

MURPA Seminar Friday 29 July 2011 at 9am: Are We Using the Wrong Metrics for Large Scale HPC Systems?

Presenter: **Bill Kramer**, Deputy Project Director at the National Center for Supercomputing Applications; he is responsible for leading the NSF Blue Waters Project

Venue: Seminar Room 135, Building 26, Monash Clayton

Abstract:

Sustained performance is the amount of useful work a system can produce in a given amount of time on a regular basis. So, why isn't such an easy concept used more in Computer Science? It is because performance evaluation is exceedingly tricky to get right (despite what others may tell you.)

How much useful work a HPC system can achieve is difficult to assess in a simple, general manner because different communities have their own views of what useful work means and because a large number of system characteristics influence its usefulness. Yet, we, as users, intuitively, and sometimes explicitly, know when a HPC system is more useful than another. On the other hand, we also know when measures do not accurately portray a system's usefulness, yet some of them are commonly used over and over again.

This talk will review the important concepts of measuring sustained performance, particularly for very large systems, and discuss different approaches for doing the measurements and point out some of the issues that prevent our community from developing common, effective measures that are effective both in assessing technology and in describing its potential to others.

Bio:

William T.C. Kramer is deputy project director at the National Center for Supercomputing Applications; he is responsible for leading the Blue Waters project, a National Science Foundation-funded project, to deploy the first general purpose, open science, sustained-petaflop supercomputer as a powerful resource for the nation's researchers. Blue Waters is an 8 year project with an overall cost of over \$500M

Previously Kramer was the general manager of the National Energy Research Scientific Computing Center (NERSC), the flagship computing facility of the Department of Energy's Office of Science at Lawrence Berkeley National Laboratory (LBNL).

Prior to Berkeley Lab, Kramer worked at the NASA Ames Research Center, where he was responsible for all aspects of operations and customer service for NASA's Numerical Aerodynamic Simulator (NAS) supercomputer center and other large computational projects as well as starting a major Air Traffic Control Program.

Blue Waters will be the 20th supercomputer Kramer deploys and manages. Several were first of their kind, including the world's first UNIX supercomputer and the first production quality massively parallel system. In addition he deployed and managed large clusters of workstations, several extremely large data repositories, some of the world's most intense

networks, and other extreme scale systems. He has also been involved with the design, creation and commissioning of six “best of class” HPC facilities.