

Workshop IV: New Architectures and Algorithms

Monday November 26, 2018

- 8:00–8:55 *Check-In/Light Breakfast (Hosted by IPAM)*
- 8:55–9:00 *Welcome and Opening Remarks*
- 9:00–9:50 **Alan Lee** (AMD)
Big Compute: Architectures and Algorithms
- 10:00–10:15 *Break*
- 10:15–11:05 **Matthias Troyer** (Microsoft Research)
A quantum future of computation
- 11:15–11:30 *Break*
- 11:30–12:20 **Gavin Jones** (IBM Research - Almaden)
Quantum Computing and IBM Q: An Introduction
- 12:30–2:30 *Lunch (on your own)*
- 2:30–3:20 **Jack Raymond** (D-Wave Systems)
Quantum simulation of condensed matter in programmable qubit lattices
- 3:30–4:00 *Break*
- 4:00–4:50 **Maya Gokhale** (Lawrence Livermore National Laboratory)
System software and architecture assist for application-directed data views
- 5:00–6:30 *Poster Session & Reception (Hosted by IPAM)*

Tuesday November 27, 2018

- 8:00–9:00 *Check-In/Breakfast (Hosted by IPAM)*
- 9:00–9:50 **Hatem Ltaief** (King Abdullah Univ. of Science and Technology (KAUST))
Leveraging Computational Astronomy Applications on Massively Parallel Hardware Systems
- 10:00–10:15 *Break*
- 10:15–11:05 **Ameet Talwalkar** (Carnegie Mellon University)
Massively Parallel Hyperparameter Tuning
- 11:15–11:30 *Break*

(Tuesday schedule continued on next page)



(Tuesday schedule continued from previous page)

- 11:30–12:20 **James Demmel** (University of California, Berkeley (UC Berkeley))
Architectural Implications of Communication-Avoiding Algorithms
- 12:30–2:30 *Lunch (on your own)*
- 2:30–3:20 **Amir Gholaminejad** (University of California, Berkeley (UC Berkeley))
Scalable Neural Network Training Using Second Order Stochastic Training
- 3:30–4:00 *Break*
- 4:00–4:50 **Yingyu Liang** (University of Wisconsin-Madison)
Learning Overparameterized Neural Networks on Structured Data

Wednesday November 28, 2018

- 8:00–9:00 *Check-In/Breakfast (Hosted by IPAM)*
- 9:00–9:50 **William Gropp** (University of Illinois at Urbana-Champaign)
Challenges in Programming Extreme Scale Systems
- 10:00–10:15 *Break*
- 10:15–11:05 **Aydin Buluc** (Lawrence Berkeley National Laboratory)
Communication-Avoiding Sparse Matrix Algorithms for Large Graph and Machine Learning Problems
- 11:15–11:30 *Break*
- 11:30–12:20 **Tobias Weinzierl** (Durham University)
Stop talking to me - some communication-avoiding programming patterns
- 12:30–2:30 *Lunch (on your own)*
- 2:30–3:20 **Michael Schulte** (AMD)
Architectures and Technologies for Extreme-Scale Computing
- 3:30–4:00 *Break*
- 4:00–4:50 **James Sexton** (IBM Research)
Architecture Challenges and Opportunities for Exascale Systems

Thursday November 29, 2018

- 8:00–9:00 *Check-In/Breakfast (Hosted by IPAM)*
- 9:00–9:50 **John Shalf** (Lawrence Berkeley Laboratory)
Computing Beyond Moore's Law
- 10:00–10:15 *Break*

(Thursday schedule continued on next page)

(Thursday schedule continued from previous page)

- 10:15–11:05 **Shang-Hua Teng** (University of Southern California (USC))
Scalable Algorithms in the Age of Big Data and Network Sciences
- 11:15–11:30 *Break*
- 11:30–12:20 **Stefano Ambrogio** (IBM Research - Almaden)
Novel Technologies for Artificial Intelligence: prospects and challenges
- 12:30–2:30 *Lunch (on your own)*
- 2:30–3:20 **Nina Balcan** (Carnegie Mellon University)
Data Driven Algorithm Design
- 3:30–4:00 *Break*
- 4:00–4:50 **Yaron Singer** (Harvard University)
Maximizing Submodular Functions Exponentially Faster

Friday November 30, 2018

- 8:00–9:00 *Check-In/Breakfast (Hosted by IPAM)*
- 9:00–9:50 **Chris Anderson** (University of California, Los Angeles (UCLA))
Generalized Aggregate Descriptor Approximations
- 10:00–10:15 *Break*
- 10:15–11:05 **Petter Bjoerstad** (University of Bergen)
Domain Decomposition and Computer Architecture - a stable relationship?
- 11:15–11:30 *Break*
- 11:30–12:20 **Edmond Chow** (Georgia Institute of Technology)
Asynchronous Iterative Methods
- 12:30–2:30 *Lunch (on your own)*
- 2:30–3:20 **Andreas Krause** (ETH Zurich)
Deep Learning Architectures with Discrete Decisions via Submodularity
- 3:30–4:00 *Break*
- 4:00–4:50 **Jeffrey Hittinger** (Lawrence Livermore National Laboratory)
Rethinking Finite Precision

