

2023 IEEE High Performance Extreme Computing Conference (HPEC 2023)

**Boston, Massachusetts, USA
25-29 September 2023**



**IEEE Catalog Number: CFP23HPE-POD
ISBN: 979-8-3503-0861-7**

**Copyright © 2023 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:	CFP23HPE-POD
ISBN (Print-On-Demand):	979-8-3503-0861-7
ISBN (Online):	979-8-3503-0860-0
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

Benchmarking Deep Learning Classifiers for SAR Automatic Target Recognition.....	1
<i>Jacob Fein-Ashley, Tian Ye, Rajgopal Kannan, Viktor Prasanna, Carl Busart</i>	
Parallel Longest Common SubSequence Analysis in Chapel.....	7
<i>Soroush Vahidi, Baruch Schieber, Zhihui Du, David A. Bader</i>	
Pruning Binarized Neural Networks Enables Low-Latency, Low-Power FPGA-Based Handwritten Digit Classification.....	13
<i>Syamantak Payra, Gabriel Loke, Yoel Fink, Joseph D. Steinmeyer</i>	
High-Level Framework for Solving Systems of the PDEs on Distributed Systems.....	21
<i>Yevhen Pankevych, Oleg Farenjuk</i>	
Opportunistic Query Execution on SmartNICs for Analyzing In-Transit Data.....	26
<i>Jianshen Liu, Carlos Maltzahn, Craig Ulmer</i>	
Modeling and Analyzing Wind Velocity at Entrance Doors to Avoid Accidents.....	33
<i>Abu Asaduzzaman, Luke Mercer, Md Raihan Uddin, Yoel Woldeyes</i>	
USAP: An Ultra-Fast Stochastic Graph Partitioner.....	38
<i>Chih-Chun Chang, Tsung-Wei Huang</i>	
In-Place Multicore SIMD Fast Fourier Transforms.....	45
<i>Benoît Dupont De Dinechin, Julien Hascoët, Orège Desrentes</i>	
SMOG: Accelerating Subgraph Matching on GPUs.....	51
<i>Zhibin Wang, Ziheng Meng, Xue Li, Xi Lin, Long Zheng, Chen Tian, Sheng Zhong</i>	
pPython Performance Study.....	58
<i>Chansup Byun, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Hayden Jananathan, Michael Jones, Anna Klein, Peter Michaleas, Lauren Milechin, Guillermo Morales, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Charles Yee, Jeremy Kepner</i>	
A Massively Parallel BWP Algorithm for Solving Large-Scale Systems of Nonlinear Equations.....	65
<i>Bruno Silva, Luiz Guerreiro Lopes</i>	
Solving Sparse Linear Systems Via Flexible GMRES with In-Memory Analog Preconditioning.....	71
<i>Vasileios Kalantzis, Mark S. Squillante, Chai Wah Wu, Anshul Gupta, Shashanka Ubaru, Tayfun Gokmen, Lior Horesh</i>	
A Look into a GraphBLAS Entry Point into an LLVM Lowering Pass, with a Precision Formatting Example.....	78
<i>Roy Gulla</i>	
Exploring Challenges Associated with Employing SmartNICs as General-Purpose HPC Accelerators.....	82
<i>Brody Williams, Yong Chen, Wendy Poole, Steve Poole</i>	
Automatic Differentiation for Inverse Problems with Applications in Quantum Transport.....	89
<i>Ivan Williams, Eric Polizzi</i>	

Towards a Flexible Hardware Implementation for Mixed-Radix Fourier Transforms	94
<i>Mario Vega, Xiaokun Yang, John Shalf, Doru Thom Popovici</i>	
Hardware Root-Of-Trust Support for Operational Technology Cybersecurity in Critical Infrastructures.....	101
<i>Alan Ehret, Peter Moore, Milan Stojkov, Michel A. Kinsky</i>	
FAST-CON: A Multi-source Approach for Efficient S- T Connectivity on Sparse Graphs	108
<i>Leonardo Fraccaroli, Rosalba Giugno, Samuele Cancellieri, Federico Busato, Nicola Bombieri</i>	
Mapping of Internet “Coastlines” Via Large Scale Anonymized Network Source Correlations.....	114
<i>Hayden Jananathan, Jeremy Kepner, Michael Jones, William Arcand, David Bestor, William Bergeron, Chansup Byun, Timothy Davis, Vijay Gadepally, Daniel Grant, Michael Houle, Matthew Hubbell, Anna Klein, Lauren Milechin, Guillermo Morales, Andrew Morris, Julie Mullen, Ritesh Patel, Alex Pentl, Sandeep Pisharody, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Tyler Trigg, Gabriel Wachman, Charles Yee, Peter Michaleas</i>	
A Composable Just-In-Time Programming Framework with LLMs and FBP	123
<i>Andy Vidan, Lars Fiedler</i>	
Decomposition Based Refinement for the Network Interdiction Problem	131
<i>Krish Matta, Xiaoyuan Liu, Ilya Safro</i>	
Tyche: A Compact and Configurable Accelerator for Scalable Probabilistic Computing on FPGA	139
<i>Yashash Jain, Utsav Banerjee</i>	
Improved Models for Policy-Agent Learning of Compiler Directives in HLS	146
<i>Robert Munafo, Hafsa Shahzad, Ahmed Sanaullah, Sanjay Arora, Uli Drepper, Martin Herboldt</i>	
Fast Spectral Graph Partitioning with a Randomized Eigensolver.....	154
<i>Heliezer J. D. Espinoza, Jennifer A. Loe, Erik G. Boman</i>	
From Words to Watts: Benchmarking the Energy Costs of Large Language Model Inference	161
<i>Siddharth Samsi, Dan Zhao, Joseph McDonald, Baolin Li, Adam Michaleas, Michael Jones, William Bergeron, Jeremy Kepner, Devesh Tiwari, Vijay Gadepally</i>	
RaftGP: Random Fast Graph Partitioning	170
<i>Yu Gao, Meng Qin, Yibin Ding, Li Zeng, Chaorui Zhang, Weixi Zhang, Wei Han, Rongqian Zhao, Bo Bai</i>	
Parallel Algorithms for Computing Jaccard Weights on Graphs Using Linear Algebra.....	177
<i>Elaheh Hassani, Md Taufique Hussain, Ariful Azad</i>	
Optimizing Compression Schemes for Parallel Sparse Tensor Algebra	184
<i>Helen Xu, Tao B. Schardl, Michael Pellauer, Joel S. Emer</i>	
ANEDA: Adaptable Node Embeddings for Shortest Path Distance Approximation.....	191
<i>Frank Pacini, Allison Gunby-Mann, Sarel Cohen, Peter Chin</i>	
Automated Indexing of TEM Diffraction Patterns Using Machine Learning	198
<i>Nathaniel Tomczak, Sanmukh Kuppanagari</i>	
A Framework for Analyzing the Robustness of Graph Models.....	205
<i>Khaled Abdelaal, Richard Veras</i>	

Hybrid Quantum-Classical Multilevel Approach for Maximum Cuts on Graphs	211
<i>Anthony Angone, Xiaoyuan Liu, Ruslan Shaydulin, Ilya Safro</i>	
Accelerating Training Data Generation Using Optimal Parallelization and Thread Counts	218
<i>Jonathan Levine, Leonard Maceachern</i>	
IRIS-DMEM: Efficient Memory Management for Heterogeneous Computing	225
<i>Narasinga Rao Miniskar, Mohammad Alaul Haque Monil, Pedro Valero-Lara, Frank Y. Liu, Jeffrey S. Vetter</i>	
Addressing Endpoint-Induced Congestion for Accelerator Scale-Out in a Medium-Scale Domain	232
<i>Timothy Chong, Venkata Krishnan</i>	
Leveraging Mixed Precision in Exponential Time Integration Methods	240
<i>Cody J. Balos, Steven Roberts, David J. Gardner</i>	
MAVR: Multi-Functional Point Cloud Annotations Using Virtual Reality	248
<i>Xiao Zhang, Zhanhong Huang, Xinming Huang</i>	
Decontentioned Stochastic Block Partition	254
<i>Ahsen J. Uppal, Thomas B. Rolinger, H. Howie Huang</i>	
Creating a Dataset for High-Performance Computing Code Translation Using LLMs: A Bridge Between OpenMP Fortran and C++	260
<i>Bin Lei, Caiwen Ding, Le Chen, Pei-Hung Lin, Chunhua Liao</i>	
Optimizing a Distributed Graph Data Structure for K - Path Centrality Estimation on HPC.....	267
<i>Lance Fletcher, Trevor Steil, Roger Pearce</i>	
Quantifying OpenMP: Statistical Insights into Usage and Adoption	274
<i>Tal Kadosh, Niranjana Hasabnis, Timothy Mattson, Yuval Pinter, Gal Oren</i>	
Accelerating Multi-Agent DDPG on CPU-FPGA Heterogeneous Platform	281
<i>Samuel Wiggins, Yuan Meng, Rajgopal Kannan, Viktor Prasanna</i>	
On the Three P's of Parallel Programming for Heterogeneous Computing: Performance, Productivity, and Portability	288
<i>Atharva Gondhalekar, Wu-Chun Feng</i>	
Robust Fine-Tuning of Vision-Language Models for Domain Generalization.....	295
<i>Kevin Vogt-Lowell, Noah Lee, Theodoros Tsiligkaridis, Marc Vaillant</i>	
TenSQL: An SQL Database Built on GraphBLAS	302
<i>Jon Roose, Miheer Vaidya, Ponnuswamy Sadayappan, Sivasankaran Rajamanickam</i>	
Towards the FAIR Asset Tracking Across Models, Datasets, and Performance Evaluation Scenarios	310
<i>Piotr Luszczek, Tokey Tahmid</i>	
ProtoX: A First Look	316
<i>Het Mankad, Sanil Rao, Phillip Colella, Brian Van Straalen, Franz Franchetti</i>	
Performance of Graph Neural Networks for Point Cloud Applications	322
<i>Dhruv Parikh, Bingyi Zhang, Rajgopal Kannan, Viktor Prasanna, Carl Busart</i>	
PaCKD: Pattern-Clustered Knowledge Distillation for Compressing Memory Access Prediction Models.....	329
<i>Neelesh Gupta, Pengmiao Zhang, Rajgopal Kannan, Viktor Prasanna</i>	

G-MAP: A Graph Neural Network-Based Framework for Memory Access Prediction.....	336
<i>Abhiram Rao Gorle, Pengmiao Zhang, Rajgopal Kannan, Viktor K. Prasanna</i>	
Optimization and Performance Analysis of Shor's Algorithm in Qiskit	343
<i>Dewang Sun, Naifeng Zhang, Franz Franchetti</i>	
Errant Beam Detection Using the AMD Versal ACAP and Vitis AI.....	350
<i>Anthony M Cabrera, Yigit A Yucesan, Frank Y Liu, Willem Blokland, Jeffrey S Vetter</i>	
Generating High-Performance Number Theoretic Transform Implementations for Vector Architectures	356
<i>Naifeng Zhang, Austin Ebel, Negar Neda, Patrick Brinich, Benedict Reynwar, Andrew G. Schmidt, Mike Franusich, Jeremy Johnson, Brandon Reagen, Franz Franchetti</i>	
Quantifying the Gap Between Open-Source and Vendor FPGA Place and Route Tools.....	363
<i>Shachi Khadilkar, Ahmed Sanaullah, Martin Margala</i>	
A Holistic Optimisation - Success Mantra for HPC Performance.....	369
<i>Ashish Bisht, Deepika H. V, Haribabu P, S A Kumar, S D Sudarsan</i>	
Exploiting Fusion Opportunities in Linear Algebraic Graph Query Engines	375
<i>Yuttapichai Kerdcharoen, Upasana Sridhar, Tze Meng Low</i>	
Scalable Deep Learning for Pilot Performance Analysis Using Multimodal Physiological Time Series	382
<i>Noah Lee, Patrick W. Moore, Laura J. Brattain</i>	
Lincoln AI Computing Survey (LAICS) Update.....	388
<i>Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, Jeremy Kepner</i>	
Parallel Quasi-Concave Set Function Optimization for Scalability Even Without Submodularity	395
<i>Praneeth Vepakomma, Yulia Kempner, Rodmy Paredes Alfaro, Ramesh Raskar</i>	
Energy Estimates Across Layers of Computing: from Devices to Large-Scale Applications in Machine Learning for Natural Language Processing, Scientific Computing, and Cryptocurrency Mining I	403
<i>Sadasivan Shankar</i>	
Multiarchitecture Hardware Acceleration of Hyperdimensional Computing	409
<i>Ian Peitzsch, Mark Ciora, Alan D. George</i>	
Selective Encryption of Compressed Image Regions on the Edge with Reconfigurable Hardware.....	416
<i>Justin Kawakami, Dominik Zajac, Miriam Leeser</i>	
Meta-Learning and Self-Supervised Pretraining for Storm Event Imagery Translation.....	422
<i>Ileana Rugina, Rumun Dangovski, Mark Veillette, Pooya Khorrami, Brian Cheung, Olga Simek, Marin Soljacic</i>	
Contextualizing Enhances Gradient Based Meta Learning for Few Shot Image Classification.....	431
<i>Evan Vogelbaum, Rumun Dangovski, Li Jing, Marin Soljacic</i>	
UNet Performance with Wafer Scale Engine (Optimization Case Study).....	444
<i>Vyacheslav Romanov</i>	

Zero Trust Architecture Approach for Developing Mission Critical Embedded Systems	450
<i>Michael Vai, David Whelihan, Eric Simpson, Donato Kava, Alice Lee, Huy Nguyen, Jeffrey Hughes, Gabriel Torres, Jeffery Lim, Ben Nahill, Roger Khazan, Fred Schneider</i>	
Deployment of Real-Time Network Traffic Analysis Using GraphBLAS Hypersparse Matrices and D4M Associative Arrays.....	455
<i>Michael Jones, Jeremy Kepner, Andrew Prout, Timothy Davis, William Arcand, David Bestor, William Bergeron, Chansup Byun, Vijay Gadepally, Michael Houle, Matthew Hubbell, Hayden Jananthan, Anna Klein, Lauren Milechin, Guillermo Morales, Julie Mullen, Ritesh Patel, Sandeep Pisharody, Albert Reuther, Antonio Rosa, Siddharth Samsi, Charles Yee, Peter Michaleas</i>	
Machine Learning Across Network-Connected FPGAs	463
<i>Dana Diaconu, Yanyue Xie, Mehmet Gungor, Suranga Handagala, Xue Lin, Miriam Leeser</i>	
Triangle Counting Through Cover-Edges.....	470
<i>David A. Bader, Fuhuan Li, Anya Ganeshan, Ahmet Gundogdu, Jason Lew, Oliver Alvarado Rodriguez, Zhihui Du</i>	
Accelerating Garbled Circuits in the Open Cloud Testbed with Multiple Network-Attached FPGAs.....	477
<i>Kai Huang, Mehmet Gungor, Suranga Handagala, Stratis Ioannidis, Miriam Leeser</i>	
Scalable and Portable Pipelines for Predicting 3D Protein Structures on Standalone and HPC Systems.....	485
<i>Adam Michaleas, Darrell O. Rieke</i>	
AOCL-Compression - A High Performance Optimized Lossless Data Compression Library	489
<i>S. Biplab Raut</i>	
Decreasing the Computing Time of Bayesian Optimization Using Generalizable Memory Pruning.....	496
<i>Alexander E. Siemenn, Tonio Buonassisi</i>	
Adaptive Sparse Deep Neural Network Inference on Resource-Constrained Cost-Efficient GPUs.....	503
<i>Ming Dun, Xu Zhang, Huawei Cao, Yuan Zhang, Junying Huang, Xiaochun Ye</i>	
Asymmetric Grouped Convolutions for Logarithmic Scale Efficient Convolutional Neural Networks	510
<i>Li Jing, Rumen Dangovski, Marin Soljacic</i>	
Manifold Transfer Networks for Lens Distortion Rectification	522
<i>Li Jing, Lay Jain, Rumen Dangovski, Marin Soljacic</i>	
Fast Triangle Counting	530
<i>David A. Bader</i>	
Parallel Clustering with Resolution Variation	536
<i>Nikos Pitsianis, Dimitris Floros, Tiancheng Liu, Xiaobai Sun</i>	
High-Level Frameworks: Effect on Transformer Inference Time and Power on Embedded GPU Devices	544
<i>Marika E. Schubert, David Langerman, Alan D. George</i>	
Image Segmentation with Topological Priors.....	552
<i>Shakir Showkat Sofi, Nadezhda Alshanova</i>	

Photonic Accelerators for Image Segmentation in Autonomous Driving and Defect Detection	558
<i>Lakshmi Nair, David Widemann, Brad Turcott, Nick Moore, Alexandra Wleklinski, Darius Bunandar, Ioannis Papavasileiou, Shihu Wang, Eric Logan</i>	
The Aggressive Oversubscribing Scheduling for Interactive Jobs on a Supercomputing System.....	567
<i>Shohei Minami, Toshio Endo, Akihiro Nomura</i>	
Dynamic Data Partitioning in the WAFL File System.....	574
<i>Jian Hu, Matthew Curtis-Maury, Vinay Devadas</i>	
Machine Learning at the Edge Using Neural Network Processor	581
<i>Edwin Lee, Michael Parker, Michael Cervantes, Ben Plotner</i>	
Accelerating GNN-Based SAR Automatic Target Recognition on HBM-Enabled FPGA	585
<i>Bingyi Zhang, Rajgopal Kannan, Viktor Prasanna, Carl Busart</i>	
Leveraging Mathworks Tools to Accelerate the Prototyping of Custom 5G Applications in Hardware	592
<i>Joshua Geyster, Karen Gettings, Paul Monticciolo, Matthew Rebholz</i>	
A GPU Parallel Algorithm for Finding a Negative Subset Disjoint Cycle in a Graph	598
<i>Piotr Sielski, Akif Çördük, Hugo Linsenmaier, Alex Fender</i>	
An Analysis of Accelerator Data-Transfer Modes in NoC-Based SoC Architectures	605
<i>Kuan-Lin Chiu, Davide Giri, Luca Piccolboni, Luca P. Carloni</i>	
Continuous Deep Equilibrium Models: Training Neural ODEs Faster by Integrating Them to Infinity.....	612
<i>Avik Pal, Alan Edelman, Chris Rackauckas</i>	
Finding Your Niche: An Evolutionary Approach to HPC Topologies	621
<i>Stephen J. Young, Joshua Suetterlein, Jesun Firoz, Joseph Manzano, Kevin Barker</i>	
Build Energy-Efficient GPU Computing Environment for Machine Learning Algorithms with Register File Packing Technique	630
<i>Xin Wang, Wei Zhang</i>	
GLARE: Accelerating Sparse DNN Inference Kernels with Global Memory Access Reduction	637
<i>Shui Jiang, Tsung-Wei Huang, Tsung-Yi Ho</i>	
Property Graphs in Arachne	644
<i>Oliver Alvarado Rodriguez, Fernando Vera Buschmann, Zhihui Du, David A. Bader</i>	
Feature-Oriented FSMs for FPGAs.....	651
<i>Justin Deters, Peyton Gozon, Max Camp-Oberhauser, Ron K. Cytron</i>	
Performance Analysis of Graph Neural Network (GNN) for Manufacturing Feature Recognition Problem	658
<i>Igor Betkier, Mateusz Oszczypala, Janusz Pobożniak, Sergiusz Sobieski</i>	
Focusing and Calibration of Large Scale Network Sensors Using GraphBLAS Anonymized Hypersparse Matrices	664
<i>Jeremy Kepner, Michael Jones, Phil Dykstra, Chansup Byun, Timothy Davis, Hayden Jananathan, William Arcand, David Bestor, William Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Anna Klein, Lauren Milechin, Guillermo Morales, Julie Mullen, Ritesh Patel, Alex Pentland, Sandeep Pisharody, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Tyler Trigg, Charles Yee, Peter Michaleas</i>	

An Analysis of Energy Requirement for Computer Vision Algorithms.....	673
<i>Daniel Edelman, Siddharth Samsi, Joseph McDonald, Adam Michaleas, Vijay Gadepally</i>	
An Integrated Approach for Accelerating Stochastic Block Partitioning	680
<i>Frank Wanye, Vitaliy Gleyzer, Edward Kao, Wu-Chun Feng</i>	
Advanced Ultra Low-Power Deep Learning Applications with Neuromorphic Computing	687
<i>Mark Barnell, Courtney Raymond, Lisa Loomis, Darrek Isereau, Daniel Brown, Francesca Vidal, Steven Smiley</i>	
Multi-Sweep-Line Algorithm for Rectangle Union on GPU and Its Application for VLSI Density Calculation.....	691
<i>Chang-Hung Wu, Che-Rung Lee</i>	
Acceleration of Synthetic Aperture Radar for On-Board Space Systems.....	698
<i>Marc Solé, Iván Rodríguez-Ferrandez, David Steenari, Leonidas Kosmidis</i>	

Author Index