging grid technology to effectively link supercomputing resources, virtual organizations, observing systems, and data archival and visualization resources.

After these two presentations by NCAR personnel, the remainder of Day Two was dedicated to the conduct of breakout sessions. In the morning, six breakout sessions were held and were organized along the following disciplinary lines: meteorology, space physics, hydrology and sedimentology, solid earth, oceanography, and paleoclimate. Each of the breakout groups was asked to focus on the discussion of HPC and collaboratory needs unique to their discipline area. The specific questions posed to each of the discussion groups were as follows:

- How can a geosciences HPC enterprise best be structured to meet the needs of your discipline?
- What will be required to make such an enterprise function effectively?
- What are the characteristics of a computing resource allocation model that will best meet the needs of your discipline?
- What types of service and support will your discipline require and what are the types of applications that will need to be supported?

Following the break for lunch, the morning breakout session leaders and reporters presented summaries of each group's deliberations and key findings in a plenary session. A summary of the key recommendations made is provided in the section of this document entitled "Major Workshop Findings and Community Recommendations."

Over the lunch period, Mr. Tom Bettge of NCAR gave a presentation outlining the NCAR model for providing supercomputing support for the atmospheric and related sciences. In his remarks to workshop participants, Mr. Bettge described how NCAR provides not only access to needed supercomputing resources but also essential support services including mass storage, research data archival (RDA), data analysis and visualization, climate research resources via the Earth System Grid (ESG), and computational science research.

After the completion of the morning breakout session reports in plenary session, workshop attendees were asked to participate in cross-disciplinary breakout sessions structured to allow for the discussion of more general, overarching collaboratory formation issues. The questions the four afternoon breakout groups were asked to address were as follows:

- How can the community best facilitate – through the development of geoscience cyberinfrastructure – better collaboration in the areas of education, outreach, training, and workforce development?
- How can the economy of scale for a geosciences HPC enterprise best be demonstrated?
- What is the best model for establishing a fair relationship between available resources, required science, and allocation of computing?
- How can a geosciences HPC enterprise best be designed in order to provide the capability for addressing emerging opportunities and for supporting "hero" computing requirements?

Day Three – Summarization and Looking Ahead

The third – and final – day of the workshop began with comments from Dr. Art Goldstein of NSF. Speaking as a representative of the NSF geosciences personnel in attendance at the workshop, Dr. Goldstein reiterated the commitment of GEO to identifying a strategy for effectively meeting the cyberinfrastructure needs of the geosciences community. In response to a question from one of the audience members about how a geosciences collaboratory might be funded from within NSF, Dr. Goldstein noted that the current thinking within GEO is that the development of such an enterprise can and should be done without inflicting financial harm on other, equally important GEO research programs.

After the comments from Dr. Goldstein, session leaders and reporters from the Day Two afternoon breakout sessions presented summaries of the various session discussions. As with the morning breakout session findings, observations and recommendations presented by the afternoon session groups are presented in the "Major Workshop Findings and Community Recommendations" section of this report.

Following the delivery of Day Two afternoon breakout session reports, a community panel consisting of representatives of the earth, ocean, and atmospheric sciences was provided with the opportunity to synthesize the workshop discussions and offer perspectives on required next steps for building up a geosciences collaboratory. Participants on the panel were Dr. Dag Nummedal (Colorado School of Mines), Dr. Tim Palmer (ECMWF), Dr. Thomas "Zack" Powell (University of California at Berkeley), Dr. Jimmy Raeder (University of New Hampshire), Dr. James Syvitski (University of Colorado at Boulder and INSTAAR Director), and Dr. Kraig Winters (Scripps Institution of Oceanography).

The workshop was officially brought to a close by Dr. Tim Killeen, who gave a brief presentation summarizing key workshop outcomes and the next steps to be taken by NSF, NCAR, and the community to maintain momentum with the geosciences collaboratory building effort.