PMAM 2020 Workshop Program

(22nd February, 2020)

09:00 – 09:10: Opening Remarks

Quan Chen, Zhiyi Huang, Min Si, Pavan Balaji

09:10 – 10:00: Keynote

Keynote Speaker: Prof. Zizhong (Jeffrey) Chen, UC Riverside, US

Title: The Interply among power, resilience, and performance

Abstract: While tremendous research have been done on improving applications' energy efficiency from the system point of view, few research exploits applications' algorithmic characteristics to maximize applications' energy efficiency. In this talk, I will discuss our recent work on algorithm-based techniques to improve applications' energy efficiency. We demonstrate that, in many cases, systems cannot make optimal decisions without application specific knowledge. We investigate the interplay among application power, resilience, and performance, and leverage the linear algebra application algorithmic characteristics to achieve higher performance with less power and energy in the widely used matrix operations.

BIO: Zizhong Chen is a professor at the University of California, Riverside. He specializes in high performance computing, numerical algorithms and software, and algorithm-based fault tolerance. Dr. Chen received a CAREER Award from the U.S. National Science Foundation and a Best Paper Award from the International Supercomputing Conference. He is a Senior Member of the IEEE and a Life Member of the ACM. Dr. Chen currently serves as Subject Area Editor for Elsevier Parallel Computing journal and Associate Editor for IEEE Transactions on Parallel and Distributed Systems.

10:00- 10:30: Morning Break

10:30 – 12:00: Session 1: Coherence and Synchronization

Session Chair: Shigang Li, ETH Zurich

- Yuxin Ren, Gabriel Parmer and Dejan Milojicic. "Bounded Incoherence: A Programming Model for Non-Cache-Coherent Shared Memory Architectures" .
- Daming Chen, Phil Gibbons and Todd Mowry. "TardisTM: Incremental Repair for Transactional Memory".

- Kenneth Lamar, Christina Peterson and Damian Dechev. "Lock-free Transactional Vector".
- Millad Ghane, Sunita Chandrasekaran and Margaret S. Cheung. "Towards A Portable Hierarchical View of Distributed Shared Memory Systems: Challenges and Solutions".

12:00 – 13:30: Lunch break

13:30 – 14:20: Keynote 2

Keynote Speaker: Prof. Aparna Chandramowlishwaran, UC Irvine, US

Title: Accelerating Turbulent Fluid Simulations

Abstract: Computational Fluid Dynamics (CFD) is the de-facto method for solving the discretized Navier-Stokes equations. It is computationally expensive especially for turbulent flows and spatially complex domains. We can accelerate these iterative simulations by reducing the time per iteration and/or by decreasing the number of iterations. In this talk, I will present ideas that fall into both of the above two camps for accelerating the Reynolds Averaged Navier Stokes equations. First, I will discuss how to accelerate CFD simulations using a combination of new parallel algorithms and principled performance engineering for reducing the time per iteration. We also ask whether CFD applications can be expressed in DSLs and whether such an implementation can deliver a sufficient combination of optimizations to compete with a hand-tuned code and what are its limitations. Then, I will introduce CFDNet, a physical simulation and deep learning coupled framework for accelerating the convergence of these simulations. It is designed to predict the primary physical properties of the fluid including velocity, pressure, and eddy viscosity using a single convolutional neural network at its core. CFDNet speeds the simulations up by a factor of 1.9 – 7.4x on both steady laminar and turbulent flows on a variety of geometries, without relaxing the convergence constraints of the physics solver showing the model's capacity for generalization and extrapolation.

BIO: Aparna Chandramowlishwaran is an assistant professor at the University of California, Irvine, in the Department of Electrical Engineering and Computer Science. Her research lab, HPC Forge, is interested in parallel algorithms, principled performance analysis, and domain-specific compilers. She is a recipient of the Google Faculty Research Award (2018), NSF CAREER award (2018), Intel Ph.D. fellowship (2012), ACM/IEEE George Michael Memorial HPC Fellowship (2010), Best Paper at IPDPS (2010), and co-recipient of the ACM Gordon Bell Prize (2010) among others. External to UCI, she serves as the associate editor of the ACM Transactions on Parallel Computing. She received her Ph.D. in Computational

Science and Engineering from Georgia Tech in 2013 and was a research scientist at MIT.

14:30 – 15:00: Afternoon break

15:00-16:30 Session2: Embedded Systems and Accelerators

Session Chair: Zizhong Chen, University of California, Riverside

- Xiaoxin Tang. "ELSE: an Efficient Link-time Static Instrumentation Tool for Embedded System".
- Patrick Finnerty, Tomio Kamada and Chikara Ohta. "Self-adjusting Task Granularity for Global Load Balancer Library on Clusters of Many-core Processors".
- Akif Rehman, Masab Ahmad and Omer Khan. "Exploring Accelerator and Parallel Graph Algorithmic Choices for Temporal Graphs".
- Sebastian Litzinger and Jörg Keller. "Generating Energy-Efficient Code for Parallel Applications Specified by Streaming Task Graphs with Dynamic Elements".

16:30 – 16:35: Closing Remarks

Quan Chen, Zhiyi Huang, Min Si, Pavan Balaji