High-End Computing in Climate Modeling

Contributing Authors:
Robert C. Malone (Los Alamos National Laboratory),
John B. Drake (Oak Ridge National Laboratory),
Philip W. Jones (Los Alamos National Laboratory),
Douglas A. Rotman (Lawrence Livermore National Laboratory)

What is Climate Science?

The earth's climate and the study of the circulations of the atmosphere and oceans is one of the most intriguing investigations in all of science. It has received new attention as the skill at predicting seasonal events, such as the El Nino Southern Oscillation (ENSO), has been proven and as the nation faces compelling questions about the role of human activities in changes of the global climate system.

The earth environment consists of several interacting parts, including the atmosphere, hydrosphere (ocean), cryosphere (ice), and biosphere (plants, animals, and land surface). The science of climate encompasses physical, chemical, and biological processes operating in these domains with the objective of understanding how these pieces interact to produce the observed distribution of temperature, precipitation, plants and chemicals such as ozone and carbon dioxide. Climate science seeks to identify the most important processes in each component, make observations to quantify the operation of those processes, develop computer models of the major components, and integrate them into a single predictive model. Models are refined and validated by comparing model results with atmospheric, oceanic and terrestrial measurements.

The fundamental processes that form the backbone of the climate system can be summarized as follows:

- Energy from the sun drives the entire system by warming the tropical latitudes. This energy arrives in the form of visible and near-visible radiation from the sun. Once absorbed by the atmospheric constituents, clouds in the atmosphere, land or ocean surfaces and converted to heat, the heat is reradiated upward as infrared radiation.
- Infrared radiation ("heat") escapes to outer space predominantly at polar latitudes, despite the fact that it is much colder at the poles than at the equator. This occurs because the water vapor concentration is much higher in the tropics and acts to absorb and reemit (partially downward) infrared radiation rising from the atmosphere below and from the earth's surface (both land and ocean). This is the so-called "greenhouse gas" effect; water vapor is a very effective greenhouse gas.
- The circulations of the atmosphere and ocean are driven by the need to transfer heat from the tropics to the poles. The circulation patterns are made complex by the non-linearity of the governing equations, by the rotation of the earth, and by interactions of the flow with the non-uniformities of the earth's topography (which underlies both the atmosphere and the ocean).

The absorption of visible light and reemission of infrared radiation, cloud formation and optical properties, and solution of the "equations of motion" for the atmosphere are encapsulated in "atmospheric general circulation models" (AGCM). This class of model has been driven to a high degree of sophistication by the demand for accurate numerical weather predictions. The DOE program "Atmospheric Radiation Measurements" (ARM) is designed to measure upwelling and down welling radiation and relate those measurements of clouds and other optically active media such as aerosols. These measurements are used to develop better "parameterizations" of atmospheric processes. The counterpart models for the ocean are called "oceanic GCMs", and are less well developed, in large part because much less data is available on the real ocean. The ocean is vast, often inhospitable, and difficult to observe. Nevertheless, great strides were made in the World Ocean Circulation Experiment (WOCE), a multi-national, decade-long program to make comprehensive measurements of many properties of the ocean.

Unlike weather prediction, which attempts to forecast the day-to-day variability of quantities like temperature and precipitation, climate science focuses on the longer term "mean" or average state of the earth system. However, a "mean" climate state includes substantial variability even at longer timescales. Climate science includes understanding seasonal variability, interannual variability like El Nino, decadal oscillations like the Pacific Decadal Oscillation or North Atlantic Oscillation, and centennial changes such as global warming. The prospect of human-induced climate change has created a strong incentive to understand the entire climate system and its variability.

The National Academy report (June 6, 2001) titled "Climate Change Science: An Analysis of Some Key Questions" identified the need to "reduce the wide range of uncertainty inherent in current model predictions of global climate change" by producing advances in "the understanding and modeling of the factors that determine atmospheric concentrations of greenhouse gases and aerosols and the so-called 'feedbacks' that determine the sensitivity of the climate system to a prescribed increase in greenhouse gases." For policy makers to have the best information from climate scientists requires that leadership-class research programs be maintained within the national laboratory and university systems.

[Picture of Earth from Space, caption describing interacting oceans, atmosphere and biosphere at full range of scales]

Climate simulations based on the mathematical description of ocean and atmospheric flows coupled with chemical and biological process models are providing scientists with new insights into earth's complex climate system. The new coupled models are able to balance the fundamental physical quantities of energy, momentum, mass, as well as freshwater, salinity and chemical species among terrestrial, ocean, ice and atmospheric pools. The modeling of the global carbon cycle to better understand the interactions, effects and feedbacks in the climate system, is one goal of present scientific research. [Cartoon of Carbon cycle, caption explaining sources and sinks of carbon]

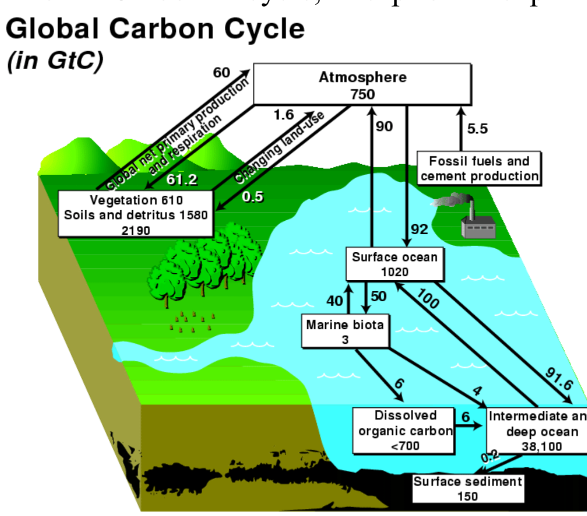

Impacts on Science and Society

An international consensus is emerging that humans are changing Earth's climate. It is expected that climate change will continue and even accelerate. Clearly, future climate change will have important impacts on many sectors of society, including agriculture, water resource management, energy production and demand, human health, and recreation. Natural ecosystems and biodiversity will also be affected. The cost of adaptation to climate change could be large and we must attempt to anticipate and quantify potential damage due to climate change. Adaptation strategies might reduce the damage but this approach will also have an associated cost.

[Picture of an ecosystem, no caption]

Greenhouse gases like carbon dioxide have long residence times, so delays in responding to climate change may dramatically increase costs and decrease effectiveness of mitigation strategies. Accurate long term predictions that include as many known feedbacks as possible will be required to evaluate the impacts of climate change and the effectiveness of emission-reduction scenarios and carbon-sequestration methods. Policy makers need such tools now. A better understanding of potential societal impacts is needed to properly weigh the costs of mitigating climate change (e.g. by developing new carbon-free energy sources or developing carbon sequestration technologies) against the costs of allowing climate change and its impacts to occur. Demonstration and implementation of carbon sequestration methodologies and new carbon-free energy production technologies will require decades to develop. Whatever policies are followed in the future, anthropogenic climate change will continue for decades and the explanation of observed changes will require high-end climate modeling.

Scientific Opportunities

Among the key questions facing climate scientists is the need to reduce the uncertainty in current model predictions of global climate change and to understanding the factors that determine the sensitivity of the climate system to a prescribed increase in greenhouse gases. Improved climate models are essential tools for gaining a more complete understanding of climate and impacts. In fact, climate models are the only means of integrating our knowledge of the components (atmosphere, ocean, land surface, and sea ice) that make up the complex climate system. And they are the only acceptable means for carrying out “experiments” on the climate system to study the projected changes and impacts of different scenarios. To be truly useful, climate models must make credible predictions on a regional spatial scale (e.g., within a state). Due to the coarse resolution and other limitations of today’s climate models, predictions are considered reliable only on continental and larger scales, but not on a regional scale. In order to make reliable region-scale predictions, climate models need greatly increased spatial resolution, improved treatments of sub-grid-scale physical phenomena (e.g. clouds), and inclusion of additional physical, chemical, and biogeochemical processes.

The Intergovernmental Panel on Climate Change (IPCC), sponsored by the United Nations, is a highly regarded multinational scientific body that performs extensive studies of potential climate change and publishes their findings on a five-year cycle. The IPCC is beginning to collect scientific results for the Fourth Assessment to be completed in 2007. The Community Climate System Model (CCSM2) has been developed as a multi-agency initiative with support from the National Science Foundation and the Office of Science in the U.S. Department of Energy. CCSM2 will be one of the primary climate models used in the next IPCC assessment. With an increase in dedicated computing resources from 10 Tflops to 100 Tflops, new studies could be performed with a higher-resolution atmospheric model providing much improved spatial detail

[**Sidebar: Why is increased resolution important?**
An important long-term objective of climate modeling is to have the spatial resolution of the atmospheric and oceanic components both at ~1/10° (~10 km resolution at the Equator). The reason for this objective for the ocean is that [3]mesoscale eddies[2] play an important role in maintaining the mean circulation. In fact, most of the energy of motion (kinetic energy) in the ocean is associated with these eddies. Simulations at various resolutions using the Los Alamos Parallel Ocean Program (POP) have demonstrated that, because such eddies have diameters ~10-200 km, the grid spacing must be < 10 km to adequately resolve the eddy spectrum. This is illustrated qualitatively in the set of four images of the sea-surface temperature (SST). Figure (a) shows a snapshot of SST from satellite observations, while the three other figures are snapshots from the POP simulations at resolutions of (b) 2°, (c) 0.28°, and (d) 0.1°. The narrow, meandering current off the coast of Japan is the Kuroshio Current; it is the North Pacific analogue of the Gulf Stream in the North

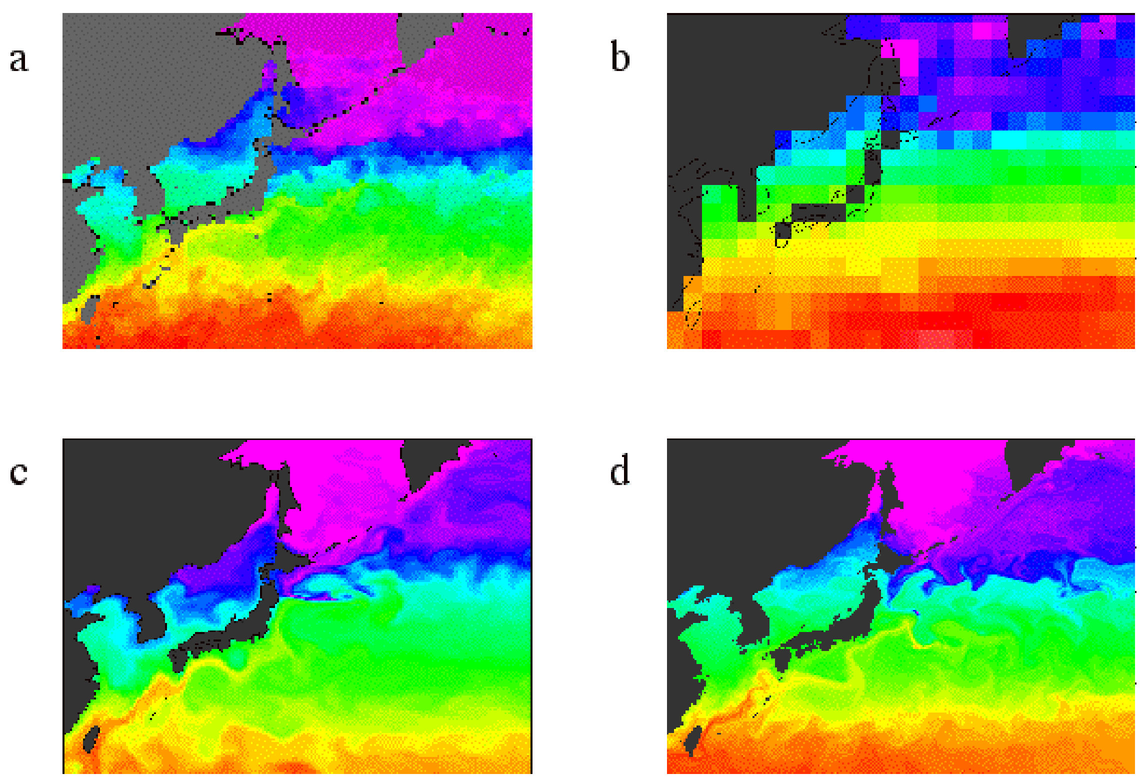

Atlantic. It is not evident that the Kuroshio even exists in the 2° case. The Kuroshio is quite evident in the 0.28 case (c), but does not exhibit the separation from the coast and extension into the North Pacific that is seen in both observations and in the 0.1° simulation(d). These results are confirmed even more strongly by comparing the variability of sea-surface height (SSH) variability between observations and the three different model resolutions (not shown).]

Research Issues

The question of the extent and significance of natural climate variability on the decadal to century time scales and our ability to accurately predict future climate states must be periodically revisited and answered with the best scientific grounding possible. The great advances that have been made in understanding and modeling the individual components of the climate system, the atmosphere, the ocean, the ice and land, are now being applied to understand feedbacks in a earth system model that couples all of the individual components. What causes the various climate oscillations, the El Nino (ENSO), the Pacific Decadal oscillation (PDO) and the North Atlantic oscillation (NAO), and how these interact with each other and anthropogenic factors, is still a matter of active research. These issues have bearing on the sensitivity of the climate system to switches between stable states and on the frequency of extreme weather events such as droughts and catastrophic storms. The predictability of climate depends in large measure on the ability of models to capture and faithfully reproduce the balance of physical processes in what mathematicians call a dynamical system.

[Picture of Ocean biogeochemistry, caption ocean color simulated chlorophyll distributions in the Pacific ocean using the LANL Parallel Ocean Program.

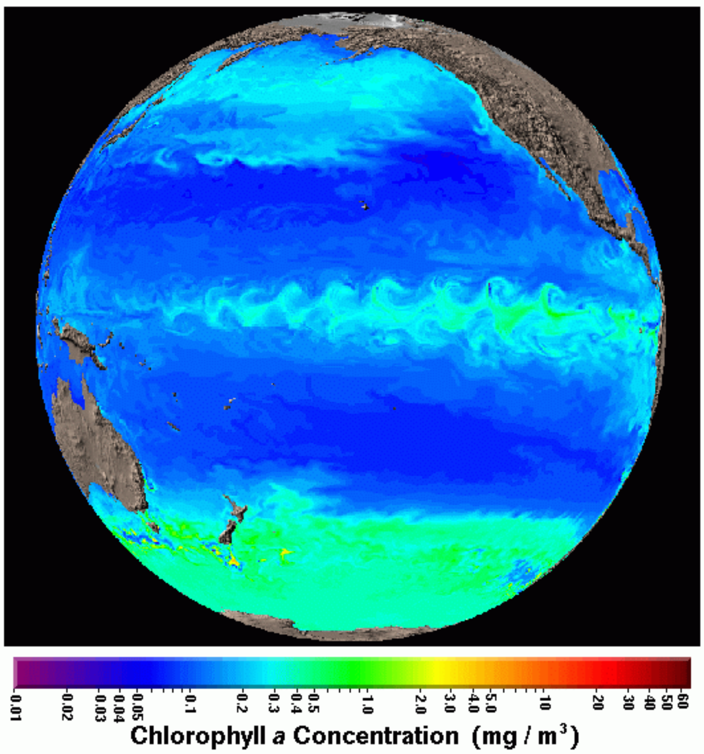

Ocean biogeochemistry includes phytoplankton growth in response to the upwelling of nutrients and temperature].

There are research issues in every facet of climate modeling, ranging from physical process parameterizations, to submodels, to the fully coupled model. Within each of these facets, research issues exist with respect to (sub)model validation by comparison with regional or global measurements and observations. It is often beneficial to reformulate some portion of a (sub)model. This involves research to find a better way to represent a given process or system. Examples are different “dynamical cores” used in atmospheric component models, such as spectral transform, semi-Lagrange, and finite-volume. Each is a different approximation to the same physical equations, and each uses different techniques to solve the equations. The suitability of these methods for reproducing the nearly two-dimensional turbulence in the atmosphere and the conservation of chemical species advected by the winds, continues as an active area with significant overlap with research in computational fluid dynamics and numerical methods.

Finally, research continues on methods by which model codes can be structured and written so that the codes can be easily ported onto a variety of computer architectures without giving up high performance. This software engineering is particularly challenging today because the return of vector- based architectures to the US market has sent code developers reeling with the quasi-orthogonality of vector-based and cache-based programming styles.

	Parameterizations	**Component submodels**	**Fully coupled climate model**
Validate with observations	Measurements from field experiments	Global observations	Global observations
Reformulate	Better "physics"	Better dynamics	Better coupling
Optimize	Better solution	Efficiency, portability	Efficiency, portability

Comprehensive datasets are needed to compare with model predictions (model validation), increase our understanding of the complicated non-linear behavior of the climate system, and provide a basis for development of better models. Several programs provide observational data which is used to validate and refine climate models. Among these are the DOE Atmospheric Radiation Program, the NASA Earth Observing System and the World Ocean Circulation Experiment. As the climate record is extended and better records of the historical forcing are constructed, modelers are able to gain confidence in the fidelity and accuracy of predictions. State of the art simulations of the historical period from 1870 to present using the actual solar variability, greenhouse gas increase and recorded volcanic eruptions, show that models are able to capture the global variation of temperature very well. The recent availability of global ocean data has also been used to initialize coupled models and produced simulations that do not drift from the present climate state. This indicates that the balance of processes in the models has improved dramatically, rendering the physics of the real system more accurately.

Technology Barriers

A number of barriers to progress in climate modeling are linked directly with the computer hardware. The performance of climate models suffers from inadequate memory bandwidth and high memory latency that can not be masked by multiple levels of cache, making it difficult to achieve more than 10% of peak performance of the processor. Performance is dramatically improved on computers with high memory bandwidth. High latency (> 5 s) interconnection networks are also a bottleneck. The poor performance in interconnection networks has severely limited scaling of climate models to hundreds of processors and constrained the algorithms and coding styles that scientists must use to make progress.

Other features of the computer architecture are also important. Climate modeling codes perform better on computers with faster, tightly integrated processors (SMP compute nodes and vector processors). For long time simulations where the time step size decreases with increasing resolution, faster processors are needed to maintain the same level of simulation throughput and scientific productivity.

The growth of computer capability and capacity must take place in a balanced fashion for the climate community to obtain maximum benefit from increased investments in computer hardware. In addition, emphasis must be placed on making dedicated resources available for the long simulations and concentrated studies required to support climate change applications. The climate community has found it difficult to obtain the resources and throughput needed for long simulations at existing computer centers, whose mission is to serve large numbers of users.

Other barriers to progress are associated with the quantity of data that must be handled in climate simulations. As the data accumulates from coupled simulations run for centuries, it has become difficult to analyze using current tools. The distributed archives and analysis centers are making good progress assembling the software tools needed for such analyses but there is a need for high performance networks and high performance switches linking the centers. As the community begins to use data assimilation techniques for climate studies, the bandwidth required to exchange data will grow dramatically.

Some of the most significant advancements in climate simulations have utilized new algorithms to improve accuracy and to achieve greater throughput. Areas of promising research are iterative solution of implicit methods, which are attractive because they may allow larger time steps. These methods suffer computationally because they require global reduction operations on each iteration. These reduction operations result in serious performance degradation in machines built with high latency interconnects.

The same considerations apply to spectral transform methods used in the atmospheric simulations, because the transforms are global in extent. Fast transform methods, such as the FFT (fast Fourier transform), still yield superior operation counts in many solution algorithms, but may scale poorly on massively parallel machines and at high resolution. New methods based on icosahedral meshes and smoothly mapped meshes show promise for better scalability.

Fundamental research into new mathematical methods for the simulation of complex, interacting fluid flows are also needed to advance climate simulation capabilities over the next decade. Specific topics of interest include, semi-implicit and operator split methods to allow long time integrations in the presence of fast moving gravity waves, Lagrangian vertical coordinate systems and conservative remapping schemes that allow accurate thermodynamic simulation of transport and moist processes, and fast methods for solution of elliptic systems.

Also of considerable interest are software engineering practices that allow a community of hundreds of computational climate scientists, mathematicians, computer scientists and numerical analysts to develop community codes. It is difficult to write codes that will run effectively on the wide range of computer architectures that are available today.

Compiler technology is not keeping pace with high performance scientific computing demands. Parallelism constructs that are stable and robust are also sorely needed. As the memory hierarchies deepen and more architecture specific layers are added, modelers require more support in dealing with software issues. A significant effort is required, and more software engineering support needed, to adapt to new architectures and maintain the pace of scientific development.

Resources Required

There are many ways in which increased computational resources could be used to improve climate system simulations. Efforts to model biogeochemical cycles and their relationship to climate require development of comprehensive coupled models on a scale unprecedented in any computational science discipline. To develop models sufficient for addressing these questions and issues will require changes to the current state of the art in the following ways:

(A) Increase the spatial resolution of the grids of the coupled model components. At present, the "standard" atmospheric (oceanic) simulations uses a grid with 300 (100) km resolution and 26 (40) vertical levels. The resolution targets are about 10 km in both the atmosphere and ocean, but for different reasons. It has been demonstrated that 10 km resolution is needed to resolve oceanic mesoscale eddies. A similar resolution is needed in the atmospheric component to obtain predictions of surface temperature and precipitation in sufficient detail to analyze the regional and local implications of climate change.

(B) Increase the completeness of the coupled model by adding to each component model important, interactive physical, chemical and biological processes that heretofore have been omitted due to their computational complexity. Inclusion of atmospheric chemistry, both tropospheric and stratospheric, and biogeochemistry in the ocean are essential to understand the ecological implications of climate change.

(C) Increase the fidelity of the model by replacing parameterizations of sub-grid physical processes by more realistic and accurate treatments as our understanding of the underlying physical processes improves, often as the result of observations field programs like the DOE Atmospheric Radiation Measurement Program.

(D) Increase the length of both control and climate-change-scenario runs. Longer control runs will reveal any tendency for the coupled model to drift, and will also improve our estimates of model variability. Longer climate-change-scenario runs will permit examination of critical issues such as the potential collapse of the global thermohaline circulation that may occur on time-scales of centuries in global warming scenarios.

(E) Increase the number of simulations in each ensemble of control runs or climate-change-scenario runs. Increase the number of climate-change scenarios investigated. These issues are both examples of perfectly parallel extensions of present-day simulations: each instance of another scenario or ensemble member is completely independent of every other instance. Ensemble members are

distinguished by small perturbations in their initial conditions, which are quickly amplified by the non-linearity of the equations. The use of ensembles provides an important measure of the range of variability of the climate system.

These are summarized in the following table.

Issue	*Motivation*	*Compute Factor*
Spatial resolution	Provide regional details	10^3-10^5
Model completeness	Add "new" science	10^2
New parameterizations	Upgrade to "better" science	10^2
Run length	Long-term implications	10^2
Ensembles, scenarios	Range of model variability	10
Total Compute Factor		10^{10}-10^{12}

It is important to note that these "compute factors" are relative to the capability of computer systems typically used *now* for coarse resolution (300 km atmosphere, 100 km ocean) simulations. Sustained performance on present-day climate simulations is from 5-20 Gflops. By contrast the Japanese Earth Simulator has demonstrated 20-30 Tflops.

An increase of capability by a factor of 3 would allow the addition of dynamic vegetation to a fine scale land model. Another factor of 2 would provide enough power to routinely include tropospheric chemistry. Ocean biogeochemistry could be included in the coupled model for an additional factor of 3-5. Interactive carbon in a full carbon cycle will require a further factor of two. Extending the atmospheric model to include a full stratosphere and increasing the vertical resolution requires another factor of 5 increase in capability. These additions may not be fully effective unless the ocean and atmospheric horizontal resolution is substantially increased. An eddy resolving ocean would require a factor of 1200. In the 10 year time frame, it will be important to include cloud-resolving atmospheric simulations in a fully coupled earth-system-modeling framework. The cumulative requirement supporting these developments is estimated to be a factor of nearly 10,000 in flops. This would take us into the peta-flop range for computing

In addition to the performance requirements listed above, storage requirements would also increase for each of the possible paths listed above. The storage requirements for sample atmosphere and ocean resolutions are shown in the table below. The table also shows the storage required for each additional tracer in an atmospheric chemistry or ocean biogeochemistry simulation. A low-resolution coupled model run for a century produces almost 1TB of output. In a configuration with a 75km atmosphere and a 0.1 degree ocean, a single century run would produce 23 TB. Storage for future cloud resolving and eddy resolving simulations is on the order of Petabytes.

Component	**Resolution**	**History output (GB/sim-yr)**	**Each tracer (GB/sim-yr)**
Atmosphere	T42 (300 km)	7.5	0.02
Atmosphere	T85 (150 km)	29	0.08
Atmosphere	T170 (75 km)	110	0.3
Ocean	1° (100 km)	1.7	0.2
Ocean	0.1° (10 km)	120	17

There is not a well-defined path from the present-day model to a hypothetical future model. The presumption is that the end-state of the model of the future (a full earth system model) does *not* depend on which path is followed to get there. An example is given in Figure 1 based on the assumption that an Earth Simulator-class machine is utilized to first increase the component-model resolutions to the target of 0.1°, then complete sub models of atmospheric chemistry and ocean biogeochemistry are added to make up a fully coupled model that is then used for regional climate studies and large ensembles. The red arrow in Figure 1 illustrates "Moore's Law" for the growth-rate of processor capability. Even if Moore's Law continues to hold, computing power will still fall short of the capability needed.

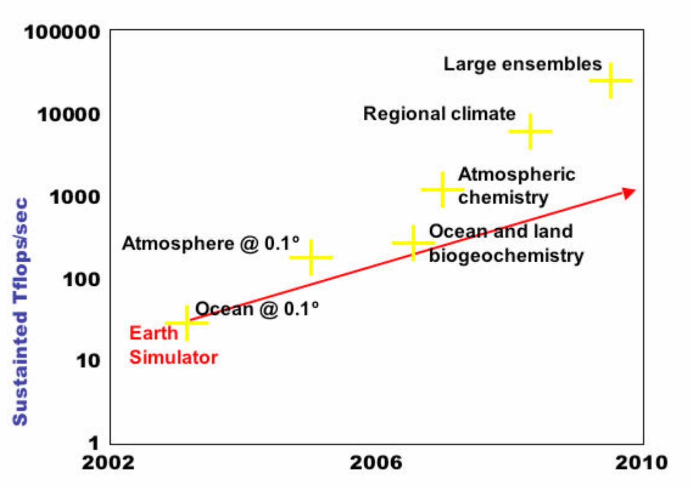

Clearly we cannot rely solely on increases in brute computer power if we are to attain our goals in a timely way. A significant role will be played by scientists finding better ways to model the climate and by continued vigilance in validating model predictions against the best observational records available. History has shown that development of new algorithms and numerical methods can reduce computing requirements by orders of magnitude. Similarly, new ways of formulating models and approximating physical, biological and chemical processes can lead to substantial reductions in cost and/or improved fidelity of the models. Such improvements are harder to forecast than the technological improvements in computer hardware, like that embodied in Moore's Law, but when they do occur they enable major breakthroughs in understanding and predictive capability.

The research objectives for climate science require a diversified investment in environmental science disciplines as well as cross-cutting, enabling technologies. The development of multi-disciplinary teams, such as those funded in the DOE SciDAC Program have been very successful in bringing resources to bear on specific problems and providing timely solutions to difficult problems.

Metrics of Success:

The objective of high end climate modeling is increased scientific understanding of the climate system and the possible effects of human induced climate change. That the science is the driver is the only thing that makes sense and so we pose metrics in terms of scientific productivity. The climate community expresses simulation throughput in simulated years per wall clock day. It has been a goal to achieve overnight turnaround of a “standard” 20-year atmospheric run in support model development. This is roughly equivalent to 15,000 times real time. Even for the standard low resolution runs we don’t achieve better than 1,000 times real time. These metrics assume that the science drives the resolution and complexity requirements. The U.S. climate modeling community is currently failing to achieve its science objectives because it cannot attain the needed throughput rate with high resolution models. Of course, the more we simulate the more data there is to process and analyze. Usually, data analysis takes longer than the run.

Another measure of success is the U.S. climate community’s ability to support climate change simulations of scenarios suggested by the IPCC Assessment and the U.S. National Assessment of Climate Change. Adequate computing capacity and networking infrastructure are crucial to meet these commitments.