Advanced Scientific Computing Advisory Committee
Biological and Environmental Advisory Committee

Report on Computational and Informational Technology Rate Limiters to the Advancement of Climate Change Science

31 October 2007

Joint ASCAC-BERAC Subcommittee

James J. Hack, NCAR (Co-chair)
Eugene Bierly, AGU (Co-chair)
Dean N. Williams, LLNL
William D. Collins , LBL
Philip Jones , LANL
John B. Drake, ORNL
Phil Colella , LBL
Dave Bader, LLNL
Ian Foster, ANL
Edward S. Sarachik, University of Washington
Brian Gross, GFDL

Executive Summary

Introduction:

At the August 2007 meeting of the Advanced Scientific Computing Advisory Committee (ASCAC) for the Office of Science (SC), United States Department of Energy (DOE), the committee was jointly charged with the Biological and Environmental Research Advisory Committee (BERAC) by Dr. Raymond Orbach, Director of the Office of Science, to identify the key computational and information technology obstacles to advancing climate change science and improving climate change projections using state-of-the-science coupled climate models. The ASCAC and BERAC co-chairs assembled a subcommittee of experts to address this charge which needed to deliver a report by the November 2007 meeting of the ASCAC. The committee held one teleconference call, and assembled in Washington on 16-17 October to complete drafting the major elements of this report.

The charge to this committee was very broad and it was necessary to agree to some ground rules in order to make progress in the extremely short period of time we were given. The issues pacing progress in climate change science are far broader than historically have been, or can be, unilaterally addressed by the DOE Office of Science. So, our discussions were focused on DOE's strengths and opportunities to leverage broader efforts in the climate change and climate modeling community. Given the time constraints, and the fact that other efforts are currently underway to address similar issues (e.g., an ongoing NRC study of *The Potential Impact of High-End Computing On Four Fields of Science and Engineering*) our goal was to produce a relatively short balanced response to the charge.

In parsing the charge the committee looked at the respective investments and roles of ASCR and BER in climate change science. These investments are complementary in many respects, with a number of clear opportunities for partnerships in specific topical areas. They include computational requirements and solutions, advanced software needs and solutions, development and application of advanced algorithm and applied math techniques to efficiently exploit evolving computational architectures, data management challenges and solutions, networking challenges and solutions, and the development and application of technologies to support scientific collaboration. There are also comparably important and unique investments made by BER in contributing to basic scientific knowledge about the climate system, in specialized observational programs, and in the exploration and enhancement of global modeling approaches and techniques. With this in mind, the committee decided to approach the charge by examining rate limiters to specific science opportunities, and how ASCR and BER might be able to best invest resources to address these pacing factors.

Climate Change Science

The climate community is facing significant challenges and opportunities in its efforts to advance basic science and the application to policy formation. With the release of the 2007 IPCC assessment and the CCSP reports, climate science is entering a new phase. Several of the classical problems -- in particular the detection, attribution, and "finger-printing" of climate change -- have essentially been resolved in these latest assessments. The global community is now faced with a new set of urgent problems, including robust projections of regional impacts; forecasts of abrupt and extreme climate change; simulation of shifts in the water cycle; and prognosis of carbon-cycle feedbacks. In order to address these issues, the community should develop and undertake a coordinated research program balanced among observation, computation, and theory. Collaboration between DOE's Advanced Scientific Computing Research (ASCR) and Biological and Environmental Research (BER) has produced world-class climate models adapted to the most advanced computing platforms in existence. Meeting future challenges in climate change science will require a qualitatively different level of scientific understanding, modeling capabilities, and computational infrastructure than is currently available to the scientific community. Many of the questions now facing climate change science will require the development of a new generation of more comprehensive climate models, most frequently referred to as Earth System Models (ESMs), which among other things will need to accurately incorporate biogeochemical cycles and feedbacks. A strategic partnership between ASCR and BER could accelerate progress on the major new challenges in climate science.

Climate science has become increasingly international in scope. The success of national and international assessments depends to a large degree on the free exchange of observations and simulations. One of the premier examples is the multi-model archive of simulations for the IPCC Fourth Assessment Report (AR4) assembled and maintained by the Program for Climate Model Diagnosis and Intercomparison (PCMDI). The global climate community has authored 324 peer-reviewed publications based upon this archive as of October 2007. These articles have been instrumental in assessing the fidelity of global climate models and in developing new initiatives to advance the science of climate simulation. The scientific productivity of the PCMDI archive suggests that the development of comparable (perhaps virtual) collections would be tremendously beneficial for future assessment activities. However, the next generation of ESMs will challenge existing frameworks for computation, communication, and analysis. A strategic partnership between ASCR and BER will be essential to insure that the U.S. can enhance its leadership in climate science through advances in these key fields shared with the international community.

Scientific Opportunities

One of the most promising pathways to improving our understanding of climate change has been the development of models that represent the complexity of interactions in the Earth system as accurately as possible. Over the last 30 years, these models have advanced considerably in spatial and temporal resolution and in the representation of key climate processes. However, forecasts of environmental and societal responses to climate change remain highly uncertain. The opportunity to credibly contribute to the current global change debate is hindered by the current limits of climate models to address regional and local-scale impacts on time scales of greatest interest to society. Examples include the ability to accurately project climate change impacts on regional space scales and decadal time scales, to project changes in extreme events (e.g., heat waves, drought, precipitation frequency and intensity, and the intensity of larger-scale synoptic weather events), to accurately anticipate changes in low-frequency climate variability, characterize changes to the water cycle (with the associated impacts on agriculture, biofuels, and human health), to quantify sea level rise, and to explore processes that might contribute to abrupt climate change. Significant investment over the next few years can lead to a quantitative improvement in the scientific community's ability to address these difficult but societal relevant questions.

One example of an immediate scientific challenge and opportunity which is already central to Office of Science activities; the incorporation of biogeochemical processes in climate models. The science surrounding the biogeochemical coupling of climate has become central to answering climate change questions as we learn more about how the coupled carbon cycle has changed in the fossil record, how it is changing in the present day, and how it might change in response to global climate change. Addressing the science issues will require new observations and methods of analysis, new theoretical understanding of the carbon cycle, and new models of the Earth system that include the interactions between human society and the environment. These models play pivotal roles in interpreting the paleoclimate records, in synthesizing and integrating measurements to study the current carbon cycle, and in projecting the future responses of human society and the natural world to evolving climate regimes.

Other scientific challenges are related to developing an understanding of the significant impacts that could follow from abrupt changes in the climate system. One of the largest uncertainties in current climate assessments is the rate of sea level rise. The melting of large ice sheets in Greenland and Antarctica can result in 12m of sea level rise (6m from Greenland and another 6m from the West Antarctic Ice Sheet). Melting of these large ice sheets was thought to be a millennial process because current models only represent the slow melting in response to surface heating. Recent observations indicate ice sheets can melt on much more rapid timescales due to dynamical processes in large outlet glaciers and ice streams within the ice sheet. Basal lubrication of ice streams and the melting of floating ice shelves that buttress outlet glaciers have been observed to increase flow rates by factors of 2-10, greatly increasing the loss of ice and the subsequent rate of sea level rise. A high priority for climate models is the inclusion of fully dynamic ice sheet models and the ice shelf-ocean interactions needed to assess the rate and magnitude of sea level rise due to rapid ice sheet melting.   Abrupt climate change can also result from thresholds and nonlinearities in the response of climate to slower climate forcing. Examples include rapid changes in ocean circulation, large scale vegetation mortality and succession, release of methane frozen in ocean and permafrost clathrates, rapid ice sheet melting and megadroughts and dust. The climate community will need to use models to identify thresholds of forcing in the climate system and explore the likelihood and impacts of such abrupt change scenarios. These are but a few of the many immediate science opportunities that can be exploited through targeted partnership investments in ASCR and BER.

Rate Limiting Issues

The community's efforts to advance climate modeling and its application to science and policy will require advances in essentially every aspect of the models' theoretical, observational, and computational foundation. These advances represent near-term opportunities for targeted investment by ASCR and BER.

Traditional projections on centennial time scales are strongly influenced by the future trajectory of anthropogenic emissions, while forecasts for decadal time scales are governed primarily by the past history of the ocean. Therefore near-term climate forecasts will require two *new* developments: multi-scale models that can explicitly resolve important meteorological systems at regional scales, and retrospective analyses of the global oceans to initialize the forecasts. The field of ocean data assimilation is still in early stages and development and exploration, and it would benefit from the transfer of adaptive assimilation methods under development for atmospheric applications. For multi-scale atmospheric models, one of the primary challenges is the reproduction of weather and climate-related phenomena with sufficient fidelity for both meteorologists and climate scientists.

Climate science is largely data limited, and the success of the research is contingent on basic measurements and observations necessary to validate, verify, and constrain Earth system models. Therefore, quantifying the uncertainties in predictions is expected to require a new level of integration between modeling and observational science. New mathematical methods and algorithmic techniques will also be required to address the fundamental challenges of multi-scale coupling of physical, dynamical, and biogeochemical processes. A flexible high-performance computing infrastructure has been and will continue to be a key factor in making these advances possible.

There are many examples of scientific challenges that broadly rely on strategic investments in scientific and computational infrastructure. Accurate projections of changes in the local frequency of climate extremes including heat waves, droughts, flood, and synoptic events will be essential for the development of robust adaptation strategies. However, extremes represent the high-order moments of the climate system, and climate models have been designed primarily to treat the low-order moments. Much more research is required to understand how simulated extremes change with increasing model resolution, increasingly sophisticated parameterized treatments of non-resolvable processes (e.g., clouds, moist convection), and how the statistics of extremes are related to large-scale lower-frequency variability. A better understanding of, and ability to reproduce, observed low-frequency variability is also critical for the detection of climate-change signals. For Earth system modeling, it has become essential to characterize the natural modes of variability in the carbon cycle, terrestrial ecosystems, and dynamic vegetation.

As suggested earlier, a large number of significant impacts could follow from abrupt changes in the climate system. These occur when the gradual increases in climate forcing trigger an abrupt transition (or fast, growing mode) of the coupled system to a new state. Potential examples of abrupt change include dynamic dissolution of the Greenland or Antarctic ice sheets and bifurcations among multiple equilibria of the ocean circulation system. Characterization of the risk of abrupt climate change requires a new paradigm for climate change modeling, one in which the models are integrated over the full range of uncertainties in forcing and parameterized physics. Exploration of this phase space will require implicit formulations of the coupled system designed for fast equilibration combined with parametric continuation techniques and sustained petascale computing.

Better understanding of low-frequency variability is critical for the detection of climate-change signals. For Earth system modeling, it has become essential to characterize the natural modes of variability in the carbon cycle, terrestrial ecosystems, and dynamic vegetation. Current understanding is limited by the length of the observational record and, more fundamentally, by open issues regarding the stationarity of climate statistics and the coupling between climate processes and variability. The wide dynamic range in relevant space and time scales complicates resolution of the coupling issues. New mathematical methods designed for multiscale systems hold promise and should be actively explored for this class of problems.

Multiscale interactions also complicate investigations of the water cycle. As with variability, process-level understanding of the water cycle is limited by the lack of basic observations. Studies of cloud formation and evolution require measurements that are not available with present technology, including the vertical velocity and supersaturation of the cloudy atmosphere. While the absence of these data still represents a barrier to progress, near-term enhancements in computational capacity would permit the resolution of fundamental phenomena involved in both weather and climate change. Continued targeted investments in observational programs like ARM would provide the necessary data to validate high-resolution process modeling studies. These include tropical storms, extratropical storms and fronts, mesoscale convective complexes, and rain shadows downwind of major mountain ranges.

The ocean is responsible for much of the inertia or “memory” in the climate system due to the ocean heat capacity and the long time scales associated with ocean circulation and ocean mixing processes. Ocean data assimilation will be necessary to provide an initial ocean state for decadal prediction that corresponds to the Earth’s recent climate history and represents a pacing item for seasonal to interannual to decadal prediction. Ocean assimilation has been hampered by a lack of data, particularly for salinity and for ocean properties at depths below 1000m. Consequently, ocean assimilation is still a relatively new field with limited experience, particularly in a climate context. Recent progress in deploying large numbers of floats and the launch of new satellites that will measure surface salinity will greatly improve our ability to effectively constrain ocean models with assimilation. Efforts to improve the data assimilation in current ocean climate models will require the adoption of assimilation methods beyond the current simple optimal interpolation approaches.

Finally, there are very important software and hardware infrastructure challenges pacing progress in climate change science. Examples include the following:

- **Scalability**

Climate modeling tools need to be able to exploit petaflop computer systems, which places a severe demand on scalability of very complex modeling systems

- **Operating systems**

Some application software can only run on specific platforms and operating systems. Even worse, application software may produce different results on different platforms.

- **Optimizing software:**

There are many aspects to application software, which covers a wide-range of software development. Combining and optimizing the use and functionality of these disparate software components at each stage of the application development use can make them more efficient and/or allow the use of fewer resources, thus reducing bottlenecks.

- **Documentation:**

Poorly documented software and finding the right help for application software creates productivity bottlenecks

- **Maintenance:**

Maintenance, porting mechanisms to new platforms, and validation procedures for application software is needed to reduce bottlenecks.

Implications for Investments

Models

The climate community needs to develop a new generation of Earth system models based upon new and expansive requirements:

- Ability to more accurately reproduce major modes of natural variability to enable predictive capabilities from intraseasonal to decadal time scales
- Functionality for decadal-scale ensemble forecasts at very high spatial resolution;

- Flexibility to incorporate new data on the physical, chemical, and ecological climate system in the form of process representation, thereby increasing the fidelity of climate simulations;
- Connectivity with user communities for adaptation and mitigation strategies; and
- Capability for two-way interactions among emissions, impacts, adaptation, and mitigation

The community will work to meet these requirements by leveraging ongoing investments in geophysical and computational science supported by the Office of Science. However, the anticipated complexity of the models and applications are sufficiently demanding that new frameworks are needed for prototyping, testing, and evaluation. Based upon current methods, it has proved challenging to attribute systematic features in the simulations to the specific aspects of the dynamical, physical, or numerical formulation of the models. For this reason, new Earth system models should be modular and hence easy to disaggregate and reassemble. Each functional module should be accompanied by test cases and the observational and/or simulation data required for rigorous and reproducible evaluation. The modules should be assembled in a flexible model superstructure that enables staged increases in process complexity. A strategic partnership between ASCR and BER could provide the new mathematical and computational frameworks required for robust and extensible model development.

At the core of BER's climate change program is the study of the earth's carbon cycle, a research endeavor that began fifty years ago and continues today with strong measurement programs, processes studies and links between the Climate Change Science Plan and the Climate Change Technology Initiative. In order to understand climate's effect on ecosystems and the feedbacks between land and ocean ecosystems, the global climate models are being extended to earth system models that include a full balancing of the carbon budget. A unique opportunity exists for DOE to lead the development and application of these models to advance the science and integrate observations to improve models. The strengths of the Terrestrial Carbon Program complement the Climate Change Prediction Program and SciDAC Program projects in this area where the challenges of assimilating flux tower data and employing process model studies to predict and design ecosystem manipulation experiments is largely untapped. The algorithmic and mathematical challenges call for new methods and optimization techniques to be found to deal with this coupling of the biological and physical systems.

While the uptake of carbon by ocean and terrestrial ecosystems is a key element of the carbon cycle, the dynamic changes in vegetation entail fine scale changes in land cover and radiative albedo. Both the carbon in the atmosphere and the ecosystems ability to hold soil moisture and alter the evaporation and transpiration link carbon processes with the hydrological cycle of the climate system. Indeed, the dynamic forcing of land cover is dependent on regional precipitation patterns of the atmospheric circulation, pointing to the need to improve the hydrologic biases in the physical climate system in tandem with extending the models to include biogeochemical cycles. Tropical biases and cloud forcing have remained important areas of research for the DOE program with the ARM program providing essential data and radiative parameterizations to the global modeling program. The column radiation model in the CCSM is one of the cleanest examples of software interoperability allowing comparative studies of new ideas and university lead process studies. This also represents another opportunity to exploit new mathematical methods and algorithmic techniques to address the fundamental challenges of multi-scale coupling of these physical, dynamical, and biogeochemical processes.

Observations

Meteorological and oceanic analyses have become an important tool for studying the mean state and variability of the current physical climate. These analyses are constructed using a model that is adjusted by incorporating observations during its integration. These analyses have proved particularly useful for understanding the relationship between observations and the underlying dynamics of the climate system. It would be especially valuable to have a comparable analysis of biogeochemical cycles that could relate local and global biogeochemical processes. However, there are no extant analyses that encompass the physical, chemical, and biogeochemical processes in the climate system. Development of these analyses will require significant investment in assimilation systems for chemical and biogeochemical observations from in situ and satellite platforms. It will also require much more advanced models to understand the error characteristics of the analysis system.

Simulation of biogeochemical cycles also requires detailed understanding of terrestrial and oceanic ecosystems; the exchange of organic and inorganic carbon compounds with other parts of the climate system; and the fluxes of energy, water, and chemical compounds (e.g., nutrients) that affect these ecosystems. The critical nutrient cycles for ocean and land ecosystems span over time scales ranging from a few days (such as nitrogen) to over thousand years. Modeling over these large time constants to fully evaluate the couplings between biogeochemical cycles and ecology will be a significant computational challenge. The spatial heterogeneity in the biosphere, below ground and above ground ecology is a fundamental issue overlying much of this science. Development of new models that can develop sensible volume/area/mass averaged and mass conserving idealizations that preserve the heterogeneity of the process and still allow for a degree of conceptualization are needed. Some other major open challenges are the sophistication of the ecological representations, the effects of high-frequency spatial and temporal variability on the carbon cycle (e.g., fronts and eddies); and the behavior of the biogeochemical cycles in coastal zones (Doney et al, 2003).

The ecosystem representations tend to be formulated as paradigms of ecological functions. The field certainly needs more mechanistic models of these ecosystems constructed at the level of individual organisms. It also needs much more detailed understanding of the nutrient networks and how these networks affect the carbon cycle. The effects of sharp gradients or rapid changes in the physical environment of the components of the carbon cycle are not well understood. With the advent of ultra-high-resolution ESMs over the next decade, it should become possible to probe the effects of rapid variability on scales much smaller than the mesoscale. Finally, the biogeochemistry in coastal zones has not been adequately studied. These regions have been challenging to simulate in global models with insufficient resolution to resolve the coastal regions, the discharges of river sediments into the regions, and other related features.

ALSO NEED WORDS IN HERE TO HIGHLIGHT THE INCREASED NEED FOR TARGETED OBSERVATIONS OF FUNDAMENTAL PROCESSES IN THE PHYSICAL CLIMATE SYSTEM (e.g., ARM).

Algorithms

There is a broad class of mathematical and numerical algorithms that need to be explored for application to the climate problem. We list several of the more obvious opportunities for enabling higher resolution simulations with shorter time to solution.

Alternative vertical discretizations

Climate modelers are beginning to introduce new vertical discretizations to better capture both boundary layer processes and isentropic/isopycnal flow outside the boundary layer. In particular, the use of quasi-Lagrangian coordinate schemes will permit better simulation of flow and minimize numerical diffusion. In the ocean, Arbitrary Lagrangian-Eulerian schemes are being introduced to maintain Lagrangian coordinates in the deep ocean while still resolving the surface mixed layer with fixed Eulerian levels. These new vertical methods require new techniques for determining and generating the optimal vertical grid based on physical properties of the simulation. In addition, methods for high-order conservative remapping of variables will be required as grids evolve in time.

Implicit time stepping

As the resolution of climate models is increased to improve accuracy and to predict regional changes, the allowable time step decreases in current explicit forward time integration methods. However, processor performance will not be increasing rapidly enough to make up for the reduction in time step size and the resulting increase in number of explicit time steps required for multi-century integrations. In addition, there are very long timescales in the climate system, like the thermohaline circulation and the long equilibration rate for terrestrial and ocean ecosystems, that will require very long simulations to provide a realistic climate state. The community will need to begin to explore the use of alternative strategies for increasing time step size, such as fully implicit models, Jacobian-Free Newton-Krylov solvers or other advanced solver techniques.

Regridding

A fundamental issue in coupled climate models is the communication of energy, water and tracer fluxes between system components in a conservative and accurate manner. Current methods work reasonable well, but will not scale for more comprehensive configurations planned for these models. Robust grid remapping algorithms that work efficiently for high resolution and future dynamic grids will be required. Higher-order regridding with monotone limiters will also be required for Earth System Models.

Assimilation
As climate modeling enters a more predictive paradigm, data assimilation will become increasingly important. While assimilation has been extensively developed and used in the weather community, the climate community will need to evaluate which assimilation methodology is best suited for climate simulation and the creation of realistic initial states for climate change scenarios. Optimal interpolation and simple methods have so far been adequate for the ocean due to sparseness of data, but with the influx of new ocean data sets, advanced techniques like ensemble Kalman filters or 4-D variational assimilation will need to be examined.

Variable resolution
For the atmospheric component of the climate model, there are strong arguments to exploiting higher-resolution variable gridding configurations. The computational demands of uniform ultra-high discretization of an atmospheric model would exceed the capacity of a petascale system. A more practical approach to dealing with resolution issues is to use a multi-resolution discretization, such as nested refinement, provided that the regions that require the finest resolution are a small fraction (10% or less) of the entire domain. In that case, the computational capability required could be reduced by an order of magnitude or more, and make the goal of computing with such ultra-high resolution models more feasible. Such multi-resolution methods have begun to be used in other areas of atmospheric modeling, such as numerical weather prediction; however, they are not yet in use in production climate models. There are a broad range of design issues that would need to be addressed for such models could be used routinely in atmospheric models, including choice of discretization methods, coupling between grids at different resolution, and dependence of sub-grid models on grid resolution.

More generally, variable resolution grids and dynamic grids will be increasingly important. The former may provide a more efficient approach to high resolution simulation in both atmospheric and ocean modeling. The latter might be required as we allow sea level to rise and ice sheets to retreat, changing the topography and coastal outlines. The climate community has some experience in this area, but could easily take advantage of the extensive expertise and developed software already developed under ASCR support.

Uncertainty quantification
In order to provide useful information to policy makers, the climate modeling community needs to better characterize the uncertainty in simulation results. Ensembles and basic statistics are currently used to assess uncertainty due to internal variability intrinsic to the climate system. More formal methods for verification, validation and uncertainty quantification are needed from the computer science, mathematics and statistical science communities. A particular challenge is the sparse nature of much of the climate data necessary to perform model validation.

Software/Middleware

The scientific challenge for future generations of climate models requires a closer integration among DOE's measurement and observational program as well as a step change in model development activities since additional data will facilitate the representation of appropriate climate processes in models. Hardware and software infrastructure is required to support these efforts and take advantage of petascale computer systems dedicated to climate science research goals. High productivity systems for multi-scale climate prediction will field hundreds of thousands of processors in a high-bandwidth, low latency interconnection fabric that seamlessly interfaces through parallel input and output operations with massive online and archival storage. The storage systems will be connected by high speed Internet connecting geographically distant centers in a computational grid that offers interactive analysis and data mining tools to an international community of researchers.

The BER investment in software for climate modeling is largely as part of the Climate Change Prediction Program (CCPP) and the SciDAC program. Since model formulation, building and testing require close coordination between climate scientists, mathematicians and computer scientists, a BER and ASCR partnership is natural and offers many opportunities for gains in scientific productivity. Software is the common currency for translating algorithmic and scientific hypothesis into computational experiments and studies. Climate models, such as the CCSM3, are sophisticated software projects that support research by a large community of scientists as well as major assessment studies such as the IPCC fourth assessment report.

The current DOE climate modeling program has deployed a *computational end station* for climate simulation at ORNL in collaboration with several DOE laboratories. Additionally, DOE is recognized as the international leader in advanced technologies to archive, distribute and analyze very large data sets generated by climate model simulations, as evidenced by the impact on the recent IPCC report by the multi-model database housed at LLNL. Other elements include a major ocean sea-ice model development effort at LANL, a model parameterization development testbed (also at LLNL), a world-leading capability in regional model development and application at PNNL, and extensive climate application simulations performed at LBNL NERSC. In addition, ANL, BNL, and PNNL competence in cloud and aerosol climate process understanding brings to bear additional required expertise.

The climate community has specific needs for the software infrastructure and facilities to support climate science. The needs flow from the day-to-day development, support for operational studies, and anticipation of future computer and software architectures. The needs associated with the day-to-day business of developing and maintaining a community code include a central coordination of the software repositories to provide version control, a distributed testing framework to ensure proper functioning on all supported systems, an open software toolset for diagnosing and evaluating problem reports, a workflow and data management infrastructure to track control simulations and results, stable programming environments from desktop to petaflop systems, robust and stable development, debugging and performance monitoring environments.

On an operational level, the software engineering discipline for high performance computing systems also requires a significant infrastructure investment to sustain production efficiency. Modern software is organized into layers with utility services and library support at the foundational layer, a middle layer that provides the tools and interfaces to an abstract modeling layer. Roughly speaking, the lowest level provides access to the machine particulars and is the source of much of the single processor performance efficiencies of the code. The middle layer implements parallel algorithms and data structures that minimize communication overhead and perform load balancing. This layer is responsible for the parallel efficiency on distributed memory computers. New numerical algorithms offer alternative parallel implementations that improve efficiency on parallel systems so an ongoing investment in algorithm research and evaluation is important to improving parallel efficiencies. The highest level, expresses the scientists modeling assumptions and formulations. Enormous modeling breakthroughs are possible by rethinking the basic assumptions. Indeed, this is necessary research as we move to incorporate multi-scale processes of clouds and biogeochemistry in an Earth system model. This layered organization cannot be strictly enforced because the highest level abstractions are not standardized (and should not be).

As we move forward to utilize tens to hundreds of thousands of processors over the next few years, the scalability challenge is forcing software architectures to expose higher degrees of parallelism. A significant additional effort is called for in this area: exposing parallelism for hundreds of thousands of processors requires re-factoring the modeling system code and using different data structures in the middle layer. Tools that support software analysis of dependencies and aid in the identification of parallelism are sorely missing. With computer languages and compliers lagging years behind advanced hardware (for example the cell processor has no scientific programming language supporting its use) the model development community has opted for reliable and portable, easily optimized languages like Fortran90, Fortran95 and gnu C++. The burden of development on scalable systems with these languages is increasingly problematic and requires a larger investment in software engineering support and personnel.

Many scientists have found the growing requirements to support the software on high performance computers as a distraction from the central scientific goals of improving climate models and answering fundamental questions about climate feedbacks and variability. This view is offset by the new scientific opportunities provided by dramatic increases in computational power. The issue is scientific productivity. What is needed is a software framework that not only scales from desktop to petascale, but also that supports multi-scale model development and process integration. The same modules that are used in a global climate simulation should be used for regional and site-specific process studies across bench to field to global spatial scales. This vision for a seamless modeling environment has only been realized in a few areas, e.g. column radiation models. As a closer connection with observational data and process studies is required to advance the science of regional climate prediction, the software must also become more closely integrated and supported across scales. Software will increasingly be required to support data assimilation and other data intensive frameworks like DOE’s CAPT activity.

These software frameworks will emerge as key bottlenecks to progress. An investment now would have important payoffs in the not very distant future.

DOE requires flexibility and rapid response of the software infrastructure to DOE specific requirements. DOE should gain a new level of direct control over the model development process and support a DOE Earth Systems Model that builds on the core physical modeling infrastructure of the Community Climate System Model, but focuses on the additional biogeochemical pieces. A significant investment to support the software development effort would be required along with an increased research emphasis in predictive, multi-scale earth system science.

There are three drivers for the development of software infrastructure in the long term. The first is the expected radical change in the nature of the computer hardware. Gains in aggregate performance going forward are expect to come from increasing the number of processing units on a chip (“multicore / manycore”), rather than increasing the clock speed of a processor. High-end systems built of such chips will not be amenable to either the shared-memory or distributed-memory programming models used to implement current production codes. The second driver is the non-incremental nature of the model and algorithm changes described above. Code development will require an aggressive and nimble exploration of the design space. This exploration will have to be done on the very high-end systems that we expect to be a moving target, since the changes in the algorithms will have to be evaluated at the high spatial and temporal resolutions that we expect them to be valid. The final driver is data. Increases in spatial and temporal resolution will also increase the sheer volume of data. In addition, the data is expected to be put to more uses by a much broader range of stakeholders.

The approach to these three problems is roughly the same: to design the high-level software tools that hide from the developer and user the low-level details of the problem, without foreclosing important design options. For dealing with the new hardware model, such tools will be new programming environments to replace MPI / OpenMP. For new algorithm development, the tools are software frameworks and libraries that provide a collection of high-level parallel algorithmic components (discretization libraries, data holders, solvers) from which new simulation capabilities can be built. A similar collection of visualization, data analysis, and data management components would provide capabilities in these areas. In all three cases, there is already a collection of methods and prototypes that have been developed by the mathematics and computer science research communities as part of the SciDAC program that would be form a starting point for such toolsets for climate modeling. These would have to be customized in collaboration with the climate community and hardened to support a production capability.

Facilities and Infrastructure

Currently, allocation of ASCR computer cycles is decided using a peer-reviewed proposal-driven process focused on breakthrough science. Such an allocation process not only requires scientists to undergo two peer reviews for their science, but can also place required program deliverables at the mercy of an external review process. One of the missions of the BER climate modeling community is to provide input to policy makers on the impacts of energy portfolio choices. This programmatic need drives a large fraction of the computer cycles used by BER climate modelers and involves assessments with ensembles of relatively coarse resolution models. These assessment products have firm deadlines and a well-defined product and simulation schedule. Future demands may require rapid turnaround in response to queries from policymakers. These everyday production simulations do not often fit into the INCITE paradigm of large-scale breakthrough science, yet are critical to programmatic deliverables. Currently, the community has been able to obtain resources for programmatic work by bundling these deliverables together with a few large-scale science simulations through a large Computational Climate End Station proposal, under which much of the assessment cycles are managed internally. While this strategy has been successful to date, there is a possibility that peer-reviewers in the future might favor large flashy results and our more routine program deliverables will be placed at risk in a proposal-driven process. ASCR management should work with other SC program offices to ensure computing capability for required programmatic work is being adequately provided.

Networking

To be useful, the data and software that underpin climate change research must be made freely available to global change researchers worldwide, in a manner that allows convenient access, analysis, evaluation, discussion, intercomparison, and application. Thus, we must necessarily plan for an infrastructure and collaborative environment that links centers, users, models, data, and resources on a global scale. The creation of such an infrastructure and environment is vital to the success of climate change research and critical for the impact sought by ASCR and BER. It demands continued investment in data management, software, networking and collaboration technologies.

The beginnings of such an effort have already begun with the Earth System Grid (ESG), which has its mission the construction of a universal high-performance seamless access point for petascale data and computing resources. This effort involves distributed data management and resources, high-bandwidth wide-area networks, and remote computing using climate data analysis tools in a highly collaborative problem-solving environment. It is already enabling 1000s of climate researchers worldwide to access >200 TB of data products from CCSM and IPCC simulations. The work of ESG can be leveraged to meet some fraction of the requirements needed to resolve current bottlenecks.

While ESG is a start, it is far from being complete. A significantly enhanced and more integrated system is needed to bring together simulations and experimental data from a variety of sources and a variety of sensors to accelerate global change studies. By so doing, we can enable a growing communicate of climate and impacts researchers to leverage these studies to gain insight into Earth science process, trends, and interactions, with the goal of answering new scientific questions.

Data management needs and current solutions

Climate model simulations are run on large supercomputers and typically produce nonstandard proprietary formatted data that are stored on tertiary storage. These data can reach terabytes in volume and are typically known only to one or two data producers who happen to be “in the know” concerning the data’s location. For other scientists, data is all too often essentially inaccessible. To obtain the data, they must first be made aware of the data’s existence, then find the right person to obtain the data. They must then download the data to their site, and convert it into the right format for their analysis tools. They must also organize and store the data for future use. These steps require far too much hand holding and manual labor: it can take literally months to obtain, process, and analyze a dataset. The process in and of itself is a bottleneck.

Harmonious data management requires first of all a standard output or community convention for processing and producing data output. Currently, a large part of the climate community modeling of physical oceanography, atmospheric sciences, and atmospheric chemistry has adopted the netCDF Climate and Forecast (CF) metadata conventions. Biogeochemistry and chemistry modelers are now working with the CF committee members to get their requirements into CF. Datasets that are CF-compliant are self-describing, in the sense that each data variable has an associated description of its standard name attributes, physical units, and spatio-temporal coordinate structure. CF metadata is highly detailed and contains information pertaining to the specific models and simulation scenarios. The file structure is regular, allowing a knowledgeable person to browse via ftp to find specific files. Quality assurance routines are also needed by the data producers to assure the data are correct and to make sure that the data are in the correct CF format.

Once data are in the right format (i.e., CF), then—regardless of which model, simulation, climate scientific application, or run—cataloging must occur. Distributed metadata catalogs must meet requirements for consistency and security of metadata and data information. These distributed catalogs must be flexible and meet the functionality and performance needs of a large and distributed user community. While metadata is stored in shared catalogs, the data itself is stored on a variety of storage systems at different sites worldwide. Individual researchers must be able to search, browse, and discover metadata and data regardless of physical location. Publication in this context means the act of putting data in “the database” (i.e., storing it somewhere, and recording its location in the catalog) and making it visible to others, while cataloging involves creating and storing the information about where a data set, file, or database entity is located in the distributed environment.

In the not too distant future, large coupled runs will produce much larger data sets. With this increased complexity of data, we must rethink our storage and retrieval paradigm. It will be impractical for most researchers to download more than a small fraction of climate simulation datasets for local analysis (indeed, it is already impractical today for many). Thus, if we want to allow any substantial use of these data, we must support new approaches such as large-scale server-side analysis, replication to multiple national or regional centers, and caching of popular simulation and derived data. We may want to include various of these processes (e.g., popular analyses, replication) in automated data generation pipelines. Overall, data creation, publication, and analysis processes must become distributed, more automated and closely integrated in terms of running models and directly archiving for immediate use.

Networking needs and current solutions.

ESG and related programs already make heavy use of ESNet and other networks, for example to transfer data from supercomputers to archives and from archives to users. Climate research demands on networks will grow yet further as data volumes increase, as systems such as ESG make data more accessible, and as data publication and analysis procedures become more automated.

Computational Facilities

The major climate modeling centers have established a modeling pipeline in which there are present generation workhorse models that are scientifically proven through peer-reviewed publications, next generation workhorse models in the process of being scientifically proven, and models being used to explore parameter space beyond the next generation workhorse. It is important for resources to be made available to provide adequate turnaround for all of these types of models, from production runs to debugging large, less-mature models. This is necessarily a mix of capability and capacity computing. Consequently, existing computational capacity continues to be inadequate in real terms, and via existing allocation mechanisms. Current demands continue to require enhancements to data management, migration and analysis mechanisms, which argues for attention to be paid to suitable storage hierarchy, bandwidth, support for workflow and analysis for climate science applications, which also provides for ways of dealing with both model and observationally generated data. Part of this involves making adjustments to optimally manage facilities for production, high-throughput debug, and analysis work. Priority needs to evolve toward providing stable environments which will enhance scientific productivity.

Data Storage Facilities

Needs help

Analysis Environments

Needs help

Collaborative tools and Technologies

As discussed earlier, climate science is necessarily distributed and collaborative. As interest in climate science continues to grow and its scope broadens to encompass issues of ecosystem and economic impacts, and the evaluation of mitigation and adaptation strategies, the number of participants will increase also. The overall productivity of researchers and the quality of the research output can likely be improved significantly by the use of advanced collaboration and workflow technologies such as the following.

Service oriented architecture and workflow
As data sources grow more diverse, so too will the range of data analysis procedures. Service oriented architecture and science workflow methods (a web application that combines data from more than one source into a single integrated tool) may play an important role as a means of enabling first the publication of diverse data sources and analysis procedures, and second the composition of services to create higher-level analysis procedures. Thus, for example, a user should be able to select climate model output data from multiple sources, compute statistics on that data, pass the original data and the computed statistics to a regional impacts model, and publish the results to a shared database. Needless to say, none of those steps are automatable at present.

Trust management and provenance
As the number of participants and the variety of computations performed grows, the need emerges to be able to document clearly the provenance (who, what, how) of derived data products. It will also be important to be able to manage who can consume what will sometimes be substantial amounts of computing, storage, and network resources. Trust management mechanisms must scale to far larger user communities than today.

Collaborative tagging
While standard metadata and data formats are vital, as noted above, human ability to describe data will always outstrip standardization. Collaborative tagging technologies have proved useful in other contexts as a means of communicating informal perspectives. They and other modern “Web 2.0” technologies are worth exploring to determine uses in climate research.

Notification
As the amount and variety of data grows, it becomes increasingly difficult for users to keep track of what is new and what has changed. Automated notifications to users with updates tailored to their research interests, as well as sophisticated discovery capabilities, will serve to enable collaboration (by allowing for the formation of communities of interest) as well as potentially reduce data management bottlenecks (by avoiding frequent checking for updates).

Visualization Technologies

Needs help

Institutional Issues

The DOE investment in software for climate modeling has been largely on the part of BER. The SciDAC2 program, which started in 2007, is the notable exception. A Scientific Application Partnership (SAP) (PI: Worley) is funded by ASCR dealing with the scalability of the Community Climate System Model (CCSM) as part of the SciDAC2 "Scalable and Extensible Earth System Model" project (PI: Drake and Jones) Since model formulation, building and testing require close coordination between climate scientists, mathematicians and computer scientists, the BER and ASCR partnership is natural and offers many opportunities for gains in scientific productivity. Software is the common currency for translating algorithmic and scientific hypothesis into computational experiments. Climate models, such as the CCSM3, are sophisticated software projects that support research by a large community of scientists as well as major assessment studies such as the 2007 IPCC AR4. The software engineering management must support scientific needs, production schedules as well as reliability and performance requirements. With a distributed team of developers and application scientists, the codes require full time coordination of gatekeepers and a scientific staff available to diagnose problems and provide solutions. The task of integration and coordination should be clearly designated and supported.

The development of new methods, especially new dynamical cores for ocean and atmosphere components, requires a concerted effort over several years by a small team. The steps for bringing new methods into consideration for production use are well delineated but difficult to traverse without becoming a climate domain expert. Mechanisms for mathematicians to be included in these joint ventures are needed.

Finally, the climate modeling enterprise in the DOE is increasingly driven by the need to obtain scientific results for policy makers in a timely fashion. In such an environment, the development of innovative models, algorithms and software must be managed as a project, as opposed to an open-ended research program, in order to have the desired impact. Some aspects of such an approach are well-understood, such as the need for planning, schedule visibility, and milestones. A more difficult problem is the potential dependence of success on delivering high-risk products in models, algorithms, and software on the required schedule. Many of these products, such as new discretization methods, or new programming models, represent non-incremental departures from the current methods used in production climate models, but may be necessary to achieve the goals of the project. Risk management in such a setting is requires careful planning and a close and continuing collaboration between the climate and math / CS communities.

SUMMARY

Need a balanced investment portfolio; no silver bullet that will immediately accelerate progress. Investment in computational infrastructure, basic science, computer science, and applied mathematics is important to the overall progress of climate change science.

Computational capability, albeit growing at a healthy rate due to ASCR investments, remains a bottleneck and should remain a high priority investment.

As the science and complexity of climate simulation grows, so will new technical and scientific challenges. Proactive investments in software, algorithms, data management, and other pacing items is strongly recommended.

Development of innovative models, algorithms and software must be managed as a project, as opposed to an open-ended research program, in order to have the desired impact. A more tightly coordinated effort in climate change science, via partnerships with ASCR and BER is highly desirable.