

# **2018 IEEE International Conference on Cluster Computing (CLUSTER 2018)**

**Belfast, United Kingdom  
10-13 September 2018**



**IEEE Catalog Number: CFP18235-POD  
ISBN: 978-1-5386-8320-0**

**Copyright © 2018 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:	CFP18235-POD
ISBN (Print-On-Demand):	978-1-5386-8320-0
ISBN (Online):	978-1-5386-8319-4
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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2018 IEEE International Conference on Cluster Computing **CLUSTER 2018**

## Table of Contents

Message from the General Co-Chairs .xiv.....	
Message from the Program Chair .xv.....	
FTS 2018 Workshop Welcome .xvi.....	
HPCMASPA 2018 Workshop Welcome .xvii.....	
WRAP 2018 Workshop Welcome .xix.....	
REV-A 2018 Workshop Welcome .xx.....	
Committees .xxi.....	
Keynotes .xxvi.....	

## IEEE Cluster Conference

### Session I: Best Papers (Areas 1 and 2)

PaSTRI: Error-Bounded Lossy Compression for Two-Electron Integrals in Quantum Chemistry .1.....	
<i>Ali Murat Gok (Northwestern University), Sheng Di (Argonne National Laboratory), Yuri Alexeev (Argonne National Laboratory), Dingwen Tao (The University of Alabama), Vladimir Mironov (Lomonosov Moscow State University), Xin Liang (University of California, Riverside), and Franck Cappello (Argonne National Laboratory, University of Illinois at Urbana-Champaign)</i>	
SALaR: Scalable and Adaptive Designs for Large Message Reduction Collectives .12.....	
<i>Mohammadreza Bayatpour (The Ohio State University), Jahanzeb Maqbool Hashmi (The Ohio State University), Sourav Chakraborty (The Ohio State University), Hari Subramoni (The Ohio State University), Pouya Kousha (The Ohio State University), and Dhableswar K. Panda (The Ohio State University)</i>	

### Paper Session II: Matrix Algorithms

Energy Analysis and Optimization for Resilient Scalable Linear Systems .24.....	
<i>Zheng Miao (Clemson University), Jon Calhoun (Clemson University), and Rong Ge (Clemson University)</i>	
Load-Balancing-Aware Parallel Algorithms of H-Matrices with Adaptive Cross Approximation for GPUs .35....	
<i>Tetsuya Hoshino (The University of Tokyo), Akihiro Ida (The University of Tokyo), Toshihiro Hanawa (The University of Tokyo), and Kengo Nakajima (The University of Tokyo)</i>	

A New Approach for Sparse Matrix Classification Based on Deep Learning Techniques .46.....  
*Juan C. Pichel (CiTIUS, Universidade de Santiago de Compostela (Spain)) and Beatriz Pateiro-López (Universidade de Santiago de Compostela (Spain))*

### **Paper Session III: Architecture and Interconnect**

Janus: A Generic QoS Framework for Software-as-a-Service Applications .55.....  
*Qingye Jiang (The University of Sydney), Young Choon Lee (Macquarie University), and Albert Y. Zomaya (The University of Sydney)*

Exploring HPC and Big Data Convergence: A Graph Processing Study on Intel Knights Landing .66.....  
*Alexandru Uta (Vrije Universiteit Amsterdam), Ana Lucia Varbanescu (Universiteit van Amsterdam), Ahmed Musaafer (Vrije Universiteit Amsterdam), Chris Lemaire (TUDelft), and Alexandru Iosup (Vrije Universiteit Amsterdam)*

### **Paper Session IV: Generating and Optimizing Applications**

Whole Program Generation of Massively Parallel Shallow Water Equation Solvers .78.....  
*Sebastian Kuckuk (Friedrich-Alexander-Universität Erlangen-Nürnberg) and Harald Köstler (Friedrich-Alexander-Universität Erlangen-Nürnberg)*

OpenACC vs the Native Programming on Sunway TaihuLight: A Case Study with GTC-P .88.....  
*Linjin Cai (Shanghai Jiao Tong University), Yi-Chao Wang (Shanghai Jiao Tong University), William Tang (Princeton University), Bei Wang (Princeton University), Stephane Ethier (Princeton Plasma Physics Laboratory, Princeton University), Zhao Liu (National Supercomputing Center in Wuxi), and James Lin (Shanghai Jiao Tong University)*

Parallel Approximation of the Maximum Likelihood Estimation for the Prediction of Large-Scale Geostatistics Simulations .98.....  
*Sameh Abdulah (King Abdullah University of Science Technology), Hatem Ltaief (King Abdullah University of Science Technology), Ying Sun (King Abdullah University of Science Technology), Marc G. Genton (King Abdullah University of Science Technology), and David E. Keyes (King Abdullah University of Science Technology)*

Modeling I/O Performance Variability Using Conditional Variational Autoencoders .109.....  
*Sandeep Madireddy (Argonne National Laboratory), Prasanna Balaprakash (Argonne National Laboratory), Philip Carns (Argonne National Laboratory), Robert Latham (Argonne National Laboratory), Robert Ross (Argonne National Laboratory), Shane Snyder (Argonne National Laboratory), and Stefan Wild (Argonne National Laboratory)*

### **Paper Session V: Filesystems and Applications**

SciDP: Support HPC and Big Data Applications via Integrated Scientific Data Processing .114.....  
*Kun Feng (Illinois Institute of Technology), Xian-He Sun (Illinois Institute of Technology), Xi Yang (Teradata Inc), and Shujia Zhou (Northrop Grumman Information Technology)*

Applying Pwrake Workflow System and Gfarm File System to Telescope Data Processing .124.....  
*Masahiro Tanaka (Keio University), Osamu Tatebe (University of Tsukuba), and Hideyuki Kawashima (Keio University)*

UniviStor: Integrated Hierarchical and Distributed Storage for HPC .134.....  
*Teng Wang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Bin Dong (Lawrence Berkeley National Laboratory), and Houjun Tang (Lawrence Berkeley National Laboratory)*

HIDStore: A Hierarchical Intermediate Data Storage System for Seismic Processing Application .145.....  
*Yida Wang (Beihang University), Changhai Zhao (BGP Inc. CNPC), Zengbo Wang (BGP Inc. CNPC), Chao Liu (Beihang University), Chao Li (Beihang University), Haihua Yan (Beihang University), and Jiamin Wen (BGP Inc. CNPC)*

## Poster Reception

Performance Evaluation of Parallel Stripmap CS-SAR Imaging on NVLink-Connected GPUs .150.....  
*Masato Gocho (Mitsubishi Electric Corporation) and Takehiro Hoshino (Mitsubishi Electric Corporation)*

C2-HyTM : An Adaptive Hybrid Transactional Memory Scheme as Efficient Synchronization Mechanism for Multicore Parallel Processing .152.....  
*Shin Young-Sung (Chonbuk National University), Jang Jin-Su (Chonbuk National University), Jang Yeon-Woo (Chonbuk National University), and Chang Jae-Woo (Chonbuk National University)*

Scalable Shared-Memory Parallelization of the Block Recursive Inversion Algorithm .154.....  
*Maria Clara Silva (Polytechnique Montreal), Iria C. S. Cosme (Instituto Federal do Rio Grande do Norte), Idalmis M. Sardina (Universidade Federal do Rio Grande do Norte), and Samuel Xavier-de-Souza (Universidade Federal do Rio Grande do Norte)*

Maximizing Throughput on Power-Bounded HPC Systems .156.....  
*Pengfei Zou (Clemson University), Derek Rodriguez (Clemson University), and Rong Ge (Clemson University)*

EDDAPS: An Efficient Data Distribution Approach for PCM-Based SSD .158.....  
*Ronnie Mativenga (Ajou University), Prince Hamandawana (Ajou University), Se Jin Kwon (Kangwon National University), and Tae-Sun Chung (Ajou University)*

PADS: Performance-Aware Dynamic Scheduling for Effective MapReduce Computation in Heterogeneous Clusters .160.....  
*Prince Hamandawana (Ajou University), Ronnie Mativenga (Ajou University), Se Jin Kwon (Kangwon National University), and Tae-Sun Chung (Ajou University)*

Reducing Tail Latency of Interactive Multi-tier Workloads in the Cloud .162.....  
*Kejiang Ye (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences) and Cheng-Zhong Xu (Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences)*

A Non-blocking Buddy System for Scalable Memory Allocation on Multi-core Machines .164.....	
	<i>Romolo Marotta (Sapienza, University of Rome), Mauro Ianni (Sapienza, University of Rome), Andrea Scarselli (Sapienza, University of Rome), Alessandro Pellegrini (Sapienza, University of Rome), and Francesco Quaglia (Università di Roma Tor Vergata)</i>
Graph Partitioning Algorithm with LSH: Poster Extended Abstract .166.....	
	<i>Weidong Zhang (Peking University) and Mingyue Zhang (Peking University)</i>

## Session VI: Best Papers - Areas 3 and 4

Neural Network Based Silent Error Detector .168.....	
	<i>Chen Wang (University of Illinois at Urbana-Champaign), Nikoli Dryden (University of Illinois at Urbana-Champaign), Franck Cappello (Argonne National Laboratory), and Marc Snir (University of Illinois at Urbana-Champaign)</i>
An Efficient Transformation Scheme for Lossy Data Compression with Point-Wise Relative Error Bound .179.	
	<i>Xin Liang (University of California, Riverside), Sheng Di (Argonne National Laboratory), Dingwen Tao (The University of Alabama), Zizhong Chen (University of California, Riverside), and Franck Cappello (Argonne National Laboratory)</i>

## Paper Session VII: Benchmarking and Modeling

A Methodology for Characterizing the Correspondence Between Real and Proxy Applications .190.....	
	<i>Omar Aaziz (Sandia National Laboratories), Jeanine Cook (Sandia National Laboratories), Jonathan Cook (New Mexico State University), Tanner Juedeman (Sandia National Laboratories), David Richards (Lawrence Livermore National Laboratory), and Courtenay Vaughan (Sandia National Laboratories)</i>
Lightweight Requirements Engineering for Exascale Co-design .201.....	
	<i>Alexandru Calotoiu (TU Darmstadt), Alexander Graf (TU Darmstadt), Torsten Hoefler (ETH Zurich), Daniel Lorenz (TU Darmstadt), Sebastian Rinke (TU Darmstadt), and Felix Wolf (TU Darmstadt)</i>
Next Stop "NoOps": Enabling Cross-System Diagnostics Through Graph-Based Composition of Logs and Metrics .212.....	
	<i>Michał Zasadziski (CA Technologies), Marc Solé (CA Technologies), Alvaro Brandon (Universitat Politecnica de Madrid), Victor Muntés-Mulero (CA Technologies), and David Carrera (Universitat Politecnica de Catalunya)</i>

## Paper Session VIII: Managing Heterogeneity and Imbalance

Cutting the Tail: Designing High Performance Message Brokers to Reduce Tail Latencies in Stream Processing .223.....	
	<i>M. Haseeb Javed (The Ohio State University), Xiaoyi Lu (The Ohio State University), and Dhableswar K. Panda (The Ohio State University)</i>

Dynamic Control of CPU Usage in a Lambda Platform .234.....	
	<i>Young Ki Kim (The University of Sydney), M. Reza HoseinyFarahabady (The University of Sydney), Young Choon Lee (Macquarie University), Albert Y. Zomaya (The University of Sydney), and Raja Jurdak (CSIRO)</i>
A Heterogeneity-Aware Task Scheduler for Spark .245.....	
	<i>Luna Xu (Virginia Tech), Ali R. Butt (Virginia Tech), Seung-Hwan Lim (Oak Ridge National Lab), and Ramakrishnan Kannan (Oak Ridge National Lab)</i>

## **Paper Session IX: Graphs and Big Data Analytics**

Computing Exact Vertex Eccentricity on Massive-Scale Distributed Graphs .257.....	
	<i>Keita Iwabuchi (Lawrence Livermore National Laboratory), Geoffrey Sanders (Lawrence Livermore National Laboratory), Keith Henderson (Lawrence Livermore National Laboratory), and Roger Pearce (Lawrence Livermore National Laboratory)</i>
A Scalable Distributed Louvain Algorithm for Large-Scale Graph Community Detection .268.....	
	<i>Jianping Zeng (University of Nebraska-Lincoln) and Hongfeng Yu (University of Nebraska-Lincoln)</i>
Optimizing Distributed Data-Intensive Workflows .279.....	
	<i>Ryan D. Friese (Pacific Northwest National Laboratory), Nathan R. Tallent (Pacific Northwest National Laboratory), Malachi Schram (Pacific Northwest National Laboratory), Mahantesh Halappanavar (Pacific Northwest National Laboratory), and Kevin J. Barker (Pacific Northwest National Laboratory)</i>

## **Paper Session X: Performance Engineering in Filesystems**

Harmonia: An Interference-Aware Dynamic I/O Scheduler for Shared Non-volatile Burst Buffers .290.....	
	<i>Anthony Kougkas (Illinois Institute of Technology), Hariharan Devarajan (Illinois Institute of Technology), Xian-He Sun (Illinois Institute of Technology), and Jay Lofstead (Sandia National Laboratories)</i>
CRUM: Checkpoint-Restart Support for CUDA's Unified Memory .302.....	
	<i>Rohan Garg (Northeastern University), Apoorve Mohan (Northeastern University), Michael Sullivan (NVIDIA Corp.), and Gene Cooperman (Northeastern University)</i>
Fixed-PSNR Lossy Compression for Scientific Data .314.....	
	<i>Dingwen Tao (The University of Alabama), Sheng Di (Argonne National Laboratory), Xin Liang (University of California, Riverside), Zizhong Chen (University of California, Riverside), and Franck Cappello (Argonne National Laboratory)</i>

GekkoFS - A Temporary Distributed File System for HPC Applications .319.....  
*Marc-André Vef (Johannes Gutenberg University Mainz), Nafiseh Moti (Johannes Gutenberg University Mainz), Tim Süß (Johannes Gutenberg University Mainz), Tommaso Tocci (Barcelona Supercomputing Center), Ramon Nou (Barcelona Supercomputing Center), Alberto Miranda (Barcelona Supercomputing Center), Toni Cortes (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya), and André Brinkmann (Johannes Gutenberg University Mainz)*

## Paper Session XI: Hierarchy and Sharing in System Architecture

Hierarchical Clock Synchronization in MPI .325.....  
*Sascha Hunold (TU Wien) and Alexandra Carpen-Amarie (Fraunhofer ITWM)*

Hybrid MPI+OpenMP Reactive Work Stealing in Distributed Memory in the PDE Framework sam(oa)^2 .337.  
*Philipp Samfass (Technical University of Munich), Jannis Klinkenberg (RWTH Aachen University), and Michael Bader (Technical University of Munich)*

Co-Scheduling HPC Workloads on Cache-Partitioned CMP Platforms .348.....  
*Guillaume Aupy (Inria, Labri, Univ. Bordeaux, CNRS, Bordeaux-INP, France), Anne Benoit (Laboratoire LIP, École Normale Supérieure de Lyon, France - Georgia Institute of Technology, Atlanta, GA, USA), Brice Goglin (Inria, Labri, Univ. Bordeaux, CNRS, Bordeaux-INP, France), Loïc Pottier (Laboratoire LIP, École Normale Supérieure de Lyon, France), and Yves Robert (Laboratoire LIP, École Normale Supérieure de Lyon, France - University of Tennessee, Knoxville TN, USA)*

## Paper Session XII: Scheduling, Elasticity and Energy

Leveraging Dependency in Scheduling and Preemption for High Throughput in Data-Parallel Clusters .359.....  
*Jinwei Liu (Clemson University), Haiying Shen (University of Virginia), Ankur Sarker (University of Virginia), and Wingyan Chung (University of Central Florida)*

Self-Consumption Optimization of Renewable Energy Production in Distributed Clouds .370.....  
*Benjamin Camus (Inria), Anne Blavette (CNRS), Fanny Dufossé (Inria), and Anne-Cécile Orgerie (CNRS)*

Elasticity in Graph Analytics? A Benchmarking Framework for Elastic Graph Processing .381.....  
*Alexandru Uta (Vrije Universiteit Amsterdam), Sietse Au (TUDelft), Alexey Ilyushkin (TUDelft), and Alexandru Iosup (Vrije Universiteit Amsterdam)*

## Paper Session XIII: Deep Learning

Efficient Training of Convolutional Neural Nets on Large Distributed Systems .392.....  
*Dheeraj Sreedhar (IBM Research), Vaibhav Saxena (IBM Research), Yogish Sabharwal (IBM Research), Ashish Verma (IBM Research), and Sameer Kumar (Google Inc)*



- Accelerating Deep Learning Frameworks with Micro-Batches .402.....  
*Yosuke Oyama (Tokyo Institute of Technology), Tal Ben-Nun (ETH Zurich), Torsten Hoefler (ETH Zurich), and Satoshi Matsuoka (RIKEN Center for Computational Science, Tokyo Institute of Technology)*
- swCaffe: A Parallel Framework for Accelerating Deep Learning Applications on Sunway TaihuLight .413.....  
*Liandeng Li (Tsinghua University), Jiarui Fang (Tsinghua University), Haohuan Fu (Tsinghua University), Jinlei Jiang (Tsinghua University), Wenlai Zhao (Tsinghua University), Conghui He (Tsinghua University), Xin You (Beihang University), and Guangwen Yang (Tsinghua University)*

## Paper Session XIV: Languages and Programming Models

- CharmPy: A Python Parallel Programming Model .423.....  
*Juan J. Galvez (University of Illinois at Urbana-Champaign), Karthik Senthil (University of Illinois at Urbana-Champaign), and Laxmikant Kale (University of Illinois at Urbana-Champaign)*
- And Now for Something Completely Different: Running Lisp on GPUs .434.....  
*Tim Süß (Johannes Gutenberg University Mainz), Nils Döring (Johannes Gutenberg University Mainz), André Brinkmann (Johannes Gutenberg University Mainz), and Lars Nagel (Loughborough University)*
- The AllScale Runtime Application Model .445.....  
*Herbert Jordan (University of Innsbruck), Thomas Heller (University of Erlangen-Nuremberg), Philipp Gschwandtner (University of Innsbruck), Peter Zangerl (University of Innsbruck), Peter Thoman (University of Innsbruck), Dietmar Fey (University of Erlangen-Nuremberg), and Thomas Fahringer (University of Innsbruck)*

## Paper Session XV: Scalable Filesystems

- PDFE: Flexible Parallel State Machine Replication for Cloud Computing .456.....  
*Lihui Wu (Sun Yat-sen University), Weigang Wu (Sun Yat-sen University), Ning Huang (Sun Yat-sen University), and Zhiguang Chen (Sun Yat-sen University)*
- IOMiner: Large-Scale Analytics Framework for Gaining Knowledge from I/O Logs .466.....  
*Teng Wang (Lawrence Berkeley National Laboratory), Shane Snyder (Argonne National Laboratory), Glenn Lockwood (Lawrence Berkeley National Laboratory), Philip Carns (Argonne National Laboratory), Nicholas Wright (Lawrence Berkeley National Laboratory), and Suren Byna (Lawrence Berkeley National Laboratory)*
- A Transparent Server-Managed Object Storage System for HPC .477.....  
*Jingqing Mu (the HDF group), Jerome Soumagne (the HDF group), Houjun Tang (Lawrence Berkeley National Laboratory), Suren Byna (Lawrence Berkeley National Laboratory), Quincey Koziol (Lawrence Berkeley National Laboratory), and Richard Warren (the HDF group)*

## Paper Session XVI: Algorithms, Applications and Performance

Efficient Algorithms for the Summed Area Tables Primitive on GPUs .482 .....	
<i>Peng Chen (Tokyo Institute of Technology), Mohamed Wahib (National Institute of Advanced Industrial Science and Technology), Shinichiro Takizawa (National Institute of Advanced Industrial Science and Technology), Ryousei Takano (National Institute of Advanced Industrial Science and Technology), and Satoshi Matsuoka (RIKEN Center for Computational Science)</i>	
Relaxing Scalability Limits with Speculative Parallelism in Sequential Monte Carlo .494 .....	
<i>Balazs Nemeth (Universiteit Hasselt), Tom Haber (Universiteit Hasselt), Jori Liesenborgs (Universiteit Hasselt), and Wim Lamotte (Universiteit Hasselt)</i>	
Predicting Performance Using Collaborative Filtering .504 .....	
<i>Shweta Salaria (Tokyo Institute of Technology), Aleksandr Drozd (Tokyo Institute of Technology), Artur Podobas (Tokyo Institute of Technology), and Satoshi Matsuoka (RIKEN Center for Computational Science)</i>	

## Workshops

### FTS'18: Fault Tolerant Systems Workshop

A Failure Prediction-Based Adaptive Checkpointing Method with Less Reliance on Temperature Monitoring for HPC Applications .515 .....	
<i>Muhammad Alfian Amrizal (Tohoku University), Pei Li (Tohoku University), Mulya Agung (Tohoku University), Ryusuke Egawa (Tohoku University), and Hiroyuki Takizawa (Tohoku University)</i>	
Challenges in Developing MPI Fault-Tolerant Fortran Applications .524 .....	
<i>Nathan Weeks (Iowa State University), Glenn Luecke (Iowa State University), Pieter Maris (Iowa State University), and James Vary (Iowa State University)</i>	

# HPCMASPA'18: Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications

Large-Scale System Monitoring Experiences and Recommendations .532.....	
<i>Ville Ahlgren (Center for Scientific Computing (CSC - IT Center for Science)), Stefan Andersson (High Performance Computing Center Stuttgart (HLRS)), Jim Brandt (Sandia National Laboratories (SNL)), Nicholas Cardo (Swiss National Supercomputing Centre (CSCS)), Sudheer Chunduri (Argonne Leadership Computing Facility (ALCF)), Jeremy Enos (National Center for Supercomputing Applications (NCSA)), Parks Fields (Los Alamos National Laboratory (LANL)), Ann Gentile (Sandia National Laboratories (SNL)), Richard Gerber (National Energy Research Scientific Computing Center (NERSC)), Michael Gienger (High Performance Computing Center Stuttgart (HLRS)), Joe Greenseid (Cray Inc.), Annette Greiner (National Energy Research Scientific Computing Center (NERSC)), Bilel Hadri (King Abdullah University of Science and Technology (KAUST)), Yun He (National Energy Research Scientific Computing Center (NERSC)), Dennis Hoppe (High Performance Computing Center Stuttgart (HLRS)), Urpo Kaila (Center for Scientific Computing (CSC - IT Center for Science)), Kaki Kelly (Los Alamos National Laboratory (LANL)), Mark Klein (Swiss National Supercomputing Centre (CSCS)), Alex Kristiansen (Argonne Leadership Computing Facility (ALCF)), Steve Leak (National Energy Research Scientific Computing Center (NERSC)), Mike Mason (Los Alamos National Laboratory (LANL)), Kevin Pedretti (Sandia National Laboratories (SNL)), Jean-Guillaume Piccinali (Swiss National Supercomputing Centre (CSCS)), Jason Repik (Sandia National Laboratories (SNL)), Jim Rogers (Oak Ridge National Laboratory (ORNL)), Susanna Salminen (Center for Scientific Computing (CSC - IT Center for Science)), Mike Showerman (National Center for Supercomputing Applications (NCSA)), Cary Whitney (National Energy Research Scientific Computing Center (NERSC)), and Jim Williams (Los Alamos National Laboratory (LANL))</i>	
Modeling Expected Application Runtime for Characterizing and Assessing Job Performance .543.....	
<i>Omar Aaziz (Sandia National Laboratories), Jonathan Cook (New Mexico State University), and Mohammed Tanash (New Mexico State University)</i>	
Improving Power Efficiency Through Fine-Grain Performance Monitoring in HPC Clusters .552.....	
<i>Mathieu Stoffel (Bull, Atos technologies) and Abdelhafid Mazouz (Bull, Atos technologies)</i>	
Characterizing Supercomputer Traffic Networks Through Link-Level Analysis .562.....	
<i>Saurabh Jha (University of Illinois at Urbana-Champaign), Jim Brandt (Sandia National Laboratories), Ann Gentile (Sandia National Laboratories), Zbigniew Kalbarczyk (University of Illinois), and Ravishankar Iyer (University of Illinois at Urbana-Champaign)</i>	
A Big Data Analytics Framework for HPC Log Data: Three Case Studies Using the Titan Supercomputer Log .571.....	
<i>Byung H. Park (Oak Ridge National Laboratory), Yawei Hui (Oak Ridge National Laboratory), Swen Boehm (Oak Ridge National Laboratory), Rizwan A. Ashraf (Oak Ridge National Laboratory), Christopher Layton (Oak Ridge National Laboratory), and Christian Engelmann (Oak Ridge National Laboratory)</i>	

Federating XDMoD to Monitor Affiliated Computing Resources .580.....	
	<i>Jeanette Sperhac (Center for Computational Research), Benjamin D. Plessinger (Center for Computational Research), Jeffrey T. Palmer (Center for Computational Research), Rudra Chakraborty (Center for Computational Research), Gregory Dean (Center for Computational Research), Martins Innus (Center for Computational Research), Ryan Rathsam (Center for Computational Research), Nikolay Simakov (Center for Computational Research), Joseph P. White (Center for Computational Research), Thomas R. Furlani (Center for Computational Research), Steven M. Gallo (Center for Computational Research), Robert L. DeLeon (Center for Computational Research), Matthew D. Jones (Center for Computational Research), Cynthia Cornelius (Center for Computational Research), and Abani Patra (Center for Computational Research)</i>

## **WRAp'18: Workshop on Representative Applications**

MiniApp for Density Matrix Renormalization Group Hamiltonian Application Kernel .590.....	
	<i>Wael Elwasif (Oak Ridge National Laboratory), Ed. D'Azevedo (Oak Ridge National Laboratory), Arghya Chatterjee (Oak Ridge National Laboratory), Gonzalo Alvarez (Oak Ridge National Laboratory), Oscar Hernandez (Oak Ridge National Laboratory), and Vivek Sarkar (Georgia Institute of Technology Atlanta)</i>
UnSNAP: A Mini-App for Exploring the Performance of Deterministic Discrete Ordinates Transport on Unstructured Meshes .598.....	
	<i>Tom Deakin (University of Bristol), Simon McIntosh-Smith (University of Bristol), Justin Lovegrove (AWE), Richard Smedley-Stevenson (AWE), and Andrew Hagues (AWE)</i>
Towards a Mini-App for Smoothed Particle Hydrodynamics at Exascale .607.....	
	<i>Danilo Guerrero (University of Basel), Rubén M. Cabezon (University of Basel), Jean-Guillaume Piccinali (Swiss National Supercomputing Centre), Aurélien Cavelan (University of Basel), Florina M. Ciorba (University of Basel), David Imbert (Nextflow Software), Lucio Mayer (University of Zürich), and Darren Reed (University of Zürich)</i>
BookLeaf: An Unstructured Hydrodynamics Mini-Application .615.....	
	<i>David Truby (University of Warwick), Steven Wright (University of York), Robert Kevis (Atomic Weapons Establishment), Satheesh Maheswaran (Atomic Weapons Establishment), Andrew Herdman (Atomic Weapons Establishment), and Stephen Jarvis (University of Warwick)</i>

## **REV-A'18: Re-Emergence of Vector Architectures Workshop**

SVE-Enabling Lattice QCD Codes .623.....	
	<i>Nils Meyer (University of Regensburg), Peter Georg (University of Regensburg), Dirk Pleiter (Forschungszentrum Juelich), Stefan Solbrig (University of Regensburg), and Tilo Wettig (University of Regensburg)</i>

Optimizations of COAWST for a Large Simulation on the Earth Simulator .629.....  
*Shivanshu Kumar Singh (NEC Technologies India), Kota Sakakura (NEC Corporation), Sourav Saha (NEC Technologies India), Raghunandan Mathur (NEC Technologies India), Chirag Sharma (NEC Technologies India), Koji Goto (NEC Corporation), Osamu Watanabe (NEC Corporation), and Akihiro Musa (NEC Corporation)*

Simulating SVE-Optimised Genomics Workloads on Gem5 .637.....  
*Javier Setoain (Arm Ltd.), Alejandro Chacon (Arm Ltd.), and Filippo Spiga (Arm Ltd.)*

Performing SVE Studies using the Arm Instruction Emulator .638.....  
*Miguel Tairum Cruz (Arm)*

Stillwater Knowledge Processing Unit (KPU^TM) .639.....  
*Theodore Omtzigt (Stillwater Supercomputing, Inc.)*

Cost Modelling for Vectorization on ARM .644.....  
*Angela Pohl (Technische Universität Berlin), Biagio Cosenza (Technische Universität Berlin), and Ben Juurlink (Technische Universität Berlin)*

The Tofu Interconnect D .646.....  
*Yuichiro Ajima (Fujitsu Limited), Takahiro Kawashima (Fujitsu Limited), Takayuki Okamoto (Fujitsu Limited), Naoyuki Shida (Fujitsu Limited), Kouichi Hirai (Fujitsu Limited), Toshiyuki Shimizu (Fujitsu Limited), Shinya Hiramoto (Fujitsu Limited), Yoshiro Ikeda (Fujitsu Limited), Takahide Yoshikawa (Fujitsu Limited), Kenji Uchida (Fujitsu Limited), and Tomohiro Inoue (Fujitsu Limited)*

**Author Index 655** .....