

SIGCSE Special Project Grant

Final Report - May 2016

Title of Project: OnRamp: An Interactive Learning Portal for Parallel Computing Environments

Applicant Name(s): Samantha Foley and Joshua Hursey

Problem Statement

Students must understand parallel and distributed computing concepts, as reflected in the ACM Computer Science Curriculum guidelines. Unfortunately, parallel computing environments (PCEs) have a high barrier to entry for even the most motivated students and educators. Historically, obtaining and configuring the hardware and software to create a PCE for educational purposes was an expensive and labor intensive task. Recent hardware and software projects targeted at the educational community have made it easier to obtain a small physical, virtual, or cloud based PCE. Concurrently, the CS education community has been creating and distributing educational modules to support educators that wish to incorporate parallel computing concepts into their curriculum. Several projects were highlighted at the “Budget Beowulf” Special Session at SIGCSE 2015 [1], including curriculum modules provided by the Shodor Foundation¹ and CSinParallel².

However, there still exists a significant barrier to entry for learning how to become productive in a PCE due to often unfamiliar and complex system software, programming interfaces, and tools. For example, users must navigate the system via a terminal, deal with complex compilation and execution procedures, negotiate resource allocation via batch submission systems, and utilize parallel programming libraries – a process that may be new and intimidating to students and educators. Our project, OnRamp to Parallel Computing, provides a web-based portal that coaches users through interactive tutorials that teach them about the software ecosystem and parallel computing while allowing them to launch parallel applications from day one. As users become more comfortable with running parallel applications on PCEs, the OnRamp portal transforms into a reference guide.

Project Description

The OnRamp portal will allow users to interact with, configure and run parallel programs in a familiar, yet informative and powerful way. The primary goal of the portal is to assist users who are new to parallel computing explore parallel computing concepts and common software tools in a layered and self-guided approach. As users move through the learning activities they develop the knowledge and confidence required to directly and effectively use a variety of PCEs. After completing those learning activities the web portal becomes a reference as they move to the native platform.

Completed Work

This grant funded the PIs and an undergraduate student to develop the core functionality of the three major pieces of the software, and two basic curriculum module for the purposes of demonstrating that the infrastructure works. This work included the development of the web frontend, web backend and database, as well as the web service that runs on the PCE to launch jobs. The bulk of this work involved the design and implementation using CherryPy as the web server and PCE service which launches jobs on behalf of the user. In addition to the ability to launch jobs, the admin capabilities of: managing connections to PCEs; creating and managing users; and organizing users, PCEs, and modules into workspaces was designed and developed.

¹<http://www.shodor.org/petascale/materials/modules/>

²<http://serc.carleton.edu/csinparallel/index.html>

During the course of this effort, three more students were recruited and contributed to the project, including one Blue Waters scholar from the University of Wyoming. These students worked together with the PIs to take the basic functionality developed in summer and fall 2015 to an alpha-release state with all major functionality complete. At the end of May 2016, the product was working and a user-study was conducted. The work was disseminated in several venues over the course of the funding period:

- Foley, S., Hursey, J. (2015). OnRamp to Parallel and Distributed Computing. EduHPC- 15: Workshop on Education for High-Performance Computing held in conjunction with SC-15: The International Conference on High Performance Computing, Networking, Storage, and Analysis. (Workshop Paper and Presentation)
- Foley, S., Hursey, J. (2016). OnRamp to Parallel and Distributed Computing: A Web-portal for teaching parallel and distributed computing. Proceedings of the 47th ACM Technical Symposium on Computing Science Education. (Demonstration)
- Foley, S., Hursey, J., Koepke, D., Ragatz, J., Regina, J. (2016). OnRamp to Parallel and Distributed Computing: Web-portal for teaching parallel and distributed computing. Proceedings of the 47th ACM Technical Symposium on Computing Science Education. (Poster)
- Foley, S., Koepke, D., Brehm, C., Ragatz, J., Regina, J. (2016). OnRamp to Parallel and Distributed Computing. Proceedings of the 49th Annual Midwest Instruction and Computing Symposium. (Demonstration)
- Regina, J., Foley, S. (2016). Finite Water-Content Module: Teaching Hydrology and HPC Using OnRamp. Proceedings of the 5th Annual XSEDE Conference. (Poster)
- Foley, S., Brehm, C., Koepke, D., Ragatz, J., Regina, J., Hursey, J. (2017). OnRamp to Parallel and Distributed Computing. *Journal of Parallel and Distributed Computing*, 105, 138–149. (Journal Article)

Continued Work

After reflecting on the implementation of the project and a user study that was conducted, we decided that the project needed to be reimplemented in a different web-framework in order to address several stability, security, and scalability issues. Due to changes in personnel for the project (students graduating) the next set of students to work on the project started to migrate to Django and MySQL for the back-end server implementation. This work has continued off and on over the next few years as students joined and left the project. A WiSys grant funded the most recent development during the 2017–2018 academic year.

The project continues to be developed and if you are interested in the latest developments, please contact Dr. Samantha Foley.

References

- [1] J. C. Adams, J. Caswell, S. J. Matthews, C. Peck, E. Shoop, and D. Toth, “Budget Beowulfs: A showcase of inexpensive clusters for teaching PDC,” in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE ’15. New York, NY, USA: ACM, 2015, pp. 344–345. [Online]. Available: <http://doi.acm.org/10.1145/2676723.2677317>