

# PMAM 2015 Workshop Program

## Day 1 (7th February)

08:45 – 09:00: Opening Remarks

Pavan Balaji, Minyi Guo, Zhiyi Huang

09:00 – 10:00: Keynote 1

John Shalf, Lawrence Berkeley National Laboratory

“Abstract Machine Models for Exascale Programming Systems”

### Abstract

For the past twenty-five years, a single model of parallel programming (largely bulk-synchronous MPI), has for the most part been sufficient to permit translation of this into reasonable parallel programs for more complex applications. In 2004, however, a confluence of events changed forever the architectural landscape that underpinned our current assumptions about what to optimize for when we design new algorithms and applications. We have been taught to prioritize and conserve things that were valuable 20 years ago, but the new technology trends have inverted the value of our former optimization targets. The time has come to examine the end result of our extrapolated design trends and use them as a guide to re-prioritize what resources to conserve in order to derive performance for future applications. This talk will describe the challenges of programming future computing systems. It will then provide some highlights from the search for durable programming abstractions more closely track emerging computer technology trends so that when we convert our codes over, they will last through the next decade.

### Biography of the speaker

John Shalf is CTO for the National Energy Research Supercomputing Center and also Department Head for Computer Science and Data Sciences at Lawrence Berkeley National Laboratory. Shalf is a co-author of over 60 publications in the field of parallel computing software and HPC technology, including three best papers and the widely cited report “The Landscape of Parallel Computing Research: A View from Berkeley” (with David Patterson and others), as well as “ExaScale Software Study: Software Challenges in Extreme Scale Systems,” which sets the Defense Advanced Research Project Agency’s (DARPA’s) information technology research investment strategy for the next decade. He was a member of the Berkeley Lab/NERSC team that won a 2002 R&D 100 Award for the RAGE robot. Before joining Berkeley Lab in 2000, he was a research programmer at the National Center for Supercomputing Applications at the University of Illinois and a visiting scientist at the Max-Planck-Institut für Gravitationsphysik/Albert Einstein Institute in

Potsdam, Germany, where he co-developed the Cactus code framework for computational astrophysics.

10:00- 10:30: Morning Break

10:30 – 12:00: Session 1: Performance Evaluation

Session Chair: Xiaoyi Lu

“Energy Efficiency and Performance Frontiers for Sparse Computations on GPU Supercomputers”, Hartwig Anzt, Stan Tomov and Jack Dongarra

“Energy-efficient Computing for HPC Workloads on Heterogeneous Manycore Chips”, Akhil Langer, Ehsan Toton, Udatta Palekar and Laxmikant Kale

“A Performance Study of Java Garbage Collectors on Multicore Architectures”, Maria Carpen Amarie, Patrick Marlier, Pascal Felber and Gael Thomas

12:00 – 14:00: Lunch break

14:00 – 15:30: Session 2: Programming Models, Languages and Tools (1)

Session Chair: Long Zheng

“Toward an Evolutionary Task Parallel Integrated MPI + X Programming Model”, Richard Barrett, Dylan Stark, Courtenay Vaughan, Ryan Grant, Stephen Olivier and Kevin Pedretti

“Design and Evaluation of a Novel DataFlow based BigData Solution”, Yao Wu, Long Zheng, Brian Heilig and Guang Gao

“Programming Support for Reconfigurable Custom Vector Architectures”, Mehmet Ali Arslan, Krzysztof Kuchcinski, Flavius Gruian and Yangxurui Liu

15:30 – 16:00: Afternoon Break

16:00 - 18:00 Session 3: Algorithms and Applications

Session Chair: Yonghong Yan

“Thread-Level Parallelization and Optimization of NWChem for the Intel MIC Architecture”, Hongzhang Shan, Samuel Williams, Wibe Jong and Leonid Oliker

“Parallelism vs. Speculation: Exploiting Speculative Genetic Algorithm on GPU”,  
Yanchao Lu, Long Zheng, Li Li and Minyi Guo

“GPU Technology Applied to Reverse Time Migration and Seismic Modeling via  
OpenACC”, Ahmad Qawasmeh, Maxime Hugues, Henri Calandra and Barbara  
Chapman

“Parallelizing a Discrete Event Simulation Application Using the Habanero-Java  
Multicore Library”, Wei-Cheng Xiao, Jisheng Zhao and Vivek Sarkar

## Day 2 (8th February)

09:00 – 10:00: Keynote 2

Barbara Chapman, University of Houston

“Portable Parallel Programming in an Age of Diversity”

### Abstract

Computer architectures are continuing to evolve in order to provide higher levels of application performance while keeping power utilization low. Core counts continue to grow. Some platforms rely on heterogeneity of cores to deliver performance; others do not. With the near-ubiquitous reliance on parallelism, it is more important than ever to provide application developers with programming models that facilitate the expression of parallelism in their code and offer portability. Yet given the scale and diversity of architectures, it is also harder than ever. Moreover, in order to exploit the features of recent and emerging computer systems, some applications will need significant restructuring. In this talk, I will discuss the status of directive-based multicore/manycore programming interfaces and the challenges that lie ahead for them. I will also ask where the porting tools are.

### Biography of the speaker

Dr. Chapman is a Professor of Computer Science at the University of Houston, TX, USA, where she also directs the Center for Advanced Computing and Data Systems. Chapman has performed research on parallel programming languages and the related implementation technology for over 20 years and has been involved in the OpenMP directive-based programming standard development since 2001. She also contributes to the OpenSHMEM and OpenACC programming standards efforts. Her research group has developed OpenUH, a state-of-the-art open source compiler that is used to explore language, compiler and runtime techniques, with a special focus on multithreaded programming. Dr. Chapman’s research also explores optimization of partitioned global address space programs, strategies for runtime code optimizations, compiler-tools interactions and high-level programming models for embedded systems.

10:00- 10:30: Morning Break

10:30 – 12:00: Session 4: Programming Models, Languages and Tools (2)

Session Chair: Abdelhalim Amer

“RaftLib: A C++ template library for high performance stream parallel processing”,  
Jonathan Beard, Peng Li and Roger Chamberlain

“A Java Util Concurrent Park Contention Tool”, Panagiotis Patros, Eric Aubanel, David Bremner and Michael Dawson

“Supporting Multiple Accelerators in High-Level Programming Models”, Yonghong Yan, Pei-Hung Lin, Chunhua Liao, Bronis R. De Supinski and Daniel J. Quinlan

12:00 – 14:00: Lunch break

14:00 – 15:30: Session 5: Runtime Support

Session Chair: CJ Newburn

“Effective Communication for a System of Cluster-on-a-Chip Processors”, Pablo Reble, Stefan Lankes, Fabian Fischer and Matthias S. Mueller

“Exploiting Communication Concurrency on High Performance Computing Systems”, Nicholas Chaimov, Khaled Ibrahim, Sam Williams and Costin Iancu

“CRA: A Dynamic Task Allocation Algorithm for Many-core Processor”, Chang Wang, Jiang Jiang, Yongxin Zhu, Xu Liu and Xing Han

15:30 – 16:00: Afternoon Break

16:00 - 17:30 Session 6: Programming Models, Languages and Tools (3)

Session Chair: TBD

“Patty: A Pattern-based Parallelization Tool for the Multicore Age”, Korbinian Molitorisz, Tobias Müller and Walter F. Tichy

“Deadlock-free Buffer Configuration for Stream Computing”, Peng Li, Jonathan Beard and Jeremy Buhler

“Debugging parallel programs using fork handlers”, Javier Alcázar

17:30 – 17:45: Closing Remarks

Pavan Balaji, Minyi Guo, Zhiyi Huang