

The Globalization of Large-Scale Computational Science

During recent trips to Europe it has become clear to me that the international scientific community is beginning to challenge the historical U.S. dominance in high-performance scientific computing. Many people in the United States have responded by saying that the nation should back off on its investments in leadership-class systems as being too expensive, while others argue that national security requirements demand that the United States go it alone in investments for high-end systems research and development. I maintain, however, that the United States should welcome the “rise of the rest” as the natural maturing of the global scientific computing community. I see this situation as an opportunity for the United States to continue to lead in high-end computing as the “trusted and honest broker”: one that can organize the global effort and forge long-term relationships that span the global community to ensure that systems will be developed, that software will be available, and—most important—that the best computational science targets will be pursued regardless of national boundaries.

Europe has formed the PRACE Consortium, whose purpose is to transform high-performance computing in Europe through the combined efforts of the European member nations. Five of these members (France, Germany, Spain, the Netherlands, and the United Kingdom) will take the lead in deploying three to five petascale systems over the next few years. The announcement of the PRACE initiative, coupled with the news of Saudi Arabia’s pursuing a petaflop capability, underscores a dramatic and likely permanent shift in the global landscape of high-performance computing (HPC). It is no longer a race dominated by the United States, occasionally punctuated with bold moves by Japan. It is becoming a global landscape rich with competition and cooperation opportunities as more nations recognize HPC’s value to their scientific research, education, and industrial enterprises.

Arguably, the majority of the hardware platforms targeted for near-term international acquisition and deployment are products of U.S.-based manufacturers, and the key software enablers (operating systems, programming models, libraries) are mostly of U.S. origin. But this situation is unlikely to hold true in the long term unless the United States embraces the globalization model for HPC, establishing a broad set of international partnerships that not only collaborate at the scientific level but co-invest in the research and development (R&D) necessary for sustaining leadership. The United States simply cannot afford to go it alone, but neither can the United States afford to forgo the benefits of leading-edge facilities at U.S. institutions. The level of R&D investment needed to reach an exascale capability is likely to approach \$200 million to \$300 million per system, per

Biography in Brief

Professor Rick Stevens is associate laboratory director for Computing, Environment, and Life Sciences at Argonne National Laboratory and a professor of computer science at the University of Chicago. His research interests are in high-performance computer architectures, and computational science, especially problems in the life sciences. He heads Argonne’s efforts in advanced computing targeting the



Professor Rick Stevens

development of exascale computing technology and applications in systems and computational biology and environmental modeling and simulation. He is a fellow of the American Association for the Advancement of Science and a co-founder and senior fellow of the Argonne/University of Chicago Computation Institute, a multidisciplinary institute aimed at connecting computing to all areas of inquiry at the University and the Laboratory.

vendor, perhaps totaling a billion dollars over 10 years. And bringing into the multicore era the existing billions of lines of scientific and engineering codes will cost even more, probably exceeding several billion dollars over the next 10 years. These sums are large, but they underpin sustaining U.S. leadership in scientific and engineering productivity.

The United States has given up its leadership in several areas of large-scale science. *HPC should not be the next.* The re-emphasis on HPC in Europe and its emergence in other regions provide a golden opportunity for the United States to work with the global community to deploy petascale systems, to co-develop and deploy exascale architectures, and to lead the global HPC software development. The potential payoff is enormous: assuring the United States access to the scientific frontiers in computational science and the best software developed worldwide, sustaining a worldwide developer base needed for HPC, and inviting a larger market for U.S.-based systems and designs. ●

Contributor: Rick Stevens, associate laboratory director for Computing, Environment, and Life Sciences at Argonne National Laboratory, and professor of computer science at the University of Chicago