

2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW 2024)

**San Francisco, California, USA
27-31 May 2024**

Pages 1-617



**IEEE Catalog Number: CFP2451J-POD
ISBN: 979-8-3503-6461-3**

**Copyright © 2024 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP2451J-POD
ISBN (Print-On-Demand):	979-8-3503-6461-3
ISBN (Online):	979-8-3503-6460-6

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

TABLE OF CONTENTS

The 33rd Heterogeneity in Computing Workshop (HCW) - Committees	1
Message from the HCW 2024 Steering Committee Co-Chairs	2
Message from the HCW 2024 General Co-Chairs	3
Message from the HCW 2024 Technical Program Committee Co-Chairs	4
HCW 2024 Keynote - Hetero: Where We've Been, Where We Are, and What Next?	5
<i>Yale Patt</i>	
Performance Portability of the Chapel Language on Heterogeneous Architectures.....	6
<i>Josh Milthorpe, Xianghao Wang, Ahmad Azizi</i>	
Towards Dynamic Autotuning of SpMV in CUSP Library	14
<i>Miroslav Demek, Jiří Filipovič</i>	
A Runtime Manager Integrated Emulation Environment for Heterogeneous SoC Design with RISC-V Cores.....	23
<i>H. Umut Suluhan, Serhan Gener, Alexander Fusco, Joshua Mack, Ismet Dagli, Mehmet Belviranli, Cagatay Edemen, Ali Akoglu</i>	
Dynamic Tasks Scheduling with Multiple Priorities on Heterogeneous Computing Systems	31
<i>Hayfa Tayeb, Béranger Bramas, Mathieu Favergé, Abdou Guermouche</i>	
PSyGS Gen a Generator of Domain-Specific Architectures to Accelerate Sparse Linear System Resolution.....	41
<i>Niccolò Nicolosi, Francesco Renato Negri, Francesco Pesce, Francesco Peverelli, Davide Conficconi, Marco Domenico Santambrogio</i>	
Toward a Holistic Performance Evaluation of Large Language Models Across Diverse AI Accelerators.....	48
<i>Murali Emani, Sam Foreman, Varuni Sastry, Zhen Xie, Siddhisanket Raskar, William Arnold, Rajeev Thakur, Venkatram Vishwanath, Michael E. Papka, Sanjiv Shamugavelu, Darshan Gandhi, Hengyu Zhao, Dun Ma, Kiran Ranganath, Rick Weisner, Jiunn-yeu Chen, Yuting Yang, Natalia Vassilieva, Bin C. Zhang, Sylvia Howland, Alexander Tsyplykhin</i>	
IRIS: Exploring Performance Scaling of the Intelligent Runtime System and Its Dynamic Scheduling Policies	58
<i>Beau Johnston, Narasinga Rao Miniskar, Aaron Young, Mohammad Alaul Haque Monil, Seyong Lee, Jeffrey S. Vetter</i>	
Heterogeneous Hyperthreading	68
<i>Mingxuan He, Fangping Liu, Sang Wook Stephen Do</i>	
31st Reconfigurable Architectures Workshop (RAW 2024)	79
<i>Jürgen Becker, Zhenman Fang, Viktor K. Prasanna, Marco Santambrogio, Ramachandran Vaidyanathan</i>	
RAW 2024 Committees	80
<i>Marco Domenico Santambrogio</i>	

RAW 2024 Monday Keynote	82
<i>Deming Chen</i>	
RAW 2024 Invited Talk-1: Auto-Generating Diverse Heterogeneous Designs	83
<i>Wayne Luk</i>	
RAW 2024 Invited Talk-2: Digital In-Memory Computing to Accelerate Deep Learning Inference on the Edge	84
<i>Stefania Perri</i>	
RAW 2024 Invited Talk-3: Self-aware Reliable and Reconfigurable Computing Systems — An Overview	85
<i>Diana Göhringer</i>	
RAW 2024 Invited Talk-4: Reconfigurable Computing: Quo Vadis?	86
<i>Dirk Stroobandt</i>	
RAW 2024 Invited Talk-5: Reconfigurable AI Processing for Embedded Systems	87
<i>Masato Motomura</i>	
RAW 2024 Invited Talk-6: Reconfigurable Architectures for High-Performance Computing	88
<i>Kentaro Sano</i>	
RAW 2024 Invited Talk-7: Efficient Cross-layer Design Flow for Multi-die FPGA	89
<i>Wei Zhang</i>	
RAW 2024 Invited Talk-8: Practical Reconfigurable Computing for Next-Generation Edge Applications	90
<i>Hayden Kwok-Hay So</i>	
RAW 2024 Invited Talk-9: Riallto: An Open-Source Exploratory Framework for Ryzen AI™	91
<i>Andrew Schmidt</i>	
FPGA Acceleration of DL-Based Real-Time DC Series Arc Fault Detection	92
<i>Yufei Mao, Roland Weiss, Yi Zhang, Yu Li, Marc Rothmann, Mario Porrmann</i>	
An Accurate Union Find Decoder for Quantum Error Correction on the Toric Code	99
<i>Federico Valentino, Beatrice Branchini, Davide Conficconi, Donatella Sciuto, Marco D. Santambrogio</i>	
Towards the Acceleration of the Sparse Blossom Algorithm for Quantum Error Correction	106
<i>Marco Venere, Valentino Guerrini, Beatrice Branchini, Davide Conficconi, Donatella Sciuto, Marco D. Santambrogio</i>	
Exploring Large Language Models for Verilog Hardware Design Generation	111
<i>Erik H. D'Hollander, Ewout Danneels, Karel-Brech Decorte, Senne Loobuyck, Arne Vanheule, Ian Van Kets, Dirk Stroobandt</i>	
Auto-Generating Diverse Heterogeneous Designs	116
<i>Jessica Vandebon, Jose G.F. Coutinho, Wayne Luk</i>	
Self-Aware Reliable and Reconfigurable Computing Systems - an Overview	124
<i>Diana Göhringer, Ariel Podlubne, Fabian Vargas, Milos Krstic</i>	
Digital In-Memory Computing to Accelerate Deep Learning Inference on the Edge	130
<i>Stefania Perri, Cristian Zambelli, Daniele Ielmini, Cristina Silvano</i>	

A Fast Scalable Hardware Priority Queue and Optimizations for Multi-Pushes	134
<i>Samuel Collinson, Allan Bai, Oliver Sinnen</i>	
FPGA-Based Implementation for Industrial Motion Control System	141
<i>Claudio Rubattu, Antonio Ledda, Francesco Ratto, Chaitanya Jugade, Dip Goswami, Francesca Palumbo</i>	
An FPGA-Based Accelerator for Graph Embedding Using Sequential Training Algorithm.....	148
<i>Kazuki Sunaga, Keisuke Sugiura, Hiroki Matsutani</i>	
TaPaSCo-AIE: An Open-Source Framework for Streaming-Based Heterogeneous Acceleration Using AMD AI Engines.....	155
<i>Carsten Heinz, Torben Kalkhof, Yannick Lavan, Andreas Koch</i>	
An Architectural Template for FPGA Overlays Targeting Data Flow Applications	162
<i>Anna Drewes, Vitalii Burtsev, Bala Gurumurthy, Martin Wilhelm, David Broneske, Gunter Saake, Thilo Pionteck</i>	
Performance Evaluation of VirtIO Device Drivers for Host-FPGA PCIe Communication.....	169
<i>Sahan Bandara, Ahmed Sanaullah, Zaid Tahir, Ulrich Drepper, Martin Herbordt</i>	
Accelerating TinyML Inference on Microcontrollers Through Approximate Kernels	177
<i>Giorgos Armeniakos, Georgios Mentzos, Dimitrios Soudris</i>	
A Case for Low Bitwidth Floating Point Arithmetic on FPGA for Transformer Based DNN Inference	178
<i>Jiajun Wu, Mo Song, Jingmin Zhao, Hayden Kwok-Hay So</i>	
HPC Systems with Reconfigurable Optical Networks: Performance and Energy Consumption Exploration	186
<i>Xianwei Cheng, Che-Yu Liu, Roberto Proietti, S. J. Ben Yoo</i>	
Balancing Intra-Die and Inter-Die Placement Optimization in 2.5D FPGA Architectures.....	187
<i>Raveena Raikar, Dirk Stroobandt</i>	
Scalable Dual-Instruction Multiple-data Processing on an Efficient Systolic-array Architecture	188
<i>Yuxi Tan, Riadh Ben Abdelhamid, BingJie Guo, Qixiang Gao, Masaru Nishimura, Yoshiki Yamaguchi</i>	
ConvMap: Boosting Convolution Throughput on FPGAs with Efficient Resource Mapping	189
<i>Shubhayu Das, Nanditha Rao, Sharad Sinha</i>	
A Reconfigurable Architecture of a Scalable, Ultrafast, Ultrasound, Delay-And-Sum Beamformer	190
<i>V. Kyriotis, G. Smaragdos, P. Kruizinga, D. Soudris, C. Strydis</i>	
ML-Based Real-Time Control at the Edge: An Approach Using Hls4ml.....	191
<i>R. Shi, S. Ogreni, J.M. Arnold, J.R. Berlioz, P. Hanlet, K.J. Hazelwood, M.A. Ibrahim, H. Liu, V.P. Nagaslaev, A. Narayanan, D.J. Nicklaus, J. Mitrevski, G. Pradhan, A.L. Saewert, B.A. Schupbach, K. Seiya, M. Thieme, R.M. Thurman-Keup, N.V. Tran</i>	
Network Adapter for Secure Networks-On-Chip	192
<i>Julian Haase, Nico Volkens, Diana Göringer</i>	
Multi-Core Multi-Rule VeBPF Firewall for Secure FPGA IoT Device Deployments.....	193
<i>Zaid Tahir, Sahan Bandara, Martin Herbordt</i>	

POCA: A PYNQ Offloaded Cryptographic Accelerator on Embedded FPGA-Based Systems	194
<i>Roberto A. Bertolini, Filippo Carloni, Davide Conficconi, Marco Domenico Santambrogio</i>	
APDCM 2024 Preface and Committee List	195
<i>Jacir L. Bordim, Koji Nakano</i>	
APDCM 2024 Keynote Talk	197
<i>Hiroki Otsuji</i>	
Application of Network Calculus Models to Heterogeneous Streaming Applications.....	198
<i>Clayton J. Faber, Roger D. Chamberlain</i>	
Data-Driven Locality-Aware Batch Scheduling	202
<i>Maxime Gonthier, Elisabeth Larsson, Loris Marchal, Carl Nettelblad, Samuel Thibault</i>	
Combining Lossy Compression with Multi-Level Caching for Data Staging Over Network	212
<i>Rei Aoyagi, Keichi Takahashi, Yoichi Shimomura, Hiroyuki Takizawa</i>	
A Scalable Secure Fault Tolerant Aggregation for P2P Federated Learning	222
<i>Yujiro Yahata, Keisuke Sugiura, Hiroki Matsutani</i>	
Accelerating BFT Database with Transaction Reconstruction	232
<i>Aoi Kida, Hideyuki Kawashima</i>	
Optimizing Aria Concurrency Control Protocol with Early Dependency Resolution	242
<i>Yusuke Miyazaki, Takashi Hoshino, Hideyuki Kawashima</i>	
Shared-Memory Parallel Algorithms for Community Detection in Dynamic Graphs.....	250
<i>Subhajit Sahu, Kishore Kothapalli, Dip Sankar Banerjee</i>	
A Probabilistic Model for Asynchronous Iterative Methods	260
<i>Pratik Nayak, Hartwig Anzt</i>	
The Logarithmic Random Bidding for the Parallel Roulette Wheel Selection with Precise Probabilities.....	270
<i>Koji Nakano</i>	
Introduction to Computational Quantum Chemistry for Computer Scientists.....	273
<i>Yasuaki Ito, Satoki Tsuji, Haruto Fujii, Kanta Suzuki, Nobuya Yokogawa, Koji Nakano, Akihiko Kasagi</i>	
AsHES 2024 Preface and Committee List.....	283
<i>Shintaro Iwasaki</i>	
Block-Based GPU Programming with Triton.....	285
<i>Philippe Tillet</i>	
Performance Versus Maintainability: A Case Study of Scream on Frontier	286
<i>James B. White</i>	
ParaGraph: Weighted Graph Representation for Performance Optimization of HPC Kernels	293
<i>Ali TehraniJamsaz, Alok Mishra, Akash Dutta, Abid M. Malik, Barbara Chapman, Ali Jannesarri</i>	
Alternative Quadrant Representations with Morton Index and AVX2 Vectorization for AMR Algorithms Within the P4est Software Library	301
<i>Mikhail Kirilin, Carsten Burstedde</i>	

Avoiding Training in the Platform-Aware Optimization Process for Faster DNN Latency Reduction	311
<i>Raúl Marichal, Ernesto Dufrechou, Pablo Ezzatti</i>	
A Comparative Study on Simulation Frameworks for AI Accelerator Evaluation	321
<i>Christoffer Åleskog, Håkan Grahn, Anton Borg</i>	
Extending the SYCL Joint Matrix for Binarized Neural Networks.....	329
<i>Zheming Jin</i>	
Message from the EduPar-24 Workshop Chairs	334
<i>Sushil K. Prasad</i>	
EduPar-24 Workshop Organization	335
<i>Sushil K. Prasad</i>	
EduPar 2024 Keynote Speaker.....	337
<i>Charles E. Leiserson</i>	
Helping Faculty Teach Software Performance Engineering.....	338
<i>John D. Owens, Bruce Hoppe</i>	
Parallel Optimization for Robotics: An Undergraduate Introduction to GPU Parallel Programming and Numerical Optimization Research.....	342
<i>Brian Plancher</i>	
Teaching Parallel Algorithms Using the Binary-Forking Model.....	346
<i>Guy E. Blelloch, Yan Gu, Yihan Sun</i>	
Peachy Parallel Assignments (EduPar 2024).....	352
<i>Alina Lazar, Ethan Scheelk, Elizabeth Shoop, David P. Bunde</i>	
Codeless PDC Modules for Early Computing Curriculum.....	357
<i>Chris Bourke, Justin Firestone</i>	
Visualizing PRAM Algorithm for Mergesort	365
<i>Cade Wiley, Grey Ballard</i>	
Integrating Interactive Performance Analysis in Jupyter Notebooks for Parallel Programming Education.....	369
<i>Lena Oden, Klaus Nölpp, Philipp Brauner</i>	
Interactive Textbooks for Parallel and Distributed Computing Across the Undergraduate CS Curriculum.....	377
<i>Elizabeth Shoop, Richard Brown, Suzanne J. Matthews, Joel C. Adams</i>	
Teaching Performance Metrics in Parallel Computing Courses	385
<i>Sandino Vargas-Pérez</i>	
Speedcode: Software Performance Engineering Education Via the Coding of Didactic Exercises.....	391
<i>Tim Kaler, Xuhao Chen, Brian Wheatman, Dorothy Curtis, Bruce Hoppe, Tao B. Schardl, Charles E. Leiserson</i>	
ESSA 2024 Message and Committees.....	395
<i>François Tessier, Weikuan Yu</i>	

The Impact of Asynchronous I/O in Checkpoint-Restart Workloads	397
<i>Hariharan Devarajan, Adam Moody, Donglai Dai, Cameron Stanavige, Elsa Gonsiorowski, Marty McFadden, Olaf Faaland, Greg Kosinovsky, Kathryn Mohror</i>	
Benchmarking Variables for Checkpointing in HPC Applications	406
<i>Xiang Fu, Xin Huang, Wubiao Xu, Weiping Zhang, Shiman Meng, Luanzheng Guo, Kento Sato</i>	
Extending the Mochi Methodology to Enable Dynamic HPC Data Services.....	414
<i>Matthieu Dorier, Philip Carns, Robert Ross, Shane Snyder, Rob Latham, Amal Gueroudji, George Amvrosiadis, Chuck Cranor, Jerome Soumagne</i>	
Adaptive Per-File Lossless Compression of Floating-Point Data	423
<i>Andrew Rodriguez, Noushin Azami, Martin Burtscher</i>	
Optimizing Forward Wavefield Storage Leveraging High-Speed Storage Media.....	431
<i>João Speglitch, Navjot Kukreja, George Bisbas, Atila Saraiva, Jan Hückelheim, Fabio Luporini, John Washbourne</i>	
The Art of Sparsity: Mastering High-Dimensional Tensor Storage.....	439
<i>Bin Dong, Kesheng Wu, Suren Byna</i>	
GrAPL 2024 Preface and Committees.....	447
<i>Nesreen K. Ahmed, Manoj Kumar</i>	
To Tile Or Not to Tile, that is the Question	449
<i>Altan Haan, Doru Thom Popovici, Koushik Sen, Costin Iancu, Alvin Cheung</i>	
Teaching Network Traffic Matrices in an Interactive Game Environment.....	459
<i>Chasen Milner, Hayden Jananthan, Jeremy Kepner, Vijay Gadepally, Michael Jones, Peter Michaleas, Ritesh Patel, Sandeep Pisharody, Gabriel Wachman, Alex Pentland</i>	
Characterizing the Performance of Emerging Deep Learning, Graph, and High Performance Computing Workloads Under Interference	468
<i>Hao Xu, Shuang Song, Ze Mao</i>	
The GraphBLAS 3.0 Project	478
<i>Raye Kimmerer, Timothy G. Mattson, Scott McMillan, Benjamin Brock, Erik Welch, Michel Pelletier, José E. Moreira</i>	
Edge-Parallel Graph Encoder Embedding.....	482
<i>Ariel Lubonja, Cencheng Shen, Carey Priebe, Randal Burns</i>	
Multi-Level GNN Preconditioner for Solving Large Scale Problems	486
<i>Matthieu Nastorg, Jean-Marc Gratien, Thibault Faney, Michele Alessandro Bucci, Guillaume Charpiat, Marc Schoenauer</i>	
STGraph: A Framework for Temporal Graph Neural Networks.....	496
<i>Joel Mathew Cherian, Nithin Puthalath Manoj, Kevin Jude Concessao, Unnikrishnan Cheramangalath</i>	
GraphBinMatch: Graph-Based Similarity Learning for Cross-Language Binary and Source Code Matching.....	506
<i>Ali TehraniJamsaz, Hanze Chen, Ali Jannesari</i>	
GraphBLAS.jl V0.1: An Update on GraphBLAS in Julia	516
<i>Raye Kimmerer</i>	

ECG: Expressing Locality and Prefetching for Optimal Caching in Graph Structures	520
<i>Abdullah T. Mughrabi, Morteza Baradaran, Ahmed Samara, Kevin Skadron</i>	
Unlocking the Potential: Performance Portability of Graph Algorithms on Kokkos Framework	526
<i>Shaikh Arifuzzaman, Hasan S. Arikan, M.A.M. Faysal, Maximilian Bremer, John Shalf, Doru Popovici</i>	
Shared-Memory Parallel Edmonds Blossom Algorithm for Maximum Cardinality Matching in General Graphs.....	530
<i>Gregory Schwing, Daniel Grosu, Loren Schwiebert</i>	
HiCOMB 2024 Preface and Committees	540
<i>Alba Cristina M. A. de Melo, Ananth Kalyanaraman</i>	
Re-Visiting the Third Pillar of Science for Synergistic (Bio)Computing	541
<i>Wu Feng</i>	
Lessons Learned Designing Irregular Genomic Algorithms on Parallel Systems and Architectures	542
<i>Giulia Guidi</i>	
Empirical Study of Molecular Dynamics Workflow Data Movement: DYAD Vs. Traditional I/O Systems.....	543
<i>Ian Lumsden, Hariharan Devarajan, Jack Marquez, Stephanie Brink, David Boehme, Olga Pearce, Jae-Seung Yeom, Michela Taufer</i>	
ZSMILES: An Approach for Efficient SMILES Storage for Random Access in Virtual Screening.....	544
<i>Gianmarco Accordi, Davide Gadioli, Giorgio Seguini, Andrea R. Beccari, Gianluca Palermo</i>	
Further Optimizations and Analysis of Smith-Waterman with Vector Extensions	561
<i>Reza Sajjadinasab, Hamed Rastaghi, Hafsa Shahzad, Sanjay Arora, Ulrich Drepper, Martin Herbordt</i>	
High Performance Binding Affinity Prediction with a Transformer-Based Surrogate Model.....	571
<i>Archit Vasan, Ozan Gokdemir, Alexander Brace, Arvind Ramanathan, Thomas Brettin, Rick Stevens, Venkatram Vishwanath</i>	
PAISE 2024 Preface and Committees	581
<i>Pete Beckman</i>	
FrameFeedback: A Closed-Loop Control System for Dynamic Offloading Real-Time Edge Inference.....	584
<i>Matthew Jackson, Bo Ji, Dimitrios S. Nikolopoulos</i>	
A Converting Autoencoder Toward Low-Latency and Energy-efficient DNN Inference at the Edge.....	592
<i>Hasanul Mahmud, Peng Kang, Kevin Desai, Palden Lama, Sushil K. Prasad</i>	
PCM Enabled Low-Power Photonic Accelerator for Inference and Training on Edge Devices.....	600
<i>Juliana Curry, Ahmed Louri, Avinash Karanth, Razvan Bunescu</i>	
Towards Accelerating k-NN with MPI and Near-Memory Processing	608
<i>Hooyoung Ahn, Seonyoung Kim, Yoomi Park, Woojong Han, Nick Contini, Bharath Ramesh, Mustafa Abduljabbar, Dhahaleswar K. Panda</i>	
CGRA4HPC 2024 Welcome Message and Committee List	616
<i>Artur Podobas</i>	

An Architecture-Agnostic Dataflow Mapping Framework on CGRA	618
<i>Jiangnan Li, Yazhou Yan, Jingyuan Li, Shaoyang Sun, Boyin Jin, Wenbo Yin, Lingli Wang</i>	
TransMap: An Efficient CGRA Mapping Framework Via Transformer and Deep Reinforcement Learning	626
<i>Jingyuan Li, Yuan Dai, Yihan Hu, Jiangnan Li, Wenbo Yin, Jun Tao, Lingli Wang</i>	
Comparative Analysis of Executing GPU Applications on FPGA: HLS Vs. Soft GPU Approaches	634
<i>Chihyo Ahn, Shinnung Jeong, Liam Paul Cooper, Nicholas Parnenzini, Hyesoon Kim</i>	
CGRA-ME 2.0: A Research Framework for Next-Generation CGRA Architectures and CAD.....	642
<i>Omar Ragheb, Stephen Wicklund, Matthew Walker, Rami Beidas, Adham Ragab, Tianyi Yu, Jason Anderson</i>	
A Scalable Mapping Method for Elastic CGRAs	650
<i>Makoto Saito, Takuya Kojima, Hideki Takase, Hiroshi Nakamura</i>	
GIM (Ghost in the Machine): A Coarse-Grained Reconfigurable Compute-In-Memory Platform for Exploring Machine-Learning Architectures	658
<i>Maya Borowicz, James Ding, Winnie Fan, Zhongqi Gao, Davis Jackson, Ares Lu, Sophia Rohlfsen, Ray Simar</i>	
HIPS 2024 Preface and Committees	664
<i>Seyong Lee, Lena Oden</i>	
Architecture and Programming of Analog In-Memory-Computing Accelerators for Deep Neural Networks	666
<i>Hsin Yu Sidney Tsai</i>	
ECC++ : A Compiler Construction Framework for Embedded Domain-Specific Languages	667
<i>Marc Gonzalez Tallada, Joel Denny, Pedro Valero-Lara, Seyong Lee, Keita Teranishi, Jeffrey S. Vetter</i>	
Comprehensive Study for Just-In-Time Pack Functions in Open MPI.....	678
<i>Yicheng Li, Joseph Schuchart, George Bosilca</i>	
Dynamic Resource Management for Elastic Scientific Workflows Using PMIx	686
<i>Rajat Bhattachari, Howard Pritchard, Sheikh Ghafoor</i>	
GrOUT: Transparent Scale-Out to Overcome UVM's Oversubscription Slowdowns	696
<i>Ian Di Dio Lavore, Davide Maffi, Marco Arnaboldi, Arnaud Delamare, Daniele Bonetta, Marco D. Santambrogio</i>	
Towards Fine-Grained Parallelism in Parallel and Distributed Python Libraries.....	706
<i>Jamison Kerney, Ioan Raicu, John Raicu, Kyle Chard</i>	
Automated Data Analysis for Defining Performance Metrics from Raw Hardware Events.....	716
<i>Daniel Barry, Anthony Danalis, Jack Dongarra</i>	
Performance Analysis of the NVIDIA HPC SDK and AMD AOCC Compilers in an HPC Cluster Using Pooled, Robust and Relative Metrics	726
<i>Yectli A. Huerta</i>	
9th IEEE International Workshop on Automatic Performance Tuning (iWAPT 2024)	738
<i>Valero Lara Pedro</i>	

IWAPT 2024 Keynote Talk: What Happens to a Dream Deferred? Chasing Automatic Offloading in Fortran 2023	740
<i>Damian Rouson</i>	
An Exploration of Global Optimization Strategies for Autotuning OpenMP-Based Codes.....	741
<i>Gregory Bolet, Giorgis Georgakoudis, Konstantinos Parasyris, Kirk W. Cameron, David Beckingsale, Todd Gamblin</i>	
Communication-Computation Overlapping for Parallel Multigrid Methods.....	751
<i>Kengo Nakajima</i>	
PML-MPI: A Pre-Trained ML Framework for Efficient Collective Algorithm Selection in MPI.....	761
<i>Mingzhe Han, Goutham Kalikrishna Reddy Kuncham, Ben Michalowicz, Rahul Vaidya, Mustafa Abduljabbar, Aamir Shafi, Hari Subramoni, Dhabaleswar K. Panda</i>	
Application-Agnostic Auto-Tuning of Open MPI Collectives Using Bayesian Optimization	771
<i>Emmanuel Jeannot, Pierre Lemarinier, Guillaume Mercier, Sophie Robert-Hayek, Richard Sartori</i>	
XAI-Based Feature Importance Analysis on Loop Optimization.....	782
<i>Toshinobu Katayama, Keichi Takahashi, Yoichi Shimomura, Hiroyuki Takizawa</i>	
Cost-Effective Methodology for Complex Tuning Searches in HPC: Navigating Interdependencies and Dimensionality.....	792
<i>Adrian P. Dieguez, Min Choi, Mahmut Okyay, Mauro Del Ben, Bryan M. Wong, Khaled Z. Ibrahim</i>	
27th Workshop on Job Scheduling Strategies for Parallel Processing; (JSSPP 2024).....	802
<i>Dalibor Klusáček, Julita Corbalán, Gonzalo P. Rodrigo</i>	
ParSocial 2024 Welcome and Committee List	803
<i>Hien Nguyen, Jeremy E. Thompson</i>	
RIMR: Reverse Influence Maximization Rank	805
<i>Jay Vap, Peter Kogge</i>	
Distributed Multi-GPU Community Detection on Exascale Computing Platforms	815
<i>Naw Safrin Sattar, Hao Lu, Feiyi Wang, Mahantesh Halappanavar</i>	
Lock-Free Computation of PageRank in Dynamic Graphs	825
<i>Subhajit Sahu, Kishore Kothapalli, Hemalatha Eedi, Sathya Peri</i>	
Proposal for a Flexible Benchmark for Agent Based Models.....	835
<i>Elizabeth Koning, William Gropp</i>	
Socio-Behavioral Influences in Epidemic Modeling: Towards a Unified Framework	839
<i>Suresh Subramanian, Vairavan Murugappan, Eunice E. Santos</i>	
Towards Improved Uncertainty Quantification of Stochastic Epidemic Models Using Sequential Monte Carlo.....	843
<i>Arindam Fadikar, Abby Stevens, Nicholson Collier, Kok Ben Toh, Olga Morozova, Anna Hotton, Jared Clark, David Higdon, Jonathan Ozik</i>	
Revolutionizing Personal Recommendations Via Federated Contrastive Transformer Learning.....	853
<i>Youcef Djenouri, Fabio Augusto de Alcantara Andrade, Gautam Srivastava, Ahmed Nabil Belbachir</i>	

High-Speed Transcript Collection on Multimedia Platforms: Advancing Social Media Research Through Parallel Processing.....	857
<i>Mert Can Cakmak, Nitin Agarwal</i>	
Parallelizing Accelerographic Records Processing.....	861
<i>Ronaldo Canizales, Luis Mixco, Jедидия McClurg</i>	
PDCO 2024 Preface and Committee List.....	870
<i>N/A</i>	
GPU-Accelerated Tree-Search in Chapel Versus CUDA and HIP	872
<i>Guillaume Helbecque, Ezhilmathi Krishnasamy, Nouredine Melab, Pascal Bouvry</i>	
Parallel Maximum Cardinality Matching for General Graphs on GPUs	880
<i>Gregory Schwing, Daniel Grosu, Loren Schwiebert</i>	
GPU-LSolve: An Efficient GPU-Based Laplacian Solver for Million-Scale Graphs.....	890
<i>Sumiaya Dabeer, Amitabha Bagchi, Rahul Narain</i>	
KOptim: Kubernetes Optimization Framework	900
<i>Tarek Menouer, Christophe Cérin, Patrice Darmon</i>	
Electric Drive Assignment Strategies Optimization for Plugin Hybrid Urban Buses on Tailored Emissions Mapping	909
<i>José Miguel Aragón-Jurado, Marina Díaz-Jiménez, Bernabé Dorronsoro, Pablo Pavón-Domínguez, Marcin Seredyński, Patricia Ruiz</i>	
DUST: Resource-Aware Telemetry Offloading with a Distributed Hardware-Agnostic Approach.....	919
<i>Mehrnaz Sharifian, Diman Zad Tootaghaj, Chen-Nee Chuah, Puneet Sharma</i>	
Understanding Multi-Dimensional Efficiency of Fine-Tuning Large Language Models Using SpeedUp, MemoryUp, and EnergyUp.....	929
<i>Dayuan Chen, Noe Soto, Jonas F. Tuttle, Ziliang Zong</i>	
Compiler-Driven SWAR Parallelism for High-Performance Bitboard Algorithms.....	938
<i>Florian Fey, Sergei Gorlatch</i>	
Multiobjective Based Strategy for Neural Architecture Search for Segmentation Task.....	947
<i>Abass Sana, Kaoutar Senhaji, Amir Nakib</i>	
A Mathematical Model and a Convergence Result for Totally Asynchronous Federated Learning	956
<i>Didier El-Baz, Jia Luo, Hao Mo, Lei Shi</i>	
State-Space Search to Find Energy-Aware Pareto-Efficient Optimal Task Schedules	964
<i>Yasith Udagedara, Andrea Raith, Oliver Sinnen</i>	
Message from the PDSEC-24 Workshop Chairs	974
<i>Sabine Roller, George Bosilca, Raphaël Couturier, Neda Ebrahimi Pour, Jean-Claude Charr, Thomas Rauber, Gudula Rünger, Laurence T. Yang</i>	
Multi-Criteria Mesh Partitioning for an Explicit Temporal Adaptive Task-Distributed Finite-Volume Solver	976
<i>Alice Lasserre, Jean-Marie Couteyen-Carpaye, Abdou Guermouche, Raymond Namyst</i>	
Graph Computing on Long Vector Architectures (Yes, it Works!)	986
<i>Pablo Vizcaino, Jesus Labarta, Filippo Mantovani</i>	

Integration of Modern HPC Performance Tools in Vlasiator for Exascale Analysis and Optimization	996
<i>Camille Coti, Yann Pfau-Kempf, Markus Battarbee, Urs Ganse, Sameer Shende, Kevin Huck, Jordi Rodriguez, Leo Kotipalo, Jennifer Faj, Jeremy J. Williams, Ivy Peng, Allen D. Malony, Stefano Markidis, Minna Palmroth</i>	
Performance Portability of Generated Cardiac Simulation Kernels Through Automatic Dimensioning and Load Balancing on Heterogeneous Nodes	1006
<i>Vincent Alba, Olivier Aumage, Denis Barthou, Raphaël Colin, Marie-Christine Counilh, Stéphane Genaud, Amina Guermouche, Vincent Loechner, Arun Thangamani</i>	
A Parallel Workflow for Polar Sea-Ice Classification Using Auto-Labeling of Sentinel-2 Imagery	1016
<i>Jurdana Masuma Iqrah, Wei Wang, Hongjie Xie, Sushil K. Prasad</i>	
Automated Calibration of Parallel and Distributed Computing Simulators: A Case Study	1026
<i>Jesse McDonald, Maximilian Horzela, Frédéric Suter, Henri Casanova</i>	
Pretraining Billion-Scale Geospatial Foundational Models on Frontier.....	1036
<i>Aristeidis Tsaris, Philipe Ambrozio Dias, Abhishek Potnis, Junqi Yin, Feiyi Wang, Dalton Lunga</i>	
Scaling Ensembles of Data-Intensive Quantum Chemical Calculations for Millions of Molecules	1047
<i>Kshitij Mehta, Massimiliano Lupo Pasini, Stephan Irle, Pilsun Yoo, Frédéric Suter, Dmitry Ganyushin, Scott Klasky</i>	
Accelerating Quantum Light-Matter Dynamics on Graphics Processing Units	1057
<i>Tauseq Mohammed Razakh, Thomas Linker, Ye Luo, Rajiv K. Kalia, Ken-Ichi Nomura, Priya Vashishta, Aiichiro Nakano</i>	
Q-CASA 2024 Preface and Committee List.....	1067
<i>Ashfaq Khokhar, Mary Eshaghian-Wilner, Robert Basili</i>	
Quantifying Performance of Wire-Based Quantum Circuit Cutting with Entanglements	1068
<i>Shiplu Sarker, Wenyang Qian, Soham Pal, Robert Basili, Mary Eshaghian-Wilner, Ashfaq Khokhar, Glenn Luecke, James P. Vary</i>	
Parallel Quantum Circuit Extraction from MBQC-Patterns	1078
<i>Marcel Quanz, Korbinian Staudacher, Karl Fürlinger</i>	
Hybrid Classical-Quantum Simulation of MaxCut Using QAOA-in-QAOA	1088
<i>Aniello Esposito, Tamuz Danzig</i>	
A Delay-Efficient Implementation of Quantum Carry Select Adders	1095
<i>Annalisa Massini, Federico Mingardi</i>	
Quantum Circuit Mapping Using Binary Integer Nonlinear Programming.....	1105
<i>Aaron Orenstein, Vipin Chaudhary</i>	
Measurement-Based Quantum Approximate Optimization.....	1115
<i>Tobias Stollenwerk, Stuart Hadfield</i>	
Image Compression and Reconstruction Based on Quantum Network	1128
<i>J. Xun, Qin Liu, Shan Huang, Andi Chen, W. Shengjun</i>	
Cutting a Wire with Non-Maximally Entangled States	1136
<i>Marvin Bechtold, Johanna Barzen, Frank Leymann, Alexander Mandl</i>	

Introduction to Poster-Elect Papers	1146
<i>N/A</i>	
A Deep Dive into Task-Based Parallelism in Python	1147
<i>William Ruys, Hochan Lee, Bozhi You, Shreya Talati, Jaeyoung Park, James Almgren-Bell, Yineng Yan, Milinda Fernando, George Biros, Mattan Erez, Martin Burtscher, Christopher J. Rossbach, Keshav Pingali, Milos Gligoric</i>	
A New Exact State Reconstruction Strategy for Conjugate Gradient Methods with Arbitrary Preconditioners	1150
<i>Viktoria Mayer, Wilfried N. Gansterer</i>	
A Stochastic Composite Model to Understand the Impact of Rare, Colossal Interference in HPC Systems.....	1153
<i>Muna Tageldin, Majeed M. Hayat, Jered Dominguez-Trujillo, Patrick G. Bridges</i>	
Accelerating Native Transaction Processing in LSM-Based Persistent Key-Value Stores.....	1156
<i>Jin Xue, Zili Shao</i>	
AdCoalescer: An Adaptive Coalescer to Reduce the Inter-Module Traffic in MCM-GPUs	1159
<i>Xu Zhang, Guangda Zhang, Lu Wang, Xia Zhao</i>	
An SR-IOV SSD Optimized for QoS-Sensitive IaaS Cloud Storage	1161
<i>Xiang Chen, Ru Ying, Haocong Ma, Yao Wang, Xianjun Meng, Guangjun Xie, Yonghui Zhan, Fenyong Yuan, Ying Yang, Ying Yang, Tao Lu, Jinqiang Wang, You Zhou, Fei Wu</i>	
Asynchrony and Failure Masking Via Pseudo-Local Process Recovery in MPI Applications.....	1164
<i>Mathew Whitlock, Hemanth Kolla, Aurelien Bouteiller, Jackson R. Mayo, Nicolas M. Morales, Keita Teranishi, George Bosilca</i>	
EDDIS: Accelerating Distributed Data-Parallel DNN Training for Heterogeneous GPU Cluster	1167
<i>Shinyoung Ahn, Hooyoung Ahn, Hyeonseong Choi, Jaehyun Lee</i>	
Efficient Multi-Processor Scheduling in Increasingly Realistic Models (Brief Summary).....	1169
<i>Pál András Papp, Georg Anegg, Aikaterini Karanasiou, Albert-Jan N. Yzelman</i>	
Energy-Aware Decentralized Learning with Intermittent Model Training	1172
<i>Martijn De Vos, Akash Dhasade, Paolo Dini, Elia Guerra, Anne-Marie Kermarrec, Marco Miozzo, Rafael Pires, Rishi Sharma</i>	
Enhancing Energy Efficiency with Multi-Site Scheduling Strategies	1175
<i>Alok Kamatar, Valerie Hayot-Sasson, Yadu Babuji, Andre Bauer, Gourav Rattihalli, Ninad Hogade, Dejan Milojevic, Kyle Chard, Ian Foster</i>	
Evaluation of Programming Models and Performance for Stencil Computation on GPGPUs.....	1178
<i>Baodi Shan, Mauricio Araya-Polo</i>	
Exploiting Tensor Cores in Sparse Matrix-Multivector Multiplication Via Block-Sparsity-Aware Clustering	1181
<i>Eunji Lee, Yoonsang Han, Gordon Euhyun Moon</i>	
FedClust: Optimizing Federated Learning on Non-IID Data Through Weight-Driven Client Clustering	1184
<i>Md Sirajul Islam, Simin Javaherian, Fei Xu, Xu Yuan, Li Chen, Nian-Feng Tzeng</i>	

FedSZ: Leveraging Error-Bounded Lossy Compression for Federated Learning Communications	1187
<i>Grant Wilkins, Sheng Di, Jon C. Calhoun, Zilinghan Li, Kibaek Kim, Robert Underwood, Richard Mortier, Franck Cappello</i>	
Integration Framework for Online Thread Throttling with Thread and Page Mapping on NUMA Systems.....	1189
<i>Janaina Schwarzrock, Arthur F. Lorenzon, Samuel Xavier de Souza, Antonio Carlos S. Beck</i>	
MDLoader: A Hybrid Model-Driven Data Loader for Distributed Deep Neural Networks Training.....	1193
<i>Jonghyun Bae, Jong Youl Choi, Massimiliano Lupo Pasini, Kshitij Mehta, Khaled Z. Ibrahim</i>	
Proactive, Accuracy-Aware Straggler Mitigation in Machine Learning Clusters.....	1196
<i>Suraiya Tairin, Haiying Shen, Anand Iyer</i>	
Scalable Node Embedding Algorithms Using Distributed Sparse Matrix Operations.....	1199
<i>Isuru Ranawaka, Ariful Azad</i>	
Scheduling and Allocation of Disaggregated Memory Resources in HPC Systems.....	1202
<i>Jie Li, George Michelogiannakis, Brandon Cook, John Shalf, Yong Chen</i>	
Shared-Memory Parallel Dynamic Louvain Algorithm for Community Detection.....	1204
<i>Subhajit Sahu, Kishore Kothapalli, Dip Sankar Banerjee</i>	
System Optimizations for Enabling Training of Extreme Long Sequence Transformer Models	1206
<i>Sam Ade Jacobs, Masahiro Tanaka, Chengming Zhang, Minjia Zhang, Reza Yazdani Aminabadi, Shuaiwen Leon Song, Samyam Rajbhandari, Yuxiong He</i>	
Toward Self-Adjusting k-Ary Search Tree Networks.....	1209
<i>Evgeniy Feder, Anton Paramonov, Pavel Mavrin, Iosif Salem, Stefan Schmid, Vitaly Aksenov</i>	
Understanding Different Transport Coexistence in Datacenter Networks.....	1212
<i>Dinghuang Hu, Dezun Dong</i>	
IPDPS 2024 PhD Forum	1214
<i>Sanmukh Kuppannagari, Tanwi Mallick</i>	

Author Index