LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-440583

HPC projects: Now's the time to look forward to exascale computing

D. L. Crawford

July 8, 2010

Scientific Computng World

Dona Crawford is the Associate Director for Computation at Lawrence Livermore National Laboratory in Livermore, Calif. She is responsible for the Laboratory's computing infrastructure and the high-performance computing systems that serve research programs from national security to materials science, energy, and climate change.

HPC projects: Now's the time to look forward to exascale computing

Computing is in Lawrence Livermore National Laboratory's DNA. Ever since the Laboratory's founding by Edward Teller and Ernest Lawrence, computing has been central to the research and development that put Lawrence Livermore on the global map. Even before the Laboratory officially opened its doors in September 1952, Teller and Lawrence ordered a Remington–Rand Universal Automatic Computer, or Univac-1, which, with its 1,000-word memory and 5,600 vacuum tubes, represented the HPC heavyweight of its time. Teller and Lawrence understood that computing was the key to transforming the Lab from an obscure second nuclear weapons lab, still in the shadow of the better-known Lawrence Berkeley Lab, into the world-class research and development institution it is today. Scientific computing has historically been a principal integrating element of the multidisciplinary team science that defines Lawrence Livermore's *modus operandi*.

The Laboratory's leadership has always taken the long view of high-performance computing in an effort to anticipate the need for the ever more powerful computing systems required to fulfill missions in national security as well as the basic science underpinning all lab research and development. The needs of research scientists for computational capabilities to solve intractable problems have always exceeded the computational power available. Consequently, developing new computing resources has long been a part of the national labs' charter. Recognizing this fact of research and development, the U.S. Department of Energy (DOE) has historically played a role in developing new computing capabilities by working in partnership with industry (that role today is played by DOE's quasi-autonomous National Nuclear Security Administration). In my time at Sandia and Livermore, the need for these capabilities has been driven by the ban on testing nuclear weapons, something we all agree is a good idea. In the absence of testing, the computational muscle required to address the technical and scientific challenges of the effort to ensure the safety and reliability of the nation’s shrinking nuclear arsenal, grows as weapons age decades beyond their intended service life.

HPC has played a vital role in making it possible to reduce the U.S. nuclear arsenal to its smallest size since the administration of President Dwight Eisenhower. The process of reducing the stockpile further, and eventually eliminating nuclear weapons as President Barack Obama envisions, will require increasingly powerful HPC systems—capabilities that can also be leveraged to address other national challenges, such as energy and climate change.

We now know that petaFLOP/s (quadrillions of floating point operations per second) systems will not be sufficient to meet the needs of the weapons program or other global challenges the Lab has taken on in security, energy, environment, and climate change. For example, exascale computing will be required to better understand nuclear boost, combustion (how to reduce emissions), improved fission energy, fusion energy, power grid modeling, carbon sequestration, and climate modeling. That is why, even though we have yet to install our first petaFLOP/s system on the floor of Livermore’s simulation facility, we are already starting to plan for exascale computing. Overcoming the many technical hurdles to exascale—including power requirements, system reliability, and writing applications for billion-way parallelism—and developing the needed tools and technologies will require years. Current forecasts put the deployment of the first exascale systems in the 2018–20 timeframe. But that can only happen if we start now.

DOE has launched a planning effort for an exascale initiative to lay the groundwork for next-generation HPC. Leaders in HPC from our Lab along with representatives from other labs, including Sandia, Los Alamos, Argonne, Oak Ridge, and Lawrence Berkeley, meet on a regular basis to discuss the path forward. I am confident that we will reach exascale, and my confidence is based on experience. In 1995 when DOE conceived the Accelerated Strategic Computing Initiative (ASCI), few computing experts believed the program would be able to take high-performance computing from approximately 50 gigaFLOP/s to 100 teraFLOP/s in under 10 years. The 100-teraFLOP/s IBM Purple system, dedicated at LLNL in 2005, represented the culmination of that initiative, though it was deployed at the same time as IBM’s first BlueGene/L system, a 360-teraFLOP/s machine, also at LLNL. BlueGene/L held the No. 1 ranking on the Top500 list seven times from 2004 to 2008. ASCI, which brought together the computing expertise of Livermore, Los Alamos, and Sandia national laboratories, established the framework for advancing computing to where it is today with each of the national labs working in partnership with industry to pursue different approaches to realizing next-generation supercomputers.

We have a lot of experience working with industry to bring new HPC technologies to fruition. In partnership with IBM, we have brought five generations of supercomputers to Livermore and are working to bring in a sixth next year—the 20-petaFLOP/s IBM BlueGene/Q system, Sequoia. Previous IBM systems include: the 3-teraFLOP/s Blue Pacific; the 12-teraFLOP/s White; the 100-teraFLOP/s Purple; the 500-teraFLOP/s BlueGene/L; and the recently installed 500-teraFLOP/s BlueGene/P system called Dawn, which is paving the way for Sequoia. With Sequoia we are making the single biggest leap ever attempted in computing power, from 500 teraFLOP/s today to 20 petaFLOP/s in 2012. To reduce the time required to adapt applications to multi-petaFLOP/s systems and to develop the network and storage these systems will require, we have again turned to industry for assistance.

Through an innovative collaborative project called Hyperion, we are working with 10 computing industry partners to address some of these issues. Hyperion brings together Dell, Intel, Supermicro, QLogic, Cisco, Mellanox, DDN, Sun, LSI, and RedHat to create a large-scale testbed (100 teraFLOP/s) for HPC technologies critical to our work maintaining the aging U.S. nuclear weapons stockpile, as well as industry's ability to make petaFLOP/s computing and storage more accessible to commerce, industry, and research and development. Hyperion represents a new way of doing business. Collectively, we have been able to build a testbed none of us could have built individually. This project is a cost-effective way to advance the state-of-the-art while benefitting both the end users, such as national labs like Livermore, and industry.

Building on these collaborations and forming new innovative partnerships will be essential to achieving exascale computing. The challenge is daunting and the distance appears beyond bridging. But to those of us at Livermore who have been part of this computational enterprise for a while, this position is not unfamiliar. Surmounting these technical and scientific obstacles is what is expected of national labs like Livermore. And we eagerly pick up the gauntlet.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. UCRL-XX-XXXXXX