2020 IEEE High Performance Extreme Computing Conference (HPEC 2020)

Waltham, Massachusetts, USA
22 – 24 September 2020

IEEE Catalog Number: CFP20HPE-POD
ISBN: 978-1-7281-9220-8

****** 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:	CFP20HPE-POD
ISBN (Print-On-Demand):	978-1-7281-9220-8
ISBN (Online):	978-1-7281-9219-2
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

TABLE OF CONTENTS

EFFICIENT SPARSE MATRIX-VECTOR MULTIPLICATION ON INTEL PIUMA ARCHITECTURE 1
Sriram Aananthakrishnan; Robert Pawlowski; Joshua Fryman; Ibrahim Hur

BIT-ERROR AWARE QUANTIZATION FOR DCT-BASED LOSSY COMPRESSION 3
Jialing Zhang; Jiaxi Chen; Aekyeung Moon; Xiaoyan Zhuo; Seung Woo Son

LESSMINE: REDUCING SAMPLE SPACE AND DATA ACCESS FOR DENSE PATTERN MINING 10
Tianyu Fu; Ziqian Wan; Guohao Dai; Yu Wang; Huazhong Yang

A HIGH THROUGHPUT PARALLEL HASH TABLE ON FPGA USING XOR-BASED MEMORY 17
Ruizhi Zhang; Sasindu Wijeratne; Yang Yang; Sanmukh R. Kuppannagari; Viktor K. Prasanna

HALF-PRECISION FLOATING-POINT FORMATS FOR PAGERANK: OPPORTUNITIES AND CHALLENGES 24
Amir Sabbagh Molahosseini; Hans Vandierendonck

AN EFFICIENT LP ROUNDING SCHEME FOR REPLICA PLACEMENT 31
Zhihui Du; Sen Zhang; David A. Bader; Jingkun Hu

HOW TO EFFICIENTLY TRAIN YOUR AI AGENT? CHARACTERIZING AND EVALUATING DEEP REINFORCEMENT LEARNING ON HETEROGENEOUS PLATFORMS 38
Yuan Meng; Yang Yang; Sanmukh Kuppannagari; Rajgopal Kannan; Viktor Prasanna

FPGAS IN THE NETWORK AND NOVEL COMMUNICATOR SUPPORT ACCELERATE MPI COLLECTIVES 45
Pouya Haghi; Anqi Guo; Qingqing Xiong; Rushi Patel; Chen Yang; Tong Geng; Justin T. Broaddus; Ryan Marshall; Anthony Skjellum; Martin C. Herbordt

BANDWIDTH ALLOCATION IN SILICON-PHOTONIC NETWORKS USING APPLICATION INSTRUMENTATION 55
Aditya Narayan; Ajay Joshi; Ayse K. Coskun

DYNAMIC COMPUTATIONAL DIVERSITY WITH MULTI-RADIX LOGIC AND MEMORY 57
Paul G. Flikkema; James Palmer; Tolga Yalcin; Bertr Cambou

LARGE-SCALE SPARSE TENSOR DECOMPOSITION USING A DAMPED GAUSS-NEWTON METHOD 63
Teresa M. Ranadive; Muthu M. Baskaran

HASH TABLE SCALABILITY ON INTEL PIUMA 71
Balasubramanian Seshasayee; Joshua Fryman; Ibrahim Hur

DISTRIBUTED NON-NEGATIVE TENSOR TRAIN DECOMPOSITION 73
Manish Bhattarai; Gopinath Chennupati; Erik Skau; Raviteja Vangara; Hristo Djidjev; Boian S. Alexandrov

AUTOMATIC MAPPING AND OPTIMIZATION TO KOKKOS WITH POLYHEDRAL COMPILATION 83
Muthu Baskaran; Charles Jin; Benoit Meister; Jonathan Springer

MULTI-TEMPORAL ANALYSIS AND SCALING RELATIONS OF 100,000,000,000 NETWORK PACKETS 90
Jeremy Kepner; Chad Meiners; Chansup Byun; Sarah McGuire; Timothy Davis; William Arcand; Jonathan Bernays; David Bestor; William Bergeron; Vijay Gadepally; Raul Harnasch; Matthew Hubbell; Micheal Houle; Micheal Jones; Andrew Kirby; Anna Klein; Lauren Mile

A FRAMEWORK FOR TASK MAPPING ONTO HETEROGENEOUS PLATFORMS 96
Ta-Yang Wang; Ajitesh Srivastava; Viktor Prasanna

DISCRETE INTEGRATED CIRCUIT ELECTRONICS (DICE) 102
Zach Fredin; Jiri Zemanek; Camron Blackburn; Erik Strand; Amira Abdel-Rahman; Premila Rowles; Neil Gershenfeld

A DEEP Q-LEARNING APPROACH FOR GPU TASK SCHEDULING 110
Ryan S. Luley; Qinru Qiu

LAYER-PARALLEL TRAINING WITH GPU CONCURRENCY OF DEEP RESIDUAL NEURAL NETWORKS VIA NONLINEAR MULTIGRID 117
Andrew Kirby; Siddharth Samsi; Michael Jones; Albert Reuther; Jeremy Kepner; Vijay Gadepally

GPU-ACCELERATED DISCONTINUOUS GALERKIN METHODS: 30X SPEEDUP ON 345 BILLION UNKNOWNS 124
Andrew C. Kirby; Dimitri J. Mavriplis

COMPUTE, TIME AND ENERGY CHARACTERIZATION OF ENCODER-DECODER NETWORKS WITH AUTOMATIC MIXED PRECISION TRAINING 131
Siddharth Samsi; Michael Jones; Mark M. Veillette

OPTIMIZING USE OF DIFFERENT TYPES OF MEMORY FOR FPGAS IN HIGH PERFORMANCE COMPUTING 137
Kai Huang; Mehmet Gungor; Stratis Ioannidis; Miriam Leeser

BENCHMARKING NETWORK FABRICS FOR DATA DISTRIBUTED TRAINING OF DEEP NEURAL NETWORKS 144
Siddharth Samsi; Andrew Prout; Michael Jones; Andrew Kirby; Bill Arcand; Bill Bergeron; David Bestor; Chansup Byun; Vijay Gadepally; Michael Houle; Matthew Hubbell; Anna Klein; Peter Michaleas; Lauren Milechin; Julie Mullen; Antonio Rosa; Charles Yee; Alb

USING RAPIDS AI TO ACCELERATE GRAPH DATA SCIENCE WORKFLOWS 150
Todd Hricik; David Bader; Oded Green

NORTHEAST CYBERTEAM - BUILDING AN ENVIRONMENT FOR SHARING BEST PRACTICES AND SOLUTIONS FOR RESEARCH COMPUTING 154
John Goodhue; Julie Ma; Adrian Del Masetro; Sia Najafi; Bruce Segee; Scott Valcourt; Ralph Zottola

EXECUTION OF COMPLETE MOLECULAR DYNAMICS SIMULATIONS ON MULTIPLE FPGAS 159
Carlo Pascoe; Lawrence Stewart; Brian W. Sherman; Vipin Sachdeva; Martin W. Herbordt

DESIGN, OPTIMIZATION, AND BENCHMARKING OF DENSE LINEAR ALGEBRA ALGORITHMS ON AMD GPUS 161
Cade Brown; Ahmad Abdelfattah; Stanimire Tomov; Jack Dongarra

ARITHMETIC AND BOOLEAN SECRET SHARING MPC ON FPGAS IN THE DATA CENTER 168
Rushi Patel; Pierre-François Wolfe; Robert Munafo; Mayank Varia; Martin Herbordt

ACCELERATOR DESIGN AND PERFORMANCE MODELING FOR HOMOMORPHIC ENCRYPTED CNN INFERENCE 176
Tian Ye; Rajgopal Kannan; Viktor K. Prasanna

A COMMUNICATION-EFFICIENT MULTI-CHIP DESIGN FOR RANGE-LIMITED MOLECULAR DYNAMICS 183
Chunshu Wu; Tong Geng; Chen Yang; Vipin Sachdeva; Woody Sherman; Martin Herbordt

FAST GRAPHLET TRANSFORM OF SPARSE GRAPHS 191
Dimitris Floros; Nikos Pitsianis; Xiaobai Sun

VARIABLE PRECISION MULTIPLICATION FOR SOFTWARE-BASED NEURAL NETWORKS 199
Richa Singh; Thomas Conroy; Patrick Schaumont

A DYNAMICALLY CONFIGURABLE NETWORK FOR SOFTWARE-DEFINED HARDWARE 206
William Butera

A CONGESTION CONTROL MECHANISM FOR SDN-BASED FAT-TREE NETWORKS 213
Haitham Ghalwash; Chun-Hsi Huang

GBTLX: A FIRST LOOK 220
Sanil Rao; Anurag Kutuluru; Paul Brouwer; Scott McMillan; Franz Franchetti

EXPLOITING GPU DIRECT ACCESS TO NON-VOLATILE MEMORY TO ACCELERATE BIG DATA PROCESSING 227
Mahsa Bayati; Miriam Leeser; Ningfang Mi

SURVEY OF MACHINE LEARNING ACCELERATORS 233
Albert Reuther; Peter Michaleas; Michael Jones; Vijay Gadepally; Siddharth Samsi; Jeremy Kepner

TOWARDS AN OBJECTIVE METRIC FOR THE PERFORMANCE OF EXACT TRIANGLE COUNT 245
Mark P. Blanco; Scott McMillan; Tze Meng Low

PROFILING AND OPTIMIZATION OF CT RECONSTRUCTION ON NVIDIA QUADRO GV100 252
Shekhar Dwivedi; Andreas Heumann

SELF-SCALING CLUSTERS AND REPRODUCIBLE CONTAINERS TO ENABLE SCIENTIFIC COMPUTING 259
Peter Z. Vaillancourt; J. Eric Coulter; Richard Knepper; Brandon Barker

TOWARDS A DISTRIBUTED FRAMEWORK FOR MULTI-AGENT REINFORCEMENT LEARNING RESEARCH 267
Yutai Zhou; Shawn Manuel; Peter Morales; Sheng Li; Jaime Pena; Ross Allen

HOMOMORPHIC ENCRYPTION BASED SECURE SENSOR DATA PROCESSING 276
Vijay Gadepally; Mihailo Isakov; Rashmi Agrawal; Jeremy Kepner; Karen Gettings; Michel A. Kinsy

SCALABLE DATA GENERATION FOR EVALUATING MIXED-PRECISION SOLVERS 283
Piotr Luszczek; Yaohung Tsai; Neil Lindquist; Hartwig Anzt; Jack Dongarra

TRIANGLE COUNTING WITH CYCLIC DISTRIBUTIONS 289
Andrew Lumsdaine; Luke Dalessandro; Kevin Deweese; Jesun Firoz; Scott McMillan

METACL: AUTOMATED "META" OPENCL CODE GENERATION FOR HIGH-LEVEL SYNTHESIS ON FPGA 297
Paul Sathre; Atharva Gondhalekar; Mohamed Hassan; Wu-Chun Feng
A COMPREHENSIVE COMPARISON AND ANALYSIS OF OPENACC AND OPENMP 4.5 FOR NVIDIA GPUS 305
R. Usha; Prachi Pandey; N. Mangala
A FEASIBILITY STUDY FOR MPI OVER HDFS 311
W. Feng; D. Zhang; J. Zhang; K. Hou; S. Pumma; H. Wang
PROCESSING OF CROWDSOURCED OBSERVATIONS OF AIRCRAFT IN A HIGH PERFORMANCE COMPUTING ENVIRONMENT 318
Andrew Weinert; Ngaire Underhill; Bilal Gill; Ashley Wicks
WORK-EFFICIENT PARALLEL ALGORITHMS FOR ACCURATE FLOATING-POINT PREFIX SUMS 324
Sean Fraser; Helen Xu; Charles E. Leiserson
MINESWEEPER: A NOVEL AND FAST ORDERED-STATISTIC CFAR ALGORITHM 331
Carl L. Colena; Michael J. Russell; Stephen A. Braun
STABILITY-OPTIMIZED HIGH ORDER METHODS AND STIFFNESS DETECTION FOR PATHWISE STIFF STOCHASTIC DIFFERENTIAL EQUATIONS 337
Chris Rackauckas; Qing Nie
KTRUSSEXPLORER: EXPLORING THE DESIGN SPACE OF K-TRUSS DECOMPOSITION OPTIMIZATIONS ON GPUS 345
Safaa Diab; Mhd Ghaith Olabi; Izzat El Hajj
SCALING GRAPH CLUSTERING WITH DISTRIBUTED SKETCHES 353
Benjamin W. Priest; Alec Dunton; Geoffrey Sanders
GRAPHCHALLENGE.ORG TRIANGLE COUNTING PERFORMANCE 360
Siddharth Samsi; Jeremy Kepner; Vijay Gadepally; Michael Hurley; Michael Jones; Edward Kao; Sanjeev Mohindra; Albert Reuther; Steven Smith; William Song; Diane Staheli; Paul Monticciolo
GRAPHCHALLENGE.ORG SPARSE DEEP NEURAL NETWORK PERFORMANCE 369
Jeremy Kepner; Simon Alford; Vijay Gadepally; Michael Jones; Lauren Milechin; Albert Reuther; Ryan Robinett; Sid Samsi
ANALYSIS OF FLOATING-POINT ROUND-OFF ERROR IN LINEAR ALGEBRA ROUTINES FOR GRAPH CLUSTERING 376
L. Minah Yang; Alyson Fox
AT-SCALE SPARSE DEEP NEURAL NETWORK INFERENCE WITH EFFICIENT GPU IMPLEMENTATION 383
Mert Hidayetoglu; Carl Pearson; Vikram Sharma Mailthody; Eiman Ebrahimi; Jinjun Xiong; Rakesh Nagi; Wen-Mei Hwu
COMBINATORIAL TILING FOR SPARSE NEURAL NETWORKS 390
Filip Pawlowski; Rob H. Bisseling; Bora Uçar; A. N. Yzelman
TRIC: DISTRIBUTED-MEMORY TRIANGLE COUNTING BY EXPLOITING THE GRAPH STRUCTURE 397
Sayan Ghosh; Mahantesh Halappanavar
POST QUANTUM CRYPTOGRAPHY(PQC) - AN OVERVIEW 403
Manoj Kumar; Pratap Pattnaik
DENIAL OF SERVICE IN CPU-GPU HETEROGENEOUS ARCHITECTURES 412
Hao Wen; Wei Zhang
PACKING NARROW-WIDTH OPERANDS TO IMPROVE ENERGY EFFICIENCY OF GENERAL-PURPOSE GPU COMPUTING 417
Xin Wang; Wei Zhang
GPU ACCELERATED ANOMALY DETECTION OF LARGE SCALE LIGHT CURVES 424
Austin Chase Minor; Zhihui Du; Yankui Sun; David A. Bader; Chao Wu; Jianyan Wei
EVALUATING CRYPTOGRAPHIC PERFORMANCE OF RASPBERRY PI CLUSTERS 431
Daniel Hawthorne; Michael Kapralos; Raymond W. Blaine; Suzanne J. Matthews
CQNN: A CGRA-BASED QNN FRAMEWORK 440
Tong Geng; Chunshu Wu; Cheng Tan; Bo Fang; Ang Li; Martin Herbordt
FAST GPU GRAPH CONTRACTION BY COMBINING EFFICIENT SHALLOW SEARCHES AND POST-CULLING 447
Roozbeh Karimi; David M. Koppelman; Chris J. Michael
OPENCL PERFORMANCE ON THE INTEL HETEROGENEOUS ARCHITECTURE RESEARCH PLATFORM 454
Steven Harris; Roger D. Chamberlain; Christopher Gill

HOMOMORPHIC ENCRYPTION FOR QUANTUM ANNEALING WITH SPIN REVERSAL TRANSFORMATIONS 463
Daniel O'Malley; John K. Golden
A GRAPHBLAS SOLUTION TO THE SIGMOD 2014 PROGRAMMING CONTEST USING MULTI-SOURCE BFS 469
Márton Elekes; Attila Nagy; Dávid Sándor; János Benjamin Antal; Timothy A. Davis; Gábor Szárnyas
FAST MAPPING ONTO CENSUS BLOCKS 476
Jeremy Kepner; Andreas Kipf; Darren Engwirda; Navin Vembar; Michael Jones; Lauren Milechin; Vijay Gadepally; Chris Hill; Tim Kraska; William Arcand; David Bestor; William Bergeron; Chansup Byun; Matthew Hubbell; Michael Houle; Andrew Kirby; Anna Klein; Ju
HUMAN BALANCE MODELS OPTIMIZED USING A LARGE-SCALE, PARALLEL ARCHITECTURE WITH APPLICATIONS TO MILD TRAUMATIC BRAIN INJURY 484
Gregory A. Ciccarelli; Michael Nolan; Hrishikesh M. Rao; Tanya Talkar; Anne O'Brien; Gloria Vergara-Diaz; Ross Zafonte; Thomas F. Quatieri; Ryan J. McKindles; Paolo Bonato; Adam Lammert
HARDWARE ACCELERATION OF NONLOCAL MEANS-BASED SPECKLE NOISE REMOVAL APPLIED TO SAR IMAGERY 492
Hector A. Li Sanchez; Alan D. George
DESIGN AND PERFORMANCE EVALUATION OF OPTIMIZATIONS FOR OPENCL FPGA KERNELS 499
Anthony M. Cabrera; Roger D. Chamberlain
ENHANCED PARALLEL SIMULATION FOR ACAS X DEVELOPMENT 506
Adam Gjersvik
CHIP-TO-CHIP OPTICAL DATA COMMUNICATIONS USING POLARIZATION DIVISION MULTIPLEXING 513
Darko Ivanovich; Chenfeng Zhao; Xuan Zhang; Roger D. Chamberlain; Amit Deliwala; Viktor Gruev
A NOVEL INFERENCE ALGORITHM FOR LARGE SPARSE NEURAL NETWORK USING TASK GRAPH PARALLELISM 521
Dian-Lun Lin; Tsung-Wei Huang
ACCELERATING DISTRIBUTED INFERENCE OF SPARSE DEEP NEURAL NETWORKS VIA MITIGATING THE STRAGGLER EFFECT 528
Mohammad Hasanzadeh Mofrad; Rami Melhem; Yousuf Ahmad; Mohammad Hammoud
STUDYING THE EFFECTS OF HASHING OF SPARSE DEEP NEURAL NETWORKS ON DATA AND MODEL PARALLELISMS 535
Mohammad Hasanzadeh Mofrad; Rami Melhem; Yousuf Ahmad; Mohammad Hammoud
FAST TRAINING OF DEEP NEURAL NETWORKS ROBUST TO ADVERSARIAL PERTURBATIONS 542
Justin Goodwin; Olivia Brown; Victoria Helus
TRAIN AND DEPLOY AN IMAGE CLASSIFIER FOR DISASTER RESPONSE 549
Jianyu Mao; Kiana Harris; Nae-Rong Chang; Caleb Pennell; Yiming Ren
IDENTIFYING EXECUTION ANOMALIES FOR DATA INTENSIVE WORKFLOWS USING LIGHTWEIGHT ML TECHNIQUES 554
Cong Wang; George Papadimitriou; Mariam Kiran; Anirban Mandal; Ewa Deelman
A SCALABLE ARCHITECTURE FOR CNN ACCELERATORS LEVERAGING HIGH-PERFORMANCE MEMORIES 561
Maarten Hattink; Giuseppe Di Guglielmo; Luca P. Carloni; Keren Bergman
A HARDWARE ROOT-OF-TRUST DESIGN FOR LOW-POWER SOC EDGE DEVICES 567
Alan Ehret; Eliakin Del Rosario; Karen Gettings; Michel A. Kinsy
PARAMETER SENSITIVITY ANALYSIS OF THE SPARTEN HIGH PERFORMANCE SPARSE TENSOR DECOMPOSITION SOFTWARE 573
Jeremy M. Myers; Daniel M. Dunlavy; Keita Teranishi; D. S. Hollman
SPARTEN: LEVERAGING KOKKOS FOR ON-NODE PARALLELISM IN A SECOND-ORDER METHOD FOR FITTING CANONICAL POLYADIC TENSOR MODELS TO POISSON DATA 580
Keita Teranishi; Daniel M. Dunlavy; Jeremy M. Myers; Richard F. Barrett
ARCHITECTURAL ANALYSIS OF DEEP LEARNING ON EDGE ACCELERATORS 587
Luke Kljucaric; Alex Johnson; Alan D. George
VYASA: A HIGH-PERFORMANCE VECTORIZING COMPILER FOR TENSOR CONVOLUTIONS ON THE XILINX AI ENGINE 594
Prasanth Chatarasi; Stephen Neuendorffer; Samuel Bayliss; Kees Vissers; Vivek Sarkar
A HYBRID-PIPELINED ARCHITECTURE FOR FPGA-BASED BINARY WEIGHT DENSENET WITH HIGH PERFORMANCE-EFFICIENCY 604
Shihao Zeng; Yihua Huang

USING GRAPHLET SPECTROGRAMS FOR TEMPORAL PATTERN ANALYSIS OF VIRUS-RESEARCH COLLABORATION NETWORKS 609
Dimitris Floros; Tiancheng Liu; Nikos Pitsianis; Xiaobai Sun

GRAPHSDH: A GENERAL GRAPH SAMPLING FRAMEWORK WITH DISTRIBUTION AND HIERARCHY 616
Jingbo Hu; Guohao Dai; Yu Wang; Huazhong Yang

ACCURACY AND PERFORMANCE COMPARISON OF VIDEO ACTION RECOGNITION APPROACHES 623
Matthew Hutchinson; Siddharth Samsi; William Arcand; David Bestor; Bill Bergeron; Chansup Byun; Micheal Houle; Matthew Hubbell; Micheal Jones; Jeremy Kepner; Andrew Kirby; Peter Michaleas; Lauren Milechin; Julie Mullen; Andrew Prout; Antonio Rosa; Albert

APPROXIMATE INVERSE CHAIN PRECONDITIONER: ITERATION COUNT CASE STUDY FOR SPECTRAL SUPPORT SOLVERS 631
M. Harper Langston; Pierre-David Letourneau; Julia Wei; Larry Weintraub; Mitchell Harris; Richard Lethin; Eric Papenhausen; Meifeng Lin

TOTAL IONIZING DOSE RADIATION TESTING OF NVIDIA JETSON NANO GPUS 639
Windy S. Slater; Nayana P. Tiwari; Tyler M. Lovelly; Jesse K. Mee

BEYOND FLOATING-POINT OPS: CNN PERFORMANCE PREDICTION WITH CRITICAL DATAPATH LENGTH 642
David Langerman; Alex Johnson; Kyle Buettner; Alan D. George

HYBRID APPROACH TO HPC CLUSTER TELEMETRY AND HARDWARE LOG ANALYTICS 651
Justin Thaler; Woong Shin; Steven Roberts; James H. Rogers; Todd Rosedahl

EVALUATING SEU RESILIENCE OF CNNS WITH FAULT INJECTION 658
Evan T. Kain; Tyler M. Lovelly; Alan D. George

COMPUTING PAGERANK SCORES OF WEB CRAWL DATA USING DGX A100 CLUSTERS 663
Seunghwa Kang; Alex Fender; Joe Eaton; Brad Rees

IMPLEMENTING SPARSE LINEAR ALGEBRA KERNELS ON THE LUCATA PATHFINDER-A COMPUTER 667
Géraud P. Krawezik; Shannon K. Kuntz; Peter M. Kogge

ACTIVE LEARNING PIPELINE FOR BRAIN MAPPING IN A HIGH PERFORMANCE COMPUTING ENVIRONMENT 673
Adam Michaleas; Lars A. Gjesteby; Michael Snyder; David Chavez; Meagan Ash; Matthew A. Melton; Damon G. Lamb; Sara N. Burke; Kevin J. Otto; Lee Kamentsky; Webster Guan; Kwanghun Chung; Laura J. Brattain

SCALABILITY OF STREAMING ON MIGRATING THREADS 679
Brian A. Page; Peter M. Kogge

ACCELERATING MRI RECONSTRUCTION ON TPUS 687
Tianjian Lu; Thibault Marin; Yue Zhuo; Yi-Fan Chen; Chao Ma

PROJECTING PERFORMANCE FOR PIUMA USING DOWN-SCALED SIMULATION 696
Stijn Eyerman; Wim Heirman; Yigit Demir; Kristof Du Bois; Ibrahim Hur

A FAULT TOLERANT IMPLEMENTATION FOR A MASSIVELY PARALLEL SEISMIC FRAMEWORK 703
Suha N. Kayum; Hussain Alsalim; Thierry-Laurent Tonellot; Ali Momin

TARGET CLASSIFICATION IN SYNTHETIC APERTURE RADAR AND OPTICAL IMAGERY USING LOIHI NEUROMORPHIC HARDWARE 711
Mark Barnell; Courtney Raymond; Matthew Wilson; Darrek Isereau; Chris Cicotta

IBENCH: A DISTRIBUTED INFERENCE SIMULATION AND BENCHMARK SUITE 717
Wesley Brewer; Greg Behm; Alan Scheinine; Ben Parsons; Wesley Emeneker; Robert P. Trevino

BEST OF BOTH WORLDS: HIGH PERFORMANCE INTERACTIVE AND BATCH LAUNCHING 723
Chansup Byun; Jeremy Kepner; William Arcand; David Bestor; Bill Bergeron; Vijay Gadepally; Michael Houle; Matthew Hubbell; Michael Jones; Andrew Kirby; Anna Klein; Peter Michaleas; Lauren Milechin; Julie Mullen; Andrew Prout; Antonio Rosa; Siddharth Samsi

PERFORMANCE STRATEGIES FOR PARALLEL BITONIC SORT ON A MIGRATORY THREAD ARCHITECTURE 730
Kaushik Velusamy; Thomas B. Rolinger; Janice McMahon

INFERENCE BENCHMARKING ON HPC SYSTEMS 737
Wesley Brewer; Greg Behm; Alan Scheinine; Ben Parsons; Wesley Emeneker; Robert P. Trevino

MULTISCALE DATA ANALYSIS USING BINNING, TENSOR DECOMPOSITIONS, AND BACKTRACKING 746
Dimitri Leggas; Thomas S. Henretty; James Ezick; Muthu Baskaran; Brendan Von Hofe; Grace Cimaszewski; M. Harper Langston; Richard Lethin

OPTIMISING AI TRAINING DEPLOYMENTS USING GRAPH COMPILERS AND CONTAINERS 753
Nina Mujkanovic; Karthee Sivalingam; Alfio Lazzaro

LEVERAGING LINEAR ALGEBRA TO COUNT AND ENUMERATE SIMPLE SUBGRAPHS 761
Vitaliy Gleyzer; Andrew J. Soszynski; Edward K. Kao
ON THE FEASIBILITY OF USING REDUCED-PRECISION TENSOR CORE OPERATIONS FOR GRAPH ANALYTICS 769
Jesun Sahariar Firoz; Ang Li; Jiajia Li; Kevin Barker
ENERGY-EFFICIENT ANALYSIS OF SYNCHROPHASOR DATA USING THE NVIDIA JETSON NANO 776
Suzanne J. Matthews; Aaron St. Leger
HIGH-THROUGHPUT IMAGE ALIGNMENT FOR CONNECTOMICS USING FRUGAL SNAP JUDGMENTS 783
Tim Kaler; Brian Wheatman; Sarah Wooders
CONSTRAINED-OPTIMIZATION APPROACH DELIVERS SUPERIOR CLASSICAL PERFORMANCE FOR GRAPH PARTITIONING VIA QUANTUM-READY METHOD 792
Uchenna Chukwu; Raouf Dridi; Jesse Berwald; Michael Booth; John Dawson; Deyung Le; Mark Wainger; Steven P. Reinhardt
MACHINE LEARNING ALGORITHM PERFORMANCE ON THE LUCATA COMPUTER 798
Paul L. Springer; Thomas Schibler; Géraud Krawezik; Jack Lightholder; Peter M. Kogge
INCREMENTAL STREAMING GRAPH PARTITIONING 805
Lisa Durbeck; Peter Athanas
Author Index