


# Stampede supercomputer enables discoveries throughout science and engineering

June 20 2014, by Aaron Dubrow

---



Volume rendering of the entropy in a full 3-D GRMHD simulation of a differentially rotating and highly magnetized progenitor to a supernova. Red colors indicate high entropy (hot) material while blue represents low entropy (cold) material. Strongly magnetized material is continuously launched from the

surface of the proto-neutron star in the center but gets severely distorted such that, instead of a clean jet observed in the ultrastrong magnetic field case, two giant polar lobes are formed. The box size for the visualization is 2000km cubed. GRMHD refers to a general relativistic magnetohydrodynamic simulation. Credit: Philipp Moesta, TAPIR, California Institute of Technology

Sometimes, the laboratory just won't cut it. After all, you can't recreate an exploding star, manipulate quarks or forecast the climate in the lab. In cases like these, scientists rely on supercomputing simulations to capture the physical reality of these phenomena—minus the extraordinary cost, dangerous temperatures or millennium-long wait times.

When faced with an unsolvable problem, researchers at universities and labs across the United States set up virtual models, determine the initial conditions for their simulations—the weather in advance of an impending storm, the configurations of a drug molecule binding to an HIV virus, the dynamics of a distant dying star—and press compute.

And then they wait as the Stampede supercomputer in Austin, Texas, crunches the complex mathematics that underlies the problems they are trying to solve.

By harnessing thousands of computer processors, Stampede returns results within minutes, hours or just a few days (compared to the months and years without the use of supercomputers), helping to answer science's—and society's—toughest questions.

Stampede is one of the most powerful supercomputers in the U.S. for open research, and currently ranks as the seventh most powerful in the world, according to the November 2013 TOP500 List. Able to perform nearly 10 trillion operations per second, Stampede is the most capable of

the high-performance computing, visualization and data analysis resources within the National Science Foundation's (NSF) Extreme Science and Engineering Discovery Environment (XSEDE).

Stampede went into operation at the Texas Advanced Computing Center (TACC) in January 2013. The system is a cornerstone of NSF's investment in an integrated advanced cyberinfrastructure, which allows America's scientists and engineers to access cutting-edge computational resources, data and expertise to further their research across scientific disciplines.

At any given moment, Stampede is running hundreds of separate applications simultaneously. Approximately 3,400 researchers computed on the system in its first year, working on 1,700 distinct projects. The researchers came from 350 different institutions and their work spanned a range of scientific disciplines from chemistry to economics to artificial intelligence.

These researchers apply to use Stampede through the XSEDE project. Their intended use of Stampede is assessed by a peer review committee that allocates time on the system. Once approved, researchers are provided access to Stampede free of charge and tap into an ecosystem of experts, software, storage, visualization and data analysis resources that make Stampede one of the most productive, comprehensive research environments in the world. Training and educational opportunities are also available to help scientists use Stampede effectively.

"It was a fantastic first year for Stampede and we're really proud of what the system has accomplished," said Dan Stanzione, acting director of TACC. "When we put Stampede together, we were looking for a general purpose architecture that would support everyone in the scientific community. With the achievements of its first year, we showed that was possible."

## Helping today, preparing for tomorrow

When the National Science Foundation (NSF) released their solicitation for proposals for a new [supercomputer](#) to be deployed in 2013, they were looking for a system that could support the day-to-day needs of a growing community of computational scientists, but also one that would push the field forward by incorporating new, emerging technologies.



At nearly 10 petaflops, TACC's Stampede supercomputer provides comprehensive advanced computing capabilities to the nation's scientists. Credit: Sean Cunningham, Texas Advanced Computing Center

"The model that TACC used, incorporating an experimental component embedded in a state-of-the-art usable system, is a very innovative choice and just right for the NSF community of researchers who are focused on

both today's and tomorrow's scientific discoveries," said Irene Qualters, division director for Advanced Cyberinfrastructure at NSF. "The results that researchers have achieved in Stampede's first year are a testimony to the system design and its appropriateness for the community."

"We wanted to put an innovative twist on our system and look at the next generation of capabilities," said TACC's Dan Stanzione. "What we came up with is a hybrid system that includes traditional Intel Xeon E5 processors and also has an Intel Xeon Phi card on every node on the system, and a few of them with two.


The Intel Xeon Phi [aka the 'many integrated core (MIC) coprocessor'] squeezes 60 or more processors onto a single card. In that respect, it is similar to GPUs (graphics processing units), which have been used for several years to aid parallel processing in high-performance computing systems, as well as to speed up graphics and gaming capabilities in home computers. The advantage of the Xeon Phi is its ability to perform calculations quickly while consuming less energy.

"The Xeon Phi is Intel's approach to changing these power and performance curves by giving us simpler cores with a simpler architecture but a lot more of them in the same size package," Stanzione said

As advanced computing systems grow more powerful, they also consume more energy—a situation that can be addressed by simpler, multicore chips. The Xeon Phi and other comparable technologies are believed to be critical to the effort to advance the field and develop future large-scale supercomputers.

"The exciting part is that MIC and GPU foreshadow what will be on the CPU in the future," Stanzione said. "The work that scientists are putting in now to optimize codes for these processors will pay off. It's not

whether you should adopt them; it's whether you want to get a jump on the future. "



Picture DNA on Facebook. The image above is a map of links between the genes of the mustard plant *Arabidopsis thaliana*. Genes involved in the same biological process are connected by lines: red for more certain links, blue for less certain links. "It's not unlike a social network," says biologist Seung Yon Rhee. Credit: Insuk Lee, Michael Ahn, Edward Marcotte, Seung Yon Rhee, Carnegie Institution for Science


Though Xeon Phi adoption on Stampede started slowly, it now

represents 10-20 percent of the usage of the system. Among the projects that have taken advantage of the Xeon Phi co-processor are efforts to develop new flu vaccines, simulations of the nucleus of the atom relevant to particle physics and a growing amount of weather forecasting.

## Built to handle to big data

The power of Stampede reaches beyond its ability to gain insight into our world through computational modeling and simulation. The system's diverse resources can be used to explore research in fields too complex to describe with equations, such as genomics, neuroscience and the humanities. Stampede's extreme scale and unique technologies enable researchers to process massive quantities of data and use modern techniques to analyze measured data to reach previously unachievable conclusions.

Stampede provides four capabilities that most data problems take advantage of. Leveraging 14 petabytes of high speed internal storage, users can process massive amounts of independent data on multiple processors at once, thus reducing the time needed for the data analysis or computation.



A snapshot from the Penn State University real-time hurricane analysis and forecast system presents the 114-hour simulated-radar reflectivity and 10-minute wind speed forecasts for Hurricane Sandy (2012) before its landfall. Credit: Fuqing Zhang and Yonghui Weng at Penn State University

Researchers can use many data analysis packages optimized to run on Stampede by TACC staff to statistically or visually analyze their results. Staff also collaborates with researchers to improve their software and make it run more efficiently in a high-performance environment.

Data is rich and complex. When the individual data computations become so large that Stampede's primary computing resources cannot handle the load, the system provides users with 16 compute nodes with one terabyte of memory each. This enables researchers to perform complex data analyses using Stampede's diverse and highly flexible computing engine.


Once data has been parsed and analyzed, GPUs can be used remotely to explore data interactively without having to move large amounts of information to less-powerful research computers.

"The Stampede environment provides data researchers with a single system that can easily overcome most of the technological hurdles they face today, allowing them to focus purely on discovering results from their data-driven research," said Niall Gaffney, TACC director of Data Intensive Computing.

Since it was deployed, Stampede has been in high demand. Ninety percent of the compute time on the system goes to researchers with grants from NSF or other federal agencies; the other 10 percent goes to



industry partners and discretionary programs.



Peering into a colloidal gel, from the research group of Roseanna N. Zia, Cornell University. Dynamic simulation of the structural evolution of 750,000 attractive hard spheres. Credit: Roseanna N. Zia, Cornell University

"The system is utilized all the time—24/7/365," Stanzone said. "We're getting proposals requesting 500 percent of our time. The demand exceeds time allocated by 5-1. The community is hungry to compute."

Stampede will operate through 2017 and will be infused with second generation Intel Xeon Phi cards in 2015.

With a resource like Stampede in the community's hands, great discoveries await.

"Stampede's performance really helped push our simulations to the limit," said Caltech astrophysicist Christian Ott who used the system to study supernovae. "Our research would have been practically impossible without Stampede."

Provided by National Science Foundation

Citation: Stampede supercomputer enables discoveries throughout science and engineering (2014, June 20) retrieved 28 April 2024 from <https://phys.org/news/2014-06-stampede-supercomputer-enables-discoveries-science.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.