Research Projects Selected for Poster Presentation

1 Kobe Bergmans (KU Leuven), *Algorithms for Parallel Shared-Memory Sparse Matrix-Vector Multiplication on Unstructured Matrices*

2 Tor Andre Haugdahl (Norwegian University of Science and Technology), *Parallelization and Scheduling Opportunities in the Modelica Advanced Research Compiler (MARCO)*

3 Nathaniel Tomczak (Case Western Reserve University), *A Graph-Processing Perspective for Attention to Accelerate Sparse Transformers*

4 Aranya Banerjee (Georgia Institute of Technology), *Advancing Algorithms for Scalable Genome Sequence Alignment*

5 Ruiyang Chen (Shanghai Jiao Tong University), *GIFTS: Efficient \underline{G}CN \underline{I}nference \underline{F}ramework on Py\underline{T}orch-CPU via Exploring the \underline{S}parsity*

6 Madhukara Kekulandara (University of Rhode Island), *A Comparative Analysis of Parallel Markov Chain Monte Carlo Algorithms for Redistricting*

7 Aishwarya R. Parab (BITS Pilani, KK Birla Goa campus, India), *A Scalable and Upgradable Framework for Decentralized Applications Leveraging Smart Contracts in Blockchain*

8 Bowen Sun (William and Mary), *PeakLife: Better Cloud Management Using a Deep Surrogate Model for VM Utilization and Lifetime*

9 Chen-Chun Chen (The Ohio State University), *Unified Designs of Multi-rail-aware MPI Allreduce and Alltoall Operations Across Diverse GPU and Interconnect Systems*

10 Durga Keerthi Mandarapu (Purdue), *Accelerating spatial queries using GPU Ray Tracing*

11 Jianxiong Li (SKLP, Institute of Computing Technology, Chinese Academy of Sciences), *Enhance the Strong Scaling of Neural-Network-Based Molecular Dynamics with Long-Range Electrostatic Interactions*

12 Joy Kitson (University of Maryland), *Scalable Epidemiological Agent-based Modeling with Dynamic Behaviors*

13 Marco D'Antonio (Queen's University Belfast), *Toward Efficient Asynchronous Single-Source Shortest Path*

14 ioanna Tasou (NTUA), *Exploring the Performance of Sparsified Transformer Inference*

15 Md Hasanur Rashid (University of Delaware), *AdapTBF: Decentralized Bandwidth Control via Adaptive Token Borrowing for HPC Storage*

16 Eman Alqudah (Iowa State University), *Optimization of Pilot Sequences Assignment for Real-Time Massive MIMO Systems using Parallel Genetic Algorithm on FPGA*

17 Ali Khan (University of North Texas), *Parallel Motif Counting for Large-Scale Dynamic Datasets*

18 Nikolaos Triantafyllis (National Technical University of Athens), *Outlining an HPC Co-Scheduler*

19 Wentao Feng (Beihang University), *Memory-Centric Profiling of Arm-based Supercomputing Systems*

20 Wahid Uz Zaman (The Pennsylvania State University), *Optimizations for Computational Storage Devices*

21 Upasana Sridhar (Carnegie Mellon University), *A Unified Abstraction for Efficient and Portable CNN Libraries*

22 Julian Bellavita (Cornell University), *Efficient Large-Scale Multi-GPU Clustering using Sparse Linear Algebra*

23 Chiang-Heng Chien (Brown University), *GPU-HC: A Paradigm Shift for Solving Large Polynomial Systems*

24 Lorenzo Carpentieri (University of Salerno), *Energy Saving By Approximate Computing*

25 Jelle van Dijk (Universiteit van Amsterdam), *Improving the Energy Efficiency of Large-Scale Scientific Simulations*

26 Saeedeh Baneshi (University of Amsterdam), *Energy Labeling of Digital Services in the Computing Continuum*