

# **2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID 2017)**

**Madrid, Spain  
14 – 17 May 2017**

**Pages 1-568**



**IEEE Catalog Number: CFP17276-POD  
ISBN: 978-1-5090-5980-5**

**Copyright © 2017 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:	CFP17276-POD
ISBN (Print-On-Demand):	978-1-5090-5980-5
ISBN (Online):	978-1-5090-6611-7

**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

# 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing

## CCGRID 2017

### Table of Contents

Welcome Message from the General Chairs.....	xvi
Welcome Message from the Program Chairs.....	xviii
Committees.....	xix
Program Committees.....	xxi

---

#### Applications and Big Data (I)

Efficient Event Correlation over Distributed Systems .....	1
<i>Long Cheng, Boudewijn F. van Dongen, and Wil M. P. van der Aalst</i>	
4CeeD: Real-Time Data Acquisition and Analysis Framework for Material-Related Cyber-Physical Environments .....	11
<i>Phuong Nguyen, Steven Konstanty, Todd Nicholson, Thomas O'Brien, Aaron Schwartz-Duval, Timothy Spila, Klara Nahrstedt, Roy H. Campbell, Indranil Gupta, Michael Chan, Kenton Mchenry, and Normand Paquin</i>	
Efficient Cache Update for In-Memory Cluster Computing with Spark .....	21
<i>Li-Yung Ho, Jan-Jan Wu, Pangfeng Liu, Chia-Chun Shih, Chi-Chang Huang, and Chao-Wen Huang</i>	
GPU in-Memory Processing Using Spark for Iterative Computation .....	31
<i>Sumin Hong, Woohyuk Choi, and Won-Ki Jeong</i>	

#### Performance Modeling and Evaluation (I)

Optimal Resource Configuration of Complex Services in the Cloud .....	42
<i>Abhinandan S. Prasad, David Koll, Jesus Omana Iglesias, Jordi Arjona Aroca, Volker Hilt, and Xiaoming Fu</i>	
Modeling Distributed Platforms from Application Traces for Realistic File Transfer Simulation .....	54
<i>Anchen Chai, Mohammad-Mahdi Bazm, Sorina Camarasu-Pop, Tristan Glatard, Hugues Benoit-Cattin, and Frédéric Suter</i>	
A Comparison of Reinforcement Learning Techniques for Fuzzy Cloud Auto-Scaling .....	64
<i>Hamid Arabnejad, Claus Pahl, Pooyan Jamshidi, and Giovani Estrada</i>	

Performance Modelling and Cost Effective Execution for Distributed Graph Processing on Configurable VMs .....	74
<i>Zengxiang Li, Bowen Zhang, Shen Ren, Yong Liu, Zheng Qin, Rick Siow Mong Goh, and Mohan Gurusamy</i>	

Fine-Grained Nested Virtual Machine Performance Analysis through First Level Hypervisor Tracing .....	84
<i>Hani Nemati, Suchakrapani Datt Sharma, and Michel R. Dagenais</i>	

## **Scheduling and Resource Management (I)**

Adaptive Hybrid Queue Configuration for Supercomputer Systems .....	90
<i>Vineetha Kondameedi and Sathish Vadhiyar</i>	

Flexible Scheduling of Distributed Analytic Applications .....	100
<i>Francesco Pace, Daniele Venzano, Damiano Carra, and Pietro Michiardi</i>	

CtrlCloud: Performance-Aware Adaptive Control for Shared Resources in Clouds .....	110
<i>Omer Adam, Young Choon Lee, and Albert Y. Zomaya</i>	

QoS-Aware Virtual Infrastructures Allocation on SDN-Based Clouds .....	120
<i>Felipe Rodrigo de Souza, Charles Christian Miers, Adriano Fiorese, and Guilherme Piegas Koslovski</i>	

## **Scheduling and Resource Management (II)**

Deploying High Throughput Scientific Workflows on Container Schedulers with Makeflow and Mesos .....	130
<i>Chao Zheng, Ben Tovar, and Douglas Thain</i>	

A New On-line Method for Scheduling Independent Tasks .....	140
<i>Giorgio Lucarelli, Fernando Mendonca, and Denis Trystram</i>	

A Two-Stage Multi-objective Optimization of Erasure Coding in Overlay Networks .....	150
<i>Nishant Saurabh, Dragi Kimovski, Francesco Gaetano, and Radu Prodan</i>	

Multi-dimensional Admission Control and Capacity Planning for IaaS Clouds with Multiple Service Classes .....	160
<i>Marcus Carvalho, Francisco Brasileiro, Raquel Lopes, Giovanni Farias, Alessandro Fook, João Mafra, and Daniel Turull</i>	

A Robust Tabu Search Heuristic for VM Consolidation under Demand Uncertainty in Virtualized Datacenters .....	170
<i>Robayet Nasim and Andreas J. Kessler</i>	

## **Mobile, Hybrid and Emerging Clouds**

Evaluation of Data Locality Strategies for Hybrid Cloud Bursting of Iterative MapReduce .....	181
<i>Francisco J. Clemente-Castello, Bogdan Nicolae, M. Mustafa Rafique, Rafael Mayo, and Juan Carlos Fernandez</i>	

Leveraging Renewable Energy in Edge Clouds for Data Stream Analysis in IoT .....	186
<i>Yunbo Li, Anne-Cécile Orgerie, Ivan Roderio, Manish Parashar, and Jean-Marc Menaud</i>	

DOTA: Delay Bounded Optimal Cloudlet Deployment and User Association in WMANs .....	196
<i>Longjie Ma, Jigang Wu, and Long Chen</i>	

## **Applications and Big Data (II)**

Mitigating YARN Container Overhead with Input Splits .....	204
<i>Wonbae Kim, Young-Ri Choi, and Beomseok Nam</i>	
Parallel Variable Selection for Effective Performance Prediction .....	208
<i>Jonathan Wang, Wucherl Yoo, Alex Sim, Peter Nugent, and Kesheng Wu</i>	
Towards Big Data Analytics across Multiple Clusters .....	218
<i>Dongyao Wu, Sherif Sakr, Liming Zhu, and Huijun Wu</i>	

## **Architecture and Networking**

APHiD: Hierarchical Task Placement to Enable a Tapered Fat Tree Topology for Lower Power and Cost in HPC Networks .....	228
<i>George Michelogiannakis, Khaled Z. Ibrahim, John Shalf, Jeremiah J. Wilke, Samuel Knight, and Joseph P. Kenny</i>	
Swift-X: Accelerating OpenStack Swift with RDMA for Building an Efficient HPC Cloud .....	238
<i>Shashank Gugrani, Xiaoyi Lu, and Dhableswar K. (DK) Panda</i>	
Offloading Communication Control Logic in GPU Accelerated Applications .....	248
<i>Elena Agostini, Davide Rossetti, and Sreeram Potluri</i>	
Preliminary Performance Analysis of Multi-rail Fat-Tree Networks .....	258
<i>Noah Wolfe, Misbah Mubarak, Nikhil Jain, Jens Domke, Abhinav Bhatele, Christopher D. Carothers, and Robert B. Ross</i>	
SynAPTIC: Secure and Persistent Connectivity for Containers .....	262
<i>Alireza Ranjbar, Miika Komu, Patrik Salmela, and Tuomas Aura</i>	

## **Scheduling and Resource Management (III)**

CloudSight: A Tenant-Oriented Transparency Framework for Cross-Layer Cloud Troubleshooting .....	268
<i>Hyunwook Baek, Abhinav Srivastava, and Jacobus Van der Merwe</i>	
Enabling Distributed Software-Defined Environments Using Dynamic Infrastructure Service Composition .....	274
<i>Moustafa Abdelbaky, Javier Diaz-Montes, Merve Unuvar, Melissa Romanus, Ivan Rodero, Malgorzata Steinder, and Manish Parashar</i>	
CBase: A New Paradigm for Fast Virtual Machine Migration across Data Centers .....	284
<i>Fei Zhang, Xiaoming Fu, and Ramin Yahyapour</i>	

## Programming Models and Runtime Systems (I)

Preemptive Software Transactional Memory .....	294
<i>Emiliano Silvestri, Simone Economo, Pierangelo Di Sanzo, Alessandro Pellegrini, and Francesco Quaglia</i>	
Supporting Fault-Tolerance in Presence of In-Situ Analytics .....	N/A
<i>Jiaqi Liu and Gagan Agrawal</i>	
Advanced Thread Synchronization for Multithreaded MPI Implementations .....	314
<i>Hoang-Vu Dang, Sangmin Seo, Abdelhalim Amer, and Pavan Balaji</i>	

## Performance Modeling and Evaluation (II)

Unveiling the Interplay Between Global Link Arrangements and Network Management	
Algorithms on Dragonfly Networks .....	325
<i>Fulya Kaplan, Ozan Tuncer, Vitus J. Leung, Scott K. Hemmert, and Ayse K. Coskun</i>	
Application-Agnostic Power Monitoring in Virtualized Environments .....	335
<i>James Phung, Young Choon Lee, and Albert Y. Zomaya</i>	
A Performance Study of UCX over InfiniBand .....	345
<i>Nikela Papadopoulou, Lena Oden, and Pavan Balaji</i>	
Performance Modelling and Verification of Cloud-Based Auto-Scaling Policies .....	355
<i>Alexandros Evangelidis, David Parker, and Rami Bahsoon</i>	

## Scheduling and Resource Management (IV)

Joint Optimization of Scaling and Placement of Virtual Network Services .....	365
<i>Sevil Drăxler, Holger Karl, and Zoltán Ádám Mann</i>	
Acyclic Partitioning of Large Directed Acyclic Graphs .....	371
<i>Julien Herrmann, Jonathan Kho, Bora Uçar, Kamer Kaya, and Ümit V. Çatalyürek</i>	
Towards Energy Budget Control in HPC .....	381
<i>Pierre-François Dutot, Yiannis Georgiou, David Glessner, Laurent Lefevre, Millian Poquet, and Issam Rais</i>	
On Estimating Minimum Bids for Amazon EC2 Spot Instances .....	391
<i>Markus Lumpe, Mohan Baruwal Chhetri, Quoc Bao Vo, and Ryszard Kowalczyk</i>	
Practical Service Placement Approach for Microservices Architecture .....	401
<i>Mennan Selimi, Llorenç Cerdà-Alabern, Marc Sánchez-Artigas, Felix Freitag, and Luís Veiga</i>	

## Security, Privacy and Reliability (I)

Combating the Bandits in the Cloud: A Moving Target Defense Approach .....	411
<i>Terry Penner and Mina Guirguis</i>	
Chrysaor: Fine-Grained, Fault-Tolerant Cloud-of-Clouds MapReduce .....	421
<i>Pedro A. R. S. Costa, Fernando M. V. Ramos, and Miguel Correia</i>	
An Anomaly Detection Fabric for Clouds Based on Collaborative VM Communities .....	431
<i>Rashid Tahir, Ali Raza, Mazhar Naqvi, Fareed Zaffar, and Matthew Caesar</i>	

LOGAIDER: A Tool for Mining Potential Correlations of HPC Log Events .....	442
<i>Sheng Di, Rinku Gupta, Marc Snir, Eric Pershey, and Franck Cappello</i>	
Designing and Modelling Selective Replication for Fault-Tolerant HPC Applications .....	452
<i>Omer Subasi, Gulay Yalcin, Ferad Zyulkyarov, Osman Unsal, and Jesus Labarta</i>	

## **Security, Privacy and Reliability (II)**

Security Implications of Blockchain Cloud with Analysis of Block Withholding Attack .....	458
<i>Deepak K. Tosh, Sachin Shetty, Xueping Liang, Charles A. Kamhoua, Kevin A. Kwiat, and Laurent Njilla</i>	
ProvChain: A Blockchain-Based Data Provenance Architecture in Cloud Environment with Enhanced Privacy and Availability .....	468
<i>Xueping Liang, Sachin Shetty, Deepak Tosh, Charles Kamhoua, Kevin Kwiat, and Laurent Njilla</i>	
T-VMI: Trusted Virtual Machine Introspection in Cloud Environments .....	478
<i>Lina Jia, Min Zhu, and Bibo Tu</i>	
Modeling Correlation between QoS Attributes for Trust Computation in Cloud Computing Environments .....	488
<i>Manel Mrabet, Yosra Ben Saied, and Leila Azouz Saidane</i>	
Crowdsourced Data Integrity Verification for Key-Value Stores in the Cloud .....	498
<i>Grisha Weintraub and Ehud Gudes</i>	

## **Performance Modeling and Evaluation (III)**

Maximum Sustainable throughput Prediction for Data Stream Processing over Public Clouds .....	504
<i>Shigeru Imai, Stacy Patterson, and Carlos A. Varela</i>	
WattsKit: Software-Defined Power Monitoring of Distributed Systems .....	514
<i>Maxime Colmant, Pascal Felber, Romain Rouvoy, and Lionel Seinturier</i>	
Predicting Cloud Performance for HPC Applications: A User-Oriented Approach .....	524
<i>Giovanni Mariani, Andreea Anghel, Rik Jongerius, and Gero Dittmann</i>	
An Approach and Case Study of Cloud Instance Type Selection for Multi-tier Web Applications .....	534
<i>Christian Davatz, Christian Inzinger, Joel Scheuner, and Philipp Leitner</i>	
Toward a Holistic Framework for Conducting Scientific Evaluations of OpenStack .....	544
<i>Ronan-Alexandre Cherrueau, Dimitri Petin, Anthony Simonet, Adrien Lebre, and Matthieu Simonin</i>	

## **Storage and I/O (I)**

Dynamic Management of In-Memory Storage for Efficiently Integrating Compute-and-Data-Intensive Computing on HPC Systems .....	549
<i>Pengfei Xuan, Feng Luo, Rong Ge, and Pradip K. Srimani</i>	
High-Performance Key-Value Store on OpenSHMEM .....	559
<i>Huansong Fu, Manjunath Gorentla Venkata, Ahana Roy Choudhury, Neena Imam, and Weikuan Yu</i>	
Energy-Efficient I/O Thread Schedulers for NVMe SSDs on NUMA .....	569
<i>Junjie Qian, Hong Jiang, Witawas Srisa-An, Sharad Seth, Stan Skelton, and Joseph Moore</i>	

## Data Centers and Cyberinfrastructure

Energy Efficient Algorithm for VNF Placement and Chaining .....	579
<i>Oussama Soualah, Marouen Mechtri, Chaima Ghribi, and Djamel Zeglache</i>	
KPI-Agnostic Control for Fine-Grained Vertical Elasticity .....	589
<i>Ewnetu Bayuh Lakew, Alessandro Vittorio Papadopoulos, Martina Maggio, Cristian Klein, and Erik Elmroth</i>	
PCSSampler: Sample-Based, Private-State Cluster Scheduling .....	599
<i>Chunliang Hao, Jie Shen, Celia Chen, Heng Zhang, Yanjun Wu, and Mingshu Li</i>	
Optimized Cloud Deployment of Multi-tenant Software Considering Data Protection Concerns .....	609
<i>Zoltán Ádám Mann and Andreas Metzger</i>	

## Programming Models and Runtime Systems (II)

AnalyzeThat: A Programmable Shared-Memory System for an Array of Processing-In-Memory Devices .....	619
<i>Sangkuen Lee, Hyogi Sim, Youngjae Kim, and Sudharshan S. Vazhkudai</i>	
Implementation and Evaluation of One-Sided PGAS Communication in XcalableACC for Accelerated Clusters .....	625
<i>Akihiro Tabuchi, Masahiro Nakao, Hitoshi Murai, Taisuke Boku, and Mitsuhsa Sato</i>	
Combining Both a Component Model and a Task-Based Model for HPC Applications: A Feasibility Study on GYSELA .....	635
<i>Olivier Aumage, Julien Bigot, H�el�ene Coullon, Christian P�erez, and J�er�ome Richard</i>	

## Storage and I/O (II)

Pattern-Directed Replication Scheme for Heterogeneous Object-Based Storage .....	645
<i>Jiang Zhou, Wei Xie, Dong Dai, and Yong Chen</i>	
A New File System I/O Mode for Efficient User-Level Caching .....	649
<i>Jiwoong Park, Cheolgi Min, and HeonYoung Yeom</i>	
COPS: Cost Based Object Placement Strategies on Hybrid Storage System for DBaaS Cloud .....	659
<i>Djillali Boukhelef, Kamel Boukhalfa, Jalil Boukhobza, Hamza Ouarnoughi, and Laurent Lemarchand</i>	

## Scale Challenge

Scalable Assembly for Massive Genomic Graphs .....	665
<i>Jintao Meng, Ning Guo, Jianqiu Ge, Yanjie Wei, Pavan Balaji, and Bingqiang Wang</i>	
mD3DOCKxb: An Ultra-Scalable CPU-MIC Coordinated Virtual Screening Framework .....	671
<i>Shaoliang Peng, Xiaoyu Zhang, Shunyun Yang, Wenhe Su, Zhiqiang Zhang, Dong Dong, Kai Lu, Yutong Lu, Xiangke Liao, Bertil Schmidt, Weiliang Zhu, and Kuan-Ching Li</i>	
Scaling a Convolutional Neural Network for Classification of Adjective Noun Pairs with TensorFlow on GPU Clusters .....	677
<i>V�ctor Campos, Francesc Sastre, Maurici Yag�es, Jordi Torres, and Xavier Gir�-I-Nieto</i>	
Scaling HDFS to More Than 1 Million Operations Per Second with HopsFS .....	683
<i>Mahmoud Ismail, Salman Niazi, Mikael Ronstr�m, Seif Haridi, and Jim Dowling</i>	



Massively Parallel Simulations of Spread of Infectious Diseases over Realistic Social Networks .....	689
<i>Abhinav Bhatele, Jae-Seung Yeom, Nikhil Jain, Chris J. Kuhlman, Yarden Livnat, Keith R. Bisset, Laxmikant V. Kale, and Madhav V. Marathe</i>	

## Doctoral Symposium

Enhancing the rCUDA Remote GPU Virtualization Framework: From a Prototype to a Production Solution .....	695
<i>C. Reaño, F. Silla, and J. Duato</i>	
Cable-Geometric Error-Prone Approach for Low-Latency Interconnection Networks .....	699
<i>Truong Thao Nguyen and Michihiro Koibuchi</i>	
Towards Distributed Software-Defined Environments .....	703
<i>Moustafa Abdelbaky, Javier Diaz-Montes, and Manish Parashar</i>	
Dynamic Resource Management Across Cloud-Edge Resources for Performance-Sensitive Applications .....	707
<i>Shashank Shekhar and Aniruddha Gokhale</i>	
Optimization of Checkpoints and Execution Model for an Implementation of OpenMP on Distributed Memory Architectures .....	711
<i>Van Long Tran, Eric Renault, and Viet Hai Ha</i>	
Secure Cloud Storage Service for Detection of Security Violations .....	715
<i>Carlos André Batista de Carvalho, Miguel Franklin De Castro, and Rossana Maria De Castro Andrade</i>	
Data-Aware Support for Hybrid HPC and Big Data Applications .....	719
<i>Silvina Caíno-Lores, Florin Isaila, and Jesús Carretero</i>	

## Posters and Research Demos

Performance Optimization by Dynamically Altering Cache Replacement Algorithm in CPU-GPU Heterogeneous Multi-core Architecture .....	723
<i>Juan Fang, Qingwen Fan, Xiaoting Hao, Yanjin Cheng, and Lijun Sun</i>	
Extending Message Passing Interface Windows to Storage .....	727
<i>Sergio Rivas-Gomez, Stefano Markidis, Ivy Bo Peng, Erwin Laure, Gokcen Kestor, and Roberto Gioiosa</i>	
TuNao: A High-Performance and Energy-Efficient Reconfigurable Accelerator for Graph Processing .....	731
<i>Jinhong Zhou, Shaoli Liu, Qi Guo, Xuda Zhou, Tian Zhi, Daofu Liu, Chao Wang, Xuehai Zhou, Yunji Chen, and Tianshi Chen</i>	
A Live Demo for Showing the Benefits of Applying the Remote GPU Virtualization Technique to Cloud Computing .....	735
<i>Javier Prades and Federico Silla</i>	
Load and Video Performance Patterns of a Cloud Based WebRTC Architecture .....	739
<i>Vamis Xhagjika, Òscar Divorra Escoda, Leandro Navarro, and Vladimir Vlassov</i>	

Lemonade: A Scalable and Efficient Spark-Based Platform for Data Analytics .....	745
<i>Walter dos Santos, Luiz F. M. Carvalho, Gustavo de P. Avelar, Átila Silva Jr, Lucas M. Ponce, Dorgival Guedes, and Wagner Meira Jr.</i>	
IBM Research Hybrid Cloud Poster .....	749
<i>Lorraine M. Herger and Carlos A. Fonseca</i>	
Using the Jetstream Research Cloud to Provide Science Gateway Resources .....	753
<i>Richard Knepper, Eric Coulter, Marlon Pierce, Suresh Marru, and Sudhakar Pamidighantam</i>	
DSA: Scalable Distributed Sequence Alignment System Using SIMD Instructions .....	758
<i>Bo Xu, Changlong Li, Hang Zhuang, Jiali Wang, Qingfeng Wang, Jinhong Zhou, and Xuehai Zhou</i>	
AURA: Recovering from Transient Failures in Cloud Deployments .....	762
<i>Ioannis Giannakopoulos, Ioannis Konstantinou, Dimitrios Tsumakos, and Nectarios Koziris</i>	
BBQ: Elastic MapReduce over Cloud Platforms .....	766
<i>Nikolaos Chalvantzis, Ioannis Konstantinou, and Nektarios Koziris</i>	
Multi-Agent Recommendation System in Internet of Things .....	772
<i>Agostino Forestiero</i>	
Massive Data Load on Distributed Database Systems over HBase .....	776
<i>Ainhoa Azqueta-Alzúaz, Marta Patiño-Martinez, Ivan Brondino, and Ricardo Jimenez-Peris</i>	
Mermaid: Integrating Vertex-Centric with Edge-Centric for Real-World Graph Processing .....	780
<i>Jinhong Zhou, Chongchong Xu, Xianglan Chen, Chao Wang, and Xuehai Zhou</i>	
Representing Variant Calling Format as Directed Acyclic Graphs to Enable the Use of Cloud Computing for Efficient and Cost Effective Genome Analysis .....	784
<i>Sanna Aizad, Ashiq Anjum, and Rizos Sakellariou</i>	
Techniques for Handling Error in User-Estimated Execution Times During Resource Management on Systems Processing MapReduce Jobs .....	788
<i>Norman Lim, Shikharesh Majumdar, and Peter Ashwood-Smith</i>	
EffiEye: Application-Aware Large Flow Detection in Data Center .....	794
<i>Binfeng Wang, Jinshu Su, Lin Chen, Jinsheng Deng, and Long Zheng</i>	

## Workshop Papers

### Workshop on Clusters, Clouds and Grids for Life Sciences (CCGrid – Life 2017)

Using the Cloud for Parameter Estimation Problems: Comparing Spark vs MPI with a Case-Study .....	797
<i>Patricia González, Xoan C. Pardo, David R. Penas, Diego Teijeiro, Julio R. Banga, and Ramón Doallo</i>	
Analog-Digital Approach in Human Brain Modeling .....	807
<i>Alexander Bogdanov, Alexander Degtyarev, Dmitriy Guschanskiy, Kirill Lysov, Nataliya Ananieva, Nataliya Zalutskaya, and Nikolay Neznanov</i>	
Fine-Grained Supervision and Restriction of Biomedical Applications in Linux Containers .....	813
<i>Michael Witt, Christoph Jansen, Dagmar Krefting, and Achim Streit</i>	

BIOPET: Towards Scalable, Maintainable, User-Friendly, Robust and Flexible NGS Data Analysis Pipelines .....	823
<i>Peter Van't Hof, Wibowo Arindrarto, Sander Bollen, Szymon Kielbasa, Jeroen Laros, and Hailiang Mei</i>	
Medical Imaging Processing on a Big Data Platform Using Python: Experiences with Heterogeneous and Homogeneous Architectures .....	830
<i>Estefania Serrano, Javier Garcia Blas, Jesus Carretero, Monica Abella, and Manuel Desco</i>	

### **3rd International Workshop on Scalable Computing for Real-Time Big Data Applications (SCRAMBL 2017)**

Scheduling Data Stream Jobs on Distributed Systems with Background Load .....	838
<i>Anca Vulpe and Marc Frincu</i>	
Cloud Resource Scaling for Big Data Streaming Applications Using a Layered Multi-dimensional Hidden Markov Model .....	848
<i>Olubisi Runsewe and Nancy Samaan</i>	
Cost Model and Analysis of Iterative MapReduce Applications for Hybrid Cloud Bursting .....	858
<i>Francisco J. Clemente-Castelló, Rafael Mayo, and Juan Carlos Fernández</i>	

### **2nd International Workshop on Distributed Big Data Management (DBDM 2017)**

Apply Block Index Technique to Scientific Data Analysis and I/O Systems .....	865
<i>Tzuhsien Wu, Jerry Chou, Norbert Podhorszki, Junmin Gu, Yuan Tian, Scott Klasky, and Kesheng Wu</i>	
Smart RDF Data Storage in Graph Databases .....	872
<i>Roberto De Virgilio</i>	
A Level-Wise Load Balanced Scientific Workflow Execution Optimization Using NSGA-II .....	882
<i>Phyo Thandar Thant, Courtney Powell, Martin Schlueter, and Masaharu Munetomo</i>	
RAPID: A Fast Data Update Protocol in Erasure Coded Storage Systems for Big Data .....	890
<i>G. J. Akash, Ojus Thomas Lee, S. D. Madhu Kumar, Priya Chandran, and Alfredo Cuzzocrea</i>	
MapReduce-Based Algorithms for Managing Big RDF Graphs: State-of-the-Art Analysis, Paradigms, and Future Directions .....	898
<i>Alfredo Cuzzocrea, Rajkumar Buyya, Vincenzo Passanisi, and Giovanni Pilato</i>	
Optimized MapFile Based Storage of Small Files in Hadoop .....	906
<i>Shalini Sheoran, Divyasikha Sethia, and Huzur Saran</i>	
An Accuracy-Aware Implementation of Two-Point Three-Dimensional Correlation Function Using Bin-Recycling Strategy on GPU .....	913
<i>Iván Méndez-Jiménez, Miguel Cárdenas-Montes, Juan José Rodríguez-Vázquez, Ignacio Sevilla Noarbe, Eusebio Sánchez Álvaro, Miguel A. Vega-Rodríguez, and David Alonso</i>	
A Big Data Architecture for Automotive Applications: PSA Group Deployment Experience .....	921
<i>Amir Haroun, Ahmed Mostefaoui, and François Dessables</i>	

## **Sixth IEEE International Workshop on Cloud Computing Interclouds, Multiclouds, Federations, and Interoperability (Intercloud 2017)**

Improving Resource Efficiency of Container-Instance Clusters on Clouds .....	929
<i>Uchechukwu Awada and Adam Barker</i>	
Hybrid Mobile Edge Computing: Unleashing the Full Potential of Edge Computing in Mobile Device Use Cases .....	935
<i>Andreas Reiter, Bernd Prünster, and Thomas Zefferer</i>	
Defining Intercloud Security Framework and Architecture Components for Multi-cloud Data Intensive Applications .....	945
<i>Yuri Demchenko, Fatih Turkmen, Mathias Slawik, and Cees de Laat</i>	

### **Short Paper**

## **2nd International Workshop on Theoretical Approaches to Performance Evaluation, Modeling and Simulation (TAPEMS 2017)**

Analyzing the Parallel I/O Severity of MPI Applications .....	953
<i>Sandra Mendez, Dolores Rexachs, and Emilio Luque</i>	
Formal Modeling and Performance Evaluation of a Run-Time Rank Remapping Technique in Broadcast, Allgather and Allreduce MPI Collective Operations .....	963
<i>Jesús M. Álvarez-Llorente, Juan C. Díaz-Martín, and Juan A. Rico-Gallego</i>	
Automatic Adaption of the Sampling Frequency for Detailed Performance Analysis .....	973
<i>Michael Wagner and Andreas Knüepfer</i>	
Performance Models for Communication in Collective I/O Operations .....	982
<i>Shweta Jha and Edgar Gabriel</i>	
Reducing Load Imbalance of Virtual Clusters via Reconfiguration and Adaptive Job Scheduling .....	992
<i>Sina Mahmoodi Khorandi, Siavash Ghiasvand, and Mohsen Sharifi</i>	
Empirical Mode Decomposition for Modeling of Parallel Applications on Intel Xeon Phi Processors .....	1000
<i>Gary Lawson, Masha Sosonkina, Tal Ezer, and Yuzhong Shen</i>	
Energy Model for Low-Power Cluster .....	1009
<i>Edson Flórez, Johnatan E. Pecero, Joseph Emeras, and Carlos J. Barrios</i>	
Near-Optimal Policies for Energy-Aware Task Assignment in Server Farms .....	1017
<i>Misikir Eyob Gebrehiwot, Samuli Aalto, and Pasi Lassila</i>	

## **Workshop on the Integration of Extreme Scale Computing and Big Data Management and Analytics (EBDMA 2017)**

An Empirical Evaluation of How the Network Impacts the Performance and Energy Efficiency in RAMCloud .....	1027
<i>Yacine Taleb, Shadi Ibrahim, Gabriel Antoniu, and Toni Cortes</i>	
On the Use of In-memory Analytics Workflows to Compute eScience Indicators from Large Climate Datasets .....	1035
<i>Alessandro D’Anca, Cosimo Palazzo, Donatello Elia, Sandro Fiore, Ioannis Bistinas, Kristin Böttcher, Victoria Bennett, and Giovanni Aloisio</i>	
Exploring Shared State in Key-Value Store for Window-Based Multi-pattern Streaming Analytics .....	1044
<i>Ovidiu-Cristian Marcu, Radu Tudoran, Bogdan Nicolae, Alexandru Costan, Gabriel Antoniu, and María S. Pérez-Hernández</i>	
Evaluation of HPC-Big Data Applications Using Cloud Platforms .....	1053
<i>Shweta Salaria, Kevin Brown, Hideyuki Jitsumoto, and Satoshi Matsuoka</i>	
A Data-Driven Approach Based on Auto-Regressive Models for Energy-Efficient Clouds .....	1062
<i>Albino Altomare and Eugenio Cesario</i>	

## **International Workshop on Assured Cloud Computing and QoS aware Big Data (WACC 2017)**

Automatic Consolidation of Virtual Machines in On-premises Cloud Platforms .....	1070
<i>Carlos de Alfonso, Ignacio Blanquer, Germán Moltó, and Miguel Caballer</i>	
A Game-Theoretic Approach for Runtime Capacity Allocation in MapReduce .....	1080
<i>Eugenio Gianniti, Danilo Ardagna, Michele Ciavotta, and Mauro Passacantando</i>	
IT Security and Privacy Standards in Comparison: Improving FedRAMP Authorization for Cloud Service Providers .....	1090
<i>Carlo Di Giulio, Read Sprabery, Charles Kamhoua, Kevin Kwiat, Roy Campbell, and Masooda N. Bashir</i>	
A Lightweight MapReduce Framework for Secure Processing with SGX .....	1100
<i>Rafael Pires, Daniel Gavril, Pascal Felber, Emanuel Onica, and Marcelo Pasin</i>	
PRIVAAaaS: Privacy Approach for a Distributed Cloud-Based Data Analytics Platforms .....	1108
<i>Tania Basso, Regina Moraes, Nuno Antunes, Marco Vieira, Walter Santos, and Wagner Meira Jr.</i>	
Evaluating the Performance of Continuous Test-Based Cloud Service Certification .....	1117
<i>Philipp Stephanow and Christian Banse</i>	
A Game Theoretic Method for VM-to-Hypervisor Attacks Detection in Cloud Environment .....	1127
<i>Amin Nezarat</i>	

## **Author Index**