National Aeronautics and Space Administration

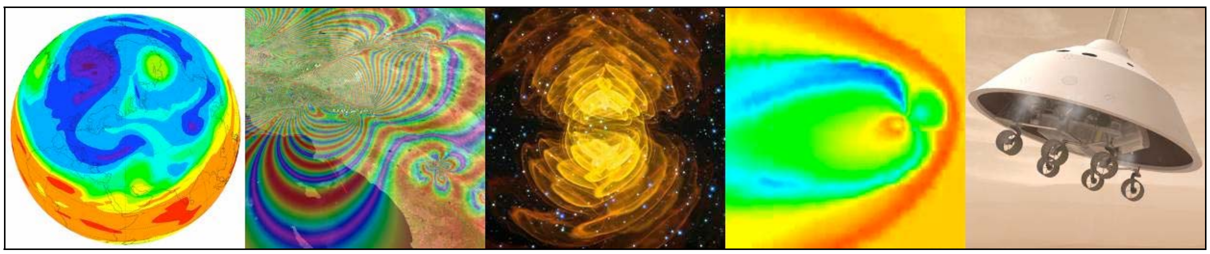

NASA Science Mission Directorate Computational Modeling Capabilities Workshop

The Inn & Conference Center
University of Maryland University College
College Park, MD
July 29–30, 2008

Final Report

Table of Contents

Executive Summary

The NASA Science Mission Directorate (SMD) Computational Modeling Capabilities Workshop was held July 29–30, 2008 at The Inn & Conference Center, University of Maryland University College. The purpose of this community workshop was to identify SMD science and engineering computing needs and evaluate them against current capabilities. The workshop was open to interested members of the NASA SMD community, but required an invitation. All SMD users of NASA High-End Computing (HEC) facilities were encouraged to request an invitation and participate. Workshop participants were organized into five discipline specific panels—Earth System Modeling and Assimilation, Solid Earth and Natural Hazards, Astrophysics, Heliophysics, and Planetary Science and Mission Engineering. Panel Chairs and Co-Chairs were recruited from among leading NASA modelers in each discipline, and charged with assembling a panel of experts from within their discipline community to identify SMD science and engineering computing needs and evaluate them against current NASA HEC capabilities. Other interested members of the NASA computing community were offered an opportunity to participate in the panels by requesting an invitation through the HEC Program website.

Participants were assigned to discipline panels based on their expertise and prior modeling experience. Each panel was asked to collect information from the groups with major HEC usage, identify the major computational applications that are driving their discipline's computing needs, and project their discipline's requirements for the 2013 timeframe. The panel reports detail these application drivers and include prioritized recommendations to the HEC Program and SMD based on their evolving computing needs. Specific recommendations are included in the panel reports. A technology committee made up of NASA HEC experts participated in the panel discussions, and was tasked to review the panel reports to extract common and crosscutting requirements. This committee developed its report based on the panel discussions and the observations captured in each panel's report.

The following sections contain highlights from the panel reports. The reader is encouraged to review the discipline panel reports individually to draw from their specific presentations of scientific rationale and mission relevance.

Earth System Modeling and Assimilation Panel

Earth science models codify our understanding, gleaned from the satellite observations, of the many processes that make up the complex Earth system. They help to quantify the interactions and the balances between the various components acting on a wide variety of scales in both space and time. NASA's model and assimilation systems support international programs such as the World Climate Research Program (WCRP), the World Weather Research Programme (WWRP), UNESCO's Intergovernmental Oceanographic Commission (IOC), the periodic Scientific Assessments of Ozone Depletion carried out on behalf of the World Meteorological Organization (WMO) and the United Nations Environment Programme (UNEP), and the assessment of climate forcing factors on behalf of the Intergovernmental Panel on Climate Change (IPCC); and national programs such as the Climate Change Science Program (CCSP), the U.S. Weather Research Program (USWRP), and the U.S. Integrated Earth Observation System that supports the international Global Earth Observation System of Systems (GEOSS).

Increased resolution and complexity in these models, and their associated assimilation systems, is driving the requirements for HEC resources. The panel projects a need for a 20-fold increase in CPU-hours over the next 5 years to address the critical problems in climate change and prediction, the implementation of high-resolution 4D variational methods in weather analysis and prediction, the contribution to the IPCC assessment using fully interacting complete atmospheric chemistry at 1/2° resolution, and the expected incorporation of non-hydrostatic dynamics and physics into the next-generation GEOS model. With this production requirement growth comes a comparable growth in data storage for both short-lived data (~1 PB/day with 30 day lifetime) and archival data products (37 PB/year). These short-term products, typically used for immediate analysis, visualization, and various post-processing data reductions, drive the need for 3 to 5 PB of fast online data storage, and substantial (40 MB/s) sustained network bandwidth to the end user of these products.

This expanding volume of data will require more analysis and visualization support at the computing center, since moving this data to individual researchers' sites is not likely to be feasible. Computing center staff will need to be available to help users post-process the model and analysis results. Remote visualization tools and analysis tools, along with tools for managing the data volume and enabling remote access to the data, will need to be supported by the computing center.

Solid Earth and Natural Hazards Panel

The next great revolution in Earth sciences will involve development of predictive models of complex, interconnected solid Earth processes. These models encompass polar science, geomagnetism, crustal deformation science, and related applications. For these models to be successful, particularly for an understanding and forecasting of hazards, high-resolution, global observations with real-time or near-real-time data streams and processing will be required. Integrating the projected huge quantities of data and information into forecast models will require that information technology resources be developed in concert with advances in sensors and detection capabilities.

The solid Earth community currently uses a spectrum of computing resources, spanning highly capable local clusters to the NASA HEC computing centers. Substantial growth in local cluster capability is projected to be needed in the next 5 years. HEC center requirements are projected to increase by the same 20-fold as Earth System Modeling. Data storage requirements are expected to increase at a similar rate, with the exception of geodynamo ensemble runs, which could generate 80 PB of archival data in the 2013 timeframe. The distributed nature of the model work envisioned by this community will also drive a substantial (30 MB/s) network bandwidth requirement between the users and the computing assets.

This community is distinct from the other community panels in this workshop in that the computing done today is highly distributed. The emerging "Cloud Computing" paradigm and an SOA (service-oriented architecture) are expected to become standard practice for much of solid Earth science.

Astrophysics Panel

Computational science provides critical support for NASA's Great Observatories and other missions. Increasingly, our understanding of the universe is codified in computational models. These models span the range from cosmological structure to galaxy and star formation to planet formation and dynamics. The simulation methods include N-body dynamics, hydrodynamics, and radiation transport, with consideration of full general relativistic physics.

The panel anticipates a five- to 10-fold increase in CPU-hour requirements for these modeling efforts over the next 5 years. This is driven largely by the need to increase resolution to adequately cover the required spatial scales. A 10-fold increase in archival storage requirements is consistent with this scale-up of resolution. This community has the same need for improved visualization and analysis capabilities as the other panels, and is particularly sensitive to the ability to access large numbers of processors for long execution times. With model complexity increasing, there is also a need for expert assistance from computing center staff to improve the runtime performance and scaling of these codes.

Heliophysics Panel

The progress in understanding the overall dynamics of the Sun-to-Earth or Sun-to-planet chain has created an increasing desire to describe the relevant physical processes quantitatively. At the same time, the growing national need to forecast and describe space weather mandates a transition from discovery and qualitative scientific description to the deep level of quantitative understanding required to forecast harmful space weather effects. Models are now widely used by the research community to assist in the scientific analysis of spacecraft-provided datasets, as well as in mission planning and conception. Furthermore, modeling has evolved into a core element of programs aiming at the development of new space weather forecasting capabilities.

From magnetic reconnection to space weather prediction, the modeling challenge spans an enormous range of space and timescales. At the high end of computation, the panel anticipates a requirement for a thousand-fold increase in sustained computing throughput for the large-scale problems (from sustained teraflops capability today to sustained petaflops capability in 2013). Support for large data volumes is a concern, both for archival storage and post-processing analysis. Visualization tasks will require tools that can handle large data volumes at the sites where the data resides, as the typical scientist's local environment is unlikely to be capable of handling these data volumes.

The panel also emphasized the need to support both high-end computation and local computing capability. The community currently relies on community models accessible to a large number of heliophysics researchers. This mode of operation needs to continue to be supported in the future.

Planetary Science and Mission Engineering Panel

In recent years, some of the most stringent requirements in engineering modeling, and the necessity of HEC to provide adequate computational capability, have arisen from certain development and operations categories of planetary science missions. In particular, flight operations considerations for high-autonomy interplanetary spacecraft during critical events are known to benefit from HEC in a real-time "go/no-go" decisional context. Such events include spacecraft approach to another planet to enter its atmosphere, autonomous rendezvous with an object or other spacecraft at another planet, or preparation to initiate maneuvers to descend to the surface of a previously unmapped object such as a comet or asteroid. For example, the Phoenix mission was significantly enabled by use of high-fidelity Entry, Descent, and Landing (EDL) simulations that ran during development on a shared institutional supercomputer at JPL, then during flight operations on a several-hundred-compute-node, mission-dedicated supercomputing cluster. This computational approach is being followed, and expanded upon, by the Mars Science Laboratory (MSL) in its final years of development before launch.

The use of supercomputing for mission engineering calculations is qualitatively different from the science uses described in the other panels. The Phoenix and MSL examples are dominated by Monte Carlo calculations that make use of commercial software and require rapid turnaround in order to support time-critical decisions during mission operations. Though the CPU capacity is small by science modeling standards, this near real-time support requirement and mission-critical operations have required dedicated access to 1,000 processors during the preparation for mission readiness review and EDL execution. With MSL as the exemplar for future applications in Mission Engineering, the panel projects an order-of-magnitude increase in the CPU capacity required for missions in the 2013 timeframe. A similar increase in storage capacity is required.

These CPU and storage requirements are small fractions of the other panels' requirements, and as such, are not drivers for HEC resource provisioning. Mission Engineering's unique requirements are in the areas of programming tools, near real-time turnaround of computations, and the availability of an equivalent capability backup system during flight operations. MATLAB, Python, Java, and Ruby are in use today, and need to be available on future platforms. The ability to migrate the current code base (even though it is evolving) to other HEC platforms requires a well-tested, configuration-controlled, integrated set of commonly available third-party software. The environment cannot be allowed to change during and leading up to flight operations. And the identical environment needs to be available on the backup capability.

Crosscut Analysis

Aggregating the CPU-hour requirements in 2013 across all of the panels yields a projected capacity requirement on the order of 1 billion CPU-hours. This requirement is dominated by Earth System Modeling (Figure 1), with the other discipline requirements being an order of magnitude smaller. Figure 2 below compares that requirement to the historical SMD usage on the NASA Advanced Supercomputing (NAS) and NASA Center for Computational Sciences (NCCS) facilities. The FY2005 utilization represents an opportunistic jump in SMD utilization that coincides with the delivery of Columbia to the NAS facility (before allocations were enforced across all of the mission directorates). Note that the vertical scale is logarithmic, exaggerating the relative proportions of the NAS and NCCS usages.

Based only on the historical trend, the projected CPU-hour requirement in 2013 might appear to be an over estimate. However, it is important to note that in the most recent year, SMD utilization was capacity limited—i.e., SMD completely used all the CPU cycles available to it. With hours requested currently running at more than twice the available capacity, it is clear that there is pent-up demand for HEC resources.

In general, the panels indicated that they have adjusted to the distributed memory parallel architectures that dominate the marketplace. No one is counting on substantial improvement in processor speed. The disciplines are prepared to live in an environment where improved performance must come from more parallelism within their applications. However, most of the panels identified the need for additional technical support from the HEC centers for scaling their applications to higher processor counts.

Each of the panels expressed an interest in experimenting with the new hybrid computing architectures that are emerging (computational accelerators such as general purpose graphics processing units integrated with traditional CPUs), but all have reservations about the difficulty of programming such architectures, and their ability to support the current code base without substantial code rewrites.

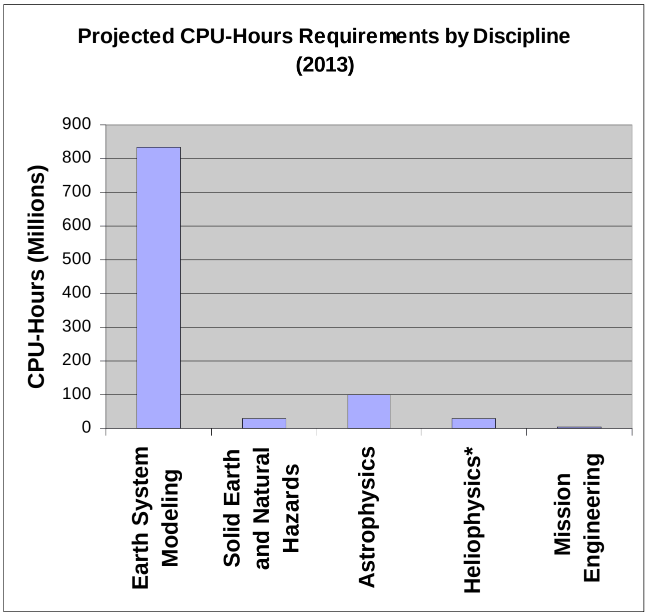

Figure 1: Projected CPU-hour requirements by discipline. * Heliophysics stated a requirement for a sustained petaflops capability for simulations lasting a week. The CPU-hours requirement is derived from this requirement. The Mission Engineering requirement is in the low millions, and does not show up on this scale.

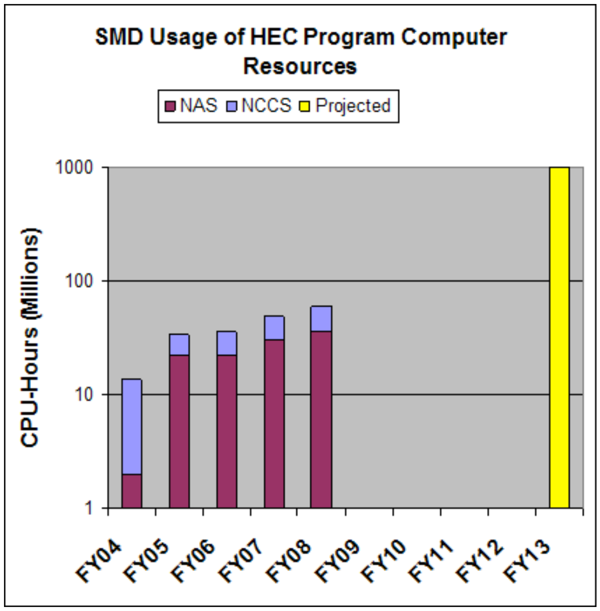

Figure 2: Comparison of historical SMD usage to projected requirement in 2013.

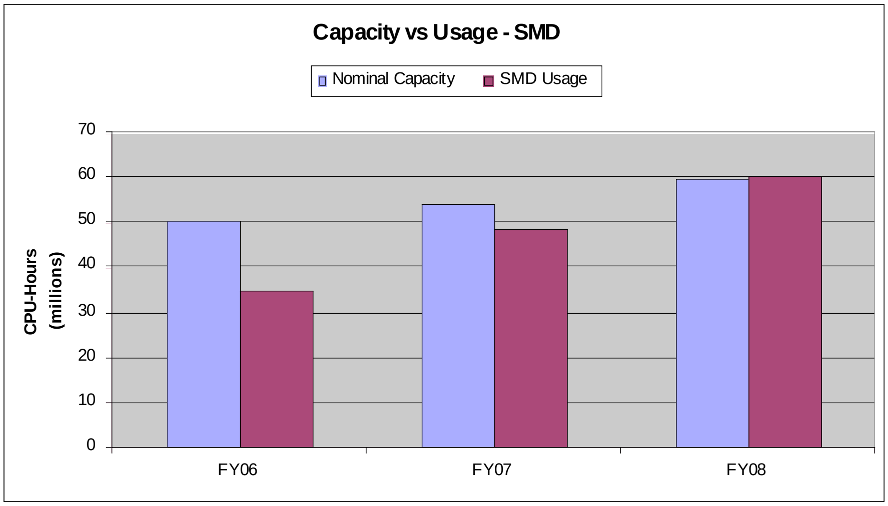

Figure 3: SMD usage of HEC resources vs. nominal capacity. Nominal capacity is now oversubscribed.

Storage requirements from the panels are also dominated by Earth System Modeling, which is again about an order of magnitude larger than the total of the other disciplines. This storage capacity is handled by tape robots today, and is expected to continue to be handled in the same manner in the future. However, the panels identified the need for substantial short-term (3 months) storage on online disks for efficient utilization of the results of the modeling. This disk capacity is projected to be in the 3 to 5 PB range. The online storage capacity today is too small to allow users to keep the files needed for short-term analysis on fast media, and is impacting the productivity of the research.

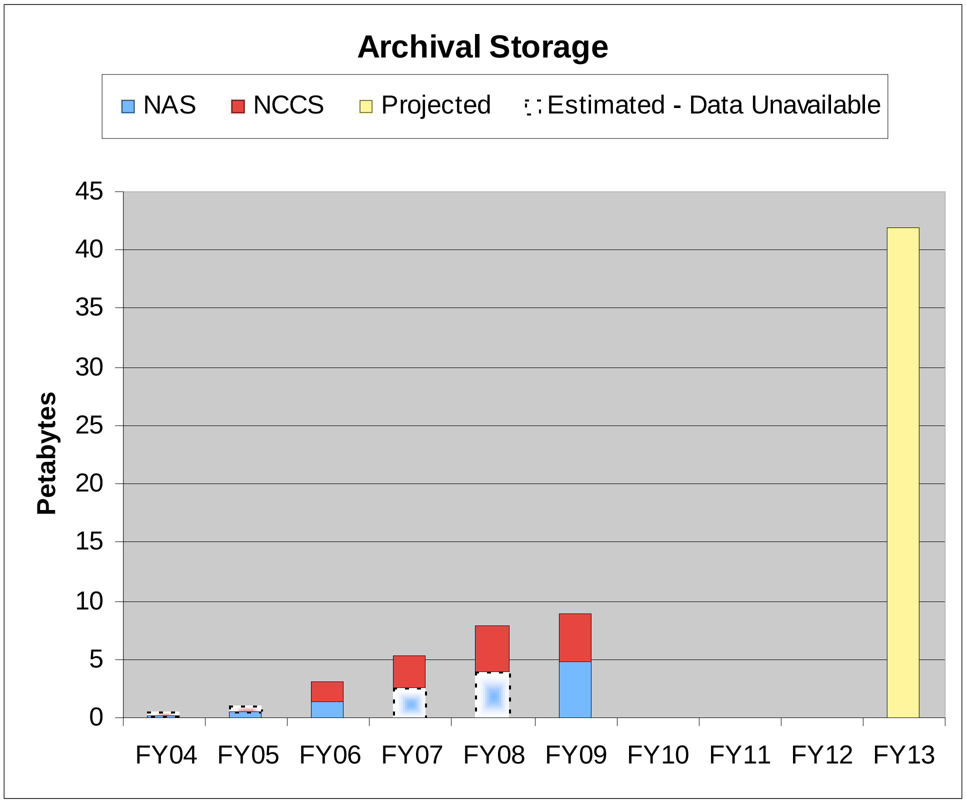

Figure 4: Comparison of historical archival storage volume to projected SMD requirement.

The Earth System Modeling and the Solid Earth disciplines both have substantial short-term and long-term archival storage requirements, and a distributed community of users of the model output. In order to effectively exploit the data, sustained Wide Area Network (WAN) bandwidth is required at about 40 MB/s (~400 Mbps) to the end user. WAN speeds of 1 Gbps are becoming common these days, but researchers indicated that they do not see this level of sustained throughput even with a 1 Gbps connection. The HEC facilities will be called upon to assist users with end-to-end performance optimization.

Besides the basic hardware requirements, the panels identified the need for productivity improvements in the capabilities and usability of the HEC user environment. The specific needs identified are in the areas of computation management (frameworks and workflow); data management (searching, subsetting, metadata management, etc.); remote data analysis and visualization; and distributed environments (managing computation that spans the computing center and the local user facilities).

Finally, the panels expressed their concerns that the complexity of NASA models is challenging their ability to optimize performance and improve scalability (in order to exploit the thousands of processors required to get performance in the future). The panels recommend that the HEC Program increase the technical assistance provided by the computing centers for optimizing code, improving application scalability, and training the user community on the available tools, best practices, and performance engineering so that users can make the best use of the facilities.

Summary

There is substantial demand today by SMD researchers for HEC computing resources, and this demand is expected to grow by an order of magnitude over the next 5 years. The SMD discipline panels all recommend continued investment in HEC capability and capacity increases *for SMD alone* that amount to approximately three to four times the current installed base *shared among all of the NASA Mission Directorates*. With this increase in hardware capability also comes the need for additional professional support to the researchers by the computing centers for code performance improvements, and enhancements in the user environment to enable the efficient management and analysis of the models being run.

The specific recommendations from each panel may be found in their respective sections in the body of this document.

Introduction

NASA's High-End Computing (HEC) Program (http://www.hec.nasa.gov) organized the NASA Science Mission Directorate (SMD) Computational Modeling Capabilities Workshop to determine the HEC capabilities and infrastructure investments required to enable the goals defined in the *Science Plan For NASA's Science Mission Directorate 2007–2016*. As was done at a predecessor workshop in 2002[1], the 2008 workshop sought community input as SMD identifies the capabilities and investments that will be necessary to enable modeling and analysis for supporting the Science Objectives and Outcomes presented in the Science Plan. The workshop was open to interested members of the NASA SMD community, but required an invitation in order to limit attendance to the capacity of the meeting room. All SMD users of HEC facilities were encouraged to request an invitation and participate. In the end, all requests for invitations were honored, and the meeting room was filled to capacity. Attendees' contributions will greatly assist the guidance and subsequent advocacy of the computational areas essential to the achievement of SMD's science goals.

Workshop Purpose

The purpose of this community workshop was to identify SMD science and engineering computing needs and evaluate them against current capabilities. Identifying and quantifying these requirements is a key step in prioritizing HEC investments in this decade and the next.

In the context of this workshop, computational capabilities and infrastructure span the hardware, software, and human capital needed to enable the effective use of remotely sensed data in complex SMD simulations and models, as well as in the engineering analysis and design of future missions. Thrusts of this effort will enable interdisciplinary science, engineering, and applications scenarios employing linked or nested modeling components.

Workshop Process

The SMD Computational Modeling Capabilities Workshop was held on July 29–30, 2008 at The Inn & Conference Center, University of Maryland University College.

Workshop participants were organized into five breakout sessions: Earth System Modeling, Solid Earth and Natural Hazards, Astrophysics, Heliophysics, and Planetary Science and Mission Engineering. Each breakout session was tasked to:

- Identify the SMD computing capabilities required to achieve the discipline objectives and outcomes in the Science Plan.
- Evaluate current SMD computing capabilities against these requirements.
- Determine and quantify what gaps in capabilities exist.
- Determine which gaps should be addressed by HEC capability and infrastructure (including software tools) investments.

[1] In 2002, the NASA Earth Science Enterprise Computational Technology Requirements Workshop assessed the Enterprise's needs for computational technology development. These results were factored into SMD strategic planning and proved valuable in prioritizing and focusing research and development efforts. Workshop results are available at: http://ct.gsfc.nasa.gov/ese_ct.results.html

Each breakout session was led by invited panelists who are experts in the session focus. Breakout sessions were intended to develop requirements and identify gaps and potential areas for investment. A plenary session on the second day summarized the results of these discussions. Invited panelists then meet in closed sessions to develop their reports and synthesize a crosscut of requirements across the panel disciplines, resulting in this document. Table I-1 identifies the panel leadership. A list of participants is included at the end of each panel's report.

Discipline Panel	Name	Organization
Earth System Modeling	Michele Rienecker	NASA/GSFC
	Bill Lapenta	NASA/MSFC
Solid Earth & Natural Hazards	Andrea Donnellan	NASA/JPL
	John LaBrecque	NASA/HQ
	Geoffrey Fox	Indiana University
Astrophysics	Joan Centrella	NASA/GSFC
	Jim van Meter	NASA/GSFC
	Michael Salamon	NASA/HQ
Heliophysics	Michael Hesse	NASA/GSFC
	Aaron Roberts	NASA/GSFC
Planetary Science and Mission Engineering	Michael Lisano	NASA/JPL
	David Skulsky	NASA/JPL

Table I-1: Panel Leadership

Following the workshop, the discipline panel leaders developed their individual panel reports and recommendations from the material gathered and discussed in each breakout session. Once these reports were completed, a technology panel reviewed the reports to extract computing technical requirements and develop a crosscutting set of recommendations that reflect the computational requirements that are common across the SMD disciplines and identify any discipline-unique requirements. This panel did not hold a separate technology session—panelists were embedded in the discipline panels so that they were fully aware of the discussions and could provide technical expertise as needed. Technology panelists were all HEC experts drawn from the NASA centers most involved with SMD HEC applications and the provisioning of HEC resources.

The following sections are the reports with findings and recommendations of the five discipline panels, and the technical panel report representing the crosscutting analysis of the discipline reports.

Workshop Organizing Committee

Tsengdar Lee, NASA/HQ, HEC Program Executive (Workshop Co-Chair)
Azita Valinia, NASA/GSFC (Workshop Co-Chair)
James Fischer, NASA/GSFC
Thomas Clune, NASA/GSFC
Robert Ferraro, NASA/JPL
Jarrett Cohen/NASA/GSFC/GST
Angela Taylor/Harris

Earth System Modeling and Assimilation Panel

Science Drivers

NASA's primary role in Earth Science is to innovate global observations of the Earth system from space. The Science Mission Directorate (SMD) pioneers the scientific use of these satellite measurements to improve understanding of the Earth system and to advance benefits to society. Earth science models codify our understanding, gleaned from the satellite observations, of the many processes that make up the complex Earth system. They help to quantify the interactions and the balances between the various components acting on a wide variety of scales in both space and time. Models and assimilation systems are the tools that synthesize the diverse array of information from many satellites and bring that information to bear on improving prediction of: weather and air quality; future climate change and its impacts; changes in atmospheric composition and terrestrial and marine ecosystems; and important phenomena that contribute to climate variability, such as changes in the water cycle, ocean circulation, and El Niño and its impacts. As a result, improvements in the accuracy of Earth science models are the end products of NASA research that most directly impact human society. Examples include narrowing the uncertainty of global warming and sea-level rise expected this century, predicting the likelihood of rapid changes in climate due to the catastrophic collapse of the major ice sheets, forecasting the likelihood of changes in water availability or extreme weather events with changes in climate, predicting the rainfall impact of an El Niño event expected to arrive a year from now, or improving the accuracy of the timing, location, and strength of hurricane landfall next week.

Most of NASA's new observations are targeted at the most poorly modeled Earth system processes. Sea-level changes are dependent on the response of polar ice sheets and the deep ocean to global warming. Climate sensitivity is dependent on unknown feedbacks in radiative energy modulated by changes in cloudiness. Long-term temperature changes are dependent on the radiative impact of aerosols. Climate sensitivity also depends on complex connected physical, chemical, and biological processes in the atmosphere, land, and ocean. Feedbacks between these processes that could produce large and undesired responses to perturbations resulting from human activities are, as yet, unknown. Many challenges remain on all timescales; examples include the ability to model Eastern Pacific stratus clouds, monsoon and intra-seasonal variations such as the Madden-Julian Oscillation, the diurnal cycle of precipitation, ocean mixing processes, and phenomena at high latitudes, where the impact of global change is most evident.

The tie between NASA modeling efforts and satellite missions provides the context and justification for the numerical experiments that define the computing requirements for Earth system modeling and assimilation. Models and assimilation systems are tools, like the instrument algorithms themselves, essential to realizing the value of the nation's investment in satellite technologies. Currently there are 15 missions in orbit. Within the timeframe addressed by this report, we expect to see the launch of OCO (January 2009), Glory (2009), Aquarius (2010), NPP (2010), LDCM (2011), SMAP (2012), NPOESS (2012), and GPM (2013); and ICESat II (2015) will be on the horizon. Planning and prioritization of other missions recommended by the National Research Council (NRC) Decadal Survey for Earth science research will also be required during this period. In addition to identifying critical missions, the NRC Decadal Survey also identified the importance of "models and data assimilation systems that allow effective use of the observations to make useful analyses and forecasts."

Modeling and assimilation are also core elements of the science programs in five of NASA's Earth Science focus areas. NASA's 2007 Science Plan identifies overarching goals for these areas (not including Earth surface and interior, which is addressed by another panel): (1) Understand and improve predictive capability for changes in the ozone layer, climate forcing, and air quality associated with changes in atmospheric composition; (2) enable improved predictive capability for weather and extreme weather events; (3) quantify global land cover change and terrestrial and marine productivity, and improve carbon cycle and ecosystem models; (4) quantify the key reservoirs and fluxes in the global water cycle and improve models of water cycle change and fresh water availability; and (5) understand the role of oceans, atmosphere, and ice in the climate system and in improving predictive capability for its future evolution. Models are key to achieving these goals and to integrating the science across the focus areas to explore, discover, and predict the Earth as it acts as a system.

NASA's model and assimilation systems support international programs such as the World Climate Research Program (WCRP), the World Weather Research Programme (WWRP), UNESCO's Intergovernmental Oceanographic Commission (IOC), the periodic Scientific Assessments of Ozone Depletion carried out on behalf of the World Meteorological Organization (WMO) and the United Nations Environment Programme (UNEP), and the assessment of climate forcing factors on behalf of the Intergovernmental Panel on Climate Change (IPCC); and national programs such as the Climate Change Science Program (CCSP), the U.S. Weather Research Program (USWRP), and the U.S. Integrated Earth Observation System that supports the international Global Earth Observation System of Systems (GEOSS).

Models and Applications

The Earth system modeling tools range from comprehensive, global whole-Earth system models to local, more process-oriented models—models of the ocean, atmosphere, land surface, sea-ice, ice-sheet, marine and terrestrial ecosystems, and atmospheric and oceanic chemistry. Assimilation tools use the theoretical approaches of estimation science to merge model predictions and observations. Today's assimilation systems use complex variational and ensemble approaches and require the highest-performance computing available.

The models and applications can generally be summarized as:

- Global atmospheric data assimilation systems to generate specialized products to support physical retrieval algorithms by NASA instrument teams; study and optimize satellite data's impact on global weather, air quality, and climate prediction; contribute to observing system science, including future mission planning and prioritization; and, through retrospective analyses, provide a climate data record of essential climate variables. Systems are emerging that include aerosols, carbon species, and reactive gases.

- Global ocean, ice, and land data assimilation systems to generate specialized products; study the role of ocean, cryosphere, and land processes in climate; initialize the slow components of the climate memory; resolve the transport of carbon species, nutrients and biota; and provide a climate data record of essential climate variables.

- Coupled ocean-atmosphere-land-sea-ice models to simulate and predict climate at subseasonal-to decadal timescales, optimize the use of satellite data in ocean and land surface models to enhance predictability, and study climate-weather interactions.
- Coupled chemistry/biogeochemistry climate models to support simulation and prediction of ozone hole recovery, investigate chemistry-climate feedbacks, investigate carbon cycle feedbacks, and provide information on climate forcings and climate change projections.
- Mesoscale regional models to perform high-resolution weather prediction and pollution transport, and understand hurricane formation and local atmosphere-land interactions.
- Prototypes of next-generation systems to resolve processes on fine scales.

Meeting the scientific challenges confronting us and improving predictive skill will require increased model resolution, development of an integrated Earth system modeling and assimilation capability, and improved assimilation techniques to optimize the use of current and new high-resolution observations. These developments will require large increases in the total computing capacity for concurrent images and also in the capability to run single large images.

The Need for Increased Resolution

Increasing computing power has allowed us to increase the spatial resolution with which models represent motions in the atmosphere and ocean and surface properties and topographic effects in the land surface. In numerical weather prediction (NWP), horizontal resolution has increased 10-fold over 40 years, from 200-km to 20-km grid sizes in the horizontal, and the benefits for prediction have been dramatic. Still further increases in resolution are needed for better predictions of, for example, hurricane intensity, and for making the most of the high-resolution data now available from satellites such as CloudSat and CALIPSO and data anticipated from future satellites such as GPM. Resolution is also needed to remove or reduce the need for the parameterization of deep convective clouds (100 m to a few kilometers), an issue that has remained one of the most contentious and unresolved problems in global models for almost 40 years.

Climate modelers have been less aggressive than the NWP community in increasing resolution, choosing instead to add more complexity and make more and longer experiments; but even in climate applications, the resolution of the atmospheric components has increased some four-fold, from 400-km to 100-km grid sizes, with a greater increase in ocean models. With the release of the Fourth Assessment Report (AR4) of the IPCC, the climate community is now focusing attention towards robust projections of regional impacts of climate change and of changes in the water cycle (in addition to other problems), requiring a more aggressive approach to increased resolution in climate models. With the increased resolution of climate models and the longer historical record now available from satellite data, there is also increased attention towards identifying the interactions between climate and extreme weather events.

Resolution is particularly important for ocean and cryosphere processes. For example, everything we know about ocean eddies suggests that their property fluxes can accumulate, changing the ocean measurably and importantly from what it would be if eddies were absent. Narrow boundary currents make major contributions to scalar property (heat, fresh water, carbon, oxygen, etc.) transports that are of central interest for climate; these boundary currents are not parameterizable and, until they are resolved, there will always be doubts that the ocean model is simulating their property transports realistically. Similarly, recent NASA observations and studies suggest that improved representation of the small-scale behavior of sea ice is critical for understanding the recent evolution of the polar oceans and for predicting its behavior in future climate change scenarios.

Arguably, the land surface shows higher levels of spatial heterogeneity than the atmosphere and ocean. Variations in topography, vegetation, and soil type, for example, significantly modify local hydrological behavior even over scales of tens of meters. Given this variability, global models will, for the foreseeable future, continue to require some parameterization of subgrid-scale land surface processes. Even so, increases in land surface spatial resolution can mitigate many of the grosser errors associated with these parameterizations and can thereby improve their overall performance.

The Need for Increased Complexity

Increasing computing power has allowed us to add complexity, such as reactive chemistry in the troposphere and stratosphere, the simulation of aerosol distributions and their effects on other parts of the system, and models of biogeochemical cycles. Detailed simulation of biogeochemical cycles requires the inclusion of dynamic vegetation, terrestrial, and oceanic ecosystems; the fluxes (energy, water, nutrients, etc) that affect these ecosystems; and the carbon transport within and across components. Such Earth System Models (ESMs) include a full treatment of the carbon cycle and are needed to model and understand carbon cycle feedbacks and links, such as those between mineral dust aerosols, land surface changes, and the eventual deposition of minerals into the oceans, where they may have an impact on biological processes and carbon dioxide balance. Feedbacks between physical, chemical, and biological processes in the atmosphere, ocean and land surface could produce large and undesired responses to perturbations resulting from human activities. The Working Group (WG1) report of the IPCC's AR4 recommends that "*models that attempt to perform reliable projections of future climate changes should account explicitly for the feedbacks between climate and the processes that determine the atmospheric concentrations of greenhouse gases, reactive gases, and aerosol particles.*"

Increases in complexity are also needed for the investigation of abrupt climate change. Current climate modeling development includes fully dynamic ice-sheet models and ice-shelf models needed to assess the rate and magnitude of sea-level rise due to rapid ice-sheet melting and dissipation due to dynamical processes in ice streams and large outlet glaciers. However, these models are still maturing, and it will be a few years before we expect them to be routinely and effectively included in climate models. Other possible sources of abrupt change are: reduced carbon absorption, methane emission, and rapid changes in circulation of the ocean and/or atmosphere.

The Importance of Assimilation

Data assimilation provides powerful constraints on predictive models by providing realistic initial conditions from which to start the predictions. In assimilation, models are used to synthesize diverse in-situ and satellite data streams into a single product (analysis) that combines the strengths of each dataset and of the model itself. The need to generate initialization fields for NWP has driven the direction of atmospheric assimilation development. However, both ocean and land data assimilation are now mature. Assimilation for atmospheric constituents and ocean biology are emerging. The pathway forward now leads us to integrated or consistent analyses across the components and within components (the physical state with biogeochemistry). This development is important for improved initialization of longer-term prediction systems and for budget studies particularly on the climate timescale. Assimilation analyses can be used to infer unobserved variables and to provide a time series of the essential climate variables. Assimilation also plays an essential role in defining observing system requirements through observing system experimentation.

An important outcome of the development of ESMs, with the separate components working together, is the ability to undertake a consistent analysis across these separate components—an Integrated Earth System Analysis (IESA). The Interagency Working Group for Climate Variability and Change has identified development of an IESA as one of the top priorities for the CCSP. The CCSP recognizes the IESA as "*a fundamental prerequisite to understanding the coupling of, and feedbacks within, the Earth system*" and "*fundamental to advancing climate prediction capabilities, whether on seasonal or multi-year to decadal timescales*" [from *Our Changing Planet 2009*].

HEC Requirements

The HEC requirements for NASA's Earth modeling and assimilation activities were estimated by collecting information from the groups that currently dominate the HEC allocations. These groups also provided their view of the challenges, current and future, to be faced in using HEC resources to meet their science goals.

For the purpose of outlining the advances in computing technology required to enable our science goals for 2013, we focus on just a couple of these applications—the problems of weather prediction (Focus Areas: Weather, Water and Energy Cycle), short-term climate prediction and climate change projections (Focus Areas: Climate Variability and Change, Water and Energy Cycle), and chemistry-climate interactions (Focus Areas: Atmospheric Composition, Carbon and Ecosystems). This approach is justified by the fact that the types of problems to be solved for these areas are not substantially different from those for other areas. The challenges and strategies for the future are highlighted in the discussion of next-generation systems, epitomized by global, high-resolution, non-hydrostatic atmospheric models. Similar challenges will pertain to global high-resolution ocean models needed to resolve mixing, convection, restratification, and transport by eddies and boundary currents.

The Weather Prediction Scenario

Assimilation analyses of the Earth's environment have many applications. One of the few ways to evaluate the quality of analyses is through their impact on prediction skill. This is the standard metric for meteorological analyses. Today, the U.S. is lagging behind Europe in the quality of our weather forecasts. The superiority of ECMWF is clearly related to the implementation of four-dimensional variational (4DVar) assimilation. One of the important aspects of 4DVar from a NASA perspective is that it enables better use of satellite data. Inclusion of the fourth (time) dimension in 4DVar is required to get the full benefit of satellite measurements with high temporal resolution. It is also key for the assimilation of precipitation (rain rate) data because the extraction of tendency information from the data requires that tendency to be simulated by the observation operator used in the assimilation. In addition to a superior analysis system, much of ECMWF's success in weather forecasting can be tied to use of a high-resolution model, which leads to a better congruence with observations and a better acceptance by the model of the observations in the assimilation and the subsequent forecasts.

Thus, the implementation of a high-resolution 4DVar atmospheric assimilation system is critical to making best use of NASA's satellite observations, including the upcoming GPM, and also to support NASA instrument team retrievals with the best analysis products available. Having a state-of-the-art system will also be important for modeling and assimilation science to be able to contribute to the design of future missions through Observing System Simulation Experiments (OSSEs). The goal is to increase the resolution over time as more processors and larger memory become available while maintaining a throughput of several days per wall-clock day.

The global 4DVar implementation planned for 2009 will be at 1/4° resolution using the cubed-sphere configuration of the GEOS-5 atmospheric general circulation model (AGCM). It will include a replay analysis of Global Modeling Initiative (GMI)-based chemistry with about 60 tracers. The expected throughput of 2.5 days per wall-clock day will require 1,024 cores. The plan for 2013 is for a 1/8° meteorological analysis with chemistry at 1/4°, requiring 4,096 cores. The expected throughput of 1.25 days per wall-clock day is at the margins of acceptability.

Currently, about 4 GB of data are input for each 6-hour analysis cycle, dominated by AIRS data. This will grow to at least 10 GB in 2009 with the inclusion of IASI and other sensors. With each new hyperspectral sensor, the data volume will grow considerably, up to about 15 GB in 2013. However, this ingest is not a rate-limiting step. Rather, the I/O associated with restarts and products places a heavy burden on throughput. For 4DVar, this increases with the need to store the model trajectory. Currently, we use about 1 GB per core, but we plan to increase to 2 GB per core where it is available.

Typically, two sets of analyses and 5-day forecasts are run daily for the entire year. One set is operational, generating real-time products for NASA instrument teams and field campaigns. The other set is a test of proposed science improvements. The throughput for real-time streams to support missions is a major consideration when deciding on the model configuration and processor request. Naturally, queues and schedulers are a critical issue. Development and science runs are usually run at coarser resolution to improve the time to solution. About 100 development, validation, and science runs are equivalent to four additional operational runs for the entire year. The HEC requirement will grow from about 40 million hours in 2009 to 225 million hours in 2013. The caveat with these estimates is that the I/O scaling (and so the I/O burden) and the scaling of 4DVar itself are not yet known, so these estimates may be very optimistic.

The temporary storage requirement will grow from about 0.5 TB per run to 2 TB per run in 2013. Online access to products for field campaigns will grow from 10 TB in 2009 to 40 TB in 2013. In addition to archival at the HEC center, products are transferred from the HEC center to the GES DISC for dissemination to instrument teams and for distribution to the community. This transfer is about 4 GB per day in 2008, and is expected to grow to 64 GB per day in 2013. The archive storage requirement will grow from an estimated 3 PB per year to 12 PB per year in 2013.

The GEOS-5 system now includes ozone, aerosols, and some tracers for air quality. The system is evolving to include other reactive gases, mainly measured by Aura, and for the carbon cycle, mainly measured by AIRS (CO, CO_2), MOPITT (CO), and OCO (CO_2). The assimilation systems are beginning to be used to simulate observations for ACE and other Decadal Survey missions (ASCENDS, GEO-CAPE). The increase in complexity drives the need for more memory, with an increase of about 50% for full chemistry.

In addition to the meteorological and air quality analyses and forecasts, we expect to undertake another retrospective analysis by 2013. This will be more comprehensive than MERRA, including aerosols and trace gases. Ocean analyses, both uncoupled and coupled, will also be undertaken with a couple of different systems in this timeframe, focused on analysis of altimetry from Jason-1, OSTM/Jason-2, gravity information from GRACE, ocean color and temperature from MODIS, and VIIRS, salinity from Aquarius, and sea-ice from ICESat, *inter alia*. Level 4 land surface products from SMAP will also be in development, with production expected in 2013. The requirements for all these are summarized in Table 1.

The Short-Term Climate Prediction Scenario

The World Modeling Summit for Climate Prediction at ECMWF in May 2008 (see http://wcrp.ipsl.jussieu.fr/Workshops/ModellingSummit/index.html) identified the importance of testing climate models in sub-seasonal and multi-seasonal prediction mode. Of course, the predictions are useful in their own right and have a very direct societal benefit.

In the 2013 timeframe, we envision that NASA will continue to contribute to the nation's multi-model seasonal predictions through forecasts and predictability experiments. The forecasts will also contribute to NOAA's multi-model Climate Testbed activities. NASA's unique emphasis is on the optimal use of satellite observations to enhance prediction skill. Ocean and land data assimilation are important to initialize the slow components of the climate memory. The initialization for the GEOS-5 Coupled Atmosphere-Ocean Model (AOGCM) will be through a weakly coupled assimilation that brings the ocean (and land) and atmosphere into a balanced initial state that reduces initialization shocks. Most of the computational burden for seasonal predictions comes from the need to undertake historical hindcasts to provide the statistical basis for calibrating the forecasts.

The current focus on 6- to 12- month prediction skill will continue through 2013. The GEOS-5 AOGCM resolution is 1° for the atmosphere and 1/2° (with equatorial refinement) for the ocean. The real-time forecasts comprise at least 10 members in the ensemble (more if additional computational resources are available) of 12-month forecasts each month. A 30-year series of historical hindcasts will be undertaken each year for 2 years. System upgrades are only expected on 5- to 7-year cycles because of the hindcast burden. For the HEC requirements presented here, we assume a frozen system. The exception is for shorter, higher-resolution coupled predictions that will be undertaken for the subseasonal timescale. These predictions will extend the weather predictions undertaken daily with the GEOS-5 AGCM, but will be undertaken only a few times a month with a few ensembles. These subseasonal forecasts will be conducted at the NWP resolution and so will vary from 1/4° in 2009 to 1/8° by 2013. The ocean resolution will remain the same as for the seasonal forecasts. In addition to the conduct of forecasts with the frozen system, developments of the next system upgrade and predictability experiments have to be undertaken. Predictability experiments are conducted once the hindcast burden has been completed. The next system development can be represented as many coupled AOGCM integrations of about 50-year duration each year.

To be useful for climate applications, single-image model performance must be roughly 1,000 simulated days per wall-clock day. For the GEOS-5 configuration above, each single-image, 12-month model integration requires 128 cores. The subseasonal forecasts, at 1/4° resolution, require 1,024 cores for a single image in 2009. The coupled initialization, the subseasonal predictions and the experimentation for the next system upgrade will require 4,096 cores in 2013. The memory requirement is defined by the ocean assimilation, which could benefit from 4 GB per core. The HEC resource requirements will be 20 million hours in 2009 and are not expected to grow during the period of this report. It is estimated that 1 PB of data will be archived per year in the 2013 timeframe.

Since global predictions are used to force offline regional models, global predictions need to be stored (and retrieved) at high spatial and temporal resolution. The use of multi-model ensembles for consensus forecasts as well as for scientific analysis requires the sharing of the large volume of ensembles of historical data with other groups undertaking seasonal forecasts. Each model is being run over multiple realizations at several physically distinct locations, and it is not realistic to transfer the entire data volume from all groups to any single location. Thus, it is necessary to enable data management and distributed access to and analysis of a virtual multi-model archive.

The Climate Change Projection Scenario

With the release of the IPCC's AR4, the climate community is now focusing attention towards the next assessment. A new set of coordinated climate model experiments, to be known as Phase Five of the Coupled Model Intercomparison Project (CMIP5), is being identified. These experiments are expected to provide most of the climate modeling information that will become the basis for the IPCC's AR5, now scheduled to be published in early 2013. To contribute to the AR5, the simulations must be conducted in the 2009–2010 timeframe. In addition to the "standard" long-term projections, the suite of proposed experiments includes (a) simulations of future climate with relatively high-resolution models and/or with fully interactive complete atmospheric chemistry and (b) experiments focused on the near-term (the first half of the 21st century), exploring the degree to which future climate states depend on the initial climate state.

Longer-Term Projections

The CMIP5 experiments include a list of mandatory experiments for contribution to the AR5. The experiment list includes control experiments for the pre-industrial era, 20th century simulations of current climate, and future scenarios. Aside from spin-ups and the pre-industrial control, most experiments are of about 150 years' duration. Fully coupled carbon climate model experiments with prescribed CO_2 emissions (rather than concentrations) are included in the list of experiments to explore the impact of the climate-carbon cycle coupling on projected climate change.

NASA's contribution to global change projections for AR5 is primarily through the Goddard Institute for Space Studies (GISS), using ModelE. Like other groups planning contributions to the next assessment, the ModelE configuration is still being finalized. Different resolution configurations are being tested and optimized—2°, 1°, and an intermediate resolution using the finite-volume cubed-sphere. The resolution for the ocean model is comparable to that for the atmosphere. Different experiments will need different tracers, about 50 for the combined aerosol/chemistry codes, and just a few for the carbon cycle experiments. Most runs will simulate a few hundred years, but there will also be some very long runs (> 1,000 years) for spin up and for simulations of the last millennium.

The coupled model configuration with a job mix equivalent to about 100 concurrent images (ensembles, parameter sensitivity sweeps, etc) has an aggregate computational resource requirement of about 9 million hours. The throughput requirement is at least 4 simulation years per wall-clock day using up to 64 cores. The memory requirement is about 20 GB in total. The storage requirement is about 5 TB per experiment, likely accumulating up to 5 PB per year.

Near-Term Projections–Decadal Prediction

Robust projections of regional impacts of climate change require a more aggressive approach to resolution in climate models. Near-term experiments, exploring the multi-decadal prediction problem, have been proposed by a WGCM/WGSIP/CLIVAR/WCRP sub-group. The experiments focus on whether we can more accurately predict the actual trajectory of future climate (including both forced and unforced change) if we initialize the models with the observed ocean and land surface states. Decadal forecasts will be conducted once per year. Experiments conducted at high resolution (1/2° for the atmosphere, 1/4° for the ocean) will help in the preparation of the next upgrade for the seasonal forecast system. The multi-decadal predictions will be conducted with aerosols and ozone chemistry. All forcings will be included as observed values for past dates, with prescribed concentrations of well-mixed greenhouse gases. For future dates, a single IPCC scenario will be used.

The series of experiments encompasses both 10-year and 30-year predictions from initialized states. In 2009, the 10-year experiments for five of the specified years (1985, 1990, 1995, 2000, 2005) will be conducted at 1° resolution with 10 ensemble members. The 1/2° experiments will be conducted for the recent period, (2001–2005) with 10 ensemble members. Three sets of 30-year experiments will be conducted at 1° resolution with three ensemble members. The aggregate requirement is 18.4 million hours and will generate about 750 TB of archive data. In subsequent years, resources will be required to continue these sets of experiments, completing the 10 case studies for the 10-year predictions, bringing the 10-year 1/2° system to near-real-time, and adding additional ensemble members for the 30-year predictions. The resource requirement will not grow much through 2013 as the AR5 experiments will be completed, but the system will continue with experimentation, near-real-time predictions, and the development of the next system upgrade. Most of these experiments can be accomplished as multiple runs of a single image running on 128 to 256 cores. The computational burden is primarily associated with the number of case studies and the use of ensembles.

As for seasonal forecasts, the climate change prediction application requires the sharing of large volumes of data with other groups undertaking these calculations and also with the science community, who will undertake analyses and comparisons that will form the basis for the AR5.

The Chemistry-Climate Interaction Scenario

Interactions between climate and atmospheric oxidants (e.g., hydroxyl [OH] and ozone) and aerosols provide important coupling mechanisms in the Earth system. The concentrations of tropospheric ozone and the emission of chemical ozone precursors (e.g., carbon monoxide, methane [CH_4], non-methane hydrocarbons, and nitrogen oxides) increase as a result of increased use of fossil fuel, more frequent biomass burning, and more intense agricultural practices. These perturbations contribute to the radiative warming of our planet. Climate change itself impacts chemistry-related processes that feed back into climate variability and change. In addition, climate change also affects the temperature and circulation of the stratosphere, and thus the recovery of stratospheric ozone.

NASA's modeling efforts have been one of the key contributors to the WMO/UNEP Scientific Assessments of Ozone Depletion. The system used for the 2006 assessment has matured significantly over the last year with the coupling of GEOS-5 with the Global Modeling Initiative (GMI) chemistry modules, resulting in the GEOS Chemistry Climate Model (CCM). This system will now be used to contribute to the next IPCC assessment as well as to WMO/UNEP Ozone Assessments.

NASA will contribute to the proposed suite of near-term experiments for the CMIP5, simulations to 2030–2040 with fully interactive complete atmospheric chemistry. These simulations will proceed with GEOS-5 coupled with the stratosphere-only chemistry (Version 2 of the GEOS CCM) and the GMI combined troposphere-stratosphere chemistry (Version 3 of the GEOS CCM). The scientific focus includes the interactions of ozone in the climate system, with a growing emphasis in 2011–2012 on radiative interactions with the ocean (in the coupled system). More vertical resolution will be included to better resolve the troposphere and tropopause. The GMI combined chemistry package is computationally demanding (currently the computational cost of the full chemistry is 10 times that of the AGCM), so much of the science integrations will be conducted at 2°. However, by 2013, experiments should be conducted at 1/2°, assuming that access to about 7,000 cores is regularly available. A single model image requires only about 256 cores, but the full computational requirement is for many experiments using ensembles. Access to platform configurations with larger memory per core (at least 2 GB) will facilitate efficient use of resources. The combined resource requirement in 2013 is about 41.4 million hours, generating about 2 PB of archive data.

Investigations are also planned with the GMI chemistry run offline in Chemistry Transport Model (CTM) mode, where the meteorology and associated fluxes are prescribed from various NWP analyses and forecasts (from GEOS-5, ECMWF, and others). These simulations will evolve from the current 2° resolution to 1° in 2010 and 1/2° by 2013. Approximately 10 years of simulation are conducted each year, focused on using Aura data for scientific analysis and for studies of air pollution and the impacts of climate change and emissions on atmospheric composition. Archive requirements will grow from about 10 TB today to about 160 TB in 2013. A large (80 TB) online disk cache or fast access into the long-term archive is needed to support these simulations. The compute resource requirement is about 28 million hours.

Related to these CTM runs is an atmospheric chemistry replay using known emissions (often climatology) and the MERRA data stream or that from its successor. Two different runs will be conducted, one with reactive chemistry, and one with carbon, nitrogen, and aerosols when the AOGCM is coupled to ocean and land biogeochemistry modules. The combined resource requirement in 2013 for an Integrated Earth System Analysis is about 147 million hours using up to about 4,096 cores.

The Next-Generation Non-Hydrostatic Global Model Scenario

Atmospheric models, especially for NWP but also for climate simulation and prediction, are undergoing dramatic increases in resolution with the need to resolve scales associated with the atmospheric branch of the hydrological cycle. As atmospheric models approach 10 km grid size, they begin to resolve most mesoscale phenomena, the most interesting of which consist of organized convective cloud systems; but at these resolutions they are still too coarse to resolve individual convective clouds. Also at this resolution, the models begin to represent motions that cannot be regarded as hydrostatic, and so they must be recoded using non-hydrostatic dynamics. From there to a resolution of roughly 1 km, another 10-fold increase, the effects of convective clouds will be partly represented—albeit crudely and in some ways incorrectly—by the dynamics, and partly parameterized. Shallow convective clouds, in particular, will continue to require parameterization. This range of resolutions might be termed *cloud-system-resolving* since it truly resolves the mesoscale organization of convective cloud systems. We hope, however, that they also provide a useful representation of some effects of the convective elements themselves, so that we may begin to remove the deep convective parameterization and might also be referred to as *cloud-permitting*. Only at resolutions below a kilometer will we be able to claim a *cloud-resolving* capability.

The next generation of the GEOS model, GEOS-6, will be capable of non-hydrostatic dynamics and physics appropriate to high resolution. The implementation uses the finite-volume dynamical core on a cubed-sphere configuration. The development timeline is for testing at 1/2º to 1/8º resolutions, comparing hydrostatic with non-hydrostatic implementations in 2010, to undertaking NWP simulations in experimental production mode in 2013. This development path involves a close collaboration with GFDL, and the computer resource requirement supports two 10-year AMIP simulations at both centers at 14-km (C720) resolution in 2010. These simulations are only feasible if partitions of 4,000 to 8,000 cores are accessible. We expect to see qualitative changes in the model's simulation capability at resolutions of 7 km (C1440) and propose to undertake simulations equivalent to three 1-year simulations at that resolution. Such a simulation will accomplish 55 days per wall-clock day on 32,000 cores. Effective testing and tuning can be conducted with this level of throughput. The resource requirement for 2010 is about 36 million hours. In this test phase, one can anticipate archiving only about 1% of the output data, so that the archive requirement would be about 144 TB.

An exciting opportunity arises with the Year of Tropical Convection (YOTC) project in 2011. We propose to contribute a 7 km (C1440) simulation for the entire year, requiring 4 million hours and access to 16,000 to 32,000 cores.

By 2013, we expect to be conducting NWP runs at 3.5-km (C2880) resolution, tied directly to the 4DVar analysis, which should be at C1440 resolution at that time. The goal is to conduct 4-day experimental forecasts with aerosols and trace gases during the hurricane season, requiring 30 million hours and 20,000 to 40,000 cores to achieve the required throughput of one forecast per day. Development and tuning of physics will continue in AMIP work at C1440 resolution. The aggregate resource requirement is about 50 million hours, generating about 2 PB of archive data.

Summary of Requirements

In addition to the above sample scenarios, there are many other applications that need HEC resources. The requirements from the major users are summarized in Table 1, grouped according to the type of application (not according to the research group).

	2008		2009		2011		2013	
Atmospheric assimilation; products for NASA instrument teams	7M	120 64GB	40M	1,024 2TB	100M	1,024 2TB	225M	4,096 8TB
Meteorological product reprocessing for instrument teams	0.5M	120 64GB	1M	1,024 2TB	20M	1,024 2TB	100M	1,024 2TB
Chemistry/carbon data assimilation	4M	120 64GB	5M	1,400 140GB	20M	1,400 2TB	35M	1,400 2TB
Ocean-ice simulations and assimilation	7M	500 0.5TB	14M	5,000 2TB	30M	5,000 2TB	90M	5,000 2TB
Land assimilaton	0.2M	32 3GB	0.5M	64 23GB	1M	128 0.1TB	1M	128 0.1TB
OSSEs for Decadal Survey Missions			5M	128 64GB	15M	1,024 2TB	60M	4,096 8TB
Reanalysis - IESA for the satellite era	5.4M	120 64GB	5.4M	120 64GB	67M	2,000 4TB	147M	4,096 8TB
Subseasonal prediction; climate-weather	7M	256 256GB	20M	1024 1TB	20M	1,024 1TB	20M	4,096 8TB
Near-term climate change projection (AR5)			18.4M	512 1TB	20M	512 1TB	20M	512 1TB
Long-term climate change projection (AR5)	4M	16 4GB	9M	64 20GB	13M	64 20GB	20M	128 40GB
Chemistry-climate simulations	4M	96 10GB	8M	140 140GB	20M	140 140GB	70M	1,024 2TB
Next-generation Earth System Model			36M	32,000 32TB	36M	32,000 64TB	50M	40,000 80TB
Mesoscale models - GCE/WRF/LIS	2.5M	128 8GB	4.3M	128 8GB	10M	128 8GB	30.5M	256 16GB
Multiscale Modeling System	1M	364 10GB	2M	364 10GB	9M	2,000 4TB	15M	10,000 20TB
Material and chemical processing models	0.3M	256 1TB	0.3M	256 1TB	0.3M	1,024 100TB	0.3M	1,024 100TB
Total	42.9M		168.9M		381.3M		883.8M	

Table 1: Summary of the major computing applications in Earth system modeling and assimilation requiring HEC resources. The evolution is shown from the current allocations through several upgrades to 2013. For each year, the left-hand column shows the processing time in hours (M = million). The right-hand column has two entries: the upper number is the high-water mark for the number of cores to be used for a single-image computation, and the lower number is the total memory requirement for that single image. For most applications, the science generally requires several runs of different sizes and potentially different processor configurations.

Storage Requirements

For the class of models and prediction systems considered here, I/O bandwidth and long-term storage capacity can be assumed to scale fairly closely with the number of operations performed, i.e., the number of wall-clock hours required. In 2002, the storage estimates were based on the rule of thumb that models typically produce 1 byte of output for every 1,000 to 10,000 floating-point operations and that only about 1–10% of the output was saved to the long-term archive. On modern HEC systems, operation counts are not done routinely. Instead, based on our experience with today's technology, we find that a good rule of thumb is that the archived data is about 40 TB per million wall-clock hours on today's processors and that this represents about 10% of the output. The estimates then are generally fairly consistent with those made in 2002, estimating an I/O bandwidth of up to 200 TB per day, and long-term archive growth of about 20 TB per day (~ 7 PB per year) in 2009. Anticipated figures for 2013 are about 1 PB per day of short-lived data and 100 TB per day (~37 PB per year) of archive data. The larger volume of short-term data is typically used only for immediate analysis, visualization, and various post-processing data reductions. However, much of this needs to be retained for at least a month, so at least 3 PB of fast disk needs to be accessible by analysis and visualization platforms. The larger this is the better.

Almost all applications consider 5 years to be a realistic lifetime of data in long-term archive. It may be that much of this data is rarely retrieved since most of it is used in immediate analysis. However, specific projects such as reanalyses will have most of the data retrieved quite often by a broad range of users in the community. These data will have a lifetime of more than 10 years. The access requirements into these datasets can be quite varied, ranging from geographical or temporal subsets of the data, to extraction of specific variables, to more complex processing such as calculation of means or other statistics or of some more complex diagnostic.

Rapid access into data volumes of hundreds of TB is needed. The ability to undertake calculations or data reduction operations on the server rather than transferring back to local machines is essential, though transfer of selected data back to local analysis machines will remain a requirement. Estimates for today are that typically a single user wants to transfer about 1 TB of data in 1 day on a sporadic basis, perhaps once a week. Thus, 10 MB/s sustained is required, though this is rarely achieved. This requirement will grow to about 40 MB/s for a single user in 2013. The aggregate requirement will be several multiples of this. A good estimate is provided by the DOE estimate for intercomparisons and download of AR5 integrations through the Earth system grid: 2–5 Gbps for the WAN [see the *Report of the Biological and Environmental Research Network Requirements Workshop* (2007) at http://www.osti.gov/bridge/servlets/purl/924773-EMYRf8/924773.PDF]. This is consistent with the requirement estimated for the seasonal forecast application in the 2002 report. Nevertheless this seems a fairly optimistic requirement given the inadequate performance today.

Two difficult issues to be addressed are the general maintenance and management of the archive—reducing the risk of loss of large sections of data that are difficult to reproduce, and identifying data that can be deleted because it has been superseded by more recent scientific advances. Data management tools continue to be inadequate.

Programming and Analysis Environment

Productivity Tools

The 2002 Workshop report recognized the importance of various tools to increase the scientific productivity and increase efficiencies from the HEC resources. The tools included the Earth System Modeling Framework (ESMF), which is now reasonably well integrated into most of the systems that dominate the HEC requirements. The report also identified the value of what could be called workflow tools:

1. A graphical user interface for model configuration and job submission and monitoring
2. A database system for tracking and annotating model experiments
3. A database that will serve as a common repository for model components
4. An interface to data services (e.g., remote data access, diagnostic packages)

The report states: "In order to be most effective, the services above must be integrated into a system that is easy to use, deploy, customize, and extend. There are requirements that the services provided must be high-performance, robust, and portable, and that there is a long-term plan for their maintenance."

Although there has been some progress towards workflow tools for some of the models used by NASA's Modeling, Analysis, and Prediction (MAP) Program, these tools are not yet mature. The absence of these tools has hampered sharing of capabilities between groups. However, it is expected that these tools will mature over the next year and become more widely used. Key issues are portability, extensibility, and maintenance. To be successful, the workflow tools must be extensible and maintainable by the developers; otherwise their utility will be short-lived.

Thus, these requirements have not changed: (a) continue investment in and maintenance of ESMF, including attention to some of the factors (I/O and adoption of new technologies; see below) limiting progress in performance, and (b) further the development of workflow tools aimed to increase productivity. Performance efficiencies could also be obtained by wider use of performance measurement tools to identify bottlenecks. The TotalView debugger is widely used to support model development.

Productivity Associated with Platform and Environment Configuration

Memory

Although 1 GB per core is adequate for many applications, several applications, particularly assimilation and chemistry-climate simulations, trade off between memory and processors in a shared-memory environment, i.e., cores are left idle while the memory is used. Often the problem size is bounded by the available memory. Approximately 90% of the work anticipated for 2013 would benefit from 2 GB per core, and a small fraction could use more than that. It is estimated that about 30% of the work *requires* 2 GB per core (i.e., there are no out-of-core workarounds). Of course if asynchronous I/O becomes more widely used, the memory requirements will increase. One of the key issues related to this trade-off between memory and cores is that the model for compute resource utilization needs to change so that idle processors do not count against an allocation and computing center performance metrics do not entirely focus on percentage utilization of the aggregate resource. To optimize the platform configuration in terms of cost-benefit, a hybrid configuration should be considered where some fraction of the nodes has large memory per core.

Queues and Schedulers

Other issues that arose at the workshop are that queue structures and scheduling policies that have to be designed for a heterogeneous job mix from multiple applications and disciplines impact the productivity of many groups. Some groups would like extremely long run times, while others are more amenable to check-pointing calculations and having a single integration broken into many shorter integrations. Particularly during development phases, rapid turn-around is critical. The difficulty of getting regular (daily) access to large numbers of cores is a perennial problem. The only way around this problem seems to be some element of partitions dedicated to particular applications, with queue structures and policies applicable to different large applications.

Analysis and Visualization

With the increases in output data volumes over the last 5 years, it has become increasingly difficult to transfer data to local platforms for analysis. This problem will only be exacerbated in the years to come. Analysis requires parallel visualization tools and large memory; so much of the analysis must be performed at the HEC center rather than on the scientist's local computing environment, although some of that will of course continue. Virtual Network Computing (VNC) has been found to be very powerful in facilitating remote visualization. Visualization tools used are generally GrADS, IDL or MATLAB. High-resolution simulations are best visualized during the integration, directly from memory rather than local disk. Thus, there is a growing need for the HEC center to include a significant analysis platform with sophisticated, parallel visualization packages. The computing center needs to have visualization experts to aid the scientists.

Once the simulation has been completed, the visualization has to be done from disk rather than memory. This will be the case for most applications. A reasonable estimate of disk capacity would be to maintain the data outputted and archived from a run online for about a month and a smaller fraction (10% is used here) for 6 months (time to publication), i.e., by 2013 there should be about 5 PB of fast-access disk available for analysis and visualization.

The network needs for visualization tools are somewhat different from those for bulk data transfer. The DOE report mentioned above provides an estimate for Earth science models: "Current large displays are 6 megapixels in size and must be updated at 30 frames per second, requiring bandwidth greater than 4 Gbps for a full color image." With several users requiring interactive analysis and visualization, the report places the network requirement in 2012 to be 5–20 Gbps LAN and 2–5 Gbps WAN. There is some optimism that this problem will be solved because of the high network bandwidth need of emerging commercial applications.

Evolution of Modeling Activities

Today's Technology

Since the early 1990s, we have been using a generally adopted computer architecture, with its associated parallel programming model. The dominant architecture has consisted of clusters of homogeneous nodes containing sockets for one or more cache-based, commodity processors. Recently, the trend has been to have each socket populated with a multi-core processor. Various levels of memory sharing have evolved, but the dominant and most economical solution has been for cores to share memory within the node and to rely on message passing software between nodes. Parallel algorithms within nodes could still rely on shared-memory strategies (OpenMP), but in fact most codes eschewed this hybrid-programming model and relied exclusively on message passing, running MPI on each core.

Our assessment is that, for the near future, this architecture and this programming mode will remain the mainstay for most science codes, particularly in the Earth sciences. The MPI-everywhere model should be able to meet all requirements through 2011 and can be expected to remain in productive use for many codes well after that. Beyond 2011, however, the very high-end codes may require different strategies.

Accelerator Technologies

For the last 20 years, we have been getting tremendous increases in computing power from increasing processor speed, as well as by increasing the number of processors working on a single problem. It is now generally accepted that processor speed will not continue to increase as it has in recent decades and that, for some time to come, increases in computing power will have to come almost exclusively from increasing processor numbers. It is also clear that processor numbers will increase primarily by increasing the number of processing cores within a single-processor chip.

It may be that the multi-core processors of the future will evolve as much more dense versions of today's homogeneous quad-core offerings, with more and more cores per chip. In that case, we will be able to retain the MPI-everywhere programming model, and existing science codes should require only minimal changes. There is the possibility, however, that the coming computer architectures will be much more heterogeneous, consisting of combinations of very different types of cores or computing elements.

A clear emerging trend in computing architecture is the development of accelerator units (IBM Cell, GPGPU, and FPGA) that contain tens to hundreds of floating-point units (FPUs) working in parallel on a single chip. Several successful products exist today that exploit this trend. These products demonstrate both the engineering feasibility and the application utility of having numerous FPUs able to run in parallel on a single chip.

Major providers of computer graphics hardware such as NVIDIA, ATI (which is now a division of Advanced Micro Devices), and Sony/IBM all currently mass-produce low-cost, modest-power-consuming graphics hardware that today have theoretical (single-precision) peak floating-point performance approaching 10^{12} floating-point operations per second and internal memory bandwidth greater than 10^{11} bytes per second. Both of these numbers are at least one order of magnitude greater than the equivalent per-processor theoretical peak performance available on a conventional high-performance cluster computer.

To date, many of the most high-profile application areas of such parallel accelerator cards have been in computer graphics and video processing, where factors of tens to hundreds in actual performance boosts have been demonstrated through a combination of aggressive hardware and software innovation. However, with some software and algorithmic work, the same hardware technology that is used to boost computer graphics performance can potentially be applied to more general large-scale, parallel simulation problems. The problem is that these technologies are complex and require explicit management of data and computation. The supported programming models/languages are technology-specific leading to codes that are non-portable.

Technical Challenges

As discussed above, the challenges to applications programmers will differ depending on the architecture that evolves. If the accelerator technologies become dominant, we can expect to have to go through a major transition in programming model, akin to the one we experienced in going from vector to parallel architectures 15 years ago. The very notable improvement in software development practices that we have seen in the Earth sciences in the last decade, including the development of frameworks such as ESMF, may make this transition less painful. Nevertheless, we can expect a period when application codes that require cutting-edge performance will have to deal directly with architectural changes, until middleware development catches up.

Even if the current parallel programming does not have to be changed, there will be other challenges.

Memory

We can expect the total memory available in a system to continue to increase dramatically, but the rapid increase in the number of processing elements made possible by multi-core processors means that it may not be economical to increase, or perhaps even to maintain, the amount of memory per available core. Current systems have 1 or 2 GB per core. More can be added, but at a cost (in terms of both dollars and power consumption) that is not usually deemed economical. This appears adequate for most codes, and those requiring more memory may need to "waste" processing cycles by idling some of the cores in a processor, or even whole processors in a node. In the end, the strategies needed to deal with these issues may not be primarily in the hands of the programmer, but in the way procurements are done. Solutions may require buying systems with more memory and empty sockets, or splitting systems so that subsets can be tailored to the different requirements of the job mix. These are notoriously difficult decisions to make, and to support them, more effort will have to be devoted to load analysis and prediction. This, in turn, will require the development of much more sophisticated benchmarks than are currently available in our field, including "models of models" and better models of load, through a collaboration between computing center managers, computational scientists, and applications experts.

Memory Bandwidth

Many of our largest computational codes, particularly the Earth system models and data assimilation systems require large-memory bandwidth (i.e., the number of memory accesses required to perform each operation is high). This is the primary reason for the very low fractions of peak performance that these codes obtain. The 2002 report noted this inefficiency and identified it as an important challenge, but the situation has not changed significantly in the intervening years. The requirement for high memory bandwidth and the inability of modern processors to provide it seem to be structural problems that are largely beyond programming strategies. It is not clear how the multi-core trend will affect this situation, but it is likely to make the bottleneck between processors and memory even worse.

I/O

It seems likely that realizable I/O bandwidth to disk will not scale with the number of cores. For codes that require such scaling, I/O—not processing power or memory—will eventually be the limiting factor in achieving the required performance. We expect that much can be done in current codes to improve their I/O performance, including making more use of parallel and asynchronous I/O. Parallel I/O, however, can be cumbersome to use, and its performance is sensitive to how the system is configured; ideal performance, therefore, can be beyond the applications programmer. Asynchronous I/O cannot be used for critical sections, such as the reading of restarts or the recording of the state at the end of a run. It can, however, be used effectively for writing diagnostic histories in the course of a run, if one has the memory to hold the fields to be output while the calculation proceeds. A similar trade-off between memory and I/O occurs for large, "out-of-core" calculations that resort to doing I/O to scratch files with temporary results that do not fit in memory. We anticipate that, for one reason or another, most applications will be seriously impacted by the I/O bottleneck, and that I/O performance will be a serious hurdle in meeting the ambitious performance goals set forth in this document.

Arithmetic and Precision

In addition to the major issues just discussed, we identified the possibility that changes in arithmetic standards and precision could require modifications in existing codes and practices. Most current scientific algorithms rely on 64-bit floating-point precision, although many calculations could be done at 32-bit precision with some redesign. Some consideration should be given to this, since most proposals for accelerators will heavily favor 32-bit arithmetic. In addition, the pressures on memory, memory bandwidth, and I/O we just discussed, which apply regardless of the architecture, can be partly relieved by storing and moving half as many bits. In the same vein, scientific programmers have come to rely on the nearly universal adoption of IEEE standards in arithmetic by today's processors. Taking advantage of most accelerator technologies will require relaxing these standards.

Other Implications of Accelerators

Finally, we should note that if accelerators become the norm, issues of memory (they will have relatively little), memory bandwidth and utilizations (the programmer may have more responsibility for memory access), and I/O (probably less bandwidth per core) will all change dramatically.

Programming Frameworks (ESMF, MAPL)

Object-oriented software frameworks are becoming an integral part of today's complex Earth system models. ESMF is the most important of these, and it is becoming widely used in coding the superstructure that couples the major components in Earth science codes. ESMF is used extensively by NASA codes, where it is supplemented by a utility layer (MAPL) that facilitates the use of ESMF and defines a coupling strategy that makes it easy to connect ESMF components that abide by its rules.

Frameworks provide general and standard ways of defining and manipulating computational objects, such as grids, fields that are defined on grids, and even entire model or coupling components. Because frameworks work primarily at the superstructure level—the most serial part of the codes—and because of their generality, using them necessarily involves a computational overhead. In current models, which typically run on a few hundred processors, this serial overhead is small, even when the framework is used extensively. It is important that the scalability of the framework be verified to tens of thousands of processors.

In addition to its use in the models' superstructure, ESMF is now beginning to be used for infrastructure tasks that involve communication between processors and processes. For this purpose, the ESMF team has developed an abstract machine model that allows the application programmer to codify these tasks in ESMF calls that the framework can support in specific computer architectures. If we can make use of ESMF's machine abstraction, the process of adapting science codes to new architectures should be greatly simplified, since problems will have to be dealt with only once, at the framework level. It will then be important that we have MAPL/ESMF support for accelerator-based processors and other architectures, even—perhaps particularly—when these are at the evaluation stage.

ESMF also provides a platform for the development of advanced I/O capabilities required by the next-generation ESMs. Such endeavor goes beyond NASA and calls for a close partnership with the other agencies integrating the ESMF community.

Language

With the new technologies being focused on commercial applications rather than science, Fortran compilers are slow to appear on these machines. The question always arises as to whether it is time to abandon Fortran as the language of choice for our models. One reason given for continuing to use Fortran is that so many of our codes are written in Fortran, and it is difficult to abandon these legacy codes. There have been recent examples of modern models being coded in C, IDL, and other languages. However, these codes have not gained any traction in the community and have generally been abandoned. Some groups test small code segments in MATLAB. However, the operational implementation within comprehensive models has been in Fortran. We expect Fortran to remain our mainstay for the foreseeable future.

Performance Modeling and Benchmarking

As the scope and sophistication of Earth system models increase, the complexity of the software implementing such systems is expected to increase at an even faster rate. As new technologies become available, our ability to efficiently evaluate and benchmark these technologies will become severely compromised if one is required to port a complex, multi-component system to the new architecture/testing center. Historically, similar difficulties have also stood in the way of partnerships with vendors and computational scientists who are not themselves experts in Earth sciences.

One solution to this problem is the systematic development of "computational analogs" of the real Earth system models, which can be easily deployed and ported to new architectures. Development of such analogs requires careful modeling of the computational profile of the real model (including scaling, memory footprint, throughput, and I/O characteristics) as well as routine calibration against a current production system. Such an effort would also provide routine benchmark, profiling, and performance analysis of the production system, helping to identify defects and bottlenecks. The implementation of this capability will require a focused activity, with a dedicated computational scientist/performance engineer working in close collaboration with the ESM developers.

Standards for Collaborative Development and Data Sharing

Data management was one of the five major technology gaps areas highlighted in the 2002 report and still remains a concern today. The sharing of data and the adoption of domain-specific metadata convention are issues that involve the scientific community at large, and ultimately the users of the particular datasets. As different data have a different audience (and some data have more than one audience), adoption of a single, uniform standard across NASA as in the past may not be in the best interest of our users. Several community-wide initiatives to develop metadata standards specific to Earth system models have emerged over the last 5–8 years (notably CF, NMM, and XML, which are now being merged within the Earth System Curator). Participation and close collaboration with these efforts is highly recommended. The application of ESMs and datasets in other disciplines should also be pursued, in particular through the dissemination of the same Earth system datasets in a form compatible with Geographical Information Systems (GIS). The current trend to make data directly accessible through the Internet using OPeNDAP, Web Map Service (WMS), Web Coverage Service (WCS), and similar technologies should be encouraged.

Model metadata standards development and management is a fundamentally important technology to assure traceability of a) model inputs, b) detailed model configurations, and c) model output datasets. The Earth System Curator project (see http://www.earthsystemcurator.org) is prototyping a software environment for "*assembling, running, and archiving information about climate models. The idea is to make it easier for scientists to perform modeling experiments, and to coordinate with each other on efforts such as Model Intercomparison Projects (MIPs) and Intergovernmental Panel on Climate Change (IPCC) assessments.*" NASA's participation in standards efforts such as the Curator project and its successors is important to ensure that such standards reflect NASA's specific requirements.

Summary and Recommendations

Our progress in Earth system modeling, assimilation, and prediction depends on high-end computing and is proportional to the resources available and their ease of use. Significant advances in many of the key research areas will depend on advances in computing technology. Some of these advances, such as increases in the number of available processors, can be anticipated as a surety, others, such as software developments to address bottlenecks or use new technologies, need investment from the Earth science community to ensure that they are realized.

It is clear that as we progress over the next 5 years, we will be going through significant transitions in both science and computer technology. There will be fundamental changes in the nature of the problems we address and the tools we use. At 3- to 5-km AGCM resolution, we will be able to undertake global cloud-permitting simulations and forecasts, and there will be qualitative changes in the nature of the solutions as we abandon parameterizations. Our atmospheric assimilation will include both meteorology and chemistry, and we will be ready for new data from new platforms and to make better use of the high-resolution data that we have already. Similar advances will be made for the ocean, ice and land surface. We will be able to quantitatively examine the connections between climate variability and change and weather extremes. The satellite data will be available, and our models will have the complexity needed to build a good carbon inventory. Our models will have more complete connections and feedbacks between the ocean, cryosphere, atmosphere, and land. We will have to take advantage of new technologies to achieve this transition, and we will have to address some significant bottlenecks, such as I/O, scaling, and data management. It is clear that close collaboration between the scientists and the computing centers and their computational science staff will be needed to achieve the scientific advances that we envision.

These advances are not extreme visions of what we would like to achieve, but rather realistic views of what is achievable without a large growth in science staff. The requirements present what it will take for NASA to remain at the forefront of the Earth science modeling and assimilation community and also be able to collaborate with our colleagues in internationally coordinated experiments with state-of-the-art systems. The view does include grand challenge problems, like non-hydrostatic model simulations, but ones that we have to undertake to make a major step forward in our ability to model and predict our environment as a system.

A brief summary of the outlook and issues discussed at the meeting follows.

Resource Requirements

- Without any attempt at "heroic" calculations, the Earth modeling and assimilation efforts need roughly a 20-fold increase in compute capacity by 2013, and an approximately 40-fold increase in capability. The latter is driven by the increases in resolution requiring a progression from jobs that typically use 100–200 cores to ones that need to have regular, assured access to about 4,000 cores. Specialized, grand challenge runs that will be state-of-the-science and keep NASA at the forefront of modeling and prediction will require access to 32,000 cores. It is important to note that the requirements are for a significant increase (four-fold) over the next year, and thereafter a growth more consistent with Moore's Law.

Recommendation 1: **Increase the HEC resources available to Earth system modeling and assimilation—a four-fold increase in capacity in the next year; at least a 20-fold increase in capacity and 40-fold capability over the next 5 years.**

- The memory configuration of NASA's HEC platforms needs greater attention. Although memory is not a limiting factor in our ability to use the compute resources, it can play a role in bounding the problem size. For many applications, trade-offs are made between memory and I/O and/or processor utilization. Approximately 90% of the workload anticipated for 2013 would benefit from 2 GB per core, and a small fraction could use more than that. A hybrid environment of partitions with different memory configurations would be a cost-effective approach to maximizing resources. A new paradigm in accounting for resource utilization needs to recognize that some cores will be underutilized in jobs that require large memory.

Recommendation 2: **Attention should be given to machine balance, specifically the trade-off between memory and compute cycles.**

- Analysis/visualization platforms are needed in HEC centers for efficiency (to limit remote data transfers). The anticipated data volumes preclude transfer of anything but a small fraction of output to local machines for analysis and visualization. HEC staff with visualization experience can help introduce new visualization techniques to users and help promote efficient tools for remote use.

Recommendation 3: **Analysis and visualization environments should be a standard part of the HEC resources with staff available to help users.**

- Data archival requirements are estimated to grow to about 37 PB per year, and the need for online fast disk for analysis and visualization is expected to grow to about 5 PB. Storage was not anticipated to be a major issue over time, however HEC centers must continue to invest in storage as part of a balanced compute environment.
- Network bandwidth remains a problem, with unreliable access to the required bandwidth between HEC and local platforms. A reliable bandwidth of least 30 MB/s sustained will be needed by 2013.

Programming and New Technologies

- We expect Fortran to remain our mainstay for core numerical computations, despite the fact that Fortran compilers are slow to emerge for new technologies.
- Efforts are needed to improve scalability, particularly for cases of fine-grained parallelism.
- I/O is a significant bottleneck for many applications, accounting for over 50% overhead in some cases. This overhead will only increase with resolution. Thus more attention is needed regarding parallel I/O and its consequences in, e.g., sharing output data.
- Frameworks (like ESMF and the MAPL toolkit) are an important element of shielding application programmers from architecture evolution. They need to play a greater role in relieving the bottleneck issues such as I/O and in the testing and adoption of new technologies.

Recommendation 4: **To achieve our science in a cost-effective way, we need to invest in exploration of new technologies, hardware, and software, including I/O technologies:**

   a. Mechanisms, such as software frameworks, are needed to facilitate adoption of new technologies.
   b. Partnerships between computing centers and application programmers need to be established or strengthened to plan for new technologies.

c. **Some (limited) investment should be made in new technology hardware to help with testing.**
d. **Partnerships should be developed with other agencies that are more aggressive in exploring new technologies.**

Data

- Sharing model output, intercomparisons with other models and data, and use of model and assimilation products in external applications are hampered by the lack of standards (format, metadata, naming conventions, ontologies) or their use by modeling groups. This is an issue for the HEC users, not the HEC centers.
- Storage capacity is not an issue, but online disk is, as is data transfer to remote machines. Online disk is needed to retain data for analysis and visualization.
- Tools to manage the archive are still lacking, making it difficult to manage obsolete data holdings.
- Data sharing between groups, between NAS and NCCS, between NASA and the external community, and even access to one's own data are hampered by network bandwidth and IT security (firewall) issues.

Recommendation 5: **Attention is required in data management, particular regarding issues of distributed access to data and interoperability standards.**

Productivity

- Partnerships between computing centers and scientists to define the compute environment (queues, schedulers), to identify performance bottlenecks and to explore options for performance improvements, to explore the potential of new technologies, to develop workflow tools, and to develop advanced visualization capabilities are needed to increase productivity and the quality of the scientific output from HEC resources.

Planning

- Our ability to efficiently evaluate and benchmark new technologies will become severely compromised if one is required to port a complex, multi-component system to the new architecture/testing center. One solution to this problem is the systematic development of "computational analogs" of the real Earth system models, which can be easily deployed and ported to new architectures. Such an effort would also provide routine benchmark, profiling, and performance analysis of the production systems, helping to identify defects and bottlenecks.

Recommendation 6: **Investment is needed in a new benchmark paradigm using computational analogs. Benchmarks should be representative of realistic computations, including the I/O burden of real applications.**

Panel Membership and Acknowledgements

Bill Lapenta, NASA/MSFC, Co-Chair
Michele Rienecker, NASA/GSFC, Co-Chair
Don Anderson, NASA/HQ
Randy Bolanos, NISN
Ilene Carpenter, SGI
Jiun-Dar Chern, UMBC/GEST
Jarrett Cohen, NASA/GSFC
Tom Clune, NASA/GSFC
Arlindo da Silva, NASA/GSFC
John Evans, NASA/GSFC/GST
Jim Fischer, NASA/GSFC
Ichiro Fukumori, NASA/JPL
Ron Gelaro, NASA/GSFC
Jing Guo, NASA/GSFC
James Hack, DOE/ORNL
Chris Hill, MIT
Michael Kalb, Science Interfaces and Assessment
Gail McConaughy, NASA/GSFC
Piyush Mehrotra, NASA/ARC
Dimitris Menemenlis, NASA/JPL
Kathy Pegion, IGES/COLA
Christa Peters-Lidard, NASA/GSFC
Harper Pryor, NASA/GSFC
Bill Putman, NASA/GSFC
H.K. Ramapriyan, NASA/GSFC
Steve Reinhardt, Interactive Supercomputing
Jose Rodriguez, NASA/GSFC
Ellen Salmon, NASA/GSFC
John Schnase, NASA/GSFC
Siegfried Schubert, NASA/GSFC
Hongbo Su, IGES/CREW
Max Suarez, NASA/GSFC
Ricardo Todling, NASA/GSFC
Michael Turner, NISN
Scott Wallace, NASA/GSFC/CSC
Bill Ward, NASA/GSFC/CSC
Phil Webster, NASA/GSFC
John Wise, NASA/GSFC
Zhujong Zhang, NASA/ARC
Shujia Zhou, NASA/GSFC

Contributions to the report were also made by Rolf Reichle (NASA/GSFC), Randy Koster (NASA/GSFC), Steven Pawson (NASA/GSFC), Gavin Schmidt (NASA/GISS), and Wei-Kuo Tao (NASA/GSFC). Their contributions are gratefully acknowledged. Jim Hack's presentation to the panel on DOE/ORNL's plans were very helpful to panel discussions.

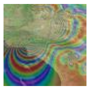

Solid Earth and Natural Hazards Panel

Science Drivers

Solid Earth science anticipates an increasing number of missions and data in the next 10 years. Space technologies will allow us to measure previously unobservable parameters and phenomena, resulting in a new understanding of complex, interconnected solid Earth processes. The next great revolution in Earth sciences will involve development of predictive models of these processes. For these models to be successful, particularly for an understanding and forecasting of hazards, high-resolution, global observations with real-time or near-real-time data streams and processing will be required. Integrating the projected huge quantities of data and information into forecast models will require that information technology resources be developed in concert with advanced sensor and detection capabilities.

The 2002 report by NASA's Solid Earth Science Working Group, *Living on a Restless Planet*, provides a 25-year research agenda for NASA's solid Earth research. The driving questions identified in the report are the following:

- What are the nature of deformation at plate boundaries and the implications for earthquake hazards?
- How is the land surface changing and producing natural hazards?
- What are the interactions among ice masses, oceans, and the solid Earth and their implications for sea-level change?
- How do magmatic systems evolve, and under what conditions do volcanoes erupt?
- What are the dynamics of the mantle and crust, and how does the Earth's surface respond?

These broad research topics are being addressed by computational modeling as well as observation, and the two are typically interconnected (Figure 1). The sub-fields of solid Earth that relate to NASA goals and must make use of computational resources include: earthquakes, volcanoes, tectonics, geodynamo, mantle dynamics, surface processes, landscape evolution, gravity, magnetic fields, cryosphere and ice modeling, ecology, hydrology, and vegetation.

Solid Earth processes often take place on scales of tens to millions of years. Even with the most advanced observational systems, the temporal sampling of such phenomena is poor. In order to fully understand these highly complex systems, simulations must be carried out concurrent with observations so that the entire system can be studied. The observational data can then be assimilated into these computational models, providing constraints and verification of the models. Because solid Earth processes occur on many different spatial and temporal scales, it is often convenient to use different models. Increasing interoperability and making use of distributed computing can enable system-level science.

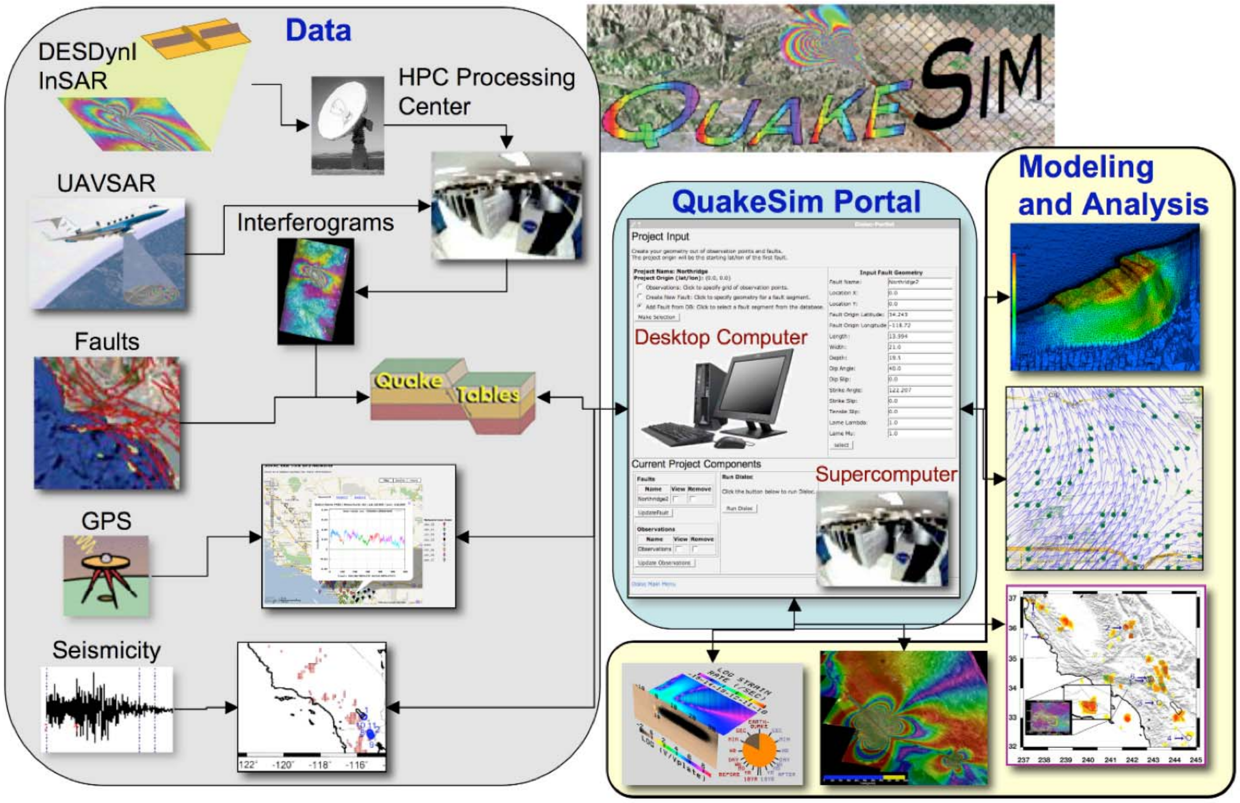

Figure 1: Interconnectivity of observational data, modeling applications, web services, and portals in solid Earth research showing the end-to-end flow.

To make the high-level drivers more concrete, we summarize below the specific scientific challenges of several prominent examples of solid Earth science activities.

Example Computational Drivers: Geodynamo and Geomagnetic Data Assimilation

The Geodynamo and Geomagnetic Data Assimilation application has the following goals:

- Understand the origin of the geomagnetic field.
- Understand the mechanisms of geomagnetic secular variation.
- Understand core-mantle interaction and the impact on global variation.
- Predict geomagnetic secular variation.

Note that all of these geodynamo/geomagnetic drivers address NASA strategic objectives and goals, including:

- Strategic Research Objectives 3A.6 "Characterize and understand Earth surface changes and variability of Earth's gravitational and magnetic fields."
- Strategic Goal 3A: "Study planet Earth from space to advance scientific understanding and meet societal needs."
- Strategic Goal 3C: "Advance scientific knowledge of the origin and history of the solar system, the potential for life elsewhere, and the hazards and resources."
- Strategic Goal 3C.1 "Learn how the Sun's family of planets and minor bodies originated and evolved."

Example Mission Driver: Deformation, Ecosystem Structure, and the Dynamics of Ice (DESDynI)

NASA and the U.S. geophysical community are preparing for DESDynI, an L-band InSAR and multi-beam LIDAR mission targeted for launch in 2014, which will produce global strain maps of Earth's deforming regions. The recently released Decadal Survey for the Earth Sciences, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*, recommends the launch of DESDynI to study surface deformation among other objectives in the 2010–2013 timeframe. The DESDynI mission will provide a dedicated, space-based source for InSAR and LIDAR imaging (Figure 2), which can be used to detect changes in polar ice sheets, ground vegetation coverage, shifts in earthquake fault zones, precursors to land slides, etc. DESDynI will generate over 600 GB of images per day, so the storage and processing alone will provide a major challenge for current NASA resources. In addition, much additional, complementary research work (e.g., modeling, data mining) will be driven by the influx of data.

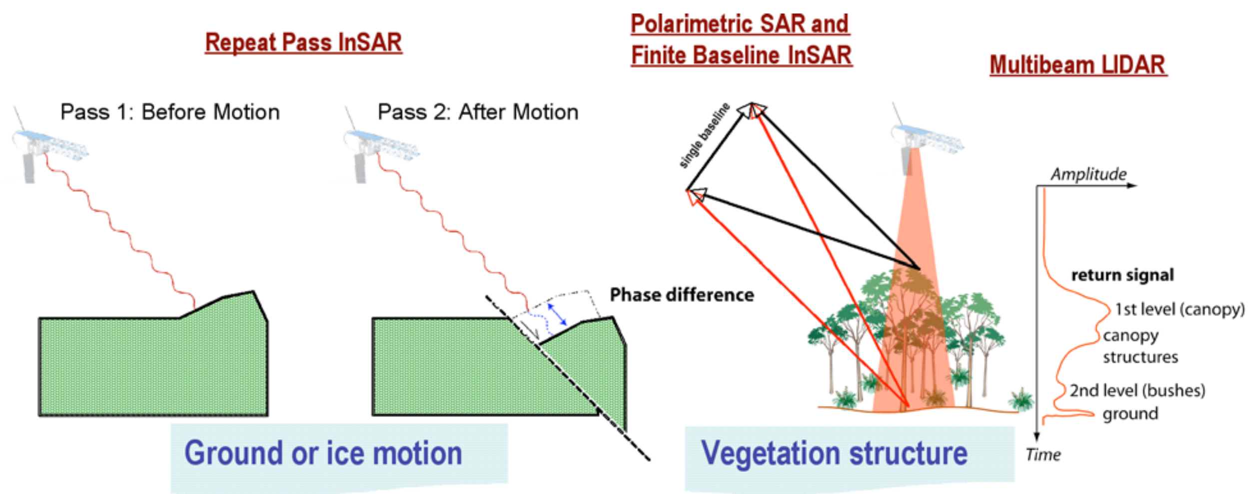

Figure 2: DESDynI will provide InSAR and LIDAR data. It is defined as an L-band InSAR and multi-beam LIDAR mission for improving our understanding of hazards, ice sheet dynamics, and ecosystems.

We summarize DESDynI's observational capabilities in Table 1. The section after the table then summarizes the associated science drivers for each of the columns.

Deformation	Ecosystems	Ice Masses	Subsurface Reservoirs
Earthquakes Probability, aftershocks, stress transfer	**Above-ground biomass** Carbon sources and sinks	**Ice sheet flow** Response of ice sheets/shelves to ocean/atmosphere	**Aquifers** Withdrawal and Recharge; Subsidence
Volcanoes Volume, depth, and migration of magma chamber	**Changes in carbon stocks** Interaction between vegetation and atmosphere	**Glaciers** Response to atmosphere	**CO_2 sequestration** Subsurface migration
Landslides Detect pre-slip	**Biodiversity** Habitat structure	**Sea ice** Interaction between ocean/atmosphere	**Oil reservoirs** Subsidence, pipe breakage

Table 1: DESDynI observational capabilities.

The DESDynI mission will Provide a Dedicated, Space-Based Source for InSAR and DESDynI Science Objectives

The science objectives for each of the major areas of DESDynI are as follows:

Deformation
- Characterize the nature of deformation at plate boundaries and the implications for earthquake hazards;
- Characterize how magmatic systems evolve to understand under what conditions volcanoes erupt; and
- Characterize landslides and detect pre-slip.

Ecosystem Structure
- Characterize global distribution of above-ground vegetation biomass;
- Quantify changes in terrestrial sources and sinks of carbon resulting from disturbance and recovery (net terrestrial carbon flux); and
- Characterize habitat structure for biodiversity assessments.

Dynamics of Ice Sheets
- Quantify the interactions of ice sheets with oceans, the atmosphere, and the solid Earth and their implications for sea level change;
- Quantify the interactions of glaciers with the atmosphere; and
- Quantify sea-ice mass balance and how it is changing.

Subsurface Reservoirs
- Characterize aquifer, hydrocarbon, and CO_2 reservoirs withdrawal and recharge.

Example Driver: UAVSAR

UAVSAR—reconfigurable, polarimetric, L-band synthetic aperture radar (SAR)—is specifically designed to acquire airborne repeat-track SAR data for differential, interferometric measurements. It serves as a testbed for DESDynI and will also augment DESDynI observations, providing improved temporal coverage, when the mission flies. Differential interferometry can provide key deformation measurements, and is important for studies of earthquakes, volcanoes, and other dynamically changing phenomena. Using precision real-time GPS and a sensor-controlled flight-management system, the system will be able to fly predefined paths with great precision. The expected performance of the flight-control system requires the flight path to be within a 10-m-diameter tube around the desired flight track.

Example Driver: GPS

NASA has already invested heavily in the development and application of technologies in the form of GPS to measure surface deformation. The 250-station, continuously operating Southern California Integrated GPS Network (SCIGN) is one example of this investment. SCIGN produces daily position time series for sites within Southern California, and the Bay Area Regional Deformation Network (BARD) does so for Northern California. EarthScope's Plate Boundary Observatory (PBO) now includes SCIGN and provides surface deformation measurements throughout the western U.S. As mentioned above, pattern recognition techniques can be used to routinely analyze the data. Time series as well as vector deformation data can be analyzed.

The growth in the number of GPS stations and the longer time series drive the need for improved analysis algorithms and compute power. Furthermore, GPS, UAVSAR, and DESDynI measurements are complementary. As such, a computational infrastructure that accesses and integrates these datasets in modeling applications will result in more robust science results.

Models and Applications

In order to access the computational infrastructure requirements for solid Earth research, we began with a sampling of applications from participants in the breakout session. Our methodology was to determine the current and projected requirements for these applications. We begin with some general remarks.

General Remarks and Requirements

Below are general observations that emerged from breakout discussions. Investments in computational infrastructure should take these factors into account.

- Modeling applications need both ensembles of modest runs (suitably run on $500,000 clusters) and fewer supercomputer runs (suitably run on a $30 million machine).
    - We will use these to guide solid Earth science requirements.
    - The amount of commodity computing power from multi-core processors will increase dramatically by 2013.
- All developers expect to continue to use MPI, and all present had existing quality MPI codes.
    - We will use all machines that support MPI and its programming model.
- Visualization is important and requires good networking.
- New instruments such as DESDynl will be very important to modeling.
    - They enable new science and modeling efforts.

    - They require data processing, mining, and assimilation.
- UAVSAR/DESDynI processing will need daily routine supercomputer runs and high bandwidth for disseminating data and data products.
- Various computational fields in solid Earth have different maturities.
    - Geodynamo is used to lots of data.
    - Earthquake and polar science are not, but large data volumes are on the horizon from UAVSAR (now) and DESDynI (soon).
- Some computational fields are real-time; others are longer-term
    - Real-Time: Weather and earthquake/volcano/landslide science.
    - Long-Term: Climate, polar science.
- There is a serious lack of computational scientists (both now and in the pipeline) to develop/understand codes/results and use computing infrastructure.
    - This makes it difficult in many cases to obtain the true level of the computational requirements.

Some Exemplary Solid Earth Applications

The first four are "modeling" class applications rather than data assimilation class applications. These applications have small input data requirements but can generate very large outputs (GB to TB) that can be compared to observation. We include DESDynl data processing as a sample of a different type of key application.

Polar Science

This area includes modeling of ice flow over Antarctica, in 2D or 3D, running for 100 years using 6-month time increments. It has a resource-constrained reality of 10% (JPL group of five researchers).

Geomagnetism, Earth Interior

Geodynamo/geomagnetic data assimilation is made up of a major five-institution group. MoSST is designed to simulate the Earth's core and planetary cores for understanding the dynamics of the core fluids and the origin of the geomagnetic and planetary magnetic fields. The MoSST_DAS is a data assimilation system based on MoSST, ensemble assimilation algorithms, and surface geomagnetic field models (CALS7K, GUFM1, and CM4), and aiming at predicting geomagnetic secular variation. The research directly addresses the NASA strategic goals 3A (Study planet Earth from space to advance scientific understanding and meet societal needs) and 3C (Advance scientific knowledge of the origin and history of the solar system, the potential for life elsewhere, and the hazards and resources).

Crustal Deformation Science

NASA applications to earthquake science focus on the interseismic, coseismic, and postseismic strain signals, because spaceborne technologies are used to detect crustal deformation. Various methods are being developed to study earthquake processes. Similar or the same tools can be used to study volcanoes, landslides, and subsidence. We focus here on earthquakes.

GeoFEST is a finite element modeling tool of elastic/viscoelastic stress and strain, including fault mechanics, stress transfer, tectonic forces, and the long-term deformation effects of earthquakes. It is one of ~3 major QuakeSim simulators. GeoFEST includes nonlinear rheology, slip on curved and wrinkled faults, and production of Coulomb Stress images. Anomalies can be detected and assessed with less approximation in the modeling. Stress transfer can be estimated for hazard assessment. These new capabilities can be more optimally presented to researchers in a portal environment.

Nonlinear parameter optimization methods, which invert InSAR images and other deformation data into parameter-driven GeoFEST models, drive improved high-performance computing needs. This application could make use of a workflow of multiple GeoFEST simulations spawned on TeraGrid or other Cloud supercomputers. GeoFEST runs can be subdivided into elastic runs and simulation time-steps that are dependent but very loosely coupled. The multiple GeoFEST runs required by the parameter optimization are fully decoupled.

Virtual California (VC) is a topologically realistic numerical simulation (boundary element code) of earthquakes occurring on the fault systems of California. It includes all the major strike-slip faults in California and has now been extended to depth-dependent boundary elements, dipping faults, as well as other new features. VC is an efficient, parallel object-oriented C++ numerical code that runs on NASA HEC's Columbia system and JPL's COSMOS computer using MPI-II protocols. VC can also compute associated surface displacements through time. From the associated simulation, InSAR interferograms can be computed for use in analyzing and interpreting signals from real interferograms obtained via radar satellites such as DESDynI. This process of comparing simulated interferograms with real interferograms will allow us to study questions relating to physics of earthquakes such as: 1) precursory failure process of major earthquakes on complex fault systems; 2) timing and statistics of major earthquakes on complex fault systems; and 3) origin of space-time correlations between major earthquakes. Another major application of VC simulations lies in earthquake forecasting. In weather forecasting, current and past observational data are routinely assimilated into numerical simulations to produce ensemble forecasts of future events in a process termed "model steering." Rundle et al. have developed a similar approach that is motivated by analyses of previous 30-year forecasts of the Working Group on California Earthquake Probabilities. By systematically comparing simulation to observed data (the variability of paleoseismic and historic data), a series of spatial probability density functions can be computed that describe the probable locations of future large earthquakes. These forecasts yield fault-based locations for the next earthquake, as well as most-probable locations for earthquakes during the next 30 years. Robust software for use in portals and services such as QuakeSim allows users to select search parameters, and to produce WGCEP-type forecasts in time spans of hours rather than the several-year time spans that represent current practice.

The Pattern Informatics (PI), Relative Intensity (RI), and RIPI methods use online seismicity catalogs to generate space-time forecast maps (Figure 3). Systematic and ongoing real-time tests of these forecasts are posted on the NASA/JPL QuakeSim website (http://quakesim.org). A real-time test of the original method was published in 2002, and a test of an updated and modified method as published in 2007. The great majority of the earthquakes have occurred on or near a colored anomaly, or "hotspot," on that map. The location of the recent magnitude-5.4 Chino Hills earthquake was successfully forecast by both maps. We propose to have continuously updated versions of a forecast map displayed on the QuakeSim portal as part of our future research.

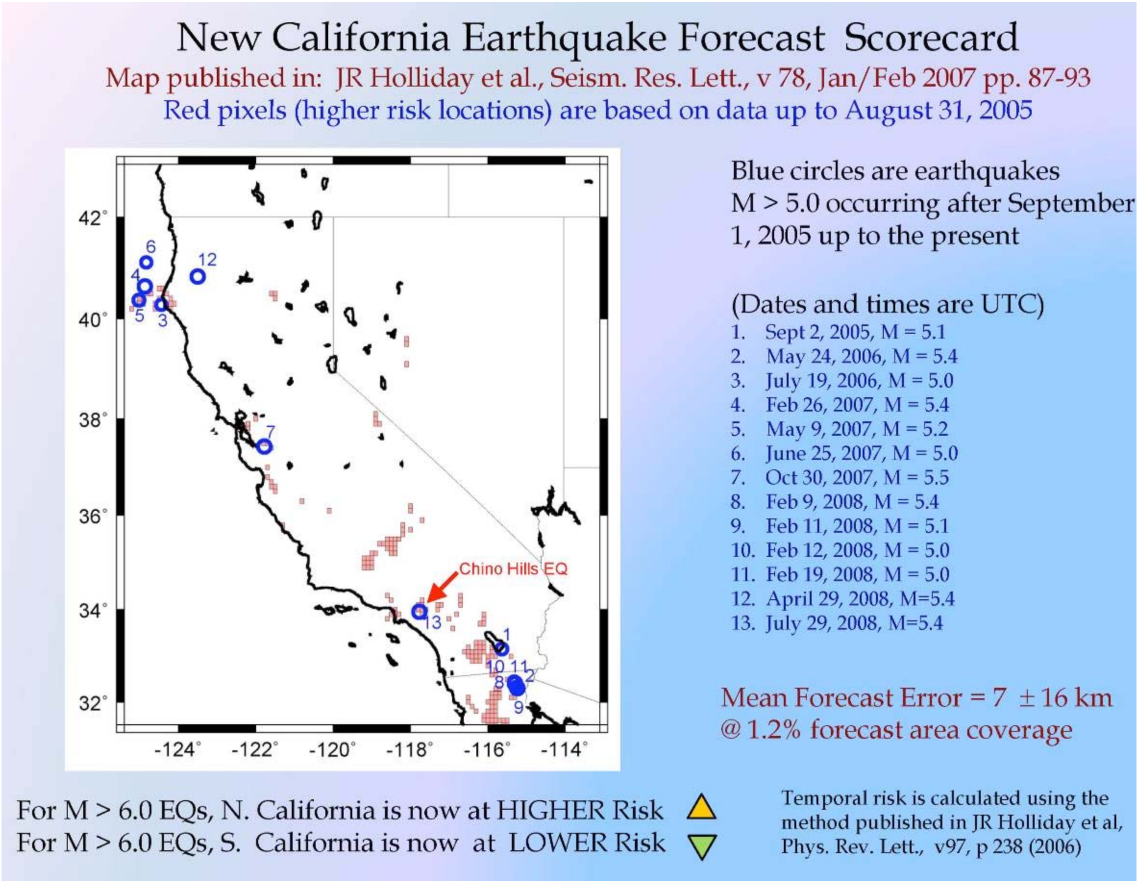

Figure 3: RIPI forecast, showing successes since the map was published in 2005. The 2002 forecast has had 22 successful hits of 25 earthquakes (http://quakesim.org/scorecard.html).

RDAHMM is another method for crustal deformation science. Signals of interest, particularly those indicating stress transfer between faults, are very subtle, are often overlain by other sorts of signals, and arise from sources as diverse as aquifer activity and atmospheric disturbances A statistical modeling approach, RDHAMM, allows one to automatically infer modes of activity within individual time series and across a network of sensors. The modeling technology allows one to be effective even in cases in which there is no model for the observed system, as well as overcome stability problems that plague standard methods. One computational challenge is that the method needs to be computationally swift enough to be applied in real-time to streaming sensor data. Current model fitting methods are iterative approaches that can take an unacceptably long time to converge. Methods such as conjugate gradient acceleration can be used to speed convergence. Runs should be made using multiple computational processors when available, and the methods should be developed so that they can be run in parallel.

Data-Fitting Techniques

Solid Earth science needs to integrate data and modeling software at the level where sensor observations are sensitive to model parameters. It is essential that new modeling components be straightforward to add by geophysicists. NASA should focus on building data-fitting core software modules such that new data or new models may be correctly combined with prior data using all the information in each and fully compatible with distributed components.

InSAR Data Processing

There is an increasing amount of InSAR data for solid Earth applications to be processed. Currently data are acquired from infrequent acquisitions. DESDynI will be producing up to 1 TB/day of SAR data. Once the raw data are processed, they can be analyzed several ways using the polarimetry and/or backscatter for ecosystems applications, speckle tracking for cryosphere objectives, or formation of interferometric pairs for crustal deformation applications. Current users of InSAR data typically process the data themselves. UAVSAR and then DESDynI will drive a paradigm for routine processing and analysis of the data and distribution to data centers. Ideally, the data will be integrated with and accessible to modeling applications, which will allow hundreds of scientists to make use of the data. The goal is to distribute the data as broadly as possible for a variety of uses. Here, we outline the steps for processing of UAVSAR data, which will be somewhat similar to DESDynI data processing. UAVSAR should serve as a testbed for automating DESDynI data processing.

Several steps are required to process UAVSAR data (Figure 4). Unlike processing of spaceborne repeat-track data, where the flight path is very smooth over time such that residual errors can be easily modeled after initial processing, airborne flight tracks are not known sufficiently well in advance. It is necessary to use the data themselves to refine the time-varying baseline between flight tracks, reprocess the data with updated flight track information, and then proceed with interferometric processing.

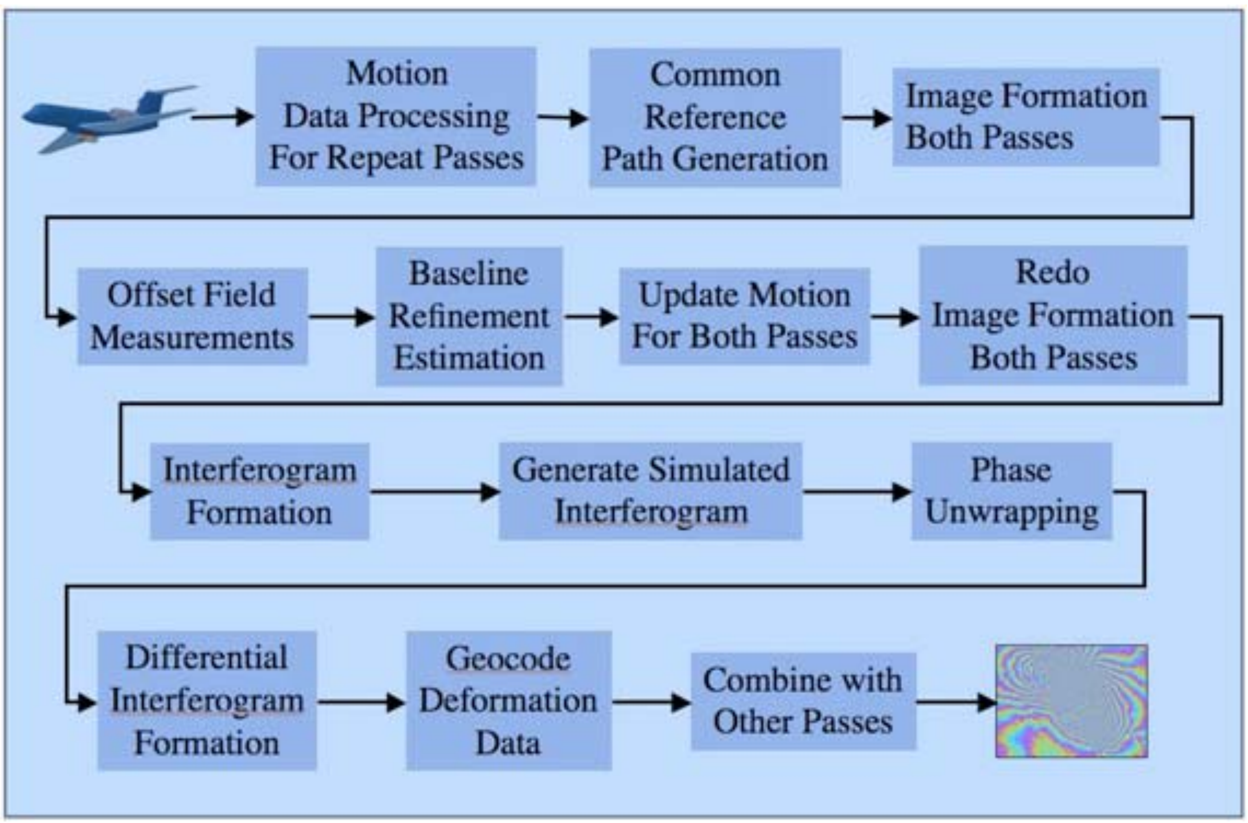

Figure 4: Processing methodology for UAVSAR repeat-track interferometry.

The ground system consists of the processing software and hardware to take the data collected from the UAVSAR and transform it into images, interferograms, and other higher-level science products. The ground processing system is based on software developed for GeoSAR and a differential interferometry processing software package developed at JPL that is distributed as ROI_PAC. The main technical challenges for the ground processor are:

- Motion processing of airborne repeat-pass data. Current metrology for obtaining the relative aircraft positions between the two flight lines is limited to about 3–10 cm accuracy. The required accuracy for repeat pass processing is about an order of magnitude better. Therefore, we must use the data themselves to solve for residual baseline errors.
- Calibration and processing of data when the antenna is actively scanning. This will require new algorithms and data analysis techniques to compensate for any systematic changes in the antenna phase center or induced phase biases that are azimuth angle dependent. Data processing techniques are being developed under UAVSAR funding.

The software is designed to be production quality, however some iteration will be required to extract the most accurate displacement signals from these data, particularly for the expected subtle signals during this short observation period. As a result, we are planning an elapsed time of 2 to 3 months for data reduction per flight campaign.

Other Applications

There are many other applications not covered here: GPS and other time series data mining and event detection, data assimilation, tsunami simulation, precision orbit determination, and pattern informatics. The needs for these applications are similar to the above applications.

Hardware Requirements Analysis

Computational Requirements and Methodology

We base our requirements analysis on the exemplary applications described in the previous section. We categorize compute needs in two categories: Ensemble and Heroic. We estimate computing needs for each of these categories for the years 2008 and 2013. We also distinguish resource usage and limitations by what is actually done and what is scientifically required. We define these terms precisely below.

- Ensemble (capacity) computing involves a $500,000 to $1 million machine. These are typically modest-sized applications of mature codes used in a production mode.
- Heroic (capability) computing involves a $30 million machine. Heroic applications use most or all of available resources to solve grand challenge problems.

We define two values for 2008 resource constraints corresponding to:

- Scientific constraints: Limited by manpower to formulate and analyze scientifically useful problems. Labeled **Science** in Tables 2 and 3.
- Resource constraints: Limited by current available hardware. Labeled **Actual** in Tables 2 and 3.

Finally, we define hardware for 2008 and 2013 in terms of computing cores:

- 2008: We assume a hypothetical cluster based on the current best hardware (four cores per socket, 1–4 GB/core). This is the current commodity "sweet spot."
- 2013: We assume 10x better machines than available in 2008, but the cores have the same performance as today.

Summary of Computational Requirements for 2008

Tables 2 and 3 summarize our survey of current computational requirements. Cluster*Years=time in years on a 512-core cluster.

2008 Ensemble Runs	Problem Size (max)	Memory Size	Number of Cores	Time per Run	Runs per Year (**Science; Actual**)	Cluster* Years
Earthquake Science	10M elements	256 GB	128	4 hours	50; 10	0.001
Polar Science	10M Elements	400 GB	256	1 day	10; 10	0.014
Geodynamo	128^3	500 GB	256	4 days	900; 400	2.2
Virtual California	128^2	10 GB	256	0.5 day	1,000; 500	0.5
UAVSAR (flight lines)	200 x 25 km^2	20 GB	10	8 hours	600; 600	5.5

Table 2: 2008 ensemble computing.

2008 Heroic Runs	Problem Size	Memory Size	Number of Cores	Time per Run	Runs per Year (**Science; Actual**)	Cluster* Years
Earthquake Science	150M elements	4 TB	2,000	12 hours	0; 2	0.01
Polar Science	20M elements	400 GB	256	16 days	20; 0	0
Geodynamo	256^3 cube	4 TB	512	>30 days	10; 0	0
Virtual California	$1{,}024^2$	10 TB	512	>7 days	5; 3	0
UAVSAR (flight lines)	200 x 25 km^2	20 GB	10	8 hours	1,800; 600	1.6

Table 3: 2008 heroic computing (essential to use Columbia or equivalent).

Note: GeoFEST (Earthquake Science application) can do larger problems than scientists have currently formulated.

Evolution of Modeling Activities: Requirements for 2013

Tables 4 and 5 summarize the extrapolated computational requirements of the exemplary applications for the year 2013. Cluster*Years=time in years on a 512-core cluster.

2013 Ensemble Runs	Problem Size	Memory Size	Number of Cores	Time per Run	Runs per Year	Cluster* Years
Earthquake Science	50M elements	1 TB	2,048	4 hours	400	0.74
Polar Science	50M elements	1 TB	512	8 days	20	0.44
Geodynamo	256^3	4 TB	512	30 days	1,800	148
Virtual California	$2{,}048^2$	20 TB	1,024	15 days	20	2
UAVSAR (flight lines)	200 x 25 km^2	20 GB	10	8 hours	600; 600	5.5
DESDynI*	2,000 x 1,000 km^2	560 GB	560	1 hour	25; 25	1.6

Table 4: 2013 ensemble computing. *DESDynI assumes launched and operational.

2013 Heroic Runs	Problem Size	Memory Size	Number of Cores	Time per Run	Runs per Year	Cluster* Years
Earthquake Science	500M elements	10 TB	5,000	30 hours	20	0.55
Polar Science	150M elements	10 TB	4,000	8 days	20	3.5
Geodynamo	512^3	40 TB	4,000	30 days	10	6.6
Virtual California	$4{,}096^2$	80 TB	2,048	30 days	30	4
UAVSAR (flight lines)	200 x 25 km^2	20 GB	10	8 hours	3,600; 1,800	3.2
DESDynI*	200M km^2	56 TB	5,600	1 hour	25; 25	16

Table 5: 2013 heroic computing. *DESDynI assumes launched and operational.

Geodynamo dominates ensemble runs in 2013. We estimate 1,000 socket *years, assuming 4–16 sockets for 512 cores (32–128 cores per chip) in 2013. In 2008, 512 cores equal 128–256 sockets.

Manpower funding constrains our ability to exploit the resources today. This will increase dramatically in 2013.

Storage Requirements

The list below gives the storage requirements (for 2008 and 2013) for our exemplary computational methods and also for solid Earth observational missions. The storage requirements for the modeling applications are typically outputs and thus are subject to some uncertainty. Note that we give the "science-driven" requirements (i.e., what the scientists would like to do; see previous tables).

GeoFEST
- 2008: 5 TB/year
- 2013: 0.5 PB/year

Polar Science
- 2008: 200 TB/year
- 2013: 1.5 PB/year (coupled with GCM)

Geodynamo
- 2008: 100 TB/year
- 2013: 80 PB/year (increased time resolution in output, increased space resolution in simulation)

Virtual California
- 2008: 50 TB/year

Instruments
- 2013: SMAP and TES-FO, 40 TB/year each
- 2013: DESDynl, .65 TB/day of raw data and 3 TB/day of products for a total of 1 PB/year (see Figure 5)

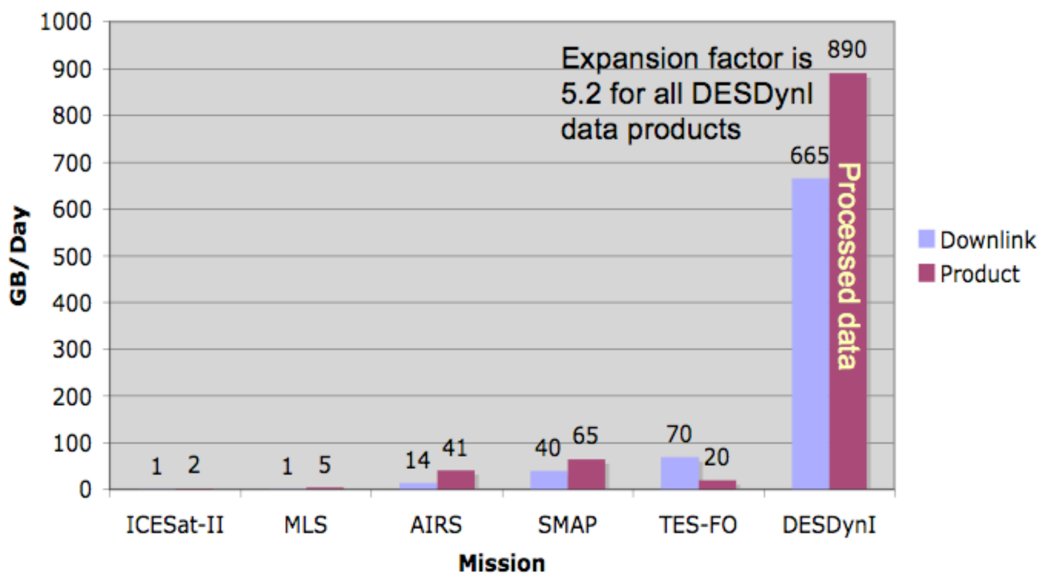

Figure 5: DESDynI approximate data volume per day for data storage requirements. DESDynI data volume is an order of magnitude more than existing and planned missions. Once downlinked, data must be moved to processing facility and distributed once processed.

Programming, Analysis, and User Environments

In addition to the hardware requirements discussed above (and implicit network requirements), we must also consider the code development, deployment, and usage environments for solid Earth.

Development: Programming Environments

All participants noted the importance of MPI to their applications; this is unlikely to change in the foreseeable future. Given the difficulty of porting codes to alternative architectures (such as Cell processors and GPUs) and insufficient manpower, it is unlikely that these alternative systems will be useful. On the other hand, we can reasonably expect commodity multi-core systems to support MPI, so this seems to be the safe strategy. We also note that many programming support tools (such as debuggers and performance tuners) will also be available on commodity multi-core systems.

Deployment: Computing and Data Clouds for Computational Science

Porting codes from development environments to multiple parallel programming systems is a notoriously time-consuming process. However, recent advances in virtualization (VMWare, Xen, etc) and Cloud computing (which uses web services to manage virtual machine and virtual data lifecycles) hold the promise for vastly simplifying this process. Instead of packaging code as tar files or RPMs, developers can package their codes and services in entire virtual operating systems or virtual computers. These are sometimes termed virtual appliances and can run in cloud hosts. The idea for virtualization has been around for many years, but the value it holds for scientific computing environments is just starting to be explored. Cloud computing also has an attractive feature for research computing centers when compared to Grid systems. Grid computing has long tried to solve the problem of connecting multiple heterogeneous computing facilities, but this has proven extremely difficult. NASA's Information Power Grid was a pioneering effort that exposed many of these problems to the Grid community. Cloud computing holds the promise of a more homogeneous collection of resources. We may term this the "Google Model" or the "Amazon Model" for resource management.

There are three common types of Cloud systems. Computing clouds such as Amazon's Elastic Computing Cluster use web services to manage Xen virtual machines. Data clouds such as the Google's File System, MapReduce, and Big Table projects provide very scalable distributed file systems with programming tools for distributed operations and queries. Open source versions of Cloud computing include Apache's Hadoop, the University of Chicago's Virtual Workspaces/Nimbus project, and the University of California, Santa Barbara's Eucalyptus project. These can be used to build locally controlled testbeds today, but we anticipate production deployments will be the norm by 2013.

Several problems currently exist when applying cloud computing to scientific problems. MPI will work on virtual clusters, but performance seems to be limited to Gbps speeds. This situation results from both the lack of drivers for higher-speed connections like InfiniBand and also limits the hypervisor used by many virtualization technologies. This also dramatically limits the ability to interact with high-performance distributed file systems such as Lustre and GPFS.

Usage: Web Services, Portals, Workflows

The last phase for all of the applications and data catalogs (hosted on Clouds or otherwise) is to be made available as online services or "appliances" to a larger community of users. This user community may be directly involved in running the codes, or they may be simply consumers of the output (such as pipelined results from InSAR processing). The standard approach for this is known as the Service Oriented Architecture (SOA), although the details of the service styles and implementations (WSDL/SOAP or REST-style web services, for example) vary. In any case, the key concept is the same: services provide remote programming interfaces for accessing applications, data, and information. The SOA approach underlies web portals, science gateways, graphical workflow engines, and related tools: all depend on online services, and SOA approaches allow flexibility when choosing web development frameworks and workflow composers.

Many of the assumptions about the best ways to build SOA, portals, and workflow composers are being challenged by Web 2.0. Social networks; rich client interfaces based on AJAX; portals made with reusable, sharable gadgets; and service mash-ups all are pushing web-enabled science in interesting new directions. These particularly hold interest to outreach efforts aimed at students and citizen scientists.

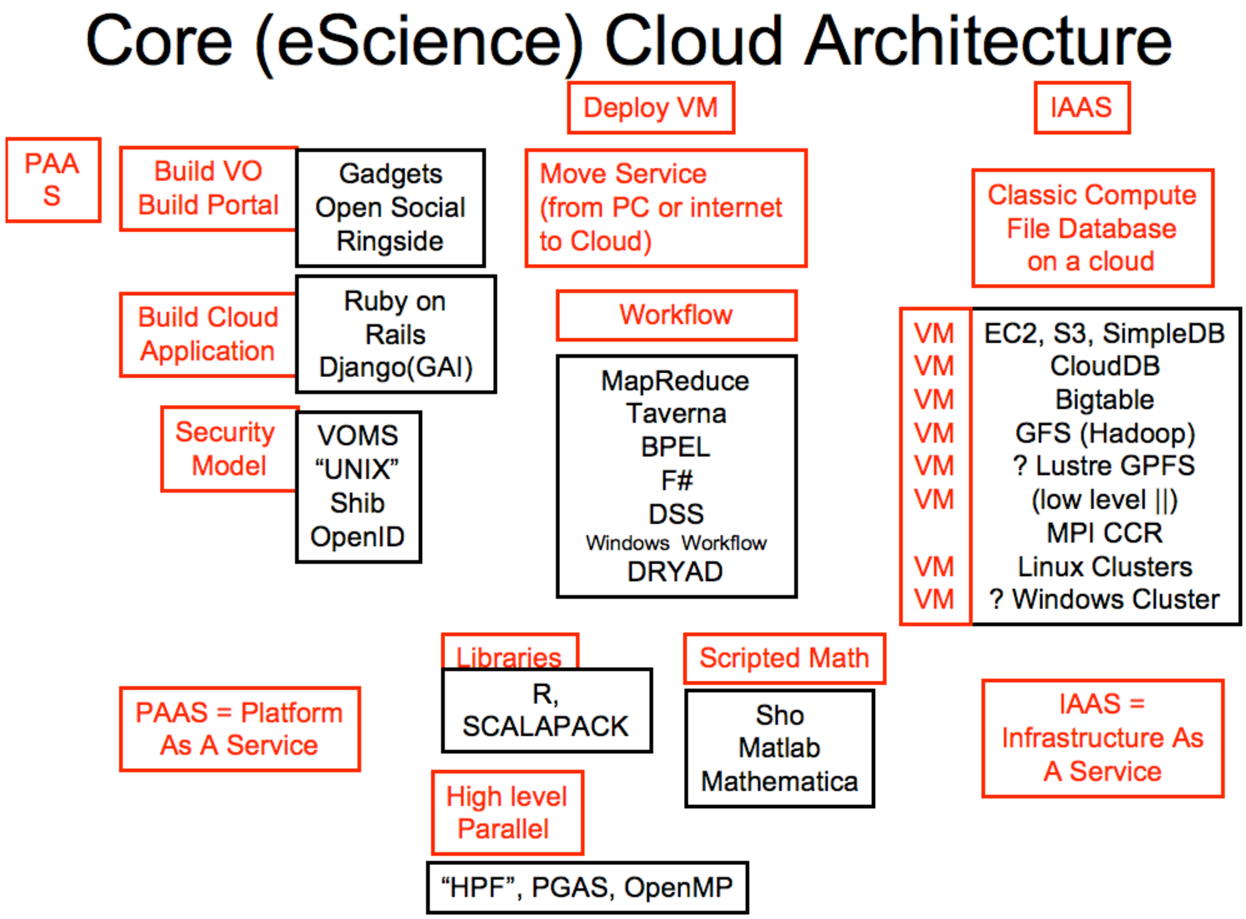

Figure 6: A summary of technologies and patterns for eScience Clouds. PAAS and IAAS (defined in the figure) are two approaches to Clouds that expose different levels of functionality to the user or developer.

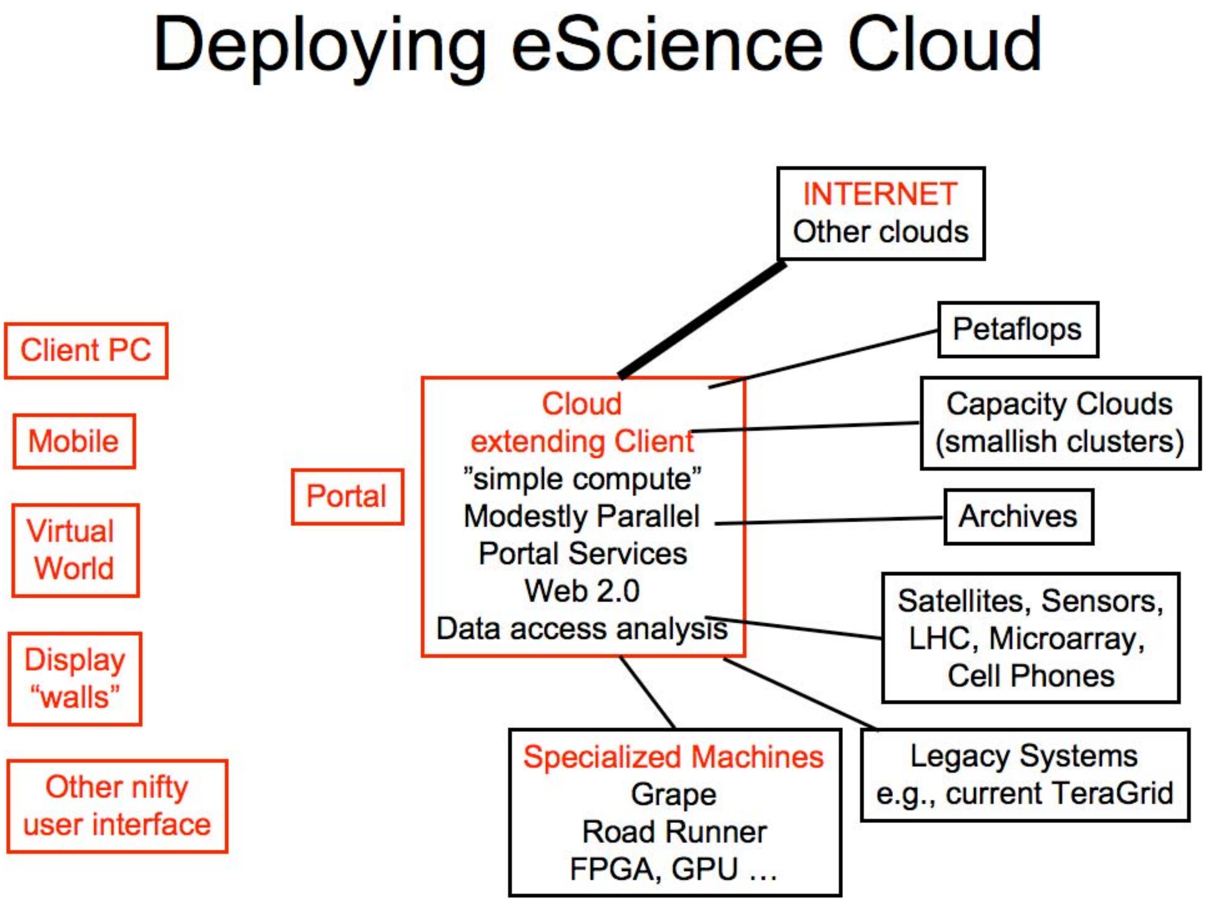

Figure 7: A diagram showing one approach to deploying an eScience Cloud.

Recommendations

Expect to use "mainstream" multi-core clusters (around 2–4 GB per core).

Support of ensemble assimilation is critical—capacity systems are required.

- **Recommendation 1:** Fund a Cloud of utility (~$1 million) clusters (six to eight just for this application here) with portals (to increase user productivity) and web services.

Solid Earth research's primary constraint is the number of scientists and developers, not hardware.

- The pipeline of computational scientists (Earth and computer scientists) needs to be improved (15 years ago, one panel member had 80 students in a parallel computing class; now, there are zero).
- It is hard to fund code development.
- **Recommendation 2:** There should be new funded positions for entering (Earth and computer) scientists that have clear attractive career paths (~10 in near term for three applications in this panel).

Remote visualization and analysis need support.

- Improve networking to the scientist.
- Approximately 100 GB/hour is needed today for Geodynamo and is not available.

- **Recommendation 3:** Fund development of web-based visualization services directed for a large user base, as well as development of visualization tools.

Panel Membership

Andrea Donnellan, NASA/JPL, Co-Chair
John LaBrecque, NASA/HQ, Co-Chair
Geoffrey Fox, Indiana Univ., Co-Chair
Robert Granat, NASA/JPL
Tony Gualtieri, NASA/GSFC
Weijia Kuang, NASA/GSFC
Eric Larour, NASA/JPL
Michael Little, NASA/LaRC
Charles Norton, NASA/JPL
Marlon Pierce, Indiana Univ.
Paul Rosen, NASA/JPL
John Rundle, Univ. of California, Davis

Astrophysics Panel

Overview

Computational science provides critical support for NASA's Great Observatories and other missions. Astrophysical simulations are required for planning future missions and data analysis strategies, understanding mission data, and generally optimizing scientific value from NASA's investment in these missions. Increasingly, our understanding of the universe is codified in computational models. Advances in observation are yielding an increasingly rich understanding of the physical processes underlying astronomical phenomena and are demanding increasingly rich astrophysical models. In turn, such models drive increasing demand for computational resources. Astrophysical simulations currently use 18.5 million CPU-hours per year of NASA's computational resources. To meet the needs of NASA's scientific goals, our primary recommendation is for continued growth in computational resources, amounting to 100 million CPU-hours within 5 years.

Science Drivers

Astrophysical simulations directly support NASA's astrophysical goals. For each of the primary astrophysics research objectives identified in the NASA Science Plan, the table below shows a few simulations in that area, along with the target missions that they assist.

NASA Science Goal	Representative Simulation Areas	Target Missions	Methods
Understand the origin and destiny of the universe, phenomena near black holes, and the nature of gravity.	Cosmology and large-scale structure	HST, WMAP, Spitzer, Planck, JWST, CMBP	N-body, hydrodynamics
	Active galactic nuclei, gamma-ray bursts	Fermi (GLAST), Chandra, HST, Swift, BHFP	GRMHD, RMHD, RPIC, GR radiative transfer
	Compact objects and gravitational radiation	LISA, Swift, Fermi (GLAST), BHFP	GR, GRMHD
Understand how the first stars and galaxies formed and how they changed over time into the objects observed in the present universe.	Galaxy formation	JWST, HST, LISA	N-body, hydrodynamics, radiation transport
	Star formation	Spitzer, JWST, SOFIA	Hydrodynamics, radiation transport
	Supernovae, supernova remnants	Swift, Chandra, HST, JDEM, Fermi (GLAST), JWST	Hydrodynamics, nuclear networks, radiation transport

NASA Science Goal	Representative Simulation Areas	Target Missions	Methods
Understand how individual stars form and how those processes ultimately affect the formation of planetary systems.	Proto-planetary disks	Spitzer, SOFIA, JWST, HST, Kepler, TPF, SIM	Hydrodynamics, MHD, radiation transport, chemical networks
	Planet formation and dynamics	Kepler, SIM, TPF, JWST, Spitzer	N-body, hydrodynamics
Create a census of extrasolar planets and measure their properties.	Exoplanet internal structure	Kepler, Spitzer, JWST	Hydrodynamics
	Exoplanet atmospheres and surface evolution	Spitzer, TPF, HST, JWST	Hydrodynamics, GCM, radiation transport, chemical networks

Given the broad range of exciting scientific work covered by astrophysical simulations, it is not possible to characterize all these areas in any detail. For the simulation areas in the table, it is convenient to provide an overview of the simulations in terms of astrophysical scales.

Cosmology, Large-Scale Structure, and the First Stars and Galaxies

At the largest scales are cosmological simulations. Simulations provide the best understanding of how the small-density perturbations now observed by missions such as the Wilkinson Microwave Anisotropy Probe (WMAP) have grown and evolved to produce the objects observed in the present-day or recent universe, such as the galaxies observed in the Hubble Space Telescope (HST) deep field view. Cosmological large-scale structure simulations explore the interaction of dense concentrations of dark matter with gas condensing to form galaxies and dynamical clusters of galaxies. The details of this process ultimately depend on the interaction these large structures have with the stars and black holes forming within them. More narrowly focused galaxy formation simulations involving dark matter, gas dynamics, and radiation relate to today's observations of galaxies as well as future observations of earlier structures that the James Webb Space Telescope (JWST) will uncover.

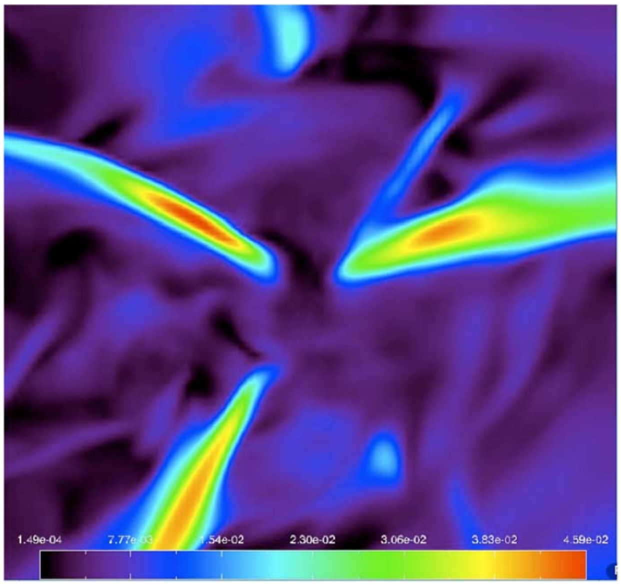

Figure 1: Still frame from a 3D MHD simulation of plasma in the cooling core of a galaxy cluster, showing a slice running through the center of the cluster. Magnetic field line conduits shown in the figure may be responsible for the heat transport in cluster cores, a process that could have a significant impact on their evolution, despite the magnetic field forces being relatively weak with respect to that of gravity. The color bar corresponds to the magnitude of the magnetic field. The size of the box corresponds to about 200 kiloparsecs. The simulation was carried out with the MHD code Athena with a resolution of 100^3.

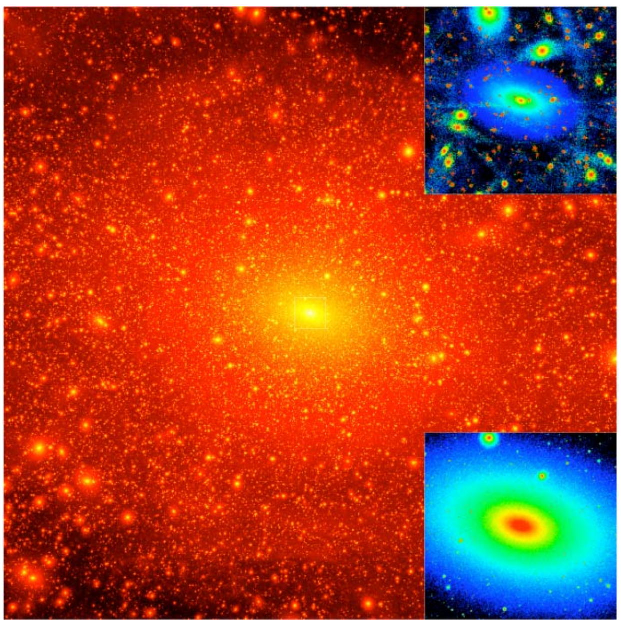

Figure 2: Projected dark matter density map of "Via Lactea II." An 800-kiloparsec cube is shown; the insets focus on an inner 40-kiloparsec cube, in local density (lower inset) and in local phase space density (upper inset). The simulation follows the growth of a Milky Way-size halo in a ΛCDM universe from redshift 100 to the present. It was performed with the parallel treecode PKDGRAV2 and samples the galaxy-forming region with 1.1 billion particles of a mass equal to 4,100 solar masses. The force resolution is 40 parsecs. Cosmological parameters were taken from the WMAP 3-year data release. The high-resolution region is embedded within a large periodic box (40 comoving megaparsecs) to account for the large-scale tidal forces. The mass within 411 kiloparsecs (the radius enclosing 200 times the mean matter density) is 2,000 billion solar masses.

On a smaller scale, within these forming galaxies dense clouds of gas condense into the formation of the first stars. We have not yet observed these stars directly, so they are presently only understood through simulations modeling gas dynamics and radiation transport. These first stars will impact their environment, changing the chemical composition of the primordial gas, releasing radiant energy that feeds back on the larger-scale processes, and altering the formation process of future generations of stars.

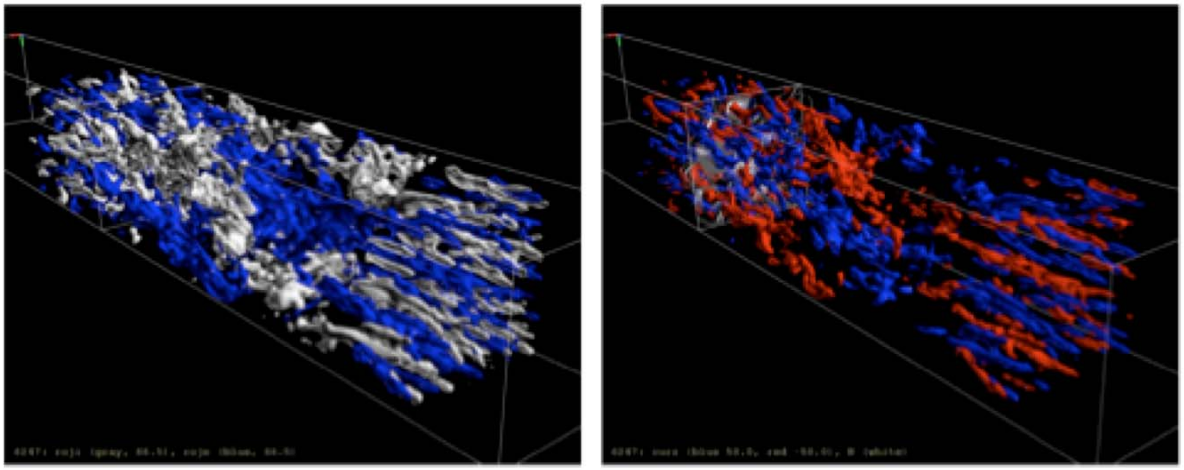

Figure 3: Snapshots viewed from the front of a jet. The left panel shows the isosurface of jet electron (blue) and positron (gray) density. The right panel shows the isosurface of the Z-component of the current density (Jz = blue, and Jz = red) with the magnetic field lines (white) in the linear stage for the case of a mono-energetic jet.

Black holes can quickly form from the collapse of these stars, or perhaps from the direct collapse of dense gas clouds. Larger black holes form and grow by accreting gas or by merging together. The largest black holes migrate to the centers of their host galaxies, where rapidly accreting gas glows brightly through frictional heating, powering some of the brightest objects of high-energy astrophysics known as Active Galactic Nuclei (AGN), which produce jets of ions. These are powerful sources for a broad range of instruments, including current missions such as HST and Chandra, as well as planned missions such as JWST and the International X-ray Observatory (IXO). Simulations of these systems, rich enough to support precision observations, require increasingly detailed physical modeling, including not only the effects of MHD but also the effects of special and general relativity, radiation transport, and more.

Compact Relativistic Systems

Typical massive galaxies in the present-day universe are observed to have black holes within their central bulges. Theoretical models suggest that many of their precursors, early galaxies and proto-galaxies, are likely to have contained large black holes even early in the structure formation process, at redshifts $5<z<15$. The building up of these structures through mergers is expected to lead to mergers of their constituent black holes. These events release the most powerful bursts of energy since the Big Bang, all in the form of gravitational waves. The Laser Interferometer Space Antenna (LISA), a planned joint mission of NASA and ESA, will directly observe these events to gain understanding about these black holes, their environments, the structure of the early universe, and the nature of gravity in its most extreme form. Physical modeling of these systems, based on Einstein's equations of general relativity, is necessary to predict the waveforms.

Systems involving star-sized black holes and neutron stars power other sources for high-energy astronomical observations, including Gamma-Ray Bursts (GRBs), which are primary sources for missions such as Swift and the Fermi Gamma-ray Telescope, formerly known as the Gamma-ray Large Area Telescope (GLAST). Likely sources for these missions also include neutron star mergers and supernovae. Simulations of these events require even more-detailed physical models, including handling networks of nuclear reactions and neutrino transport.

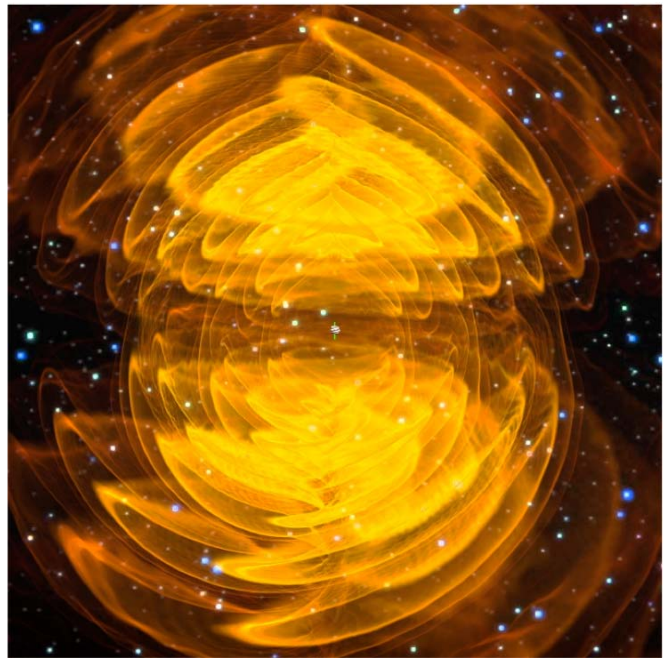

Figure 4: Numerical relativity simulations of the gravitational radiation produced by merging black holes will form part of the basis of LISA's science operations.

Exoplanetary Science

Stars are formed by the gradual collapse of dense clouds of gas and dust. Once a star ignites in the center of the disk, the remaining material may form a proto-planetary disk. While it continues to interact with its star, portions of the disk may coalesce or agglomerate, forming planets. Once formed, these planets may migrate through interactions with the remaining disk before settling into long-lasting stable orbits. Simulations of these processes are needed to understand extrasolar planet observations and to predict the abundance and conditions of possible Earth-like planets. On a still finer scale, understanding the conditions for potential life on these planets requires detailed modeling of planetary surface evolution and atmospheres.

Models

Computing capability requirements for science modeling are driven by several complementary factors. In simulating astrophysical scenarios there is always a need for higher resolution, both spatial and temporal, in order to increase the fidelity of the simulations and match improvements in data resolution. Continually improving physical modeling, incorporating more microphysical processes, calculating larger reaction networks, and including additional physics all drive the need for more computing capability, even at fixed resolution. It is also necessary to match the qualitative improvement in current data collection by missions and to anticipate and model future mission data capabilities.

Methods

Astrophysical simulation studies typically require the evolution of data on a 3D spatial domain through time, founded on basic models for the underlying physics governing the system. For example, in some cases the system can be described in terms of a large number of astrophysical objects that move through the spatial domain, with their motion and interaction governed by the basic physical equations. N-body dynamical studies of galaxy formation are an example of this type of approach. Realistic simulations do not allow a single computational body for, say, each star in a galaxy, so there is a limit to the level of detail that can be treated with physical realism. In general, greater physical realism is achieved by using a larger number of bodies, up to 1 billion or more, increasing the computational cost. Physical treatment of the propagation of light through the computational domain, known as radiation transport, can be handled with a similar approach.

Alternatively, astrophysical systems can be modeled by fields of numerical values covering the entire computational domain. The values describe crucial physical properties at each point, such as the density of matter or the strength of gravitational or electromagnetic fields, with physical models providing equations for how the field values evolve in time. Many astrophysical phenomena are modeled with a variation of the basic theory for the dynamics of gas-like materials known as hydrodynamics. Depending on the circumstances, additional details of physics may be crucial to understanding an astrophysical system. If charged particles are present, it is often crucial to consider the effect of magnetic fields, in a class of models known as magnetohydrodynamics (MHD). When velocities are large, special relativity must be taken into account as well, requiring relativistic MHD (RMHD). Strong gravitational fields, such as those governing black hole interactions are described by general relativity (GR). Where matter is present around black holes, or in strongly gravitational objects such as neutron stars, both GR and MHD effects may be important, requiring the combined modeling technique of GRMHD.

For such systems, the physical space might be divided into a regular 3D grid. Physical realism requires a sufficient density of grid points to represent the crucial details of the system. Often there are localized regions where important phenomena, such as hydrodynamic shocks, must be treated with high-resolution grids, requiring techniques such as adaptive mesh refinement (AMR) to achieve sufficient realism over a sufficiently large domain. Even so, today's simulations often require 10 million grid points or more, with increasing demand for higher fidelity, and thus, larger computational domains.

In cases where there are different types of material interacting, it can be necessary to also model how materials may be involved in chemical or nuclear reactions. A model describing networks of possible reactions must be included in such cases. These chemical or nuclear networks are then coupled with one of the dynamical models discussed above, adding complexity to the computational problem.

The study of exoplanets may involve these general components of astrophysical models. The surface physics involve phenomena similar to those studied in great detail in Earth science, so these studies may draw from specific Earth science models such as Global Circulation Models (GCMs).

Case Studies

Project 1: Gravitational Waves from Binary Black Hole Mergers (GSFC Numerical Relativity Group)

Simulations: Inspiral, merger, and ringdown of black hole binaries, of various spins and mass ratios. These simulations typically evolve ~10 million grid points through ~200,000 timesteps.

Relevance: Binary black hole mergers are expected to be a key gravitational wave source for LISA. Understanding LISA's observations, as well as understanding how well LISA will be able to measure the characteristics of these systems, will depend on precise predictions based on Einstein's theory of gravity through supercomputer simulations, known as numerical relativity. NASA's past supercomputing investment has had a direct and profound impact on bringing this field to maturity, so that precision calculations of the gravitational wave signatures from binary black hole mergers are now being produced around the world. LISA's observations will be of especially high value if electromagnetic counterparts can be expected. This opportunity motivates the study, using new GRMHD simulations, of how gas may behave in the vicinity of these systems. These simulations have also had an unanticipated impact on interpretation of current observations with the discovery that asymmetric gravitational wave emissions may expel black hole from their host galaxies.

Current resources used: Typically 400 to 2,000 processors for up to 100 wall-clock hours.

Demand drivers: There remains a broad parameter space to study, including important unstudied areas such as systems with large differences in the sizes of two black holes. The most likely systems may have mass ratios in the vicinity of 20:1, but such simulations are not yet tractable with current resources and techniques. LISA science will require large numbers of longer-lasting simulations with higher resolution than are currently common. Increases in computational needs are motivated by the need to conduct broader studies of binary black hole parameter space, eventually including systems with mass ratios perhaps as great as 100:1. Even under ideal scaling assumptions, such simulations would be expected to be about 100 times more costly than current simulations. While investments in developing new methodology can be expected to result in improved efficiency, these are nonetheless likely to be computationally expensive calculations.

Future requirements: Expect to increase by a factor of 10.

Project 2: Cosmological Reionization (Renyue Cen, Princeton University)

Simulations: Cosmological reionization simulations including concurrent treatment of star formation with a high-resolution N-body code for dark matter (current dynamical range 10^{10}), high-resolution real hydrodynamics for the intergalactic medium, and accurate 3D radiative transfer with ray tracing.

Relevance: These simulations are timely and will help set up a solid framework to properly interpret the proliferating observational database at $z>5$ from major NASA missions (HST, WMAP, Spitzer, and Planck, among others), in conjunction with ground-based facilities (Keck, SDSS, 21-centimeter missions, and others), and to maximize the scientific returns of these missions. Equally important, they will provide urgently needed feedback to future major missions, in particular JWST, with regard to observations of high-redshift galaxies.

Current resources used: 512 to 1,024 processors run for 150 to 300 wall-clock hours.

Demand drivers: In order to cover the necessary dynamic range, 10^{12} particles are needed, which is about a factor of 40 higher than what is doable today. With this capability, a box of size 100h-1 megaparsecs could be simulated, and mini-halos of a mass equal to 10^6 million suns could be resolved, i.e., such a simulation will be able to resolve radiation sources and sinks and, at the same time, capture large-scale structures.

Future requirements: 20,000 CPUs, 200 TB of RAM and 1 to 10 PB of disk space would optimally be used. An estimated wall-clock time for such a simulation is 1,000 hours. An SMP architecture and Fortran support are preferred.

Project 3: Galaxy Formation and Evolution (Renyue Cen, Princeton University)

Simulations: When and how did galaxies form? How did galaxies evolve? Do cold dark matter model predictions agree with observations? To answer these grand questions requires cosmological simulations that have unprecedented dynamic ranges and physical sophistication. Current simulations have a dynamical range of dark matter particles of 10^{10}, which is just marginal for resolving the interstellar medium at a redshift of z=3.

Relevance: These simulations will provide both the light production history and mass assembly history of galaxies, to be confronted with observations from major NASA missions (HST, Spitzer, JWST).

Resources currently used: 1,024 to 2,048 processors for 600 to 1,200 wall-clock hours.

Demand drivers: A spatial dynamic range of 2^7 will be needed, with an AMR hydrocode.

Future requirements: 100,000 CPUs, 200 TB of RAM, and 1 to 10 PB of disk space will be required. An SMP or multi-core cluster architecture is good, and Fortran support is desirable.

Project 4: Dark Matter Halo of the Milky Way (Piero Madau, University of California, Santa Cruz)

Simulations: A recent groundbreaking simulation evolved 1 billion particles over a physical period of 14 billion years.

Relevance: Will help interpret data from the Fermi/GLAST mission.

Resources currently used: This simulation ran on 3,000 processors of the Oak Ridge National Laboratory Jaguar system, for roughly 1 million CPU-hours.

Demand drivers: In order to resolve the dark matter better than 1,000 solar masses per particle, it will be necessary to scale up by a factor of at least five. Note that a simulation of 3 billion particles was recently performed at Germany's Max Planck Institute.

Future requirements: These simulations of several million CPU-hours cannot practically be performed on less than several thousand processors. Correspondingly, several hundred TB of storage are required.

Project 5: Extrasolar Gas Giants (Doug Lin, University of California, Santa Cruz)

Simulations: Atmospheric dynamics of extrasolar gas giants have been simulated with radiative hydrodynamics.

Relevance: Explains temperatures of gas giants inferred from Spitzer data and informs models of planet formation.

Resources currently used: Currently running on 64 processors.

Demand drivers: Sufficient accuracy will require a 10-fold increase in resolution.

Future requirements: Ideally 640+ processors.

Recommendations

Recommendation 1

Our highest-priority recommendation is to increase the total number of available processors. The total SMD allocation for astrophysics applications for the current year is 18.5 million CPU-hours. We found that most researchers polled need to scale up their simulations by five- to 10-fold over the next 5 years in order to cover adequate spatial scales with adequate resolution. In addition, the breadth of astrophysical phenomena being fully interpreted with the aid of significant computational modeling continues to grow. Over the next 5 years, we conservatively predict the demand for astrophysics computation time to increase to roughly 100 million CPU-hours per year. Also note that 2 GB or more of RAM per processor is preferred for most of these applications; otherwise processor-demand might be doubled just to accommodate memory requirements.

Recommendation 2

Correspondingly, archive memory will need to be increased by five- to 10-fold. Our non-exhaustive poll indicates that well over 2 PB of archival storage will be required. Along with this requirement comes the need for faster storage retrieval or remote interactive tools. Alternatively, in some cases storage demands might be partly mitigated by development of concurrent visualization.

Recommendation 3

Our next highest-priority recommendation is that of funding for manpower. This recommendation applies to both facility staff support and training of students and postdocs. Training for the latter would include both developing and running code. In some cases, other needs include help in making code more scalable. Also of benefit would be expansion of facility staff support services into software "clinics," where staff experts can perform profiling on codes and suggest specific code changes to take advantage of various hardware architectures more efficiently. Staff support for development of concurrent visualization may also be valuable.

Recommendation 4

Computing centers need to address difficulties with current queue scheduling protocols. Queue structures are currently such that an 8- to 24-hour wall-clock limit is typically imposed on runs. Restarting capabilities are ubiquitous, but a frequently encountered problem is the long wait-time between restart segments. Of course, increasing the number of available processors, as suggested above, would mitigate this particular problem. Secondary problems include limited storage space for checkpoint files, as well as wasted overlap time between restarts. The need for frequent checkpoints makes adequate storage space, already discussed above, all the more critical. Our only specific recommendations in regards to queue scheduling are for queue-administrators to poll users on the wall-clock duration of typical simulations and to consider greater queue flexibility. Another suggestion to consider is to somehow distinguish restart jobs and give them higher priority than fresh-start jobs, so as to reduce the wait-time between restarts.

Recommendation 5

Our final recommendation is to consider small investments in processors with new or unconventional architectures for experimental purposes. For example, it is possible that the superior speed of the GPU would be of benefit to the astrophysics community. However, its lack of Fortran support and significant RAM make its utility questionable. Thus, a small GPU "laboratory" system would be useful to determine whether and how the astrophysics community can make use of GPUs. Another experimental architecture is that of OpenMP clusters, which could enable those groups using OpenMP instead of MPI to do so more efficiently.

Panel Membership

Joan Centrella, NASA/GSFC, Co-Chair
Jim van Meter, NASA/GSFC, Co-Chair
Michael Salamon, NASA/HQ, Co-Chair
John Baker, NASA/GSFC
Renyue Cen, Princeton Univ.
Akshay Kularni, Cornell Univ.
Doug Lin, Univ. of California, Santa Cruz
Steve Lubow, Space Telescope Science Institute
Piero Madau, Univ. of California, Santa Cruz
Scott Michael, Indiana Univ.
Ken-Ichi Nishikawa, NASA/MSFC
Steinn Sigurðsson, Penn State
Stephen Simms, Indiana Univ.
John Wise, NASA/GSFC

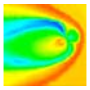

Heliophysics Panel

The progress in understanding the overall dynamics of the Sun-to-Earth or Sun-to-planet chain has created an increasing desire to describe the relevant physical processes quantitatively. At the same time, the growing national need to forecast and describe space weather mandates a transition from discovery and qualitative scientific description to the deep level of quantitative understanding required to forecast harmful space weather effects.

As a consequence, modeling, and, in particular, numerical simulations have grown in importance to the Heliophysics science enterprise. Models are now widely used by the research community to assist in the scientific analysis of spacecraft-provided datasets, as well as in mission planning and conception. Furthermore, modeling has evolved into a core element of programs aiming at the development of new space weather forecasting capabilities. These developments have exposed scientific computing as an important foundation of the Heliophysics enterprise.

Heliophysics modeling falls into two categories. In the first category, individuals and groups of researchers require large computational resources to address forefront scientific problems. The requirements associated with these calculations are typically too demanding for a large group of researchers to participate—yet these calculations push the forefront of knowledge. Therefore, there is a definite need for large, centralized computational support.

The second category includes the much larger community consisting of researchers interested in more limited problems, specialized data analysis needs, users of model results executed on request, and generally researchers interested in the development of space weather forecasting capabilities. The large number of required calculations, the tailoring of data stream processing, the required rapid turn-around, or, at times, the protection of intellectual property typically mandate that these calculation be performed on smaller systems. In order to optimize access and tailoring, these systems are best operated at the scientist's site.

Following the team charter, this report focuses on larger computations. It is, however, important to recognize that evolution of larger-scale computing at the expense of smaller, distributed computing would be detrimental to the overall scientific progress in Heliophysics.

The following section presents a sample of forefront scientific problems that would benefit from new capabilities in large-scale computing. A subset of these problems, which have computational requirements characteristic of other science problems, is discussed in greater level of detail. The following sections discuss space weather modeling, data analysis, and community model access. Further sections analyze technological and support needs, and the final section sums up our recommendations.

Science Drivers and Computational Needs

This section discusses a selection of science drivers for forefront computing. The goal here is not to be comprehensive—instead, the problems listed below serve as examples for a considerably larger class of scientific problems, which benefit from modeling. It should also be noted that the computational requirements listed here pertain only to the largest, leading-edge-type calculations.

For these as well as many other scientific problems, there is an ongoing need for smaller computations, which are best executed on smaller computational platforms right at the modeler site.

Magnetic Reconnection

Magnetic reconnection is arguably the most important plasma transport and energy conversion process in space and astrophysical plasmas. Magnetic reconnection enables plasma transport across magnetic barriers, and it facilitates the conversion of stored magnetic energy into particle kinetic energy. Magnetic reconnection provides the energy release in solar eruptions, the energy entry into the magnetospheres of the Earth and of other planets, and energy conversion inside these magnetospheres. Magnetic reconnection also directly accelerates particles to sufficiently high energies to be of concern for humans and their assets in space.

In addition, magnetic reconnection is believed to be an important process in astrophysical plasmas, such as occur in pulsar winds; as a heating mechanism in the galactic halo; and in astrophysical jets.

Owing to its importance, NASA has embarked on a space-based study aimed at understanding magnetic reconnection. NASA's Magnetospheric MultiScale (MMS) mission will have the instrumentation to study in-situ magnetic reconnection within the Earth's magnetosphere. MMS encompasses a modeling program; however, available resources limit the realism of the approximations adopted in reconnection modeling.

The universal importance of magnetic reconnection renders a realistic, kinetic model of magnetic reconnection a high priority for numerical modeling, and this problem is therefore adopted as a sample, grand challenge problem.

It is illustrative to consider magnetic reconnection modeling in two spatial dimensions only, with a proton and electron plasma of realistic mass ratio, in a limited system of approximately 3,000 x 1,500 km in the magnetosphere. In order to model this system, the following is needed:

	2008	**2013**
Resolution	9,600 x 4,800	9,600 x 4,800 x 1,000
Particle number (100/cell)	5×10^{9}	5×10^{13}
RAM requirement	2×10^{11} bytes	2×10^{14} bytes
Floating point ops/time-step	4×10^{11}	5×10^{14}
Total number of time-steps	6×10^{5}	6×10^{5}
Floating-point operations per second (flops) requirement for execution w/in 1 week	1 time-step/second 4×10^{11} sustained (1 teraflops sustained)	1 time step/second 5×10^{14} sustained (1 petaflops sustained)
Storage requirement for entire run	6×10^{12} bytes	6×10^{15} bytes

Physics of Shocks, Particle Acceleration, and Rotational Discontinuities

After Voyager 1 and 2 crossed the termination shock in 2004 and 2007, respectively, it has become clear that the physical processes at shocks and in sheaths are much more complex than previously expected. For example, recent observations indicate that the termination shock is a particle-mediated shock. The temperature in the heliosheath was significantly lower than the value expected, with the remaining energy possibly in pick-up and suprathermal ions and waves. Neither spacecraft observed the near-power-law spectral shape expected for anomalous cosmic rays (ACRs) at the time of shock crossing. As possible explanations, ACR source locations either in the shock flank or tail, or deeper in the heliosheath, have been proposed. Observations also contain indications for heliospheric asymmetries, for example, the termination shock crossing by Voyager 2 by 10 AU closer to the Sun than Voyager 1, pre-shock ion anisotropies, and east-west asymmetries in energetic neutral atoms observed with STEREO. The interpretations of these puzzles are complicated by solar cycle effects. Voyager 1 crossed the shock near solar maximum and Voyager 2 near solar minimum, making the separation of spatial and temporal effects challenging. With the launch of IBEX, additional diagnostics becomes available through global maps in energetic neutral atoms. The recent observations point to a richness in the shock and heliosheath physics yet to be explored.

In order to fully explore these regions, sophisticated numerical modeling needs to work in close proximity with observations. The numerical modeling community was driven to increased sophistication using 3D MHD with adaptive mesh refinement (AMR). Because the processes involved are both spatial and temporally disparate, a complete model that will be able to tackle this region will have to utilize both a fluid and a kinetic approach, especially in the vicinity of the shock. For example, termination shock particles were observed ahead of the shock 3 to 4 years streaming in opposite directions to Voyager 1 and 2. In order to fully explore the acceleration of the particles in the shock, cross field diffusion and turbulence need to be included in a local kinetic model, while global simulation needs to be present to capture the overall geometry of the shock. Only then will we be able to fully understand the complexity of the shock physics that will have consequences not only for the termination shock, but for other heliospheric as well as astrophysics shocks.

Usual Solar System-Interstellar Medium Run
Resolution: Minimum delta x 0.2 AU near the boundary.
Iterations: Can go to 100,000 to reach a steady state of 4 million cells.
Calculation (similar to a CME): Requires 20,000 processor-hours and runs in 76 hours clock-time on 256 processors at NASA Ames Research Center (ARC) (this is for one steady state).

This is just a stationary run, where no solar cycle effects were included. This is also only to resolve the global structure of the heliosphere; the current sheet and its tilt with solar cycle are neglected. (To estimate, for example, the "memory of the heliosheath," several solar cycles will be needed.)

Ideally, it is important to estimate the effect of Global Merging Interacting Regions, and this will require coupling with the inner heliosphere Sun-Earth modules, which will dramatically increase the computational requirements. For these runs as well as the one to resolve the heliospheric current sheet, a resolution at least of 0.1 AU is needed, and the number of cells will increase beyond 1 billion.

Turbulent Dissipation

The turbulent evolution of magnetic and velocity fields in space plasmas is known to play an important role in the heliosphere and may be important in heating the solar corona and accelerating the solar wind. Turbulence also plays a role in the Earth's magnetosphere, in particular in the acceleration and heating of particles in the magnetotail. The cascade of energy from large to small scales due to nonlinear fluid-like interactions eventually reaches a scale where the detailed motion of the particles is important. In ordinary fluids, the energy heats the particles, but in plasmas the dissipation process may also result in the acceleration of particles. The production of high-energy tails on the distribution may also be essential to understanding problems such as solar wind acceleration. The processes will involve complex wave-particle interactions. The computational requirements for understanding the dissipation of turbulence are at least as great as those involved in reconnection, in that both ion and electron scales are important. Furthermore, it will be important to run the code for many ion timescales, to follow the turbulent cascade, while following the electron motion in detail. Thus, this problem is of deep significance to heliospheric physics, and will involve the same level of computational resources as the intensive simulations of reconnection.

Initiation of Flares and Coronal Mass Ejections

Coronal Mass Ejections (CMEs) and solar flares are the principal drivers of all strong space weather disturbances. Despite many years of study, the physical processes active in the initiation of CMEs have not been clearly identified. Direct observation of coronal magnetic field evolution would clarify many controversial issues. However, the techniques for high-quality coronal field observations are still in their infancy, so modeling of these regions is essential.

Theoretical ideas about the way in which the CME and flare energy are stored in the corona and about the trigger mechanisms that initiate its release are tested using large 3D ideal or resistive MHD codes. The dynamic range of scales important in these simulations is extremely large. The pre-event magnetic configurations can be of order 50,000 km in size. CMEs can span a solar radius in size by the time they reach 10 solar radii from the Sun. Reconnection processes in current sheets of scale 1–10 km may be responsible for the event onset. Sophisticated AMR techniques are often used in an effort to reproduce a small part of this dynamic range in scale. Resolving the thinnest current sheets may never be feasible. The long-term hope is to supplement these adaptively refining models with appropriate sub-grid parameterizations, guided by insights from an improved understanding of fundamental reconnection physics. By extensive experimentation and comparison with high-resolution observations, these models will clarify the basic processes and evolutions of these events.

In the short term, the most appropriate advances will be to push the model resolutions toward the resolution supported by the latest solar observations. To model a typical active region with the 0.6-arcsec resolution supported by the Solar Dynamics Observatory (SDO) and an implicit 3D ideal MHD code will require: a 3D grid of size 500 x 500 x 500, a time-step of order 1 sec., a model run duration of 1 day, and 2,000 floats per time-step per grid cell, for a total of (500 x 500 x 500 x 24 x 60 x 60 x [2,000–10,000 depending on algorithm complexity]=) 2.16–10.8 x 10^{16} floats. A machine sustaining 1 teraflops will complete this calculation in 6 to 30 hours. With 200 data-words per cell, this requires (500 x 500 x 500 x 200 x 8=) 0.2 TB of memory. Given that our principal tool for exploring ideas on CME initiation is numerical modeling, many runs of this type would be required.

CME Evolution

The structure and evolution of Interplanetary CMEs (ICMEs) is still poorly understood. Direct observation of their structure is challenging because it usually relies on measurements from a single spacecraft. When multiple spacecraft observe the same ICME, they almost always do so at different times. As a result, numerical models are almost always required to develop physically coherent interpretations of the observations. These models must capture the overall scale and structure of the ICME, while also being capable of resolving the leading-edge shocks with sufficient detail to support particle acceleration models. They will need to support a multi-fluid description. They must support field-line connections to the Sun and couple with non-local or kinetic models of collisionless electron streaming.

They must also couple with models of the CME initiation at the Sun in order to adequately represent the ICME's internal magnetic structure. As noted below, turbulent evolution of magnetic and velocity fields plays a major role in shaping the interplanetary medium and so will also strongly influence models of ICMEs as they propagate. How to include the effects of this turbulence in ICME models remains an open question. The ICMEs also must be coupled with a realistic description of a background-structured solar wind.

It should be clear at this point that the wealth of new physics and algorithms to be added to these models as they are developed guarantees that they will demand extensive computational resources. Current single-fluid, adaptively refining, 3D MHD models, which capture rather than resolve the edge shocks, complete in approximately 1 week on 128 processors.

Usual CME Run
Resolution: Minimum delta x 3E-3Rs near the boundary and Sun-Earth 0.1Rs.
Iterations: Can go to 100,000 to reach a steady state of 4 million cells.
Calculation (with CME): Requires 20,000 processor-hours and runs in 76 hours clock-time on 256 processors at NASA ARC.

This run is just to resolve the "nose" of the shock, not the flanks, and to under-resolve the active region and the filament.

Photosphere-Corona Connection

Modeling the solar atmosphere to the sub-surface represents a very challenging, but essential step towards a comprehensive, first-principles space weather model. High-fidelity models of the chromosphere are required to achieve this goal. For example, as magnetic fields emerge through the photospheric surface, the physics of the very thin chromosphere determines which field-lines close locally and which expand as they emerge to contribute open magnetic flux in the corona and heliosphere.

However our understanding of the physics of the chromosphere is still so limited that all models connecting the solar sub-surface to the corona belong firmly to the research domain. Only in the last 2 years have MHD models appeared that are capable of including the upper convection zone through the low corona in a single computational domain. These models are restricted to a very small surface patch, and they make severe simplifying assumptions in their treatments of key physics, most notably radiation transport, ionization balance, and the equation of state. The simple radiative transport treatment limits the model's ability to compute observed emissions and make direct comparison with observations. The modest spatial resolution limits the model's resolution of important current sheet structures. The limited depth of the computational domain constrains the possible structuring of coronal loops and reduces the range of convection zone parameters that can be explored.

Even with these simplifications, these model runs are computationally expensive (see table below). With algorithmic improvement and better treatment of these processes, the calculations will be able to absorb all the machine cycles to which the model developers have access.

Recent (Hinode) and imminent (SDO) high-resolution observations, in combination with these model improvements, promise to lead to a substantial improvement in our understanding of this region of the Sun. This scientific payoff can be accelerated with a concomitant investment in computing support.

	2008	**2013**
Grid	750 x 380 x 140	3,000 x 750 x 300
Scales	75 x 38 x 14 Mm^3	150 x 38 x 15 Mm^3
Resolution	100 km	50 km
Processor-hours	128 processors (on NCCS Discover) x 500 hours = 64,000 processor-hours	1,024 processors x 500 hours = 512,000 processor-hours

The table illustrates the current computational performance and compares it with a run we would hope could be supported within 5 years. This "future" run is just one of a very large range of model scenarios that might be tackled over the next 5 years. These model scenarios, which can be enabled with expected growth in computational resources, include:

- Higher-resolution runs.
- Larger-volume runs (pushing the bottom boundary deeper into the convection zone requires shorter integration time-stepping).
- Runs with more complex sub-surface fields to be emerged.
- Detailed parameter studies (influence of different properties of the convection zone, degree of twist of emerging flux, complexity of the pre-existing coronal field into which new flux is emerged, etc.).
- Higher-fidelity radiation transport.
- Transient ionization balance.

Plasma Environments of Moons, Plasma Environments Near Spacecraft

Moons

The study of the interaction between plasma subsonic and supersonic flow and moons is a fundamental problem of space physics. First of all, the interaction of the solar wind with the moon may result in sputtering of small fragments from the surface and their incorporation in the plasma-flow-like pick-up ions. The dynamics of these fragments may help to make a decision about chemical composition of the lunar surface. The study of the plasma environment of the outer planetary moons (Io, Titan, Europa) may provide information about mass, energy, and momentum exchange between Saturnian and Jovian magnetospheric plasma and the moons' atmospheres. For the first time, hybrid code model runs have accounted for the transition of plasma characteristics that range from a finite-gyro-radius description far from the moon to the collisional plasma of the ionosphere. One of the essential results of this modeling effort is the description of the ion-phase space distribution and their moments in the interaction region of Titan. The 3D hybrid simulation of this type of interaction needs about a 1-month run with 16 processors and 32 GB RAM on the NCCS SGI Altix Explore system for computational models with a grid 301 x 301 x 301 and a few particles per cell. For a simulation with 10 to 20 particles per cell, we need a run of 10 months using 320 GB of memory.

Spacecraft

Another fundamental problem of space plasma is the interaction of solar wind with spacecraft, in particular for the "Solar Probe+" project. This project studies the spatial distributions of the thermal and energetic particles in the plasma environment (plasma wake) of spacecraft (Solar Probe), their velocity distributions, and the electromagnetic waves. The model results are essential in developing the design strategy of the plasma wave instrument, including orientation of the antenna and particle instruments planned for the future Solar Probe mission. First results demonstrate the formation of the whistler and at later time the Alfven wing directed by the interplanetary magnetic field. Current work focuses on further development of the 3D model, which may include the kinetic description for electron dynamics. A 3D hybrid simulation of this problem needs a 1-month run with 16 processors and 24 GB RAM on Explore.

Tangential and Rotational Discontinuities

The magnetopause (MP) is a critical region of geospace, since it controls the transfer of energy and momentum from the solar wind into the magnetosphere. Observations show that the MP is a finite thickness discontinuity, which separates the post-shocked solar wind from the magnetosphere. Depending on orientation (northward or southward) of the interplanetary magnetic field, the MP has been modeled as a rotational discontinuity, and as a tangential discontinuity, respectively. A 2D hybrid simulation of the tangential (rotational) discontinuity with realistic proton/electron mass ratio (grid 201 x 201 and 106 macro-particles) needs 500 MB memory and 1 processor for 1 week. A 3D hybrid simulation (grid 201 x 201 x 201 and 108 macro-particles) already needs 16 GB memory and 16 processors for 1 month.

Ring Current and Radiation Belt Dynamics

The radiation belts and ring current consist of energetic electrons (up to several MeV) and ions (up to several hundred MeV) trapped in the terrestrial magnetosphere in the vicinity of L shells from ~1.3 to 8. During geomagnetic storms and substorms, the particle fluxes of radiation belts and ring current can increase more than an order of magnitude over the quiet-time levels. The enhancements of these energetic particles have great impacts on the radiation environment in space. Moderate energy (~10 to 100 keV) electrons and ions can cause surface charging effects and relativistic (~0.1 to 5 MeV) particles can cause deep-dielectric charging on space systems. The intensification of the ring current and the associated strong electric field and current in the ionosphere could also lead to uncontrolled current surges in high-latitude power grids, causing damage to power grids and transformers. Therefore, understanding the physical processes that are controlling the development of the radiation belts during active periods and being able to predict their variability have both scientific and practical significance.

Numerical models have been developed to understand the dynamics of the radiation belts and ring current and how these energetic particle populations respond to various solar wind conditions. Simulation results are also used to interpret satellite and ground-based measurements. Most of the radiation belts and ring current models are kinetic models that solve the distribution functions of particle species. The NASA Radiation Belt Storm Probe (RBSP) mission was designated to understand, ideally to the point of predictability, how relativistic electrons and ions form and change in response to the variable inputs of energy from the Sun. RBSP will carry particle detectors in the ring current as well as radiation belt energy range and electric and magnetic field instruments. There is a separate program for theory and modeling to support the RBSP mission. The announcement of opportunity will be soon available to the space science community.

The radiation belts, ring current, plasmasphere, and ionosphere are strongly coupled systems. To accurately model the dynamics of the energetic populations, the configuration of the core plasmas and the feedbacks from the ionosphere have to be considered as well. Wave-particle interactions are found to be very important in the acceleration and decay of the relativistic particles. Simulation codes have been developed to self-consistently model the radiation belts, ring current, plasmasphere, magnetospheric electric field, and plasma waves. Owing to the complexity of interacting systems, this modeling effort requires large RAM and intensive computational resources. For example, distribution functions in the 2D spatial space (bounce-averaged) and 2D velocity space (gyro-averaged) require a grid of 50 × 50 × 50 × 50 × (3 species of electrons H+ and O+) and 8×10^8 bytes of RAM. If the computation resources are improved in the next few years, finer grid resolution is desirable to improve the performance of the modeling codes.

	2008	2013
Resolution:	50 × 50 × 50 × 50	100 × 100 × 100 × 100
Number of species	3	3
RAM requirement	8×10^8 bytes	1×10^{10} bytes
Total number of time-steps	70,000	70,000
Storage requirement for entire run	1 GB	16 GB

Currently, the codes are run on single-processor PCs. One day of a real-time run requires about 1 week of CPU time. In order to perform faster calculation or real-time prediction, the codes must be run on supercomputers with MPI and at least 16 CPUs.

Cross-Scale and Inter-Regional Coupling: Micro-Macro Interactions, "Sun-to-Mud"

Space plasmas are typically a multi-scale phenomenon. On scales comparable to the size of magnetospheres or of major solar features, plasmas behave like fluids, which are typically described by MHD or multi-fluid approaches. However, larger-scale systems are permeated by boundaries, which separate different plasma regimes and involve current sheets, shocks, or magnetic reconnection processes. The structure and physics of these layers are determined by kinetic processes, i.e., by processes that involve interactions between particles and electromagnetic fields.

The overall dynamics of the system, such as the evolution of the magnetic field above a solar active region, is determined by the interplay of processes on these very different scales. Processes on the small scales, which can be many orders of magnitude smaller, can provide the release of energy originally stored by large-scale convection and large-scale dynamics; in turn, these larger-scale phenomena will set up small-scale processes such as collisionless shocks. Thus, the large-scale dynamics sets up the conditions for small-scale processes, and the latter enable further large-scale evolution.

Modeling of these interactions not only poses substantial physics problems but also high demands on computational capabilities. For example, including kinetic shock or reconnection physics in a large-scale heliospheric model in principle requires executing two or more very large problems simultaneously, with data exchanges between the subsystems. Therefore, the demands for physically realistic cross-scale coupling models are in excess of those for each individual sub-problem.

Cradle-to-grave modeling of a solar eruption from its initiation to its impact in terrestrial or other planetary environments naturally requires both resolution of each domain and coupling of sub-domains across their boundaries. Accordingly, calculations of this kind entail not only the requirements for modeling of each domain but also the added complications of information transfer and execution synchronization. The often vastly different time and spatial scales render the combined problem mathematically very stiff. As a result, a minimum computational requirement for such forefront calculations involves the sum of all coupled domain calculations.

Space Weather

Space weather analysis covers the same topics as general heliophysics science, in particular taking CME, particle, and flare events (particularly shocks) as they evolve across the different scales of the inner and outer heliosphere through the solar wind and interplanetary medium all the way to the Earth (or other target body). Cross-scale coupling is essential for understanding the time evolution and impact of events. Space weather science falls within the Living With a Star (LWS) Program's goal for prediction of solar events and their impact on Earth.

Space weather scientific analysis uses validated HEC methods in near-real-time, real-time, or faster-than-real-time predictions of significant solar events that will affect Earth, near-Earth, or planetary systems. Such MHD, plasma/fluid, and kinetic modeling require multi-scale methods to fully model each single event.

While a typical HEC scientific analysis looks at just a few events at the highest spatial, temporal, and spectral resolution available within the limits of the hardware, space weather has a more rapid analysis requirement and therefore will frequently use lower resolutions as a trade-off in covering a significantly larger number of events.

The need to run in at least near-real-time requires that the HEC method be validated and verified so as to produce valid results across a wide range of possible input cases. In usage, many event cases are run in order to generate predictions. Predictions can then be compared with Earth science measurements (for Earth-incident events) to improve our theoretical understanding of space weather events or used in advance to provide early-warning capability of hazards.

Running many cases often means that space weather analysis requires dedicated machines and service levels, first for the high levels of testing required to validate, then for use in prediction. In addition, rapid data assimilation into the HEC system is required, either through direct data handling, preprocessing, and filtering or creation of sparse datasets. This process can incur network bandwidth issues between the data archive and the HEC system.

Understanding and predicting space weather is part of NASA Strategic Goal 3, specifically Sub-goal 3B, "Understand the Sun and its effects on Earth and the solar system" and Sub-goal 3C, "Advance scientific knowledge of ... the hazards present as humans explore space." LWS mandates "large modeling activities that address coupling across traditional science domains in the Sun-Earth chain specifically be included as strategic capabilities."

Current space weather prediction in ionospheric science is well advanced, with codes such as GAIM running in real time. Ring current and radiation belt work show promise. Particle events are well measured, and statistical methods are improving the prediction windows. CME studies are a rapidly advancing field that still has large error bars for predicted arrival times, in part as a full heliosphere-interstellar medium-ionosphere multi-scale model does not yet currently exist.

Data Handling and Analysis

Heliophysics will increasingly deal with very large volumes of data, both from observations and from simulations. (We will use the term "data" for simulation output as well as for measurements from physical observatories and spacecraft.) For example, the SDO mission will generate data at a rate comparable to many Earth science missions (~1 TB/day), and the processing of these data to produce physically meaningful results will require significant computational resources. The standard processing pipeline for SDO will involve calculating the vector magnetic field near the Sun at modest time resolution, and this will be performed by a dedicated 512-processor system. In addition, however, higher time resolution and other products for specific modeling tasks will require comparable resources, although not continuously. In addition to the SDO data, large heliophysics simulation codes will often produce terabytes of output in a single run.

These large data volumes pose both storage problems and difficulties for data analysis and assimilation. Whatever the source of the data, we will want to be able to perform calculations (e.g., finding currents from magnetic fields, calculating forces) and perform data mining tasks such as identifying active regions or reconnection sites using algorithms that in many cases need to be developed. Often, it will be useful or perhaps essential to have the processing done at the site where the data reside. This will also be increasingly true for visualization tasks that will require more computational power and memory than a typical scientist will have on a local system.

Given the above considerations, we see a significant evolution in the way NASA will need to handle data. We recommend that NASA continue to support adequate large storage systems; this will at least be partially be eased by the rapid evolution in storage technology. In addition, it will be important to be able to find and easily access the data and this requires that NASA support metadata standards such as the SPASE data model in heliophysics and the Virtual Observatories based on that data model. Finally, new tools will be needed to exploit increasingly large data volumes, and support for the development of such tools is essential for efficient use of the data.

Community Model Access

The majority of the heliophysics scientific community does not engage in model development. Instead, scientists desire to use models for analysis of primarily spacecraft but also ground-based measurements, to prepare for measurement campaigns, or to provide context for the latter. These scientists form the user base of a modeling service, akin to what is provided today by the Community Coordinated Modeling Center (CCMC).

This service consists of the execution of routine, runs-on-request for the scientific user. These runs are executed in a timely manner by CCMC staff using models provided by the scientific modeling community. In order to provide a rapid return service, permit run tailoring and monitoring, and preserve the intellectual property or sensitivity of the models, these runs must be executed on dedicated computational platforms. Run results are converted and visualized by dedicated servers, and disseminated to the user via a tailored web interface. At the present time, the CCMC executes hundreds of such runs every year.

It is worth noting that these runs typically are not forefront scientific modeling calculations as described above. Instead, the vast majority of these calculations execute smaller problems tailored to a specific scientific question. Calculations of this nature, combined with post-processing and analysis, are best performed on local, dedicated compute servers. This need will continue in years to come.

Technology and User Support Needs

The heliophysics science community is following with great interest the development of new computations technologies, such as Cell processors, GPUs, and Fogs. With the exception of one scientist, these technologies have not found their way into heliophysics computing yet, primarily due to the complications associated with the programming environment. Instead, the primary computational platform employed is the cluster, based on state-of-the-art dual- or quad-core processors. These platforms are used at all levels, from smaller, desktop systems, to medium-size, special-purpose clusters, and to high-performance computational platforms. The predominant programming language is Fortran, and this is expected to remain so in the future. For these reasons, development of multi-purpose, cluster-type technologies is likely to yield the maximum near- to mid-term benefit.

However, the benefit of new computational technologies must be recognized. Better, standardized programming environments, such as MPI, need to be developed to bring these platforms to bear on the bulk of heliophysics science problems. Since the development of new computational platforms results essentially from game industry spin-off, the biggest benefit technology program would focus on the development of a multi-architecture, standardized programming environment. Ideally, such an environment would be usable on all platform sizes of interest. In addition to the existing user support, the establishment of a group charged with developing this environment is highly recommended. The integration of modern research codes into this programming environment should, initially, be conducted at a small number of testbeds, which could partially be located at distributed sites. Finally, the establishment of a model conversion test case would be beneficial.

Recommendations

This section sums up the recommendations of the Heliophysics panel.

Recommendation 1

The ability to execute and analyze forefront science calculations is a top priority, with an estimated increase of a factor of 10^3 over the next 5 years required. However, as these calculations will only involve an albeit spectacular minority of science problems, this must not be conducted at the expense of the vast multitude of computationally smaller problems, which are essential for a healthy heliophysics modeling activity.

Recommendation 2

There is a need for both large, shared parallel platforms for forefront science calculations and also for smaller, distributed units for model development, smaller science runs, data analysis, and community model access. Most modern calculations are done on clusters.

Recommendation 3

Data handling and analysis benefit from support for data standards. In addition, adequate storage systems for grand-challenge problems are important. Furthermore, the magnitude of data streams from science missions such as SDO call for the development of new analysis tools, which, among others, support data mining.

Recommendation 4

The present model for community model access has been highly successful. It should continue in order to foster the maximum science community benefit from modern space science models.

Recommendation 5

Within the areas of technology needs and benefits and user support needs, support for the adaptation of new computational technologies is highly desirable. This support can take the form of the development of Fortran-compatible, parallel libraries and establishment of testbeds, as well as leveraging and enhancing staff support to transition models to HEC systems.

Panel Membership

Scientists
Michael Hesse, NASA/GSFC, Co-Chair
Aaron Roberts, NASA/HQ, Co-Chair
Alexander Antunes, NRL
Mei-Ching Fok, NASA/GSFC
Lika Guhathakurta, NASA/HQ
George Khazanov, NASA/MSFC
Alexander Lipatov, NASA/GSFC
Peter MacNeice, NASA/GSFC
Merav Opher, George Mason Univ.

Technology Committee Support
Dan Duffy, NASA/GSFC
Hamid Oloso, NASA/GSFC/AMTI

Planetary Science and Mission Engineering Panel

Capability Drivers

Some of the most stressing model development challenges, and also high-end computing (HEC) usages, within the Science Mission Directorate (SMD) stem from the development and operations of planetary science missions. To be clear, this is not to exclude considerations of non-planetary missions. Planetary and non-planetary science missions alike feature complex instruments or subsystems—for instance, the optics of the James Webb Space Telescope, the deployment of the Mars Science Laboratory (MSL) parachute in a supersonic flow, or the control-structure interactions for the large antenna of the Soil Moisture Active Passive (SMAP) spacecraft—that strongly benefit from, or require, many of the same architectures, parallel software advances, and compute-node/memory scales, discussed in other parts of this report.

However, in recent years, some of the most stringent requirements in engineering modeling, and the necessity of HEC to provide adequate computational capability, have arisen from certain development and operations categories of planetary science missions. In particular, flight operations considerations for high-autonomy interplanetary spacecraft during critical events are known to benefit from HEC in a real-time "go/no-go" decisional context. Such events include spacecraft approach to another planet to enter its atmosphere, autonomous rendezvous with an object or other spacecraft at another planet, or preparation to initiate maneuvers to descend to the surface of a previously unmapped object such as a comet or asteroid.

As discussed below, this situation is illustrated by recent experience in robustly engineering and performing the Entry, Descent, and Landing (EDL) of the Phoenix Mars Lander mission. The Phoenix effort was significantly enabled by use of high-fidelity EDL simulations that ran during development on a shared institutional supercomputer at JPL, then during flight operations on several-hundred-compute-node, mission-dedicated supercomputing clusters. This computational approach is being followed, and expanded upon, by MSL in its final years of development before launch. And intensive computation is foreseen for some of the most demanding future missions specified in the National Research Council (NRC) Decadal Survey, *New Frontiers in the Solar System*[1], and the amending report, *Opening New Frontiers in Space: Choices for the Next New Frontiers Announcement of Opportunity*[2].

Models and Applications

This section provides descriptions of the application of HEC to both the Phoenix mission, which landed on Mars on May 25, 2008, and the MSL mission, which is in post-Critical Design Review development at the time of writing. Storage, archival, and distribution of data for these missions are described in the next section.

<u>Phoenix Mars Lander</u>

The Phoenix Mars Lander spacecraft was launched on August 4, 2007, and entered the Martian atmosphere and safely landed in the northern arctic reaches on May 25, 2008. The Phoenix engineering team began making heavy use of HEC resources at JPL when concerns arose over the safe performance of landing radar during Phoenix terminal descent. These concerns were due to discovery in the winter of 2006, via simulation and tests, of a subtle high-probability failure mechanism in the radar's detailed ground search-and-acquisition firmware design[3]. Engineers made intensive use of a high-fidelity simulation of the Phoenix EDL dynamics, incorporating a very-high-fidelity physics-and-firmware-modes model of the Phoenix radar, in assessing options to fix the problem. The simulation was hosted on JPL's 1,000-processor "Cosmos" Linux supercomputer, which initially was a resource that the Phoenix team shared with the rest of the JPL scientific and engineering HEC customer community.

It was determined at an institutional level that JPL Cosmos supercomputer usage priority went first to Phoenix runs to evaluate a radar firmware evaluation and close out the validation of the modification before the spacecraft went into final assembly (including incorporation of the flight radar unit) in June 2007. The firmware modification devised by the Phoenix team was shown—via high-fidelity Monte Carlo simulations of the Phoenix EDL, executed on Cosmos—to greatly reduce the risk of catastrophic EDL failure to an acceptably low risk level. The vendor modified the flight radar with the firmware change in time for its integration and testing at NASA Kennedy Space Center in June 2007.

In the course of engineering the radar firmware fix and determining the EDL performance margins and risk of failure, the Phoenix team used more than 20 node-years of processing time on Cosmos and generated data requiring more than 2 TB of storage. The simulation team also established an ability to turn around the solution of a 2,000-case Monte Carlo simulation of the Phoenix EDL using the very-high-fidelity radar model in less than 24 hours, including comprehensive statistical and plotting post-processing of data for systems engineering assessments. Importantly, these FY 2007 experiences of the Phoenix EDL team in leveraging HEC resources for high-fidelity simulation and performance analysis of the EDL system also led to envisioning, then establishing, a plan for FY 2008 use of supercomputing and simulations during Phoenix post-launch flight operations.

During final Mars approach, multiple 2,000-case Monte Carlo simulations of the EDL trajectory were re-generated almost daily (with new inputs based on updated spacecraft tracking trajectories from the Navigation team). These simulations were used to compute critical performance margins of EDL subsystems with respect to their hard-failure limits and evaluate the need to update any parameters in the onboard EDL software. In this timeframe, HEC resources were also used to perform daily calculations to assess landing site safety (LSS), based on the latest navigation results and high-resolution images (from the HiRise camera of the Mars Reconnaissance Orbiter [MRO]) of the terrain in the landing target region.

	Phoenix EDL	Phoenix LSS	MSL EDL	MSL LSS
Number of Particles	2,000 trajectories	100 x 10,000 ellipses	8,001 trajectories	100 x 100,000 ellipses
Resolution	0.1 sec.	100 m/pixel	10/200/1,000 Hz	1 m/pixel
Runtime for a Single Particle	6 hours	< 1 sec.	10 min. @ 200 Hz	Minutes
CPUs	1/trajectory	1/10,000 ellipses	1/trajectory	1/100,000 ellipses
Storage per Run	75 GB	5 MB	800 GB	~50 MB
Issues and Bottlenecks	(1) High-fidelity radar model; (2) post-processing speed; (3) preventing and/ or catching input errors	Availability of institutional MATLAB licenses during operations	Same as (1) and (2) from Phoenix EDL	Real-time processing of maps and MATLAB license availability issues during operations
Other Requirements	Need for a backup supercomputer during mission-critical flight operations	Need to produce landing site safety assessment in operations timeframe (easily met with < 50 processors)	Rapid-solution turnaround during operations; backup supercomputer	Need to turn around operations in minutes; may require 60–100 processors
Future Improvements	200 Hz dynamics instead of 10 Hz; faster runtime; visualizing differences between Monte Carlo sets	Ability to process higher-resolution surface image maps	N/A	N/A

Table 1: HEC Metrics from Mars Phoenix Lander During Operations and as Projected for Mars Science Laboratory (MSL). The table compares HEC usage metrics for the Phoenix mission EDL and LSS computations during flight operations (August 2007–May 2008), along with approximate projections for the MSL EDL and LSS equivalent activities. The MSL assumptions are discussed in the next sub-section, titled "Mars Science Laboratory and Other Future Missions."

For the period of Phoenix flight operations (which went from August 4, 2007 launch until May 25, 2008 landing at Mars), JPL purchased and operated two new 500-processor Linux-cluster supercomputers, which were dedicated entirely for Phoenix usage. The two machines—named Galaxy and Nebula—were identical, with uninterruptible power supplies, 7 TB of storage capacity apiece, and RAID 5 backup capabilities. The reason for providing two such machines was to satisfy a Phoenix flight operations requirement that during the final month of spacecraft approach to

Mars—when daily 2000-EDL-trajectory Monte Carlo simulations would be run—there would be a primary supercomputer and a backup supercomputer (located physically off the premises of JPL) that could still operate during mission-critical events, even in the contingency of a local power grid failure at JPL.

The Phoenix supercomputers were used throughout the interplanetary cruise phase to handle a variety of computationally intensive EDL simulation tasks:

- Tuning of numerous key onboard software parameters affecting EDL and terminal descent; performing robustness analyses to accurately confirm expected safety margins during EDL;
- Supporting multiple-day, operational readiness tests that exercised the EDL engineering team in operations protocols and procedures in real-time; and
- Developing a "survivability book" of performance sensitivities to off-nominal conditions (e.g., extreme high or low atmospheric density, out-of-spec entry flight path angle) that could be used to send "survival" parameters to the spacecraft, a few hours pre-entry, in an extreme late-breaking contingency scenario.

Mars Science Laboratory and Other Future Missions

The MSL HEC requirements given in Table 1 for EDL engineering are based on existing simulations of that spacecraft's Mars entry, which are being used during MSL development as of Summer 2008. The MSL spacecraft EDL approach differs radically from the Phoenix EDL design in that the MSL spacecraft maneuvers to precisely guide its trajectory during hypersonic entry (Phoenix entry had no closed-loop guidance) and its "sky crane" deployment system lowers the rover to the ground (Phoenix descended to the surface with its landing rockets). Compared with Phoenix, the MSL simulation must be executed at a high frequency (200 or 1,000 Hz) to capture detailed dynamics and performance of the EDL system. The desire to capture time history data for diagnosing problems and evaluating system and flight software performance increases the per-run storage requirements significantly. Moreover, given the complexity of the system dynamics to achieve its sampling error requirements, MSL requires 8,001 trajectory samples per Monte Carlo run (vs. Phoenix's requirement for 2,000 trajectory samples per Monte Carlo). These factors combine to increase the requirement for number of nodes (for a given throughput) and also storage capacity (a large fraction of 1 TB per Monte Carlo run), compared with Phoenix.

Future Mars missions such as MSL, or other pinpoint landing missions such as Mars Sample Return (MSR) or an Astrobiology Rover, for example, will require hazard maps with very high resolution for LSS analysis. Given the demonstrated performance of MRO HiRise camera imagery, these maps should be anticipated to be on the order of 1 m/pixel compared to the 100 m/pixel maps used on Phoenix. Processing these maps in a real-time environment like operations could prove to be very difficult without HEC resources. While new techniques could help to improve the processing time, the amount of analysis that can be done in real-time will be limited without the use of parallel processing. Even with multi-processor supercomputing available, the current MATLAB-based software could also be limited by licensing issues. In order to keep the analysis at least at the level of what was performed for Phoenix, high-end resources will be a necessity, and the resources to use those processors must be available as well.

Storage Requirements

Table 1 in the section above captures some of the more stressing requirements for data storage for mission engineering of planetary science missions, i.e., multi-terabyte storage is required to support the engineering campaigns for EDL-intensive missions such as Phoenix and MSL.

Further, more detailed requirements for mission engineering storage of large datasets include:

- A large-capacity (several TB) data storage server, with controlled access limited to team members and system administrators, to archive and share the team's large datasets;
- Ability for a geographically diverse, multi-organizational engineering team to access data on the storage server. This entails maintaining institutional security firewalls, but enabling key team members from different organizations (e.g., in the case of Phoenix EDL, the team came from NASA Langley Research Center, Lockheed Martin in Denver, and JPL) to have access to the necessary data in a timely manner; and
- Frequent, automated backup capability for the data on the storage server.

Programming and Analysis Environment

Third-party software sets that act as a powerful coding-support infrastructure on supercomputers are in common use across the engineering community, including planetary science mission engineering. The availability of these third-party software sets enables the fullest possible range of porting options for existing or expanding engineering codes, which are implemented using such languages as Fortran, C, C++, Java, Python, Perl, MATLAB, and Ruby, and often operated using Unix/Linux C-shell scripts. They also provide the necessary languages, compilers, and dependency libraries (e.g., Scientific Python is used increasingly in Python-language scripting, which depends on the GNU Scientific Library for precision linear algebraic and differential-equation solving, which in turn depends on the latest update of the gfortran compiler). This is likely to continue to be true for the foreseeable future, at least the next 5 years, in the mission engineering context of leveraging HEC resources.

While sophistication in integrated development environments and configuration management tools is likely to continue to increase and be available at a range of prices for mission engineering modeling and coding activities, the common underpinning of software sets used in much of mission engineering is the set of languages and operating systems listed above. Experience in the application of HEC to mission engineering has shown that the presence of a well-tested, integrated set of commonly available (but not always *integrated-out-of-the-box*) third-party software on HEC platforms can significantly facilitate the transition of engineering codes from workstation development areas—or even archival storage—into operational use on supercomputers. Experience has also shown the converse to be true.

As a HEC facility upgrades operating system versions, maintenance of key, common coding-support infrastructure—including regular and robust verification of unbroken dependency chains in it—is an essential task that the facility should provide to its engineering and scientific coding customer community. Also, continuation of powerful, modest-maintenance-cost Unix/Linux/Unix-derivative operating systems is an important architectural consideration for scientific and engineering computing in planetary science mission engineering.

Evolution of Modeling Activities

The future of deep space science missions over the next 5 years consists of options being evaluated over a variety of mission classes. In particular, the planetary science missions that are responsive to the NRC's Decadal Survey, which have not yet been designated as missions, and are called out in the NRC's *New Frontiers* documents, are:

- Lunar South Pole/Aitken Basin Sample Return
- Venus In-Situ Explorer
- Comet Surface Sample Return
- Asteroid Surface Sample Return
- Mars Sample Return

Note that each of these missions entails EDL, either at Earth or at another planetary body (Mars or Venus). Therefore, like the Mars Exploration Rovers, Phoenix, and MSL, they will entail intensive computation related to EDL systems, e.g., CFD modeling of hypersonic flows and parachute modeling, Monte Carlo analyses of EDL system performance. All missions will likely benefit from leveraging HEC resources for mission engineering, to at least the extent of Phoenix and MSL.

Moreover, the panel is of the opinion that the MSR mission and the "Small Body" (Comet or Asteroid) Surface Sample Return missions present particularly stressing uses for HEC—likely with greater throughput and memory requirements than even MSL— as the missions are developed and flown. The rationale for making this statement is as follows:

- *MSR* involves 1) two EDLs, one for the sample collection spacecraft at Mars, and another for the sample return spacecraft at Earth; 2) an autonomous first-time flight of a Mars Ascent Vehicle with the Mars sample stowed aboard; and 3) per presently-tendered architectures, autonomous rendezvous of a sample-return spacecraft with the sample vessel—including mechanical capture of the sample—in orbit above Mars, prior to Earth return with the sample. The existence of four distinct autonomous critical events (versus one EDL for MSL) indicates a strong potential for HEC requirements for MSR, greater than even MSL's, in order to provide thorough model-based analyses and tests of the autonomous subsystems prior to the events.

- Similarly, a *Small Body Surface Sample Return (SBSSR)* mission will consist of not less than two of three possible autonomous mission-critical events: 1) at least one autonomously-triggered or timer-based (but with substantial delays with controllers on Earth) descent to the small body surface to acquire a sample; 2) an ascent from the surface into some safe orbit of the small body; and 3) EDL at Earth of the capsule containing the sample or samples. Sensors on board an SBSSR spacecraft that are used in the descent are likely to be some combination of optical (cameras, lasers) and radar instruments, and the engineering of the mission will likely entail Monte Carlo analyses of the descent—over a wide variety of surface types to account for the unknown comet or asteroid surface slope, roughness, rockiness, cratering, and albedos—to assess the detailed functioning of the spacecraft onboard guidance, navigation, and control software based on these sensors. The combination of analyzing the surface-sample acquisition campaign (with complex sensors, over an unknown surface) and also the usual rigors of analyzing an (Earth-return) EDL indicates at least MSL-level requirements for HEC, if not greater.

Hybrid or "Exotic" Processing Architectures

For both the MSR and a SBSSR missions, there are aspects of both development design and operations decision-making that could potentially benefit from an advanced ability to rapidly perform simulations of closed-loop control based upon high-fidelity simulated imaging data, "flash LIDAR" imaging for hazard avoidance during descent, or terminal autonomous rendezvous. The future mission HEC usage scenarios presented here are not different in principle from the EDL development or operations of Phoenix or MSL, but high-performance optical or LIDAR sensors will add complexity to design and operations decisions. To mitigate risk stemming from sensor complexity in these future mission scenarios, the HEC goal is to leverage the maximum-fidelity sensor models available (instead of using simplified parametric noise models) to permit considerably more nuanced understanding of physical effects that can affect performance and risk, and to determine how these effects are changing day-to-day, or even hour-to-hour, during a mission's critical events—such as MSR rendezvous and capture of the orbiting sample container for Earth return and SBSSR preparations to permit the spacecraft to descend autonomously to the asteroid or comet surface to retrieve a surface sample.

Thus, from the standpoint of mission engineering for planetary science missions, further research and experimentation using "hybrid" architectures that accelerate processing speed at each individual compute node is encouraged, as these could potentially benefit certain classes of calculations done during design development cycles.

However, if the ability to accelerate a broad range of calculations (e.g., not solely ray-tracing or image rendering) by an order of magnitude or more is not afforded by these hybrid-processor architectures versus standard CPU architectures, the benefits of exotic technologies might not pay off for mission engineering development, at least in the context of the development or operations of guidance, navigation and control of missions such as MSL, MSR, or SBSSR. Concerns of porting code to these architectures could outweigh implementing engineering solutions on hybrid processors if these hybrids do not significantly outperform HEC systems using standard processors—increased cost and development schedule impacts, as well as the risk of limiting future portability of sophisticated, but hybrid-processor-dedicated codes, being foremost among the concerns. These performance and cost-schedule-risk concerns should be addressed by those responsible for developing or evaluating hybrid-architecture processors for prospective future mission engineering usage.

Recommendations

Recommendation 1

From the Planetary Science and Mission Engineering perspective, the highest-priority HEC need for the next 5 years is a sustained and increased mission-dedicated access to moderate-to-large (hundreds to thousands of compute nodes) parallel-processing architectures, with multiple-terabyte networked storage capacities and state-of-the-art system administration practices. Moreover, as demonstrated on Phoenix, having a dedicated parallel machine during flight operations, of even a modest (500-node) scale, enables robust engineering decision-making during mission-critical events (such as EDL at Mars) entailing spacecraft autonomous operation at great distances from Earth. Therefore, our panel recommends that NASA SMD create and maintain a state-of-the-art parallel processing capability for mission-critical flight operations of such missions.

Recommendation 2

The panel also recommends that NASA evaluate how to best provide a backup parallel processing capability for mission-critical flight operations that use parallel processing, including multi-center backup approaches. One example is using the NAS Facility "Columbia" system to back up JPL flight operations parallel processing facilities in the event of an emergency outage of the primary (JPL) machines during a mission-critical analysis period.

Recommendation 3

Finally, while there is no (currently apparent) requirement for hybrid or Cell-based architectures for MSL, MSR, or SBSSR missions of the future, if such architectures could supply order-of-magnitude or greater acceleration of complex calculations (e.g., high-speed ray-tracing calculations for radar or optical sensor models in a closed-loop EDL or small-body surface-descent simulation), they could enable greater data throughput to inform systems engineering design or operations decisions.

Panel Membership

Mike Lisano, NASA/JPL, Co-Chair
Eli D. Skulsky, NASA/JPL, Co-Chair *in absentia*
Gene Bonfiglio, NASA/JPL
Dan Burkhart, NASA/JPL
Mike Rilee, TopQuadrant
Larry Roelofs, GST

References

1) *New Frontiers in the Solar System: An Integrated Exploration Strategy*. The National Academies Press, Washington, DC, 2003.
2) *Opening New Frontiers in Space: Choices for the Next New Frontiers Announcement of Opportunity*. National Research Council, Prepublication Copy, April 2008.
3) Skulsky, E.D., Shaffer, S., Bailey, E., Chen, C., Cichy, B., and Shafter, D. *Under the Radar: Trials and Tribulations of Landing on Mars*. Paper AAS 08-033, presented at 31st AAS Guidance and Control Conference, Breckenridge, CO, Feb. 2008.

Technology Panel

Introduction

The purpose of the technology panel was to extract computing technical requirements from the discipline panel reports and develop a crosscutting set of recommendations reflecting the computational requirements that are common across the Science Mission Directorate (SMD) disciplines. This panel did not hold a separate technology session—panelists were embedded in the discipline panels so that they were fully aware of the discussions and could provide technical expertise as needed. Technology panelists were all high-end computing (HEC) experts drawn from the NASA centers most involved with SMD HEC applications and the provisioning of HEC resources.

After the discipline panel reports were completed, the technology panelists reviewed the discussions and recommendations to capture common technology requirements across the SMD disciplines. This section summarizes those requirements that arose in all or most of the discipline panels. It is filtered through their expert knowledge of current HEC facilities and practice, and the likely evolution of HEC technology over the next 5 years. In some cases, particular requirements identified by one discipline were found to be subsumed by a larger or more broadly stated requirement from another discipline. This section extracts only major findings and recommendations from the discipline discussions that would drive SMD HEC investments over the next 5 years. This section also notes any discipline-unique requirements that arose from the panel reports and discussions.

Computing Platform Architectures

This section addresses the characteristics of computing platforms that will be required to meet the computational capabilities and throughputs identified in the gap analysis by the various science panels. In addition to raw floating-point operations per second (flops), this section also addresses related computational platform characteristics such as amount of memory, scalability, and degree of centralization implied by the requirements.

Architectures

Although the panels were somewhat varied as to the precise balance of centralized versus local computing platforms, a general consensus exists that mainstream commodity cluster technology will be sufficient for most, if not all, requirements through 2013. It is generally understood that continued Moore's Law progress in performance will be in the form of increasing number of computational cores within a processor (socket). This exponential trend in the number of cores arises at the end of an era of exponentially increasing processor clock speeds. These unavoidable technology trends in parallelism and clock speed imply that an ever-increasing level of parallelism within SMD applications is required to exploit future computing platforms. Fortunately, such parallelism is largely made possible by corresponding increases in spatial resolution.

Although most communities recognize the potential benefits from any of several accelerator technologies that are beginning to target HEC, the lack of portability and the daunting programming challenges were cited as justification for not making major investments. Example accelerator technologies include the IBM Cell Broadband Engine, various Graphics Processing Units (GPUs) such as NVIDIA, and Field Programmable Gate Array (FPGA) processors. Specialization enables these technologies to greatly exceed commodity processors in terms of peak performance, memory bandwidth, and power requirements. Further, the gap compared to commodity technology is expected to significantly widen over the next 5 years. Programming these processors generally requires specialized data and process management and may require major, non-portable modifications of core algorithms. Most vendors of this technology provide no Fortran support at this time, though there is some possibility of better support in the near future.

Memory

Perhaps the largest gap between industry trends and the science panel reports is that of memory, or more specifically memory-per-core. Although the memory-per-node within clusters continues to rise exponentially in a Moore's Law manner, the onset of multi-core technology has flattened this growth when measured per core. To some degree, the panels were inconsistent with their expectation when one compares aggregate memory requirements, expected scalability, and per-core requirements. The cost of memory is expected to continue to be a major portion of the cost for computing platforms and could easily come to dominate the cost. If a balanced degree of parallel scalability cannot be obtained in major applications, NASA's aggregate computational capabilities may be severely constrained as ever-larger fractions of resources are invested in memory subsystems.

Memory bandwidth continues to be a major performance bottleneck for many applications and is expected to gradually worsen over time. Computational capabilities of processors generally continue to improve more rapidly than bandwidth to main memory. In some algorithms cache can be used to reduce the impact of this discrepancy. One positive aspect of the flattening of clock speeds is that the impact of flat latencies on main memory is now stable.

Communication Fabric

Scalability of many applications is limited by either the bandwidth or the latency of the communication network between nodes of a cluster. Due to various limitations, many of the competing technologies in this arena have now converged. In particular, latencies less than ~1 microsecond are generally not obtainable between nodes, which leads to some performance limitations in various applications. No major requirements for communications were expressed by the various science panels, which probably indicate an understanding of these limitations.

Capability and Capacity Computing

HEC often distinguishes two broad categories of computation for convenient analysis. "Capability" computing is computing that requires extreme values for one or more platform characteristics, and typically is driven by applications with large computational workload and/or memory for a single instance. Within NASA, capability computing is usually driven by requirements for high-spatial resolution within a given model. "Capacity" computing is related to requirements that are driven by aggregation of requirements from large numbers of smaller instances (perhaps involving various unrelated applications). A common driver for capacity computing within NASA is that of model ensembles. No sharp line exists between capability and capacity computing and NASA continues to have a spectrum of requirements that span the two categories.

Demand for capacity computing can be met by a variety of mechanisms ranging from large centralized systems to federated clusters to "cloud" computing where elements are highly distributed. Although cloud computing is generally the least expensive, some SMD algorithms cannot tolerate the relatively high latencies associated with such approaches. The Solid Earth panel endorses cloud computing for a significant portion of their compute requirements, whereas other panels are relatively silent on distinctions among capacity requirements. A geographically distributed set of mid-range clusters should be adequate for all of the identified capacity requirements.

The need for capacity computing within SMD for the next 5 years is dominated by Earth system models, while capacity requirements among the Earth System, Heliophysics, and Astrophysics panels are all in the neighborhood of 1 petaflops. The table below summarizes some of the computing requirements from these panels. Missing entries reflect the absence of requirements that directly address the given parameter. Note that teraflops (TF) here refers to the peak capability of the *platform* that is *inferred* from existing application performance data.

Requirements — 2013					
Discipline	**Capacity (CPU-hours)**	**Capability (TF)**	**Memory (TB)**	**Memory per core**	**Total Processors**
Earth System	883M	300	100	2 GB	30,000
Heliophysics		1,000			100,000
Astrophysics	100M	10	200	2 GB	100,000
Solid Earth	30M	40	40	4 GB	4,000
Mission Eng.					100's
Combined	**> 1,000M**	**1,000**	**200**	**4 GB**	**100,000**

Table 1.1: Requirements — 2013 for Each SMD Discipline and Combined

Findings and Recommendations

Finding 1.1: Hardware accelerators may prove to be a disruptive technology for NASA scientific modeling but are currently too difficult to use for most mainstream efforts.

Recommendation 1.1: NASA should maintain a low-level effort to investigate the potential benefits and costs of introducing accelerator technology within major SMD applications. As the advantages and ease-of-use improve, NASA should be ready for an intense effort to adopt this technology.

Finding 1.2: Many models are ill prepared to exploit the shift in the implication of Moore's Law from faster clocks to higher parallelism (multi-core).

Recommendation 1.2: NASA should make staffing investments within the computing centers to assist modelers in improving the scalability of major applications. These improvements are needed both to achieve the various capability goals and also to reduce the fraction of computational resources required for memory-per-core. Further, such improvements should also improve cache-reuse, and therefore the serial performance, within each core.

Storage and I/O

In contrast to the 2002 workshop, all panels identified storage and I/O as significant concerns. Currently, processing architectures do not scale storage and I/O linearly with processor capacity. It is not uncommon, in scientific applications, to find that many or most of the processors assigned to an application sit idle while a small set of processors access input files or write to the file system. There are a number of issues that create these limitations. Some obstacles are a function of the communications path between processors, storage, and end users. Others are related to the storage hardware and file systems themselves.

SMD HEC applications vary in the amount of input data required—from relatively small sets of input parameters to large volumes of ingested observational data or large restart files. They all have large output requirements, since they typically produce snapshots of the state of the model or analysis at regular intervals and, for long running models, restart files that contain the entire state of the model. The model output is rarely the end product for the science investigation. Further analysis of some or all of the output is typically required. This analysis (including visualization) may be done local to the model output (i.e., at the HEC center) or the output may be moved to an end user's computer for additional processing. The useful lifetime of the output data also varies. In some cases (such as ESM reanalysis), the output products may need to be archived for several years and be accessible to a community beyond the investigator that ran the model. In many cases, however, the output has a useful lifetime of a month or two and can be discarded afterwards. The I/O and storage capabilities must satisfy all of these usage scenarios.

The storage and I/O capabilities must also support fast access to the application output. The data must be accessible not only by the supercomputer, but also by other processing systems for data analysis and visualization. In addition, the creation of appropriate metadata and its management is necessary to assist the end users in identifying those particular files need for further processing (both locally and remotely). Standards groups have arisen to address these issues and NASA is participating to reflect the needs of the SMD scientific and engineering communities.

Technology Description

Historically, tape systems were the primary archive. When data was needed, it was staged onto the disk so that codes could access it. The decrease in cost of high-performance disk systems with large capacities has done much to speed the process of scientific analysis in SMD mission areas. Newly emerging solid-state storage systems, such as flash drives, have created opportunities at the desktop-scale machine. Over the course of the next 5 years, these devices may grow into competitors with the current 1 TB disk drives, creating an opportunity for significant improvements in storage by integrating solid-state, disk, and tape systems into a single hierarchical storage system.

Today, the central focus of storage and I/O technologies are the files themselves and making them appear to the codes when they are needed, not after a lengthy setup process in staging them. The availability of data directly online in contrast to requiring it to be staged for even the smallest interaction has been a major step forward since the 2002 workshop.

File system technologies today perform two vital functions that older systems did not. Parallel file systems access permits the I/O operations to occur must faster for larger files since they are written simultaneously to several controllers and disks. Advances in shared file systems have emerged to permit fast and reliable access by multiple computers to the same sets of data without making multiple copies. Most parallel file systems perform best in a specific hardware environment, usually tied to a specific vendor, so mixing processor and storage vendors in a given file system is challenging and often has mixed success in reliability. The predominant shared file systems all handle parallel access to multiple disk systems and include Lustre, GPFS, and StorNext. These parallel file systems need substantial improvements in the reliability of their systems. Several generations of redesign have yielded improvements in robustness, but acceptable reliability of these systems remains a challenge to long-term data stewardship. Again, potential benefits from the emergence of solid-state storage devices of sufficient size and performance may be realized in their integration into the existing storage management systems and applying a hierarchical model to act as large-scale caches with disk providing the depth of capacity needed.

Analysis of the Requirements

In the SMD science and engineering community, there are some common storage system and I/O characteristics that are generally needed across the board for modeling and simulation capabilities to move ahead. These are described in detail in each community's section of this report, but are summarized in Table 2.1 for the 2013 timeframe.

Characteristic	ESM	Solid Earth	Astro	Helio	Mission Engineering
Throughput	200 TB/day	100 GB/hr			
Increase Archive (PB/year)	37	3 (80[1])	2	>0.5	
Online Disk per App (PB)			1–10	6	.001
Online Disk (PB)	3–5				
WAN Throughput to the End User (MB/s)	40	20			
Improve Data Management, Archive Management Tools		Yes		Yes	
Data Standards	Yes			Yes	
Interactive Data Analysis		Yes	Yes	Yes	

[1] Geodynamo projected requirement.

Table 2.1: Panel Needs Impacting Storage and Communications

A review of the respective panel findings indicates a substantial need for improvements to capacity and performance of storage systems.

Findings and Recommendations

Finding 2.1: All projections for modeling efforts depend upon the availability of substantially more online and archival storage, both for working purposes and for comparison studies between runs and different types of models. Each discipline area projects that the size of data storage requirements will radically increase.

Recommendation 2.1: A combination of market forces and directed research is needed. Clearly, the current trend in solid state and disk system performance will continue driving down the cost of disk and increasing the size and speed of these devices; they are sufficiently widespread that investment by NASA will be largely inconsequential. Partnership with the major system providers in the form of acquisition will help ride those market forces to yield the best possible results.

Finding 2.2: Access to model output is currently slow and cumbersome for large datasets. This is a result of the relative storage capacity vs. cost of the installed tape archive storage and high-speed disk storage. All disciplines report that the available disk storage capacity at the computing centers is too small to allow them to keep needed files on fast media for short-term analysis.

Recommendation 2.2: Local communications within the range of Fibre Channel or InfiniBand will also improve the local access to the data held within the data center's archive; as users realize the inherent weaknesses of TCP/IP for high volume data access, they will relocate close enough to the storage facility to permit installation of this type of connection-oriented path to their workstations. Two major shortcomings do not seem to have sufficient market forcing behind them to predict a suitable outcome: wide-area network (WAN) communications and solidification of metadata model standards for the various SMD communities of interest. Both of these areas could benefit from investment and focus within NASA.

Finding 2.3: Remote access to large data files via long-haul networks is slow and unreliable. WAN throughput to the end user on the order of 400 Mbps (40 MB/s) is required to effectively move model output from the computing center to the end user site, but is not typically available today. Even at these speeds, the volume of model output will continue to preclude the wholesale movement of the output from the computing center to the end user site.

Scientists need more rapid, direct access to the data for applying their visualization and analysis tools, which requires the data to be available on disk instead of on tape. Where the HEC centers provide user services personnel to work with scientists, scientists report an increase in productivity and the need for far more support of this type.

Recommendation 2.3: Wide-area communications improvements in both bandwidth and latency are essential to permit remote researchers to access their own data at the data center and to use it with an effectiveness that will permit the rapid maturation of the scientific and engineering models and their application to SMD missions.

Finding 2.4: The need for metadata and data file interoperability standards is growing as more data becomes available and it is retained for longer periods of time and re-analyzed by its originator and new scientists. Describing the nature, format, and pedigree of the data becomes essential for any data to be useful over longer lifetimes.

Recommendation 2.4: A concerted effort to develop and obtain the consensus on metadata models is also a worthwhile effort in which SMD's research program could yield significant gains.

User Interfaces and Environments

Technology Description

The user interface for HEC resources is typically a command line or text window, characterized by typing individual commands (with arcane arguments) and waiting for/reacting to any error conditions. With such an information-poor interface, users may struggle for a long time to understand and effectively use available resources and services. Thus, even though HEC is a universally potent tool for the advancement of science, the standard command-line interface may make the barrier to initial use too high for many new users and may drastically reduce the productivity of existing users.

This section summarizes the interest or requirements of SMD users for higher-level (e.g., more intuitive, intelligent, automated, and integrated) interfaces to HEC resources and services. Such a portal would seek to dramatically reduce user effort and interaction requirements in completing the HEC-related tasks in their scientific workflow, such as that depicted in Figure 3.1. For example, the user environment might provide a graphical interface to domain-, application-, and user-specific information, tools, models, workflows, and results; and major portions of the workflow might be initiated and robustly completed with a single click. A notional example of such a user interface is shown in Figure 3.2 (as presented at the SMD workshop).

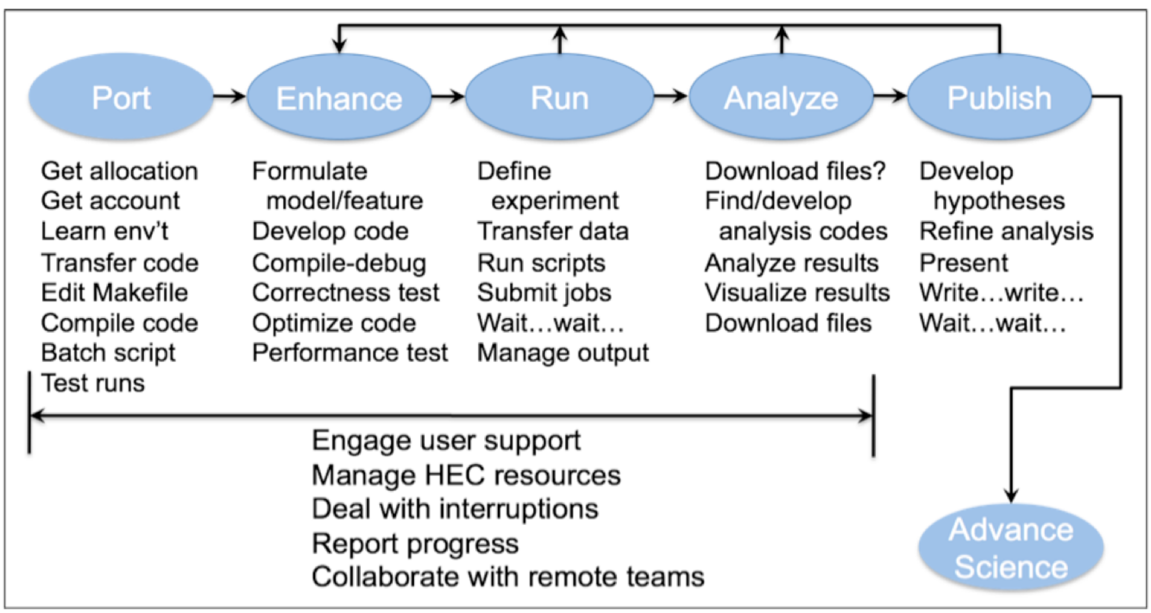

Figure 3.1: Typical (simplified) SMD HEC user's science workflow. Scientific tasks (e.g., formulating models, defining experiments, developing hypotheses) are often drastically slowed by the many low-level, user-intensive tasks required to use high-end computers. A higher-level user interface could dramatically accelerate the scientific workflow.

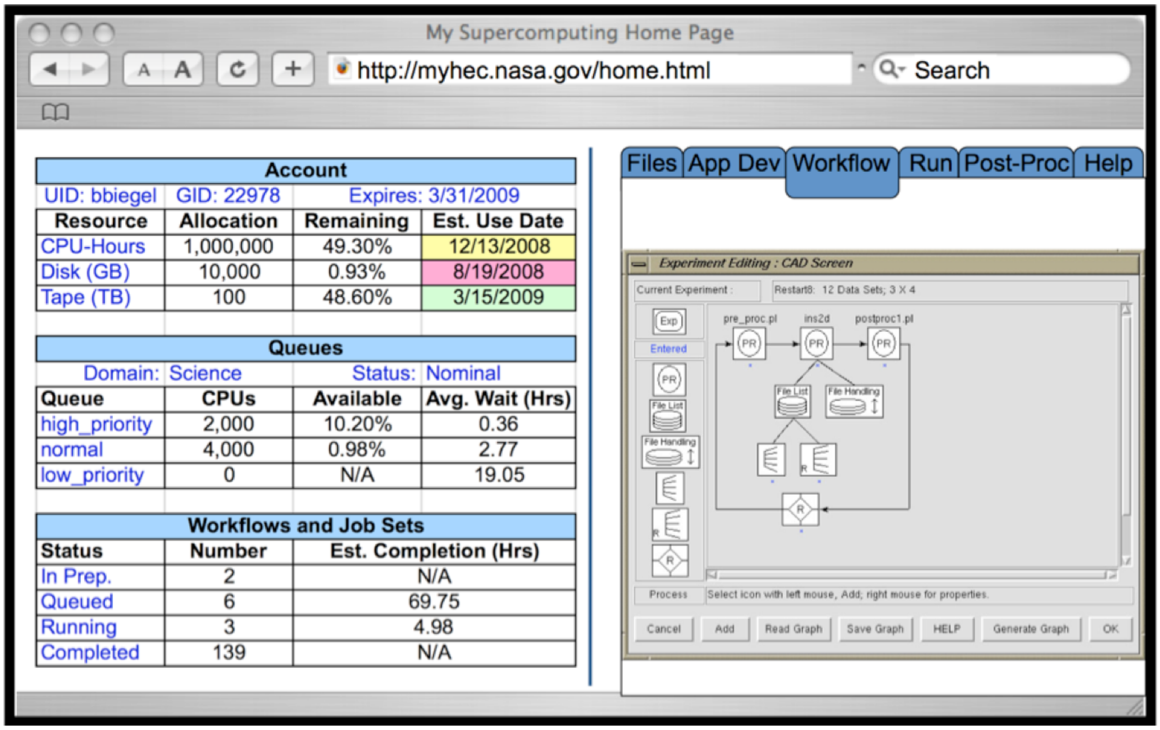

Figure 3.2: Notional example of an information-rich, intuitive user interface. It provides information, tools, workflows, and results of interest to the user.

The few existing high-level interfaces to HEC resources and services generally assume that computational jobs are small or embarrassingly parallel (easy to distribute across poorly connected processors), which is exactly the opposite of the typically large and tightly coupled SMD computational jobs. Also, existing high-level interfaces typically require a web-based, service-oriented architecture (SOA) or Grid computing interface, which are not compatible with NASA IT security policies. Thus, providing high-level interfaces to SMD users would generally require NASA to create or commission them. Such an investment must be supported by substantial user demand and an analysis predicting a sufficiently high benefit to NASA. In this section, we summarize SMD user demand for high-level tools and user environments, as expressed at the SMD requirements workshop, and in the last subsection of this technology area, we provide findings and recommendations to the NASA HEC Program regarding these tools and environments.

User Requirements Summary

The primary requirement of NASA's SMD HEC users is for more CPU time. However, each SMD discipline panel expressed interest or challenges in their HEC workflow that would benefit from improvements to the capabilities and usability of their HEC user environment. The most prominent areas of interest are summarized below, including enhancements in: (1) computation management, (2) data management, (3) data analysis and visualization, and (4) distributed environments.

Computation Management: Capabilities to automate complex and tedious computational workflow tasks also generated wide interest. Users requested support for modeling frameworks (e.g., ESMF), including work to improve and maintain their performance. Other desired tools include graphical model configuration, job submission and monitoring, and tracking and annotating model experiments. All SMD panels needed to do ensemble computations in order to achieve their scientific goals, whether through parameter, sensitivity, optimization, time-series, Monte Carlo, grid resolution, or statistical studies. These studies may involve tens, hundreds, thousands, or even millions of computations and associated data analyses. Other tools requested would enable systematic comparison of simulation and observational data. Users also indicated a need for tools that shepherd long computations through checkpoint-restart cycles, since queue time limits may never be long enough for many applications to complete in a single run.

Data Management: Improvements in HEC data management capabilities were widely requested. The Earth Science panel described a need to archive and access some data for 5 years or more, with the ability to selectively supersede data. Because today's simulation and observation datasets are so large, data provenance needs to be maintained, and complex analyses need to be done (and re-done), users feel that they need a massive data archive that is tightly coupled to a HEC facility and yet highly accessible to users. Specific capabilities requested included high-throughput, sophisticated data search, subsetting, ingest, analyses, reduction, and rendering at the HEC facility before transferring the result to the user. Also requested was the implementation of domain-specific data models and metadata standards, as well as Internet data sharing standards (e.g., OPeNDAP), to support collaboration and multi-site modeling and analysis.

Data Analysis and Visualization: Users from most panels requested more options for analysis and visualization of their HEC modeling and simulation results and associated observational data. Both Earth Science and Astrophysics are interested in concurrent visualization (CV), which involves analysis and rendering the results of each simulation time step as it is computed, rather than storing and post-processing typically only every hundredth or thousandth time step. Several users from these communities have already benefited from CV, and their enthusiasm seems to be spreading. Users also cited the very large, high-dimensional, and increasingly multi-site simulation and observational datasets, and they seek new data mining, analysis, and visualization services, tools, and techniques to help derive full insight from this data. The Solid Earth panel further requested development of web-based visualization services and tools.

Distributed Environments: The Solid Earth panel advocated for increased support of distributed environments (a.k.a. portals) and web services on NASA supercomputers, to enhance scientific productivity. QuakeSim is a good example of an application-specific portal that includes many of the workflow tasks in Figure 3.1. At present, the IT security regimen of NASA supercomputers seems to prevent allowing access by such Grid portals. If this were not the case, possibly many existing and emerging Grid/cloud computing and web services capabilities could be drawn upon to improve the productivity of NASA's HEC user environment.

Findings and Recommendations

SMD user requests and suggestions for improving their interface to NASA's HEC resources and services are many and diverse. However, a few prominent opportunities emerged, as summarized in the following findings and recommendations.

Finding 3.1: The productivity of SMD's HEC users would benefit substantially from a more intuitive, integrated, and intelligent (e.g., learning, reliable, able to express and automate complex workflows) graphical user interface to NASA's HEC resources.

Recommendation 3.1: NASA's HEC Program should explore options and conduct tests to provide HEC users with intuitive, integrated, and intelligent graphical interfaces to its HEC resources. The primary goal should be to reduce user effort in completing the increasingly sophisticated tasks of their HEC workflows, balanced with the effort to implement and maintain these interfaces by the HEC Program.

Finding 3.2: The productivity of SMD users would increase substantially if NASA's HEC user interface facilitated submission and reliable completion of these complex computational workflows. Computation workflow tasks that could be facilitated and/or automated include model configuration, job submission and monitoring, checkpoint-restart and ensemble computation shepherding, managing input and output files, and performing data analysis and rendering.

Recommendation 3.2: The NASA HEC Program should explore options and provide tools and support services to facilitate the submission and robust completion of complex computational workflows while requiring minimal user interaction.

Finding 3.3: SMD users face various challenges in their HEC data management that could be mitigated with a higher-level, intelligent interface that transparently couples data archive and HEC resources. Features of value to SMD users include high-throughput, sophisticated data search, subsetting, ingest, analyses, reduction, and rendering at the HEC facility before transferring results to the user. Also of value to SMD users would be HEC facility hosting of selected data for community access using data sharing standards.

Recommendation 3.3: NASA's HEC Program should explore options and provide SMD users with a high-level interface that transparently couples data archive and HEC resources, enabling sophisticated operations on SMD datasets. In addition, using community data sharing standards, HEC Program archive interfaces should enable users to easily publish designated datasets for community use and use datasets published at non-NASA archives as if they were located in the NASA archive. SMD users should be supported in modifying their codes to implement associated metadata standards.

Finding 3.4: The standard model of NASA's HEC facilities for supporting data analysis—letting users download and analyze the data on their local system—is not feasible for many of SMD's HEC users. New data analysis and visualization models are needed that better leverage NASA's HEC resources and expertise, and that are community-aware. VNC (Virtual Network Computing), multi-stream concurrent visualization (CV), and web services are examples of models that have proved effective for SMD users.

Recommendation 3.4: NASA's HEC facilities should expand and advertise options, and support user requests, for on-site data analysis and rendering, with real-time interaction from and streaming of results to user desktops. These capabilities should support both continuous visualization and post-processing models.

Finding 3.5: NASA's IT security policies appear to substantially limit the types of services that NASA's HEC facilities are able to provide to SMD HEC users and their collaborators. For example, SMD users are not able to attach Grid, Cloud, or web services portals to NASA's HEC resources, or to leverage the vast array of associated technologies and capabilities. This also creates a barrier between NASA's HEC capabilities and any publicly accessible NASA or non-NASA data archive, making it difficult for SMD users to share data or participate in collaborative activities as freely as users at other HEC facilities. *Thus, NASA's SMD HEC users are at a competitive disadvantage to their peers who use other resources.*

Recommendation 3.5: The NAS HEC Program should seek technical and administrative solutions that meet NASA's IT security requirements while providing SMD users with interfaces to computational and data resources that will enable them to compete scientifically on an equal footing with their peers.

Programming Paradigms and Software Tools

Technology Description

Multi-core chips, driven by energy consumption and other factors have become ubiquitous in high-end computing systems raising significant challenges for the programmers. Even though deeper memory hierarchies and more complex inter-connection communication networks are being utilized, the increase in the computing power of a multi-core chip has not been accompanied by similar gains in the memory, cache and communication performance leading to a decrease in per-core capabilities. This imbalance is further exacerbated by the introduction of GPGPUs and other add-on accelerators that raise the issues of process and data management across hybrid subsystems.

Programming Paradigms

The focus of designers of programming models and languages has always been to provide high-level abstractions to programmers allowing them to express their applications using features as close to the users' domains as possible while providing enough information to the compilers and runtime systems to effectively exploit the underlying architectural features. This tension and trade-off between portability and maintainability on the one hand and performance on the other has not been effectively resolved.

The Message Passing Interface (MPI), introduced in 1993, has become the de-facto standard for programming distributed-memory machines in particular and parallel systems in general. MPI supports a tasking model in which each process has access to its own local data only and data sharing is achieved via explicit inter-process messages. The OpenMP standard, introduced in 1997, provided support for programming shared address space systems with program directives to exploit loop level parallelism and to control access to shared data. OpenMP has been effective for small, shared-memory systems; however, in most cases these applications do not scale on large systems. Programmers have also effectively exploited a hybrid approach, using OpenMP directives within MPI tasks, to target multiple levels of parallelism available in clusters of shared-memory nodes. Programming paradigms utilized by programmers for high-end multi-processor systems have not evolved for more than a decade. Even though both MPI and OpenMP approaches have added features since their initial introductions, these features have mostly added new capabilities, such as task parallelism in OpenMP, rather than provide support for overcoming performance challenges and issues raised by modern complex high-end systems.

Alternative parallel programming approaches have been proposed over the years but most have not made significant inroads. One major example is the Partitioned Global Address Space (PGAS) model supported by languages such as Unified Parallel C (UPC), Co-Array Fortran, and Titanium. Using this model, the programmers have to specify the distribution of data explicitly. However, the computation code is written at a high level using a global array index space, allowing the compilers to optimize the actual data movement based on the target architecture. These languages have achieved success in small pockets but have not seen widespread use mainly because they do not solve many of the challenges noted above even though they do provide higher levels of usability.

DARPA's High-Productivity Computing Systems (HPCS) project has focused on the design and development of the next-generation highly parallel hardware along with the software environments required for such systems. As part of HPCS, three new languages are under development, Chapel by Cray Inc., X10 by IBM, and Fortress by Sun. These language efforts have focused more on productivity issues, such as raising the abstraction level of expressing the computation and parallelism relying more heavily on the compilers to extract the performance. Given the fact that these languages and the support systems are still under development, the level of the performance they will ultimately deliver is not clear.

There have also been other language efforts such as CUDA, Brook, and RapidMind that have been successful in specific domains. However, they have mainly targeted specialized architectures and hence are not portable.

Software Tools

The increase in the complexity of target architectures and programming models raises the demand for tools that ease the burden on programmers tasked with producing correct, efficient codes. These same complexity issues, however, also complicate the implementation of effective tools themselves, which must track the changes to both the programming models and the target architectures. Furthermore, those changes often require development of new tool techniques to address issues newly posed by the model or architecture innovations. For example, introducing NUMA memory and OpenMP results in new types of programming problems, such as non-local memory usage, that require the help of a tuning tool.

While the debugging side of parallel tools seems reasonably well in hand with TotalView and DDT (from Allinea), on the program-tuning side the situation is much more balkanized. There are a number of tuning tools, but none has really achieved the level of success or recognition that TotalView has on the debugging side. Some are limited in scope to a single computer vendor's platforms (e.g., VTUNE). Others attempt to be more widely applicable but then suffer from a lack of depth. Also as the size of HEC systems steadily increases, scaling these tools to handle thousands of processes has become a significant challenge.

User Requirements Summary

Most SMD applications use Fortran as the programming language of choice and MPI for expressing parallelism. There are a few projects that utilize other languages such as C, particularly for infrastructure routines, but the overwhelming majority depends on Fortran. Similarly, there is a sprinkling of codes using OpenMP on shared-memory clusters, however, most projects use MPI due to its portability across a variety of HEC systems. Since both Fortran and MPI are standard approaches in the wider high-performance computing community, there is no real concern that they will not be available on future HEC systems.

The panels did not express a need for any new and esoteric programming paradigms since there is nothing on the horizon that promises a big enough increase in performance to overcome the significant effort required to recode and validate user applications. On the other hand, most panels expressed interest in exploring the potential performance boost available with emerging technologies such as add-on accelerators. However, the specialized programming environments available with such systems, leading to non-portable codes, were raised as a significant impediment in utilizing these systems.

A common theme across the panels was the lack of expertise among the scientists in improving the overall performance of their applications. Most projects recommended increased support from facilities staff in analyzing their codes in order to enhance their performance. In addition, the panels expressed interest in increased availability of performance analysis tools, including training on how to utilize these tools for NASA's HEC systems.

Findings and Recommendations

Finding 4.1: As the computational models used by NASA scientists become more complex, porting and scaling the application codes on the increasingly advanced HEC systems will raise many challenges. Software tools for analyzing and enhancing the performance of their codes will become increasingly beneficial. Domain scientists will require more support from the HEC facilities staff in utilizing these tools to effectively exploit the target architectures.

Recommendation 4.1a: NASA's HEC Program should provide increased assistance in porting and scaling of application codes on current and future HEC systems. The Program should explore establishing partnerships between optimization experts from the HEC centers and domain scientists to study issues in scaling applications on parallel systems.

Recommendation 4.1b: The HEC program should hold periodic "software clinics" focused on transferring knowledge from the experts at the HEC facilities to the scientists. These clinics should include apprising the scientists of the best practices in utilizing the machines and on code structuring techniques to extract the best performance, along with providing training on using the software debugging and analysis tools available on the systems.

Finding 4.2: Alternative architectural approaches, including hybrid systems with add-on accelerators such as FPGAs, GPGPUs, and IBM's Cell Broadband Engine, have the capability of providing a significant computational boost. However, theses systems generally have specialized programming environments that require a significant porting effort and lead to targeted non-portable codes.

Recommendation 4.2: NASA's HEC Program should establish projects to study the utility of alternative architectural approaches, including their software support environments for SMD applications. The program should explore the option of establishing path-finding technology labs that acquire and deploy emerging technologies for experimentation by joint teams of computer and domain scientists.

Panel Membership

Thomas Clune, NASA/GSFC, Co-Chair
Robert Ferraro, NASA/JPL, Co-Chair
Bryan Biegel, NASA/ARC
Christopher Catherasoo, NASA/JPL
Dan Duffy, NASA/GSFC
Michael Little, NASA/LaRC
Piyush Mehrotra, NASA/ARC
Hamid Oloso, NASA/GSFC/AMTI
L. Harper Pryor, NASA/GSFC/SAIC

Appendix A: Acronyms

ACE	Aerosol-Cloud Ecosystems
ACRs	Anomalous Cosmic Rays
AGCM	Atmospheric General Circulation Model
AGN	Active Galactic Nuclei
AIRS	Atmospheric Infrared Sounder
AJAX	Asynchronous JavaScript and XML
AMR	Adaptive Mesh Refinement
AOGCM	Atmospheric-Oceanic General Circulation Model
ARC	Ames Research Center
AR4	Fourth Assessment Report
AR5	Fifth Assessment Report
ASCENDS	Active Sensing of CO_2 Emissions Over Nights, Days, and Seasons
AU	Astronomical Units
BARD	Bay Area Regional Deformation Network
BHFP	Black Hole Finder Probe
CALIPSO	Cloud-Aerosol Lidar and Infrared Pathfinder Satellite
CCM	Chemistry Climate Model
CCMC	Community Coordinated Modeling Center
CMBP	Cosmic Microwave Background Polarization
CME	Coronal Mass Ejection
CMIP5	Coupled Model Intercomparison Project, Phase 5
CCSP	Climate Change Science Program
CF	Climate and Forecast
CFD	Computational Fluid Dynamics
CH_4	Methane
CLIVAR	Climate Variability and Predictability
CO	Carbon Monoxide
CO_2	Carbon Dioxide
COLA	Center for Ocean-Land-Atmosphere Studies
CPU	Central Processing Unit
CREW	Center for Research on Environment and Water
CSC	Computer Sciences Corporation
CTM	Chemistry Transport Model
DARPA	Defense Advanced Research Projects Agency
DAS	Data Assimilation System
DDT	Distributed Debugging Tool
DESDynI	Deformation, Ecosystem Structure, and the Dynamics of Ice
DOE	Department of Energy
ECMWF	European Centre for Medium-range Weather Forecasts
EDL	Entry, Descent, and Landing
ESA	European Space Agency
ESE	Earth Science Enterprise
ESM	Earth System Model
ESMF	Earth System Modeling Framework
FLOPS	Floating-Point Operations per Second

FPGA	Field-Programmable Gate Array
FPU	Floating-Point Unit
FY	Fiscal Year
GAIM	Global Assimilation Ionosphere Model
GB	Gigabyte
Gbps	Gigabits per second
GCE	Goddard Cumulus Ensemble
GCM	Global Circulation Model
GeoFEST	Geophysical Finite Element Simulation Tool
GeoSAR	Geographic Synthetic Aperture Radar
GFDL	Geophysical Fluid Dynamics Laboratory
GEO-CAPE	Geostationary Coastal and Air Pollution Events
GEOS	Goddard Earth Observing System
GEOS-5	Goddard Earth Observing System Model, Version 5
GEOSS	Global Earth Observation System of Systems
GES DISC	Goddard Earth Science Data and Information Center
GIS	Geographical Information Systems
GISS	Goddard Institute for Space Studies
GLAST	Gamma-ray Large Area Space Telescope (renamed the Fermi Gamma-ray Telescope)
GMI	Global Modeling Initiative
GNU	GNU's Not Unix
GPFS	General Parallel File System
GPGPU	General-Purpose computation on Graphics Processing Units
GPM	Global Precipitation Measurement
GPS	Global Positioning System
GPU	Graphics Processing Unit
GR	General Relativity
GrADS	Grid Analysis and Display System
GRACE	Gravity Recovery and Climate Experiment
GRB	Gamma-Ray Burst
GRMHD	General Relativistic Magnetohydrodynamics
GSFC	Goddard Space Flight Center
GST	Global Science and Technology, Inc.
HEC	High-End Computing
HPCS	High Productivity Computing Systems
HQ	Headquarters
HST	Hubble Space Telescope
Hz	Hertz
IAAS	Infrastructure As A Service
IBEX	Interstellar Boundary Explorer
ICESat	Ice, Cloud, and land Elevation Satellite
ICME	Interplanetary Coronal Mass Ejection
IDL	Interactive Data Language
IEEE	Institute of Electrical and Electronics Engineers
IESA	Integrated Earth System Analysis
InSAR	Interferometric Synthetic Aperture Radar
I/O	Input/Output

IOC	Intergovernmental Oceanographic Commission
IPCC	Intergovernmental Panel on Climate Change
IT	Information Technology
IXO	International X-ray Observatory
JDEM	Joint Dark Energy Mission
JPL	Jet Propulsion Laboratory
JWST	James Webb Space Telescope
K	Thousand
keV	Kilo electron Volt
LaRC	Langley Research Center
LAN	Local-Area Network
ΛCDM	Lambda Cold Dark Matter
LDCM	Landsat Data Continuity Mission
LIDAR	Light Detection and Ranging
LIS	Land Information System
LISA	Laser Interferometer Space Antenna
LSS	Landing Site Safety
LWS	Living With A Star
M	Million
MAP	Modeling, Analysis, and Prediction
MAPL	Modeling Analysis and Prediction Layer
MB	Megabyte
MB/s	Megabytes per second
Mbps	Megabits per second
MeV	Mega electron Volt
MERRA	Modern Era Retrospective-analysis for Research and Applications
MHD	Magnetohydrodynamics
MIPs	Model Intercomparison Projects
Mm	Megameter
MMS	Magnetospheric MultiScale
MODIS	Moderate Resolution Imaging Spectroradiometer
MOPITT	Measurements of Pollution in the Troposphere
MoSST	Modular, Scalable, Self-consistent, and Three-dimensional
MP	Magnetopause
MPI	Message Passing Interface
MRO	Mars Reconnaissance Orbiter
MSFC	Marshall Space Flight Center
MSL	Mars Science Laboratory
MSR	Mars Sample Return
NAS	NASA Advanced Supercomputing
NASA	National Aeronautics and Space Administration
NCCS	NASA Center for Computational Sciences
NG	Northrop Grumman
NMM	Numerical Model Metadata
NOAA	National Oceanic and Atmospheric Administration
NPOESS	National Polar-orbiting Operational Environmental Satellite System
NPP	NPOESS Preparatory Project

NRC	National Research Council
NRL	Naval Research Laboratory
NUMA	Non-Uniform Memory Access
NWP	Numerical Weather Prediction
OCO	Orbiting Carbon Observatory
OH	Hydroxyl
OPeNDAP	Open-source Project for a Network Data Access Protocol
OpenMP	Open Multi-Processing
OSSEs	Observing System Simulation Experiments
ORNL	Oak Ridge National Laboratory
OSTM	Ocean Surface Topography Mission
PAAS	Platform As A Service
PB	Petabyte
PBO	Plate Boundary Observatory
PGAS	Partitioned Global Address Space
PI	Pattern Informatics
R	Radius
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RBSP	Radiation Belt Storm Probe
RDAHMM	Regularized Deterministic Annealing (Expectation-Maximization) (Algorithm for Fitting) Hidden Markov Models
REST	Representational State Transfer
RI	Relative Intensity
RIPI	Relative Intensity Pattern Informatics
RMHD	Relativistic Magnetohydrodynamics
ROI_PAC	Repeat Orbit Interferometry Package
RPIC	Relativistic Particle-in-Cell (kinetic simulation)
RPM	Redhat Package Manager
SAR	Synthetic Aperture Radar
SBSSR	Small Body Surface Sample Return
SCIGN	Southern California Integrated GPS Network
SDO	Solar Dynamics Observatory
SDSS	Sloan Digital Sky Survey
SIM	Space Interferometry Mission
SMAP	Soil Moisture Active Passive
SMD	Science Mission Directorate
SMP	Symmetric Multiprocessing
SOA	Service-Oriented Architecture
SOFIA	Stratospheric Observatory for Infrared Astronomy
SPASE	Space Physics Archive Search and Extract
STSI	Space Telescope Science Institute
TB	Terabyte
TES-FO	Tropospheric Emission Spectrometer-Follow On
TF	Teraflops
TPF	Terrestrial Planet Finder
UAVSAR	Uninhabited Aerial Vehicle Synthetic Aperture Radar
UNEP	United Nations Environment Programme

UNESCO	United Nations Educational, Scientific, and Cultural Organization
UPC	Unified Parallel C
USWRP	U.S. Weather Research Program
VC	Virtual California
VIIRS	Visible/Infrared Imager/Radiometer Suite
VM	Virtual Machine
VNC	Virtual Network Computing
WAN	Wide-Area Network
WCRP	World Climate Research Program
WCS	Web Coverage Service
WG1	Working Group 1
WGCEP	Working Group On California Earthquake Probabilities
WGCM	Working Group on Coupled Modeling
WGSIP	Working Group on Seasonal to Interannual Prediction
WMAP	Wilkinson Microwave Anisotropy Probe
WMO	World Meteorological Organization
WMS	Web Map Service
WRF	Weather and Research Forecasting
WSDL/SOAP	Web Services Description Language/Simple Object Access Protocol
WWRP	World Weather Research Programme
XML	Extensible Markup Language
YOTC	Year of Tropical Convection

Appendix B: Participants

Legend: ^ Workshop Co-Chair, * Panel Co-Chair, + Organizing Committee

Name	Affiliation	Panel
Don Anderson	NASA/HQ	Earth System Modeling
Sandy Antunes	NRL	Heliophysics
John Baker	NASA/GSFC	Astrophysics
Robert Bauer	NASA/GSFC	NASA Observer
Bryan Biegel	NASA/ARC	Technology
Randy Bolanos	NASA/ARC	Earth System Modeling
Gene Bonfiglio	NASA/JPL	Planetary Sci. & Mission Eng.
Joe Bredekamp	NASA/HQ	NASA Observer
Dan Burkhart	NASA/JPL	Planetary Sci. & Mission Eng.
Ilene Carpenter	SGI	Earth System Modeling
Chris Catherasoo	NASA/JPL	Technology
Renyue Cen	Princeton Univ.	Astrophysics
Joan Centrella*	NASA/GSFC	Astrophysics
Jiun-Dar Chern	NASA/GSFC	Earth System Modeling
Tom Clune*+	NASA/GSFC	Technology/Earth System Modeling
Jarrett Cohen+	NASA/GSFC/GST	Earth System Modeling
Arlindo da Silva	NASA/GSFC	Earth System Modeling
Manuel de la Torre Juarez	NASA/JPL	Planetary Sci. & Mission Eng.
Andrea Donnellan*	NASA/JPL	Solid Earth
Dan Duffy	NASA/GSFC	Technology/Heliophysics
William Emanuel	NASA/HQ	NASA Observer
Jared Entin	NASA/HQ	NASA Observer
John Evans	NASA/GSFC/GST	Earth System Modeling
Robert Ferraro*+	NASA/JPL	Technology
Jim Fischer+	NASA/GSFC	Earth System Modeling
Mei-Ching Fok	NASA/GSFC	Heliophysics
Geoffrey Fox*	Indiana Univ.	Solid Earth
Michael Freilich	NASA/HQ	NASA Observer
Ichiro Fukumori	NASA/JPL	Earth System Modeling
Ron Gelaro	NASA/GSFC	Earth System Modeling
Robert Granat	NASA/JPL	Solid Earth
Anthony Gualtieri	NASA/GSFC	Solid Earth
Lika Guhathakurta	NASA/HQ	Heliophysics
Jing Guo	NASA/GSFC/SAIC	Earth System Modeling
Jim Hack	ORNL	Earth System Modeling
Michael Hesse*	NASA/GSFC	Heliophysics
Chris Hill	MIT	Earth System Modeling
Michael Kalb	Science Interfaces and Assessment	Earth System Modeling
George Khazanov	NASA/MSFC	Heliophysics
Weijia Kuang	NASA/GSFC	Solid Earth
Akshay Kulkarni	Cornell Univ.	Astrophysics
John LaBrecque*	NASA/HQ	Solid Earth
Bill Lapenta*	NASA/MSFC	Earth System Modeling
Eric Larour	NASA/JPL	Solid Earth
Tsengdar Lee^+	NASA/HQ	NASA Observer
Doug Lin	Univ. of California, Santa Cruz	Astrophysics

Name	Affiliation	Panel
Alexander Lipatov	NASA/GSFC/GEST	Heliophysics
Mike Lisano*	NASA/JPL	Planetary Sci. & Mission Eng.
Michael Little	NASA/LaRC	Technology/Solid Earth
Steve Lubow	STSI	Astrophysics
Peter MacNeice	NASA/GSFC	Heliophysics
Piero Madau	Univ. of California, Santa Cruz	Astrophysics
Gail McConaughy	NASA/GSFC	Earth System Modeling
Piyush Mehrotra	NASA/ARC	Technology/Earth System Modeling
Dimitris Menemenlis	NASA/JPL	Earth System Modeling
Scott Michael	Indiana Univ.	Astrophysics
Duncan Niblett	NASA/HQ	NASA Observer
Ken-Ichi Nishikawa	NASA/MSFC	Astrophysics
Charles Norton	NASA/JPL	Solid Earth
Hamid Oloso	NASA/GSFC/AMTI	Technology/Heliophysics
Merav Opher	George Mason Univ.	Heliophysics
Kathy Pegion	COLA	Earth System Modeling
Christa Peters-Lidard	NASA/GSFC	Earth System Modeling
Marlon Pierce	Indiana Univ.	Solid Earth
L. Harper Pryor	NASA/GSFC/SAIC	Technology/Earth System Modeling
Bill Putman	NASA/GSFC	Earth System Modeling
Rama Ramapriyan	NASA/GSFC	Earth System Modeling
Steve Reinhardt	Interactive Supercomputing, Inc.	Earth System Modeling
Michele Rienecker*	NASA/GSFC	Earth System Modeling
Mike Rilee	TopQuadrant	Planetary Sci. & Mission Eng.
Aaron Roberts*	NASA/GSFC	Heliophysics
Jose Rodriguez	NASA/GSFC	Earth System Modeling
Larry Roelofs	GST	Planetary Sci. & Mission Eng.
Paul Rosen	NASA/JPL	Solid Earth
John Rundle	Univ. of California, Davis	Solid Earth
Michael Salamon*	NASA/HQ	Astrophysics
Ellen Salmon	NASA/GSFC	Earth System Modeling
John Schnase	NASA/GSFC	Earth System Modeling
Siegfried Schubert	NASA/GSFC	Earth System Modeling
Steinn Sigurðsson	Penn State Univ.	Astrophysics
Stephen Simms	Indiana Univ.	Astrophysics
Eli D. Skulsky*	NASA/JPL	Planetary Sci. & Mission Eng.
Hongbo Su	CREW	Earth System Modeling
Max Suarez	NASA/GSFC	Earth System Modeling
Angela Taylor+	Harris Corp.	
Ricardo Todling	NASA/GSFC	Earth System Modeling
Michael Turner	NASA/MSFC	Earth System Modeling
Azita Valinia^+	NASA/GSFC	NASA Observer
Jim van Meter*	NASA/GSFC	Astrophysics
Scott Wallace	NASA/GSFC/CSC	Earth System Modeling
Bill Ward	NASA/GSFC/CSC	Earth System Modeling
Phil Webster	NASA/GSFC	Earth System Modeling
John Wise	NASA/GSFC	Earth System Modeling, Astro
Zhiyong Zhang	Eloret, Inc.	Earth System Modeling
Shujia Zhou	NASA/GSFC/NG	Earth System Modeling