

Supercomputing Goes Back to School

Aaron Dubrow, TACC



It's back-to-school season and, here at the Texas Advanced Computing Center, we're back to school, too. For us, that means teaching scientific computing classes at The University of Texas at Austin to eager computational scientists-in-training, and leading in-person and [Virtual Workshops](#) [1] in parallel programming and other facets of high-performance computing for thousands of researchers across the nation.

TACC is a leader in this respect, offering more than 50 training sessions each year in addition to teaching five courses at The University of Texas at Austin that are required to earn a certificate in [Statistics and Scientific Computing](#) [2]. TACC's experts in hardware, software, visualization and data-driven science teach these courses and training classes, which span a wide range of levels. Workshops range from a few hours to a few days to the weeklong Summer Supercomputing Institute, now in its seventh year.

On September 19-20, TACC offered its 10th training session focused on TACC's newest supercomputing system, Stampede, which came online in January 2013 and has already completed more than 1.5 million jobs on over 2,000 research projects. Stampede is currently the sixth largest computer in the world and the largest available to academic researchers in the U.S.

The "*Parallel Computing on Stampede*" workshop is intended for beginner and intermediate HPC users and anyone interested in using Stampede in the future. The two-day session was offered in-person and via webcast to accommodate users across the country.

"We teach our users everything they need to know to get started on the system quickly and to effectively and efficiently use it in the support of their research

Supercomputing Goes Back to School

Published on Scientific Computing (<http://www.scientificcomputing.com>)

projects," said Chris Hempel, director of user services at TACC. "Training the user community how to use HPC resources is a contributing factor to the success of TACC."

The training is free to the UT Austin, UT System, TACC partners, and Texas higher education research communities, as well as to participants in the National Science Foundation-funded XSEDE (Extreme Science and Engineering Discovery Environment) project. Individuals and groups that engage in collaborative research and development activities with TACC, and other public sector organizations including TACC's industrial affiliates, are welcome to participate as well. TACC training classes are also available for a fee to interested persons who are part of the private sector.

For most researchers, Stampede is larger, faster and more comprehensive than any computing system they have used before. Much about computing on Stampede is familiar to scientists who have used Ranger, Lonestar, Longhorn and other national supercomputers in the past, but the system also incorporates a new processing architecture — the Intel Xeon Phi many integrated core (MIC) coprocessor — that is a novel element even to experienced supercomputer users. The Xeon Phi is programmed using the same computer languages as traditional CPUs, but it has some key differences that require basic instruction and an understanding of the new paradigm. For new users, or users looking to take advantage of the new capabilities Stampede offers, some training is required.

Eric Fahrenthold, a mechanical engineering professor at The University of Texas at Austin and a long-time supercomputing user, attended a Stampede workshop on the MIC architecture last spring. Attending the session provided an on-ramp for using the MIC on Stampede — he didn't have to teach himself how to use the MIC by studying the web or white papers.

"TACC explained the MIC architecture and showed me how to test my code's performance when using the MIC on Stampede," he said. "I thought it was great. I found out very quickly what the code performance issues are, and the various ways in which I can access Stampede's coprocessor cores."

TACC began offering Stampede training last November with a special session at the Supercomputing'12 conference in Salt Lake City. At the [SIAM Conference on Computational Science & Engineering](#) [3] in March, TACC led a mini-symposium for 50 researchers. Later training sessions at the U.S. National Conference on Computational Mechanics, the International Supercomputing Conference (ISC) in Germany, and at the XSEDE'13 conference significantly expanded the number of researchers TACC trained.

It's been a whirlwind tour for many of the Stampede experts at TACC. But it isn't only staff in Texas who have been participating in the effort. The Cornell University Center for Advanced Computing, an academic partner in the Stampede project, began offering training for local and regional researchers, as well as remote participants, in Fall 2012. Cornell specializes in the development and dissemination of training materials and offers up to 64 hours of Stampede-specific training each

Supercomputing Goes Back to School

Published on Scientific Computing (<http://www.scientificcomputing.com>)

year, reaching hundreds of researchers in a variety of fields.

“Biology, applied math, nutritional sciences, civil engineering, earth and atmospheric science, mechanical engineering, pharmacology — our attendees are very diverse,” said Susan Mehringer, assistant director of the Cornell University Center for Advanced Computing (CAC), who leads the effort there. “We explain the concepts that are generally applicable to high performance computing, but all of the specifics they’re learning pertain to Stampede.”

The workshops allow the local and regional community to access one of the largest HPC systems in the country, even if they have little experience with the nation’s cyberinfrastructure.

“When researchers’ needs exceed their local resources, there are national large computing resources available,” Mehringer said. “It’s very helpful for people to learn locally while using a national resource. Within 30 minutes of being at the workshop, they’re logged into Stampede.”

The training is designed to introduce researchers to the system and the new Intel Xeon Phi architecture in a practical manner. Staff present lectures and hands-on exercises to get users acquainted with the new platform and to explore different execution modes. The training also helps researchers learn about parallelization and optimization through example testing and reports. Users are invited to bring their own codes to compile for the Xeon Phi.

For Fahrenthold, who has been using Stampede to study impact physics on body armor and spacecraft materials, learning from knowledgeable HPC experts means he can focus more of his efforts on the science.

“Even though I have a lot of experience with HPC, learning about the MIC system’s hardware and software was new,” he said. “I like to take advantage of TACC consulting and tutorials, because it allows me to concentrate on modeling the physics — when I need to understand system architecture and system software issues I rely on TACC support.”

To learn about future workshops at TACC, Cornell, and elsewhere, visit the [TACC training page](#) [4] and sign up for training updates.

Source URL (retrieved on 01/31/2015 - 11:11am):

<http://www.scientificcomputing.com/news/2013/10/supercomputing-goes-back-school>

Links:

[1] <https://www.cac.cornell.edu/VW/>

[2] <http://ssc.utexas.edu/>

[3] <http://www.siam.org/meetings/cse13/>

[4] <http://www.tacc.utexas.edu/user-services/training>

Supercomputing Goes Back to School

Published on Scientific Computing (<http://www.scientificcomputing.com>)
