

2019 IEEE High Performance Extreme Computing Conference (HPEC 2019)

**Waltham, Massachusetts, USA
24 – 26 September 2019**



**IEEE Catalog Number: CFP19HPE-POD
ISBN: 978-1-7281-5021-5**

**Copyright © 2019 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:	CFP19HPE-POD
ISBN (Print-On-Demand):	978-1-7281-5021-5
ISBN (Online):	978-1-7281-5020-8
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

DEEP LEARNING-BASED NUCLEI SEGMENTATION OF CLEARED BRAIN TISSUE	1
<i>Pooya Khorrami ; Kevin Brady ; Mark Hernandez ; Lars Gjestebj ; Sara N. Burke ; Damon G. Lamb ; Matthew A. Melton ; Kevin J. Otto ; Laura J. Brattain</i>	
PROTOTYPE CONTAINER-BASED PLATFORM FOR EXTREME QUANTUM COMPUTING ALGORITHM DEVELOPMENT	3
<i>Patrick Dreher ; Madhuvanti Ramasami</i>	
LOW OVERHEAD INSTRUCTION LATENCY CHARACTERIZATION FOR NVIDIA GPGPUS	10
<i>Yehia Arafa ; Abdel-Hameed A. Badawy ; Gopinath Chennupati ; Nandakishore Santhi ; Stephan Eidenbenz</i>	
OPTIMIZING THE VISUALIZATION PIPELINE OF A 3-D MONITORING AND MANAGEMENT SYSTEM	18
<i>Rebecca Wild ; Matthew Hubbell ; Jeremy Kepner</i>	
SCALABLE SOLVERS FOR CONE COMPLEMENTARITY PROBLEMS IN FRICTIONAL MULTIBODY DYNAMICS	23
<i>Saibal De ; Eduardo Corona ; Paramsothy Jayakumar ; Shraavan Veerapaneni</i>	
ARTIFICIAL NEURAL NETWORK AND ACCELERATOR CO-DESIGN USING EVOLUTIONARY ALGORITHMS	30
<i>Philip Colangelo ; Oren Segal ; Alex Speicher ; Martin Margala</i>	
A SURVEY ON HARDWARE SECURITY TECHNIQUES TARGETING LOW-POWER SOC DESIGNS	38
<i>Alan Ehret ; Karen Gettings ; Bruce R. Jordan ; Michel A. Kinsky</i>	
A DATA-DRIVEN FRAMEWORK FOR UNCERTAINTY QUANTIFICATION OF A FLUIDIZED BED	46
<i>V M Krushnarao Kotteda ; Anitha Kommu ; Vinod Kumar</i>	
BLAST: BLOCKCHAIN-BASED TRUST MANAGEMENT IN SMART CITIES AND CONNECTED VEHICLES SETUP	53
<i>Farah Kandah ; Brennan Huber ; Amani Altarawneh ; Sai Medury ; Anthony Skjellum</i>	
QXSQA: GPGPU-ACCELERATED SIMULATED QUANTUM ANNEALER WITHIN A NON-LINEAR OPTIMIZATION AND BOLTZMANN SAMPLING FRAMEWORK	60
<i>Dan Padilha ; Serge Weinstock ; Mark Hodson</i>	
ON COMPUTING WITH DIAGONALLY STRUCTURED MATRICES	68
<i>Shahadat Hossain ; Mohammad Sakib Mahmud</i>	
EMERGING APPLICATIONS OF 3D INTEGRATION AND APPROXIMATE COMPUTING IN HIGH-PERFORMANCE COMPUTING SYSTEMS: UNIQUE SECURITY VULNERABILITIES	74
<i>Pruthvy Yellu ; Zhiming Zhang ; Mohammad Mezanur Rahman Monjur ; Ranuli Abeysinghe ; Qiaoyan Yu</i>	
COMET: A DISTRIBUTED METADATA SERVICE FOR FEDERATED CLOUD INFRASTRUCTURES	81
<i>Cong Wang ; Komal Thareja ; Michael Stealey ; Paul Ruth ; Ilya Baldin</i>	
LOSSLESS COMPRESSION OF INTERNAL FILES IN PARALLEL RESERVOIR SIMULATION	88
<i>Marcin Rogowski ; Suha N. Kayum ; Florian Mannuss</i>	
OPTIMAL RESOURCE ALLOCATION FOR PARALLEL RESERVOIR SIMULATION	91
<i>Suha N. Kayum ; Marcin Rogowski</i>	
EVALUATION OF THE IMBALANCE EVOLUTION IN PARALLEL RESERVOIR SIMULATION	95
<i>Marcin Rogowski ; Suha N. Kayum</i>	
FFTX FOR MICROMECHANICAL STRESS-STRAIN ANALYSIS	102
<i>Anuva Kulkarni ; Daniele G. Spampinato ; Franz Franchetti</i>	
SCALABLE LAZY-UPDATE MULTIGRID PRECONDITIONERS	104
<i>Majid Rasouli ; Vidhi Zala ; Robert M. Kirby ; Hari Sundar</i>	
GRAPH ALGORITHMS IN PGAS: CHAPEL AND UPC++	111
<i>Louis Jenkins ; Jesun Sahariar Firoz ; Marcin Zalewski ; Cliff Joslyn ; Mark Raugas</i>	
MULTI-SPECTRAL REUSE DISTANCE: DIVINING SPATIAL INFORMATION FROM TEMPORAL DATA	117
<i>Anthony M. Cabrera ; Roger D. Chamberlain ; Jonathan C. Beard</i>	
TOWARDS IMPROVING RATE-DISTORTION PERFORMANCE OF TRANSFORM-BASED LOSSY COMPRESSION FOR HPC DATASETS	125
<i>Jialing Zhang ; Aekyeung Moon ; Xiaoyan Zhuo ; Seung Woo Son</i>	

C TO D-WAVE: A HIGH-LEVEL C COMPILATION FRAMEWORK FOR QUANTUM ANNEALERS	132
<i>Mohamed W. Hassan ; Scott Pakin ; Wu-Chun Feng</i>	
SURVEY OF ATTACKS AND DEFENSES ON EDGE-DEPLOYED NEURAL NETWORKS	140
<i>Mihailo Isakov ; Vijay Gadepally ; Karen M. Gettings ; Michel A. Kinsy</i>	
AUXILIARY MAXIMUM LIKELIHOOD ESTIMATION FOR NOISY POINT CLOUD REGISTRATION	148
<i>Cole Campton ; Xiaobai Sun</i>	
INCREASING ACCURACY OF ITERATIVE REFINEMENT IN LIMITED FLOATING-POINT ARITHMETIC ON HALF-PRECISION ACCELERATORS	155
<i>Piotr Luszczek ; Ichitaro Yamazaki ; Jack Dongarra</i>	
CONCURRENT KATZ CENTRALITY FOR STREAMING GRAPHS	161
<i>Chunxing Yin ; Jason Riedy</i>	
TARGET-BASED RESOURCE ALLOCATION FOR DEEP LEARNING APPLICATIONS IN A MULTI-TENANCY SYSTEM	167
<i>Wenjia Zheng ; Yun Song ; Zihao Guo ; Yongchen Cui ; Suwen Gu ; Ying Mao ; Long Cheng</i>	
APPLICATION OF APPROXIMATE MATRIX MULTIPLICATION TO NEURAL NETWORKS AND DISTRIBUTED SLAM	174
<i>Brian Plancher ; Camelia D. Brumar ; Iulian Brumar ; Lillian Pentecost ; Saketh Rama ; David Brooks</i>	
AN FPGA DECISION TREE CLASSIFIER TO SUPERVISE A COMMUNICATION SOC	181
<i>Abdelrahman Elkanishy ; Derrick T. Rivera ; Abdel-Hameed A. Badawy ; Paul M. Furth ; Z. M. Saifullah ; Christopher P. Michael</i>	
HYPERSPARSE NEURAL NETWORK ANALYSIS OF LARGE-SCALE INTERNET TRAFFIC	187
<i>Jeremy Kepner ; Kenjiro Cho ; Kc Claffy ; Vijay Gadepally ; Peter Michaleas ; Lauren Milechin</i>	
INTRODUCING DYMONDS-AS-A-SERVICE (DYMAAS) FOR INTERNET OF THINGS	198
<i>Marija Ilic ; Rupamathi Jaddivada</i>	
A GPU IMPLEMENTATION OF THE SPARSE DEEP NEURAL NETWORK GRAPH CHALLENGE	207
<i>Mauro Bisson ; Massimiliano Fatica</i>	
FAST STOCHASTIC BLOCK PARTITIONING VIA SAMPLING	215
<i>Frank Wanye ; Vitaliy Gleyzer ; Wu-Chun Feng</i>	
H-INDEX: HASH-INDEXING FOR PARALLEL TRIANGLE COUNTING ON GPUS	222
<i>Santosh Pandey ; Xiaoye Sherry Li ; Aydin Buluc ; Jiejun Xu ; Hang Liu</i>	
SPARSE DEEP NEURAL NETWORK GRAPH CHALLENGE	229
<i>Jeremy Kepner ; Simon Alford ; Vijay Gadepally ; Michael Jones ; Lauren Milechin ; Ryan Robinett ; Sid Samsi</i>	
ACCELERATING DNN INFERENCE WITH GRAPHBLAS AND THE GPU	236
<i>Xiaoyun Wang ; Zhongyi Lin ; Carl Yang ; John D. Owens</i>	
FAST TRIANGLE COUNTING ON GPU	242
<i>Chuangyi Gui ; Long Zheng ; Pengcheng Yao ; Xiaofei Liao ; Hai Jin</i>	
MULTITHREADED LAYER-WISE TRAINING OF SPARSE DEEP NEURAL NETWORKS USING COMPRESSED SPARSE COLUMN	249
<i>Mohammad Hasanzadeh Mofrad ; Rami Melhem ; Yousuf Ahmad ; Mohammad Hammoud</i>	
PERFORMANCE OF TRAINING SPARSE DEEP NEURAL NETWORKS ON GPUS	255
<i>Jianzong Wang ; Zhangcheng Huang ; Lingwei Kong ; Jing Xiao ; Pengyu Wang ; Lu Zhang ; Chao Li</i>	
ACCELERATING SPARSE DEEP NEURAL NETWORKS ON FPGAS	260
<i>Sitao Huang ; Carl Pearson ; Rakesh Nagi ; Jinjun Xiong ; Deming Chen ; Wen-Mei Hwu</i>	
FAST BFS-BASED TRIANGLE COUNTING ON GPUS	267
<i>Leyuan Wang ; John D. Owens</i>	
WRITE QUICK, RUN FAST: SPARSE DEEP NEURAL NETWORK IN 20 MINUTES OF DEVELOPMENT TIME VIA SUITESPARSE:GRAPHBLAS	273
<i>Timothy A. Davis ; Mohsen Aznaveh ; Scott Kolodziej</i>	
UPDATE ON TRIANGLE COUNTING ON GPU	279
<i>Carl Pearson ; Mohammad Almasri ; Omer Anjum ; Vikram S. Mailthody ; Zaid Qureshi ; Rakesh Nagi ; Jinjun Xiong ; Wen-Mei Hwu</i>	
UPDATE ON K-TRUSS DECOMPOSITION ON GPU	286
<i>Mohammad Almasri ; Omer Anjum ; Carl Pearson ; Zaid Qureshi ; Vikram S. Mailthody ; Rakesh Nagi ; Jinjun Xiong ; Wen-Mei Hwu</i>	
DISTTC: HIGH PERFORMANCE DISTRIBUTED TRIANGLE COUNTING	293
<i>Loc Hoang ; Vishwesh Jatala ; Xuhao Chen ; Udit Agarwal ; Roshan Dathathri ; Gurbinder Gill ; Keshav Pingali</i>	
LINEAR ALGEBRA-BASED TRIANGLE COUNTING VIA FINE-GRAINED TASKING ON HETEROGENEOUS ENVIRONMENTS : (UPDATE ON STATIC GRAPH CHALLENGE)	300
<i>Abdurrahman Yasar ; Sivasankaran Rajamanickam ; Jonathan Berry ; Michael Wolf ; Jeffrey S. Young ; Ümit V. Çatalyürek</i>	

SCALABLE TRIANGLE COUNTING ON DISTRIBUTED-MEMORY SYSTEMS	304
<i>Seher Acer ; Abdurrahman Yasar ; Sivasankaran Rajamanickam ; Michael Wolf ; Ümit V. Catalyürek</i>	
EXPLORATION OF FINE-GRAINED PARALLELISM FOR LOAD BALANCING EAGER K-TRUSS ON GPU AND CPU	309
<i>Mark Blanco ; Tze Meng Low ; Kyungjoo Kim</i>	
SCALABLE INFERENCE FOR SPARSE DEEP NEURAL NETWORKS USING KOKKOS KERNELS	316
<i>J. Austin Ellis ; Sivasankaran Rajamanickam</i>	
ONE QUADRILLION TRIANGLES QUERIED ON ONE MILLION PROCESSORS	323
<i>Roger Pearce ; Trevor Steil ; Benjamin W. Priest ; Geoffrey Sanders</i>	
SCALING AND QUALITY OF MODULARITY OPTIMIZATION METHODS FOR GRAPH CLUSTERING	328
<i>Sayan Ghosh ; Mahantesh Halappanavar ; Antonino Tumeo ; Ananth Kalyanarainan</i>	
DISTRIBUTED DIRECTION-OPTIMIZING LABEL PROPAGATION FOR COMMUNITY DETECTION	334
<i>Xu Liu ; Jesun Sahariar Firoz ; Marcin Zalewski ; Mahantesh Halappanavar ; Kevin J. Barker ; Andrew Lumsdaine ; Assefaw H. Gebremedhin</i>	
IMPROVING PARALLELISM OF BREADTH FIRST SEARCH (BFS) ALGORITHM FOR ACCELERATED PERFORMANCE ON GPUS	340
<i>Hao Wen ; Wei Zhang</i>	
HETEROGENEOUS CACHE HIERARCHY MANAGEMENT FOR INTEGRATED CPU-GPU ARCHITECTURE	347
<i>Hao Wen ; Wei Zhang</i>	
OVERCOMING LIMITATIONS OF GPGPU-COMPUTING IN SCIENTIFIC APPLICATIONS	353
<i>Connor Kenyon ; Glenn Volkema ; Gaurav Khanna</i>	
IMPROVING SCHEDULING FOR IRREGULAR APPLICATIONS WITH LOGARITHMIC RADIX BINNING	362
<i>James Fox ; Alok Tripathy ; Oded Green</i>	
FAST AND SCALABLE DISTRIBUTED TENSOR DECOMPOSITIONS	369
<i>Muthu Baskaran ; Thomas Henretty ; James Ezick</i>	
AN INTERACTIVE LIDAR TO CAMERA CALIBRATION	376
<i>Yecheng Lyu ; Lin Bai ; Mahdi Elhousni ; Xinming Huang</i>	
SINGULARITY FOR MACHINE LEARNING APPLICATIONS - ANALYSIS OF PERFORMANCE IMPACT	382
<i>Bruce R. Jordan ; David Barrett ; David Burke ; Patrick Jardin ; Amelia Littrell ; Paul Monticciolo ; Michael Newey ; Jean Piou ; Kara Warner</i>	
PROACTIVE CYBER SITUATION AWARENESS VIA HIGH PERFORMANCE COMPUTING	388
<i>Allan Wollaber ; Jaime Peñna ; Benjamin Blease ; Leslie Shing ; Kenneth Alperin ; Serge Vilvovsky ; Pierre Trepagnier ; Neal Wagner ; Leslie Leonard</i>	
A PARALLEL SIMULATION APPROACH TO ACAS X DEVELOPMENT	395
<i>Adam Gjersvik ; Robert J. Moss</i>	
BREADTH-FIRST SEARCH ON DYNAMIC GRAPHS USING DYNAMIC PARALLELISM ON THE GPU	401
<i>Dominik Tödling ; Martin Winter ; Markus Steinberger</i>	
HARDWARE IP CLASSIFICATION THROUGH WEIGHTED CHARACTERISTICS	408
<i>Brendan McGeehan ; Flora Smith ; Thao Le ; Hunter Nauman ; Jia Di</i>	
MANY-TARGET, MANY-SENSOR SHIP TRACKING AND CLASSIFICATION	414
<i>Leonard Kosta ; John Irvine ; Laura Seaman ; Hongwei Xi</i>	
CYBER BASELINING: STATISTICAL PROPERTIES OF CYBER TIME SERIES AND THE SEARCH FOR STABILITY	421
<i>Alexia Schulz ; Ethan Aubin ; Pierre Trepagnier ; Allan Wollaber</i>	
LARGE SCALE ORGANIZATION AND INFERENCE OF AN IMAGERY DATASET FOR PUBLIC SAFETY	428
<i>Jeffrey Liu ; David Strohschein ; Siddharth Samsi ; Andrew Weinert</i>	
DEEP-LEARNING INFERENCE WITH HIGH-PERFORMANCE HARDWARE ACCELERATORS	434
<i>Luke Kljucaric ; Alan D. George</i>	
EFFICIENT IMPLEMENTATION OF SPARSE MATRIX-SPARSE VECTOR MULTIPLICATION FOR LARGE SCALE GRAPH ANALYTICS	441
<i>Mauricio J. Serrano</i>	

LOW POWER COMPUTING AND SIMULTANEOUS ELECTRO-OPTICAL/RADAR DATA PROCESSING USING IBM'S NS16E 16-CHIP NEUROMORPHIC HARDWARE	448
<i>Mark Barnell ; Courtney Raymond ; Daniel Brown ; Matthew Wilson ; Eric Cote</i>	
IP CORES FOR GRAPH KERNELS ON FPGAS	453
<i>Sanmukh R. Kuppannagari ; Rachit Rajat ; Rajgopal Kannan ; Aravind Dasu ; Viktor K. Prasanna</i>	
DESIGN AND IMPLEMENTATION OF KNOWLEDGE BASE FOR RUNTIME MANAGEMENT OF SOFTWARE DEFINED HARDWARE	460
<i>Hongkuan Zhou ; Ajitesh Srivastava ; Rajgopal Kannan ; Viktor Prasanna</i>	
STREAMING 1.9 BILLION HYPERSPARSE NETWORK UPDATES PER SECOND WITH D4M	467
<i>Jeremy Kepner ; Vijay Gadepally ; Lauren Milechin ; Siddharth Samsi ; William Arcand ; David Bestor ; William Bergeron ; Chansup Byun ; Matthew Hubbell ; Michael Houle ; Michael Jones ; Anne Klein ; Peter Michaleas ; Julie Mullen ; Andrew Prout ; Antonio Rosa ; Charles Yee ; Albert Reuther</i>	
OPTIMIZING XEON PHI FOR INTERACTIVE DATA ANALYSIS	473
<i>Chansup Byun ; Jeremy Kepner ; William Arcand ; David Bestor ; William Bergeron ; Matthew Hubbell ; Vijay Gadepally ; Michael Houle ; Michael Jones ; Anne Klein ; Lauren Milechin ; Peter Michaleas ; Julie Mullen ; Andrew Prout ; Antonio Rosa ; Siddharth Samsi ; Charles Yee ; Albert Reuther</i>	
IDPRISM: RAPID ANALYSIS OF FORENSIC DNA SAMPLES USING MPS SNP PROFILES	479
<i>Darrell O. Ricke ; James Watkins ; Philip Fremont-Smith ; Adam Michaleas</i>	
MULTISTART METHODS FOR QUANTUM APPROXIMATE OPTIMIZATION	484
<i>Ruslan Shayduln ; Ilya Safro ; Jeffrey Larson</i>	
APPLYING NEUROMORPHIC COMPUTING TO COMPRESSIVE SENSING	492
<i>Ronald Scrofano ; Douglas P. Enright ; George C. Valley</i>	
MESSAGE SCHEDULING FOR PERFORMANT, MANY-CORE BELIEF PROPAGATION	494
<i>Mark Van Der Merwe ; Vinu Joseph ; Ganesh Gopalakrishnan</i>	
SPACELAND EMBEDDING OF SPARSE STOCHASTIC GRAPHS	501
<i>Nikos Pitsianis ; Alexandros-Stavros Iliopoulos ; Dimitris Floros ; Xiaobai Sun</i>	
FPGA-ACCELERATED SPREADING FOR GLOBAL PLACEMENT	509
<i>Shounak Dhar ; Love Singhal ; Mahesh A. Iyer ; David Z. Pan</i>	
SECURING HPC USING FEDERATED AUTHENTICATION	516
<i>Andrew Prout ; William Arcand ; David Bestor ; Bill Bergeron ; Chansup Byun ; Vijay Gadepally ; Michael Houle ; Matthew Hubbell ; Michael Jones ; Anna Klein ; Peter Michaleas ; Lauren Milechin ; Julie Mullen ; Antonio Rosa ; Siddharth Samsi ; Charles Yee ; Albert Reuther ; Jeremy Kepner</i>	
SYNTHESIS OF HARDWARE SANDBOXES FOR TROJAN MITIGATION IN SYSTEMS ON CHIP	523
<i>Christophe Bobda ; Taylor Whitaker ; Joel Mandebi Mbongue ; Sujan Kumar Saha</i>	
MEXT: A FLOW FOR MULTIPROCESSOR EXPLORATION	529
<i>Christophe Bobda ; Harold Ishebabi ; Philipp Mahr ; Joel Mandebi Mbongue ; Sujan Kumar Saha</i>	
USING CONTAINER MIGRATION FOR HPC WORKLOADS RESILIENCE	536
<i>Mohamad Sindi ; John R. Williams</i>	
COMBINING TENSOR DECOMPOSITIONS AND GRAPH ANALYTICS TO PROVIDE CYBER SITUATIONAL AWARENESS AT HPC SCALE	546
<i>James Ezick ; Tom Henretty ; Muthu Baskaran ; Richard Lethin ; John Feo ; Tai-Ching Tuan ; Christopher Coley ; Leslie Leonard ; Rajeev Agrawal ; Ben Parsons ; William Glodek</i>	
EMBEDDED PROCESSOR-IN-MEMORY ARCHITECTURE FOR ACCELERATING ARITHMETIC OPERATIONS	553
<i>Richard Muri ; Paul Fortier</i>	
DISTRIBUTED DEEP LEARNING FOR PRECIPITATION NOWCASTING	560
<i>Siddharth Samsi ; Christopher J. Mattioli ; Mark S. Veillette</i>	
TAPIRXLA: EMBEDDING FORK-JOIN PARALLELISM INTO THE XLA COMPILER IN TENSORFLOW USING TAPIR	567
<i>Tao B. Scharndl ; Siddharth Samsi</i>	
DEPLOYING AI FRAMEWORKS ON SECURE HPC SYSTEMS WITH CONTAINERS	575
<i>David Brayford ; Sofia Vallecorsa ; Atanas Atanasov ; Fabio Baruffa ; Walter Riviera</i>	
COMBINATORIAL MULTIGRID: ADVANCED PRECONDITIONERS FOR ILL-CONDITIONED LINEAR SYSTEMS	581
<i>M. Harper Langston ; Mitchell Tong Harris ; Pierre-David Letourneau ; Richard Lethin ; James Ezick</i>	
SKIP THE INTERSECTION: QUICKLY COUNTING COMMON NEIGHBORS ON SHARED-MEMORY SYSTEMS	588
<i>Xiaojing An ; Kasimir Gabert ; James Fox ; Oded Green ; David A. Bader</i>	
FAST LARGE-SCALE ALGORITHM FOR ELECTROMAGNETIC WAVE PROPAGATION IN 3D MEDIA	595
<i>Mitchell Tong Harris ; M. Harper Langston ; Pierre-David Letourneau ; George Papanicolaou ; James Ezick ; Richard Lethin</i>	

AUTOMATIC PARALLELIZATION TO ASYNCHRONOUS TASK-BASED RUNTIMES THROUGH A GENERIC RUNTIME LAYER	602
<i>Charles Jin ; Muthu Baskaran ; Benoit Meister ; Jonathan Springer</i>	
PROGRESSIVE OPTIMIZATION OF BATCHED LU FACTORIZATION ON GPUS	613
<i>Ahmad Abdelfattah ; Stanimire Tomov ; Jack Dongarra</i>	
AN EFFICIENT AND COMPOSABLE PARALLEL TASK PROGRAMMING LIBRARY	619
<i>Chun-Xun Lin ; Tsung-Wei Huang ; Guannan Guo ; Martin D. F. Wong</i>	
EXPLORING THE EFFICIENCY OF OPENCL PIPE FOR HIDING MEMORY LATENCY ON CLOUD FPGAS	626
<i>Arnab A Purkayastha ; Sai Raghavendran ; Jhanani Thiagarajan ; Hamed Tabkhi</i>	
TRAINING BEHAVIOR OF SPARSE NEURAL NETWORK TOPOLOGIES	633
<i>Simon Alford ; Ryan Robinett ; Lauren Milechin ; Jeremy Kepner</i>	
EMBEDDED GPU CLUSTER COMPUTING FRAMEWORK FOR INFERENCE OF CONVOLUTIONAL NEURAL NETWORKS	639
<i>Evan Kain ; Diego Wildenstein ; Andrew C. Pineda</i>	
GARBLED CIRCUITS IN THE CLOUD USING FPGA ENABLED NODES	646
<i>Kai Huang ; Mehmet Gungor ; Xin Fang ; Stratis Ioannidis ; Miriam Leeser</i>	
A NOVEL DESIGN OF ADAPTIVE AND HIERARCHICAL CONVOLUTIONAL NEURAL NETWORKS USING PARTIAL RECONFIGURATION ON FPGA	652
<i>Mohammad Farhadi ; Mehdi Ghasemi ; Yezhou Yang</i>	
LARGE SCALE PARALLELIZATION USING FILE-BASED COMMUNICATIONS	659
<i>Chansup Byun ; Jeremy Kepner ; William Arcand ; David Bestor ; Bill Bergeron ; Vijay Gadepally ; Michael Houle ; Matthew Hubbell ; Michael Jones ; Anna Klein ; Peter Michaleas ; Julie Mullen ; Andrew Prout ; Antonio Rosa ; Siddharth Samsi ; Charles Yee ; Albert Reuther</i>	
ECG FEATURE PROCESSING PERFORMANCE ACCELERATION ON SLURM COMPUTE SYSTEMS	666
<i>Michael Nolan ; Mark Hernandez ; Philip Fremont-Smith ; Albert Swiston ; Kajal Claypool</i>	
SURVEY AND BENCHMARKING OF MACHINE LEARNING ACCELERATORS	670
<i>Albert Reuther ; Peter Michaleas ; Michael Jones ; Vijay Gadepally ; Siddharth Samsi ; Jeremy Kepner</i>	
Author Index	