Features:

HPC User Forum Wrap-Up

by John E. West
Senior Fellow, DoD HPC Modernization Program

Industry research group IDC hosts five to six User Forum meetings around the world each year. About 100 people participated in the most recent meeting, representing government, industry and academia, as well as all the major HPC vendors. Each User Forum has a theme; this one focused on the use of HPC in the energy industry.

What follows is a subjective selection of highlights and topics of potential general interest from this meeting, which took place in Santa Fe, N.M. on Sept. 26-27.

The Keynote

The meeting keynote was delivered by Victor Reis, Senior Advisor, Office of the Secretary, Department of Energy. He has primary responsibility for the Global Nuclear Energy Partnership, part of President George W. Bush's Advanced Energy Initiative, and he is also a member Strategic Advisory Group of the U.S. Strategic Command. Reis was the Director, Defense Research and Engineering when the DoD's High Performance Modernization Program started, and he was a senior official at DOE when it began the ASCI (Accelerated Strategic Computing Initiative) program.

Reis reviewed the history of ASCI and what it has accomplished to date, and then discussed a potential new DOE program involving physics-based design of nuclear reactors for peaceful energy production. He feels that the timing is correct for instituting a new HPC program for this purpose and is gathering information to support such a program. He mentioned several potential modeling efforts that would contribute to the program, such as optimization of the nuclear reactor fuel cycle, design and qualification of new nuclear fuels, detailed modeling of new reactor designs, and environmental effects on nuclear reactors, particularly earthquakes. Several DOE talks followed which discussed modeling of fission reactors and the status of nuclear fusion research.

Energy-Related Discussions

The theme of this meeting was HPC in energy, so naturally there were several discussions of advanced energy research in addition to coverage

in the keynote.

Keith Gray of BP discussed their seismic imaging research and development, which is designed to improve the information content of seismic images by processing with HPC capabilities. He specified several basic computational challenges and requirements: large-memory nodes for development work, easier parallel tools, effective use of emerging multicore systems, and bigger and better file systems.

Mark Nimlos of the National Renewable Energy Laboratory discussed the status of various forms of alternative energy sources and concentrated on his work in the biofuels program, which has a goal of replacing 30 percent of current transportation fuels with biofuels by 2030. He is carrying out sophisticated molecular dynamics computations of how one of the key enzymes breaks down cellulose into sugars, with the intent of understanding how to optimize the process.

Pratul Agarwal of Oak Ridge National Laboratory, working in the same overall program, discussed the multiscale nature of biofuel processing and the need for collaborative efforts between experimental and computational work. His group is considering the use of new HPC technologies such as FPGAs and GPUs to accelerate the computation of the enzymatic pathways involved in the conversions of cellulose to sugars. He noted that the follow-on processing of sugars to alcohols (fermentation) was well understood, at least at the production level, because of the many thousands of years of experimentation by human beings in this process.

HPC Acquisition and Architectures

In addition to the domain-focused fare there were also several discussions of recent HPC acquisitions and new HPC architectures.

Rupak Biswas of NASA-Ames discussed their Columbia system and efforts underway to procure a replacement for it. He provided information on NASA's HPC requirements as part of the discussion, highlighting growth of those requirements across several NASA directorates.

Richard Walsh of IDC, and formerly of the Army High Performance Computing Research Center (AHPCRC), provided his taxonomy of processor architectures and applications, stating that there were significant drivers toward heterogeneous processors in the near future.

John Daly of Los Alamos National Laboratory (LANL) discussed issues of running applications at large scale, including how to handle interrupts and how often to write out checkpoint/restart files. He showed data that indicated that even small jobs on really big systems might be at significant risk of an interrupt because of the dependence of mean-time-between-application-interrupt on numbers of job processors; instead of a linear dependence, the low-number-of-processor behavior of this interrupt time is considerably less than linear. This leads to alternative scheduling policies that emphasize large-number-of-processor jobs at the expense of long running times.

John Gustafson of ClearSpeed Technology provided some impressive speed-ups on a variety of application codes with their accelerator technology and also presented a rule of thumb estimate of current HPC system power per volume. According to him, this turns out to be about 70 watts per liter.

Finally, your very own John West gave a talk discussing what HPC systems will look like approximately ten years in the future. After talking to multiple industry sources, I envision a multicore future with general-purpose HPC systems of the 2017 era comprised of chips containing hundreds of computational cores per chip (not thousands), among other interesting features. This presentation was, of course, fascinating.

University Panel

One of the panels at the forum involved representatives from several universities involved in HPC and university-affiliated computer centers: Penn State, Ohio Supercomputer Center, Pittsburgh Supercomputing Center, San Diego Supercomputer Center, the University of Minnesota, the University of Nevada at Las Vegas, the University of Tennessee, Utah State University, and Virginia Tech.

There was extensive discussion of the role of university computing centers versus NSF national computing centers. As expected, the university computing centers would like NSF support for their niche in the overall structure. Many of these universities offer unique degree programs in computational science. One suggestion arising from the discussion was to develop a partnership among the federal agencies (NSF, DoD and DOE) to promote and support these educational programs in computational science.

Data Intensive Computing Environment (DICE)

Roger Panton of Avetech, the executive director of the Data Intensive Computing Environment (DICE) program, provided the history and status of that program. DICE was motivated by the HEC/RTF report of several years ago and is a partnership among DoD, NASA and DOE. The goal of the program is to set up a testbed to evaluate data management technologies that could improve data accessibility over geographically distributed sites. Current organization partners include Advanced Simulation and Computing (ASC), Ohio Supercomputer Center (OSC), NASA-Goddard, and Avetech; the High Performance Computing Modernization Program (HPCMP) participates through ASC. New sites will include the Pacific Northwest National Laboratory and Sandia National Laboratories. Roger discussed future partners and projects.

The next two U.S. meetings of the HPC User Forum are scheduled for April 14-16, 2008 in Norfolk, Va., and Sept. 8-10, 2008 in Tucson, Ariz. You can find out more about these events by clicking over to http://hpcuserforum.com.