

# Computer scientists collect computing tools for next-generation machines

February 15 2012, by Eric Gedenk

---

(PhysOrg.com) -- Researchers using the OLCF's resources can foresee substantial changes in their scientific application code development in the near future.

The OLCF's new [supercomputer](#), a Cray XK6 named Titan with an expected peak speed of 10-20 petaflops (10-20 thousand trillion calculations per second), will use a [hybrid architecture](#) of conventional, multipurpose central [processing units](#) (CPUs) and highperformance graphics processing units (GPUs) which, until recently, primarily drove modern [video game graphics](#). Titan is set to be operational by early 2013. The machine will supplant the OLCF's current [fastest supercomputer](#), Jaguar, a Cray XT5 using an entirely CPU-based platform.

With Titan's arrival, fundamental changes to computer architectures will challenge researchers from every scientific discipline. Members of the OLCF's Application Performance Tools (APT) group understand the challenge. Their goal is to make the transition as smooth as possible.

"The effort necessary to glean insight from large-scale computation is already considerable for scientists," computational astrophysicist Bronson Messer said. "Anything that tool developers can do to reduce the burden of porting codes to new architectures, while ensuring performance and correctness, allows us to spend more time obtaining scientific results from simulations."

The APT group is working to ensure that researchers receiving allocations on leadership-class [computing resources](#) will not have to spend huge amounts of time learning how to effectively use their codes as the OLCF shifts to hybrid computing architectures.

An application tool can do anything from translating one computer programming language into another to finding ways to optimize performance. "We decide which tools—pieces of software that enable application scientists to perform their simulations in an effective manner—are of interest to us, and whether or not these current tools are sufficient; and by sufficient I mean in the context of our production environment, Jaguar," said Richard Graham, group leader of the APT group at the OLCF. "If they are not sufficient, we need to understand if we can enhance the current tool set, or if we need to go out and see if there is something else out there."

Graham explained that many of the same tools that helped Jaguar advance beyond the petascale threshold will also play imperative roles in getting Titan off and running toward the exascale—a thousand-fold increase over petascale computing power. However, just like scientific application codes, software tools will have to be scaled up for a larger machine.

## **Hammering out details**

The group has expanded the capabilities of debugging software to meet the needs of large, leadership-class systems. Debuggers are capable of helping to identify programming glitches in users' application codes. Until recently, though, debugger software did not scale up to supercomputer architectures well. The APT group has worked with British company Allinea on the DDT debugger for several years to bring debugging up to scale. "People claimed that you could not do debugging beyond several hundred to several thousand processes," Graham said.

"Now, my group routinely debugs parallel code at more than 100,000 processes."

Another major issue facing supercomputers is collecting and analyzing performance data. The tools group is collaborating with the German Technische Universität Dresden on the "Vampir" suite of trace-based performance analysis tools. These tools look at the performance of an application run from both macroscopic and microscopic lenses. In addition to examining the program counter, which directs the traffic of the functions a computer should be executing, they are capable of tracing back through all of the call stacks, which help organize the sequence of operations for the measured performance. Just as the group scaled debugging to new heights, they have done the same with performance analysis. "Basically, we've been able to run trace-based applications at 200,000 processes, and the previous record was on the order of 100,000," Graham said.

Perhaps the largest challenge related to hybrid architectures comes from shifts in programming languages. Both Graham and computer scientist Greg Koenig agree that compilers were one of the most important tools for scaling computer applications to the petascale, and they will undoubtedly play an imperative role in integrating scientific applications into hybrid architectures. Compilers play the role of translators, taking programming languages such as C++ and Fortran and converting them to serial-based languages—binaries—that a computer can understand. "One of the major challenges of computer science is compiling codes from one category to another," Koenig said. Graham explained that for the last two years, the APT group has been collaborating with CAPS Enterprise, a French company producing the Hybrid Multi-Core Parallel Processor (HMPP) compiler. Two years' worth of collaboration has started to yield results. "The work has resulted in significant capabilities being added to the compiler that help us incrementally transition our existing applications to accelerator-based computers and has started to lead to

some performance enhancements," he said.

National laboratories such as Oak Ridge could help foster the transition to hybrid computing. Koenig noted that businesses often have harsh deadlines for turning a profit, and academic institutions must often publish findings before they begin heavily collaborating with other institutions.

"I like to think that between the two of these, the national labs fit." Koenig said. "You have a lot of unique capabilities in the national lab system that you do not get in other places. You can think a little more abstractly than you can than if you were in a business setting. We have the luxury of being able to look further into the future, and we don't have to turn a profit immediately. We are doing research, but it is research that can be applied for solving real problems, and that's particularly true in the Application Tools group, because we are trying to address real needs of computational scientists."

Provided by Oak Ridge National Laboratory

Citation: Computer scientists collect computing tools for next-generation machines (2012, February 15) retrieved 20 April 2024 from <https://phys.org/news/2012-02-scientists-tools-next-generation-machines.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.