


Report from the Earth Science Enterprise Computational Technology Requirements Workshop


April 30 - May 1, 2002




Observations



Data Management




Visualization



**Problem
Solving
Environments
and Tools**



**Computing
Platforms and
System Management**



Algorithms

Prediction

Table of Contents

Executive Summary	i
I. Introduction	1
II. Science Drivers and Capability Requirements	3
A. Weather	3
B. Climate	7
C. Solid Earth	9
III. Technology Cross-Cut of Gaps Identified	12
A. Computing Platforms	12
B. Data Management	15
C. Programming Environments and Tools	17
D. Algorithms	19
E. Visualization and Data Mining	20
F. Distributed Computing	20
G. Computing Platforms System Management	21
IV. Summary with Priorities	22
A. Gaps Expected to be Filled Without Investment	22
B. Recommended Investments	23
C. Identification of Infrastructure Issues	26
Appendix A: Weather Panel Report	A-1
Appendix B: Climate Panel Report	B-1
Appendix C: Solid Earth Panel Report	C-1
Appendix D: Workshop Attendees	D-1

Executive Summary

The ESE Computational Technology Requirements Workshop was held April 30 and May 1 at the Washington Plaza Hotel in Washington, DC. It was open to all interested members of the community and was announced widely. Three discipline specific panels — Weather, Climate, and Solid Earth — were charged with synthesizing the community input collected during the 2 days of the workshop discussions and preparing final reports and roadmaps at its conclusion.

Weather Computational Requirements Driver

The NASA role in weather forecasting is specifically a role of enabling. Therefore, NASA's activities are dependent upon the activities in other agencies.

To succeed in enabling the Nation's weather prediction goals, NASA research activities need to be integrated with NOAA and Navy operational requirements. These include:

- 1) Techniques developed within NASA have to fit into the operational time allowance for real-time weather forecasting. This requires the end-to-end forecast-assimilation system to run in less than 2 hours of wall time.
- 2) Attributes of NASA algorithms, e.g., resolution, data usage, parameterizations, must be consistent with the state-of-the-art operational systems. Otherwise, the impact of NASA research will be diminished.

The computational requirements for the anticipated forecast-assimilation system for 2010 (see Appendix A) are currently driven by data volume. A data volume of 10^{11} observations per day is chosen based on projections of satellite observing systems that are planned to be operational and relevant to weather forecasting. As a baseline, the current horizontal resolution is order 100 km, with 50 vertical levels, and current data usage is order 10^6 observations per day. Using the presently measured sensitivities of the model and assimilation algorithm to resolution and data usage, the projected computer requirements for 2010 are 10^9 times the current requirement. The primary driver of this requirement is data usage. Such a requirement is unrealistic, and therefore, to address the data usage requirements new strategies are needed to compact the information and, effectively, reduce the data volume, prior to assimilation. These strategies require consideration of all aspects of the system.

Climate Computational Requirements Driver

Climate (or more generally Earth system) research, including that for improved prediction, relies heavily on high-end computing resources. Requirements come from a variety of subdisciplines – seasonal-to-interannual (S-I) variability and prediction, decadal variability and prediction, global change assessments, atmospheric/ocean/land assimilation for analyses (as well as that for forecast initialization), and coupled physical-biogeochemistry modeling, to name a few. For the purpose of this report the panel focused on the problems of seasonal-to-interannual prediction and decadal prediction because the types of problems to be solved and the models and resolutions used for these areas are similar to those used for the other subdisciplines.

Current S-I climate prediction systems consist of ocean-atmosphere-land-sea-ice models coupled to ocean and land data assimilation systems. To be useful for climate applications, single-image

model performance must be roughly 1,000 simulated days per wall-clock day, and since climate forecasts need to be based on ensembles of runs, typical systems run 8 to 10 concurrent coupled images (or 15 to 20 uncoupled images) now. The panel's assessment of requirements for future computer platforms is based on the assumption that as computer power increases the size of the calculation and the number of concurrent images will also increase while roughly maintaining these performance goals (1,000 d/d single image, 20,000–30,000 d/d aggregate throughput).

During the next 10 years, the magnitude of the calculation will increase due to the greater model complexity and increases in resolution. The panel has assumed that the latter will be the main driver of computer performance. The panel also used the atmospheric model as characteristic of the entire problem.

A single image of an atmospheric configuration of 1° resolution with 100 layers (including the full stratosphere) and 40 on-line chemical tracers would require a throughput of about 5 TFLOPS; with a coupled ocean of 50 layers at $1/2^\circ$ resolution the requirement is 6.6 TFLOPS. The coupled model configuration with a job mix equivalent to 100 concurrent images (ensembles, parameter sensitivity sweeps, etc) would then require 660 TFLOPS.

Solid Earth Computational Requirements Driver

The field of solid Earth science is currently mission and data poor. The panel anticipates, however, an increasing number of missions and data in the next 10 years. These include gravity field determination missions such as GRACE, GRACE follow-on, GOCE; radar interferometry missions such as ECHO; CHAMP and other magnetic field determination missions, ICESat, a laser altimeter mission; GPS data collection from hundreds of ground stations (SCIGN/PBO) as well as space platforms; and a wide variety of hyperspectral, high-resolution visual, multispectral thermal, lidar imaging missions, such as VCL.

Space technologies will allow measurement of 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 huge quantities of data and information to be collected into forecast models will require that information technology resources be developed in concert with advanced sensor and detection capabilities.

Taking ECHO as an example, there will be 100 GB of data per day, which should be fully utilized for forecasting events on the global fault system. Accuracy may be refined in regions of interest by applying finite element techniques using known variations in rheology and requiring 1 TFLOPS sustained. For forecasting, a dynamic model is required that can span scales larger than a finite element code. Boundary element methods combined with good fault-friction models can fulfill this need in principle, although substantial testing and development remain. Such a system can be coupled efficiently with the finite element regional models. Operations count can be estimated from assumptions of regional cost, plus assumptions of multipole coupling between regions; most of the cost is then in the regional interactions. So with 80 regions covered at 10 per day, 10^5 fault patches per region with hence 10^{10} interactions, 100 operations per interaction, 1,000 update steps for data-driven corrections (including recent history), and 0.1 computational efficiency, the requirement is 1 TFLOPS sustained. For better forecasting, eigenpattern rate

techniques are needed. These pattern rates will be used to daily update an earthquake hazard map. End to end, processing ECHO for earthquake forecasting comes to about 2 TFLOPS sustained rate.

Technology Cross-Cut Gaps Identified

The specific technology capabilities required to address the ESE prediction goals have been identified, based on the capability needs identified by the panels. To the extent possible, those capabilities are quantified and the driving science application responsible for the “high water marks” identified.

Sustained compute throughput in the 10 TFLOPS range for single applications is common to all of the panels. In all cases, this capability is driven by resolution requirements. Current model resolutions are grossly insufficient for addressing the prediction goals. The weather model throughput target is the stressing mark — it is higher than required for climate due to higher resolution and the need to meet real-time processing deadlines in an operational environment.

Individual application performance is called out as an issue by both the weather and climate panels. Currently, applications running on high-end computers typically realize only about 10% of the peak rated performance of the constituent processors. If that efficiency continues to hold, achievement of teraFLOPS throughput performance will require 10,000’s of processors. Increasing single processor performance efficiency to 40% will result in a 4x reduction in processors required, and thus a 4x cost reduction in the computing platforms NASA must purchase to meet these requirements. Scaling to 1,000’s of processors for many key applications is essential but remains untested.

All three panels specify the capability to manage petabytes (10^{15} bytes) of data as a requirement in 2010. The increases in model resolution over the decade will directly translate into dramatic increases in input data, output data products, archival storage, and transported data volume. It is the number one issue for the weather panel, and a driving concern for climate and solid Earth. With projected data products produced over the course of a year being in the 100’s to 1,000’s of petabytes, these drive a data warehousing and distribution requirement that is quantitatively different than the data archive center problem being addressed in other technology programs. Indeed, the mission data archive centers will be required to feed data into the modeling activities at rates at least as fast as the data is collected. The rapid increase in data (both numerical and observational) is expected to swamp existing visualization tools.

With the sources of data widely varied in the future, any application that needs to access data from a variety of sources cannot possibly be expected to deal with the idiosyncrasies of each. The data management system must present a uniform interface to all of the data under its purview, even to the point of transparently translating among a set of standard formats.

The data movement implied by the data volumes required will drive requirements on network layers between computing assets and storage assets, and end user clients into the 100’s of Gbytes/s range.

In this era of ever-increasing complexity of scientific software, the ability to develop new software components and assemble existing components into new applications is viewed as a critical need across all three panels. Two of the panels (Climate and Weather) directly cite the ESTO/CT Earth System Modeling Framework (ESMF) as a key technology to address such capabilities in the coming decade. Sustained government support for maintenance and future

development of the ESMF beyond its 3-year development project is imperative. The Solid Earth panel envisions virtually identical issues driving the creation of a Problem Solving Environment (PSE), which is directed towards the solid Earth modeling problems.

Portability without significant loss of performance is a capability required by all three panels. A key enabler of portability is the availability of standard common libraries across all of the platforms of interest. Libraries that are optimized to the specific platform while maintaining the same interface across platforms are required to amortize the code development investment in each application. The ability for the average code developer to evaluate application performance on thousands of processors and understand what to do to improve it is also required.

Across all disciplines, potential improvements in algorithms offer hope for significant reduction in resource requirements or alternatively more/better scientific results from fixed resources. Different panels predictably identified different areas in which algorithmic improvement is most critical. The most extreme case is in the numerical weather prediction community, where the volume of observational data anticipated by 2010 is expected to overwhelm hardware improvements by many orders of magnitude. History shows that as much improvement in throughput is obtained from algorithm advancements as is obtained from hardware improvements.

Conclusion

NASA science requires major advances in computational technology. NASA's unique driver is the data. Science applications that are key to the achievement of the prediction goals have already been identified that will require new developments in computational technologies. Continued, focused investment in a science driven technology development program is required for success in the ESE.

I. Introduction

Workshop Purpose

Fulfilling NASA's Earth Science Enterprise (ESE) prediction goals for 2010 and beyond requires state of the art computational technologies. As an initial step to plan for future requirements, the high-end computing workshop, sponsored by ESE, was conceived to

- define Earth system scientific drivers needed to achieve prediction goals
- evaluate existing capabilities
- perform a gap analysis identifying required computational technologies
- prioritize identified requirements for possible Enterprise technology development investment.

Computational technologies in this context spans both software and hardware needed to enable efficient ingestion of large amounts of remotely sensed data into Earth system models, execute those models in a timely manner, and understand the output of those models. Many thrusts of this effort will enable interdisciplinary science and applications scenarios requiring linked or nested modeling components.

Workshop Process

The ESE Computational Technology Requirements Workshop was held April 30 and May 1 at the Washington Plaza Hotel in Washington, DC. It was open to all interested members of the community and was announced via email to an extensive list of NASA affiliated Earth science researchers identified via the SYSEFUS database.

The workshop was organized into three discussion sessions under the topics of Weather, Climate, and Solid Earth. Well-recognized contributors to the identified scientific disciplines as well as technologists representing areas pertinent to the science were specially invited to participate as panel members. The three discipline specific panels were charged with synthesizing the community input collected during the two days of the workshop discussions and preparing final reports and roadmaps at its conclusion.

Three documents were made available to all participants in advance of the workshop for study and preparation:

- Understanding Earth System Change, NASA's Earth Science Enterprise Research Strategy for 2000-2010
- Earth System Observations and Modeling Implementation Plan (Chapter 7 of the Enterprise strategy document)
- Earth Science Prediction Goals for 2010

At the opening session, participants were welcomed by Dr. Azita Valinia (HQ/YS). Mr. George Komar, gave an overview of the Earth Science Technology Office. Dr. Jack Kaye gave an overview of NASA's ESE Research and Modeling strategy. Workshop participants were then charged by Dr. Robert Ferraro (JPL) to carry out the workshop charter.

At the end of each day the participants were convened at the plenary session to discuss reports of the day's results in each panel. At the end of the second day the non-panelists were thanked for their input and excused.

On the third day, May 2, the panelists convened in executive session and began assembling their reports. The teams returned final reports of their respected panel within 7 to 10 days. ESTO/CT project staff led by Mr. Jim Fischer and Dr. Robert Ferraro along with Dr. Azita Valinia (HQ/YS) created this integrated document based on the panels' original reports. This integrated report was briefed to the Associate Administrator for NASA's Earth Science Enterprise, Dr. Ghassem Asrar, on May 23, 2002.

The body of this integrated report synthesizes the three original panel reports and provides an executive summary of the workshop findings. Following the introduction, it highlights the primary scientific drivers of the identified technology gaps, tabulates the gap technologies from all the science panels, and then prioritizes the gap technologies. Appendices A, B, and C provide the reports submitted by Weather, Climate, and Solid Earth panels respectively and in their entirety. Appendix D lists the names of the 123 workshop attendees.

Summarized Charter

The elements of the ESE Computational Technology Requirements Workshop charter are:

- 1) Identify the Earth Science Enterprise capabilities required to achieve the Earth System Prediction Goals for 2010.
- 2) Evaluate current Earth Science Enterprise capabilities against these requirements.
- 3) Determine and quantify what gaps in capabilities exist.
- 4) Determine which gaps can be addressed by advances in computational technologies
- 5) Identify and quantify the advancements in computational technology capabilities that require NASA investment in order to bridge these gaps.
- 6) Prioritize these capability advancements in terms of their likelihood to enable the Earth System Prediction Goals for 2010.
- 7) Create a roadmap for each unique capability advancement.
- 8) Create a final report.

II. Science Drivers and Capability Requirements

The full reports submitted by each Weather, Climate, and Solid Earth panels appear in Appendices, A, B, and C, respectively. In this section, we summarize the science drivers, and capability requirements based on the reports.


A. Weather

Science Drivers

ESE Prediction Goals for Weather

The given prediction goals from the science plan are:

Today's Capability	2010+ Capability
3-day forecast at 93%*	5-day forecast at > 90%*
7-day forecast at 62%*	7–10-day forecast at 75%*
3-day rainfall forecast not achievable	3-day rainfall forecast routine
Hurricane landfall +/- 400Km at 2–3 days	Hurricane landfall +/- 100Km at 2–3 days
Air quality day by day	Air quality forecast at 2 days



The diagram on the previous page represents the high-level functions that must be addressed by the forecast-assimilation system, highlighting the development areas and interfaces that most directly impact high-performance computational technologies. In general computational requirements are driven by resolution, data usage, comprehensiveness or coupling (running processes concurrently rather than sequentially), robustness (explicit representation of processes modeled as approximations or parameterization), and data volume and data access. (These estimates are given in the ensuing table.)

	2002 System	2010+ System	
Resolution • Horizontal • Vertical levels • Time step • Observations ○ Ingested ○ Assimilated	100 km 55 30 minutes 10^7 / day 10^5 / day	10 km 100 6 minutes 10^{11} / day 10^8 / day	
System Components:	• Atmosphere • Land-surface • Data assimilation	• Atmosphere • Land-surface • Ocean • Sea-ice • Next-generation data assimilation • Chemical constituents (100)	
Computing: • Capability (single image system) • Capacity (includes test, validation, reanalyzes, development)	10 GFLOPS 100 GFLOPS	Must Have 20 TFLOPS 400 TFLOPS	Important 50 TFLOPS 1 PFLOPS
Data Volume: • Input (observations) • Output (gridded)	400 MB / day 2 TB / day	1 TB / day 10 PB / day	
Networking/Storage • Data movement ○ Internal ○ External • Archival	4 TB / day 5 GB / day 1 TB / day	20 PB / day 10 TB / day 10 PB / day	

The forecast-assimilation system is represented (in the diagram) by the checked square and oval with an interface between them. These two components traditionally demand much of the scientific and computational development activities. The interface is specifically noted as the frequency of communication between the modeling component and assimilation/analysis component profoundly impacts computational requirements. Projecting a target resolution provides the most straightforward parameter to estimate requirements growth in the model. A horizontal resolution for a global atmospheric model of 10 km is chosen based on the target resolution of both NOAA and the NAVY. This resolution would resolve many features well enough for assessment of the regional impact of changing climate-scale phenomena. A vertical resolution using 100 levels from the ground to the top of the mesosphere is chosen. The

computational requirement for the assimilation/analysis algorithm is currently driven by data volume. A data volume of 5×10^{10} observations per day is chosen based on projections of satellite observing systems that are planned to be operational and relevant to weather forecasting. As a baseline, the current horizontal resolution is order 100 km, with 50 vertical levels, and current data usage is order 10^6 observations per day. Using the presently measured sensitivities of the model and assimilation algorithm to resolution and data usage, the projected computer requirements for 2010 are 10^9 times the current requirement. The primary driver of this requirement is data usage. Such a requirement is unrealistic, and therefore, to address the data usage requirements new strategies are needed. These strategies require consideration of all aspects of the system.

Panel Recommendations

1. Data and Information Management

Improved data and information management is critical. There are two major aspects of this. The first aspect concern the development of computational systems that effectively acquire, order, and prepare data for use by either other algorithms or scientific users.

The second concerns methods to improve the extraction of information from observational data prior to assimilation must be developed. This is termed the optimization of the observing system, and activities are far ranging. They range from methods advocated in computer science such as data mining and hierarchical segmentation, to sampling and data selection techniques used in current data assimilation systems, to increased utility of observing system simulation experiments (OSSE's) and observing system experiments (OSE's). These final approaches run end-to-end assimilation systems to predict and validate data impact. Increased interaction between the modeling and assimilation components of the data assimilation system to target which observations might have the most impact provides a potentially fruitful approach with broad implications for computational systems. New data assimilation algorithms must be developed that are quasi-independent of data volume.

2. Systems Software and Systems Engineering to Support High-End Computing

In order to assure that the increases in technological capability that are achieved by the field as a whole benefits NASA applications, direct investment in systems software and systems engineering is needed. This investment in such a specialized field requires significant planning for organization and integration. The cross organization development of standards to support vendors, applications and discipline software is essential.

3. Application-Specific Software Tools

The development of new application-specific software environments like the Earth System Modeling Framework (ESMF) is an example of what is needed for the future. The weather development community and the Earth science community in general is moving toward the open standards efforts to support, data frameworks, visualization tools, and modular code environments. The current tools like IDL will be overwhelmed by the tremendous increase in data volumes of the next decade. The recasting of applications into truly parallel data warehouses with supporting tools is essential for the next level in data assimilation, modeling, and forecasting. We strongly recommend the design basis require the remove all order one bottlenecks in the end-to-end processing flows.

4. Analysis-Specific Software Tools

The analysis of product by users will be directly dependent the provided frameworks. These tool needs will encompass all the issues for hardware and software, and the development of tool products from a user centric position is essential for success. So the user groups must be involved in the design, development, and re-design of the analysis specific tools. Again, the lack of a commodity marketplace for developers will require NASA-specific investments to support discipline needs.

Visualization is an area where leveraged NASA investments can advance software and hardware computational technology for the pursuit of the 2010 ESE weather prediction goals. As mentioned, the data volume involved in weather prediction will increase by many orders of magnitude by 2010. As a result, the development of weather prediction software systems will inherently involve the analysis of huge data sets. This analysis will be interactive during the early development stages; the search for useful analysis or assimilation techniques must be conducted by domain experts. Humans have a sensory system that is highly evolved to transfer massive data sets to the brain: vision. Thus, visualization technologies will be central to developing ESE computation and data assimilation algorithms. In all stages of the program, visualization is a critical enabling technology for transforming experimental data produced by coupled Earth system models to information and knowledge, and an essential tool for the communication of complex interactions to both experts, decision makers, and the public.

B. Climate

Climate (or more generally Earth system) research, including that for improved prediction, relies heavily on high-end computing resources. Requirements for high-end computing resources come from a variety of sub-disciplines — seasonal-to-interannual (S-I) variability and prediction, decadal variability and prediction, global change assessments, atmospheric/ocean/land assimilation for analyses (as well as that for forecast initialization), coupled physical-bio-geochemistry modeling, to name a few. During the workshop, the climate panel focused on the problems of S-I and decadal prediction.

Science Drivers

The Seasonal-to-Interannual Prediction Scenario

In the 2010 timeframe, we envision that NASA will continue to contribute to the nation's multi-model seasonal predictions through forecasts and predictability experiments conducted by NSIPP. NSIPP will be one of approximately five national groups conducting routine experiments, sharing forecasts, simulations, and analyses. NASA's unique emphasis is on the optimal use of satellite observations to enhance prediction skill. The current focus is on 6-month prediction skill and all groups undertaking such predictions with coupled general circulation models regard them as experimental. ESE's goal for 2010 is for 6–12 month routine seasonal prediction and 12–24 month experimental prediction. Increased emphasis will be placed on larger ensembles as we attempt to extract more skill and shorter (intraseasonal) timescales and look to predict the likelihood of extreme event (e.g., flood) occurrences. With the expected increase in processor speed will come the desired increase in model resolution: 25 km for the atmosphere and 10km for the ocean. The need to quantify forecast reliability will lead to larger ensembles at longer lead forecast time. This may be accommodated by multiple-scale ensembles, with degraded resolution at longer leads. Reliable estimates of forecast spread will require the uses of multi-model ensembles. The models will have expanded physics to account for influences on forecast skill outside the tropical Pacific — e.g., high latitude processes, stratospheric processes, changes in land cover. As in other disciplines, attention to better cloud representation will be essential. The coupled initialization will continue to focus on the ocean and land. Since global predictions will be used to force offline regional hydrologic and crop models, global predictions will need to be stored (and retrieved) at high spatial and temporal resolution. We anticipate a successful transition to complete utilization of NASA's ESMF and this will enable flexibility in model component mix for ensembles.

The Decadal Experimental Prediction Scenario

To this point predictions on the decade-to-century scale time frame have focused primarily on global or hemispheric scales; for practical purposes, regional scale predictions are required. Although confidence in regional predictions will probably increase in the timeframe to 2010, it is to be expected that predictions will remain on an experimental track for some time. NASA's contributions to these goals will emerge from GISS, with plans for a 1/2° ocean model, a 1° atmospheric model resolving both the troposphere and stratosphere, and air chemistry with about 40 tracers if computing power is available. The same model will also be used for even longer global change integrations, possibly at reduced resolution.

Required Capabilities

Several developments are needed to satisfy the high-end computing requirements for climate prediction. These are summarized as follows:

1. Computing Platforms

- Internode communications and shared memory utilization
- I/O bandwidth and storage
- Fortran compilers to improve single processor performance

2. Problem Solving Environment

- Commitment for sustainability of the ESMF
- Increased functionality of the ESMF, especially in terms of data structures, operations to support assimilation, and integrated tools to enhance model development and utilization

3. Data Management

- High-performance networks
- High-performance analysis platform optimized for I/O
- Distributed data archive
- Standardized software for storage and analysis
- Parallel extraction tools for data mining
- Parallel and distributed analysis and visualization tools

C. Solid Earth

The field of solid Earth science is currently mission and data poor. However, it is anticipated that the number of solid Earth missions and hence related data volume will be increasing rapidly in the next 10 years. The sub-fields of solid Earth that relate to NASA's prediction 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, and ecology, hydrology, and vegetation.


To fully understand and forecast solid hazards, high-resolution, global observations with real-time or near-real-time data streams and processing will be required. However, even with the most advanced observational systems, the temporal sampling of such phenomena is poor. In order to understand fully these highly complex systems, simulations must be carried out concurrently 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. Integrating the huge quantities of data and information to be collected into forecast models will require that information technology resources be developed in concert with advanced sensor and detection capabilities.

The quantification of the computational needs is timely in light of the recent activities of the NASA Solid Earth Science Working Group (SESWG). This group is putting together a vision for solid Earth science for the next 25 years (<http://solidearth.jpl.nasa.gov>).

Science Drivers

The physical processes associated with the solid Earth take place on many scales of space and time. Simulations and theory must account for how these many scales interact. The solid Earth is complex, nonlinear, and self-organizing. Recent work suggests strong correlations in both space and time resulting in observable space-time patterns. Recent advances in computational science and numerical simulations enables studies of the complex solid Earth system making it possible to address the following critical scientific questions:

- 1) How can the study of strongly correlated solid Earth systems be enabled by space-based data sets?
- 2) What can numerical simulations reveal about the physical processes that characterize these systems?
- 3) How can modern Information Technology enable new understanding of the basic physics?
- 4) How do interactions in these systems lead to space-time correlations and patterns?
- 5) What are the important feedback loops that mode-lock the system behavior?
- 6) How do processes on a multiplicity of different scales interact to produce the emergent structures that are observed?
- 7) Do the strong correlations allow the capability to forecast the system behavior in any sense?



The physical processes associated with the solid Earth take place on many scales of space and time. Simulations and theory must account for how these many scales interact.

To fully address these questions, three-dimensional modeling of Earth's gravity and geomagnetic field is needed. In the next 10 years, information from multiple sources on multiple scales needs to be incorporated into an integrated analysis. For this purpose, worldwide computational systems supporting the gathering, integration, visualization, simulation and interpretation of several petabytes of data per year is needed.

Required Capabilities

1. Developing a Virtual Observatory

An important aspect of data collection is to create distributed centers for storing unique data sets and developing the infrastructure to compare, analyze, and ingest complementary data sets, such as ice topography and sea level changes. Currently, cooperative federated databases do not exist; and the data are heterogeneous and widely distributed in a variety of formats. Because of the high volumes of data, now and in the future, data mining and other approaches for interacting with the data must be developed (possibly including more onboard processing). A Solid Earth Research Virtual Observatory (SERVO) is recommended to address the critical need for seamless access to large distributed volumes of data. Such a system would have the following characteristics: distributed data at centers and ground stations; thousands of sites with volumes of 1 TB to 1 PB; a multi-tier architecture for staging of the data; middleware to control integrity and versioning; support standards developed within the community; working within data grid projects; customized for key NASA solid Earth needs; and high-performance access of 100 GB files within 40 TB datasets within 5 minutes to user and program-to-program communication in

milliseconds using staging, streaming, and advanced cache replication. A roadmap for the virtual observatory project is given in the full report (Appendix C).

2. Developing an Earth Science Problem Solving Environment

Creating realistic simulations of solid Earth processes requires complex models. The complexity of the problem requires high-performance computers to realize these integrated models. The solid Earth panel recommends the following to address this need (in order of priority):

- 1) Develop a solid Earth science Problem Solving Environment to support 10 solid Earth sub-fields. This includes a prototyping environment for developing a model in a month; a modular framework for solid Earth applications; a collaboratory, which includes teams of scientists and computational experts; visualization and data analysis tools, seamless computer access; and integration of data and multicomponent models.
- 2) Development of improved parallel algorithms/applications (on the order of 100) with a scaled efficiency of at least 50% on large clusters; models and model assessment methods; and data assimilation techniques.
- 3) The computational needs required to address model complexity include 100 TFLOPS sustained rate capability per model, 5 TB total memory per model, and 10^4 petaFLOPS throughput (3 TFLOPS per second sustained) per sub-field per year.

A roadmap for the problem solving environment project is given in the complete panel report (Appendix C).

3. Improving Computational Environment

Many of the solid Earth problems are extremely computationally expensive, beyond the capability of existing supercomputers. Therefore, the capacity of high-end computers must be increased to petaFLOPS with terabytes of RAM. Distributed and cluster computers should be available for decomposable problems, for rapid development and for cost performance. The capacity of networks must be increased and investments should be made in the development of Grid technologies. Computational analysis would also be enabled by more open and seamless access to computational resources.

A roadmap for the computational environment project is given in the complete panel report (Appendix C).

III. Technology Cross-Cut of Gaps Identified

In this section, we identify the specific technology capabilities required to address the ESE prediction goals, based on the capability needs identified by the panels. To the extent possible, we quantify those capabilities and identify the driving science requirement responsible for the “high water mark”. Some technology capabilities are not easily quantified — in these cases, we describe a use scenario that can be used to bound the capability requirement.

The technology areas are arranged in order of decreasing constituency. Not all panels identified the capabilities described towards the end of this section as needs. However, this is not an indication that these are less important than those described first. Each technology listed will be an important contributor to achieving the ESE prediction goals

A. Computing Platforms

Sustained Throughput

Sustained compute throughput in the teraFLOPS range for single applications is common to all of the panels. In all cases, this capability is driven by resolution requirements. Current model resolutions are grossly insufficient for addressing the prediction goals. The weather and climate panels identified specific model resolution targets that they feel are necessary for their prediction goals, and solid Earth has developed its throughput requirement in terms of the expected operations count extrapolated from current analysis codes. The weather model throughput target is the stressing mark — it is higher than required for climate due to higher resolution and the need to meet real-time processing deadlines in an operational environment. The climate and solid Earth throughputs are less time critical, and are based on expectations of reasonable turn-around times (1,000 days/wall clock day for climate, 1-day end-to-end processing for solid Earth earthquake prediction).

	Model	Single Image Throughput	Estimated Capacity Required
Weather	10-Day Forecast Atmosphere: 10 km horizontal, 100 levels vertical 10^{10} observations	20 TFLOPS	400 TFLOPS
Climate	S-I Prediction Atmosphere: 25 km horizontal Ocean: 6 km horizontal	5 TFLOPS	100's TFLOPS
Solid Earth	Earthquake Fault Slip 16M finite elements 100k boundary elements	2 TFLOPS	10's – 100 TFLOPS

Sustained Throughput and Capacity Requirements

III. Technology Cross-Cut of Gaps Identified

Capacity measures the aggregated throughput capability required in the day-to-day environment. The predication goals cannot be accomplished by single executions. Weather prediction, for example, is expected to be a continuous process of updates every 2 hours, model development, testing, validation, and reanalysis of past observations. Climate prediction is based on ensembles of model executions (10 to 20 executions), and solid Earth advocates a scenario with many mission data sets being analyzed by many research organizations. Capacity estimates are often hard to justify outside an operational environment (what is “enough”?). The weather prediction scenario has the most precision, since it is a reasonable extrapolation based on current operational capabilities in the DAO. These numbers should be taken as order of magnitude estimates only. But it is clear that total system throughput must be substantially larger than the single application requirement.

Capacity is an infrastructure capability, and could be provided by federating many individual platforms together to meet the aggregate throughput demand. The solid Earth panel advocates just such an approach, even to the extent that the resources are geographically distributed. An operational environment for weather prediction is on the other extreme. The data volumes ingested and the data volumes output and archived (see data management below) will drive data transport networking requirements in the direction of a dedicated facility. The tradeoff between single (or a few) large platforms capable of 100’s of teraFLOPS sustained throughput and 100’s of multi-teraFLOPS throughput systems is going to be discipline specific.

Single Application Performance

Individual application performance is called out as an issue by both the weather and climate panels. The following table illustrates the problem, based on the throughput requirement for the weather forecast system. Moore’s law is expected to govern the improvement in single-processor peak performance throughout the rest of the decade. Assuming that processor speeds double every 18 to 20 months, the single processor peak performance in 2010 will be in the 20 to 50 gigaFLOPS range. Currently, applications running on high-end computers typically realize only about 10% of the peak rated performance of the constituent processors. If that efficiency continues to hold, achievement of teraFLOPS throughput performance will require thousands of processors.

Single Application Throughput	Projected Single Processor Performance	Application Efficiency	Total Processors Required	Total Application Memory
20 TFLOPS	20–50 GFLOPS	10%	4,000–10,000	4 TB

Single Application Processor Count and Memory Requirements.

Current application efficiency is dominated by the inability of compilers to maximize cache memory reuse on these types of applications. Thus many applications run at memory access speeds. Memory speed is not improving as rapidly as processor speed, so a major concern for the future is a continuous degradation in the execution efficiency of codes even though processor speeds continue to increase. ***Compiler technology, memory-to-processor bandwidth, or program performance optimization tools will need to improve substantially over the next decade if we are to reach 20 TFLOPS sustained performance on these applications.***

III. Technology Cross-Cut of Gaps Identified

Increasing single processor performance efficiency to 40% will result in a 4x reduction in processors required, and thus a 4x cost reduction in the computing platforms NASA must purchase to meet these requirements.

Application scaling to thousands of processors is also a concern. NASA's current installed high-end computer base first reached that number of processors on a single machine at the beginning of this decade. Some heroic applications efforts have demonstrated performance on a thousand processors, but that is not the rule in the current environment. NSIPP is typically running on 16 to 64 processors. The DAO is operational on 100's of processors. Scaling to 1,000's of processors for these applications remains untested. The resolution increase projected over the next decade for the weather and climate applications is 1,000x, which will substantially help mitigate the scaling issue in their cases, provided processor interconnect bandwidth continues to track processor speed increases. The solid Earth modeling applications examined, however, have a different compute complexity than do the other areas, and the panel raised the concern that scaling of some of these techniques is going to require investment in new algorithms.

Scaling relies on continuing improvement in processor interconnect technology, but is dominated in practice by inefficiencies in the software layer that manages communication among processors, or memory conflicts on shared memory machines. ***Improvements in the data transport layers on both distributed and shared memory parallel computers (both hardware and software) at the same rate as processor speed increases is required if scalability of these applications to 1,000's of processors is to be achieved.***

I/O and Online Storage

Input/Output capabilities and online data storage volume are often overlooked requirements when application performance is analyzed. Indeed, none of the panels raised these as issues. We include the capability requirements analysis here for completeness.

Data volume input required for application initialization has the most immediate impact on application performance, since the application cannot proceed until the initialization is completed. Application output typically takes place throughout execution and could be overlapped with computation if the platform and operating system support this paradigm. If output cannot be done simultaneously with computation, then it will also impact throughput. Estimates of data transfer rates required for the weather forecast, climate S-I prediction, and earthquake prediction applications are gathered below:

	Input Data Volume	Output Data Volume	Sustained Bandwidth
Weather Forecast	1 TB/day	10 PB/day	200 GB/s
S-I Prediction	10's GB/run	100 TB/day	2 GB/s
Earthquake	100 GB/day	10 PB/run	200 GB/s ?

Online Data Storage and I/O Bandwidth Requirements

Again, the weather forecast system has the most stressing requirements. These have not been adjusted for the fact that the anticipated forecast cycle time is two hours, and assume that I/O can be overlapped with computation. These would increase by a factor of 10 if I/O were sequential to computation (assuming, then, that 10% of execution time is all that one would be willing to devote to waiting for input and output). These numbers also imply that online data storage (local disks) must be capable of managing tens of petabytes of data.

B. Data Management

All three panels specify the capability to manage petabytes (10^{15} bytes) of data as a requirement in 2010. The data results from NASA missions and from the modeling applications themselves. The overall factor of a 1,000-fold increase in model resolution over the decade directly translates into a similar increase in input data, output data products, archival storage, and transported data volume. It is the number one issue for the weather panel, and a driving concern for climate and solid Earth. The anticipated data explosion from ESE missions over the next decade presents both an exciting opportunity and a dreaded challenge for all three areas.

Each panel describes different use scenarios and constraints for their data problem. Weather forecasting has a data ingestion problem on a fixed time scale coupled with an output data products storage, analysis, and dissemination problem. Climate prediction has short-term data storage and analysis requirements, but is not stressing for archival purposes. Solid Earth anticipates a data rich research environment, but with sources widely distributed geographically. Each specifies the need to manage petabytes of data on a routine basis, with weather forecasting again being the most stressing.

Volume Data Management

The archive and retrieval problem is one aspect of the data management capability requirements. With projected data products produced over the course of a year being in the 100's to 1,000's of petabytes, these drive a data warehousing and distribution requirement that is quantitatively different than the data archive center problem being addressed in other technology programs. Indeed, the mission data archive centers will be required to feed data into the modeling activities at rates at least as fast as the data is collected. For weather forecasting, the data will stream in real time, possibly directly from the collection points. Climate modeling uses archived data sources\ but will need reasonable access speeds to prevent duplicating the data sources at the modeling site. The solid Earth research environment envisioned is completely distributed among federated servers with service brokers and transparent client access. To enable this vision, the data servers will have to provide low latency access. Output data generated is orders of magnitude greater than the inputs — but has different lifetimes depending on the discipline. Climate modeling output is mostly reprocessed almost immediately, with little long-term archiving requirement. Weather forecasts get archived and catalogued for later re-analysis, with a non-trivial data stream distributed to other users. In the solid Earth environment, the output data is ingested back into the data archives and catalogued for later re-access by other analysis applications.

Visualization of the catalogued data is another access mode that places unique access requirements on the data management system. Visualization for analysis purposes requires indexed, low-latency, high-bandwidth access to subsets of the data in a data warehouse system.

III. Technology Cross-Cut of Gaps Identified

Data rates for visualization purposes are discussed later but are not bandwidth drivers on the data management system.

Seamless, uniform, and standardized access to these data volumes is a major requirement.

With the sources of data so widely varied, any application that needs to access data from a variety of sources cannot possibly be expected to deal with the idiosyncrasies of each. The data management system must present a uniform interface to all of the data under its purview, even to the point of transparently translating among a set of standard formats. Other implementation approaches are possible — including imposing a standard representation format on the entire system.

	Observational Data	Access Modes Rates	Output Data	Storage Term/ Re-access Mode
Weather Forecast	1 TB/day Multiple Sources Continuous	Streamed input 20 GB/s	10 PB/day – Archival 10 TB/day – external distribution	Medium – Long Catalogued
Climate Modeling	10's of GB from archival sources	Data archive request 2 GB/s (latency tolerant)	100's TB/day	50% Short term - Immediate analysis 50% Medium term - Catalogued
Solid Earth Research	100's of GB/day Distributed sources	Distributed archives – low latency access	1 PB/day – ingested into distributed archives	Medium – Long Catalogued access

Data management system capabilities required

Network Transport Requirements

The data movement implied by the data access methods enumerated and the data volumes required will drive requirements on network layers between computing assets and storage assets, and end user clients. The order of magnitude bandwidth requirements are easily derived from the ingest and output data rates. There is an additional capability requirement that derives from the data volumes themselves, and is often overlooked. With petabytes per day of data moving between compute assess and storage asset, and terabytes per day moving through external networks, ***reliable and transparent transport services are required***. Failures must be automatically corrected. Faults must be detected and handled without user intervention. Current transport protocols do not scale to these data volumes.

Compute System to Short-Term Storage	Compute System/Short-Term Storage to Archival Storage	External Network Interfaces (
300 GB/s	150 GB/s	150 MB/s

Network bandwidth capacity requirements derived from the Weather Forecast strawman system.

Although the network bandwidth requirements appear daunting, these are aggregate capacity requirements. Multiple parallel data streams will likely be needed to accommodate these requirements. As long as the transport mechanism is transparent to the user/application, the details can be left to the implementation.

C. Programming Environments and Tools

Problem Solving Environments

In this era of ever-increasing complexity of scientific software, the ability to develop new software components and assemble existing components into new applications is viewed as a critical need across all three panels. Applications of the future will be composed from many separate models and data sources. It will no longer be the case that a single person can know every line of code in the application and know its operation in detail. Maintaining and upgrading such applications will require that the software be engineered and componentized and that it makes extensive use of re-usable components and libraries. Applications may be composed of geographically distributed services (e.g, data servers, compute servers) that are assembled for the task. The researcher will need to manage the complexity of developing applications within an environment that automates the process of design, implementation, testing, and optimization.

Two of the panels (Climate and Weather) directly cite the ESTO/CT Earth System Modeling Framework (ESMF) as a key technology to address such capabilities in the coming decade. Currently under development, the ESMF is expected to provide a software framework to enable high-performance interoperability of climate and weather applications from diverse institutions in a robust manner. ***Sustained government support for maintaining and development of the ESMF beyond its 3-year development project is imperative.*** The climate panel further stressed the need to extend the ESMF, especially in terms of data structures, operations to support assimilation, and integrated tools to enhance model development and utilization.

The DOE Common Component Architecture is cited as another example of the emphasis placed on this kind of technology.

Although the ESMF is not directly addressing the needs of the solid Earth (SE) community, the SE panel envisions virtually identical issues driving the creation of some sort of Problem Solving Environment (PSE), which is perhaps a bit broader in scope than the ESMF. The PSE include a prototyping environment for developing a model in a month; a modular framework for solid Earth applications; a collaboratory, which includes teams of scientists and computational experts; visualization and data analysis tools, seamless computer access; and integration of data and multicomponent models.

Portability

Achieving high performance on a single platform is limited in value if that performance cannot be maintained when an application is moved to a different platform. Hardware architectures continue to change and evolve on short timescales, and applications move from platform to platform as they move from researcher to researcher, or researcher to operational system. ***Thus, portability without significant loss of performance is a capability required by all three panels.*** Efforts such as the ESMF are expected to be major improvements in the modeler's ability to move applications among platforms, but this effort is not considered sufficient for even the climate modeling community.

Portability is inherently built on standards (programming language, messaging interfaces, and common libraries). The current major impediment to portability is that optimization done on one platform does not carry over to a different platform. This problem has been with us since the beginning of scientific programming, but is more pronounced today because of the wide range of parallel computing architectures in common use.

Libraries

A key enabler of portability is the availability of standard common libraries across all of the platforms of interest. Libraries that are optimized to the specific platform while maintaining the same interface across platforms are required to amortize the code development investment in each application. The current diversity of architectures — especially the difference between the tightly coupled global shared memory systems and the loosely coupled cluster computers — makes the common optimized library problem a difficult one. The best example currently available is ScaLAPACK, the parallel version of the popular linear algebra package LAPACK. It is built on the universally available message passing programming model, and is thus available on almost all platforms. However, it is not necessarily optimized for any platform.

Performance Tuning Tools

There is a general concern among the panels that the gap between peak and sustained performance on commodity technology is widening and significantly lower than on previous vector architectures. One panel even went so far as to identify a risk that vendor focus on peak performance may result in zero growth in sustained performance. Realignment of peak and sustained performance would be of broad benefit across many disciplines and would directly decrease the hardware cost for scientific investigations.

Given the variation in processor architectures, memory organization, and processor interconnect bandwidths and latencies, the application developer is currently left to his own devices to design and optimize his application. Current performance assessment tools are primitive at best. The result is that many applications run without any optimization at all, inefficiently utilizing the resources, and driving up the overall cost of maintaining the required throughput. Applications that are optimized have been made so by an expert — which is an added cost to the development and maintenance of the application. ***The ability for the average code developer to evaluate application performance on thousands of processors and understand what to do to improve it is required.*** With regard to parallel performance tools, the climate panel specifically suggested tools to (1) provide specialized applications greater control over shared memory usage and (2) reduce latency by providing specialized access to switching hardware.

Execution Management Tools

The complexity of scientific applications continues to grow not only in terms of development, but also in terms of utilization. Development of tools to enable effective use of complex models was suggested by all three teams. Solid Earth includes such activity within their PSE, while the climate panel suggests multiple tools for this issue. The climate panel calls for a graphical user interface for model configuration and job submissions/monitoring, and further indicates the need for a database system for tracking and annotating model experiments. With the transition to operational research environments where multiple application executions on several different platforms are coupled with output analysis of terabytes of data across those executions, the complexity of orchestrating the entire process becomes a productivity limiting process.

An interface that standardizes the process of assembling the required inputs at the required computing assets, automates the execution of the application (which may consist of multiple executables on multiple platforms), and manages the transport of the application output to its ultimate storage destination will be required to alleviate the burden on the researcher. This capability may need to be integrated into the problem solving environment, but would also be useful as a stand-alone environment as well.

D. Algorithms

Across all disciplines, potential improvements in algorithms offer hope for significant reduction in resource requirements or, alternatively, more/better scientific results from fixed resources. Different panels predictably identified different areas in which algorithmic improvement is most critical. The most extreme case is in the numerical weather prediction community, where the volume of observational data anticipated by 2010 is expected to overwhelm hardware improvements by many orders of magnitude. “Using the presently measured sensitivities of the model and assimilation algorithm to resolution and data usage, the projected computer requirements for 2010 are 10^9 times the current requirement.” Effective strategies for selecting pertinent data from the envisioned flood are imperative. History shows that as much improvement in throughput is obtained from algorithm advancements as is obtained from hardware improvements.

Algorithms are often thought of as science or mathematics instead of as technology. This is not the case today. The complexities of modern processor architectures, memory structures, and interconnect structures have direct influence on the implementation of any numerical representation of a physical model or simulation. The panel reports identify scalability of basic numerical algorithms like linear equation solvers, finite volume schemes, and spectral methods as capabilities needing to be addressed. The technology challenge is the design and implementation of numerical algorithms in a manner that allows them to execute efficiently within the application structure and also execute optimally on a variety of computing architectures.

The algorithmic improvements singled out by the climate community include improved parallelization on clusters, data assimilation techniques, and model assessment methods. The climate panel emphasizes the need to develop algorithms that minimize shared-memory conflicts. This is particularly relevant to the basic numerical methods used in solving differential equations: finite difference, finite volume, finite elements, particle methods, and spectral methods. *New implementations of these methods that can be adapted in individual applications to specific processor architectures and memory hierarchies (including shared and distributed memory implementations) while maintaining efficiency are required.*

E. Visualization and Data Mining

All three panels expect that significant technological advances must be made in the coming decade to maintain the key scientific ability to readily render new data in a form suitable for human consumption. The rapid increase in data (both numerical and observational) is expected to swamp existing, primarily serial, volume rendering tools. Although this issue is closely related to that of data management described above, there are some capabilities specific to visualization that are discussed. The solid Earth panel provided numerous specific target capabilities for the 2010 time frame: real-time volume rendering of 10^9 elements, visual data fusion from multiple datasets, co-registration and layering for 1 TB datasets, as well as support for distributed visualization of heterogeneous data from many (10^4) sources.

The current set of commercial visualization tools do not scale to terabyte datasets. And they are implemented with the assumption that the data to be visualized is local to the visualization engine. ***With the explosion of data, it becomes necessary to have visualization capabilities that can ingest terabytes of data from remote data sources, render arbitrary subsets and combinations of these data in real time, and deliver the rendered images interactively to the researcher at their local site.*** The networking capability of this requirements is expected to be less strenuous than the capability required for data management (see above), but will be a specific driver for “last mile” infrastructure at the point where the researcher views the rendered product. Migration to higher and higher display resolution capabilities is required over the next decade, perhaps by a factor of 100x increase in pixel count. HDTV was specifically mentioned by one panel, though the specific resolution capability required was not quantified.

Closely related to visualization is the capability to do data mining. ***Targets for data mining include pattern analysis for order 10^6 dimensions, wavelet analysis on the order of 10^6 scales, and inversion techniques involving 10^6 parameters.*** These requirements are mentioned specifically by the solid Earth modeling community. They are based on the need to analyze phenomena that span many orders of magnitude in scale. The climate and weather communities would also benefit from such capabilities, but have not stated explicitly any target capability requirements.

F. Distributed Computing

Distributed computing is expected to be driven by the need to coordinate data movement from many sources through many services to numerous consumers. The solid Earth panel is the main proponent of this technology, since their community consists of individual researchers who are widely scattered geographically. The Solid Earth Virtual Observatory concept is based on the availability of a distributed computing infrastructure that can provide transparent access and coordination of data sources (in the form of data servers), compute engines (which may in fact be distributed processing sites themselves), access portals, and a national (even international) network of high capacity and low latency. ***Seamless and uniform access, including reservation, allocation, and scheduling of geographically distributed assets, through a layer that enables low latency request servicing, is required if a useful distributed computing environment is to be enabled.*** Specific research and production applications would be built on top of such a layer. Currently, there is a national effort in Grid-enabled computing that will likely be the first step towards this capability.

G. Computing Platforms System Management

The weather panel specifically identified the problem currently faced by production-oriented computing centers in designing, implementing, and managing the infrastructure that will support real-time forecasting at the rate and resolution envisioned for 2010. Currently, this process is an ad hoc collection of expert individuals, hardware vendors and their sales staff, and center staff charged with keeping the place running. There are apparently no system design tools that can be employed in this process.

However, the field of system engineering is quite mature in other engineering endeavors. Any large production-oriented facility, from chemical and automobile manufacturing to airports and shopping malls, has a system or facility engineering process that is used to integrate disparate complex subsystems into a unified whole. A similar systems management approach to the design and operation of a computing facility that supports production computing like weather prediction is required if the operational goals of 2010 are to be realized. It may be possible to borrow and tailor tools from these other system engineering application to meet this goal, but some study is required to determine if this is feasible.

IV. Summary with Priorities

In this section, we summarize the capability requirements derived from the panel reports, along with their recommendations. The panels were asked to prioritize their recommendations into high, medium, and low categories (or must have/enabling, important to have/enhancing, or nice to have, but will not have any major impact). This prioritization was not done uniformly across the panels. The basic message from the workshop plenary sessions was that any capability identified in their reports was at least of medium priority.

Another general message uniformly communicated by each of the panels was that they would do their work on commercially available hardware that has vendor support, and that they do not believe that NASA has any influence on the direction that the commercial computing industry will take in the future. The hardware market, including the device technology that is integral to its capabilities, will be driven by forces outside of the government, and that researchers will do what they can to utilize the technology as it evolves. The community has many pressing science issues to address as their first priority. They require stable, commercially viable computing resources in order to address these science issues.

However, given that the scientific computing community does not drive any market, the panels recognize that some software technology investments will be necessary to enable them to use what the industry provides. This recognition is evident in the capabilities that have been identified as enabling or enhancing, which are, for the most part, software-enabled capabilities.

A. Gaps Expected to be Filled Without Investment

All three panels expressed the opinion that no amount of investment by NASA in any hardware technology would have any influence on the commercially available computing platforms of the future. Moore's Law is expected to continue governing the advancement in processor capability for the next 10 years. It also governs the increase in network bandwidth at the electronics layer. Thus, there is an expectation that the baseline component technology will provide a reliable 20x–50x in performance improvement during the next decade.

This is not the 1,000x in computing capability increase required for the stressing applications identified by the panels. But there is also a nearly uniform expectation that the commercial computing industry will turn the 1,000-plus processor systems that are experimental today into routinely delivered systems of the future — at least as a custom configurable option. A similar expectation for storage volume exists. These advancements are driven entirely by the commercial services market place, where transaction processing and data mining are driving the industry.

It is not possible to predict the future of computing technology with any accuracy. Video-on-demand has been touted as the next “killer application” for the Internet — but it has yet to get off the ground in any meaningful implementation. If such a market place develops, then high-bandwidth networking and large-volume, low-latency data storage access will naturally fall out of that market segment. If it does not, no amount of NASA investment will make it happen. (Some other commercial interest might spur the development, though).

Based on this analysis, it is more reasonable to list the technology areas that will progress driven by market forces and not by government investment:

Processors: 20x–50x improvement

Memory and memory bandwidth: 20x–50x density improvement, 10x bandwidth improvement

Network technology: 10x–20x bandwidth improvement, essentially no hardware latency improvement

Commercial parallel and cluster computers: built upon commercial components and optimized for non-scientific computing (probably transaction processing and data mining), configurable into systems with 1000's of processors and 100's of teraFLOPS peak performance.

Information display technology: Not predictable, outside of the gradual convergence to HDTV formats

Data storage: commercial magnetic media will continue to follow a doubling in density every two years, but bandwidth and latency will not significantly improve.

Given these expectations, the workshop recommendation is not to invest in these technologies directly. These should be assumed as the baseline, and investments should be made to enable the profitable use of these expected commercial offerings.

B. Recommended Investments

The required capabilities identified in Section III are listed below, along with technologies that potentially enable those capabilities. The technologies lists may not be exhaustive – new approaches to achieving required capabilities may arise over time. It is important to view specific technologies in the context of their ability to enable the requirements. Technologies should not be an end in themselves. This workshop recommends that any investment in individual technologies be made in a manner that assures that the technology development stays connected to and driven by the science enabling capability to be provided. Conducting the technology development in an environment that requires direct infusion into prototype science applications is the strategy recommended.

Sustained Throughput for Individual Applications — 20 TeraFLOPS

Vendors are expected to offer 20 to 50 GFLOPS processors, platforms with 10,000 processors, plenty of memory and storage. Gaps are in achievable applications performance.

Needs:

- Single-processor application performance at some significant fraction of peak
- Application scalability to thousands of processors
- I/O performance that scales with the application performance

Technology required:

- Performance optimization tools
- Compilers that achieve significant fractions of peak architecture performance (Programming language/paradigm continues to impact this ability)
- Scalable, portable algorithms
- Scalable operating systems
- Low-latency, high-bandwidth interprocessor communications

- Low-latency, high-bandwidth parallel I/O
- Memory management tools
- Performance assessment tools
- Efficient parallel programming paradigms
- Low overhead Operating Systems
- Thread management tools
- Highly scalable, highly efficient numerical libraries
- Low-latency remote data access (messaging or shared memory)

Seamless Uniform Access to and Management of Petabytes of Data

Vendors are expected to provide physical storage solutions. Gaps are in management, access, and distribution of the large data volumes.

Needs:

- Uniform, location independent service for identifying, managing, and accessing metadata and raw data
- Data transport performance that scales to consumer requirements
- Low-latency random access

Technology required:

- Internal computing center bandwidths @ 300 GB/s
- WAN bandwidths that scale to provide 150 MB/s for *every producer/consumer simultaneously*
- Reliable, location-independent data transport services
- Intelligent data caching
- New data organization and management applications
- Next-generation database tools for petabytes of science data
- Location-independent view and control of data
- Low-latency network transport protocols
- Transparently parallel data streaming

Programming/Problem Solving Environments

No specific vendor offerings are expected in this area.

Needs:

- Application frameworks/composable component architectures
- Platform independent program design and execution environment
- Highly efficient applications that scale to 1,000's of processors without heroic effort

Technology required:

- Discipline-specific frameworks and components (e.g. ESMF, SERVO)
- Performance tuning tools
- Portable, scalable components from which to build applications
 - Numerical libraries
 - PDE component toolkits
 - Data, information, and knowledge components
- Standardized program execution environment independent of execution assets

Distributed Computing

Vendor offerings in this area in 2010 are unpredictable. Currently, there is a multi-agency investment (NSF, DOE, NASA) in this area (Grid computing).

Needs:

- Uniform, seamless, transparent access and programming environment

Technology required:

- Universal standard middleware layer for all distributed assets
- Transparent data caching mechanisms
- Reliable high-bandwidth, standardized data transport layer
 - With automatic data format translation
- User single-entry point with global application execution control
- Distributed application program composition tools
- Distributed performance evaluation tools spanning the data sources, application services, clients, and networks

Scalable Algorithms

No vendor offerings are expected in this area. New implementation paradigms for algorithms are required that:

- Allow transparent incorporation into applications
- Maintain efficient execution on 1,000's of processors
- Automate latency tolerance
- Are completely portable

Technology required:

- Frameworks/component architectures
- New technologies not currently envisioned

Real-Time visualization of Terabytes of Data

Current commercial offerings will not scale to these data volumes. Future visualization systems must use vendor provided display technology. We expect a 10x display resolution over 10 years.

Visualization Applications must:

- Ingest terabytes of data from geographically distributed sources
- Render arbitrary combinations of diverse data sets in real time
- Deliver rendered product to the end user interactively at their local site

Technology required:

- Networks/data transport layers similar to the data management requirements
- Distributed visualization applications that scale to hundreds (perhaps thousands) of processors
- WANs with 150 MB/s bandwidth deliverable to the end user
- Low-latency service request protocols

Scalable Data Mining

This technology is currently in its infancy. New data mining applications for extracting science information require development.

Computing Platform Systems Management

Currently this is an *ad hoc* process. System engineering practice needs to be integrated into the computing center infrastructure design and operation. Building discipline specific computing centers needs to be supported by integration tools and systems management best practices. The only specific technology identified is:

- Performability analysis tools

C. Identification of Infrastructure Issues (not technology related)

The issue raised uniformly across all of the panels is the need for a commitment to infrastructure in addition to technology investment. All three panels project a need for installed sustained production computing throughput capacity in the hundreds of teraFLOPS for each discipline area — aggregated Enterprise wide, the required infrastructure capacity will need to exceed a petaFLOPS throughput. ***This is perhaps four orders of magnitude greater than the Enterprise has installed today.*** A similar increase in data storage and delivery infrastructure is also required, along with the networking bandwidth to support the volume.

The weather and climate panels both identified the need to support and expand the infrastructure that supplies production computing cycles and storage to the researchers who will develop the applications to achieve the prediction goals. The availability of these resources is one of the items critical to their success.

Appendix A: Weather Panel Report

Computational Requirements for Weather Modeling and Prediction

Robert Atlas, NASA GSFC, co-chair
 Ricky Rood, NASA GSFC, co-chair
 Tim Hogan, NRL
 Shian-Jiann Lin, NASA GSFC/DAO
 Wei-Kuo Tao, NASA GSFC

Arlindo da Silva, NASA GSFC
 Jim Dodge, NASA HQ
 Bill LaPenta, NASA MSFC
 Jeff McQueen, NOAA

Report from
 ESE Computational Technology Requirements Workshop
 April 30 – May 2, 2002

Introduction and Charge

The Panel was charged with evaluating the adequacy of computational capabilities needed for NASA to address its 2010 goals for enabling weather forecasts. *De facto* the charge of the panel was extended to a broader Enterprise scale, but focus was maintained on problems of weather and environmental hazards associated with weather. In addition, data assimilation to support research in other areas, such as climate and chemistry, were considered. The numbers presented in the tables are specifically targeted at the 2010 weather prediction goals, and these related activities would require computational resources to support them, but are not believed to require different capabilities. That is, the related activities require increasing the capacity of similar capabilities, perhaps with the addition of formal interfaces to transfer software and data, either observed or simulated, across the interfaces.

The following tenets are assumed:

- 1) There is no single specific programmatic or disciplinary investment that stands independent of other programs or disciplines. Namely, a balanced investment is needed in science, information technology, observing systems, *etc.* to achieve the prediction goals.
- 2) That many fundamental aspects of computational capability, processor speed, network performance, mass storage capabilities, memory, memory bandwidth, *etc.* will be determined by the evolution of the field as a whole and that NASA investment in the evolution of the capabilities will be limited to specific Agency-centric goals.
- 3) That significant, additional Agency investment in the infrastructure to support applications and system software is needed to assure that high-performance computational platforms are available to support Agency modeling and assimilation activities. The prediction goals addressed by this ESTO workshop represent only a subset of Agency goals that require high-end computing.

ESE Prediction Goals for Weather

The given prediction goals from the science plan are:

Today's Capability	2010+ Capability
3-day forecast at 93%*	5-day forecast at > 90%*
7-day forecast at 62%*	7–10-day forecast at 75%*
3-day rainfall forecast not achievable	3-day rainfall forecast routine
Hurricane landfall +/- 400Km at 2–3 days	Hurricane landfall +/- 100Km at 2–3 days
Air quality day by day	Air quality forecast at 2 days

The NASA role in weather forecasting is specifically a role of enabling. Therefore, NASA's activities are dependent upon the activities in other agencies. NOAA and the Navy both provide operational products. A formal mechanism linking NASA with NOAA and, soon, the Navy, is the Joint Center for Satellite Data Assimilation. The goals listed above were verified to be consistent with NOAA and Navy goals. These goals imply a link with both global and regional activities. The panel considers the given goals ambitious and believes that successful achievement of the goals depends on numerous scientific and technological developments whose success is not assured.

To succeed in enabling the Nation's weather prediction goals, NOAA and Navy operational requirements need to be integrated with NASA research activities. These include:


- 1) Techniques developed within NASA have to fit into the operational time allowance for real-time weather forecasting. This requires the end-to-to forecast-assimilation system to run in less than 2 hours of wall time.
- 2) Attributes of NASA algorithms, e.g., resolution, data usage, parameterizations, must be consistent with the state-of-the-art operational systems. Otherwise, the impact of NASA research will be diminished.

A number of specific computational technology issues are listed as priorities due to requirements to transition from NASA research to, specifically, NOAA operations. These include:

- 1) Testbed facilities: Results derived from NASA research activities need to be transferred with a minimum of retesting from NASA to the operational agencies. This requires sharing of software, and ultimately, experimentation in the operational systems. The current NOAA plan calls for the development of testbed facilities that serve to support experimentation and testing with the operational systems or candidate operational systems.
- 2) Modeling framework: The successful implementation of testbed facilities relies on the success, evolution, and maintenance of the Earth System Modeling Framework (ESMF) to provide a common model infrastructure.
- 3) Portability and tests: The ability to port applications software from one computational platform to another, and tests to assure that the portability is successful, is required. Otherwise, duplicate computational systems would be required.

Analysis of Problem: Forecasting System in 2010

In order to organize the discussion the following diagram will be used. This diagram represents the high-level functions that must be addressed by the forecast-assimilation system, highlighting the development areas and interfaces that most directly impact high-performance computational technologies. In general computational requirements are driven by resolution, data usage, comprehensiveness of coupling (running processes concurrently rather than sequentially), robustness (explicit representation of processes modeled as approximations or parameterizations), and data volume and data access.



The forecast-assimilation system is represented by the checked square and oval with an interface between them. These two components traditionally demand much of the scientific and computational development activities. The interface is specifically noted, as the frequency of communication between the modeling component and assimilation/analysis component profoundly impacts computational requirements. Projecting a target resolution provides the most straightforward parameter to estimate requirements growth in the model. A horizontal resolution for a global atmospheric model of 10 km is chosen based on the target resolution of both NOAA and the Navy. This resolution would resolve many features well enough for assessment of the regional impact of changing climate-scale phenomena. A vertical resolution using 100 levels from the ground to the top of the mesosphere is chosen. The computational requirement for the assimilation/analysis algorithm is currently driven by data volume. A data volume of 5×10^{10} observations per day is chosen based on projections of satellite observing systems that are planned to be operational and relevant to weather forecasting. As a baseline, the current horizontal resolution is order 100 km, with 50 vertical levels, and current data usage is order 10^6 observations per day. Using the presently measured sensitivities of the model and assimilation algorithm to resolution and data usage, the projected computer requirements for 2010 are 10^9 times the current requirement. The primary driver of this requirement is data usage. Such a requirement is unrealistic, and therefore, to address the data usage requirements new strategies are needed. These strategies require consideration of all aspects of the system.

Before these new strategies are discussed, the computational attributes of the other parts of the system are summarized. In order to incorporate already substantiated impacts of land-surface and oceanic processes on weather forecasts, coupling of the atmospheric models to other geophysical models is required. To assure that models address increasingly important issues of environmental security (for example, flash floods, coastal flooding, and air quality), the global models must be able to provide initial and boundary information for higher resolution models. These higher resolution models have more robust representation of local processes such as runoff and stream flow. Finally, since a number of weather-related products sit at the foundation of investigations of climate change, the ability to deliver consistent reanalysis of observations and provide regional climate predictions is needed. All of these applications require high-end computational resources, including application software support and formal interfaces between data systems and research groups. These interfaces are symbolized in the figure.

As a matter of emphasis, the throughput requirements of the application software are also noted in the figure. Because of the specific product requirements of weather and climate simulation and assimilation, if strict throughput is not maintained, then the results of NASA research is minimized. The throughput requirements for weather forecasting are easy to substantiate: the forecast must be useful in near real time. For climate research, the validation and provision of data sets to support regularly scheduled mandated assessments pushes the time criticality. Computational capability is highly sensitive to these time requirements.

Finally, the requirements that are discussed here are based on the idea of deterministic forecasting. Increasingly, probabilistic forecasts are being used and are likely to be a key part of achieving NASA prediction goals. Probabilistic forecasts rely on an ensemble of forecasts, generally performed at lower resolution than the deterministic forecasts described above. The requirements for probabilistic forecasts are shifted away from capability towards total capacity. However, the need to analyze perhaps 50 to 100 forecasts in a near-real-time environment increases the demand on networking and analysis software.

It was determined above that the data usage requirements, taken independently of any of the other requirements, provided an insurmountable obstacle unless new techniques are developed to utilize data and extract the information from the observations. This is needed to compact the information and, effectively, reduce the data volume, prior to assimilation. This is primarily scientific research and development. In addition, to the increase in data volume, there will be an increase in the complexity of the data environment. Historically, observations to support weather forecasting have been collected and distributed centrally through a small number of data hubs. As more use is made of research observations in the operational environment, observations will be obtained directly from the data systems that support the particular instruments. Three issues are the necessity to:

- utilize the increasing volume of observational data
- analyze and interpret the increasing volumes of simulation and assimilated data
- interface with more and more diverse input sources and user output

These issues lead the panel to the conclusion that data management is the greatest risk facing the meeting of the goals to enable weather prediction.

Numerical Estimates of Requirements

In the table below, numerical estimates of requirements are presented. These were derived by extrapolating the current DAO production system to future resolution and data use parameters. Considerations are made for increasing the comprehensiveness and robustness of system. Given that a straightforward extension of the current DAO assimilation system to future data volumes would yield an intractable problem, it is assumed that data and information management will yield a system in which the assimilation/analysis algorithm grows no faster than the model requirements. NOTE: These requirements do not express the entirety of the requirements to support all of the high-resolution applications or possible new initiatives that would build off of the global modeling activity to support the weather prediction goals.

	2002 System	2010+ System	
Resolution • Horizontal • Vertical levels • Time step • Observations ○ Ingested ○ Assimilated	100 km 55 30 minutes 10^7 / day 10^5 / day	10 km 100 6 minutes 10^{11} / day 10^8 / day	
System Components:	▪ Atmosphere ▪ Land-surface ▪ Data assimilation	▪ Atmosphere ▪ Land-surface ▪ Ocean ▪ Sea-ice ▪ Next-generation data assimilation ▪ Chemical constituents (100)	
Computing: • Capability (single image system) • Capacity (includes test, validation, reanalyzes, development)	10 GLOPS 100 GLOPS	Must Have 20 TFLOPS 400 TFLOPS	Important 50 TFLOPS 1 PFLOPS
Data Volume: • Input (observations) • Output (gridded)	400 MB / day 2 TB / day	1 TB / day 10 PB / day	
Networking/Storage • Data movement ○ Internal ○ External • Archival	4 TB / day 5 GB / day 1 TB / day	20 PB / day 10 TB / day 10 PB / day	

Panel Recommendations

The Panel's recommendations are made on the following background. First, for NASA to enable the weather prediction goals, priority must be given to partnerships with other agencies. This requires the algorithms developed through NASA research to ultimately fit into the real-time production window. In addition, credibility requires resolution and physical parameterization to be directly comparable to those used in the operational environment.

Second, it is tacit that the computational requirements to meet the weather forecast goals prescribe a set of systems attributes that are consumptive of capabilities that will be available

from the vendors. These requirements are broadly consistent with those of numerous other disciplines, though the real-time requirements may demand higher capability. In addition, the field will tailor the development of advanced algorithms to fit into the computational systems that are available.

Third, it is simple to make cogent arguments of systems that overwhelm projections of technological advancement. Numerous programs are in place to accelerate that advancement, notably the Department of Energy's Accelerated Strategic Computing Initiative. NASA should coordinate any activity it might plan to assure that NASA investments are effective. In addition, NASA needs to define its role in enabling weather prediction goals relative to the role of other agencies before research programs are initiated.

The computational requirements table above indicates that the greatest challenge facing weather prediction is data management and assimilation. We address the data management and assimilation challenge in Recommendation 1 below. The second greatest challenge is getting high performance (high peak performance and high sustained percentage of peak) from of the HPC systems that NASA purchases. The panel determined that achieving high performance will require investment in the software infrastructure aspect of computational technology, including systems software and engineering (Rec. 2), application-specific software tools (Rec. 3), and visualization tools (Rec. 4). Finally, it is apparent that the path to meeting the weather prediction goals will not be straight or easy.

Recommendation 1: Data and Information Management

Improved data and information management is critical. There are two major aspects to this. First, there are issues that would normally fall under the purview of information technology. These issues concern the development of computational systems that effectively acquire, order, and prepare data for use by either other algorithms or scientific users. Already, data management issues are coming to the forefront of issues that need to be faced by computational centers that support Earth science research. The increasing complexity projected for the future only magnifies the need to formally embrace the data management issues. Parameters that must be considered include:

- increasing number of sources of observational data
- distributed user communities involved in validation and application
- heterogeneity of observational types
- geographic or geophysical registration of observations
- data volume
- analysis in a fixed amount of time

Second, methods to improve the extraction of information from observational data prior to assimilation must be developed. This is termed the optimization of the observing system, and activities are far ranging. They range from methods advocated in computer science such as data mining and hierarchical segmentation, to sampling and data selection techniques used in current data assimilation systems, to increased utility of observing system simulation experiments (OSSE's) and observing system experiments (OSE's). These final approaches run end-to-end assimilation systems to predict and validate data impact. Increased interaction between the modeling and assimilation components of the data assimilation system to target which observations might have the most impact provides a potentially fruitful approach with broad implications for computational systems.

Third, new data assimilation algorithms must be developed that are quasi-independent of data volume.

Several processes could be used to elucidate the systems requirements more quantitatively. These include:

- the data ingest system of the DAO assimilation system
- examination of the validation process for a candidate forecast-assimilation system
- development of a scenario for a real-time ensemble prediction system
- development of a scenario to provide global products to a series of high-resolution regional applications

Recommendation 2: Systems Software and Systems Engineering to Support High-End Computing

In order to assure that the increases in technological capability that are achieved by the field as a whole benefits NASA applications, direct investment in systems software and systems engineering is needed. This investment in such a specialized field requires significant planning for organization and integration. The cross organization development of standards to support vendors, applications and discipline software is essential. Numerous standards efforts are in place and more are needed. Given the requirements, insufficient resources are expected, and much will depend on cooperative development with government and industry. Further, the significant funding levels will require continuity to ensure useful products. A real concern is that these expenditures will not represent influences on the marketplace. Unlike commodity hardware, the use of discipline-based software requires a focused design and implementation basis with extensive planning.

While some of these tools may be generic in nature, *i.e.*, discipline-independent, experience shows that adaptation to the simulation environment is required. The magnitude of needs for the development of software tools and systems engineering tools in the next decade far exceeds a simple enumeration of products. The supply of current technology tools is restricted to a few viable commercial suppliers and the open source code community. This limited availability of tools is not expected to change and needs addressing by the science agencies. The HPC vendors currently lack the willingness to support very large machine tools for financial reasons.

The capacity computing requirements will place burdens on the systems engineering tools needed to build teraFLOPS to petaFLOPS codes. If we assume the capacity growth is biased toward increasing numbers of processors, then we must realize the need for control and operation of thousands of processors for a single job. The list of useful support systems is quite long: an operating system, a compiler, a debugger, a mass file system, a visualization system, a program control system, a network analysis tool, a message passing data analysis tool, performance analysis, and an effective job submission and control program. The list of attributes for these tools is quite lengthy, and we shall not enumerate them, except to note the tremendous increase in data and comment on the need for interactive visualization. The large number of processors in use will require direct display of all the machine and user views.

Just as specific systems software investments will improve the performance of many or all ESE applications, so too will targeted systems engineering investments allow ESE to optimize the design and configuration of HPC hardware for the benefit of the entire ESE simulation environment. Determining the optimal configuration for the ESE environment will require

investigation for each major acquisition. This configuration investigation and optimization cannot be ignored or minimized. For example, ESE systems will need to balance communication to cache, between processors, and to mass storage.

Recommendation 3: Application-Specific Software Tools

The development of new application-specific software environments like the ESMF framework is an example of what is needed for the future. The weather development community, and the Earth science community in general, is moving toward the open standards efforts to support, data frameworks, visualization tools, and modular code environments. Current tools like IDL will be overwhelmed by the tremendous increase in data volumes of the next decade. The recasting of applications into truly parallel data warehouses with supporting tools is essential for the next level in data assimilation, modeling, and forecasting. We strongly recommend the design basis require the removal of all order-one bottlenecks in the end-to-end processing flows.

A final system software technology that will be needed for achieving the ESE weather prediction goals is the ESMF. The ESMF uses a common model infrastructure in the NOAA language, allowing applications to use a defined interface to interact with each other. A key driver of the ESMF is enhancing coordination of effort at the national level, enabling the U.S. to maintain leadership and in climate/weather prediction. Specifically, the ESMF will allow all of the Agencies participating in the U.S. climate modeling program to more easily divide and share the effort to advance the U.S. climate and weather prediction capability. IT will also facilitate a smooth technology transfer from research (such as that at NASA) to operations (such as those at NOAA).

Recommendation 4: Analysis-Specific Software Tools

The analysis of products by users will directly depend on the provided frameworks. These tool needs will encompass all the issues for hardware and software, and the development of tool products from a user centric position is essential for success. So, the user groups must be involved in the design, development, and re-design of the analysis-specific tools. Again, the lack of a commodity marketplace for developers will require NASA-specific investments to support discipline needs.

Visualization is an area where leveraged NASA investments can advance software and hardware computational technology for the pursuit of the 2010 ESE weather prediction goals. As mentioned, the data volume involved in weather prediction will increase by many orders of magnitude by 2010. As a result, the development of weather prediction software systems will inherently involve the analysis of huge data sets. This analysis will be interactive during the early development stages; the search for useful analysis or assimilation techniques must be conducted by domain experts. Humans have a sensory system that is highly evolved to transfer massive data sets to the brain: vision. Thus, visualization technologies will be central to developing ESE computation and data assimilation algorithms. In all stages of the program, visualization is a critical enabling technology for transforming experimental data produced by coupled Earth system models to information and knowledge. Visualization is also an essential tool for the communication of complex interactions to experts, decision makers, and the public.

Appendix B: Climate Panel Report

Climate Requirements for Advances in Computing Technology

Michele Rienecker, NASA GSFC, co-chair
 Al Kellie, NCAR/SCD, co-chair
 Brian Gross, GFDL
 Jeffrey Jonas, NASA GISS
 Tong Lee, JPL
 C.R. Mechoso, UCLA

Dave Bader, DOE
 Cecelia DeLuca, NCAR
 Ming Ji, NOAA/OGP
 Jim Kinter, COLA
 Tsengdar Lee, NASA HQ
 Max Suarez, NASA GSFC

Report from
 ESE Computational Technology Requirements Workshop
 April 30 – May 2, 2002

Climate (or more generally Earth system) research, including that for improved prediction, relies heavily on high-end computing resources. Significant advances in many of the key research areas will depend on advances in computing technology. Some of these advances, such as improvements in processor speed, can be anticipated as a surety. Others, such as software developments, need investment from the Earth science community to ensure that they are realized. Requirements for high-end computing resources come from a variety of subdisciplines: seasonal-to-interannual (S-I) variability and prediction, decadal variability and prediction, global change assessments, atmospheric-ocean-land assimilation for analyses (as well as that for forecast initialization), coupled physical-biogeochemistry modeling, to name a few.

For the purpose of outlining the advances in computing technology required to enable the Enterprise's science goals for 2010, we focus on the problems of S-I prediction and decadal prediction. This is justified by the fact that the types of problems to be solved and the models and resolutions used for these areas are not substantially different from those used for the other subdisciplines. The burden of initialization for these predictions lies primarily with the ocean and land surface, and since these components have far fewer observations than are available or needed for atmospheric initialization for numerical weather prediction (NWP), the computing technology requirements can be assessed primarily by those needed for the model simulations and predictions. The computing requirements for the other subdisciplines will be a (perhaps significant) fraction of those estimated for these prediction problems. They need to be taken into account when the Enterprise production computing resources are estimated.

The Seasonal-to-Interannual Prediction Scenario

In the 2010 timeframe, we envision that NASA will continue to contribute to the nation's multi-model seasonal predictions through forecasts and predictability experiments conducted by NSIPP. NSIPP will be one of approximately five national groups conducting routine experiments, sharing forecasts, simulations, and analyses. NASA's unique emphasis is on the optimal use of satellite observations to enhance prediction skill. The current focus is on 6-month prediction skill, and all groups undertaking such predictions with coupled general circulation models regard them as experimental. ESE's goal for 2010 is for 6–12 month routine seasonal prediction and 12–24 month experimental prediction. Increased emphasis will be placed on

larger ensembles as we attempt to extract more skill and shorter (intraseasonal) timescales and look to predict the likelihood of extreme event (e.g., flood) occurrences. With the expected increase in processor speed will come the desired increase in model resolution: 25 km for the atmosphere and 10 km for the ocean. The need to quantify forecast reliability will lead to larger ensembles at longer lead forecast time. This may be accommodated by multiple-scale ensembles, with degraded resolution at longer leads. Reliable estimates of forecast spread will require the uses of multi-model ensembles. The models will have expanded physics to account for influences on forecast skill outside the tropical Pacific – e.g., high latitude processes, stratospheric processes, and changes in land cover. As in other disciplines, attention to better cloud representation will be essential. The coupled initialization will continue to focus on the ocean and land. Since global predictions will be used to force offline regional hydrologic and crop models, global predictions will need to be stored (and retrieved) at high spatial and temporal resolution. We anticipate a successful transition to complete utilization of NASA’s Earth System Modeling Framework (ESMF) and this will enable flexibility in model component mix for ensembles.

The Decadal Experimental Prediction Scenario

To this point, predictions on the decade-to-century scale time frame have focused primarily on global or hemispheric scales; for practical purposes, regional scale predictions are required. Although confidence in regional predictions will probably increase in the timeframe to 2010, it is to be expected that predictions will remain on an experimental track for some time. NASA’s contributions to these goals will emerge from GISS, with plans for a $1/2^\circ$ ocean model, a 1° atmospheric model resolving both the troposphere and stratosphere, and air chemistry with about 40 tracers, if computing power is available. The same model will also be used for even longer global change integrations, possibly at reduced resolution.

Technology Areas

The advances in computing technology required to support the above prediction goals are summarized according to groupings that have emerged as common themes from other disciplines:

- 1) Computing platforms – the technology facet over which we have little control.
- 2) Problem Solving Environment – software developments that ease or optimize use of high-end computing architectures, or facilitate advances in coupled models.
- 3) Data management -- both hardware and software that have received little community attention to date, but which have the potential for significant development efforts. Without attention, this technology facet promises to be a significant impediment to realizing scientific advances from the other two.

We now further explore these areas and identify technology gaps that impact our requirements.

1. Computing Platforms

Current S-I climate prediction systems consist of ocean-atmosphere-land-sea-ice models coupled to ocean and land data assimilation systems. To be useful for climate applications, single image model performance must be roughly 1000 simulated days per wall-clock day, and since climate forecasts need to be based on ensembles of runs, typical systems run 8 to 10 concurrent coupled images (or 15 to 20 uncoupled images) now. Our assessment of requirements for future computer platforms is based on the assumption that as computer power increases, the size of the calculation and the number of concurrent images will also increase while roughly maintaining these performance goals (1,000 d/d single image, 20,000–30,000 d/d aggregate throughput).

Under the modest assumption of a doubling of processor speed every 30 months (i.e., slower than Moore's Law), the next 10 years should give us a 16x increase in processor performance. The actual increase we are likely to attain should be considerably higher than this because over the next 10 years we will be running higher resolution models that will scale to many more processors.

During the next 10 years, the magnitude of the calculation will increase due to the greater model complexity and increases in resolution. We have assumed that the latter will be the main driver of computer performance — this has been the case in both NWP and climate modeling for the last 30 years. Resolution may be increased in a variety of ways: by increased use of nested regions, by using embedded cloud resolving models, or simply by uniformly increasing resolution in global models. We assume that all of these approaches will result in comparable scaling properties and so take the last (a uniform global increase) to do a simple projection of performance for the next 10 years (see Table 1). We will also use the atmospheric model as characteristic of the entire problem. This is well justified for the ocean, land, and other components, such as chemistry, that will be a bigger fraction of the calculation in the future. It is not well justified for data assimilation systems, which depend more on data availability than on the resolution of the physics. Furthermore, we can be quite sure that data assimilation will play an increased role in climate prediction during the next 10 years. Nevertheless, we have ignored the different scaling of models and assimilation systems, partly for simplicity, and partly because we feel that computational requirements for assimilation will not rise to be the driver in climate problems that they are in the NWP problem — at least not in the next 10 years.

Resolution	Scaling PEs	Current Scaled Perf	5 years 4x	10 years 16x
2 deg	16	1000 2.5Gf		
1 deg	64	400 10Gf	1600 40Gf	
1/2 deg	256	160 40Gf	640 160Gf	2500 640Gf
25 km	1000	64 160Gf	256 640Gf	1000 2.5Tf
12	4K	25 640Gf	100 2.5Tf	400 10Tf
6	16K	10 2.5Tf	40 10Tf	160 40Tf
3	32K	4 10Tf	16 40Tf	64 160Tf

Table 1: Single Image Performance Scaling for the NSIPP AGCM

Table 1, then, represents our performance projections for this prototype calculation: an atmospheric GCM with uniformly increasing global resolution. The range of resolutions considered goes from that currently used for routine experimental prediction at NSIPP to a 3-km global model. This range takes us from models that resolve well the synoptic scales of weather phenomena to models that resolve the mesoscale. It also brings us to models that, while not cloud resolving, are sufficiently fine to require very different parameterization approaches. This range also takes us from the hydrostatic models we use today to ones that must represent non-hydrostatic dynamics. The hydrostatic–non-hydrostatic transition occurs around 12 km, while at 25 km we begin to resolve mesoscale phenomena, a transition that may result in qualitative improvements in AGCM results.

For these reasons, a realistic goal for S-I for the next 10 years is to begin to resolve the mesoscale, at or near the hydrostatic limit. We take this to be the 25-km model. Similar arguments can be made for ocean models reaching resolution of $1/10^\circ$ to $1/16^\circ$ — an increase comparable to that of the atmospheric model.

The requirements for decade-to-century prediction are similar in scaling, although more substantial in throughput because of increased complexity from the inclusion of atmospheric chemistry and resolution of the stratosphere. Although the atmospheric model will be slightly different from that used for S-I climate, the model and simulations will need to include realistic variability at interannual and decadal time scales. A single image of an atmospheric configuration of 1° resolution with 100 layers (including the full stratosphere) and 40 online chemical tracers would require a throughput of about 5

TFLOPS; with a coupled ocean of 50 layers at $1/2^\circ$ resolution the requirement is 6.6 TFLOPS. The coupled model configuration with a job mix equivalent to 100 concurrent images (ensembles, parameter sensitivity sweeps, etc.) would then require 660 TFLOPS.

As we can see from the table, the $1/4^\circ$ -model at S-I climate performance (1000 days/day) becomes available (under the table’s assumptions) at the 10-year mark. The conservative Moore’s Law assumption on processor speed (doubling every 30 months) should be met easily.

The other important assumption in these projections is that we be able to scale processor number with problem size out to about 1,000 processors even at the faster (16x) processor speeds. Assuming current architectural trends in the U.S. continue, this will require either much faster memories with better shared memory utilization within SMP nodes, or better communication between nodes, or — more likely — some combination of the two that allows scaling to keep up with processor speed.

The assumption that we can thus maintain scaling is quite optimistic given current trends. We should therefore view it as the largest technology gap in Compute Platforms and make it our highest priority. But what can we do to drive developments toward this goal?

Better shared memory utilization

Shared memory utilization has, of course, been a major goal of computer architects for many years. Developments such as ccNUMA show that it is still a major architectural goal, but one that we are unlikely to be able to drive. We can, however, concentrate on other parts of the problem. We can:

- 1) Promote the development of software tools that give specialized applications greater control over shared memory usage.
- 2) Organize our codes to minimize shared memory conflicts.

Most of our calculations require that only a tiny fraction of the memory be shared -- most memory accesses are local. Most shared memory usage is a matter of convenience, not necessity. (If this were not true, message passing codes would not be as successful as they are.) Our codes require a lot of memory accesses; if these are to a single shared memory, it results in a lot of memory conflicts and rapid scaling degradation. But if memory accesses can be segregated into shared and “private”, conflicts could be reduced, perhaps by orders of magnitude, and scaling greatly improved.

Better internode communication

Better communication can most easily be achieved with faster switches. Again we should do what we can to promote this development, but without any illusions of our ability to impact it.

Interconnect speed is measured in bandwidth and latency. Both will improve, but it is unlikely that improvements in switching hardware alone will match increases in processor speed. Latency improvement is generally deemed the more intractable problem. The latency impacting the scaling, however, includes both hardware and software delays, and the latter, which presumably we can do something about, still represent a large fraction of total latency for most applications.

Our priority in this area, therefore, should be the development of better software communication tools to reduce latency. This probably has to go beyond work on optimizing general-purpose message passing software, such as MPI, to providing specialized access to switching hardware for use by the highly optimized codes used for weather and climate prediction.

Better single-processor performance

In addition to communication, current throughput is limited by single processor performance. Most Earth system codes achieve only 10% of the manufacturer's stated peak processor performance even on a single node. In comparison, codes with even only minor hand optimization achieved 30 to 40% of peak on vector processors. Clearly this is a performance gap that needs to be addressed with better compilers for single-processor optimization.

Storage and bandwidth requirements

For the class of 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. Models typically produce 1 byte of output for every 1,000 to 10,000 floating-point operations. More precise estimates within this range will, of course, depend on the particular application, the model's resolution, etc.

For rough quantitative estimates of bandwidth and storage, we can take the most data intensive end of the range (1 byte / 1000 ops.). Under this assumption, the GFLOPS numbers presented in Table 1 can be reinterpreted as MB/second of output I/O bandwidth. Thus, the climate model ten years from now, running at 2.5 TFLOPS will produce 2.5 GB/second (or 200 TB/day) of I/O bandwidth at the high end of the range and 20 TB/day at the low end.

This volume of data is typically used only for immediate analysis, visualization, and various post-processing data reductions. The fraction retained for long-term storage is typically only between 1% and 10% of this. From this and a typical lifetime of 6 months in archival, we can estimate a total storage requirement of 500 Tbytes per image or several petabytes per system.

2. Technology for the Problem Solving Environment

NASA's Earth system prediction goals require research and operational communities that are capable of working efficiently and systematically to configure, run, and monitor simulations, exchange and introduce model components and parameterizations, track model experiments, and analyze, store, and retrieve data. The computing environment impacts the time and effort required by scientists to develop new models and evaluate and compare physical parameterizations. This environment is already daunting due to the complexities of computer systems, which we can expect to persist, and the complexities of applications, which we can expect to increase. The latter will be affected by the addition of new processes and components to existing modeling applications and enhanced component interoperability delivered by efforts such as the NASA Earth System Modeling Framework and the DOE Common Component Architecture project. Thus a computing environment that enables easy experimentation accelerates progress towards prediction goals.

The ESMF will work towards improving the Earth modeling environment by delivering software infrastructure that will buffer researchers against architectural uncertainty, increase interchangeability, and promote reuse. Assuming that the ESMF software is adopted by significant portions of the modeling community, it is essential that the ESMF be maintained and extended to support related community needs. It is essential that one or more agencies commit to long-term ownership and support of the ESMF in order to provide a stable software environment. Given that the major climate modeling groups requiring high-end computing technology cut across different agencies, it is important that essential infrastructures like the

ESMF receive interagency commitment to ensure that the framework longevity is adequate to warrant full implementation in codes.

Some of the basic computing environment capabilities that are currently lacking and that no current NASA program is committed to provide are:

- 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)
- 5) Extension of interoperability tools to physics parameterizations and new components
- 6) Extension of the framework to include differential operators, solvers, and other mathematical operations as well as operators relevant to innovation calculations for assimilation

In order to be most effective, the above services 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.

3. Technology for Data Management

The requirement for advances in data management can be easily articulated through the S-I prediction application. The complexity of data management for S-I prediction will be compounded in the 2010 time frame by the increasing dimensionality of the problem. Useful S-I climate prediction is necessarily probabilistic and therefore requires an ensemble of realizations to measure the uncertainty in the model integrations and hence the reliability of the forecasts. The use of large ensembles can also provide an indication of the risk of extreme events. Furthermore, there is a requirement to make use of multiple dynamic models with different physics to fully characterize the probability density function of observed states of seasonal climate. Finally, S-I predictions are being made by several groups within the U.S., so the dynamic model output will be stored at several computing facilities across the country. In order to be able to compare, analyze, and assess the forecasts produced by multiple models, each being run over multiple realizations at several physically distinct locations, it is necessary to enable data management and distributed access to and analysis of this virtual multi-model archive. The data management problem becomes significantly more challenging by 2010 by the increasing resolution and complexity of the individual models.

We have estimated that, by 2010, each group making experimental S-I predictions will generate about 500 TB of data for each image and 5 PB per 10-member ensemble experiment. This estimate is based on the discussion above on I/O bandwidth and the expected increase in model resolution, as well as the fact that forecasts must be verified over several decades of historical cases to quantify the level of forecast skill. The total amount of data per experiment to be shared among, say, five prediction research groups, is about 25 PB. A data set of this size will be generated over a period of about 6 months and will serve as a viable research data set for about 3 years. A typical problem, for example, the evaluation of extreme events over a single season, amounts to about 30 TB data volume for each group, but volumes could be much higher. Data from each group must be available for analysis and comparison to all the other prediction

research groups as well as to forecast evaluation groups and collaborating research groups elsewhere. The required capability to transfer 30 TB of data from any of the five centers to any of the others in less than one day represents a significant technology gap.

This anticipated evolution of S-I prediction systems leads to a technological requirement for high-performance analysis platforms, optimized for I/O performance, at each location where model output data sets are stored and a high-performance network connecting the participating centers. Some combination of distributed analysis capability and high-performance data transport capability will be needed. In addition, there will be a requirement for stable software that enables distributed access and analysis to the data set. Here, stable software is intended to mean software issued in stable upward compatible and tested releases that only require standard compilers, do not depend on a large number of potentially unstable and unsupported software packages from a variety of sources, and conform to community standards for code and data structures. Development of complete database documentation, automated catalogs that provide a minimum level of performance and reliability, automated metadata generation, and parallel data exploration and extraction tools will be required.

The capability for distributed model output data set access and analysis must be supported for the long term. This implies support for the storage, computational capability to do analysis, and network access at each site. It also implies that whatever strategy is followed for stable software, interfaces, catalogs, etc., has the same community-wide applicability and support as does, for example, the ESMF and an interagency commitment for support.

Summary: Technology Gaps of Highest Priority

Several developments are needed to satisfy the high-end computing requirements for climate prediction. Some of these can be achieved by (in fact require) an investment by NASA. Those of highest priority are:

Computing Platforms

- Internode communications and shared memory utilization
- I/O bandwidth and storage
- Fortran compilers to improve single-processor performance

Problem Solving Environment

- Commitment for sustainability of the ESMF
- Increased functionality of the ESMF, especially in terms of data structures, operations to support assimilation, and integrated tools to enhance model development and utilization

Data Management

- High-performance networks
- High-performance analysis platform optimized for I/O
- Distributed data archive
- Stable software for storage and analysis
- Parallel extraction tools for data mining
- Parallel and distributed analysis and visualization tools

Appendix C: Solid Earth Panel Report

Solid Earth Summary

Andrea Donnellan, NASA JPL, co-chair
John Ries, UT Austin, co-chair
John Rundle, UC Davis, co-chair
Geoffrey Fox, Indiana U.
Jay Parker, JPL

Robert Crippen, NASA JPL
Eric DeJong, NASA JPL
Ben Chao, NASA GSFC
Weija Kuang, NASA GSFC

Report from
ESE Computational Technology Requirements Workshop
April 30 – May 2, 2002

State of Solid Earth Science

The field of solid Earth science is currently mission and data poor. We anticipate, however, 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 huge quantities of data and information to be collected into forecast models will require that information technology resources be developed in concert with advanced sensor and detection capabilities.

The quantification of the computational needs is timely in light of the recent activities of the NASA Solid Earth Science Working Group (SESWG). This group is putting together a vision for solid Earth science for the next 25 years (see <http://solidearth.jpl.nasa.gov>). The group has identified the following six science questions:

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

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, and 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 concurrently 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.

Computational Model for Earth Observations

One of the major problems facing scientists today is that the scientific data volumes are increasing at a faster rate than computational power, challenging both the analysis and the modeling of observations. Resources must be put into improved algorithms to simplify processing and to approximate complex phenomena to allow researchers to handle the large volumes of data as well as find the dominant physics in a given data set. Another promising approach to handling large data volumes is to use pattern recognition to focus attention and point out subtle features in the data.

Earth Science is experiencing a major paradigm shift as it enters an era consisting of recursive cycles of: “Monitor-Model-Assess-Predict.” Each new Earth Science mission contributes datasets that increase the value of those previously acquired. New missions will generate 10 to 20 TB per week. The scientific legacy and long-term value of these data results will depend on the ability to handle complex queries over data sets from multiple instruments and multiple missions.

We propose the development of a Solid Earth Research Virtual Observatory (SERVO), which will enable investigators to seamlessly merge new multiple data sets and models; and create new queries. SERVO will support new discoveries and predictions for decades to come. SERVO will benefit from technology developed for the National Virtual Observatory (NVO) and a post-Human Genome Project “Digital Life portal.” Simulation data must be archived together with analysis/animation tools and the original simulation code.

Because of the complexity of the solid-Earth system, high-performance computers are essential requirements for scientific progress. Computations of the systems being studied, from the geodynamo to interacting fault systems, take weeks to months to run on even the most capable of current workstations. Considerable computational technology development will be required to make modeling an effective tool.

The new solid Earth missions will continuously measure parameters in space and use preprocessing to reduce data as much as possible for download. The downlink data transfer requires massive online data storage. The collaborative processing of this Earth science data with Earth-based measurements will transform the measurements into useful derived products using assimilation processes similar to those used for weather and climate.

The data storage, the science based modeling building, the analysis of data, and the timely use of derived products encompass the computational technology needed to achieve the goals set forth by the SESWG. SERVO, the Solid Earth Research Virtual Observatory, will support the many numerical laboratories proposed for meeting the solid Earth science objectives. These numerical laboratories are problem focused, and have basic computational resource requirements. The primary requirements are: data storage, computation, cognitive visualization, and computational codes. The needs statements for these numerical laboratories are focused on vendor hardware offering for resources. These expected vendor capability improvements have been and will be driven by fundamental industrial processes. The expected gains in the next 10 years approach one and a half orders of magnitude. However, to meet the science goals the hardware tools fall short of the program needs, and we must address the science goals by committing intellectual resources to attack the computational problems. History supports the notion that there is as much to gain from program and code improvements as from faster hardware.




Figure 1. Solid Earth Research Virtual Observatory (SERVO) and associated computational framework for supporting solid Earth science missions, data, analysis, and modeling.

Numerical Laboratories in the Support of Solid Earth Sciences

The physical processes associated with the solid Earth take place on many scales of space and time. Simulations and theory must account for how these many scales interact. The solid Earth is complex, nonlinear, and self-organizing. Recent work suggests strong correlations in both space and time resulting in observable space-time patterns. Recent advances in computational science and numerical simulations enable studies of the complex solid Earth system making it possible to address the following critical scientific questions:

- 1) How can the study of strongly correlated solid Earth systems be enabled by space-based data sets?

- 2) What can numerical simulations reveal about the physical processes that characterize these systems?
- 3) How do interactions in these systems lead to space-time correlations and patterns?
- 4) What are the important feedback loops that mode-lock the system behavior?
- 5) How do processes on a multiplicity of different scales interact to produce the emergent structures that are observed?
- 6) Do the strong correlations allow the capability to forecast the system behavior in any sense?




Figure 2. The physical processes associated with the solid Earth take place on many scales of space and time. Simulations and theory must account for how these many scales interact.

Three-dimensional solid Earth modeling for the gravity, geomagnetic, high-resolution visual, and multispectral thermal data is compute intensive. The development of useful problem solving environments (PSE's) should focus systems approaches to data management, computing and visualization. The management of data must meet the accessing time requirements and the long-term storage requirements. The maximum daily input data flow of 100 GB defines several important parameters. The two most important examples are the retention storage and the routine analysis for quality control.

In 2010, we anticipate that multiple solid Earth missions will be flying. Petabytes of data will be gathered every year in a widely distributed fashion. Widely distributed scientists using widely distributed computational resources will analyze these data. There will be a growing need for the incorporation of information from multiple sources on multiple scales into an integrated analysis. The goal is to have worldwide computational systems supporting the gathering, integration, visualization, simulation, and interpretation of several petabytes of data per year.

Achieving the above goal requires the following:

1) Developing a Virtual Observatory

Such a virtual observatory will allow for seamless access to large distributed volumes of data. This is essential for handling the numerous distributed heterogeneous real-time datasets. Data handling and archiving will be part of the framework. It will include tools for visualization, data mining and pattern recognition, and data fusion.

2) Developing an Earth Science Problem Solving Environment

The problem-solving environment will allow for model and algorithm development and testing, visualization, and data assimilation addressing the NASA specific challenges of multiscale modeling. Problems developed within the environment will be scalable to workstations or supercomputers, depending on the size of the problem to be solved. It will have compatible numerical libraries existing within a compatible framework. Algorithms within the framework will include PDE solvers, adaptive mesh generators, inversion tools, fast spherical harmonic transforms, wavelet analysis, particle dynamics, ray tracing and visualization preparation, and image processing and spectral analysis. It shall also provide a mechanism to facilitate teams of scientists (within and across disciplines) and IT specialists on framework design and development.

3) Improving Computational Environments

Many of the solid Earth problems are extremely computationally expensive, beyond the capability of existing supercomputers. Therefore, the capacity of high-end computers must be increased to petaFLOPS with terabytes of RAM. Distributed and cluster computers should be available for decomposable problems, rapid development and cost performance. The capacity of networks must be increased, and investments should be made in the development of Grid technologies. Computational analysis would also be enabled by more open and seamless access to computational resources.

Current and Anticipated Space Missions

The global nature of solid Earth science requires a wide variety of observation types, from *in situ* seismic arrays along faults to constellations of satellites. Increasingly we recognize that any understanding of the Earth will be incomplete if it fails to characterize, understand, and predict the interactions among the different components of the Earth system. Only with an interdisciplinary and global focus can we hope to succeed in understanding the dynamics of our planet. Space observations are uniquely suited to capturing infrequent events in their full richness, including consequences and possible precursors. Major earthquakes and large volcanic eruptions may be relatively rare in any given region, but on a global basis they occur comparatively frequently. To understand such events, which can have enormous impact in terms of both human life and economic disruption, continuous observations on a global scale are required so that these events may be studied completely. Space-based observations enable measurements not only of catastrophic events, but also of the more gradual motions that lead up to them. By measuring the full cycle (of earthquake, volcanoes, and other solid-Earth processes), a complete picture of the solid Earth is drawn.

Several missions are anticipated within the next decade to address these needs, which will produce greatly increased volumes of data (see Table 1). These include gravity field

determination missions such as GRACE, GRACE follow-on, GOCE; radar interferometry missions such as ECHO; CHAMP and other magnetic field determination missions, ICESat, a laser altimeter mission; GPS data collection from hundreds of ground stations (SCIGN/PBO) as well as space platforms; and a wide variety of hyperspectral, high-resolution visual, multispectral thermal, and lidar imaging missions, such as VCL.

In order to estimate computer resources needed for meeting the goals of advancing each science area and substantially enhancing forecast of natural hazards, we lay out a strawman data processing scenario for one of these missions. The components of this scenario are not at operational readiness, but are topics of promising research and serve as examples of methods under current consideration that suggest plausible computational loads.

For ECHO, there will be 100 GB of data per day, which should be fully utilized for forecasting events on the global fault system. We assume missions operations has processed raw data, precise orbits, optimal focusing, and phase connection to produce data in the form geographically registered images of reflectance and phase, and particularly of line-of-sight deformation data, affected by water vapor variability and other errors. Hence before inversion there needs to be a preprocessing stage using multiple image averaging (to deal with water vapor variability), including local weather, resulting in deformation with some characterization of error covariance. This stage is not highly demanding of operations, (roughly 32 global “day” images, 10^{10} numbers per image, and 20 operations, come to 0.1 GFLOPS sustained computing), but taxing for I/O and data transfer rates.

Given the complexity of full forward modeling codes, it will be advantageous to make a quick and rough inversion in terms of earthquake fault slip and change to other deformation sources. An elastic half-space inversion for these sources naively might entail the full daily data set, 2000 iterations, 10^4 sources and 100 greens-function related operations, hence 200 T FLOPS. But the iteration process can proceed from initially averaged data and other approximations, and focus on regions of most interest, so this may be reduced to 20 GFLOPS (sustained).

Accuracy may then be further refined in regions of interest by applying finite element techniques using known variations in rheology. Current work suggests 16 million finite elements can adequately characterize a complex region of high interest, such as southern California. Multiple runs involving iterative optimization techniques indicate the need for 0.1 TFLOPS rate per region of interest, which we take to be 80; but a single day of observation would cover 10 of these on average. Hence this stage requires 1 T FLOPS sustained.




Figure 3. Processor time increases linearly with mesh refinement. Adding complexity increases the unknowns and solution time is quadratic in unknowns.

For forecasting, we require a dynamic model that can span scales larger than a finite element code. Boundary element methods combined with good fault-friction models can fulfill this need in principle, although substantial testing and development remain. Such a system can be coupled efficiently with the finite element regional models; we do not estimate the coupling overhead. Operations count can be estimated from assumptions of regional cost, plus assumptions of multipole coupling between regions; most of the cost is then in the regional interactions. So, with 80 regions covered at 10 per day, 10^5 fault patches per region with hence 10^{10} interactions, 100 operations per interaction, 1,000 update steps for data-driven corrections (including recent history), and 0.1 computational efficiency, the requirement is 1 TFLOPS sustained.




Figure 4. A Northridge class simulation with 100,000 unknowns, and 4000 time steps takes 8 hours on a high-end workstation. Modeling the southern California system requires 4 million unknowns thus taking 12,800 processor hours for 1 km resolution. 0.5 km resolution would take 100,000 processor hours or 400 hours (17 days) on a dedicated 256-processor machine.

For better forecasting, eigenpattern rate techniques are needed. For each region, the 10^5 fault patches imply 10^{15} operations to diagonalize local interactions. This may be reduced by focusing on the primary modes, suggesting Lanczos methods. Direct methods indicate 0.1 TFLOPS sustained.

Finally, these pattern rates will be used to daily update an earthquake hazard map. This may be done with vector-oriented updates, and probably does not amount to more than 0.1 TFLOPS.

End to end, processing ECHO for earthquake forecasting comes to about 2 TFLOPS sustained rate, implying 10^4 petaFLOPS/year.

Table 1. Data Volumes from Observations

Mission	Onboard Volume	Derived Product
GRACE	50 MB/day onboard	8 GB/day derived product
ECHO	100 GB/day onboard	
SRTM	12 TB raw data (mission total)	
ICESat	1 GB/day onboard	2 GB/day derived
SCIGN	250 MB daily; 7.5 GB/day for real time	
Airborne observations	LIDAR	
VCL	2 GB/day onboard	4 GB/day derived
Hyperspectral imagery	100 GB/day raw	
Imaging LIDAR	> 20 GB/day	> 40 GB/day

Table 2. Data Volumes Derived from Models

Model	Today	2010
Geodynamo model	10 GB of storage for one model run	1 PB/run with minimal need of 10 runs
General earthquake/lithospheric models	1 TB/run	10 PB/run (multiple scales combined in many regions)
Gravity	100 GB/run	2 TB/run
Mantle convection models	1 TB/run	10 PB/run
Geomagnetic field models	32 GB/run	300 GB/run

Table 3. Computational Environment Metrics

Item	Today	2010
Processes	Currently none on supercomputers	1 modeling run/day
Dedicated networks	1Gb/s sustained	50Gb/s sustained
Performance	100's GigaFLOPS	PetaFLOPS
Memory (RAM)	40 GB	5 TB

Distributed Data Access

An important aspect of data collection is to create distributed centers for storing unique data sets and developing the infrastructure to compare, analyze, and ingest complementary data sets, such as ice topography and sea level changes. Currently, cooperative federated databases do not exist; the data are heterogeneous and widely distributed in a variety of formats. Because of the high volumes of data, now and in the future, data mining and other approaches for interacting with the data must be developed (possibly including more onboard processing). A Solid Earth Research Virtual Observatory (SERVO) is recommended to address this critical need. Such a system would have the following characteristics: distributed data at centers and ground stations; thousands of sites with volumes of 1 TB to 1 PB; a multi-tier architecture for staging of the data; middleware to control integrity and versioning; support standards developed within the community; working within data grid projects; customized for key NASA solid Earth needs; and high-performance access of 100 GB files within 40 TB datasets within 5 minutes to user and program-to-program communication in milliseconds using staging, streaming, and advanced cache replication.

This system may be modeled on known centers, such as EOSDIS, the Physical Oceanography Distributed Active Archive Center (PODAAC), the National Virtual Observatory (NVO), and the World Organization of Volcano Observatories (WOVO) systems.

Modeling Complexity

Creating realistic simulations of solid Earth processes requires complex models. Multiscale and multicomponent modeling requires that algorithms be developed and/or reengineered for current environments and platforms so that various components can interoperate. At present, no appropriate environment exists for testing, calibration, visualization, and development. Because software programs have been developed on different platforms, it is difficult to merge them to integrate scales and components into a more comprehensive model. Real-time modeling poses a problem of both data and computational latency. The complexity of the problem requires high-performance computers to realize these integrated models. Both models and algorithms must be developed to efficiently make use of high-performance computers. Some algorithms are difficult to parallel requiring an investment in development of algorithms and computational methods. Adjusting the size of the problem to the resource is also difficult. We propose the development of an environment that recognizes the size of the problem to be solved that then farms the job out to the appropriately sized computer.

Our recommendations, in priority order, are to:

- 1) Develop a solid Earth science Problem Solving Environment to support 10 solid Earth sub-fields. This includes a prototyping environment for developing a model in a month; a modular framework for solid Earth applications; a collaboratory, which includes teams of scientists and computational experts; visualization and data analysis tools, seamless computer access; and integration of data and multicomponent models.
- 2) Development of improved parallel algorithms/applications (on the order of 100) with a scaled efficiency of at least 50% on large clusters; models and model assessment methods; and data assimilation techniques.

- 3) The computational needs required to address model complexity include 100 TFLOPS sustained rate capability per model, 5 TB total memory per model, and 10^4 petaFLOPS throughput (3 TFLOPS per second sustained) per sub-field per year.

Understanding the Data

Understanding the data is critical to meeting the science objective set forth by the SESWG. It is difficult to explore, visualize, and understand the data, and difficult to gain meaningful knowledge out of the large amounts of data, derived products, and model results. The multidimensionality of the data makes interpretation difficult. Another bottleneck lies in registering, map projecting, layering, and displaying the data.

We recommend an investment in visualization tools, methods, and standards for parallel platforms. These tools will include volume rendering in real time (10^9 volume elements), data fusion, coregistration and layering for 1 TB datasets, and support of distributed visualization of heterogeneous data from about 10^4 sources. Rendering the data will require specialized parallel graphics hardware/software to provide stereo HDTV at 120 frames/s. The visualization tools will make it possible to sweep through large amounts of three-dimensional data and model output. Three-dimensional models are required to model the very heterogeneous solid Earth including plate tectonics, interacting fault systems, volcanoes, mantle processes, and the geodynamo.

Visualization allows a researcher to build up an intuition of the processes being studied. The development of computational approaches for data mining and data understanding allow for routine tasks to be accomplished efficiently. They also can be used to point out subtle features in the data, and redirect the scientist's attention, allowing for a more thorough understanding of the physics. The data mining approaches include pattern analysis of on the order of 10^6 dimensions, wavelet analysis on the order of 10^6 scales, and inversion techniques involving 10^6 parameters.

We recommend a combined approach of using both traditional and computerized data analysis. Traditional analysis techniques useful in cases where lots of science knowledge is required, and human intuition and larger understanding are valuable. Computerized data understanding is useful in cases where the important information is subtle and may evade even an unaided human expert, or there is too much data for human "digestion." Putting the strengths of these two sides together results in the creation of computerized data analysis that can be used as a tool for scientists or data exploration and analysis, or can create a loop where scientists feed data into a system for analysis, look at results, decide on a new analyses to perform, feed data in again, etc.

Risks

A number of hardware and software risks were identified. Not all of these may be realistic, but they could not be casually dismissed. How should NASA build on the open software model? Will the marketplace continue to support fast floating-point operations on chips? Will performance FORTRAN support survive and if so, is that a good thing? Continuation of performance FORTRAN can result in the development of efficient codes as well as the support of legacy code. However, if the software packages are written in FORTRAN it may be difficult to attract the younger generation of well-trained computer programmers, who have been schooled in object-oriented programming such as JAVA and C++. Another concern is that coordinating a large team of researchers and technologists is difficult. Any future programs should foster collaboration among solid Earth and computational technologists.

Because we anticipate large data volumes in the next decade, it is possible that the data won't be fully utilized unless there is a concerted effort in creating tools to understand the large data volumes. Bandwidth limitations for both ground and space systems may result in critical data being lost. We are only now learning, for example, that an earthquake on one fault can have far reaching effects on the regional deformation field and on other faults. We must be able to freely explore all of the data to discover new processes.

There may be bottlenecks that limit performance but no tools to help identify limiting factors. Performance of future computation platforms is also a risk. It is possible that peak, but not sustained, performance will increase. Memory performance is constrained and is anticipated to be a bottleneck. To address these risks, it is important to develop high-performance algorithms that have a high reuse of memory, and also encourage innovation in the marketplace.

Roadmaps

Developing the computational environment to address the science goals of the SESWG requires a three-pronged but integrated approach. Investments over the next decade must be made to create a solid Earth research virtual observatory (SERVO), create a rich problem solving environment, and an efficient computational environment.




Figure 5. Roadmap for the Solid Earth Virtual Observatory (SERVO).

Problem Solving Environment Project




Figure 6. Roadmap for the problem solving environment.

Computational Environment




Figure 7. Roadmap for the computational environment.

Appendix D: Workshop Attendees

Workshop Organizers

Azita Valinia, NASA HQ
Robert Ferraro, NASA JPL
Jim Fischer, NASA GSFC

Workshop Attendees

Doug Muchoney, Conservation International
Karyn Tabor, Conservation International
Tim Hogan, DOD NRL
Dave Bader, DOE HQ
Forrest Hoffman, DOE ORNL
Roger Mitchell, Earth Satellite Corp
Massood Towhidnejad, Embry-Riddle Aero. University
Jerry Garegnani, ESRI
Yuechen Chi, George Mason University
Paul Schopf, George Mason University
Larry Fishtahler, GRSS
Bill Kneisly, IBM Corporation
Lloyd Treinish, IBM Thomas J. Watson Research Center
Alexey Voinov, IEE UMCES
Geoffrey Fox, Indiana University
Eungchun Cho, Kentucky State University
Phil Merkey, MichiganTech
Avichal Mehra, Mississippi State University
Stuart Frye, Mitretek Systems
Tom Ganger, Mitretek Systems
Thomas Passin, Mitretek Systems
Rob Van der Wijngaart, NASA ARC CSC
Douglas Ford, NASA ARC AMTI
Bryan Biegel, NASA ARC
David Gambrel, NASA ARC
Chuck Niggley, NASA ARC
Bob Shiner, NASA ARC
James Stegeman, NASA GRC
Joe Schmidt, NASA GSFC AMTI
Phil Bording, NASA GSFC CSC
Udaya Ranawake, NASA GSFC GEST
Gilberto Vicente, NASA GSFC GMU
Brice Womack, NASA GSFC Northrop Grumman/TASC
Shujia Zhou, NASA GSFC Northrop Grumman/TASC
William Putman, NASA GSFC SAIC
Leonid Zaslavsky, NASA GSFC SAIC

Thomas Clune, NASA GSFC SGI
 Spencer Swift, NASA GSFC SGI
 Rodger Abel, NASA GSFC
 Robert Atlas, NASA GSFC
 Ben Chao, NASA GSFC
 Arlindo Da Silva, NASA GSFC
 John Dorband, NASA GSFC
 Igor Eberstein, NASA GSFC
 Jim Fischer, NASA GSFC
 Watson Gregg, NASA GSFC
 Milton Halem, NASA GSFC
 Jeffrey Jonas, NASA GSFC
 Gi-Kong Kim, NASA GSFC
 George Komar, NASA GSFC
 Weijia Kuang, NASA GSFC
 Nand Lal, NASA GSFC
 Vladimir Lumelsky, NASA GSFC
 Karen Moe, NASA GSFC
 John Moisan, NASA GSFC
 Amidu Oloso, NASA GSFC
 Christa Peters-Lidard, NASA GSFC
 Michele Rienecker, NASA GSFC
 Richard Rood, NASA GSFC
 Ellen Salmon, NASA GSFC
 Siegfried Schubert, NASA GSFC
 Max Suarez, NASA GSFC
 Wei-Kuo Tao, NASA GSFC
 William Webster, NASA GSFC
 Warren Wiscombe, NASA GSFC
 Robert Wolfe, NASA GSFC Raytheon ITSS
 Latonya Alexander, NASA HQ
 Jim Dodge, NASA HQ
 Jack Kaye, NASA HQ
 Tsengdar Lee, NASA HQ
 Gran Paules, NASA HQ
 Jonathan Root, NASA HQ
 Azita Valinia, NASA HQ
 Robert Crippen, NASA JPL
 Eric DeJong, NASA JPL
 Andrea Donnellan, NASA JPL
 Robert Ferraro, NASA JPL,
 Daniel S. Katz, NASA JPL
 Tony Lee, NASA JPL
 Charles Norton, NASA JPL
 Y.Tony Song, NASA JPL
 Wendy Tang, NASA JPL

Tuyet-lan Tran, NASA JPL
 Michael Little, NASA LaRC CERES
 Jassim Al-Saadi, NASA LaRC
 Kuan-Man Xu, NASA LaRC
 Bill LaPenta, NASA MSFC
 Kirstin Riesbeck, National Computational Science Institute
 Jeffrey McQueen, National Weather Service
 Al Kellie, NCAR
 Cecelia DeLuca, NCAR
 Brian Gross, NOAA GFDL
 Ming Ji, NOAA
 Louis Kouvaris, NPP
 Soundararajan Ezekiel, Ohio Northern University
 Kevin Furlong, Penn State University
 Genyuan Li, Princeton University
 Nazmi Saleous, Raytheon
 Troy Anselmo, SAIC
 Michael Jones, SAIC
 Karl Lalonde, SD School of Mines and Technology
 Brian Hahn, SGI Federal
 Bill Ryan, SGI Federal
 Ilene Carpenter, SGI
 Bruce Kiracofe, Shenandoah Mountain Geographics
 Rafael Ameller, STORMCENTER Communications
 Chung-Min Chen, Telcordia Technologies
 C.Roberto Mechoso, UCLA
 Xu Liang, University of California at Berkeley
 John Rundle, University of Colorado
 Jim Kinter, COLA
 Antonio Busalacchi, University of Maryland, ESSIC
 Tim Dixon, University of Miami
 Allan Falconer, University of Mississippi
 K Palaniappan, University of Missouri-Columbia
 Gary Winchester, University of Pittsburgh
 Nazario Ramirez, University of Puerto Rico
 John Ries, University of Texas at Austin/CSR
 Byron Tapley, University of Texas
 Zhemping Liu, Virginia Tech