May 29, 2002

From: Advanced Scientific Computing Research Advisory Committee, Office of Science

To: Dr. Raymond Orbach, Office of Science, Department of Energy

Subject: Advanced Scientific Computing Initiative

In April 2002, headlines in every leading newspaper announced that the Earth Simulator, a supercomputer funded by the Japanese government, was now the fastest computer in the world. The Earth Simulator project itself, initiated in 1997 by the Japanese government and described at many subsequent scientific meetings, was not news. What was surprising in the short term, even to Earth Simulator researchers, was that 87% of the theoretical peak speed was achieved at such an early stage. Taking a long view, however, the Earth Simulator's major and lasting success will derive not from a new record for sustained speed, but rather---far more importantly---from the scientific accomplishments that it will enable.

The Earth Simulator was developed for a scientific purpose: to make accurate predictions of global climate change. Because earth-scale experiments with climate are impossible, our understanding of climate change derives mostly from computer models. The best climate models, built from complex mathematics and sophisticated numerical methods, ultimately depend on high-end computers to perform the needed calculations. Several factors limit the quality of climate model predictions, including available computing power, the level of detail in the underlying mathematics, and the accuracy with which equations can be solved. Current U.S. climate models achieve an overall resolution of approximately 100 kilometers, which is too coarse for analysis of small-scale features involving mountain and coastal effects, river flow, cloud and storm systems, or hurricane storm predictions.

The state of the art in climate change prediction has changed overnight because of the Earth Simulator. For the 2006 International Panel on Climate Change Assessment, the Earth Simulator project is expected to lead the world by providing simulations of climate change with resolutions in the range of 10-30 kilometers. This new leadership position will follow not only from a 40- to 50-fold increase in capability compared to the most advanced U.S. machines, but also from the Earth Simulator project's insistence on extremely close interactions among teams working on experiments, data collection, mathematical models, computational algorithms, and software development.

The characteristics just described---reliance on mathematical models, numerical methods, software tools, computing, and multidisciplinary teams---are by no means unique to climate science, but exist in many other disciplines, including nanoscience, accelerator design, astrophysics, combustion, materials science, fusion energy research, and biology. In science of the twenty-first century, simulation and high-end computation are equal partners with theory and experiment.

The lesson to be drawn from the Earth Simulator is that remarkable results can follow from a sustained commitment to advanced scientific computing research in the service of science. The Department of Energy has a long history of leadership in this domain, and the Advanced Scientific Computing Advisory Committee believes that the time is right for a major new multidimensional initiative in computing from the Office of Science. Such an initiative should be grounded on strong collaborations among application scientists, computer scientists, and mathematicians; it should include research on problem formulation, mathematical modeling, numerical algorithms, visualization, software tools, computer architectures, networking, and data analysis, storage, and management. In addition, research on advanced communications is essential to provide infrastructure needed by scientists working with distributed and scalable information resources, instruments, and computational facilities. The right *mix* and *balance* of activities will be crucial to the success of such an initiative: each component must be individually important as well as demonstrate a visible contribution to the others. Since science is the ultimate driver, the initiative should have a set of urgent, challenging, and exciting scientific goals that build from the strengths of the base programs.

Although much of the fallout from the Earth Simulator's publicity has focused on questions of computer architecture, we stress that, in our view, attempts at the outset of the initiative to determine the best future architectures would be distracting and divisive. Funding for research on high-end architecture design has dwindled to almost nothing in recent years, in large part because of a reluctance to depart from machines built with commodity parts; equally, the hardware needs in consumer-based information technology do not necessarily match well with scientific needs at the highest end. Decades of experience have demonstrated that no single architecture is optimal for every application, especially when new algorithms constantly rewrite the rules about how problems should be solved. The important thing is to support enough diversity so that our options are not unnecessarily limited.

To ensure results of the highest quality, the committee recommends peer review of every aspect of the initiative, particularly the most costly elements. Further, the Office of Science is not the only agency or office with expertise in advanced scientific computing, and the committee believes that any new initiative should be based on maximal cooperation and leverage.
Within the Department of Energy, the NNSA has had substantial experience with a multi-faceted program in advanced scientific computing. The National Science Foundation is committed to development of a computational grid for support of a broad base of scientific applications. The Department of Defense has supported development of high-performance architectures with an eye toward scientific applications.

As a committee, we are unanimously enthusiastic about the prospects for a challenging initiative in advanced scientific computing to be undertaken by the Office of Science. Such an initiative would not only enable all of U.S. science at the highest end into the foreseeable future, it would also signal to the world our unquestioned leadership in and commitment to scientific computing.