

8:30-10:00 Session I: Parallel Programming Models and Languages	Session Chair: Prof. Xiaofei Liao Huazhong University of Science and Technology
Kokkos Array Performance-Portable Manycore Programming Model <i>H. Carter Edwards and Daniel Sunderland</i>	
An Hybrid Model for Very High Level Threads <i>Jafar Al-Gharaibeh, Clinton Jeffery and Kostas Oikonomou</i>	
Function Flow: Making Synchronization Easier in Task Parallelism <i>Xuepeng Fan, Hai Jin, Liang Zhu, Xiaofei Liao and Chengcheng Ye</i>	
New Strategy for Coarse Grid Solvers in Parallel Multigrid Methods using OpenMP/MPI Hybrid Programming Models <i>Kengo Nakajima</i>	
Networks Beat Pipelines: The Design of FG 2.0 <i>Peter Johnson and Thomas Cormen</i>	
10:00-10:30 Break	
10:30-12:00 Session II: Cache, Memory and IO	Session Chair: Prof. Dongrui Fan Institute of Computing Technology, Chinese Academy of Sciences, China
AGC: Adaptive Global Clock in STMs <i>Ehsan Atoofian and Amir Ghanbari Bavarsad</i>	
Efficient Execution of Time-step Computations with Pipelined Parallelism and Inter-thread Data Locality Optimizations <i>Apan Qasem</i>	
Revisiting Shared Virtual Memory Systems for Non-Coherent Memory-Coupled Cores <i>Stefan Lankes, Pablo Reble, Carsten Clauss and Oliver Sinnen</i>	
Efficient Memory Management of a Hierarchical and a Hybrid Main Memory for MN-MATE Platform <i>Kyu Ho Park, Sung Kyu Park, Hyunchul Seok, Woomin Hwang, Dong-Jae Shin, Jong Hun Choi and Ki-Woong Park</i>	
Massively Parallel Breadth First Search using a Tree-Structured Memory Model <i>Tom St. John, Jack Dennis and Guang Gao</i>	
12:00-1:30 Lunch	
13:30-15:00 Session III: Parallel Algorithm Optimization	Session Chair: Prof. Uzi Vishkin University of Maryland
PMA: Pixel-based Multi-Anchor Algorithm for Image Recognition on Multi-core Systems <i>Xiaoxin Tang, Long Zheng, Jun Ma, Yao Shen, Li Li and Minyi Guo.</i>	
Better Speedups Using Simpler Parallel Programming for Graph Connectivity and Biconnectivity <i>James Edwards and Uzi Vishkin</i>	
Shared Work List: Hacking Amorphous Data Parallelism in UPC <i>Shixiong Xu and Li Chen.</i>	
A Case for Secure and Scalable Hypervisor using Safe Language <i>Haibo Chen and Binyu Zang</i>	
Semi-sparse Algorithm Based on Multi-layer Optimization for Recommendation System	

Hu Guan, Huakang Li and Minyi Guo

15:00-15:30 Break

15:30-17:00

Session IV: GPU Acceleration

Session Chair: Prof. Ehsan Atoofian

Lakehead University, Canada

GHOST: GPGPU-Offloaded High Performance Storage I/O Deduplication for Primary Storage System

Chulmin Kim, Ki-Woong Park and Kyu Ho Park

Exploring Parallelism in Volume Ray Casting: Understanding the Programming Issues of Multithreaded Accelerators

Guilherme Cox, Cleomar Silva, Leandro Cupertino, Cristiana Bentes and Ricardo Farias

Techniques for the Parallelization of Unstructured Grid Applications on Multi-GPU Systems

Lizandro Solano-Quinde, Brett Bode and Arun Somani

Chestnut: A GPU Programming Language for Non-Experts

Andrew Stromme, Ryan Carlson and Tia Newhall