


Trillion particle simulation on hopper honored with best paper

June 3 2013



Credit: Oliver Rubel, Berkeley Lab

An unprecedented trillion-particle simulation, which utilized more than 120,000 processors and generated approximately 350 terabytes of data, pushed the performance capability of the National Energy Research Scientific Computing Center's (NERSC's) Cray XE6 "Hopper" supercomputer to its limits.

In addition to shedding new light on a long-standing astrophysics mystery, the successful run also allowed a team of computational researchers from the Lawrence Berkeley National Laboratory (Berkeley Lab) and [Cray Inc.](#) to glean valuable insights that will help thousands of scientists worldwide make the most of current petascale systems like Hopper, which are capable of computing quadrillions of calculations per second, and future exascale supercomputers, which will compute quintillions of calculations per second.

The team described their findings in "[Trillion Particles, 120,000 cores, and 350 TBs: Lessons Learned From a Hero I/O Run on Hopper](#)," which won best paper at the 2013 Cray User Group conference in Napa Valley, California.

"When production applications use a significant portion of a supercomputing system, they push its computation, memory, network, and parallel I/O (input/output) subsystems to their limits. Successful execution of these apps requires careful planning and tuning," says Surendra Byna, a research scientist in Berkeley Lab's Scientific Data Management Group and the paper's lead author. "Our goal with this project was to identify parameters that would make apps of this scale successful for a broad base of science users."

For this particular run, the team simulated more than two trillion particles for nearly 23,000 time steps with VPIC, a large-scale [plasma physics](#) application. The simulation used approximately 80 percent of Hopper's computing resources, 90 percent of the available memory on each node, and 50 percent of the Lustre scratch file system. In total, 10 separate trillion-particle datasets, each ranging between 30 to 42 terabytes in size, were written as HDF5 files on the scratch file system at a sustained rate of approximately 27 gigabytes per second.

"This is the largest I/O job ever undertaken by a NERSC application. It

is quite a feat when you consider that even the smallest bottleneck in a production I/O stack can degrade performance at scale," says Prabhat, a researcher in Berkeley Lab's Scientific Visualization Group and co-author of the paper. "If we had attempted a run like this three years ago, I would have been unsure about the level of performance that we could get from HDF5. But thanks to substantial progress made over the course of the ExaHDF5 project, we are now able to demonstrate that HDF5 can scale to petascale platforms like Hopper and achieve near peak I/O rates."

ExaHDF5 is a Department of Energy funded collaboration, led by Prabhat, to develop high performance I/O and analysis strategies for future exascale computers. HDF5 is an I/O library used by more than 150 institutions worldwide. The HDF5 Group at University of Illinois Research Park is currently maintaining the library to ensure ongoing accessibility of HDF-stored data.

According to Prabhat, many of the exascale research projects today primarily look at the I/O performance of a few select applications. But down the line, a broader base of users needs to be able to use the exascale machine. By scaling up the I/O infrastructure for the widely utilized HDF5 library, the ExaHDF5 team and HDF5 Group hope to put researchers from many different scientific disciplines on the exascale trajectory.

In addition to performing the largest-ever simulation at NERSC, the team also helped physicists Homa Karimabadi of the University of California, San Diego, and Bill Daughton, of the Los Alamos National Laboratory, identify how energetic particles are generated in magnetic reconnection. Reconnection is the mechanism behind the aurora borealis (a.k.a. northern lights) and solar flares, as well as fractures in Earth's protective magnetic field—fractures that allow energetic solar particles to seep into our planet's magnetosphere and wreak havoc in electronics,

power grids and space satellites. The scientists also discovered a power-law distribution in the particles' energy spectrum—a property that was expected to exist, but had not been validated before analysis of these simulation results. (A power law is a mathematical relationship between two quantities.)

"The outcome of this work was truly ideal," says Prabhat. "We ran a state-of-the-art simulation code at scale, which wasn't possible before, using the best computing resources and expertise, and this effort produced a first-time science result that no one had ever seen before. Computer science researchers always hope for such an outcome, but rarely do things come together in this fashion."

For future work, Byna hopes to apply the lessons learned from this hero run on Hopper to parallel I/O auto-tuning tools and techniques that will ensure that HDF5 effectively and efficiently scales up for the scientific applications that will use it most at large scale.

"We are really grateful for the critical support provided by NERSC consultants and system staff, as well as Cray engineers who assisted us with their large-scale runs," says Byna. "This work would not have been possible without DOE's Office of Advanced Scientific Research funding for the ExaHDF5 project, which has brought together a collaborative team of researchers in the HDF Group. We are also thankful for valuable support from Berkeley Lab's Visualization, Scientific Data Management, and Future Technologies groups."

Read more about the VPIC science results [Sifting Through a Trillion Electrons](#).

Provided by Lawrence Berkeley National Laboratory

Citation: Trillion particle simulation on hopper honored with best paper (2013, June 3) retrieved 9 April 2024 from <https://phys.org/news/2013-06-trillion-particle-simulation-hopper-honored.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.