

2024 IEEE High Performance Extreme Computing Conference (HPEC 2024)

**Wakefield, Massachusetts, USA
23-27 September 2024**



**IEEE Catalog Number: CFP24HPE-POD
ISBN: 979-8-3503-8714-8**

**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:	CFP24HPE-POD
ISBN (Print-On-Demand):	979-8-3503-8714-8
ISBN (Online):	979-8-3503-8713-1
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

Improving Regression in Spiking Neural Networks for Oceanographic Data Analysis	1
<i>Alissa J. Kane, Yuchou Chang</i>	
Cycle-Stealing in Load-Imbalanced HPC Applications	6
<i>Po H. Chen, Akshaya Bali, Shining Yang, Pouya Haghi, Carlton Knox, Benjamin Li, Amr A. Abouelmagd, Anthony Skjellum, Martin Herbordt</i>	
Batch-Parallel Compressed Sparse Row: A Locality-Optimized Dynamic-Graph Representation	14
<i>Brian Wheatman, Randal Burns, Helen Xu</i>	
VF2-PS: Parallel and Scalable Subgraph Monomorphism in Arachne.....	22
<i>Mohammad Dindoost, Oliver A. Rodriguez, Sounak Bagchi, Palina Pauliuchenka, Zhihui Du, David A. Bader</i>	
Breakthrough Edge AI Inference Performance using Northpole in 3U Vpx Form Factor.....	31
<i>Filipp Akopyan, William P. Risk, John V. Arthur, Andrew S. Cassidy, Michael V. Debole, Carlos O. Otero, Jun Sawada, Evan Colgan, Michael E. Criscolo, Phillip V. Mann, Heinz Baier, Kai Schleupen, Arnon Amir, Alexander Andreopoulos, Rathinakumar Appuswamy, Deepika Bablani, Peter J. Carlson, Pallab Datta, Steven K. Esser, Myron D. Flickner, Rajamohan Gandhasri, Guillaume J. Garreau, Megumi Ito, Jennifer L. Klamo, Jeffrey A. Kusnitz, Nathaniel J. McClatchey, Neil McGlohon, Jeffrey L. McKinstry, Yutaka Nakamura, Tapan K. Nayak, Jay Sivagnaname, Daniel F. Smith, Rafael Sousa, Brian Taba, Ignacio Terrizzano, Takanori Ueda, Dharmendra S. Modha</i>	
Investigating Resilience of Loops in HPC Programs: A Semantic Approach with LLMs.....	36
<i>Hailong Jiang, Jianfeng Zhu, Bo Fang, Chao Chen, Qiang Guan</i>	
Breakthrough Low-Latency, High-Energy-Efficiency LLM Inference Performance using NorthPole	46
<i>Rathinakumar Appuswamy, Michael V. Debole, Brian Taba, Steven K. Esser, Andrew S. Cassidy, Arnon Amir, Alexander Andreopoulos, Deepika Bablani, Pallab Datta, Jeffrey A. Kusnitz, Nathaniel J. McClatchey, Neil McGlohon, Jeffrey L. McKinstry, Tapan K. Nayak, Daniel F. Smith, Rafael Sousa, Ignacio Terrizzano, Filipp Akopyan, Peter J. Carlson, Rajamohan Gandhasri, Guillaume J. Garreau, Nelson M. Gonzalez, Megumi Ito, Jennifer L. Klamo, Yutaka Nakamura, Carlos O. Otero, William P. Risk, Jun Sawada, Kai Schleupen, Jay Sivagnaname, Matthew Stallone, Takanori Ueda, Myron D. Flickner, John V. Arthur, Rameswar Panda, David D. Cox, Dharmendra S. Modha</i>	
High Performance Im2win and Direct Convolutions using Three Tensor Layouts on SIMD Architectures	54
<i>Xiang Fu, Xinpeng Zhang, Jixiang Ma, Peng Zhao, Shuai Lu, Xu T. Liu</i>	
What is Normal? A Big Data Observational Science Model of Anonymized Internet Traffic.....	61
<i>Jeremy Kepner, Hayden Jananthan, Michael Jones, William Arcand, David Bestor, William Bergeron, Daniel Burrill, Aydin Buluc, Chansup Byun, Timothy Davis, Vijay Gadepally, Daniel Grant, Michael Houle, Matthew Hubbell, Piotr Luszczek, Lauren Milechin, Chasen Milner, Guillermo Morales, Andrew Morris, Julie Mullen, Ritesh Patel, Alex Pentland, Sandeep Pisharody, Andrew Prout, Albert Reuther, Antonio Rosa, Gabriel Wachman, Charles Yee, Peter Michaleas</i>	

Ultra Low Latency Hardware Optimised Radix-4 FFT for Optical Wireless FPGA Transceivers via Hermitian Symmetry Characteristics.....	68
<i>Michael Codd, Ciara McDonald, Yiyue Jiang, Chunan Chen, Holger Claussen, Miriam Leeser, John Dooley</i>	
Constant-Memory Graph Coarsening.....	75
<i>George M. Slota, Christopher Brissette</i>	
GLITCHES: GPU-FPGA LLM Inference Through a Collaborative Heterogeneous System.....	82
<i>Fan Yang, Xinhao Yang, Hongyi Wang, Zehao Wang, Zhenhua Zhu, Shulin Zeng, Yu Wang</i>	
Efficient Eigenvalue Computation of Parahermitian Matrices using Neural Networks	89
<i>Diyari A. Hassan, Yunus Egi, Soydan Redif</i>	
Optimization Strategies to Accelerate BLAS Operations with ARM SVE.....	95
<i>Aniket P. Garade, Sushil P. Singh, Juliya James, H. V. Deepika, P. Haribabu, S. A. Kumar, S. D. Sudarsan</i>	
pc-COP: An Efficient and Configurable 2048-p-Bit Fully-Connected Probabilistic Computing Accelerator for Combinatorial Optimization.....	102
<i>Kiran Magar, Shreya Bharathan, Utsav Banerjee</i>	
Enhanced Knowledge Graph Attention Networks for Efficient Graph Learning	109
<i>Fernando V. Buschmann, Zhihui Du, David Bader</i>	
Benchmarking the Performance of Large Language Models on the Cerebras Wafer Scale Engine	116
<i>Zuoning Zhang, Dhruv Parikh, Youning Zhang, Viktor Prasanna</i>	
IRIS-MEMFLOW: Data Flow-Enabled Portable Memory Orchestration in IRIS Runtime for Diverse Heterogeneity	123
<i>Mohammad A. H. Monil, Narasinga R. Miniskar, Seyong Lee, Beau Johnston, Pedro Valero-Lara, Aaron Young, Keita Teranishi, Jeffrey S. Vetter</i>	
A Performance Analysis of GPU-Aware MPI Implementations Over the Slingshot-11 Interconnect.....	132
<i>Michael Beebe, Rahul Kumar Gayatri, Kevin Gott, Adam Lavelly, Muhammad Haseeb, Brandon Cook, Yong Chen</i>	
Extracting TCPIP Headers at High Speed for the Anonymized Network Traffic Graph Challenge.....	139
<i>Zhaoyang Han, Andrew Briasco-Stewart, Michael Zink, Miriam Leeser</i>	
Anonymized Network Sensing Graph Challenge.....	145
<i>Hayden Jananthan, Michael Jones, William Arcand, David Bestor, William Bergeron, Daniel Burrill, Aydin Buluc, Chansup Byun, Timothy Davis, Vijay Gadepally, Daniel Grant, Michael Houle, Matthew Hubbell, Piotr Luszczek, Peter Michaleas, Lauren Milechin, Chasen Milner, Guillermo Morales, Andrew Morris, Julie Mullen, Ritesh Patel, Alex Pentland, Sandeep Pisharody, Andrew Prout, Albert Reuther, Antonio Rosa, Gabriel Wachman, Charles Yee, Jeremy Kepner</i>	
Multiplication of Sparse Matrices and Their Transpose using Compressed Sparse Diagonals	153
<i>Sardar A. Haque, Mohammad T. Parvez, Shahadat Hossain</i>	
Dynamic Task Scheduling with Data Dependency Awareness using Julia.....	159
<i>Rabab Alomairy, Felipe Tome, Julian Samaroo, Alan Edelman</i>	
Comparison of Vectorization Capabilities of Different Compilers for X86 and ARM CPUs.....	166
<i>Nazmus Sakib, Tarun Prabhu, Nandakishore Santhi, John Shalf, Abdel-Hameed A. Badawy</i>	

A Neural Network Based GCC Cost Model for Faster Compiler Tuning	173
<i>Hafsah Shahzad, Ahmed Sanaullah, Sanjay Arora, Ulrich Drepper, Martin Herbordt</i>	
A Dynamic Weighting Strategy to Mitigate Worker Node Failure in Distributed Deep Learning	182
<i>Yuesheng Xu, Arielle Carr</i>	
Scalable Approach for Analytic Polynomial Subspace Projection Matrices for a Space-Time Covariance Matrix	188
<i>Faizan A. Khattak, Mohammed Bakhit, Ian K. Proudler, Stephan Weiss</i>	
Benchmarking Edge AI Platforms for High-Performance ML Inference	193
<i>Rakshith Jayanth, Neelesh Gupta, Viktor Prasanna</i>	
Power Efficient Deep Learning Acceleration using Intel Xeon Processors	200
<i>Xiaofei Jiang, Mona Minakshi, Rajesh Poornachandran, Shamima Najnin</i>	
Persistent and Partitioned MPI for Stencil Communication	206
<i>Gerald Collom, Jason Burmark, Olga Pearce, Amanda Bienz</i>	
A Highly Scalable Parallel Design for Data Compression	213
<i>S. Biplab Raut</i>	
Quantum Machine Learning in the Cognitive Domain: Alzheimer's Disease Study	218
<i>Emine Akpinar</i>	
MonoCoder: Domain-Specific Code Language Model for HPC Codes and Tasks	224
<i>Tal Kadosh, Niranjana Hasabnis, Vy A. Vo, Nadav Schneider, Neva Krien, Mihai Capota, Abdul Wasay, Guy Tamir, Ted Willke, Nesreen Ahmed, Yuval Pinter, Timothy Mattson, Gal Oren</i>	
A Graph-Based Algorithm for Optimizing GCC Compiler Flag Settings	231
<i>Reza Sajjadinasab, Sanjay Arora, Ulrich Drepper, Ahmed Sanaullah, Martin Herbordt</i>	
Fully Transparent Client-Side Caching for Key-Value Store Applications using FPGAs	239
<i>Sahan Bandara, Noah Cherry, Martin Herbordt</i>	
Elucidating US Import Supply Chain Dynamics: A Spatial-Temporal Graph Neural Network Approach	246
<i>Nikolay Aristov, Ziyang Li, Thomas Koch, Elenna R. Dugundji</i>	
Towards Just-in-Time Instruction Generation for Accelerated Sparse Matrix-Matrix Multiplication on GPUs	252
<i>Seth D. Kay, H. Howie Huang</i>	
A Framework to Enable Algorithmic Design Choice Exploration in DNNs	259
<i>Timothy L. Cronin, Sanmukh Kuppannagari</i>	
Graphical Learning Optimization and Dimensionality Reduction with Geometric Multi-Resolution Analysis	266
<i>Felicia Schenkelberg, Allison Gunby-Mann, Emma Graham, Shuoxuan Li, Peter Chin</i>	
Community Detection for Large Graphs on GPUs with Unified Memory	274
<i>Emre Dinçer, Isil Öz</i>	
Exploring Sparse Inference with SuiteSparse:GraphBLAS	281
<i>Deepak Suresh, Tim Davis</i>	

Privacy-Preserving AI for Document Understanding with Controlled Unclassified Information	287
<i>Scott M. Sawyer</i>	
A Multilevel Approach for Solving Large-Scale QUBO Problems with Noisy Hybrid Quantum Approximate Optimization	294
<i>Filip B. Maciejewski, Bao G. Bach, Maxime Dupont, P. Aaron Lott, Bhuvanesh Sundar, David E. B. Neira, Ilya Safro, Davide Venturelli</i>	
Binary Bleed: Fast Distributed and Parallel Method for Automatic Model Selection	304
<i>Ryan Barron, Maksim E. Eren, Manish Bhattarai, Ismael Boureima, Cynthia Matuszek, Boian S. Alexandrov</i>	
A Deployment Tool for Large Scale Graph Analytics Framework Arachne	312
<i>Garrett Gonzalez-Rivas, Zhihui Du, David A. Bader</i>	
Intel Xeon Optimization for Efficient Media Workload Acceleration	319
<i>Karan Puttannaiah, Rajesh Poornachandran</i>	
MESM: A Query-Agnostic and Memory-Efficient Parallel Subgraph Matching Algorithm	324
<i>Shubhashish Kar, Shaikh Arifuzzaman</i>	
Enhancing Code Translation in Language Models with Few-Shot Learning via Retrieval- Augmented Generation	331
<i>Manish Bhattarai, Javier E. Santos, Shawn Jones, Ayan Biswas, Boian Alexandrov, Daniel O'Malley</i>	
VeBPF Many-Core Architecture for Network Functions in FPGA-Based SmartNICs and IoT	339
<i>Zaid Tahir, Ahmed Sanaullah, Sahan Bandara, Ulrich Drepper, Martin Herbordt</i>	
P-YOLOv8: Efficient and Accurate Real-Time Detection of Distracted Driving	348
<i>Mohamed R. Elshamy, Heba M. Emara, Mohamed R. Shoaib, Abdel-Hameed A. Badawy</i>	
Transformers: A Graph Processing Perspective	354
<i>Manish S. S. S. Routhu, Sai D. Yanduru, Nathaniel Tomczak, Sanmukh Kuppannagari</i>	
Distributed-Memory Sparse Deep Neural Network Inference using Global Arrays	361
<i>Bruce Palmer, Sayan Ghosh, Andrés Márquez</i>	
Indexed Binary Operations in the GraphBLAS	368
<i>Timothy G. Mattson, Manaswinee Bezbaruah, Matthias Maier, Scott McMillan, Michel Peletier, Erik Welch, Timothy A. Davis</i>	
An Efficient Multi-Core Parallel Implementation of SSSP Algorithm with Decreasing Delta- Stepping	375
<i>Rakibul Hassan, Shaikh Arifuzzaman</i>	
Accelerating Multi-Agent DDPG Training on Multi-GPU Platforms	382
<i>Samuel Wiggins, Viktor Prasanna</i>	
Analyzing an In-Line Compression on the Matrix Matrix Multiplication Kernel	387
<i>Steven Platt, Jon C. Calhoun</i>	
Compressed Cannon's Algorithm	398
<i>Louis Jencka, Amanda Bienz</i>	
A High-Performance Curve25519 and Curve448 Unified Elliptic Curve Cryptography Accelerator	404
<i>Aniket Banerjee, Utsav Banerjee</i>	

A Framework for Analyzing the Performance of Sparse Matrix and Graph Operations	411
<i>Khaled Abdelaal, Richard Veras</i>	
Parallel Online Directed Acyclic Graph Exploration for Atlasing Soft-Matter Assembly Configuration Spaces.....	418
<i>Rahul Prabhu, Amit Verma, Meera Sitharam</i>	
LLM Inference Serving: Survey of Recent Advances and Opportunities.....	426
<i>Baolin Li, Yankai Jiang, Vijay Gadepally, Devesh Tiwari</i>	
The Genomic Computing Revolution: Defining the Next Decades of Accelerating Genomics	434
<i>Harisankar Sadasivan, Artur Klauser, Juergen Hench, Yatish Turakhia, Gagandeep Singh, Alberto Zeni, Sarah Beecroft, Satish Narayanasamy, Jeff Nivala, Bob Robey, Onur Mutlu, Kristof Denolf, Sriranjani Sitaraman</i>	
Multilevel Diffusion Based Spectral Graph Clustering	443
<i>Malik Lechekhab, Dimosthenis Pasadakis, Olaf Schenk</i>	
GPU Accelerated Construction of Time Respecting Data Structure for Temporal Graphs	450
<i>Animan Naskar, Venkata K. Tavva, Subhasis Banerjee</i>	
Characterization and Optimization of the Fitting of Quantum Correlation Functions	457
<i>Pi-Yueh Chuang, Niteya Shah, Patrick Barry, Ian Cloët, Emil M. Constantinescu, Nobuo Sato, Jian-Wei Qiu, Wu-Chun Feng</i>	
On the Scalability of Computing Genomic Diversity using SparkLeBLAST: A Feasibility Study	465
<i>Ritvik Prabhu, Bernard Moussad, Karim Youssef, Emil Vatai, Wu-Chun Feng</i>	
Experiences with VITIS AI for Deep Reinforcement Learning	473
<i>Nabayan Chaudhury, Atharva Gondhalekar, Wu-Chun Feng</i>	
MST in Incremental Graphs Through Tree Contractions	480
<i>Akanksha Dwivedi, Sameer Sharma, Dip S. Banerjee</i>	
Augmenting HPC Profilers with Analysis Capabilities	487
<i>Abhishek Patil, K. V. Shamjith, R. K. Senthil Kumar, S. D. Sudarsan</i>	
Comparative Analysis of GCC and LLVM for Performance Optimization on Aarch64	494
<i>Mriganka Bezbaruah, Samruddhi Dhakulkar, Prachi Pandey, P. Haribabu, S. A. Kumar, S. D. Sudarsan</i>	
Sans: Streaming Anonymized Network Sensing	500
<i>Ketai Zhao, Yuhang Zhou, Hongxu Pan, Zhibin Wang, Sheng Zhong, Chen Tian</i>	
BB-CVXOPT: Basic Block Execution Count Estimation and Extrapolation using Constrained Convex Optimization.....	507
<i>Youssef Aly, Atanu Barai, Nandakishore Santhi, Abdel-Hameed A. Badawy</i>	
Spike-Driven YOLO: Ultra Low-Power Object Detection with Neuromorphic Computing.....	514
<i>Mark Barnell, Courtney Raymond, Lisa Loomis, Francesca Vidal, Daniel Brown, Darrek Isereau</i>	
Beyond Zero Trust: Reclaiming Blue Cyberspace with AI	519
<i>Scott A. Weed</i>	

Exploring the Trade-Off Between Repair Time and Reliability in Large Scale Cluster Computers: A Simulation-Based Approach	526
<i>Leslie A. Horace, Craig S. Walker, William M. Jones, Nathan A. Debardeleben, Vivian E. Hafener, Steven T. Senator</i>	
Accelerating Sensor Fusion in Neuromorphic Computing: A Case Study on Loihi-2.....	535
<i>Murat Isik, Karn Tiwari, Muhammed B. Eryilmaz, I. Can Dikmen</i>	
Application of Virtual Client for Azure Hardware Qualification.....	542
<i>Anna M. Mathew, Bryan Deyoung, Michael Chhor, Sharjil Khan</i>	
Tightly-Coupled FPGA Accelerator for Molecular Dynamics Simulation: Hardware-Software Co-Design and Fine-Grained Task Management.....	548
<i>Zekang Cheng, Zerong He, Xi Jin</i>	
Predictive Performance of Photonic SRAM-Based In-Memory Computing for Tensor Decomposition.....	555
<i>Sasindu Wijeratne, Sugeet Sunder, M. A.-A. Kaiser, Akhilesh Jaiswal, Clynn Mathew, Ajey P. Jacob, Viktor Prasanna</i>	
Composable Mission-Critical Embedded System Architecture for High Assurance.....	560
<i>Michael Vai, Eric Simpson, Alice Lee, Huy Nguyen, Jeffrey Hughes, Ben Nahill, Jeffery Lim, Roger Khazan, Sean O'Melia, Fred Schneider</i>	
LLMs for Closed-Library Multi-Document Query, Test Generation, and Evaluation.....	566
<i>Claire Randolph, Adam Michaleas, Darrell O. Ricke</i>	
Task-Level Parallelism for the Multifrontal Method in Tightly Coupled CPU-FPGA Architectures	576
<i>Zerong He, Zekang Cheng, Zhongguang Xu, Xi Jin</i>	
Hybrid Computing Architecture Based on Analog Phase-Change Memory Chips for Deep Neural Network Training	583
<i>Zhenhao Jiao, Xiaogang Chen, Tao Hong, Weibang Dai, Chengcai Tu, Shunfen Li, Houpeng Chen, Zitang Song</i>	
JACC.shared: Leveraging HPC Metaprogramming and Performance Portability for Computations that Use Shared Memory GPUs.....	589
<i>Pedro Valero-Lara, William F. Godoy, Keita Teranishi, Jeffrey S. Vetter</i>	
Hardware Trojan Detection Utilizing Graph Neural Networks and Structural Checking.....	596
<i>Hunter Nauman, Jia Di</i>	
OCO-GAT: An Accelerator for Graph Attention Network with Optimized Calculation Order	603
<i>Qi Liu, Wenjin Huang, Wenlu Peng, Yihua Huang</i>	
An Efficient Multi-DNN Accelerator Based on Multiple Systolic Arrays	610
<i>Jianjun Chen, Han Jiao, Wenjin Huang, Yihua Huang</i>	
HBM-Based Hardware Accelerator for GNN Sampling and Aggregation	617
<i>Yuchen Gui, Qizhe Wu, Wei Yuan, Huawen Liang, Xiaotian Wang, Xi Jin</i>	
Syndeo: Portable Ray Clusters with Secure Containerization	624
<i>William Li, Rodney S. L. Mercado, Jaime D. Peña, Ross E. Allen</i>	
IOS: A Low Cost Defense to Mitigate Meltdown and Spectre Like Attacks.....	635
<i>Xin Wang, Wei Zhang</i>	

MERCURY: Efficient Subgraph Matching on GPUs with Hybrid Scheduling.....	640
<i>Zhiheng Lin, Changjie Xu, Ke Meng, Guangming Tan</i>	
Algebraic Vertex Ordering of a Sparse Graph for Adjacency Access Locality and Graph Compression.....	647
<i>Dimitris Floros, Nikos Pitsianis, Xiaobai Sun</i>	
Quantum Computing for Data Calibration in Parallel Magnetic Resonance Imaging Reconstruction.....	654
<i>Girish B. Reddy, Gulfam A. Saju, Yi Liu, Yuchou Chang</i>	
Direct RF FPGAs Built with Multi-Chip Packaging Overcome Technology Challenges	660
<i>Marjorie Catt, Dustin Henderson</i>	
Understanding the Efficacy of Power Profiles: A Case Study of AMD Instinct MI100 GPU	665
<i>Ghazanfar Ali, Mert Side, Sridutt Bhalachandra, Tommy Dang, Alan Sill, Yong Chen</i>	
A Run-Time Configurable NTT Architecture for Homomorphic Encryption Based on 3D Algorithm.....	672
<i>Weicon Lu, Xiaojie Chen, Dihui Chen, Tao Su</i>	
ModelGauge: Inference Profiling of Deep-Learning Models.....	678
<i>Calvin B. Gealy, David Langerman, Alan D. George</i>	
Optimizing FPGA Memory Allocation for Matrix-Matrix Multiplication using Bayesian Optimization.....	685
<i>Mehmet Gungor, Stratis Ioannidis, Miriam Leeser</i>	
HPC Network Simulation Tuning via Automatic Extraction of Hardware Parameters	692
<i>Joshua Suetterlein, Stephen J. Young, Jesun Firoz, Joseph Manzano, Nathan Tallent, Ryan D. Friese, Kevin Barker, Timothy Stavenger</i>	
HENNC: Hardware Engine for Artificial Neural Network-Based Chaotic Oscillators	702
<i>Mobin Vaziri, Shervin Vakili, M. M. Rahimifar, J. M. Pierre Langlois</i>	
Evaluating One-Sided Communication on Graph500 with MPI-RMA and OpenSHMEM	707
<i>Jefferson Boothe, Alan D. George</i>	
FPGA Acceleration for Scalable High-Resolution OPIR Target Detection.....	714
<i>Daniel C. Stumpp, Alan D. George</i>	
Hunting the Needle - The Potential of Innovation in Architecture	721
<i>Peter M. Kogge, Janice McMahon, Timothy J. Dysart</i>	
Supercomputer 3D Digital Twin for User Focused Real-Time Monitoring	727
<i>Bill Bergeron, Matthew Hubbell, Daniel Mojica, Albert Reuther, William Arcand, David Bestor, Daniel Burrill, Chansup Byun, Vijay Gadepally, Michael Houle, Hayden Jananthan, Michael Jones, Piotr Luszczek, Peter Michaleas, Lauren Milechin, Julie Mullen, Andrew Prout, Antonio Rosa, Charles Yee, Jeremy Kepner</i>	
Reinforcement Learning-Generated Topological Order for Dynamic Task Graph Scheduling	735
<i>Cheng-Hsiang Chiu, Chedi Morchdi, Yi Zhou, Boyang Zhang, Che Chang, Tsung-Wei Huang</i>	
Mobile-Optimized Real-Time Vessel Detection for Ultrasound-Guided Surgical Procedures	742
<i>Mateusz Wolak, Fin Amin, Nancy Delosa, Brian Telfer, Benjamin Roop, Lars Gjestebj</i>	
Towards Faster Graph Partitioning via Pre-Training and Inductive Inference	749
<i>Meng Qin, Chaorui Zhang, Yu Gao, Yibin Ding, Weipeng Jiang, Weixi Zhang, Wei Han, Bo Bai</i>	

Benchmarking Thread Block Cluster	756
<i>Tim Lühnen, Tobias Marschner, Sohan Lal</i>	
LLoad: An Easy-to-Use HPC Utilization Tool	763
<i>Chansup Byun, Albert Reuther, Julie Mullen, Latoya Anderson, William Arcand, Bill Bergeron, David Bestor, Alexander Bonn, Daniel Burrill, Vijay Gadepally, Michael Houle, Matthew Hubbell, Hayden Jananthan, Michael Jones, Piotr Luszczek, Peter Michaleas, Lauren Milechin, Guillermo Morales, Andrew Prout, Antonio Rosa, Charles Yee, Jeremy Kepner</i>	
Energy Efficiency Scaling for 2 Decades (EES2) Roadmap for Computing.....	772
<i>T. Kaarsberg, J. Atulasimha, J. Baniecki, P. Fischer, S. Pawlowski, S. Misra, A. Bhavnagarwala, E. Salman, M. Ahmed, N. Li, R. Aggarwal, B. Hirano, T. Shah, C. Green, J. Booth, P. Sharps, T. McDonald, J. Ballard, Y. Chen, P. Nagapurkar, W. Huang, D. Kudithipudi, A. Paramonov, F. Musso, A. K. Ziabari, J. Luo, A. K. Petford-Long, D. Gopman, C. Gotama, T. Wei, S. Shaheen, Y. Zhang, I. Lu, K. Shimizu, E. Taylor, N. Johnson, R. Jones, S. Shankar</i>	
GPU Sharing with Triples Mode	775
<i>Chansup Byun, Albert Reuther, Latoya Anderson, William Arcand, Bill Bergeron, David Bestor, Alexander Bonn, Daniel Burrill, Vijay Gadepally, Michael Houle, Matthew Hubbell, Hayden Jananthan, Michael Jones, Piotr Luszczek, Peter Michaleas, Lauren Milechin, Guillermo Morales, Julie Mullen, Andrew Prout, Antonio Rosa, Charles Yee, Jeremy Kepner</i>	

Author Index