

SC20: International Conference for High Performance Computing, Networking, Storage and Analysis

**Virtual Event
9-19 November 2020**

Pages 1-707



**IEEE Catalog Number: CFP20SUP-POD
ISBN: 978-1-7281-9999-3**

**Copyright © 2020 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:	CFP20SUP-POD
ISBN (Print-On-Demand):	978-1-7281-9999-3
ISBN (Online):	978-1-7281-9998-6

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

A 1024-MEMBER ENSEMBLE DATA ASSIMILATION WITH 3.5-KM MESH GLOBAL WEATHER SIMULATIONS	1
<i>Hisashi Yashiro; Koji Terasaki; Yuta Kawai; Shuhei Kudo; Takemasa Miyoshi; Toshiyuki Imamura; Kazuo Minami; Hikaru Inoue; Tatsuo Nishiki; Takayuki Saji; Masaki Satoh; Hirofumi Tomita</i>	
PROCESSING FULL-SCALE SQUARE KILOMETRE ARRAY DATA ON THE SUMMIT SUPERCOMPUTER	11
<i>Ruonan Wang; Rodrigo Tobar; Markus Dolensky; Tao An; Andreas Wicenec; Chen Wu; Fred Dulwich; Norbert Podhorszki; Valentine Anantharaj; Eric Suchyta; Baoqiang Lao; Scott Klasky</i>	
TOWARD REALIZATION OF NUMERICAL TOWING-TANK TESTS BY WALL-RESOLVED LARGE EDDY SIMULATION BASED ON 32 BILLION GRID FINITE-ELEMENT COMPUTATION	23
<i>Chisachi Kato; Yoshinobu Yamade; Katsuhiro Nagano; Kiyoshi Kumahata; Kazuo Minami; Tatsuo Nishikawa</i>	
ACCELERATING LARGE-SCALE EXCITED-STATE GW CALCULATIONS ON LEADERSHIP HPC SYSTEMS	36
<i>Mauro Del Ben; Charlene Yang; Zhenglu Li; Felipe H. Da Jornada; Steven G. Louie; Jack Deslippe</i>	
PUSHING THE LIMIT OF MOLECULAR DYNAMICS WITH AB INITIO ACCURACY TO 100 MILLION ATOMS WITH MACHINE LEARNING	47
<i>Weile Jia; Han Wang; Mohan Chen; Denghui Lu; Lin Lin; Roberto Car; E. Weinan; Linfeng Zhang</i>	
SCALABLE KNOWLEDGE GRAPH ANALYTICS AT 136 PETAFL0P/S	61
<i>Ramakrishnan Kannan; Piyush Sao; Hao Lu; Drahomira Herrmannova; Vijay Thakkar; Robert Patton; Richard Vuduc; Thomas Potok</i>	
A PARALLEL FRAMEWORK FOR CONSTRAINT-BASED BAYESIAN NETWORK LEARNING VIA MARKOV BLANKET DISCOVERY	74
<i>Ankit Srivastava; Sriram P. Chockalingam; Srinivas Aluru</i>	
RECURRENT NEURAL NETWORK ARCHITECTURE SEARCH FOR GEOPHYSICAL EMULATION	89
<i>Romit Maulik; Romain Egele; Bethany Lusch; Prasanna Balaprakash</i>	
MESHFREEFLOWNET: A PHYSICS-CONSTRAINED DEEP CONTINUOUS SPACE-TIME SUPER-RESOLUTION FRAMEWORK	103
<i>Chiyu Jiang; Soheil Esmaeilzadeh; Kamyar Azzadenesheli; Karthik Kashinath; Mustafa Mustafa; Hamdi A. Tchelepi; Philip Marcus; Prabhat; Anima Anandkumar</i>	
IMPROVING ALL-TO-MANY PERSONALIZED COMMUNICATION IN TWO-PHASE I/O	118
<i>Qiao Kang; Robert Ross; Robert Latham; Sunwoo Lee; Ankit Agrawal; Alok Choudhary; Wei-Keng Liao</i>	
TAMING I/O VARIATION ON QOS-LESS HPC STORAGE: WHAT CAN APPLICATIONS DO?	131
<i>Zhenbo Qiao; Qing Liu; Norbert Podhorszki; Scott Klasky; Jiayang Chen</i>	
BORA: A BAG OPTIMIZER FOR ROBOTIC ANALYSIS	144
<i>Jian Zhang; Tao Xie; Yuzhuo Jing; Yanjie Song; Guanzhou Hu; Si Chen; Shu Yin</i>	
DENSITY MATRIX QUANTUM CIRCUIT SIMULATION VIA THE BSP MACHINE ON MODERN GPU CLUSTERS	159
<i>Ang Li; Omer Subasi; Xiu Yang; Sriram Krishnamoorthy</i>	
EFFICIENT 2D TENSOR NETWORK SIMULATION OF QUANTUM SYSTEMS	174
<i>Yuchen Pang; Tianyi Hao; Annika Dugad; Yiqing Zhou; Edgar Solomonik</i>	
VERITAS: ACCURATELY ESTIMATING THE CORRECT OUTPUT ON NOISY INTERMEDIATE-SCALE QUANTUM COMPUTERS	188
<i>Tirthak Patel; Devesh Tiwari</i>	
ACCELERATING SPARSE DNN MODELS WITHOUT HARDWARE-SUPPORT VIA TILE-WISE SPARSITY	204
<i>Cong Guo; Bo Yang Hsueh; Jingwen Leng; Yuxian Qiu; Yue Guan; Zehuan Wang; Xiaoying Jia; Xipeng Li; Minyi Guo; Yuhao Zhu</i>	
SPARSE GPU KERNELS FOR DEEP LEARNING	219
<i>Trevor Gale; Matei Zaharia; Cliff Young; Erich Elsen</i>	
SPTFS: SPARSE TENSOR FORMAT SELECTION FOR MTTKRP VIA DEEP LEARNING	233
<i>Qingxiao Sun; Yi Liu; Ming Dun; Hailong Yang; Zhongzhi Luan; Lin Gan; Guangwen Yang; Depei Qian</i>	
SCALING DISTRIBUTED DEEP LEARNING WORKLOADS BEYOND THE MEMORY CAPACITY WITH KARMA	247
<i>Mohamed Wahib; Haoyu Zhang; Truong Thao Nguyen; Aleksandr Drozd; Jens Domke; Lingqi Zhang; Ryousei Takano; Satoshi Matsuo</i>	

ZERO: MEMORY OPTIMIZATIONS TOWARD TRAINING TRILLION PARAMETER MODELS	262
<i>Samyam Rajbhandari; Jeff Rasley; Olatunji Ruwase; Yuxiong He</i>	
KRAKEN: MEMORY-EFFICIENT CONTINUAL LEARNING FOR LARGE-SCALE REAL-TIME RECOMMENDATIONS	278
<i>Minhui Xie; Kai Ren; Youyou Lu; Guangxu Yang; Qingxing Xu; Bihai Wu; Jiazhen Lin; Hongbo Ao; Wanhong Xu; Jiwu Shu</i>	
CELL-LIST BASED MOLECULAR DYNAMICS ON MANY-CORE PROCESSORS: A CASE STUDY ON SUNWAY TAIHULIGHT SUPERCOMPUTER	295
<i>Xiaohui Duan; Ping Gao; Meng Zhang; Tingjian Zhang; Hongsong Meng; Yuxuan Li; Bertil Schmidt; Haohuan Fu; Lin Gan; Wei Xue; Weiguo Liu; Guangwen Yang</i>	
EVALUATION OF A MINIMALLY SYNCHRONOUS ALGORITHM FOR 2:1 OCTREE BALANCE	307
<i>Hansol Suh; Tobin Isaac</i>	
DISTRIBUTED-MEMORY DMRG VIA SPARSE AND DENSE PARALLEL TENSOR CONTRACTIONS	319
<i>Ryan Levy; Edgar Solomonik; Bryan K. Clark</i>	
TAGO: RETHINKING ROUTING DESIGN IN HIGH PERFORMANCE RECONFIGURABLE NETWORKS	333
<i>Min Yee Teh; Yu-Han Hung; George Michelogiannakis; Shijia Yan; Madeleine Glick; John Shalf; Keren Bergman</i>	
ARCHITECTURE AND PERFORMANCE STUDIES OF 3D-HYPER-FLEX-LION FOR RECONFIGURABLE ALL-TO-ALL HPC NETWORKS	349
<i>Gengchen Liu; Roberto Proietti; Marjan Fariborz; Pouya Fotouhi; Xian Xiao; S. J. Ben Yoo</i>	
FATPATHS: ROUTING IN SUPERCOMPUTERS AND DATA CENTERS WHEN SHORTEST PATHS FALL SHORT	365
<i>Maciej Besta; Marcel Schneider; Marek Konieczny; Karolina Cynk; Erik Henriksson; Salvatore Di Girolamo; Ankit Singla; Torsten Hoefler</i>	
SCALANA: AUTOMATING SCALING LOSS DETECTION WITH GRAPH ANALYSIS	383
<i>Yuyang Jin; Haojie Wang; Teng Yu; Xiongchao Tang; Torsten Hoefler; Xu Liu; Jidong Zhai</i>	
ZEROSPY: EXPLORING SOFTWARE INEFFICIENCY WITH REDUNDANT ZEROS	397
<i>Xin You; Hailong Yang; Zhongzhi Luan; Depei Qian; Xu Liu</i>	
DRCCTPROF: A FINE-GRAINED CALL PATH PROFILER FOR ARM-BASED CLUSTERS	411
<i>Qidong Zhao; Xu Liu; Milind Chabbi</i>	
RLSCHEDULER: AN AUTOMATED HPC BATCH JOB SCHEDULER USING REINFORCEMENT LEARNING	427
<i>Di Zhang; Dong Dai; Youbiao He; Forrest Sheng Bao; Bing Xie</i>	
ALITA: COMPREHENSIVE PERFORMANCE ISOLATION THROUGH BIAS RESOURCE MANAGEMENT FOR PUBLIC CLOUDS	442
<i>Quan Chen; Shuai Xue; Shang Zhao; Shanpei Chen; Yihao Wu; Yu Xu; Zhuo Song; Tao Ma; Yong Yang; Minyi Guo</i>	
HPC I/O THROUGHPUT BOTTLENECK ANALYSIS WITH EXPLAINABLE LOCAL MODELS	455
<i>Mihailo Isakov; Eliakin Del Rosario; Sandeep Madireddy; Prasanna Balaprakash; Philip Carns; Robert B. Ross; Michel A. Kinsy</i>	
A HIERARCHICAL AND LOAD-AWARE DESIGN FOR LARGE MESSAGE NEIGHBORHOOD COLLECTIVES	468
<i>S. Mahdieh Ghazimirsaeed; Qinghua Zhou; Amit Ruhela; Mohammadreza Bayatpour; Hari Subramoni; Dhableswar K. Dk Panda</i>	
AN IN-DEPTH ANALYSIS OF THE SLINGSHOT INTERCONNECT	481
<i>Daniele De Sensi; Salvatore Di Girolamo; Kim H. McMahon; Duncan Roweth; Torsten Hoefler</i>	
CAB-MPI: EXPLORING INTERPROCESS WORK-STEALING TOWARDS BALANCED MPI COMMUNICATION	495
<i>Kaiming Ouyang; Min Si; Atsushi Hori; Zizhong Chen; Pavan Balaji</i>	
PETASCALE XCT: 3D IMAGE RECONSTRUCTION WITH HIERARCHICAL COMMUNICATIONS ON MULTI-GPU NODES	510
<i>Mert Hidayetoglu; Tekin Bicer; Simon Garcia De Gonzalo; Bin Ren; Vincent De Andrade; Doga Gursoy; Raj Kettimuthu; Ian T. Foster; Wen-Mei W. Hwu</i>	
MULTI-NODE MULTI-GPU DIFEOMORPHIC IMAGE REGISTRATION FOR LARGE-SCALE IMAGING PROBLEMS	523
<i>Malte Brunn; Naveen Himthani; George Biros; Miriam Mehl; Andreas Mang</i>	
SEGALIGN: A SCALABLE GPU-BASED WHOLE GENOME ALIGNER	540
<i>Sneha D. Goenka; Yatish Turakhia; Benedict Paten; Mark Horowitz</i>	
TOSS-2020: A COMMODITY SOFTWARE STACK FOR HPC	553
<i>Edgar A. León; Trent D'Hooge; Nathan Hanford; Ian Karlin; Ramesh Pankajakshan; Jim Foraker; Chris Chamberau; Matthew L. Leininger</i>	

GPU LIFETIMES ON TITAN SUPERCOMPUTER: SURVIVAL ANALYSIS AND RELIABILITY	568
<i>George Ostrouchov; Don Maxwell; Rizwan A. Ashraf; Christian Engelmann; Mallikarjun Shankar; James H. Rogers</i>	
IRIS: ALLOCATION BANKING AND IDENTITY AND ACCESS MANAGEMENT FOR THE EXASCALE ERA	582
<i>Gabor Torok; Mark R. Day; Rebecca J. Hartman-Baker; Cory Snively</i>	
OPTIMIZING DEEP LEARNING RECOMMENDER SYSTEMS TRAINING ON CPU CLUSTER ARCHITECTURES	593
<i>Dhiraj Kalamkar; Evangelos Georganas; Sudarshan Srinivasan; Jianping Chen; Mikhail Shiryayev; Alexander Heinecke</i>	
HERRING: RETHINKING THE PARAMETER SERVER AT SCALE FOR THE CLOUD	608
<i>Indu Thangakrishnan; Derya Cavdar; Can Karakus; Piyush Ghai; Yauheni Selivonchyk; Cory Pruce</i>	
GEMS: GPU-ENABLED MEMORY-AWARE MODEL-PARALLELISM SYSTEM FOR DISTRIBUTED DNN TRAINING	621
<i>Arpan Jain; Ammar Ahmad Awan; Asmaa M. Aljuhani; Jahanzeb Maqbool Hashmi; Quentin G. Anthony; Hari Subramoni; Dhableswar K. Panda; Raghu Machiraju; Anil Parwani</i>	
EXPERIMENTAL EVALUATION OF NISQ QUANTUM COMPUTERS: ERROR MEASUREMENT, CHARACTERIZATION, AND IMPLICATIONS	636
<i>Tirthak Patel; Abhay Potharaju; Baolin Li; Rohan Basu Roy; Devesh Tiwari</i>	
CO-DESIGN FOR A64FX MANYCORE PROCESSOR AND "FUGAKU"	651
<i>Mitsuhsa Sato; Yutaka Ishikawa; Hirofumi Tomita; Yuetsu Kodama; Tetsuya Odajima; Miwako Tsuji; Hisashi Yashiro; Masaki Aoki; Naoyuki Shida; Ikuo Miyoshi; Kouichi Hirai; Atsushi Furuya; Akira Asato; Kuniki Morita; Toshiyuki Shimizu</i>	
CHRONICLES OF ASTRA: CHALLENGES AND LESSONS FROM THE FIRST PETASCALE ARM SUPERCOMPUTER	666
<i>Kevin Pedretti; Andrew J. Younge; Simon D. Hammond; James H. Laros; Matthew L. Curry; Michael J. Aguilar; Robert J. Hoekstra; Ron Brightwell</i>	
PLINER: ISOLATING LINES OF FLOATING-POINT CODE FOR COMPILER-INDUCED VARIABILITY	680
<i>Hui Guo; Ignacio Laguna; Cindy Rubio-González</i>	
TUNING FLOATING-POINT PRECISION USING DYNAMIC PROGRAM INFORMATION AND TEMPORAL LOCALITY	694
<i>Hugo Brunie; Costin Iancu; Khaled Z. Ibrahim; Philip Brisk; Brandon Cook</i>	
SCALABLE YET RIGOROUS FLOATING-POINT ERROR ANALYSIS	708
<i>Arnab Das; Ian Briggs; Ganesh Gopalakrishnan; Sriram Krishnamoorthy; Pavel Panchekha</i>	
RDMP-KV: DESIGNING REMOTE DIRECT MEMORY PERSISTENCE BASED KEY-VALUE STORES WITH PMEM	722
<i>Tianxi Li; Dipti Shankar; Shashank Gugnani; Xiaoyi Lu</i>	
COMPILER-BASED TIMING FOR EXTREMELY FINE-GRAIN PREEMPTIVE PARALLELISM	736
<i>Souradip Ghosh; Michael Cuevas; Simone Campanoni; Peter Dinda</i>	
OMPRACER: A SCALABLE AND PRECISE STATIC RACE DETECTOR FOR OPENMP PROGRAMS	751
<i>Bradley Swain; Yanze Li; Peiming Liu; Ignacio Laguna; Giorgis Georgakoudis; Jeff Huang</i>	
PREEMPT: SCALABLE EPIDEMIC INTERVENTIONS USING SUBMODULAR OPTIMIZATION ON MULTI-GPU SYSTEMS	765
<i>Marco Minutoli; Prathyush Sambaturu; Mahantesh Halappanavar; Antonino Tumeo; Ananth Kalyananaraman; Anil Vullikanti</i>	
C-SAW: A FRAMEWORK FOR GRAPH SAMPLING AND RANDOM WALK ON GPUS	780
<i>Santosh Pandey; Lingda Li; Adolffy Hoisie; Xiaoye S. Li; Hang Liu</i>	
NEWTON-ADMM: A DISTRIBUTED GPU-ACCELERATED OPTIMIZER FOR MULTICLASS CLASSIFICATION PROBLEMS	795
<i>Chih-Hao Fang; Sudhir B. Kylasa; Fred Roosta; Michael W. Mahoney; Ananth Grama</i>	
FAST STENCIL-CODE COMPUTATION ON A WAFER-SCALE PROCESSOR	807
<i>Kamil Rocki; Dirk Van Essendelft; Ilya Sharapov; Robert Schreiber; Michael Morrison; Vladimir Kibardin; Andrey Portnoy; Jean Francois Dietiker; Madhava Syamlal; Michael James</i>	
FBLAS: STREAMING LINEAR ALGEBRA ON FPGA	821
<i>Tiziano De Matteis; Johannes De Fine Licht; Torsten Hoefler</i>	
MASSIVE PARALLELIZATION FOR FINDING SHORTEST LATTICE VECTORS BASED ON UBIQUITY GENERATOR FRAMEWORK	834
<i>Nariaki Tateiwa; Yuji Shinano; Satoshi Nakamura; Akihiro Yoshida; Shizuo Kaji; Masaya Yasuda; Katsuki Fujisawa</i>	

COST-AWARE PREDICTION OF UNCORRECTED DRAM ERRORS IN THE FIELD	849
<i>Isaac Boixaderas; Darko Zivanovic; Sergi Moré; Javier Bartolome; David Vicente; Marc Casas; Paul M. Carpenter; Petar Radojkovic; Eduard Ayguadé</i>	
TASK BENCH: A PARAMETERIZED BENCHMARK FOR EVALUATING PARALLEL RUNTIME PERFORMANCE	864
<i>Elliott Slaughter; Wei Wu; Yuankun Fu; Legend Brandenburg; Nicolai Garcia; Wilhem Kautz; Emily Marx; Kaleb S. Morris; Qinglei Cao; George Bosilca; Seema Mirchandaney; Wonchan Leek; Sean Treichlerk; Patrick McCormick; Alex Aiken</i>	
SMART-PGSIM: USING NEURAL NETWORK TO ACCELERATE AC-OPF POWER GRID SIMULATION	879
<i>Wenqian Dong; Zhen Xie; Gokcen Kestor; Dong Li</i>	
SEFEE: LIGHTWEIGHT STORAGE ERROR FORECASTING IN LARGE-SCALE ENTERPRISE STORAGE SYSTEMS	894
<i>Amirhessam Yazdi; Xing Lin; Lei Yang; Feng Yan</i>	
LIVE FORENSICS FOR HPC SYSTEMS: A CASE STUDY ON DISTRIBUTED STORAGE SYSTEMS	908
<i>Saurabh Jha; Shengkun Cui; Subho S. Banerjee; Tianyin Xu; Jeremy Enos; Mike Showerman; Zbigniew T. Kalbarczyk; Ravishankar K. Iyer</i>	
INEC: FAST AND COHERENT IN-NETWORK ERASURE CODING	924
<i>Haiyang Shi; Xiaoyi Lu</i>	
WAITING GAME: OPTIMALLY PROVISIONING FIXED RESOURCES FOR CLOUD-ENABLED SCHEDULERS	941
<i>Pradeep Ambati; Noman Bashir; David Irwin; Prashant Shenoy</i>	
METIS: LEARNING TO SCHEDULE LONG-RUNNING APPLICATIONS IN SHARED CONTAINER CLUSTERS AT SCALE	955
<i>Luping Wang; Qizhen Weng; Wei Wang; Chen Chen; Bo Li</i>	
BATCH: MACHINE LEARNING INFERENCE SERVING ON SERVERLESS PLATFORMS WITH ADAPTIVE BATCHING	972
<i>Ahsan Ali; Riccardo Pinciroli; Feng Yan; Evgenia Smirni</i>	
REDUCING COMMUNICATION IN GRAPH NEURAL NETWORK TRAINING	987
<i>Alok Tripathy; Katherine Yelick; Aydin Buluç</i>	
FEATGRAPH: A FLEXIBLE AND EFFICIENT BACKEND FOR GRAPH NEURAL NETWORK SYSTEMS	1001
<i>Yuwei Hu; Zihao Ye; Minjie Wang; Jiali Yu; Da Zheng; Mu Li; Zheng Zhang; Zhiru Zhang; Yida Wang</i>	
GE-SPMM: GENERAL-PURPOSE SPARSE MATRIX-MATRIX MULTIPLICATION ON GPUS FOR GRAPH NEURAL NETWORKS	1014
<i>Guyue Huang; Guohao Dai; Yu Wang; Huazhong Yang</i>	
ALIAS-FREE, MATRIX-FREE, AND QUADRATURE-FREE DISCONTINUOUS GALERKIN ALGORITHMS FOR (PLASMA) KINETIC EQUATIONS	1026
<i>Ammar Hakim; James Juno</i>	
DISTRIBUTED-MEMORY PARALLEL SYMMETRIC NONNEGATIVE MATRIX FACTORIZATION	1041
<i>Srinivas Eswar; Koby Hayashi; Grey Ballard; Ramakrishnan Kannan; Richard Vuduc; Haesun Park</i>	
DISTRIBUTED MANY-TO-MANY PROTEIN SEQUENCE ALIGNMENT USING SPARSE MATRICES	1055
<i>Oguz Selvitopi; Saliya Ekanayake; Giulia Guidi; Georgios A. Pavlopoulos; Ariful Azad; Aydin Buluç</i>	
RUNTIME-GUIDED ECC PROTECTION USING ONLINE ESTIMATION OF MEMORY VULNERABILITY	1069
<i>Luc Jaulmes; Miquel Moretó; Mateo Valero; Mattan Erez; Marc Casas</i>	
CRAC: CHECKPOINT-RESTART ARCHITECTURE FOR CUDA WITH STREAMS AND UVM	1083
<i>Twinkle Jain; Gene Cooperman</i>	
ANT-MAN: TOWARDS AGILE POWER MANAGEMENT IN THE MICROSERVICE ERA	1098
<i>Xiaofeng Hou; Chao Li; Jiacheng Liu; Lu Zhang; Yang Hu; Minyi Guo</i>	
SCALABLE HETEROGENEOUS EXECUTION OF A COUPLED-CLUSTER MODEL WITH PERTURBATIVE TRIPLES	1112
<i>Jinsung Kim; Ajay Panyala; Bo Peng; Karol Kowalski; P. Sadayappan; Sriram Krishnamoorthy</i>	
A SUBMATRIX-BASED METHOD FOR APPROXIMATE MATRIX FUNCTION EVALUATION IN THE QUANTUM CHEMISTRY CODE CP2K	1127
<i>Michael Lass; Robert Schade; Thomas D. Kühne; Christian Plessl</i>	
SCALING THE HARTREE-FOCK MATRIX BUILD ON SUMMIT	1141
<i>Giuseppe M. J. Barca; David L. Poole; Jorge L. Galvez Vallejo; Melisa Alkan; Colleen Bertoni; Alistair P. Rendell; Mark S. Gordon</i>	

MOHA: A COMPOSABLE SYSTEM FOR EFFICIENT IN-SITU ANALYTICS ON HETEROGENEOUS HPC SYSTEMS	1155
<i>Haoyuan Xing; Gagan Agrawal; Rajiv Ramnath</i>	
FORESIGHT: ANALYSIS THAT MATTERS FOR DATA REDUCTION	1171
<i>Pascal Grosset; Christopher M. Biber; Jesus Pulido; Arvind T. Mohan; Ayan Biswas; John Patchett; Terece L. Turton; David H. Rogers; Daniel Livescu; James Ahrens</i>	
JOB CHARACTERISTICS ON LARGE-SCALE SYSTEMS: LONG-TERM ANALYSIS, QUANTIFICATION, AND IMPLICATIONS	1186
<i>Tirthak Patel; Zhengchun Liu; Raj Kettimuthu; Paul Rich; William Allcock; Devesh Tiwari</i>	
PENCIL: A PIPELINED ALGORITHM FOR DISTRIBUTED STENCILS	1203
<i>Hengjie Wang; Aparna Chandramowlishwaran</i>	
SPEEDING UP SPMV FOR POWER-LAW GRAPH ANALYTICS BY ENHANCING LOCALITY & VECTORIZATION	1219
<i>Serif Yesil; Azin Heidarshenas; Adam Morrison; Josep Torrellas</i>	
EFFICIENT TILED SPARSE MATRIX MULTIPLICATION THROUGH MATRIX SIGNATURES	1234
<i>Sireyya Emre Kurt; Aravind Sukumaran-Rajam; Fabrice Rastello; P. Sadayappan</i>	
GPU-TRIDENT: EFFICIENT MODELING OF ERROR PROPAGATION IN GPU PROGRAMS	1248
<i>Abdul Rehman Anwer; Guanpeng Li; Karthik Pattabiraman; Michael Sullivan; Timothy Tsai; Siva Kumar Sastry Hari</i>	
GVPROF: A VALUE PROFILER FOR GPU-BASED CLUSTERS	1263
<i>Keren Zhou; Yueming Hao; John Mellor-Crummey; Xiaozhu Meng; Xu Liu</i>	
AN EFFICIENT AND NON-INTRUSIVE GPU SCHEDULING FRAMEWORK FOR DEEP LEARNING TRAINING SYSTEMS	1279
<i>Shaoqi Wang; Oscar J. Gonzalez; Xiaobo Zhou; Thomas Williams; Brian D. Friedman; Martin Havemann; Thomas Woo</i>	
PREPARING NUCLEAR ASTROPHYSICS FOR EXASCALE	1292
<i>Max P. Katz; Ann Almgren; Maria Barrios Sazo; Kiran Eiden; Kevin Gott; Alice Harpole; Jean M. Sexton; Don E. Willcox; Weiqun Zhang; Michael Zingale</i>	
A PERFORMANCE-PORTABLE NONHYDROSTATIC ATMOSPHERIC DYCORE FOR THE ENERGY EXASCALE EARTH SYSTEM MODEL RUNNING AT CLOUD-RESOLVING RESOLUTIONS	1304
<i>Luca Bertagna; Oksana Guba; Mark A. Taylor; James G. Foucar; Jeff Larkin; Andrew M. Bradley; Sivasankaran Rajamanickam; Andrew G. Salinger</i>	
ACCELERATION OF FUSION PLASMA TURBULENCE SIMULATIONS USING THE MIXED-PRECISION COMMUNICATION-AVOIDING KRYLOV METHOD	1318
<i>Yasuhiro Idomura; Takuya Ina; Yussuf Ali; Toshiyuki Imamura</i>	
CONVOLUTIONAL NEURAL NETWORK TRAINING WITH DISTRIBUTED K-FAC	1331
<i>J. Gregory Pauloski; Zhao Zhang; Lei Huang; Weijia Xu; Ian T. Foster</i>	
BIQGEMM: MATRIX MULTIPLICATION WITH LOOKUP TABLE FOR BINARY-CODING-BASED QUANTIZED DNNS	1343
<i>Yongkweon Jeon; Baeseong Park; Se Jung Kwon; Byeongwook Kim; Jeongin Yun; Dongsoo Lee</i>	
TERM QUANTIZATION: FURTHERING QUANTIZATION AT RUN TIME	1357
<i>H. T. Kung; Bradley McDanel; Sai Qian Zhang</i>	
COMPILING GENERALIZED HISTOGRAMS FOR GPU	1373
<i>Troels Henriksen; Sune Hellfritsch; Ponnuswamy Sadayappan; Cosmin Oancea</i>	
CCAMP: AN INTEGRATED TRANSLATION AND OPTIMIZATION FRAMEWORK FOR OPENACC AND OPENMP	1387
<i>Jacob Lambert; Seyong Lee; Jeffrey S. Vetter; Allen D. Malony</i>	
HIGH-PERFORMANCE PARALLEL GRAPH COLORING WITH STRONG GUARANTEES ON WORK, DEPTH, AND QUALITY	1401
<i>Maciej Besta; Armon Carigiet; Kacper Janda; Zur Vonarburg-Shmaria; Lukas Gianinazzi; Torsten Hoeftler</i>	
GRAPHPI: HIGH PERFORMANCE GRAPH PATTERN MATCHING THROUGH EFFECTIVE REDUNDANCY ELIMINATION	1418
<i>Tianhui Shi; Mingshu Zhai; Yi Xu; Jidong Zhai</i>	
ROCKET: EFFICIENT AND SCALABLE ALL-PAIRS COMPUTATIONS ON HETEROGENEOUS PLATFORMS	1432
<i>Stijn Heldens; Pieter Hijma; Ben Van Werkhoven; Jason Maassen; Henri Bal; Rob Van Nieuwpoort</i>	
Author Index	