

Information Science and Technology Seminar Speaker Series



Kate Keahey

Argonne National Laboratory

Chameleon: A Deeply Reconfigurable, Large Scale Instrument for Computer

Wednesday, July 12, 2017

2:00pm - 3:00pm

TA-3, Bldg. 1690, Room 102 (CNLS Conference Room)

Abstract: The last few decades saw an enormous amount of progress in developing the potential of distributed computing: cloud computing, edge computing, and the Internet of Things all inspire new ideas across the stack, from systems solutions to innovative applications. However, complex and distributed systems of this type – or their components – are hard to study using theoretical approaches alone. To study them properly, scientists require a deeply reconfigurable platform where they can explore phenomena in controlled environments, across a range of configurations, and at scales needed to justify their result. Like physicists, biologists, and chemists, computer scientists need an experimental instrument that can serve as a vehicle of discovery.

The Chameleon project has built such an instrument. The testbed, deployed at the University of Chicago and TACC, consists of roughly 15,000 cores, 5PB of total disk space, and leverages 100 Gbps connection between the sites. The backbone of the testbed is configured as a large homogenous partition to support experiments at scale -- on this framework were grafted heterogeneous elements including Infiniband networking, high-bandwidth I/O nodes, storage hierarchy nodes, and GPUs to support a broad diversity of research projects -- and that diversity is further enhanced by smaller clusters of FPGAs, ARM and Atom processors. To support Computer Science experiments, ranging from operating system and virtualization research to innovative applications, Chameleon provides a configuration system giving users full control of the software stack: provisioning of bare metal, support for custom kernel reboot, and console access — but also a fully functioning cloud environment to support educational projects and cloud development.

This talk will describe the goals, the design strategy, and existing and future capabilities of the testbed. I will also present a selection of the research and education projects conducted by our users and describe how they leveraged the features of the testbed. I then discuss the potential that having a well-defined Computer Science instrument creates – in particular, how it can be used to advance reproducibility and sharing for the computer science community.

Biography: Kate Keahey is one of the pioneers of infrastructure cloud computing. She created the first open source Infrastructure-as-a-Service implementation, and currently works on convergence projects adapting cloud computing concepts to scientific applications and HPC environment. Kate leads the Chameleon project, which provides a distributed large-scale experimental testbed for Computer Science research. She is also a co-founder and co-Editor-in-Chief of SoftwareX, a journal devoted to publishing scientific software. Kate is a Scientist at Argonne National Laboratory and a Senior Fellow at the Computation Institute at the University of Chicago.

For more information contact the technical host Curt Canada, 505-665-7453, cvc@lanl.gov.

Hosted by the Information Science and Technology Institute (ISTI)