2020 IEEE International Conference on Cluster Computing (CLUSTER 2020)

Kobe, Japan
14 – 17 September 2020

IEEE

IEEE Catalog Number: CFP20235-POD
ISBN: 978-1-7281-6678-0

****** 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:	CFP20235-POD
ISBN (Print-On-Demand):	978-1-7281-6678-0
ISBN (Online):	978-1-7281-6677-3
ISSN:	1552-5244

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

2020 IEEE International Conference on Cluster Computing (CLUSTER)

CLUSTER 2020

Table of Contents

Message from the General Co-Chairs xv
Message from the Program Chairs xvii
HPCMASPA 2020 Workshop Welcome xix
EE HPC SOP 2020 Workshop Welcome xxi
EAHPC 2020 Workshop Welcome xxiv
CLUSTER 2020 Committees xxvi

IEEE Cluster 2020 Conference

Best Papers

Efficient Process-to-Node Mapping Algorithms for Stencil Computations 1
Konrad von Kirchbach (TU Wien, Austria), Markus Lehr (TU Wien, Austria), Sascha Hunold (TU Wien, Austria), Christian Schulz (University of Vienna, Austria), and Jesper Larsson Träff (TU Wien, Austria)

CuVPP: Filter-based Longest Prefix Matching in Software Data Planes 12
Minseok Kwon (Rochester Institute of Technology), Krishna Prasad Neupane (Rochester Institute of Technology), John Marshall (Cisco Systems, Inc.), and M. Mustafa Rafique (Rochester Institute of Technology)

HAN: A Hierarchical AutotuNed Collective Communication Framework 23
Xi Luo (University of Tennessee, USA), Wei Wu (Los Alamos National Laboratory, USA), George Bosilca (University of Tennessee, USA), Yu Pei (University of Tennessee, USA), and Qinglei Cao (University of Tennessee, USA)

DelveFS – An Event-driven Semantic File System for Object Stores 35
Marc-André Vef (Johannes Gutenberg University Mainz, Germany), Rebecca Steiner (Johannes Gutenberg University Mainz, Germany), Reza Salkhordeh (Johannes Gutenberg University Mainz, Germany), and Jörg Steinkamp (Johannes Gutenberg University Mainz, Germany)

Performance Characterization and Scheduling

NeoMPX: Characterizing and Improving Estimation of Multiplexing Hardware Counters for PAPI .47
Yi-Chao Wang (High Performance Computing Center, Shanghai Jiao Tong University, China), Jie Wang (High Performance Computing Center, Shanghai Jiao Tong University, China), Jin-Kun Chen (High Performance Computing Center, Shanghai Jiao Tong University, China), Si-Cheng Zuo (High Performance Computing Center, Shanghai Jiao Tong University, China), Xiao-Ming Su (High Performance Computing Center, Shanghai Jiao Tong University, China), and James Lin (High Performance Computing Center, Shanghai Jiao Tong University, China)

Grade10: A Framework for Performance Characterization of Distributed Graph Processing .57
Tim Hegeman (VU Amsterdam), Animesh Trivedi (VU Amsterdam), and Alexandru Iosup (VU Amsterdam)

Evaluating Worksharing Tasks on Distributed Environments .69
Marcos Maroñas (Barcelona Supercomputing Center (BSC)), Xavier Teruel (Barcelona Supercomputing Center (BSC)), J. Mark Bull (EPCC, University of Edinburgh), Eduard Ayguadé (Barcelona Supercomputing Center (BSC), Universitat Politècnica de Catalunya (UPC)), and Vicenç Beltran (Barcelona Supercomputing Center (BSC))

Resilient Scheduling of Moldable Jobs on Failure-Prone Platforms .81
Anne Benoit (Laboratoire LIP, France), Valentin Le Fèvre (Laboratoire LIP, France), Lucas Perotin (Laboratoire LIP, France; Vanderbilt University, USA), Padma Raghavan (Vanderbilt University, USA), Yves Robert (Laboratoire LIP, France; University of Tennessee, USA), and Hongyang Sun (Vanderbilt University, USA)

Modeling the Performance of Scientific Workflow Executions on HPC Platforms with Burst Buffers .92
Loïc Pottier (USC Information Sciences Institute, USA), Rafael Ferreira da Silva (USC Information Sciences Institute, USA), Henri Casanova (University of Hawai'i at Manoa, USA), and Ewa Deelman (USC Information Sciences Institute, USA)

Co-scheML: Interference-aware Container Co-scheduling Scheme using Machine Learning Application Profiles for GPU Clusters .104
Sejin Kim (Sookmyung Women's University, Korea) and Yoonhee Kim (Sookmyung Women's University, Korea)

Architecture and Network Support for HPC Workloads

Opportunities and Limitations of Quality-of-Service in Message Passing Applications on Adaptively Routed Dragonfly and Fat Tree Networks .109
Jeremiah J. Wilke (Scalable Modeling and Analysis, Sandia National Laboratories, USA) and Joseph P. Kenny (Scalable Modeling and Analysis, Sandia National Laboratories, USA)

MonSTer: An Out-of-the-Box Monitoring Tool for High Performance Computing Systems .119
Jie Li (Texas Tech University, USA), Ghazanfar Ali (Texas Tech University, USA), Ngan Nguyen (Texas Tech University, USA), Jon Hass (Dell EMC, Inc., USA), Alan Sill (Texas Tech University, USA), Tommy Dang (Texas Tech University, USA), and Yong Chen (Texas Tech University, USA)

Dynamic Kernel Fusion for Bulk Non-contiguous Data Transfer on GPU Clusters 130
Ching-Hsiang Chu (The Ohio State University, USA), Kawthar Shafie Khorassani (The Ohio State University, USA), Qinghua Zhou (The Ohio State University, USA), Hari Subramoni (The Ohio State University, USA), and Dhabaleswar K. Panda (The Ohio State University, USA)

Autoscaling High-Throughput Workloads on Container Orchestrators 142
Chao Zheng (University of Notre Dame), Nathaniel Kremer-Herman (University of Notre Dame), Tim Shaffer (University of Notre Dame), and Douglas Thain (University of Notre Dame)

SSP: Speeding up Small Flows for Proactive Transport in Datacenters 153
Yang Bai (National University of Defense Technology, China), Dezun Dong (National University of Defense Technology, China), Shan Huang (National University of Defense Technology, China), Zejia Zhou (National University of Defense Technology, China), and Xiangke Liao (National University of Defense Technology, China)

Quantifying the Impact of Network Congestion on Application Performance and Network Metrics 162
Yijia Zhang (Boston University, USA), Taylor Groves (National Energy Research Scientific Computing Center, USA), Brandon Cook (National Energy Research Scientific Computing Center, USA), Nicholas J. Wright (National Energy Research Scientific Computing Center, USA), and Ayse K. Coskun (Boston University, USA)

Analysis of Cooling Water Temperature Impact on Computing Performance and Energy Consumption 169
Jorji Nonaka (RIKEN R-CCS, Japan), Toshihiro Hanawa (The University of Tokyo / JCAHPC, Japan), and Fumiyoshi Shoji (RIKEN R-CCS, Japan)

Framework for Data and Storage

E2Clab: Exploring the Computing Continuum through Repeatable, Replicable and Reproducible Edge-to-Cloud Experiments 176
Daniel Rosendo (University of Rennes, Inria, France), Pedro Silva (Hasso-Plattner Institut, University of Potsdam, Germany), Matthieu Simonin (University of Rennes, Inria, France), Alexandru Costan (University of Rennes, Inria, France), and Gabriel Antoniu (University of Rennes, Inria, France)

Streaming File Transfer Optimization for Distributed Science Workflows 187
Davut Ucar (University of Nevada Reno) and Engin Arslan (University of Nevada Reno)

Exploring the Potential of Fast Delta Encoding: Marching to a Higher Compression Ratio 198
Haoliang Tan (Harbin Institute of Technology, China), Zhiyuan Zhang (Harbin Institute of Technology, China), Xiangyu Zou (Harbin Institute of Technology, China), Qing Liao (Harbin Institute of Technology, China; Cyberspace Security Research Center, Peng Cheng Laboratory, China), and Wen Xia (Harbin Institute of Technology, China; Cyberspace Security Research Center, Peng Cheng Laboratory, China)

Staging Based Task Execution for Data-driven, In-Situ Scientific Workflows 209
Zhe Wang (Rutgers Discovery Informatics Institute, Rutgers University), Pradeep Subedi (Rutgers Discovery Informatics Institute, Rutgers University), Matthieu Dorier (Argonne National Laboratory), Philip E. Davis (Rutgers Discovery Informatics Institute, Rutgers University), and Manish Parashar (Rutgers Discovery Informatics Institute, Rutgers University)

Flexible Data Redistribution in a Task-Based Runtime System 221
Qinglei Cao (University of Tennessee, USA), George Bosilca (University of Tennessee, USA), Wei Wu (Los Alamos National Laboratory, USA), Dong Zhong (University of Tennessee, USA), Aurelien Bouteiller (University of Tennessee, USA), and Jack Dongarra (University of Tennessee, USA; University of Manchester, UK)

Programming, System Software and Container

DeepClone: Lightweight State Replication of Deep Learning Models for Data Parallel Training 226
Bogdan Nicolae (Argonne National Laboratory, USA), Justin M. Wozniak (Argonne National Laboratory, USA), Matthieu Dorier (Argonne National Laboratory, USA), and Franck Cappello (Argonne National Laboratory, USA)

Exploring Non-Volatility of Non-Volatile Memory for High Performance Computing Under Failures 237
Jie Ren (University of California, Merced), Kai Wu (University of California, Merced), and Dong Li (University of California, Merced)

HCL: Distributing Parallel Data Structures in Extreme Scales 248
Hariharan Devarajan (Illinois Institute of Technology), Anthony Kougkas (Illinois Institute of Technology), Keith Bateman (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)

Predicting MPI Collective Communication Performance using Machine Learning 259
Sascha Hunold (TU Wien, Austria), Abhinav Bhatele (University of Maryland, USA), George Bosilca (Innovative Computing Laboratory, University of Tennessee, USA), and Peter Knees (TU Wien, Austria)

Decomposing MPI Collectives for Exploiting Multi-lane Communication 270
Jesper Larsson Träff (TU Wien, Austria) and Sascha Hunold (TU Wien, Austria)

Power Budgeting of Big Data Applications in Container-based Clusters 281
Jonatan Enes (Universidade da Coruña, CITIC, Spain), Guillaume Fieni (University of Lille / Inria, France), Roberto R. Expósito (Universidade da Coruña, CITIC, Spain), Romain Rouvoy (University of Lille / Inria, France; IUF, France), and Juan Touriño (Universidade da Coruña, CITIC, Spain)

Estimating Power Consumption of Containers and Virtual Machines in Data Centers 288
Xusheng Zhang (Jiangsu Key Laboratory of Big Data Security & Intelligent Processing, Nanjing University of Posts and Telecommunications), Ziyu Shen (Jiangsu Key Laboratory of Big Data Security & Intelligent Processing, Nanjing University of Posts and Telecommunications), Bin Xia (Jiangsu Key Laboratory of Big Data Security & Intelligent Processing, Nanjing University of Posts and Telecommunications), Zheng Liu (Jiangsu Key Laboratory of Big Data Security & Intelligent Processing, Nanjing University of Posts and Telecommunications), and Yun Li (Jiangsu Key Laboratory of Big Data Security & Intelligent Processing, Nanjing University of Posts and Telecommunications)

HPC Applications

Fast Scalable Approximate Nearest Neighbor Search for High-dimensional Data 294
Renga Bashyam K G (Indian Institute of Science, India) and Sathish Vadhiyar (Indian Institute of Science, India)

A Hybrid MPI+PGAS Approach to Improve Strong Scalability Limits of Finite Element Solvers 303
Niclas Jansson (PDC Center for High Performance Computing, KTH Royal Institute of Technology, Sweden)

Towards Data-Flow Parallelization for Adaptive Mesh Refinement Applications 314
Kevin Sala (Barcelona Supercomputing Center (BSC), Spain), Alejandro Rico (Arm Research, USA), and Vicenç Beltran (Barcelona Supercomputing Center (BSC), Spain)

Towards End-to-End SDC Detection for HPC Applications Equipped with Lossy Compression 326
Sihuan Li (University of California, Riverside), Sheng Di (Argonne National Laboratory, USA), Kai Zhao (University of California, Riverside), Xin Liang (Oak Ridge National Laboratory, USA), Zizhong Chen (University of California, Riverside), and Franck Cappello (Argonne National Laboratory, USA; University of Illinois at Urbana-Champaign, USA)

Efficient Execution of Dynamic Programming Algorithms on Apache Spark 337
Mohammad Mahdi Javanmard (Stony Brook University, USA), Zafar Ahmad (Stony Brook University, USA), Jaroslaw Zola (University at Buffalo, USA), Louis-Noël Pouchet (Colorado State University, USA), Rezaul Chowdhury (Stony Brook University, USA), and Robert Harrison (Stony Brook University, USA)

IO, Visualization, and Machine Learning

ECS2: A Fast Erasure Coding Library for GPU-Accelerated Storage Systems With Parallel & Direct IO 349
Chan Jung Chang (National Tsing Hua University, Taiwan), Jerry Chou (National Tsing Hua University, Taiwan), Yu-Ching Chou (H3 Platform Inc., Taiwan), and I-Hsin Chung (IBM T. J. Watson, Research Center, USA)

tf-Darshan: Understanding Fine-grained I/O Performance in Machine Learning Workloads 359
Steven W. D. Chien (Division of Computational Science and Technology, KTH Royal Institute of Technology, Sweden), Artur Podobas (Division of Computational Science and Technology, KTH Royal Institute of Technology, Sweden), Ivy B. Peng (Lawrence Livermore National Laboratory, USA), and Stefano Markidis (Division of Computational Science and Technology, KTH Royal Institute of Technology, Sweden)

Extending High-Level Synthesis with High-Performance Computing Performance Visualization 371
Jens Huthmann (Riken Center for Computational Science, Japan), Artur Podobas (Royal Institute of Technology, KTH, Sweden), Lukas Sommer (Embedded Systems and Applications Group, TU Darmstadt, Germany), Andreas Koch (Embedded Systems and Applications Group, TU Darmstadt, Germany), and Kentaro Sano (Riken Center for Computational Science, Japan)

Parallel Particle Advection Bake-Off for Scientific Visualization Workloads 381
Roba Binyahib (University of Oregon), David Pugmire (Oak Ridge National Laboratory), Abhishek Yenpure (University of Oregon), and Hank Childs (University of Oregon)

Data Life Aware Model Updating Strategy for Stream-based Online Deep Learning 392
Wei Rang (UNC Charlotte, USA), Donglin Yang (UNC Charlotte, USA), Dazhao Cheng (UNC Charlotte, USA), Kun Suo (Kennesaw State University, USA), and Wei Chen (Nvidia Corporation, USA)

Optimizing GPU Memory Transactions for Convolution Operations 399
Gangzhao Lu (Computer Science and Technology, Harbin Institute of Technology, China), Weizhe Zhang (Computer Science and Technology, Harbin Institute of Technology, China), and Zheng Wang (University of Leeds, United Kingdom)

Posters

System-Level vs. Application-Level Checkpointing 404
Jonas Posner (University of Kassel, PLM, Germany)

An HPC-based Prediction on the Practicality of Long-distance Quantum Key Distributions 406
Hoon Ryu (Korea Institute of Science and Technology Information, Republic of Korea) and Ji-Hoon Kang (Korea Institute of Science and Technology Information, Republic of Korea)

Performance Evaluation of Supercomputer Fugaku using Breadth-First Search Benchmark in Graph500 408
Masahiro Nakao (RIKEN Center for Computational Science, Japan), Koji Ueno (Fixstars Corporation, Japan), Katsuki Fujisawa (Kyushu University, Japan), Yuetsu Kodama (RIKEN Center for Computational Science, Japan), and Mitsuhisa Sato (RIKEN Center for Computational Science, Japan)

OctCNN: An Energy-Efficient FPGA Accelerator for CNNs using Octave Convolution Algorithm 410
Wenqi Lou (University of Science and Technology of China), Chao Wang (University of Science and Technology of China), Lei Gong (University of Science and Technology of China), and Xuehai Zhou (University of Science and Technology of China)

ChOWDER: A New Approach for Viewing 3D Web GIS on Ultra-High-Resolution Scalable Display …. 412
Tomohiro Kawanabe (Center for Computational Science, RIKEN, Japan), Kazuma Hatta (Digital Reality Lab., IMAGICADIGITALSCAPE Co., Ltd., Japan), and Kenji Ono (Research Institute for Information Technology, Kyushu University, Japan)

An FPGA-based Sound Field Rendering System .414……
Yiyu Tan (RIKEN Center for Computational Science, Japan) and Toshiyuki Imamura (RIKEN Center for Computational Science, Japan)

The Case for Better Integrating Scalable Data Stores and Stream-Processing Systems .416……
Antonis Papaioannou (Institute of Computer Science, Foundation for Research and Technology – Hellas, Greece; University of Crete, Greece), Chrysostomos Zeginis (Institute of Computer Science, Foundation for Research and Technology – Hellas, Greece), and Kostas Magoutis (Institute of Computer Science, Foundation for Research and Technology – Hellas, Greece; University of Crete, Greece)

Prompt Report on Exa-scale HPL-AI benchmark .418……
Shuhei Kudo (Center for Computational Science, RIKEN, Japan), Keigo Nitadori (Center for Computational Science, RIKEN, Japan), Takuya Ina (Center for Computational Science, RIKEN, Japan), and Toshiyuki Imamura (Center for Computational Science, RIKEN, Japan)

Implementing a Comprehensive Networks-on-Chip Generator with Optimal Configurations .420..
Hao Zhang (Center for Biosystems, Dynamics Research, RIKEN), Itta Ohmura (Center for Biosystems, Dynamics Research, RIKEN), and Makoto Taiji (Center for Biosystems, Dynamics Research, RIKEN)

Toward OpenACC-enabled GPU-FPGA Accelerated Computing .422……
Norihisa Fujita (University of Tsukuba, Japan), Ryohei Kobayashi (University of Tsukuba, Japan), Yoshiki Yamaguchi (University of Tsukuba, Japan), Kohji Yoshikawa (University of Tsukuba, Japan), Makito Abe (University of Tsukuba, Japan), and Masayuki Umemura (University of Tsukuba, Japan)

Workshops

HPCMASPA 2020 Workshop

PIKA: Center-Wide and Job-Aware Cluster Monitoring .424……
Robert Dietrich (Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden, Germany), Frank Winkler (Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden, Germany), Andreas Knüpfer (Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden, Germany), and Wolfgang Nagel (Center for Information Services and High Performance Computing (ZIH), Technische Universität Dresden, Germany)

HPC System Data Pipeline to Enable Meaningful Insights through Analysis-Driven Visualizations 433
Benjamin Schwaller (Sandia National Laboratories, USA), Nick Tucker (Open Grid Computing, USA), Tom Tucker (Open Grid Computing, USA), Benjamin Allan (Sandia National Laboratories, USA), and Jim Brandt (Sandia National Laboratories, USA)

MAP: A Visual Analytics System for Job Monitoring and Analysis 442
Ashish Pal (IIT, India) and Preeti Malakar (IIT, India)

Towards Workload-adaptive Scheduling for HPC Clusters 449
Alexander V. Goponenko (University of Central Florida), Ramin Izadpanah (University of Central Florida), Jim M. Brandt (Sandia National Laboratories, Albuquerque, NM), and Damian Dechev (University of Central Florida)

Democratizing Parallel Filesystem Monitoring 454
Richard Todd Evans (Texas Advanced Computing Center, University of Texas at Austin, USA)

LDMS Monitoring of EDR InfiniBand Networks 459
Benjamin A. Allan (Sandia National Laboratories, USA), Michael Aguilar (Sandia National Laboratories, USA), Benjamin Schwaller (Sandia National Laboratories, USA), and Steven Langer (Lawrence Livermore National Laboratory, USA)

EE HPC SOP 2020 Workshop

Energy Optimization and Analysis with EAR 464
Julita Corbalan (BSC, UPC, Spain), Lluis Alonso (BSC, Spain), Jordi Aneas (BSC, Spain), and Luigi Brochard (Energy Aware Solutions (EAS), France)

Toward an End-to-End Auto-tuning Framework in HPC PowerStack 473
Xingfu Wu (Argonne National Laboratory, The University of Chicago, USA), Aniruddha Marathe (Lawrence Livermore National Laboratory, USA), Siddhartha Jana (Intel Corp., USA), Ondrej Vysocky (IT4Innovations National Supercomputing Center, Czech Republic), Jophin John (Technical University of Munich, Germany), Andrea Bartolini (University of Bologna, Italy), Lubomir Riha (IT4Innovations National Supercomputing Center, Czech Republic), Michael Gerndt (Technical University of Munich, Germany), Valerie Taylor (Argonne National Laboratory, The University of Chicago, USA), and Sridutt Bhalachandra (Lawrence Berkeley National Laboratory, USA)

Evaluation of Power Management Control on the Supercomputer Fugaku 484
Yuetsu Kodama (RIKEN Center for Computational Science), Tetsuya Odajima (RIKEN Center for Computational Science), Eishi Arima (Information Technology Center, The University of Tokyo), and Mitsuhisa Sato (RIKEN Center for Computational Science)

HUD-Oden: A Practical Evaluation Environment for Analyzing Hot-Water Cooled Processors 494
Jorji Nonaka (HPC Usability Development Unit, RIKEN Center for Computational Science, Japan) and Fumiyoshi Shoji (Operations and Computer Technologies Division, RIKEN Center for Computational Science, Japan)

Global Experiences with HPC Operational Data Measurement, Collection and Analysis 499
Michael Ott (Leibniz Supercomputing Centre), Woong Shin (Oak Ridge National Laboratory), Norman Bourassa (Lawrence Berkeley National Laboratory), Torsten Wilde (Hewlett Packard Enterprise), Stefan Ceballos (Oak Ridge National Laboratory), Melissa Romanus (Lawrence Berkeley National Laboratory), and Natalie Bates (Energy Efficient HPC Working Group)

A Study of Operational Impact on Power Usage Effectiveness using Facility Metrics and Server Operation Logs in the K Computer 509
Masaaki Terai (RIKEN Center for Computational Science, Japan), Fumiyoshi Shoji (RIKEN Center for Computational Science, Japan), Toshiyuki Tsukamoto (RIKEN Center for Computational Science, Japan), and Yukihiro Yamochi (RIKEN Center for Computational Science, Japan)

A Supercomputing Center Experience with Cooling Control Design 514
Michael Kercher (Computational Information, Systems Lab (CISL), National Center for Atmospheric Research (NCAR), USA) and Gary New (Computational Information, Systems Lab (CISL), National Center for Atmospheric Research (NCAR), USA)

Investigative Report on Electrical Commissioning in HPC Data Centers 519
Joseph Prisco (IBM), Grant Stewart (Los Alamos National Laboratory), Herbert Huber (Leibniz Supercomputing Center (LRZ)), Randy Rannow (Silverdraft Supercomputing), Jason Hick (Los Alamos National Laboratory), Dave Martinez (Sandia National Laboratory), Brandon Hong (Lawrence Livermore National Laboratory), and Aditya M. Deshpande (Microsoft Corporation)

EAHPC 2020 Workshop

Preliminary Performance Evaluation of the Fujitsu A64FX Using HPC Applications 523
Tetsuya Odajima (RIKEN Center for Computational Science, Japan), Yuetsu Kodama (RIKEN Center for Computational Science, Japan), Miwako Tsuji (RIKEN Center for Computational Science, Japan), Motohiko Matsuda (RIKEN Center for Computational Science, Japan), Yutaka Maruyama (RIKEN Center for Computational Science, Japan), and Mitsuhisa Sato (RIKEN Center for Computational Science, Japan)

The Effects of Wide Vector Operations on Processor Caches 531
Andrei Poenaru (University of Bristol, United Kingdom) and Simon McIntosh-Smith (University of Bristol, United Kingdom)

CoreNEURON: Performance and Energy Efficiency Evaluation on Intel and Arm CPUs 540
Joel Criado (Barcelona Supercomputing Center, Spain), Marta Garcia-Gasulla (Barcelona Supercomputing Center, Spain), Pramod Kumbhar (Blue Brain Project, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland), Omar Awile (Blue Brain Project Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland), Ioannis Magkanaris (Blue Brain Project Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland), and Filippo Mantovani (Barcelona Supercomputing Center, Spain)

Investigating Applications on the A64FX 549
Adrian Jackson (EPCC, The University of Edinburgh, United Kingdom), Michèle Weiland (EPCC, The University of Edinburgh, United Kingdom), Nick Brown (EPCC, The University of Edinburgh, United Kingdom), Andrew Turner (EPCC, The University of Edinburgh, United Kingdom), and Mark Parsons (EPCC, The University of Edinburgh, United Kingdom)

Porting Applications to Arm-based Processors 559
Bine Brank (Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany), Stepan Nassyr (Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany), Fatemeh Pouyan (Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany), and Dirk Pleiter (Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany)

Performance Evaluation of ParalleX Execution Model on Arm-based Platforms 567
Nikunj Gupta (IIT Roorkee, India; The STE, AR Group), Rohit Ashiwal (IIT Roorkee, India), Bine Brank (JSC, Jülich Research Centre, Germany), Sateesh K. Peddoju (IIT Roorkee, India), and Dirk Pleiter (JSC, Jülich Research Centre, Germany)

On the Usage of the Arm C Language Extensions for a High-Order Finite-Element Kernel 576
Sylvain Jubertie (Univ. d'Orléans, INSA CVL, France), Guillaume Quintin (Agenium Scale, France), and Fabrice Dupros (Arm Sophia-Antipolis, France)

Author Index 581