UNIVERSITY *of* WISCONSIN
LA CROSSE

Distinguished Lecture Series in Computer Science
Monday, Oct. 21

SCHEDULE OF EVENTS

10:30 a.m. **Registration** | Cleary Alumni & Friends Center | UW-L Campus

11 a.m. **SYMPOSIUM**
Avoiding, Hiding and Managing Communication

Future computing system designs will be constrained by power density and total system energy, and will require new programming models and implementation strategies. Data movement in the memory system and interconnect will dominate running time and energy costs, making communication cost reduction the primary optimization criteria for compilers and programmers. Communication cost can be divided into latency costs, which are per communication event, and bandwidth costs, which grow with total communication volume. The trends show growing gaps for both of these relative to computation, with the additional problem that communication congestion can conspire to worsen both in practice.

In this talk she will discuss prior work on open problems in optimizing communication, starting with PGAS languages. This involves reducing the communication cost, through overlap, and the frequency through caching and aggregation. Much of the compile-time work in this area was done in the Titanium language, where strong typing and data abstraction aid in program analysis, while UPC compilers tend to use more dynamic optimizations. There are still many open questions on the design of languages and compilers, especially as the communication hierarchies become deeper and more complex.

Bandwidth reduction often requires more substantial algorithmic transformations, although some techniques, such as loop tiling, are well known. These can be applied as hand-optimizations, through code generation strategies in autotuned libraries, or as fully automatic compiler transformations. Less obvious techniques for communication avoidance have arisen in the so-called "2.5D" parallel algorithms, which she will describe more generally as ".5D" algorithms. These ideas are applicable to many domains, from scientific computations to database operations. In addition to having provable optimality properties, these algorithms also perform well on large-scale parallel machines. Yelick will end by describing some recent work that lays the foundation for automating transformations to produce communication optimal code for arbitrary loop nests.

4:30 p.m. **Registration** | Cleary Alumni & Friends Center | UW-L Campus

5 p.m. **KEYNOTE**
More Data, More Science and … Moore's Law?

In the same way that the Internet has combined with web content and search engines to revolutionize every aspect of our lives, the scientific process is poised to undergo a radical transformation based on the ability to access, analyze, and merge large, complex data sets. Scientists will be able to combine their own data with that of other scientists, validating models, interpreting experiments, re-using and re-analyzing data, and making use of sophisticated mathematical analyses and simulations to drive the discovery of relationships across data sets. This "scientific web" will yield higher quality science, more insights per experiment, an increased democratization of science and a higher impact from major investments in scientific instruments.

What does this "big science data" view of the world have to do with HPC? The terms "high performance computing" and "computational science" have become nearly synonymous with modeling and simulation, and yet computing is as important to the analysis of experimental data as it is to the evaluation of theoretical models. Due to the exponential growth rates in detectors, sequencers and other observational technology, data sets across many science disciplines are outstripping the storage, computing, and algorithmic techniques available to individual scientists. Along with simulation, experimental analytics problems will drive the need for increased computing performance, although the types of computing systems and software configurations may be quite different.

Yelick will describe some of the opportunities and challenges in extreme data science and its relationship to high performance modeling and simulation. Including her research in the development of high performance, high productivity programming models. Her current research is largely focused on the problem of avoiding and minimizing the cost of communication, as data movement both to memory systems and between processors a major barrier to scalability and energy efficiency. In both simulation and analytics, programming models are the "sandwich topic," squeezed between application needs and hardware disruptions, yet often treated with some suspicion, if not outright disdain. But programming model research is, or at least should be, an exemplar of interdisciplinary science, requiring a deep understanding of applications, algorithms and computer architecture in order to map the former to the latter. She will use this thread to talk about her research interests, how she selected various research topics, and what she sees as current open problems in parallel programming model design and implementation.

Katherine Yelick

is the Associate Laboratory Director for Computing Sciences at Lawrence Berkeley National Laboratory and a Professor of Electrical Engineering and Computer Sciences at the University of California at Berkeley. She was Director of the National Energy Research Scientific Computing Center (NERSC) from 2008 to 2012 and she co-invented the UPC and Titanium languages as well as techniques for self-tuning sparse matrix kernels. She earned her Ph.D. in EECS from MIT and has been a professor at UC Berkeley since 1991 with a joint appointment at LBNL since 1996. She has received multiple teaching and research awards, including the recent ACM Athena Award. She also serves on many review and planning committees, and is a member of the California Council on Science and Technology and the National Academies Computer Science and Telecommunications Board.

For further information about the lecture contact:

Steve Senger, Ph.D.
Computer Science Department
UW-La Crosse
1725 State St.
La Crosse, WI 54601
608.785.6805
email: *compsci@uwlax.edu*

Co-sponsored by the

University of Wisconsin-La Crosse Foundation Inc.

Department of Computer Science

College of Science and Health