Bright Computing

Engineering Materials Using GPU-based HPC Clusters at Tokyo University of Agriculture and Technology

CASE STUDY

Researcher scientists use powerful HPC clusters to perform advanced materials research.

"As scientists, we always are time constrained. Our priority is our research, not managing our clusters. Bright is intuitive to use and, with it, I can effectively manage my cluster without wasting time writing scripts, or synchronizing management tool revisions. Provisioning is fast and easy. I prefer this approach over open source toolkits."

—Dr. Akinori Yamanaka, Associate Professor, Institute of Engineering

Overview

Very strong, yet lightweight materials are crucial to automotive and aerospace manufacturers, particularly if they work in extreme temperatures. Developing such materials requires a deep understanding of their underlying properties. And understanding those properties requires time-consuming computational processes on a massive scale.

Dr. Akinori Yamanaka, an associate professor with the Institute of Engineering at the Tokyo University of Agriculture and Technology (TUAT), models and simulates engineering materials, for a variety of applications, using High Performance Computing (HPC). Materials science research is computationally demanding, even for leadership-class supercomputers such as Japan's Tokyo-tech Supercomputer and Ubiquitously Accessible Mass-storage Environment (TSUBAME) at the Tokyo Institute of Technology (TITECH). As a consequence, Dr. Yamanaka and his collaborators continue to make the design and implementation of algorithms for GPU-based HPC a significant component of their research. In 2011 together with then TITECH colleagues Drs. Shimokawabe, Aoki and Matsuoka, Dr. Yamanaka's innovative use of HPC for materials science research was recognized through the Association for Computing Machinery's award of the prestigious Gordon Bell Prize for their paper titled, "Peta-scale Phase-Field Simulation for Dendritic Solidification on the TSUBAME 2.0 Supercomputer."

The Challenge

For Dr. Yamanaka, multiple time-evolution equations of order parameters (partial differential equations for phase-field variables) must be solved, but performing realistic three-dimensional simulations comes with a very high computational cost.

Even for researchers with a reputation for award-winning results, full-scale access to supercomputers like TSUBAME is limited, as demand for HPC resources well outstrips supply. This places an additional burden of responsibility onto researchers to carry out their modeling and simulations efficiently. Typical of many leadership-class supercomputers, TSUBAME is a hybrid-architecture system that makes use of both CPUs and GPUs. Each release of development-enhancing tools and utilities (e.g., the HPC compilers, debuggers, and profilers plus CUDA-enabled libraries available from NVIDIA's CUDA toolkit) introduces new possibilities for those developing algorithms for modeling and simulating various phenomena in materials.

Constrained access to an evolving computational platform can be challenging, and it may be tempting to bypass TSUBAME-scale supercomputers using an in-house HPC system. Although unlimited access to a controlled configuration platform sounds appealing, management of even the smallest of HPC clusters takes valuable time away from academic research.

The Solution

Bright Computing partner HPCTech Corporation worked with Dr. Yamanaka's research associates in order to maximize the group's potential for materials science results using a scaled-down version of TSUBAME managed by Bright Cluster Manager. This enabled the group to take advantage of both leadership-class supercomputers and a functional HPC cluster.

Bright Cluster Manager is a critical component of the overall solution—provisioning, monitoring, and managing the HPC cluster in Dr. Yamanaka's lab. Researchers now focus on their work instead of managing IT infrastructure.

Bright provides full support for HPC via GPUs. From drivers that render GPUs programmable via CUDA, to the CUDA-enabled tools, to CUDA-optimized libraries critical to materials science research, Bright Computing delivers a comprehensive platform that is research ready. Maintaining the lab's cluster is simple and quick with Bright taking care of change management and all the complexity inherent in hybrid-architecture systems.

Using Bright Cluster Manager, Dr. Yamanaka and his team have support for the latest Kepler-architecture GPUs, as well as CUDA 6.5, through routine YUM updates. Researchers making use of the lab cluster easily can switch between different versions of the CUDA toolkit using Environment Modules—open-source software bundled with Bright that allows researchers to craft and maintain highly customized development environments.

> "Bright is particularly beneficial for program developers—especially CUDA developers."
>
> — Hiroshi Sato, HPCTech Engineering Manager

The Results

Time is a researcher's most precious commodity. Bright Cluster Manager allows Dr. Yamanaka and his team to focus on the engineering of materials through high-resolution, three-dimensional phase-field modeling using GPU-based HPC—not spending time managing IT infrastructure.

Because the lab cluster at TUAT is managed with Bright Cluster Manager, researchers can devote their time to prototyping their latest algorithms, models, and simulations. Using their in-house cluster, researchers now can devote more time to:

- Eliminating defects in their code (aka bugs) through use of ad hoc techniques, as well as debugging tools from NVIDIA and others;
- Improving run-time performance through tools that allow them to profile from individual algorithms to entire models and simulations;
- Benchmarking scaling characteristics of algorithms in isolation, or entire models and simulations.

Time saved with Bright Cluster Manager translates directly to greater quality time spent on the supercomputer, as the Yamanaka group's proven prototypes run efficiently at the full scale of TSUBAME.

The engineering accomplishments achieved by this two-stage approach are impressive:

- Results from the Yamanaka Lab computational models and simulations compare well with those obtained from physical experiments. This, of course, is absolutely critical, as there is a need to make use of proven models and simulations in parameter spaces not easily accessible through physical experiments;
- The models and simulations developed by the Yamanaka Lab are extremely efficient. In making use of GPUs for HPC, in some cases, they exhibit logarithmic scalability. Not only does this continue the tradition that won Dr. Yamanaka and his collaborators the Gordon Bell Prize in 2011, this achievement underscores the ongoing effort expended in algorithm development.