BLUE WATERS

Blue Waters Education Allocations

Final Report

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

As a component of your application for a Blue Waters education allocation, you agreed to provide a report at the end of your project. Now that your project is complete, we would appreciate your submission of the following information within two weeks. Please send this report, and any supporting documents to https://bluewaters.ncsa.illinois.edu/eduallocsubmit.

This information will be shared with the Blue Waters team and the National Science Foundation. Portions of the report (we will omit names of participants) will be posted on the Blue Waters portal for public access.

Project Information

Project Name	**Practical experience and workforce development on sustained peta-scale computing systems: A workshop exploring scalable multi-physics software for multi-thousand CPU applications**

Names of project staff (instructors, TAs, etc) and their department and institutions	**David Ackerman, Department of Mechanical Engineering, Iowa State University** **Baskar Ganapathysubramanian, Department of Mechanical Engineering, Iowa State University**
URL for the project	N/A
Provide links to or attach materials made available to participants (e.g. slides, articles, exercises, etc.) that may be made publicly available	
Provide links to or attach any photos (with captions to describe activities)	
Start date	August 5, 2014
Completion date	August 31, 2014

Information about the Participants

# Participants	# Faculty or staff	# students	# other (e.g. industry)	# under-represented (e.g. women, minorities)	# institutions represented by participants
11	5	5	1	2	4

Participants included faculty/visiting faculty (Iowa State, U Buffalo, Czestochowa University, Poland), Post Docs (Iowa State), PhD students (Iowa State) and Undergraduate Students who were part of a NSF REU site (Grinnell College, Iowa State). This included two women (Post Doc, and undergraduate student).

Please describe the scope and purpose of this project. Also, please indicate if there were any changes implemented from the original proposed plan, and briefly describe why they were made.

This workshop was designed for students with an existing background and experience with MPI programming an opportunity to explore scaling onto large numbers of processors using an in-house open-source FEM library that is based on PETSC. The goal of the two week workshop was to give students with existing scientific and computing background an understanding of, and practical experience with, scaling software to large numbers of processors. The resources used are significantly beyond what they would have access to within the context of a college parallel programming course. They will gain hands-on

experience in developing applications that scale to thousands of nodes as well as the ability to test existing applications on a petascale system. Following this workshop, students will be better prepared to design applications to efficiently utilize large scale computing resources for topics including computation fluid dynamics, finite element analysis, and polymer science. This experience will better prepare them to use petascale resources in the future.

The dates of the workshop were changed due to participant offices and equipment moving during the original planned time. We also asked for a time extension because some of the very large runs were stuck in the queue behind someone using most of the machine for a couple large, long-running, jobs.

Please describe the learning outcomes of the participants. How did this project enhance the learning of the participants? What did the participants learn as a result of the use of Blue Waters system that they could not have learned using other systems?

Learning Outcomes:

1) Knowledge and training in remote login to HPC resources

2) Ability to compile open source software libraries like PETSc on HPC resources

3) Configure, compile, and run in-house MPI based parallel FEM library on BW

4) Learn and explore various compilation and optimization flags on a HPC machine like BW

5) Perform scaling studies on selected test problems on BW

6) Explore code-optimization and scaling for a set of individual problems.

Each of the learning outcomes was accomplished though specific tasks (given in the proposal).

While local resources can allow computations using 100s to 1000s of CPUs, there are no resources available for larger scale runs in the range of tens of thousands of CPUs. Awareness of the scaling considerations on these large numbers of processors is important for research the students will be pursuing in the future. Without the knowledge and experience at that scale, they will lack the ability to efficiently utilize current and future HPC resources for very large problems. Use of Blue Waters has allowed a practical application of wide scaling that is difficult or impossible to achieve using other available HPC resources. Participants learned how to evaluate scaling behavior of code on very large numbers of nodes. They are better prepared to write efficient and scalable code targeting tens of thousands of processors. This level of scaling would not have been possible with a large system such as Blue Waters. Participants also learned to use the Blue Waters system and learned more about optimizing code for performance on large scale systems.

Please describe lessons learned from the project. What would you do differently next time?

1) We learned that some of the code we used did not work correctly on Cray Optimized compilers. Unfortunately we didn't discover that until the start of the workshop which lead to a delay in starting. In the future, we would test the code in detail on the system prior to the workshop to sort out any system specific problems. Furthermore, this has now been resolved and will not be an issue with future workshops.

2) BW has a wide array of compiler and optimization options. This was different from the other HPC machines we have used. We were able to systematically test and perform timing, memory and scaling tests using a variety of these flags to identify options that were the most effective.

What would you recommend that the Blue Waters team do to enhance the success of education projects in the future?

The overall experience with BW was exceedingly positive. In the context of small SU allocation education grants, we have two suggestions:

1) The usage information was only rarely updated – it seemed once or twice a week. For an education workshop which runs only a couple weeks, this makes it very difficult to ensure that we are using the expected amount of node hours for each part of the workshop. Having the usage updated at least daily would be a great help.

2) It would also be useful if there was a way to place limits on how many node hours specific individuals can use. With a small allocation, one user's mistake could burn through most of the allocation and ruin the education project for the rest of the participants.

Please provide a summary of any surveys or evaluations you conducted of the participants. Feel free to attach any related documents.

Participants completed a survey of their experiences following the workshop. All participants agreed that the Blue Waters system was easy to use and a very good fit for the goals of this workshop. Their experience using Blue Waters compared favorably to experiences with other HPC systems they have previously used. Most users consulted the online documentation and found it very informative, especially the description of the module system and how to run jobs. However, several people noted that the section on compiler usage and optimization was much less complete which made it difficult to deal with problems compiling third party libraries. The only other concern raised was the wait times in the queue. For most of the workshop, jobs ran promptly when placed in the queue. However, for several days the queue was stalled due to multiple consecutive, 12 hour jobs which used 20,000 of XE nodes. Users expressed frustration that these jobs effectively blocked usage of the system for several days.

Please provide any anecdotal stories we may share with NSF and the public.

N/A

How would you rank the overall experience?

	Excellent	Very Good	Good	Fair	Poor	N/A
Education allocations process	X					
Blue Waters support	X					
Blue Waters computing system	X					
Blue Waters documentation		X				
Blue Waters training			X			

Do you plan to request an education allocation for other future events that will use BW?

Please describe the plans for future events, including the frequency (each semester, yearly, etc.).

Yes, the feedback from the attendee's was overwhelmingly positive. This workshop served as a good forum to train computational scientists on the use of top HPC resources.

We plan continue organizing this workshop yearly. This will enable new and incoming students, post-docs and REU students within the ISU community to explore HPC resources for simulation based science and engineering.

Please provide any other comments or suggestions.