

2025 IEEE High Performance Extreme Computing Conference (HPEC 2025)

**Wakefield, Massachusetts, USA
15-19 September 2025**



**IEEE Catalog Number: CFP25HPE-POD
ISBN: 979-8-3315-7845-9**

**Copyright © 2025 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:	CFP25HPE-POD
ISBN (Print-On-Demand):	979-8-3315-7845-9
ISBN (Online):	979-8-3315-7844-2
ISSN:	2377-6943

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

CURRAN ASSOCIATES INC.
proceedings
.com

TABLE OF CONTENTS

Degree Matrix Comparison for Graph Alignment.....	1
<i>Ashley Wang, Peter Chin</i>	
Evaluation of Habitat Robotics Using Large Language Models	8
<i>William Li, Lei Hamilton, Kaise Al-Natour, Sanjeev Mohindra</i>	
Strategic Cyber Defense Via Reinforcement Learning-Guided Combinatorial Auctions.....	16
<i>Mai Pham, Vikrant Vaze, Peter Chin</i>	
InfraredGP: Efficient Graph Partitioning Via Spectral Graph Neural Networks with Negative Corrections	23
<i>Meng Qin, Weihua Li, Jinqiang Cui, Sen Pei</i>	
AGCRS: An Adaptive Generalized Storage Scheme for Large Sparse Tensors	30
<i>Md Mehrab Hossain Opi, K. M. Azharul Hasan</i>	
Partition-Based Surface Code Compilation.....	36
<i>Hanjing Xu, Xiaoyuan Liu, Ankit Kulshrestha, Hayato Ushijima-Mwesigwa</i>	
Encoded Time-Series Model Training with U-Net Running on Wafer Scale Engine	44
<i>Vyacheslav Romanov</i>	
Enhancing Graph Partitioning with Reinforcement Learning-Based Initialization	49
<i>Chedi Morchdi, Cheng-Hsiang Chiu, Wan Luan Lee, Tsung-Wei Huang, Yi Zhou</i>	
Easy Acceleration with Distributed Arrays.....	57
<i>Jeremy Kepner, Chansup Byun, Latoya Anderson, William Arcand, David Bestor, William Bergeron, Alex Bonn, Daniel Burrill, Vijay Gadepally, Ryan Haney, Michael Houle, Matthew Hubbell, Hayden Jananthan, Michael Jones, Piotr Luszczek, Lauren Milechin, Guillermo Morales, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Charles Yee, Peter Michaleas</i>	
Towards a Framework for Etymological Analysis of Literary and AI Works.....	65
<i>Mike Hm Teodorescu, Cecilia Speranta Bolea, Horia-Nicolai L. Teodorescu</i>	
QAOA Parameter Transferability for Maximum Independent Set Using Graph Attention Networks.....	71
<i>Hanjing Xu, Xiaoyuan Liu, Alex Pothén, Ilya Safro</i>	
Scaling Triangle Counting and K-Truss on the UpDown Architecture	78
<i>Jiya Su, Alexander Fell, Andronicus Rajasukumar, David F. Gleich, Andrew A. Chien</i>	
SANST: Sensing Anonymized Network Via Sorted Triplets.....	85
<i>Jianyu Wang, Wenzhi Tang, Chenglong Shi, Zhe Zhang, Dan Chen, Miaojiang Chen, Wenjing Xiao</i>	
Comparative Analysis of RISC-V Softcore and Hardcore Processors for Space Computing.....	92
<i>Ni Nyoman Dhinar Gayatri, Alan D. George</i>	
Accelerating Supercomputing: AI-Hardware-Driven Innovation for Speed and Efficiency	99
<i>Jack Dongarra, John Gunnels, Harun Bayraktar, Azzam Haidar, Dan Ernst</i>	

Combining Performance and Productivity: Accelerating the Network Sensing Graph Challenge with GPUs and Commodity Data Science Software	106
<i>Siddharth Samsi, Dan Campbell, Emanuel Scoullos, Oded Green</i>	
Predicting HPC Job Run Time with Realistic Data Using Application Input Parameters	113
<i>Kenneth Lamar, Benjamin A. Allan, M. Scot Swan, James M. Brandt, Damian Dechev</i>	
Enhancing Sentiment Classification of E-Commerce Reviews for Actionable Insights Using LLMs	122
<i>Kevin Power, Peter Harding, Jose Lopez, Alex Carroll, Elenna Dugundji</i>	
UniPar: A Unified LLM-Based Framework for Parallel and Accelerated Code Translation in HPC	129
<i>Tomer Bitan, Tal Kadosh, Erel Kaplan, Shira Meiri, Le Chen, Peter Morales, Niranjana Hasabnis, Gal Oren</i>	
Performance and Numerical Aspects of Decompositional Factorizations with FP64 Floating-Point Emulation in INT8.....	138
<i>Piotr Luszczek, Vijay Gadepally, Latoya Anderson, William Arcand, David Bestor, William Bergeron, Alex Bonn, Daniel J. Burrill, Chansup Byun, Michael Houle, Matthew Hubbell, Michael Jones, Peter Michaleas, Guillermo Morales, Julia Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Charles Yee, Jeremy Kepner</i>	
Lincoln AI Computing Survey (LAICS) and Trends.....	145
<i>Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Jeremy Kepner</i>	
Sustainably Modeling a Sustainable Future Climate	157
<i>Rabab Alomairy, Sameh Abdulah, Qinglei Cao, Marc G. Genton, David E. Keyes, Hatem Ltaief</i>	
Scalable Graph Algorithms on Distributed UpDown Accelerators	165
<i>Brian Wheatman, Andrew A. Chien</i>	
SARComp: High-Performance Algorithms for SAR Processing from FFT Kernels to Matched Filtering	173
<i>Maron Schlemmon, Martin Schulz, Rolf Scheiber</i>	
Detecting Expert-Written Comments in Stack Exchange.....	180
<i>Himani Musku, Alea Ritchie, Nour Jedidi, Rohan Leekha, Courtland Vandam</i>	
AOCL-CST - A New CPU Stress Test Library for AMD CPUs.....	186
<i>S. Biplab Raut</i>	
Evaluating AMD-Xilinx Frameworks for Deep-Learning Acceleration on Versal.....	190
<i>Peter Drum, Alan D. George</i>	
Hardware-Accelerated Transformer Framework for Real-Time Battery SoH Estimation.....	198
<i>Talha Coskun, Hiruna Vishwamith, Murat Isik, I. Can Dikmen</i>	
Exploring Neuromorphic Computing with Loihi-2 for High-Performance CFD Simulations	204
<i>Talha Coskun, Hiruna Vishwamith, Murat Isik, Jonathan Naoukin, I. Can Dikmen</i>	
Scalable Bayesian Nonparametric Ensemble (BNE) for Spatio-Temporal Air Pollution Predictions Using High-Performance Computing	211
<i>Vijay Kumar, Jaime Benavides, Carlos Carrillo-Gallegos, Gil Speyer, Marianthi-Anna Kioumourtzoglou</i>	

GCN-Driven CUDA Parameter Optimization for Parallel Triangle Counting in Graphs.....	218
<i>Hasan Serdar Arikan, Rakibul Hassan, Shubhashish Kar, Doru Thom Popovici, Shaikh Arifuzzaman</i>	
Performance Evaluation of LAPACK Using SVE Optimized BLAS Kernels	225
<i>Aniket P. Garade, Sushil Pratap Singh, Vishal Rayala, Deepika H V, Haribabu P, S A Kumar, S D Sudarsan</i>	
A Variable-Precision Implementation of the ADER-DG Algorithm.....	232
<i>Marc Marot-Lassauzaie, Michael Bader</i>	
Multi-Stage Stochastic Programming for Heavy-Duty Electric Truck Routing Under Public Charging Congestion Uncertainty	239
<i>Ziyan Li, Nikolay Aristov, Antoine Germain, Elenka R. Dugundji</i>	
GraphBLAS Mathematical Opportunities: Parallel Hypersparse, Matrix Based Graph Streaming, and Complex-Index Matrices	245
<i>Hayden Jananathan, Jeremy Kepner, Michael Jones, Vijay Gadepally, Michael Houle, Peter Michaleas, Chasen Milner, Alex Pentland</i>	
LLMs in Crisis Triage: Benchmarking Zero-Shot Classification of Social Media.....	256
<i>Emma L. McDaniel, Samuel Scheele, Jeffrey Liu</i>	
Accelerating Temporal Triangle Counting and Betweenness Centrality on GPUs.....	266
<i>Tuteja Trimansingh Parvindersingh, Venkata Kalyan Tavva, Subhasis Banerjee, Chiranjib Sur</i>	
Analysis and Optimization of Spiking Neural Network Simulations on GPUs.....	273
<i>Dogu Kocatepe, Isil Öz</i>	
Weighted Histogram Matching for Improved Automated Synthetic Aperture Radar (SAR)-Optical Image Registration.....	280
<i>Kerri Prinos, Aditi Mungale, Karen Gettings, Tair Akhmejanov, Devanshu Mehta</i>	
Prism: Practical In-Memory Acceleration for Subgraph Matching at Scale.....	287
<i>Deteng Chen, Yu Huang, Yi Huang, Binbin Lin, Yi Zhang, Long Zheng, Xiaofei Liao, Hai Jin</i>	
A Qualifiable GPU Sharing Approach for AI Workloads in Critical Systems.....	294
<i>Marc Solé I Bonet, Jannis Wolf, Aridane Álvarez Suárez, Leonidas Kosmidis</i>	
Predicting LLM Inference Server Request Capacity	301
<i>Daniel J. Burrill, Latoya Anderson, William Arcand, David Bestor, William Bergeron, Alex Bonn, Chansup Byun, Michael Houle, Matthew Hubbell, Michael Jones, Piotr Luszczek, Peter Michaleas, Guillermo Morales, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Charles Yee, Vijay Gadepally</i>	
Adaptive Spectral Block Floating Point for Discontinuous Galerkin Methods.....	307
<i>Shivam Sundriyal, Markus Büttner, Christoph Alt, Tobias Kenter, Vadym Aizinger</i>	
A Parallel Push-Relabel Maximum Flow Algorithm in LAGraph and GraphBLAS.....	315
<i>Darin A. Peries, Timothy A. Davis</i>	
Fast FPGA-Based Implementation of the QP-Dyn Stream Cipher Using High-Level Synthesis.....	321
<i>Paolo Palazzari, Luigi Accardi, Antonio Mastrandrea, Pasquale Tommasino, Alessandro Trifiletti</i>	
Julia GraphBLAS with Nonblocking Execution	328
<i>Pascal Costanza, Timothy G. Mattson, Raye Kimmerer, Benjamin Brock</i>	

Differentiable Graph Centrality	335
<i>Georgios Kollias, Vassilis Kalantzis</i>	
On the Landscape of Scientific Computing Libraries in Python	342
<i>Niteya Shah, Pi-Yueh Chuang, Paul Sathre, Wu-Chun Feng</i>	
Load Imbalance in HPC Applications: Improved Profiling and New Ways to Use Wasted Cycles	351
<i>Shining Yang, Xiteng Yao, Grace Nansamba, Amr Akmal Abouelmagd, Anthony Skjellum, Martin Herbordt</i>	
Accelerating Multi-Party Computation Using Heterogeneous Systems	360
<i>Xiteng Yao, Shining Yang, Mayank Varia, Martin Herbordt</i>	
HiPerMotif: Novel Parallel Subgraph Isomorphism in Large-Scale Property Graphs	370
<i>Mohammad Dindoost, Oliver Alvarado Rodriguez, Bartosz Bryg, Ioannis Koutis, David A. Bader</i>	
Advancing AI Challenges for the United States Department of the Air Force*	377
<i>Christian Prothmann, Vijay Gadepally, Jeremy Kepner, Koley Borchard, Luca Carlone, Zachary Folcik, J. Daniel Griffith, Michael Houle, Jonathan P. How, Nathan Hughes, Ifueko Igbinedion, Hayden Jananathan, Tejas Jayashankar, Michael Jones, Sertac Karaman, Binoy G. Kurien, Alejandro Lancha, Giovanni Lavezzi, Gary C. F. Lee, Charles E. Leiserson, Richard Linares, Lindsey McEvoy, Peter Michaleas, Chasen Milner, Alex Pentland, Yury Polyanskiy, Jovan Popovich, Jeffrey Price, Tim W. Reid, Stephanie Riley, Siddharth Samsi, Peter Saunders, Olga Simek, Mark S. Veillette, Amir Weiss, Gregory W. Wornell, Daniela Rus, Scott T. Ruppel</i>	
Stable Iterative Solvers for Ill-Conditioned Linear Systems and Least Squares	385
<i>Vasileios Kalantzis, Mark S. Squillante, Chai Wah Wu</i>	
Towards -O _{mL} : A Deep Learning Based Approach for Outperforming Compiler Defaults	392
<i>Hafsah Shahzad, Ahmed Sanaullah, Sanjay Arora, Ulrich Drepper, Martin Herbordt</i>	
Generalized Methodology for Determining Numerical Features of Hardware Floating-Point Matrix Multipliers: Part I	400
<i>Faizan A. Khattak, Mantas Mikaitis</i>	
A Hybrid Classical-Quantum Model for QSAR-Based Biodegradability Prediction	406
<i>Batuhan Hangun, Oguz Altun, Onder Eyecioglu</i>	
A Scalable Quantum Dynamical Approach for Calculating Collisional Molecular Properties	412
<i>Prajwal Niraula, Laurent Wiesenfeld, Julien De Wit, Iouli Gordon, Robert Hargreaves, Jeremy Kepner, Deborah Woods, Cooper Loughlin</i>	
Optimizing Local Computation in Secure Matrix Multiplication for Outsourced Neural Networks	416
<i>J. Parker Diamond, Andrea Lin, R. Nicholas Cunningham, Soamar Homs, John Darby Mitchell, Asemit Pandey, Emily Shen</i>	
Towards Automated Reasoning Chains for Verification of LLM-Generated Scientific Code	423
<i>Quentin Oschatz, Naifeng Zhang, Mike Franusich, Franz Franchetti</i>	
The NorthPole Validator: A Cycle-Accurate Simulator for HW/SW Codesign of a Prescheduled Neural Inference Accelerator	430
<i>Alexander Andreopoulos, Michael V. Debole, Jeffrey A. Kunitz, Nathaniel J. McClatchey, Tapan K. Nayak, Daniel F. Smith, Brian Taba, Filipp Akopyan, Rathinakumar Appuswamy, John V. Arthur, Andrew S. Cassidy, Pallab Datta, Carlos Ortega Otero, William P. Risk, Jun Sawada, Myron D. Flickner, Dharmendra S. Modha</i>	

Automating Harmonized System (HS) Code Classification from Unstructured Shipping Manifests Using Large Language Models.....	441
<i>Thomas Koch, Kevin Power</i>	
Anonymized Network Sensing Using C++26 Std::execution on GPUs.....	446
<i>Michael Mandulak, Sayan Ghosh, S M Ferdous, Mahantesh Halappanavar, George Slota</i>	
When Structure is Silent: Opportunities for Algorithmic Dispatch in Linear Algebra	453
<i>Emmanuel Lujan, Alan Edelman</i>	
Characterization of Sparsity-Aware Parallelization of Jaccard Similarity in Graph Datasets	460
<i>Atharva Gondhalekar, Paul Sathre, Wu-Chun Feng</i>	
A Locality Sensitive Hashing Based Algorithm to Accelerate Neighborhood Search in Graph Neural Operators	467
<i>Mariam Hassan, Sanmukh Kuppannagari</i>	
Interactive Trillion Packet Anonymized Network Analysis with the GraphBLAS.....	474
<i>Chasen Milner, Michael Houle, Hayden Jananathan, Michael Jones, Jeremy Kepner, Peter Michaleas, Inna Voloshchuk, Alex Pentland</i>	
AI-Powered Network Energy Optimization Machine Learning Approaches to Reducing Network Power Consumption	480
<i>Het Mehta, Sambu Patach Arrojula</i>	
2D Distributed Label Propagation on 400 GPUs.....	487
<i>George M. Slota, Michael Mandulak, Ujwal Pandey, Anthony Fabius</i>	
CRAMP: Categorizing Classifiers and Regressors for Scalable Parallelism on Distributed and Multicore Systems.....	495
<i>Baidya Nath Saha, Pavan Sarvaiya, Wali Mohammad Abdullah, Md. Morshedul Islam</i>	
Neuro-Inspired Enhancing Spiking Graph Convolutional Networks	501
<i>Fernando Vera Buschmann, Horacio Rotstein, Vincent Oria</i>	
Performance Analysis of Inline Compression in pySDC.....	508
<i>Emily Lattanzio, Sansriti Ranjan, Robert Underwood, Thomas Saupe, Robert Speck, Jon C. Calhoun</i>	
NTT-SAA: Exploring NTT Acceleration with 2-D Systolic Array Architecture on FPGAs	516
<i>Ashwajit Singh, Zhihan Xu, Viktor K. Prasanna</i>	
SYCL for Performance Portability: Application Experience with Coupled Cluster Formalism in Quantum Chemistry on Exascale Systems	523
<i>Abhishek Bagusetty, Ajay Panyala, Álvaro Vázquez-Mayagoitia, John K. Holmen, Kevin Harms</i>	
Accelerating Dynamic Image Graph Construction on FPGA for Vision GNNs.....	531
<i>Anvitha Ramachandran, Dhruv Parikh, Viktor Prasanna</i>	
BUG: Balanced DFS-Based Subgraph Matching with a ReUse Strategy on GPUs.....	538
<i>Zicang Xu, Lei Zou</i>	
Physics-Informed Neural Networks for Low-Power Real-Time Edge Biosensing Applications	545
<i>Soheli Farhana</i>	

Adaptive Policy Synchronization for Scalable Reinforcement Learning	551
<i>Rodney Lafuente-Mercado</i>	
Implementation of Tensor Network Simulation TN-Sim Under NWQ-Sim	557
<i>Aaron C. Hoyt, Jonathan S. Bersson, Sean Garner, Chenxu Liu, Ang Li</i>	
Geans: A GPU-Accelerated Framework for Efficient End-To-End Anonymized Network Sensing	564
<i>Jun Mai, Qinggang Wang, Yu Huang, Pengcheng Yao, Long Zheng, Xiaofei Liao, Hai Jin</i>	
System-Level Performance Modeling of Photonic In-Memory Computing.....	571
<i>Jebacyril Arockiaraj, Sasindu Wijeratne, Sugeet Sunder, Md Abdullah-Al Kaiser, Akhilesh Jaiswal, Ajey P. Jacob, Viktor Prasanna</i>	
Improving Statistical Characterization of Data Tensors with the Generalized Canonical Polyadic Tensor Decomposition.....	578
<i>Matthew D. Merris, Tim Andersen</i>	
Towards an Algorithm-Based Approach for Soft Error Tolerance Using Interval Arithmetic.....	585
<i>Larry Tang, Varun Kumar, Matthew Ngaw, Siddharth Singh, Devdutt Nadkarni, Lohith Tummala, Ken Mai, Franz Franchetti</i>	
Scaling Regime-Aware Forecasting: Distributed Shifting Seasonal Matrix Factorization	592
<i>Jacob Munson, Breschine Cummins</i>	
SCoDa: Scalable Community Detection in Data Streams	599
<i>Akanksha Dwivedi, Prashant Srivastav, Dip Sankar Banerjee</i>	
Performance-Energy Characterization of ML Inference on Heterogeneous Edge AI Platforms	606
<i>Palash Kohli, Rakshith Jayanth, Neelesh Gupta, Haoyang Fan, Viktor Prasanna</i>	
Leveraging Caliper and Benchmark to Analyze MPI Communication Patterns: Insights from AMG2023, Kripke, and Laghos	613
<i>Grace Nansamba, Evelyn Namugwanya, David Boehme, Dewi Yokelson, Riley Shipley, Derek Schafer, Michael McKinsey, Olga Pearce, Anthony Skjellum</i>	
A Fully Adaptive Radau Method for the Efficient Solution of Stiff Ordinary Differential Equations at Low Tolerances.....	623
<i>Shreyas Ekanathan, Oscar Smith, Christopher Rackauckas</i>	
Secure Virtual Network Embedding Through Fully Homomorphic Encryption	632
<i>David Bruce Cousins, Carlo Pascoe, Erik Kline</i>	
Accelerating Push-Relabel Algorithm on GPU Via Two-Level Parallelism Paradigm and Efficient CSR Designs.....	639
<i>Chou-Ying Hsieh, Po-Chieh Lin, Sy-Yen Kuo</i>	
Employing High-Performance PETSc Network Simulation for Business Profit Analysis	646
<i>Abu Asaduzzaman, Nowshin Nawal</i>	
DNN-Driven Task Scheduling for High Performance Edge-Cloud Heterogeneous Systems.....	653
<i>Md Raihan Uddin, Abu Asaduzzaman, Fairuz Nawar, Christian Thompson</i>	
Performance Modeling of Heterogeneous Edge-Cloud Systems with Machine Learning	660
<i>Md Raihan Uddin, Abu Asaduzzaman, Sonu Gangadhar Gowda</i>	

Investigating the Impact of Algorithms and Hardware on Machine Learning Models in HPC Systems.....	668
<i>Christian C. Thompson, Abu Asaduzzaman, Md Raihan Uddin, Fairuz Nawar</i>	
BanglaDocAtlas: A Multi-Class Annotated Dataset for Complex Bangla Document Layout Analysis	674
<i>Md Safayat Hossain, Jannatul Ferdous, Md Raihan Uddin, K. M. A. Hossain, M. I. Ahmed, M. A. Rahman, Asif Shahriyar Sushmit, Farig Sadeque, Swakkhar Shatabda, Abu Asaduzzaman</i>	
Evaluating Efficiency and Novelty of LLM-Generated Code for Graph Analysis.....	681
<i>Atieh Barati Nia, Mohammad Dindoost, David A. Bader</i>	
Optimizing Sparse Matrix-Vector Multiplication on GPUs Using the Mathematics of Arrays.....	688
<i>Stephen Thomas, Lenore Mullin</i>	
Power Iteration with Probabilistic Updates for Systems with Heterogeneous Performance	695
<i>Soumyadip Ghosh, Lior Horesh, Vassilis Kalantzis, Georgios Kollias, Yingdong Lu, Tomasz Nowicki, Shashanka Ubaru</i>	
CTAM Tool for Hyperscaler Qualification.....	702
<i>Tommy Yan, Rajat Madhusudan, Anna Mary Mathew, Vani Pulendra</i>	
Accelerating Sparse Deep Learning Via Multi-Layer Tensor Reordering and Partitioning.....	708
<i>Gunduz Vehbi Demirci, Cagatay Dikici, Tim Atherton</i>	
Artificial Intelligence Performance and Radiation Effects in Neuromorphic NorthPole Hardware.....	716
<i>Victor M. Vergara, Francisco O. Viramontes, Amanda E. Romero, Matthew E. Spear, Evan T. Kain, Windy S. Slater, Heather M. Quinn, Qing Wu</i>	
Benchmarking Deep Learning with Representative ONNX Subgraphs.....	722
<i>Marika E. Schubert, David Langerman, Evan W. Gretok, Ian Peitzsch, Calvin B. Gealy, Jefferson Boothe, Alan D. George</i>	
P4OMP: Retrieval-Augmented Prompting for OpenMP Parallelism in Serial Code	729
<i>Wali Mohammad Abdullah, Azmain Kabir</i>	
DBOS Network Sensing: A Web Services Approach to Collaborative Awareness	735
<i>Sophia Lockton, Jeremy Kepner, Michael Stonebraker, Hayden Jananathan, Latoya Anderson, William Arcand, David Bestor, William Bergeron, Alex Bonn, Daniel Burrill, Chansup Byun, Timothy Davis, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Piotr Luszczek, Peter Michaleas, Lauren Milechin, Chasen Milner, Guillermo Morales, Julie Mullen, Michel Pelletier, Alex Poliakov, Andrew Prout, Albert Reuther, Antonio Rosa, Charles Yee, Alex Pentland</i>	
TRACER Software Switching at Waveform Timescales.....	743
<i>Connor Imes, Dong In D. Kang, Matthew French, John Paul Walters</i>	
RAILS: Retrieval-Augmented Intelligence for Learning Software Development	750
<i>Wali Mohammad Abdullah, Md. Morshedul Islam, Devraj Parmar, Happy Hasmukhbhai Patel, Sindhuja Prabhakaran, Baidya Saha</i>	
Balancing Performance and Productivity: A Comparative Study of Apache Arrow Vs. MPI	756
<i>Ritvik Prabhu, Wu-Chun Feng</i>	
A Framework for the Iterative Solution of Sparse Linear Systems on Hybrid Architectures Using Homomorphic Encryption.....	763
<i>Lior Horesh, Vassilis Kalantzis, Barry M. Trager, Shashanka Ubaru</i>	

An Improved Izhikevich Neuron Model with Time Delay and Adaptive Parameters for High-Performance Spiking Neural Networks	771
<i>Alissa Kane, Felipe Marcelino, Anton Spirkin, Yuchou Chang</i>	
Web-Based Intelligent Decision Support System for Real-Time Toll Plaza Management and AI-Driven Operational Optimization	777
<i>Pattarapon Klaykul, Wilaipom Lee, Kanabadee Srisomboon, Luepol Pipanmekaporn, Akara Prayote</i>	
Enhancing the Real-Time Solutions of Parametric Linear Systems on a GPU Through Hybrid Coarse-Grained Transprecision Computing	785
<i>Hamid Noori, Hans Vandierendonck, Roger Woods, Nick Polydorides</i>	
AIMS: An Adaptive Intelligent Multi-Objective Scheduler Powered by Digital Twins.....	792
<i>Kyrian Adimora, Hongyang Sun</i>	
Performance Analysis of the Parallel Shared-Memory Sparse Matrix-Vector Multiplication on Unstructured Matrices	800
<i>Kobe Bergmans, Karl Meerbergen, Raf Vandebril</i>	
A Time-Aware Sliding Window-Based Hotel Recommendation Framework Using Multi-Stage BERT-MRC	807
<i>Md. Nazirul Hasan Shawon, K. M. Azharul Hasan</i>	
On the Adaptation of Mixed-Radix Fast Fourier Transform for Resource-Constrained Environments	813
<i>Atharva Gondhalekar, Paul Sathre, Wu-Chun Feng</i>	
Improving Uncertainty Based Dataset Pruning with Density Estimation for Noisy Edge Environments.....	821
<i>Mike Soricelli, Yuchou Chang, Christopher J Hixenbaugh</i>	
GPU-Accelerated, Mixed Precision GMRES(m) with Varied Restarts	826
<i>Abir Haque, Suzanne M. Shontz, Xuemin Tu</i>	
A Scalable Code Generation Flow for Heterogeneous Parallel RTL Simulation Using MLIR	833
<i>Jie Tong, Zhengxiong Li, Umit Yusuf Ogras, Tsung-Wei Huang</i>	
pdGRASS: A Fast Parallel Density-Aware Algorithm for Graph Spectral Sparsification	842
<i>Tiancheng Zhao, Zekun Yin, Huihai An, Xiaoyu Yang, Zhou Jin, Jiasi Shen, Helen Xu</i>	
Predicting Ports Congestion by Utilizing State Space Models	849
<i>Nikolay Aristov, Elenka R. Dugundji</i>	
Mitigation of Applied Load Using Machine Learning for Adaptive Motor Control	856
<i>Kourosh Rahnamai, Jacob Rollins, Luke Moisan, Ryan Kayfus</i>	
Designing Parallel Algorithms for Community Detection Using Arachne.....	861
<i>Fuhuan Li, Zhihui Du, David A. Bader</i>	
Towards Efficient Sparse Deep Neural Network Inference Via Multi-Level Concurrency Orchestration	868
<i>Ming Dun, Jie Zhou, Huawei Cao, Shuhan Song, Yiming Sun, Mingyu Yan, Xiaochun Ye</i>	

Author Index