

48th International Conference on Parallel Processing



August 5(Mon)-8(Thu), 2019, Kyoto Research Park, Kyoto, Japan

Conference host




University of Tsukuba



University of Tsukuba
Center for
Computational Sciences



Orchestrating a brighter world



This program is supported by a subsidy from Kyoto City and the Kyoto Convention & Visitors Bureau.

08:30-09:00 **Registration Open**

Workshop day 09:00-17:30

09:00-12:30 **SRMPDS:**
and
14:00-17:30 **International Workshop on Scheduling and Resource Management
for Parallel and Distributed Systems**
Room-A

09:00-12:30 **AWASN:**
Room-B **Applications of Wireless Ad hoc and Sensor Networks**

14:00-17:30 **EMS:**
Room-B **International Workshop on Embedded Multicore Systems**

09:00-12:30 **PDML: Parallel and Distributed Machine Learning**
Room-C

14:00-17:30 **P2S2:**
Room-C **International Workshop on Parallel Programming Models and
Systems Software for High-End Computing**

09:00-12:30 **EE HPC SOP:**
and
14:00-17:30 **Energy Efficient HPC State of the Practice**
Room-G

12:30-14:00 **Lunch**
PATIO Please take lunch at the restaurant PATIO using a coupon.

18:00-19:30 **Welcome Reception**
PATIO

Tuesday, August 6th

08:00-08:30 **Registration Open**

B1F Foyer

08:30-09:00 **Welcome and Introduction**

Buzz Hall

09:00-10:00 **Keynote-1**

Chair: Kengo Nakajima

Buzz Hall

Depei Qian: Major issues in exascale computing in China

10:00-10:40 **Best Paper for ICPP 2019**

Chair: Kengo Nakajima, Martin Schulz

Buzz Hall

(10:00-10:10 Award Ceremony, 10:10-10:40 Presentation)

A Parallel Graph Algorithm for Detecting Mesh Singularities in Distributed Memory Ice Sheet Simulations

Ian A. Bogle (Rensselaer Polytechnic Institute), Karen Devine (Sandia National Laboratories), Mauro Perego (Sandia National Laboratories), Siva Rajamanickam (Sandia National Laboratories), George M. Slota (Rensselaer Polytechnic Institute)

10:40-11:10 **Coffee Break**

Banquet Hall

11:10-12:50 **Session T1A**

Chair: Yuetsu Kodama

Buzz Hall

Memory Architectures

MAC: Memory Access Coalescer for 3D-Stacked Memory

Xi Wang (Texas Tech University), Antonino Tumeo (Pacific Northwest National Laboratory), John D. Leidel (Tactical Computing Labs, Texas Tech University), Jie Li (Texas Tech University), Yong Chen (Texas Tech University)

Machine Learning for Fine-Grained Hardware Prefetcher Control

Jason Hiebel (Michigan Technological University), Laura E. Brown (Michigan Technological University), Zhenlin Wang (Michigan Technological University)

AVR: Reducing Memory Traffic with Approximate Value Reconstruction

Albin Eldstal-Damlin (Chalmers University of Technology), Pedro Moura Trancoso (Chalmers University of Technology), Ioannis Sourdis (Chalmers University of Technology)

Near-Data Processing-Enabled and Time-Aware Compaction Optimization for LSM-tree-based Key-Value Stores

Hui Sun (Anhui University), Wei Liu (Anhui University), Jianzhong Huang (Huazhong University of Science and Technology), Song Fu (University of North Texas), Zhi Qiao (University of North Texas), Weisong Shi (Wayne State University)

11:10-12:50 **Session T1B**

Chair: Daisuke Takahashi

Room-1

Workflow and Data Analysis Systems

Incorporating Probabilistic Optimizations for Resource Provisioning of Data Processing Workflows

Amelie Chi Zhou (Shenzhen University), Yao Xiao (Shenzhen University), Bingsheng He (National University of Singapore), Shadi Ibrahim (Inria), Reynold Cheng (University of Hong Kong)

ECOST: Energy-Efficient Co-Locating and Self-Tuning MapReduce Applications

maria malik (George Mason University, Intel), Hassan Ghasemzadeh (Washington State University), Tinoosh Mohsenin (University of Maryland, Baltimore County), Rosario Cammarota (Intel), Liang Zhao (George Mason University), Avesta Sasan (George Mason University), Houman Homayoun (George Mason University), Setareh Rafatirad (George Mason University)

Tuesday, August 6th

Stage Delay Scheduling: Speeding up DAG-style Data Analytics Jobs with Resource Interleaving

Wujie Shao (East China Normal University), Fei Xu (East China Normal University), Li Chen (University of Louisiana at Lafayette), Haoyue Zheng (East China Normal University), Fangming Liu (Huazhong University of Science and Technology)

Solving All-Pairs Shortest-Paths Problem in Large Graphs Using Apache Spark

Frank Schoeneman (University at Buffalo), Jaroslaw Zola (University at Buffalo)

11:10-12:50

Room-2

Session T1C

Chair: Ryusuke EGAWA

Data Centers

Unleashing the Scalability Potential of Power-Constrained Data Center in the Microservice Era

Xiaofeng Hou (Shanghai Jiao Tong University), Jiacheng Liu (Shanghai Jiao Tong University), Chao Li (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University)

Dynamic Load Balancing in Hybrid Switching Data Center Networks with Converters

Jiaqi Zheng (Nanjing University), Qiming Zheng (Shanghai Jiao Tong University), Xiaofeng Gao (Shanghai Jiao Tong University), Guihai Chen (Nanjing University)

Nested Virtualization Without the Nest

Mathieu Bacou (IRIT, Universite de Toulouse, CNRS, Toulouse Atos Integration, Toulouse), Gregoire Todeschi (IRIT, Universite de Toulouse, CNRS, Toulouse), Alain Tchana (I3S, Universite Sophia Antipolis, CNRS, Nice), Daniel Hagimont (IRIT, Universite de Toulouse, CNRS, Toulouse)

When Power Oversubscription Meets Traffic Flood Attack: Re-Thinking Data Center Peak Load Management

Xiaofeng Hou (Shanghai Jiao Tong University), Mingyu Lliang (Shanghai Jiao Tong University), Chao Li (Shanghai Jiao Tong University), Wenli Zheng (Shanghai Jiao Tong University), Quan Chen (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University)

12:50-14:10

Lunch

Please take a lunch box at the registration desk and have lunch at Banquet Hall or the conference rooms.

14:10-15:50

Buzz Hall

Session T2A

Chair: Tasuku Hiraishi

Memory Optimizations

LFOC: A Lightweight Fairness-Oriented Cache Clustering Policy for Commodity Multicores

Adrian Garcia-Garcia (UCM), Juan Carlos Saez (UCM), Fernando Castro (UCM), Manuel Prieto-Matias (UCM)

DICER: Diligent Cache Partitioning for Efficient Workload Consolidation

Konstantinos Nikas (National Technical University of Athens), Nikela Papadopoulou (National Technical University of Athens), Dimitra Giantsidi (National Technical University of Athens), Vasileios Karakostas (National Technical University of Athens), Georgios Goumas (National Technical University of Athens), Nectarios Koziris (National Technical University of Athens)

EMBA: Efficient Memory Bandwidth Allocation to Improve Performance on Intel Commodity Processor

Yaocheng Xiang (Peking University), Chencheng Ye (Huazhong University of Science and Technology), Xiaolin Wang (Peking University), Yingwei Luo (Peking University), Zhenlin Wang (Michigan Technological University)

CP-pf: a prefetch aware LLC partitioning approach

Jun Xiao (University of Amsterdam), Andy Pimentel (University of Amsterdam), Xu Liu (College of William and Mary)

Tuesday, August 6th

14:10-15:50

Room-1

Session T2B

Chair: Guillaume Pallez

Parallel Systems Algorithms

TLB: Traffic-aware Load Balancing with Adaptive Granularity in Data Center Networks *Best Paper Candidate

Jinbin Hu (Central South University), Jiawei Huang (Central South University), Wenjun Lv (Central South University), Weihe Li (Central South University), Jianxin Wang (Central South University), Tian He (University of Minnesota)

HyperPRAW: Architecture-Aware Hypergraph Restreaming Partition to Improve Performance of Parallel Applications Running on High Performance Computing Systems

Carlos Fernandez Musoles (The University of Sheffield), Daniel Coca (The University of Sheffield), Paul Richmond (The University of Sheffield)

A Network-aware and Partition-based Resource Management Scheme for Data Stream Processing

Yidan Wang (RMIT University), Zahir Tari (RMIT University), Xiaoran Huang (The University of Melbourne), Y. Zomaya Albert (The University of Sydney)

PhSIH: A Lightweight Parallelization of Event Matching in Content-based Pub/Sub Systems

Zhengyu Liao (Shanghai Jiao Tong University), Shiyu Qian (Shanghai Jiao Tong University), Jian Cao (Shanghai Jiao Tong University), Yanhua Cao (Shanghai Jiao Tong University), Guangtao Xue (Shanghai Jiao Tong University), Jiadi Yu (Shanghai Jiao Tong University), Yanmin Zhu (Shanghai Jiao Tong University), Minglu Li (Shanghai Jiao Tong University)

14:10-15:50

Room-2

Session T2C

Chair: Eishi Arima

NVRAM and SSD

A Read-leveling Data Distribution Scheme for Promoting Read Performance in SSDs with Deduplication

Mengting Lu (Huazhong University of Science and Technology), Fang Wang (Wuhan National Laboratory for Optoelectronics, School of Computer Science and Technology, Huazhong University of Science and Technology), Dan Feng (Wuhan National Laboratory for Optoelectronics, School of Computer Science and Technology, Huazhong University of Science and Technology), Yuchong Hu (Wuhan National Laboratory for Optoelectronics, School of Computer Science and Technology)

RFPL: A Recovery Friendly Parity Logging Scheme for Reducing Small Write Penalty of SSD RAID

Gaoxiang Xu (Wuhan National Laboratory for Optoelectronics, School of Computer Science, Huazhong University of Science and Technology), dan Feng (Huazhong University of Science and Technology), Zhipeng Tan (Huazhong University of Science and Technology), Xinyan Zhang (Huazhong University of Science and Technology), Jie Xu (Huazhong University of Science and Technology), Yifeng Zhu (University of Maine), Xi Shu (Huazhong University of Science and Technology)

TEA: A Traffic-efficient Erasure-coded Archival Scheme for In-memory Stores

Bin Xu (Huazhong University of Science and Technology), Jianzhong Huang (Huazhong University of Science and Technology), Qiang Cao (Huazhong University of Science and Technology), Xiao Qin (Auburn University)

CostPI: Cost-Effective Performance Isolation for Shared NVMe SSDs

Jiahao Liu (Huazhong University of Science and Technology), Fang Wang (Huazhong University of Science and Technology), Dan Feng (Huazhong University of Science and Technology)

15:50-16:20

Banquet Hall

2F Foyer

Coffee Break

Tuesday, August 6th

16:20-18:00

Buzz Hall

Session T3A

Chair: Toshio Endo

Parallel Architectures

SaC: Exploiting Execution-Time Slack to Save Energy in Heterogeneous Multicore Systems

Muhammad Waqar Azhar (Chalmers University of Technology), Miquel Pericas (Chalmers University of Technology), Per Stenstrom (Chalmers University of Technology)

Express Link Placement for NoC-Based Many-Core Platforms

Yunfan Li (Oregon State University), Di Zhu (University of Southern California), Lizhong Chen (Oregon State University)

Modeling the Performance of Atomic Primitives on Modern Architectures

Fazeleh Hoseini (Chalmers University of Technology), Aras Atalar (Chalmers University of Technology), Philippas Tsigas (Chalmers University of Technology)

The Case for Water-Immersion Computer Boards

Michihiro Koibuchi (National Institute of Informatics, The Graduate University of Advanced Studies), Ikki Fujiwara (National Institute of Informatics), Naoya Niwa (Keio University), Tomohiro Totoki (Keio University), Shoichi Hirasawa (National Institute of Informatics)

16:20-18:00

Room-1

Session T3B

Chair: Masahiro Nakao

Scheduling

JobPacker: Job Scheduling for Data-Parallel Frameworks with Hybrid Electrical/Optical Datacenter Networks

Zhuozhao Li (University of Chicago), Haiying Shen (University of Virginia)

Holistic Slowdown Driven Scheduling and Resource Management for Malleable Jobs

Marco D'Amico (Barcelona Supercomputing Center (BSC)), Ana Jokanovic (Barcelona Supercomputing Center (BSC)), Julita Corbalan (Universitat Politecnica de Catalunya)

Speculative Scheduling for Stochastic HPC Applications

Ana Gainaru (Vanderbilt), Guillaume Pallez (Inria), Hongyang Sun (Vanderbilt), Padma Raghavan (Vanderbilt)

Cooperative Job Scheduling and Data Allocation for Busy Data-Intensive Parallel Computing Clusters

Guoxin Liu (Clemson University), Haiying Shen (University of Virginia), Haoyu Wang (University of Virginia)

16:20-18:00

Room-2

Session T3C

Chair: Hiroya Matsuba

I/O Systems

N-Code: An Optimal RAID-6 MDS Array Code for Load Balancing and High I/O Performance

Ping Xie (The Computer College of Qinghai Normal University), Zhu Yuan (The Computer College of Qinghai Normal University), JianZhong Huang (Huazhong University of Science & Technology), Xiao Qin (Auburn University)

BPP: A Realtime Block Access Pattern Mining Scheme for I/O Prediction

Chunjie Zhu (Wuhan National Lab for Optoelectronics, Huazhong University of Science and Technology, Wuhan), Fang Wang (Wuhan National Lab for Optoelectronics, Huazhong University of Science and Technology, Wuhan), Binbing Hou (Department of Computer Science)

DeepHash: An End-to-End Learning Approach for Metadata Management in Distributed File Systems

Yuaning Gao (Shanghai Jiao Tong University), Xiaofeng Gao (Shanghai Jiao Tong University), Guihai Chen (Shanghai Jiao Tong University)

AdaM: An Adaptive Fine-Grained Scheme for Distributed Metadata Management

Shiyi Cao (Shanghai Jiao Tong University), Yuanning Gao (Shanghai Jiao Tong University), Xiaofeng Gao (Shanghai Jiao Tong University), Guihai Chen (Shanghai Jiao Tong University)

18:10-20:00

Banquet Hall

Poster Reception

Chair: Ryusuke Egawa

Poster 1

Performance improvement of MODYLAS using Remote Direct Memory Access on the K computer

Masahiro Nakao (RIKEN R-CCS), Hitoshi Murai (RIKEN R-CCS), Mitsuhsa Sato (RIKEN R-CCS) Yoshimichi Andoh (Nagoya University), Susumu Okazaki (Nagoya University)

Poster 2

Performance evaluation of kinetic code on scalar processors

Takayuki Umeda (Nagoya University)

Poster 3

Can Local Binary Convolutions Make Neural Networks Models Smaller?

Haoyu Zhang (Tokyo Institute of Technology), Mohamed Wahib (AIST), Peng Chen (Tokyo Institute of Technology, AIST), Satoshi Matsuoka (Tokyo Institute of Technology, RIKEN R-CCS)

Poster 4

Data centers are a software development challenge

Luca Bortot (ENI), Walter Nardelli(ENI), Peter Seto(EE HPC WG)

Poster 5

Efficient VNF Migration and Recovery in Edge Computing Environments

Yu Chen (Nanjing University), Sheng Zhang (Nanjing University), Yanchao Zhao (Nanjing University of Aeronautics and Astronautics), Zhuzhong Qian (Nanjing University), Yu Liang (Nanjing University), Jidong Ge (Nanjing University), Sanglu Lu (Nanjing University)

Poster 6

A Case for Co-scheduling for Hybrid Memory Based Systems

Eishi Arima (The University of Tokyo), Carsten Trinitis (Technical University of Munich)

Poster 7

Improving Strong Scalability Limits of Finite Element Based Solvers

Niclas Jansson (KTH Royal Institute of Technology)

Poster 8

AutoEncoder based Active Signal Map Crowdsourcing

Chengyong Liu (Nanjing University of Aeronautics and Astronautics), Yanchao Zhao (Nanjing University of Aeronautics and Astronautics), Kun Zhu (Nanjing University of Aeronautics and Astronautics), Sheng Zhang (Nanjing University), Jie Wu (Temple University)

Poster 9

A Relaxed Balanced Non-Blocking Binary Search Tree

Manish Singh (Wellington Institute of Technology), Lindsay Groves (Victoria University of Wellington), Alex Potanin (Victoria University of Wellington)

Poster 10

Implementation of Partitioning of Hierarchical Matrices using Task Parallel Languages

Zhengyang Bai (Kyoto University), Tasuku Hiraishi (Kyoto University), Akihiro Ida (The University of Tokyo), Masahiro Yasugi (Kyushu Institute of Technology)

Poster 11

Toward Training a Large 3D Cosmological CNN with Hybrid Parallelization

Yosuke Oyama(Tokyo Institute of Technology, Lawrence Livermore National Laboratory), Naoya Maruyama (Lawrence Livermore National Laboratory), Nikolii Dryden (University of Illinois at Urbana-Champaign, Lawrence Livermore National Laboratory), Peter Harrington (Lawrence Berkeley National Laboratory), Jan Balewski (Lawrence Berkeley National Laboratory), Satoshi Matsuoka(RIKEN R-CCS, Tokyo Institute of Technology), Marc Snir (University of Illinois at Urbana-Champaign), Peter Nugent (Lawrence Berkeley National Laboratory), and Brian Van Essen (Lawrence Livermore National Laboratory)

Poster 12

Maintaining Connectivity in Parallel Graph Partitioning

Christopher I. Jones (Rensselaer Polytechnic Institute Troy), Ian Bogle (Rensselaer Polytechnic Institute Troy), George M. Slota (Rensselaer Polytechnic Institute Troy)

Poster 13

Simple DSL for Power-Performance Modeling with Segmented Linear Models

Yuan He (Shenyang University of Technology), Yasutaka Wada (Meisei University), Guanqin Pan (Shenyang University of Technology), Masaaki Kondo (The University of Tokyo, RIKEN)

Poster 14

Enabling Data Processing under Erasure Coding in the Fog

Jad Darrous (Univ. Lyon, Inria, CNRS, ENS de Lyon), Shadi Ibrahim (Inria, IMT Atlantique, LS2N)

Poster 15

Single Precision Calculation of Iterative Refinement of Pairs of a Real Symmetric-Definite Generalized Eigenproblem by Using a Filter Composed of a Single Resolvent

Hiroshi Murakami (Tokyo Metropolitan University)

Poster 16

Performance Improvement of Deep Learning Training on Large-scale Manycore Cluster

Toshihiro Hanawa (The University of Tokyo), Kohei Tamura (The University of Tokyo)

Poster 17

Understanding the Overheads of Launching CUDA Kernels

Lingqi Zhang (Tokyo Institute of Technology), Mohamed Wahib (AIST), Satoshi Matsuoka (RIKEN R-CCS, Tokyo Institute of Technology)

Poster 18

An Operations Monitoring and Notification Infrastructure (OMNI) for Exascale Data Center Operations

Melissa Romanus (Rutgers University, Lawrence Berkeley National Laboratory), Elizabeth Bautista (Lawrence Berkeley National Laboratory), Thomas Davis (Lawrence Berkeley National Laboratory), Cary Whitney (Lawrence Berkeley National Laboratory)

Wednesday, August 7th

09:00-09:30 **Registration Open**

B1F Foyer

09:30-10:30 **Keynote-2**

Chair: Taisuke Boku

Buzz Hall

Satoshi Sekiguchi: Transforming SMEs and Manufacturing industry business into AI ready empowered by the ABCI

10:30-11:00 **Coffee Break**

Banquet Hall

11:00-12:40 **Session T4A**

Chair: Atsushi Hori

Buzz Hall

On Node Optimization

Optimized Execution of Parallel Loops via User-Defined Scheduling Policies

Seonmyeong Bak (Georgia Institute of Technology), Yanfei Guo (Argonne National Laboratory), Pavan Balaji (Argonne National Laboratory), Vivek Sarkar (Georgia Institute of Technology)

Data and Thread Placement in NUMA Architectures: A Statistical Learning Approach

Nicolas Denoyelle (Argonne National Laboratory), Brice Goglin (Inria), Emmanuel Jeannot (Inria), Thomas Ropars (LIG)

HPAS: An HPC Performance Anomaly Suite for Reproducing Performance Variations

Emre Ates (Boston University), Yijia Zhang (Boston University), Burak Aksar (Boston University), Jim Brandt (Sandia National Laboratories), Vitus J. Leung (Sandia National Laboratories), Manuel Egele (Boston University), Ayse K. Coskun (Boston University)

Reducing Kernel Surface Areas for Isolation and Scalability

Daniel Zahka (Washington University in St. Louis), Brian Kocoloski (Washington University in St. Louis), Kate Keahey (Argonne National Laboratory)

11:00-12:40 **Session T4B**

Chair: Osni Marques

Room-1

Parallel Algorithms 1

Accelerating All-Edge Common Neighbor Counting on Three Processors

Yulin Che (Hong Kong University of Science and Technology), Zhuohang Lai (Hong Kong University of Science and Technology), Shixuan Sun (Hong Kong University of Science and Technology), Qiong Luo (Hong Kong University of Science and Technology), Yue Wang (Hong Kong University of Science and Technology)

Tessellating Star Stencils

Liang Yuan (ICT, CAS), Shan Huang (ICT, CAS), Yunquan Zhang (ICT, CAS), Hang Cao (ICT, CAS)

Parallel Algorithms for Evaluating Matrix Polynomials

Sivan Toledo (Tel-Aviv University), Amit Waisel (Tel-Aviv University)

A 2D Parallel Triangle Counting Algorithm for Distributed-Memory Architectures

Ancy Sarah Tom (University of Minnesota), George Karypis (University of Minnesota)

11:00-12:40 **Session T4C**

Chair: Satoshi Imamura

Room-2

Communication Architectures

Network Congestion-aware Online Service Function Chain Placement and Load Balancing

Xiaojun Shang (Stony Brook University), Zhenhua Liu (Stony Brook University), Yuanyuan Yang (Stony Brook University)

Wednesday, August 7th

Breaking Band: A Breakdown of High-performance Communication

Rohit Zambre (University of California, Irvine), Megan Grodowitz (Arm Research), Aparna Chandramowliswaran (University of California, Irvine), Pavel Shamis (Arm Research)

Cartesian Collective Communication

Jesper Larsson Traff (TU Wien (Vienna University of Technology)), Sascha Hunold (TU Wien (Vienna University of Technology))

Design Exploration of Multi-tier Interconnection Networks for Exascale Systems

Javier Navaridas (University of Manchester), Josh Lant (University of Manchester), Jose A. Pascual (University of the Basque country), Mikel Lujan (University of Manchester), John Goodacre (University of Manchester)

12:40-14:00 **Lunch**

Please take a lunch box at the registration desk and have lunch at Banquet Hall or the conference rooms.

14:00-15:40 **Session T5A**

Chair: Miquel Pericas

Buzz Hall

System Software for GPUs

Adaptive Routing Reconfigurations to Minimize Flow Cost in SDN-Based Data Center Networks

Akbar Majidi (shanghai jiaotong university), Xiaofeng Gao (Shanghai Jiao Tong University Shanghai), Shunjia Zhu (Brown University), Nazila Jahanbakhsh (Shanghai Jiao Tong University Shanghai), Guihai Chen (Shanghai Jiao Tong University Shanghai)

A Specialized Concurrent Queue for Scheduling Irregular Workloads on GPUs

David Troendle (University of Mississippi, Department of Information and Computer Science), Tuan Ta (Cornell University, School of Electrical and Computer Engineering), Byunghyun Jang (University of Mississippi, Department of Information and Computer Science)

Predictable GPUs Frequency Scaling for Energy and Performance

Kaijie Fan (Technical University Berlin), Biagio Cosenza (Technical University Berlin), Ben Juurlink (Technical University Berlin)

Compiler-Assisted GPU Thread Throttling for Reduced Cache Contention

Hyunjun Kim (Sungkyunkwan University), Sungin Hong (Sungkyunkwan University), Hyeonsu Lee (Sungkyunkwan University), Euseong Seo (Sungkyunkwan University), Hwansoo Han (Sungkyunkwan University)

14:00-15:40 **Session T5B**

Chair: Jakub Kurzak

Room-1

Parallel Algorithms 2

SAFE: Service Availability via Failure Elimination Through VNF Scaling

Rui Xia (Nanjing University), Haipeng Dai (Nanjing University), Jiaqi Zheng (Nanjing University), Rong Gu (Nanjing University), Xiaoyu Wang (Nanjing University), Guihai Chen (Nanjing University)

On Max-min Fair Resource Allocation for Distributed Job Execution

Yitong Guan (Nanyang Technological University), Chuanyou Li (Southeast University), Xueyan Tang (Nanyang Technological University)

Improving Short Job Latency Performance in Hybrid Job Schedulers with Dice

Wei Zhou (University of Virginia), K. Preston White (University of Virginia), Hongfeng Yu (University of NebraskaLincoln)

A Practical, Scalable, Relaxed Priority Queue

Tingzhe Zhou (Lehigh University), Maged Michael (Facebook), Michael Spear (Lehigh University)

Wednesday, August 7th

14:00-15:40

Room-2

Session T5C

Chair: Toshihiro Hanawa

Networking

Network Congestion Avoidance through Packet-chaining Reservation

Ke Wu (NUDT), Dezun Dong (NUDT), Cunlu Li (NUDT), Shan Huang (NUDT), Yi Dai (NUDT)

A Tale of Two (Flow) Tables: Demystifying Rule Caching in OpenFlow Switches

Rui Li (School of Computer Science, Fudan university), Yu Pang (School of Computer Science, Fudan university), Jin Zhao (School of Computer Science, Fudan university), Xin Wang (Shanghai Key Laboratory of Intelligent Information Processing, Fudan university)

Artemis: A Practical Low-latency Naming and Routing System

Xuebing Li (Fudan University), Bingyang Liu (Huawei), Yang Chen (Fudan University), Yu Xiao (Aalto University), Jiixin Tang (Fudan University), Xin Wang (Fudan University)

Fast Recovery Techniques for Erasure-coded Clusters in Non-uniform Traffic Network

Yunren Bai (Department of Computer Science and Technology Beijing National Research Center for Information Science and Technology, Tsinghua University), Zihan Xu (Department of Computer Science and Technology Beijing National Research Center for Information Science and Technology, Tsinghua University), Haixia Wang (Beijing National Research Center for Information Science and Technology, Tsinghua University), Dongsheng Wang (Department of Computer Science and Technology Beijing National Research Center for Information Science and Technology, Tsinghua University)

15:40-16:10

Banquet Hall

2F Foyer

Coffee Break

16:10-17:50

Buzz Hall

Session T6A

Chair: Franz Franchetti

Accelerator Applications

Gravitational Octree Code Performance Evaluation on Volta GPU

Yohei Miki (The University of Tokyo)

Gossip: Efficient Communication Primitives for Multi-GPU Systems

Robin Kobus (Johannes Gutenberg University Mainz), Daniel Junger (Johannes Gutenberg University Mainz), Christian Hundt (NVIDIA AI Technology Center), Bertil Schmidt (Johannes Gutenberg University Mainz)

Controlled Asynchronous GVT: Accelerating Parallel Discrete Event Simulation on Many-Core Clusters

Ali Eker (Binghamton University), Barry Williams (Binghamton University), Kenneth Chiu (Binghamton University), Dmitry Ponomarev (Binghamton University)

Distributed Join Algorithms on Multi-GPU Clusters with GPUDirect RDMA

Chengxin Guo (Renmin University of China), Hong Chen (Renmin University of China), Feng Zhang (Renmin University of China), Cuiping Li (Renmin University of China)

16:10-17:50

Room-1

Session T6B

Chair: Allen Malony

Fault Tolerance

Transfer Learning based Failure Prediction for Minority Disks in Large Data Centers of Heterogeneous Disk Systems

Ji Zhang (Huazhong university of science and technology), Ke Zhou (Huazhong university of science and technology), Ping Huang (Temple university), Xubin He (Temple university), Zhili Xiao (Tencent Inc.), Bin Cheng (Tencent Inc.), Yongguang Ji (Tencent Inc.), Yinhu Wang (Tencent Inc.)

Wednesday, August 7th

How to Make the Preconditioned Conjugate Gradient Method Resilient Against Multiple Node Failures

Carlos Pachajoa (University of Vienna), Markus Levonyak (University of Vienna), Wilfried N. Gansterer (University of Vienna), Jesper Larsson Traff (TU Wien)

COMBFT: Conflicting-Order-Match based Byzantine Fault Tolerance Protocol with High Efficiency and Robustness

Yingyao Rong (Sun Yat-sen University), Weigang Wu (Sun Yat-sen University), Zhiguang Chen (Sun Yat-sen University)

Lightweight Fault Tolerance in Pregel-Like Systems

Da Yan (The University of Alabama at Birmingham), James Cheng (CUHK), Hongzhi Chen (CUHK), Cheng Long (Nanyang Technological University), Purushotham Bangalore (The University of Alabama at Birmingham)

16:10-17:50

Room-2

Session T6C

Chair: Kengo Nakajima

Applications 1 - Simulations

diBELLA: Distributed Long Read to Long Read Alignment

Marquita Ellis (University of California at Berkeley, Lawrence Berkeley National Lab), Giulia Guidi (University of California at Berkeley, Lawrence Berkeley National Lab), Aydin Buluc (Lawrence Berkeley National Lab), Leonid Oliker (Lawrence Berkeley National Lab), Katherine Yelick (University of California at Berkeley, Lawrence Berkeley National Lab)

Accelerating Long Read Alignment on Three Processors

Zonghao Feng (Hong Kong University of Science and Technology), Shuang Qiu (Hong Kong University of Science and Technology), Lipeng Wang (Hong Kong University of Science and Technology), Qiong Luo (Hong Kong University of Science and Technology)

Refactoring and Optimizing WRF Model on Sunway TaihuLight

Kai Xu (Shandong University, National Supercomputing Center in Wuxi), Zhenya Song (Laboratory for Regional Oceanography and Numerical Modeling, QNLM First Institute of Oceanography, MNR), Yuandong Chan (Shandong University), Shida Wang (Shandong University, National Supercomputing Center in Wuxi), Xiangxu Meng (Engineering Research Center of Digital Media Technology, Ministry of Education, Shandong University), Weiguo Liu (Shandong University, National Supercomputing Center in Wuxi), Wei Xue (Tsinghua University, National Supercomputing Center in Wuxi)

Improved Unconstrained Energy Functional Method for Eigensolvers in Electronic Structure Calculations *Best Paper Candidate

Mauro Del Ben (Lawrence Berkeley National Laboratory), Osni Marques (Lawrence Berkeley National Laboratory), Andrew Canning (Lawrence Berkeley National Laboratory)

18:20-20:20

Atrium

Banquet

Thursday, August 8th

09:00-09:30 **Registration Open**

B1F Foyer

09:30-10:30 **Keynote-3**

Buzz Hall

Chair: Martin Schulz

Richard Vuduc: Performance engineering for sparse matrix, tensor, and graph computations

10:30-11:00 **Coffee Break**

Banquet Hall

11:00-12:40 **Session T7A**

Buzz Hall

Chair: Balazs Gerofi

Programming Systems and Runtimes

Efficient Data-Parallel Primitives on Heterogeneous Systems

Zhuohang LAI (Hong Kong University of Science and Technology), Qiong LUO (Hong Kong University of Science and Technology), Xiaolong XIE (Alibaba Inc.)

Accelerated Work Stealing

D. Brian Larkins (Rhodes College), John Snyder (Duke University), James Dinan (Intel Corporation)

Runtime Adaptive Task Inlining on Asynchronous Multitasking Runtime Systems

Bibek Wagle (Louisiana State University), Mohammad Alaul Haque Monil (University Of Oregon), Kevin Huck (University Of Oregon), Allen D. Malony (University Of Oregon), Adrian Serio (Louisiana State University), Hartmut Kaiser (Louisiana State University)

HOPE: A Parallel Execution Model Based on Hierarchical Omission

Masahiro Yasugi (Kyushu Institute of Technology), Daisuke Muraoka (Kyushu Institute of Technology), Tasuku Hiraishi (Kyoto University), Seiji Umatani (Kanagawa University), Kento Emoto (Kyushu Institute of Technology)

11:00-12:40 **Session T7B**

Room-1

Chair: Toshio Endo

Performance Modeling

Performance, Energy, and Scalability Analysis and Improvement of Parallel Cancer Deep Learning CANDLER Benchmarks

Xingfu Wu (Argonne National Laboratory, University of Chicago), Valerie Taylor (Argonne National Laboratory, University of Chicago), Justin M. Wozniak (Argonne National Laboratory, University of Chicago), Rick Stevens (Argonne National Laboratory, University of Chicago), Thomas Brettin (Argonne National Laboratory), Fangfang Xia (Argonne National Laboratory)

Adaptive Learning for Concept Drift in Application Performance Modeling

Sandeep Madireddy (Argonne National Laboratory), Prasanna Balaprakash (Argonne National Laboratory), Philip Carns (Argonne National Laboratory), Robert Latham (Argonne National Laboratory), Glenn K. Lockwood (Lawrence Berkeley National Laboratory), Robert Ross (Argonne National Laboratory), Shane Snyder (Argonne National Laboratory), Stefan M. Wild (Argonne National Laboratory)

I/O Characterization and Performance Evaluation of BeeGFS for Deep Learning

Fahim Tahmid Chowdhury (Florida State University), Yue Zhu (Florida State University), Todd Heer (Lawrence Livermore National Laboratory), Saul Paredes (Florida State University), Adam Moody (Lawrence Livermore National Laboratory), Robin Goldstone (Lawrence Livermore National Laboratory), Kathryn Mohror (Lawrence Livermore National Laboratory), Weikuan Yu (Florida State University)

Performance Models for Data Transfers: A Case Study with Molecular Chemistry Kernels

Suraj Kumar (Pacific Northwest National Laboratory), Lionel Eyraud-Dubois (Inria Bordeaux Sud-Ouest), Sriram Krishnamoorthy (Pacific Northwest National Laboratory)

11:00-12:40

Room-2

Session T7C

Chair: Richard Vuduc

Simulation Techniques

OSP: Overlapping Computation and Communication in Parameter Server for Fast Machine Learning

Haozhao Wang (Huazhong University of Science and Technology, PolyU: The Hong Kong Polytechnic University), Song Guo (The Hong Kong Polytechnic University), Ruixuan Li (Huazhong University of Science and Technology)

Automatic Differentiation for Adjoint Stencil Loops *Best Paper Candidate

Jan Huckelheim (Imperial College London), Navjot Kukreja (Imperial College London), Sri Hari Krishna Narayanan (Argonne National Laboratory), Fabio Luporini (Imperial College London), Gerard Gorman (Imperial College London), Paul Hovland (Argonne National Laboratory)

Spatially-aware Parallel I/O for Particle Data

Sidharth Kumar (University of Alabama at Birmingham), Steve Petruzza (University of Utah), Will Usher (University of Utah), Valerio Pascucci (University of Utah)

The Communication-Overlapped Hybrid Decomposition Parallel Algorithm for Multi-Scale Fluid Simulations

Yi Liu (National University of Defense Technology), Xiao-Wei Guo (National University of Defense Technology), Chao Li (National University of Defense Technology), Canqun Yang (National University of Defense Technology), Xinbiao Gan (National University of Defense Technology), Peng Zhang (National University of Defense Technology), Yi Wang (National University of Defense Technology), Ran Zhao (National University of Defense Technology), Sijiang Fan (National University of Defense Technology)

12:40-14:00

Lunch

Please take a lunch box at the registration desk and have lunch at Banquet Hall or the conference rooms.

14:00-15:40

Buzz Hall

Session T8A

Chair: Tetsuya Hoshino

Deep Learning

Cynthia: Cost-Efficient Cloud Resource Provisioning for Predictable Distributed Deep Neural Network Training

Haoyue Zheng (East China Normal University), Fei Xu (East China Normal University), Li Chen (University of Louisiana at Lafayette), Zhi Zhou (Sun Yat-sen University), Fangming Liu (Huazhong University of Science and Technology)

FlowCon: Elastic Flow Configuration for Containerized Deep Learning Applications

Wenjia Zheng (Fordham University), Michael Tynes (Fordham University), Henry Gorelick (Fordham University), Ying Mao (Fordham University), Long Cheng (University College Dublin), Yantian Hou (Boise State University)

DLBooster: Boosting End-to-End Deep Learning Workflows with Offloading Data Preprocessing Pipelines

Yang Cheng (Tsinghua University, Microsoft Research), Dan Li (Tsinghua University), Zhiyuan Guo (Microsoft Research, Beihang University), Binyao Jiang (Microsoft Research, Shanghai Jiao Tong University), Jiaxin Lin (Microsoft Research, Beihang University), Xi Fan (Microsoft Research, Shanghai Jiao Tong University), Jinkun Geng (Tsinghua University), Xinyi Yu (Microsoft Research, Shanghai Jiao Tong University), Wei Bai (Microsoft Research), Lei Qu (Microsoft Research), Ran Shu (Microsoft Research), Peng Cheng (Microsoft Research), Yongqiang Xiong (Microsoft Research), Jianping Wu (Tsinghua University)

swATOP: Automatically Optimizing Deep Learning Operators on SW26010 Many-Core Processor

Wei Gao (Tsinghua University), Jiarui Fang (Tsinghua University), Wenlai Zhao (Tsinghua University), Jinzhe Yang (Imperial College London), Long Wang (Baidu), Lin Gan (Tsinghua University, National Supercomputing Center in Wuxi), Haohuan Fu (Tsinghua University, National Supercomputing Center in Wuxi), Guangwen Yang (Tsinghua University, National Supercomputing Center in Wuxi)

14:00-15:40

Room-1

Session T8B

Chair: Taisuke Boku

Tools and Their Use

A Plugin Architecture for the TAU Performance System

Allen D. Malony (University of Oregon), Srinivasan Ramesh (University of Oregon), Kevin Huck (University of Oregon), Nicholas Chaimov (ParaTools, Inc.), Sameer Shende (University of Oregon)

FancyTuner: Auto-tuning Scientific Applications With Per-loop Compilation *Best Paper Candidate

Tao Wang (North Carolina State university), Nikhil Jain (Lawrence Livermore National Laboratory), David Beckingsale (Lawrence Livermore National Laboratory), David Boehme (Lawrence Livermore National Laboratory), Frank Mueller (North Carolina State university), Todd Gamblin (Lawrence Livermore National Laboratory)

Massively Parallel Automated Software Tuning

Jakub Kurzak (University of Tennessee), Yaohung Tsai (University of Tennessee), Mark Gates (University of Tennessee), Ahmad Abdelfattah (University of Tennessee), Jack Dongarra (University of Tennessee)

Exploiting Vector Processing in Dynamic Binary Translation

Chih-Min Lin (National Taiwan University), Sheng-Yu Fu (National Taiwan University), Ding-Yong Hong (Academia Sinica), Yu-Ping Liu (National Taiwan University), Jan-Jan Wu (Academia Sinica), Wei-Chung Hsu (National Taiwan University)

14:00-15:40

Room-2

Session T8C

Chair: Rio Yokota

Applications 2 - Emerging Applications

Cosin: Controllable Social Influence Maximization and Its Distributed Implementation in Large-scale Social Networks

Jingya Zhou (Soochow University, State Key Laboratory for Novel Software Technology), Jianxi Fan (Soochow University, Provincial Key Laboratory for Computer Information Processing Technology), Jin Wang (Soochow University, Provincial Key Laboratory for Computer Information Processing Technology)

Approximate Code: A Cost-Effective Erasure Coding Framework for Tiered Video Storage in Cloud Systems

Huayi Jin (Shanghai Jiao Tong University), Chentao Wu (Shanghai Jiao Tong University), Xin Xie (Shanghai Jiao Tong University), Jie Li (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University), Hao Lin (The Alibaba Group), Jianfeng Zhang (The Alibaba Group)

VScan: Efficiently Analyzing Surveillance Videos via Model-joint Mechanism

Chen Zhang (Huazhong University of Science and Technology), Qiang Cao (Huazhong University of Science and Technology), Jie Yao (Huazhong University of Science and Technology), Yuanyuan Dong (Alibaba Group), Puyuan Yang (Alibaba Group)

Faster parallel collision detection at high resolution for CNC milling applications

Xin Chen (Georgia Tech), Dmytro Konobrytskyi (Uber Advanced Technologies Group), Thomas M. Tucker (Tucker Innovations Inc.), Thomas R. Kurfess (Georgia Tech), Richard W. Vuduc (Georgia Tech)

Thursday, August 8th

15:40-16:00 **Coffee Break**

Banquet Hall

2F Foyer

16:00-17:15 **Session T9A**

Chair: Toshihiro Hanawa

Buzz Hall

Neural Networks

An Efficient Design Flow for Accelerating Complicated-connected CNNs on a Multi-FPGA Platform

Deguang Wang (National University of Defense Technology), Junzhong Shen (National University of Defense Technology), Mei Wen (National University of Defense Technology), Chunyuan Zhang (National University of Defense Technology)

A Unified Optimization Approach for CNN Model Inference on Integrated GPUs

Leyuan Wang (Amazon, UC Davis), Zhi Chen (Amazon), Yizhi Liu (Amazon), Yao Wang (Amazon), Lianmin Zheng (Shanghai Jiao Tong University), Mu Li (Amazon), Yida Wang (Amazon)

Massively Parallel ANS Decoding on GPUs

Andre Weissenberger (Johann Wolfgang Goethe University), Bertil Schmidt (Johannes Gutenberg University of Mainz)

16:00-17:15 **Session T9B**

Chair: Sivan Toledo

Room-1

Parallel Data Structures

Building Scalable NVM-based B+tree with HTM

Mengxing Liu (Tsinghua University), Jiankai Xin (Tsinghua University), Kang Chen (Tsinghua University), Yongwei Wu (Tsinghua University)

BCL: A Cross-Platform Distributed Data Structures Library

Benjamin A. Brock (University of California, Berkeley Lawrence Berkeley National Laboratory), Aydn Buluc (Lawrence Berkeley National Laboratory University of California, Berkeley), Katherine Yelick (Lawrence Berkeley National Laboratory University of California, Berkeley)

On Integration of Appends and Merges in Log-Structured Merge Trees

Caixin Gong (Alibaba Group), Shuibing He (Zhejiang University), Yili Gong (Wuhan University), Yingchun Lei (Daowoo Times Tech. Co.)

16:00-17:15 **Session T9C**

Chair: Yohei Miki

Room-2

IoT and Edge Computing

NFV-Enabled Multicasting in Mobile Edge Clouds with Resource Sharing

Zichuan Xu (Dalian University of Technology), Yutong Zhang (Dalian University of Technology), Weifa Liang (The Australian National University), Qiufen Xia (Dalian University of Technology), Omer Rana (Cardiff University), Alex Galis (University College London), Guowei Wu (Dalian University of Technology), Pan Zhou (Huazhong University of Science and Technology)

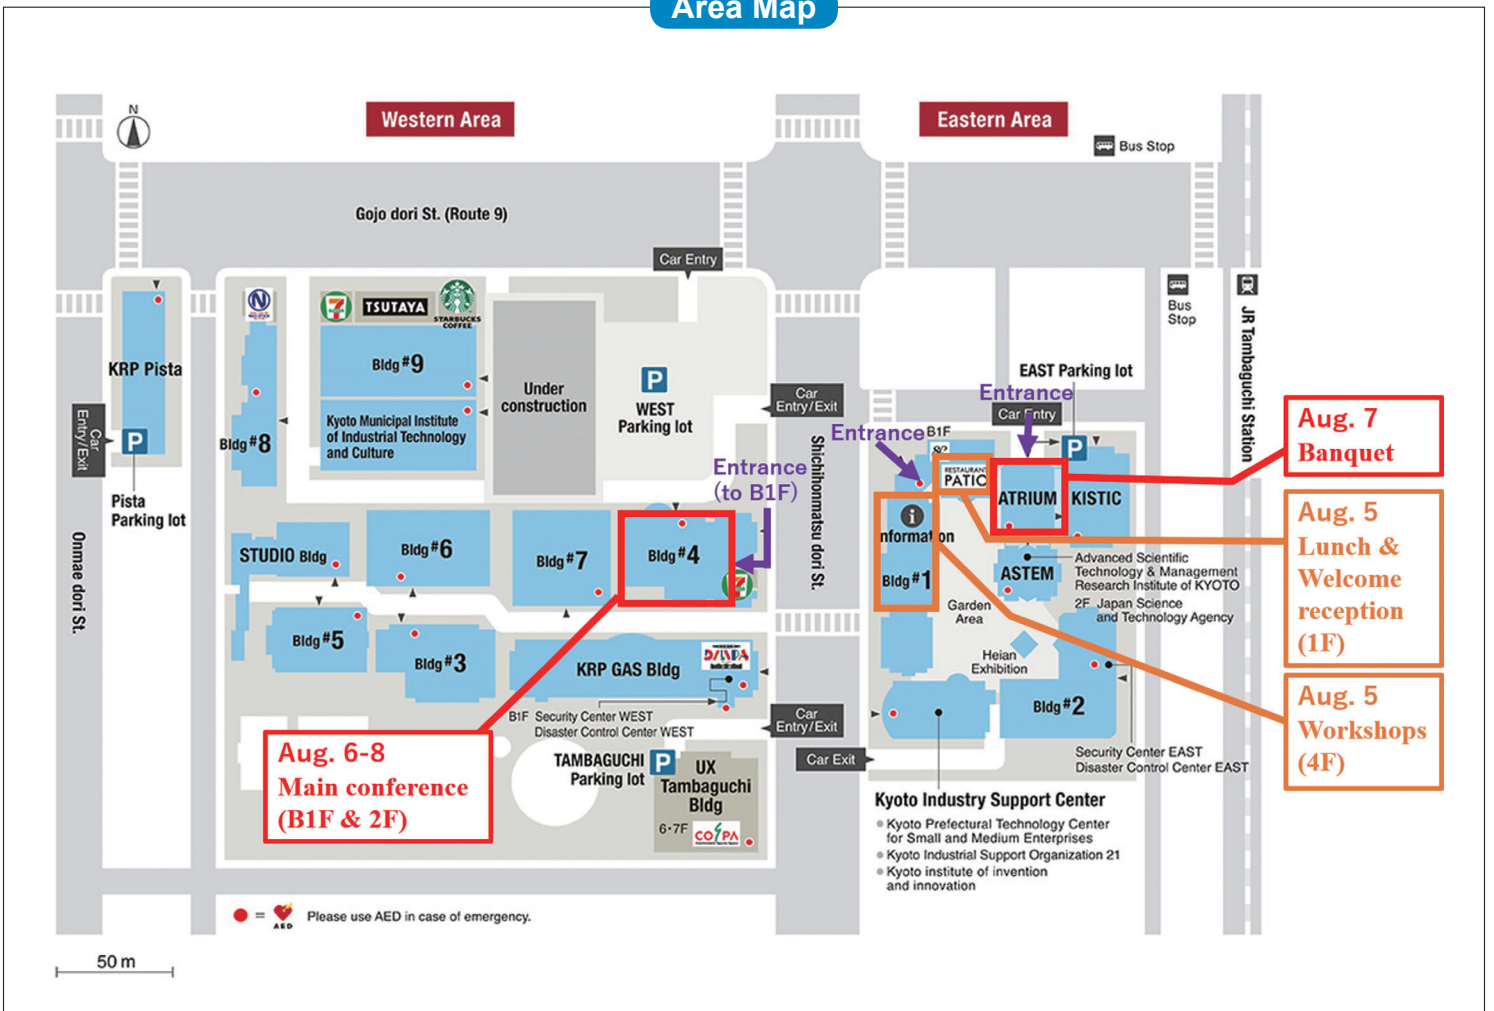
QLEC: A Machine-Learning-Based Energy-Efficient Clustering Algorithm to Prolong Network Lifespan for IoT in High-Dimensional Space

Ke Li (Shanghai Jiao Tong University), Haowei Huang (Shanghai Jiao Tong University), Xiaofeng Gao (Shanghai Jiao Tong University), Fan Wu (Shanghai Jiao Tong University), Guihai Chen (Shanghai Jiao Tong University)

Multi-Objective Reinforcement Learning for Reconfiguring Data Stream Analytics on Edge Computing

Alexandre da Silva Veith (Ecole Normale Supérieure de Lyon, Inria), Felipe Rodrigo de Souza (Ecole Normale Supérieure de Lyon, Inria), Marcos Dias de Assuncao (Ecole Normale Supérieure de Lyon, Inria), Laurent Lefevre (Ecole Normale Supérieure de Lyon, Inria), Julio Cesar Santos dos Anjos (Federal University of Rio Grande do Sul)

Area Map



Floor Map

