

ICPP 2019 Accepted Papers List

Algorithms Track

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)

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)

Predictable GPUs Frequency Scaling for Energy and Performance

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

Tessellating Star Stencils

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

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)

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)

On Max-min Fair Resource Allocation for Distributed Job Execution

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

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)

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)

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 Nebraska-Lincoln)

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

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

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)

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)

Parallel Algorithms for Evaluating Matrix Polynomials

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

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)

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)

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)

A Practical, Scalable, Relaxed Priority Queue

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

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)

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)

Applications Track

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)

ECoS_T: 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)

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.)

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)

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)

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

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

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)

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)

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)

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)

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)

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)

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)

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)

Massively Parallel ANS Decoding on GPUs

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

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)

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)

Improved Unconstrained Energy Functional Method for Eigensolvers in Electronic Structure Calculations

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

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)

diBELLA: Distributed Long Read to Long Read Alignment

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

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)

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

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

Architecture Track

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)

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)

Building Scalable NVM-based B+tree with HTM

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

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)

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)

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)

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)

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)

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)

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)

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)

TLB: Traffic-aware Load Balancing with Adaptive Granularity in Data Center Networks

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)

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)

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)

Network Congestion Avoidance through Packet-chaining Reservation

Ke Wu (NUDT), Dezun Dong (NUDT)

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)

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.)

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

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

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), Jiabin 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)

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), Jiabin Tang (Fudan University), Xin Wang (Fudan University)

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)

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)

Performance Track

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)

Gravitational Octree Code Performance Evaluation on Volta GPU

Yohei Miki (The University of Tokyo)

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)

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)

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), Paul Richmond (The University of Sheffield), Daniel Coca (The University of Sheffield)

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)

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)

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)

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)

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)

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)

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)

Performance, Energy, and Scalability Analysis and Improvement of Parallel Cancer Deep Learning CANDLE 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)

Automatic Differentiation for Adjoint Stencil Loops

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)

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), Hartmut Kaiser (Louisiana State University), Allen D. Malony (University Of Oregon), Adrian Serio (Louisiana State University)

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)

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.)

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)

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)

Software Track

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)

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)

FuncyTuner: Auto-tuning Scientific Applications With Per-loop Compilation

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)

Cartesian Collective Communication

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

Machine Learning for Fine-Grained Hardware Prefetcher Control

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

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)

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)

Accelerated Work Stealing

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

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

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

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)

Speculative Scheduling for Stochastic HPC Applications

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

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)

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)

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)

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)

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)

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)

Holistic Slowdown Driven Scheduling and Resource Management for Malleable Jobs

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

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)

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)

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)