

2022 IEEE High Performance Extreme Computing Conference (HPEC 2022)

**Virtual Conference
19 – 23 September 2022**



**IEEE Catalog Number: CFP22HPE-POD
ISBN: 978-1-6654-9787-9**

**Copyright © 2022 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:	CFP22HPE-POD
ISBN (Print-On-Demand):	978-1-6654-9787-9
ISBN (Online):	978-1-6654-9786-2
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

Performance Estimation for Efficient Image Segmentation Training of Weather Radar Algorithms	1
<i>Joseph McDonald, James M. Kurdzo, Phillip M. Stepanian, Mark Veillette, David Bestor, Michael Jones, Vijay Gadepally, Siddharth Samsi</i>	
Site-Wide HPC Data Center Demand Response	8
<i>Daniel C. Wilson, Ioannis Ch. Paschalidis, Ayse K. Coskun</i>	
Systolic Array Based FPGA Accelerator for Yolov3-Tiny	15
<i>Prithvi Velicheti, Sivani Pentapati, Suresh Purini</i>	
pPython for Parallel Python Programming.....	17
<i>Chansup Byun, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Hayden Jananathan, Michael Jones, Kurt Keville, Anna Klein, Peter Michaleas, Lauren Milechin, Guillermo Morales, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Charles Yee, Jeremy Kepner</i>	
A High Throughput Hardware Accelerator for FFTW Codelets: A First Look	24
<i>Larry Tang, Siyuan Chen, Keshav Harisrikanth, Guanglin Xu, Ken Mai, Franz Franchetti</i>	
Computing In-Place FFTs with SIMD Lane Slicing	31
<i>Benoît Dupont De Dinechin</i>	
FAST: A Scalable Subgraph Matching Framework Over Large Graphs	38
<i>Jiezhong He, Zhouyang Liu, Yixin Chen, Hengyue Pan, Zhen Huang, Dongsheng Li</i>	
Ultra Low-Power Deep Learning Applications at the Edge with Jetson Orin AGX Hardware.....	45
<i>Mark Barnell, Courtney Raymond, Steven Smiley, Darrek Isereau, Daniel Brown</i>	
Design and Implementation of a Real-Time Parallel FFT for a Direction-Finding System on an FPGA.....	49
<i>Bheema Lakshmi Pradeep, Rishu Anand, Pavan Vadakattu, Syed Azemuddin, Aquibuddin Ahmed</i>	
An Evaluation of Low Overhead Time Series Preprocessing Techniques for Downstream Machine Learning	56
<i>Matthew L. Weiss, Joseph McDonald, David Bestor, Charles Yee, Daniel Edelman, Michael Jones, Andrew Prout, Andrew Bowne, Lindsey McEvoy, Vijay Gadepally, Siddharth Samsi</i>	
Resource-Constrained Optimizations for Synthetic Aperture Radar On-Board Image Processing	61
<i>Maron Schlemmon, Martin Schulz, Rolf Scheiber</i>	
Large Scale Enrichment and Statistical Cyber Characterization of Network Traffic	69
<i>Ivan Kawaminami, Arminda Estrada, Youssef Elsakkary, Hayden Jananathan, Aydin Buluç, Tim Davis, Daniel Jones, Chad Meiners, Andrew Morris, Sandeep Pisharody, Jeremy Kepner</i>	
An SSD-Based Accelerator for Singular Value Decomposition Recommendation Algorithm on Edge.....	77
<i>Wei Wu, Letian Zhao, Qizhe Wu, Xiaotian Wang, Teng Tian, Xi Jin</i>	

Hypersparse Network Flow Analysis of Packets with GraphBLAS.....	82
<i>Tyler Trigg, Chad Meiners, Sandeep Pisharody, Hayden Jananathan, Michael Jones, Adam Michaleas, Timothy Davis, Erik Welch, William Arcand, David Bestor, William Bergeron, Chansup Byun, Vijay Gadepally, Micheal Houle, Matthew Hubbell, Anna Klein, Peter Michaleas, Lauren Milechin, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Doug Stetson, Charles Yee, Jeremy Kepner</i>	
Powering Practical Performance: Accelerated Numerical Computing in Pure Python	89
<i>Matthew Penn, Chris Milroy</i>	
Efficient Calculation of Triangle Centrality in Big Data Networks	94
<i>Wali Mohammad Abdullah, David Awosoga, Shahadat Hossain</i>	
The Viability of Using Online Prediction to Perform Extra Work While Executing BSP Applications.....	101
<i>Po Hao Chen, Pouya Haghi, Jae Yoon Chung, Tong Geng, Richard West, Anthony Skjellum, Martin C. Herbordt</i>	
Machine Learning for Accurate and Fast Bandgap Prediction of Solid-State Materials	108
<i>Shomik Verma, Shivam Kajale, Rafael Gomez-Bombarelli</i>	
Optimal GPU Frequency Selection Using Multi-Objective Approaches for HPC Systems.....	110
<i>Ghazanfar Ali, Sridutt Bhalachandra, Nicholas J. Wright, Mert Side, Yong Chen</i>	
Parallelizing Explicit and Implicit Extrapolation Methods for Ordinary Differential Equations	117
<i>Utkarsh, Chris Elrod, Yingbo Ma, Konstantin Althaus, Christopher Rackauckas</i>	
Modeling the Energy Efficiency of GEMM Using Optical Random Access Memory	126
<i>Bingyi Zhang, Akhilesh Jaiswal, Clynn Mathew, Ravi Teja Lakkireddy, Ajey P. Jacob, Sasindu Wijeratne, Viktor Prasanna</i>	
Constructing Optimal Contraction Trees for Tensor Network Quantum Circuit Simulation	133
<i>Cameron Ibrahim, Danylo Lykov, Zichang He, Yuri Alexeev, Ilya Safro</i>	
Achieving Speedups for Distributed Graph Biconnectivity	141
<i>Ian Bogle, George M. Slota</i>	
Surrogate ML/AI Model Benchmarking for FAIR Principles' Conformance.....	148
<i>Piotr Luszczek, Cade Brown</i>	
RaiderSTREAM: Adapting the STREAM Benchmark to Modern HPC Systems	153
<i>Michael Beebe, Brody Williams, Stephen Devaney, John Leidel, Yong Chen, Steve Poole</i>	
Unsupervised Adaptation of Spiking Networks in a Gradual Changing Environment.....	160
<i>Zaidao Mei, Mark Barnell, Qinru Qiu</i>	
Distributed Out-of-Memory SVD on CPU/GPU Architectures	167
<i>Ismael Boureima, Manish Bhattarai, Maksim E. Eren, Nick Solovyev, Hristo Djidjev, Boian S. Alexandrov</i>	
Kalman Filter Driven Estimation of Community Structure in Time Varying Graphs.....	175
<i>Lisa J. K. Durbeck, Peter Athanas</i>	
Demystifying the Nvidia Ampere Architecture Through Microbenchmarking and Instruction-Level Analysis.....	182
<i>Hamdy Abdelkhalik, Yehia Arafa, Nandakishore Santhi, Abdel-Hameed A. Badawy</i>	

Optimizing Performance and Storage of Memory-Mapped Persistent Data Structures	190
<i>Karim Youssef, Abdullah Al Raqibul Islam, Keita Iwabuchi, Wu-Chun Feng, Roger Pearce</i>	
Quantum Netlist Compiler (QNC)	197
<i>Shamminuj Aktar, Abdel-Hameed A. Badawy, Nandakishore Santhi</i>	
Exploring the Impacts of Software Cache Configuration for In-Line Compressed Arrays.....	204
<i>Sansriti Ranjan, Dakota Fulp, Jon C. Calhoun</i>	
On the Characterization of the Performance-Productivity Gap for FPGA	211
<i>Atharva Gondhalekar, Thomas Twomey, Wu-Chun Feng</i>	
Edge-Connected Jaccard Similarity for Graph Link Prediction on FPGA	219
<i>Paul Sathre, Atharva Gondhalekar, Wu-Chun Feng</i>	
Improved Distributed-Memory Triangle Counting by Exploiting the Graph Structure	229
<i>Sayan Ghosh</i>	
AutoPager: Auto-Tuning Memory-Mapped I/O Parameters in Userspace.....	235
<i>Karim Youssef, Niteya Shah, Maya Gokhale, Roger Pearce, Wu-Chun Feng</i>	
Optimizing Open-Source FPGA CAD Tools	242
<i>Shachi Khadilkar, Martin Margala</i>	
Sparse Deep Neural Network Inference Using Different Programming Models	246
<i>Hyungro Lee, Milan Jain, Sayan Ghosh</i>	
Predicting Ankle Moment Trajectory with Adaptive Weighted Ensemble of LSTM Networks.....	252
<i>Emilia Grzesiak, Jennifer Sloboda, Ho Chit Siu</i>	
Task-Parallel Programming with Constrained Parallelism.....	259
<i>Tsung-Wei Huang, Leslie Hwang</i>	
Enhancing the Performance Portability of Heterogeneous Circuit Analysis Programs.....	266
<i>Tsung-Wei Huang</i>	
Proposed Empirical Assessment of Remote Workers' Cyberslacking and Computer Security Posture to Assess Organizational Cybersecurity Risks	268
<i>Ariel Luna, Yair Levy, Gregory Simco, Wei Li</i>	
Benchmarking Resource Usage for Efficient Distributed Deep Learning.....	270
<i>Nathan C. Frey, Baolin Li, Joseph McDonald, Dan Zhao, Michael Jones, David Bestor, Devesh Tiwari, Vijay Gadepally, Siddharth Samsi</i>	
Enabling Transformers to Understand Low-Level Programs.....	278
<i>Zifan Carl Guo, William S. Moses</i>	
GPU-Accelerated High-Bandwidth Radar Centroiding.....	287
<i>David J. Brigada, Maximilian Merfeld, Kara Warner</i>	
Towards a Generic UVM	293
<i>Kholoud Mahmoud, Randa Ahmed, Karim Ayman, Mostafa Aymau, Waleed Taie, Yasser Ibrahim, Hassan Mostafa, Khaled Salah</i>	
Generating Permutations Using Hash Tables	299
<i>Oded Green, Corey Nolet, Joe Eaton</i>	

<i>HashTag: Fast Lookup in a Persistent Memory File System</i>	306
<i>Matthew Curtis-Maury, Yash Trivedi</i>	
Online Detection and Classification of State Transitions of Multivariate Shock and Vibration Data.....	313
<i>Nicklaus Przybylski, William M. Jones, Nathan Debardeleben</i>	
C2QA - Bosonic Qiskit	320
<i>Timothy J Stavenger, Eleanor Crane, Kevin C Smith, Christopher T Kang, Steven M Girvin, Nathan Wiebe</i>	
Parallel Computing with DNA Forensics Data.....	328
<i>Adam Michaleas, Philip Fremont-Smith, Chelsea Lennartz, Darrell O. Ricke</i>	
Evaluation of a Novel Scratchpad Memory Through Compiler Supported Simulation	335
<i>Essa Imhmed, Jonathan Cook, Abdel-Hameed Badawy</i>	
Real-Time Software Architecture for EM-Based Radar Signal Processing and Tracking	342
<i>Alan Nussbaum, Byron Keel, William Dale Blair, Umakishore Ramachandran</i>	
Deep Gaussian Process with Multitask and Transfer Learning for Performance Optimization	349
<i>Wissam M. Sid-Lakhdar, Mohsen Aznaveh, Piotr Luszczek, Jack Dongarra</i>	
Towards Hardware Accelerated Garbage Collection with Near-Memory Processing	356
<i>Samuel Thomas, Jiwon Choe, Ofir Gordon, Erez Petrank, Tali Moreshet, Maurice Herlihy, R. Iris Bahar</i>	
HuGraph: Acceleration of GCN Training on Heterogeneous FPGA Clusters with Quantization.....	362
<i>Letian Zhao, Qizhe Wu, Xiaotian Wang, Teng Tian, Wei Wu, Xi Jin</i>	
Optimizing Designs Using Several Types of Memories on Modern FPGAs	369
<i>Mehmet Gungor, Kai Huang, Stratis Ioannidis, Miriam Leeser</i>	
GraphBLAS on the Edge: Anonymized High Performance Streaming of Network Traffic.....	376
<i>Michael Jones, Jeremy Kepner, Daniel Andersen, Aydin Buluç, Chansup Byun, K Claffy, Timothy Davis, William Arcand, Jonathan Bernays, David Bestor, William Bergeron, Vijay Gadepally, Micheal Houle, Matthew Hubbell, Hayden Jananthan, Anna Klein, Chad Meiners, Lauren Milechin, Julie Mullen, Sandeep Pisharody, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Jon Sreekanth, Doug Stetson, Charles Yee, Peter Michaleas</i>	
SHARP: Software Hint-Assisted Memory Access Prediction for Graph Analytics.....	384
<i>Pengmiao Zhang, Rajgopal Kannan, Xiangzhi Tong, Anant V. Nori, Viktor K. Prasanna</i>	
Hardware Software Codesign of Applications on the Edge: Accelerating Digital PreDistortion for Wireless Communications.....	392
<i>Zhaoyang Han, Yiyue Jiang, Rahul Mushini, John Dooley, Miriam Leeser</i>	
Scalable Interactive Autonomous Navigation Simulations on HPC.....	398
<i>Wesley Brewer, Joel Bretheim, John Kaniarz, Peilin Song, Burhman Gates</i>	
Performance Modeling Sparse MTTKRP Using Optical Static Random Access Memory on FPGA.....	405
<i>Sasindu Wijeratne, Akhilesh Jaiswal, Ajey P. Jacob, Bingyi Zhang, Viktor Prasanna</i>	
Flexible Hardware Accelerator Design Generation with Spiral	412
<i>Guanglin Xu, James C. Hoe, Franz Franchetti</i>	

A High-Performance Deployment Framework for Pipelined CNN Accelerators with Flexible DSE Strategy.....	419
<i>Conghui Luo, Wenjin Huang, Dehao Xiang, Yihua Huang</i>	
A Multi-GPU Parallel Genetic Algorithm for Large-Scale Vehicle Routing Problems	427
<i>Marwan Abdelatti, Manbir Sodhi, Resit Sendag</i>	
Explicit Ordering Refinement for Accelerating Irregular Graph Analysis	435
<i>Michael Mandulak, Ruo Chen Hu, George Slota</i>	
Kv2vec: A Distributed Representation Method for Key-Value Pairs from Metadata Attributes	443
<i>Chenxu Niu, Wei Zhang, Suren Byna, Yong Chen</i>	
DASH: Scheduling Deep Learning Workloads on Multi-Generational GPU-Accelerated Clusters	450
<i>Baolin Li, Tirthak Patel, Vijay Gadepally, Karen Gettings, Siddharth Samsi, Devesh Tiwari</i>	
AI and ML Accelerator Survey and Trends	457
<i>Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, Jeremy Kepner</i>	
Hardware Design and Implementation of Classic McEliece Post-Quantum Cryptosystem Based on FPGA.....	467
<i>Shaofen Chen, Haiyan Lin, Wenjin Huang, Yihua Huang</i>	
HTC: Hybrid Vertex-Parallel and Edge-Parallel Triangle Counting.....	473
<i>Li Zeng, Kang Yang, Haoran Cai, Jinhua Zhou, Rongqian Zhao, Xin Chen</i>	
Hardware Design and Implementation of Post-Quantum Cryptography Kyber	480
<i>Qingru Zeng, Quanxin Li, Baoze Zhao, Han Jiao, Yihua Huang</i>	
Processing Particle Data Flows with SmartNICs	486
<i>Jianshen Liu, Carlos Maltzahn, Matthew L. Curry, Craig Ulmer</i>	
Accelerating Sparse Deep Neural Network Inference Using GPU Tensor Cores	494
<i>Yufei Sun, Long Zheng, Qinggang Wang, Xiangyu Ye, Yu Huang, Pengcheng Yao, Xiaofei Liao, Hai Jin</i>	
Enabling Novel In-Memory Computation Algorithms to Address Next-Generation Throughput Constraints on SWaP-Limited Platforms.....	501
<i>Jessica Ray, Chad R. Meiners</i>	
Towards Fast GPU-Based Sparse DNN Inference: A Hybrid Compute Model	508
<i>Shaoxian Xu, Minkang Wu, Long Zheng, Zhiyuan Shao, Xiangyu Ye, Xiaofei Liao, Hai Jin</i>	
Apple Silicon Performance in Scientific Computing	515
<i>Connor Kenyon, Collin Capano</i>	
Fast Graph Algorithms for Superpixel Segmentation.....	525
<i>Dimitris Floros, Tiancheng Liu, Nikos Pitsianis, Xiaobai Sun</i>	
How to Prevent a Sick ASIC	533
<i>William F. Ellersick</i>	
Performance Speedup of Quantum Espresso Using Optimized AOCL-FFTW.....	539
<i>S. Biplab Raut</i>	

Distributed Hardware Accelerated Secure Joint Computation on the COPA Framework.....	543
<i>Rushi Patel, Pouya Haghi, Shweta Jain, Andriy Kot, Venkata Krishnan, Mayrmk Varia, Martin Herbordt</i>	
Trends in Energy Estimates for Computing in AI/Machine Learning Accelerators, Supercomputers, and Compute-Intensive Applications	550
<i>Sadasivan Shankar, Albert Reuther</i>	
Challenges Designing for FPGAs Using High-Level Synthesis.....	558
<i>Clayton J. Faber, Steven D. Harris, Zhili Xiac, Roger D. Chamberlain, Anthony M. Cabrera</i>	
Analyzing Multi-Trillion Edge Graphs on Large GPU Clusters: A Case Study with PageRank	565
<i>Seunghwa Kang, Joseph Nke, Brad Rees</i>	
Towards Fast Crash-Consistent Cluster Checkpointing	572
<i>Andrew Wood, Moshik Hershcovitch, Ilias Ennmouri, Weiyu Zong, Saurav Chennuri, Sarel Cohen, Swaminathan Sundararaman, Daniel Waddington, Peter Chin</i>	
A Scalable Inference Pipeline for 3D Axon Tracing Algorithms	580
<i>Benjamin Fenelon, Lars A. Gjestebj, Webster Guan, Juhyuk Park, Kwanghun Chung, Laura J. Brattain</i>	
FPGA Acceleration of Fully Homomorphic Encryption Over the Torus	586
<i>Tian Ye, Rajgopal Kannan, Viktor K. Prasanna</i>	
Im2win: Memory Efficient Convolution on SIMD Architectures.....	593
<i>Shuai Lu, Jun Chu, Xu T. Liu</i>	
Python Implementation of the Dynamic Distributed Dimensional Data Model.....	600
<i>Hayden Jananthan, Lauren Milechin, Michael Jones, William Arcand, William Bergeron, David Bestor, Chansup Byun, Michael Houle, Matthew Hubbell, Vijay Gadepally, Anna Klein, Peter Michaleas, Guillermo Morales, Julie Mullen, Andrew Prout, Albert Reuther, Antonio Rosa, Siddharth Samsi, Charles Yee, Jeremy Kepner</i>	
A Hierarchical Jacobi Iteration for Structured Matrices on GPUs Using Shared Memory	608
<i>Mohammad Shafaet Islam, Qiqi Wang</i>	

Author Index