

Research Alliance in Math and Science

<http://www.csm.ornl.gov/Internships/RAMS.html>

Debbie McCoy, Computing and Computational Sciences Directorate
Oak Ridge National Laboratory

Summary

The Research Alliance in Math and Science (RAMS) program is carried out through the Computing and Computational Sciences Directorate at the Oak Ridge National Laboratory (ORNL) for the U. S. Department of Energy's Office of Advanced Scientific Computing Research. The RAMS program continues to provide unique, hands-on educational experience through innovative approaches to underrepresented students majoring in mathematics, computer science, engineering technology, and the computational sciences.

RAMS internship 'really challenged my intellect,' CSM intern says

Tara McQueen's first experience in the Research Alliance in Math and Science (RAMS) program at the Oak Ridge National Laboratory (ORNL) in summer, 2008 helped her to put it all together — her lifelong interest in computing and her hopes for a future career in that field.

"This top level internship allowed me to take all that I have learned over the years and direct it to a specific goal. It really challenged my intellect," said

Ms. McQueen, who has just completed her junior year at Delaware State University, where she majors in Information Technology, with a minor in criminal justice.

In fact, she got so much from the program, she is returning for another session in summer, 2009 and she is bringing some of her classmates with her. "I enjoyed every bit of my RAMS experience, which is why I have introduced friends into the program and am coming back for another term."

Ms. McQueen interned in the Computer Science and Mathematics (CSM), where she was mentored by Dr. Stephen Scott. "It was important in my life because I learned so much," she said. For her research subject in 2008, she worked on integrating new software for the science applications which researchers run on high performance computing platforms. One of the key challenges in high performance computing is to ensure that the software which researchers use to simulate and solve their problems on the computer performs efficiently. The diversity of the parallel applications software and the rapidly increasing size of the computing platforms on which they run is a major challenge for evaluating performance—and such performance evaluations have to be revised, each time there is a different mapping of application and platform.

The primary goal when carrying out a full performance evaluation experiment is to improve the user's productivity. Until now, the whole process of performance evaluation continues to be carried out manually. In this project, researchers at CSM proposed an

* (865) 574-6185, mccoydd@ornl.gov

easy-to-use framework which would capture the essentials of the infrastructure which are needed for a given evaluation. The performance evaluation process mainly consists of preparing the required environment, starting a performance diagnostic tool, executing a given application, and analyzing the results.

Automating the Run-time Infrastructure for the Performance Evaluation Process
Tara L. McQueen
Delaware State University
Research Alliance in Math and Science
Computer Science and Mathematics Division
Oak Ridge National Laboratory
Mentors: Dr. Stephen L. Scott, Dr. Geoffrey Vallee, Anand Tikolekar

Motivation

- Large effort is required to manually carry out the necessary tasks for the performance evaluation process
- Available tools lack extensive documentation on integrating new benchmark tools
- Effective system management/monitoring is required to ensure that performance evaluations can be carried out
- There is a growing disconnect between performance evaluation software and management/monitoring software
- How can this performance evaluation process be streamlined?

Introduction

- Application performance evaluation
- Prepare benchmark execution environment
- Benchmark execution
- Performance report generation/analysis
- System management/monitoring
- Allocate resources
- Node management
- Job scheduling

System Management - OSCAR

- Open Source Cluster Application Resources (OSCAR)
- Cluster suite for high performance computing
- Simplifies install, build, and operation of cluster
- Cluster management via Graphical User Interface (GUI)
- OSCAR extended via user contributed packages
- Would be beneficial to have a performance evaluation OSCAR package

Performance Evaluation - Cbench

- Facilitate cluster benchmarking and testing toolkit
- Collection of tests, benchmarks, applications and utilities
- Currently provided benchmark tests
- IPCC
- Linpack
- NPB
- Automates tasks required for performance evaluation
- Generation/resolution of jobs
- Data parsing
- Provides interpretable reports for new benchmarks
- Adds resource allocation, system monitor, and job scheduler
- No extensive documentation for integration of new benchmarks

Goals

- Integrate a benchmark tool (Cbench) and system management framework together (OSCAR)
- Add new benchmarks into Cbench and documentation

Prototyping

- Selected LAMMPS (large-scale atomic/molecular massively parallel simulator) for integration into Cbench
- Modified generic template scripts to generate jobs, start jobs and parse output of the LAMMPS benchmark
- Parser is unique to the specific benchmark
- Requires additional custom logic
- Created how-to guide for instructions of parser development
- Tested on a cluster with 18 nodes
- Required effort is minimal

Results

- Integration of LAMMPS into Cbench
- Structured How To guide for new benchmark integration into Cbench
- Benchmarking framework for OSCAR
- Provided OSCAR users with adequate scaffolding and benchmarking resources
- Integrated further benchmark integration into OSCAR by documentation

“It was very gratifying to be able to share the impact that RAMS has made on us, as well as some of the funny stories. We got to compare our internship experience with some other individuals who had internships in California, Pittsburgh, and Alaska. “Those interns said that the RAMS program seemed better than the programs they participated in, and that Tennessee seems to provide a sufficient amount of activities for away from work play—yet still has little distraction to work performance.

“It was during this time that I realized the level of competition out there and the severity of seeking available opportunities and going above and beyond to better myself through education and experience,” Ms. McQueen said.

“I got to see some of the major vendors associated with the area in which I did my research,” she said. “Working the outreach booth for ORNL really made me see the uniqueness of Tennessee, the lab, and the opportunities that RAMS has provided us.”

For further information on this subject contact:

Mr. George Seweryniak, Program Manager
Office of Advanced Scientific Computing Research
Phone: 301-903-0071
seweryni@er.doe.gov

After her RAMS summer research, Ms. McQueen attended the SuperComputing 2008 conference in Austin, Texas last November, where she and RAMS intern Jessica Traverso staffed the ORNL booth with RAMS program director, Debbie McCoy.

“SC08 was great. Jessica and I helped in the ORNL booth, where students would visit to gather information about future job opportunities and internships. A lot of the students asked us about our experience with the RAMS program—how the internship went and what it did for us.