Comments on "Bolstering U.S. Supercomputing" by S. L. Graham, M. Snir, and C. A. Patterson

Thom H. Dunning, Jr.
National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

I commend Susan Graham, Marc Snir and Cynthia Patterson for clearly laying out the current supercomputing situation in the U.S. Advances in science and engineering require access to forefront computing capabilities. These capabilities are needed by researchers analyzing increasingly voluminous data sets as well as those involved in modeling and simulation and the future promises many new applications that will require supercomputers. National supercomputing resources are currently provided by the NSF and DOE Office of Science (DOE-SC) supercomputing centers, although flat or declining budgets have limited their ability to satisfy the growing needs of the research community. The authors call for increased federal investment, as well as coordination, in this area is much needed.

The authors' thorough analysis of computing technologies is very timely. NCSA has been successfully providing "killer micros" for scientific computing for nearly a decade. However, a number of important scientific and engineering applications are not well suited for this architecture. Further, the architecture of "killer micros" is changing as chip designers face the problems associated with high chip frequencies (the traditional means of increasing computer power). Now is the time to reconsider the design of supercomputers for scientific and engineering applications, realizing that matching application to architecture will maximize scientific productivity and minimize cost. The high energy physics community is already exploring custom supercomputer architectures for their applications (*e.g.*, QCDOC).

In designing a new generation of supercomputers we must not be misled by the Top500 list. U.S. computer vendors (IBM and SGI) hold the top three spots on the most recent list, but this ranking is not a reliable indicator of performance for many applications critical to advancing scientific discovery and the state of the art in engineering. Further, to achieve the stated performance levels for the top two entries (both IBM BlueGene/L computers), applications must run efficiently on 40-65,000 processors—today, few applications scale to more than a few thousand processors. The Japanese Earth Simulator, which follows the supercomputer design philosophy articulated by the late Seymour Cray, is still considered by many to be the world's preeminent supercomputer for real-world applications. It is ranked as #4 on the Top500 list and achieves that performance level with just 5,120 processors.

Supercomputing is more than hardware. To realize the benefits of supercomputers requires new software—computing systems software to enable thousands of processors to effectively work together and scientific applications that achieve high performance levels and scale to ten thousand or more processors. In 2000, the DOE-SC created the Scientific Discovery through Advanced Computing (SciDAC) Program that was targeted at just this problem. Although funding for the SciDAC Program was relatively modest, its recent 5-year program review clearly illustrates the remarkable advances that can be made by teaming disciplinary computational scientists, computer scientists, and applied mathematicians. NSF's Information Technology Research (ITR) program supported similar activities, although it is now being ramped down. Research, development and deployment of supercomputing software must also be supported if we are to realize the full potential of supercomputing.