

Time **Friday 01 Sep, 2023**

08:30 **Registration**

09:00 – 10:00 **Plenary Session Keynote 3**

[Jahna Otterbacher](#)

10:00 – 10:30 **Coffee Break**

10:30 – 12:30 **Parallel Sessions**

T06: HPC applications

23 SWSPH: A Massively Parallel SPH Implementation for Hundred-Billion-Particle Simulation on New Sunway Supercomputer

[Ziyu Zhang](#), Junshi Chen, Zhanming Wang, Yifan Luo, Jineng Yao, Hong An and Shenghong Huang

161 Hercules: scalable and network portable in-memory ad-hoc file system for data-centric and high-performance applications

[Javier Garcia Blas](#), Genaro Sanchez-Gallegos, Cosmin Petre, Alberto Riccardo Martinelli, Marco Aldinucci and Jesus Carretero

186 Optimizing Distributed Tensor Contractions using Node-Aware Processor Grids

[Andreas Irmeler](#), Raghavendra Kanakagiri, Sebastian Ohlmann, Edgar Solomonik and Andreas Grüneis

110 FedCML: Federated Clustering Mutual Learning Framework with non-IID Data

Zekai Chen, Fuyi Wang, Shengxing Yu, [Ximeng Liu](#) and Zhiwei Zheng

T02: Resource management in the Cloud-Edge continuum

136 Geofence-Based Service Discovery in the Computing Continuum
[Kurt Horvath](#), Dragi Kimovski, Christoph Uran, Helmut Wöllik and Radu Prodan

218 Assessing power needs to run a workload with quality of service on green datacenters

Louis-Claude Canon, [Damien Landré](#), Laurent Philippe, Jean-Marc Pierson and Paul Renaud-Goud

145 SparkEdgeEmu: An Emulation Framework for Edge-enabled Apache Spark Deployments

[Moysis Symeonides](#), Demetris Trihinas, George Pallis and Marios D. Dikaiakos

42 MetaLive: Meta-Reinforcement Learning Based Collective Bitrate Adaptation for Multi-Party Live Streaming

[Yi Yang](#), Xiang Li, Yeting Xu, Wenzhong Li, Jiangyi Hu, Taishan Xu, Xiancheng Ren and Sanglu Lu **MetaLive:**

12:30 – 14:00 **Lunch**

14:00 – 15:30 **Parallel Sessions**

T03: Architectures and Accelerators (1)

69 Lock-free Bucketized Cuckoo Hashing

Wenhai Li, [Zhiling Cheng](#), Yuan Chen, Ao Li and Lingfeng Deng

T06: Various

35 Transactional-Turn Causal Consistency

Benoît Martin, Laurent Prospero and Marc Shapiro

49 A Multi-level Parallel Integer/Floating-Point Arithmetic Architecture for Deep Learning Instructions

Hongbing Tan, Jing Zhang, [Libo Huang](#), Xiaowei He, Dezun Dong, Yongwen Wang and Liqun Xiao

185 Computational Storage for an Energy-Efficient Deep Neural Network Training System

[Shiju Li](#), Kevin Tang, Jin Lim, Chul-Ho Lee and Jongryool Kim

197 Parallel Cholesky Factorization for Banded Matrices using OpenMP Tasks

Felix Liu, Albin Fredriksson and Stefano Markidis

155 Efficient Concurrent Transaction Execution Framework for Blockchains

Manaswini Piduguralla, Saheli Chakraborty, Anjana Parwat Singh and Sathya Peri

15:30 – 16:00

Coffee Break

16:00 – 17:30

Parallel Sessions

T03: Architectures and Accelerators (2)

18 Improving Utilization of Dataflow Architectures through Software and Hardware Co-Design

[Fan Zhihua](#), Wenming Li, Shengzhong Tang, Xuejun An, Xiaochun Ye and Dongrui Fan

144 BitHist: A Precision-Scalable Sparse-Awareness DNN Accelerator Based on Bit Slices Products Histogram

[Zhaoteng Meng](#), Long Xiao, Xiaoyao Gao, Lin Shu and Jie Hao

T04: Data Analytics and AI

90 Model-Agnostic Federated Learning

[Gianluca Mittone](#), Walter Riviera, Iacopo Colonnelli, Robert Birke and Marco Aldinucci

76 Auto-Divide GNN: Accelerating GNN Training with Subgraph Division

[Hongyu Chen](#), Zhejiang Ran, Keshi Ge, Zhiquan Lai, Jingfei Jiang and Dongsheng Li

43 FedGM: Heterogeneous Federated Learning via Generative Learning and Mutual Distillation

Chao Peng, [Yiming Guo](#), Yao Chen, Qilin Rui, Zhengfeng Yang and Chenyang Xu