Benefits of an Earth Simulator Class Machine for US Climate Science

Improving climate prediction using scientifically grounded simulation models is a complex endeavor involving fundamental research on physical, chemical and biological processes, ground based and satellite observing systems, and computational methods research. The simulation model is the integrating fabric that brings together theory, data and computational resources. The success of numerical weather prediction employing high-end computing is replicated in climate prediction where long time scales demand even greater computing capability and the need to evaluate many future scenarios quantifying the statistical uncertainties requires large computing capacity.

Climate research has stressed the largest computers available in the past and will continue this trend into the foreseeable future. The Japanese Earth Simulator offers a much higher capability than the currently available US computers (by a factor of 5 to 20 in climate performance metrics). Since the mid-90's climate research has benefited from access to NEC and Fujitsu vector computers available through collaborations with our foreign colleagues. The expansion of availability to large IBM SP cache based systems within the US has also benefited climate research as the cycles have become more freely available to researchers. In the last few years, US scientists were able to use ensemble forecasting methods with climate models, a practice that had already become standard in European centers with access to the large vector machines.

Several avenues open for climate model development when increased computational power becomes available, like that of an Earth Simulator class machine. The model resolution can be increased to provide regional detail in the context of a coupled ocean, atmosphere, ice and land model. The increase in resolution in the ocean allows for the transport of energy by eddy currents that are accurately modeled. Similar increases of the regional detail in the atmospheric model allow for more accurate representation of the jet streams and the effect on the winds of mountainous terrain. Of great interest to DOE is the development of comprehensive carbon accounting within coupled models. This requires significantly more computation to calculate chemical balances in the atmosphere and the inclusion of dynamic, interactive plant growth and succession. To maintain the development cycle of new capabilities in climate prediction and to maintain the solid scientific grounding of the US based models requires access to computational resources comparable to those available outside the US.

The application of climate models to answer the "what if" questions posed by energy managers and policy planners is also impacted by access to adequate computational resources. Current efforts to explain the interaction of cloud composition with the Indian monsoon, or the dependence of the global warming response of models on key processes, require many exploratory simulations. The inherent variability of the climate system, with El Ninos, North Atlantic oscillations, Pacific Decadal oscillations, creates additional complexity for evaluating the human impacts of climate change. To properly represent and predict extreme weather events in a changing climate is pushing us toward petabyte size output data sets that require analysis by geographically distributed teams of researchers. Only concerted support of the network and computational infrastructure will allow scientific projects of this magnitude and complexity to be successful.

The Japanese Earth Simulator consists of 640 nodes, each with 8 vector processors. The peak performance is rated at 40 TFLOP/s, thus the designation ES40. The nodes and processors within a node are linked with a high-performance switching fabric capable of delivering 16 Gbytes/s in each direction. The vector performance of processors together with the highly engineered distributed-memory system set this system apart from available US computers. The IBM and SGI platforms available and widely used within the US are cache based systems based on commodity processors. These systems continue to advance processor speed on a Moore's law curve. Unfortunately, the memory bandwidth and latency of remote memory access have not kept pace. While large, fast cache and high-speed interconnects are one solution to the climate codes heavy memory access patterns, the recent years of code restructuring and optimization has revealed the hardware limits of this organization. Vector processors with shared memory interconnects, specially designed to support high data volumes, demonstrate higher efficiencies for climate and weather problems. A spectral atmospheric code operating at T1279L96 (10km resolution) on all 640 nodes of the ES40 sustained 26 TFLOP/s, with efficiencies in excess of 50%. This contrasts with the best in class of 10% of peak on the cache based-cluster machines. The peak numbers are really not what concerns us, rather it is the throughput on climate simulations which impacts our science.

The benefits to US science are that we will identify the issues and quantify the problems that must be faced in a successful response to the Japanese Earth Simulator and work with vendors to solve the problems. The project will strengthen the US high-performance computing capability by bounding the problem and providing information to hardware designers and developers early in the design/production pipeline. This will focus the US response in terms of climate prediction and important scientific capabilities of relevance to the DOE.