

**2015 15th IEEE/ACM  
International Symposium on  
Cluster, Cloud and Grid  
Computing**

**(CCGrid 2015)**

**Shenzhen, China  
4 – 7 May 2015**

**Pages 1-616**



**IEEE Catalog Number: CFP15276-POD  
ISBN: 978-1-4799-8007-9**

# 2015 IEEE/ACM 15th International Symposium on Cluster, Cloud and Grid Computing

## CCGrid 2015

### Table of Contents

Message from the CCGrid 2015 General Co-Chairs.....	xix
Message from the Program Co-Chairs.....	xxi
Organizing Committee.....	xxii
Program Committee.....	xxiv
Steering Committee.....	xxx

---

### 2015 IEEE/ACM 15th International Symposium on Cluster, Cloud and Grid Computing

#### Best Paper Nominees

Analyzing the Impact of CPU Pinning and Partial CPU Loads on Performance and Energy Efficiency .....	1
<i>Andrej Podzimek, Lubomir Bulej, Lydia Y. Chen, Walter Binder, and Petr Tuma</i>	
Deferred Lightweight Indexing for Log-Structured Key-Value Stores .....	11
<i>Yuzhe Tang, Arun Iyengar, Wei Tan, Liana Fong, Ling Liu, and Balaji Palanisamy</i>	
Cowic: A Column-Wise Independent Compression for Log Stream Analysis .....	21
<i>Hao Lin, Jingyu Zhou, Bin Yao, Minyi Guo, and Jie Li</i>	
Dependable Horizontal Scaling Based on Probabilistic Model Checking .....	31
<i>Athanasios Naskos, Emmanouela Stachtari, Anastasios Gounaris, Panagiotis Katsaros, Dimitrios Tsoumakos, Ioannis Konstantinou, and Spyros Sioutas</i>	

#### Cloud Computing I

Towards Efficient Work-Stealing in Virtualized Environments .....	41
<i>Yaqiong Peng, Song Wu, and Hai Jin</i>	
Traffic-Sensitive Live Migration of Virtual Machines .....	51
<i>Umesh Deshpande and Kate Keahey</i>	

Checkpointing as a Service in Heterogeneous Cloud Environments .....	61
<i>Jiajun Cao, Matthieu Simonin, Gene Cooperman, and Christine Morin</i>	

MVAPICH2 over OpenStack with SR-IOV: An Efficient Approach to Build HPC Clouds .....	71
<i>Jie Zhang, Xiaoyi Lu, Mark Arnold, and Dhabaleswar K. (DK) Panda</i>	

## **Improving and Modeling IO Performance**

F/M-CIP: Implementing Flash Memory Cache Using Conservative Insertion and Promotion .....	81
<i>Jing Yang and Qing Yang</i>	

A Resource Allocation Model for Hybrid Storage Systems .....	91
<i>Hui Wang and Peter Varman</i>	

Triple-H: A Hybrid Approach to Accelerate HDFS on HPC Clusters with Heterogeneous Storage Architecture .....	101
<i>Nusrat Sharmin Islam, Xiaoyi Lu, Md. Wasi-ur-Rahman, Dipti Shankar, and Dhabaleswar K. (DK) Panda</i>	

PIONEER: A Solution to Parallel I/O Workload Characterization and Generation .....	111
<i>Weiping He, David H.C. Du, and Sai B. Narasimhamurthy</i>	

## **Architecture**

Revisiting ILP Designs for Throughput-Oriented GPGPU Architecture .....	121
<i>Ping Xiang, Yi Yang, Mike Mantor, Norm Rubin, and Huiyang Zhou</i>	

Non-Blocking PMI Extensions for Fast MPI Startup .....	131
<i>Sourav Chakraborty, Hari Subramoni, Adam Moody, Akshay Venkatesh, Jonathan Perkins, and Dhabaleswar K. (DK) Panda</i>	

Reconfigurations for Processor Arrays with Faulty Switches and Links .....	141
<i>Jigang Wu, Longting Zhu, Peilan He, and Guiyuan Jiang</i>	

## **Virtualization**

Network-Constrained Packing of Brokered Workloads in Virtualized Environments .....	149
<i>Christine Bassem and Azer Bestavros</i>	

A Virtual Machine Placement Taxonomy .....	159
<i>Fabio López Pires and Benjamín Barán</i>	

Optimizing Soft Real-Time Scheduling Performance for Virtual Machines with SRT-Xen .....	169
<i>Kun Cheng, Yuebin Bai, Rui Wang, and Yao Ma</i>	

## **Advanced Cyberinfrastructure**

Integrating Software Defined Networks within a Cloud Federation .....	179
<i>Ioan Petri, Mengsong Zou, Ali Reza Zamani, Javier Diaz-Montes, Omer Rana, and Manish Parashar</i>	
Partition-Aware Routing to Improve Network Isolation in Infiniband Based Multi-tenant Clusters .....	189
<i>Feroz Zahid, Ernst Gunnar Gran, Bartosz Bogdanski, Bjørn Dag Johnsen, and Tor Skeie</i>	
A Novel Query Caching Scheme for Dynamic InfiniBand Subnets .....	199
<i>Evangelos Tasoulas, Ernst Gunnar Gran, Bjørn Dag Johnsen, and Tor Skeie</i>	
Discovering and Leveraging Content Similarity to Optimize Collective on-Demand Data Access to IaaS Cloud Storage .....	211
<i>Bogdan Nicolae, Andrzej Kochut, and Alexei Karve</i>	

## **Performance Modeling**

Predicting and Mitigating Jobs Failures in Big Data Clusters .....	221
<i>Andrea Rosà, Lydia Y. Chen, and Walter Binder</i>	
Modeling Cross-Architecture Co-Tenancy Performance Interference .....	231
<i>Wei Kuang, Laura E. Brown, and Zhenlin Wang</i>	
Taming Latency in Data Center Networking with Erasure Coded Files .....	241
<i>Yu Xiang, Vaneet Aggarwal, Yih-Farn R. Chen, and Tian Lan</i>	
Adding Storage Simulation Capacities to the SimGrid Toolkit: Concepts, Models, and API .....	251
<i>Adrien Lebre, Arnaud Legrand, Frédéric Suter, and Pierre Veyre</i>	

## **Programmability and Fault Tolerance**

Power-Check: An Energy-Efficient Checkpointing Framework for HPC Clusters .....	261
<i>Raghunath Raja Chandrasekar, Akshay Venkatesh, Khaled Hamidouche, and Dhabaleswar K. (DK) Panda</i>	
An Efficient Silent Data Corruption Detection Method with Error-Feedback Control and Even Sampling for HPC Applications .....	271
<i>Sheng Di, Eduardo Berrocal, and Franck Cappello</i>	
Log-Structured Global Array for Efficient Multi-Version Snapshots .....	281
<i>Hajime Fujita, Nan Dun, Zachary A. Rubenstein, and Andrew A. Chien</i>	
Towards a High Level Programming Paradigm to Deploy e-Science Applications with Dynamic Workflows on Large Scale Distributed Systems .....	292
<i>Mohamed Ben Belgacem and Nabil Abdennadher</i>	

## BigData Applications

Parallel In Situ Detection of Connected Components in Adaptive Mesh Refinement Data .....	302
<i>Xiaocheng Zou, Kesheng Wu, David A. Boyuka II, Daniel F. Martin, Suren Byna, Houjun Tang, Kushal Bansal, Terry J. Ligocki, Hans Johansen, and Nagiza F. Samatova</i>	
A Multi-GPU Hitting Set Algorithm for GRNs Inference .....	313
<i>Danilo Carastan-Santos, Raphael Yokoingawa de Camargo, David Corrêa Martins-Jr., Siang Wun Song, Luiz Carlos Silva Rozante, and Fabrizio Ferreira Borelli</i>	
Parallel Clustering of High-Dimensional Social Media Data Streams .....	323
<i>Xiaoming Gao, Emilio Ferrara, and Judy Qiu</i>	
Optimizing the Bayesian Inference of Phylogeny on Graphic Processors .....	333
<i>Cheng Ling, Chunbao Zhou, Arong Luo, Guoguang Zhao, Tsuyoshi Hamada, and Xiaoyan Zhu</i>	

## Storage Reliability and Security

SIRF-1: Enhancing Reliability of Single Flash SSD through Internal Mirroring for Mission-Critical Mobile Applications .....	343
<i>Michael S. MacFadden, Richard Shelby, and Tao Xie</i>	
CloudSky: A Controllable Data Self-Destruction System for Untrusted Cloud Storage Networks .....	352
<i>Lingfang Zeng, Yang Wang, and Dan Feng</i>	
Full Integrity and Freshness for Outsourced Storage .....	362
<i>Hao Jin, Hong Jiang, Ke Zhou, Ronglei Wei, Dongliang Lei, and Ping Huang</i>	
General Functional Regenerating Codes with Uncoded Repair for Distributed Storage System .....	372
<i>Qing Liu, Dan Feng, Zhan Shi, and Min Fu</i>	

## Performance and Throughput Optimizations

Lark: Bringing Network Awareness to High Throughput Computing .....	382
<i>Zhe Zhang, Brian Bockelman, Dale W. Carder, and Todd Tannenbaum</i>	
Confuga: Scalable Data Intensive Computing for POSIX Workflows .....	392
<i>Patrick Donnelly, Nicholas Hazekamp, and Douglas Thain</i>	
Assessing Safe Task Parallelism in SPEC 2006 INT .....	402
<i>Tongxin Bai, Chen Ding, and Pengcheng Li</i>	
Optimal Footprint Symbiosis in Shared Cache .....	412
<i>Xiaolin Wang, Yechen Li, Yingwei Luo, Xiameng Hu, Jacob Brock, Chen Ding, and Zhenlin Wang</i>	

## Performance Evaluation

An Empirical Performance Evaluation of GPU-Enabled Graph-Processing Systems .....	423
<i>Yong Guo, Ana Lucia Varbanescu, Alexandru Iosup, and Dick Epema</i>	
Towards Latency-Optimal Distributed Relay Selection .....	433
<i>Yongquan Fu, Yijie Wang, and Xiaoqiang Pei</i>	
Evaluating the Effectiveness of Replication for Tail-Tolerance .....	443
<i>Zhan Qiu and Juan F. Pérez</i>	
ProRenaTa: Proactive and Reactive Tuning to Scale a Distributed Storage System .....	453
<i>Ying Liu, Navaneeth Rameshan, Enric Monte, Vladimir Vlassov, and Leandro Navarro</i>	

## Datacenters

Statistical Characterization of Business-Critical Workloads Hosted in Cloud Datacenters .....	465
<i>Siqi Shen, Vincent van Beek, and Alexandru Iosup</i>	
CloudSimSDN: Modeling and Simulation of Software-Defined Cloud Data Centers .....	475
<i>Jungmin Son, Amir Vahid Dastjerdi, Rodrigo N. Calheiros, Xiaohui Ji, Young Yoon, and Rajkumar Buyya</i>	
DualVisor: Redundant Hypervisor Execution for Achieving Hardware Error Resilience in Datacenters .....	485
<i>Xin Xu and H. Howie Huang</i>	
An Availability-on-Demand Mechanism for Datacenters .....	495
<i>Siqi Shen, Alexandru Iosup, Assaf Israel, Walfredo Cirne, Danny Raz, and Dick Epema</i>	

## Cloud Computing II

Performance-Based Service Differentiation in Clouds .....	505
<i>Ewnetu Bayuh Lakew, Cristian Klein, Francisco Hernandez-Rodriguez, and Erik Elmroth</i>	
Service Clustering for Autonomic Clouds Using Random Forest .....	515
<i>Rafael Brundo Uriarte, Sotirios Tsafaris, and Francesco Tiezzi</i>	
ProvErr: System Level Statistical Fault Diagnosis Using Dependency Model .....	525
<i>Peng Chen and Beth A. Plale</i>	
Quantitative Musings on the Feasibility of Smartphone Clouds .....	535
<i>Chen Chen, Moussa Ehsan, and Radu Sion</i>	

## Cloud and Cluster Applications

Risk-Driven Framework for Decision Support in Cloud Service Selection .....	545
<i>Smrati Gupta, Victor Muntès-Mulero, Peter Matthews, Jacek Dominiak, Aida Omerovic, Jordi Aranda, and Stepan Seycek</i>	
Architecture Aware Resource Allocation for Structured Grid Applications: Flood Modelling Case .....	555
<i>Vaibhav Saxena, Thomas George, Yogish Sabharwal, and Lucas Villa Real</i>	
Running MAP Inference on Million Node Graphical Models: A High Performance Computing Perspective .....	565
<i>Chen Jin, Qiang Fu, Huahua Wang, William Hendrix, Zhengzhang Chen, Ankit Agrawal, Arindam Banerjee, and Alok Choudhary</i>	
A Parallel Algorithm for Clipping Polygons with Improved Bounds and a Distributed Overlay Processing System Using MPI .....	576
<i>Satish Puri and Sushil K. Prasad</i>	

## Scheduling

Contiguity and Locality in Backfilling Scheduling .....	586
<i>Giorgio Lucarelli, Fernando Mendonca, Denis Trystram, and Frederic Wagner</i>	
A Priority-Based Scheduling Heuristic to Maximize Parallelism of Ready Tasks for DAG Applications .....	596
<i>Wei Zheng, Lu Tang, and Rizos Sakellariou</i>	
Scheduling Workloads of Workflows with Unknown Task Runtimes .....	606
<i>Alexey Ilyushkin, Bogdan Ghit, and Dick Epema</i>	
A Scheduler-Level Incentive Mechanism for Energy Efficiency in HPC .....	617
<i>Yiannis Georgiou, David Glesser, Krzysztof Rządca, and Denis Trystram</i>	

## MapReduce

GERBIL: MPI+YARN .....	627
<i>Luna Xu, Min Li, and Ali R. Butt</i>	
YARNsim: Simulating Hadoop YARN .....	637
<i>Ning Liu, Xi Yang, Xian-He Sun, Johnathan Jenkins, and Robert Ross</i>	
Towards Provenance-Based Anomaly Detection in MapReduce .....	647
<i>Cong Liao and Anna Squicciarini</i>	

## Distributed Resource Management

Joint Scheduling of Data and Computation in Geo-Distributed Cloud Systems .....	657
<i>Lingyan Yin, Jizhou Sun, Laiping Zhao, Chenzhou Cui, Jian Xiao, and Ce Yu</i>	

Platform and Co-Runner Affinities for Many-Task Applications in Distributed Computing Platforms .....	667
<i>Seontae Kim, Eunji Hwang, Tae-kyung Yoo, Jik-Soo Kim, Soonwook Hwang, and Young-ri Choi</i>	
Coordinated Resource Management for Large Scale Interactive Data Query Systems .....	677
<i>Wei Yan and Yuan Xue</i>	
A Multi-objective Biogeography-Based Optimization for Virtual Machine Placement .....	687
<i>Qinghua Zheng, Rui Li, Xiuqi Li, and Jie Wu</i>	

## **Doctoral Symposium Short Talks**

Techniques for Enabling Highly Efficient Message Passing on Many-Core Architectures .....	697
<i>Min Si, Pavan Balaji, and Yutaka Ishikawa</i>	
Runtime Support for Irregular Computation in MPI-Based Applications .....	701
<i>Xin Zhao, Pavan Balaji, and William Gropp</i>	
A Framework to Accelerate Protein Structure Comparison Tools .....	705
<i>Ahmad Salah, Kneli Li, and Tarek F. Gharib</i>	
Improving Application Performance by Efficiently Utilizing Heterogeneous Many-core Platforms .....	709
<i>Jie Shen, Ana Lucia Varbanescu, and Henk Sips</i>	
Modeling Gather and Scatter with Hardware Performance Counters for Xeon Phi .....	713
<i>James Lin, Akira Nukada, and Satoshi Matsuoka</i>	
MIC-Tandem: Parallel X!Tandem Using MIC on Tandem Mass Spectrometry Based Proteomics Data .....	717
<i>Pinjie He and Kenli Li</i>	
Optimize Parallel Data Access in Big Data Processing .....	721
<i>Jiangling Yin and Jun Wang</i>	
mD3DOCKxb: A Deep Parallel Optimized Software for Molecular Docking with Intel Xeon Phi Coprocessors .....	725
<i>Qian Cheng, Shaoliang Peng, Yutong Lu, Weiliang Zhu, Zhijian Xu, and XinBen Zhang</i>	
mAMBER: Accelerating Explicit Solvent Molecular Dynamic with Intel Xeon Phi Many-Integrated Core Coprocessors .....	729
<i>Xin Liu, Shaoliang Peng, Canqun Yang, Chengkun Wu, Haiqiang Wang, Qian Cheng, Weiliang Zhu, and Jinan Wang</i>	
Parallel Solving Method of SOR Based on the Numerical Marine Forecasting Model .....	733
<i>Renbo Pang, Jianliang Xu, and Yunquan Zhang</i>	
Towards Self Adaptable Security Monitoring in IaaS Clouds .....	737
<i>Anna Giannakou, Louis Rilling, Jean-Louis Pazat, Frederic Majorczyk, and Christine Morin</i>	



Understanding Unsuccessful Executions in Big-Data Systems .....	741
<i>Andrea Rosà, Lydia Y. Chen, and Walter Binder</i>	
Improving Energy Efficiency of Web Servers by Using a Load Distribution Algorithm and Shutting Down Idle Nodes .....	745
<i>Kai Chen, Jörg Lenhardt, and Wolfram Schiffmann</i>	
Towards Context-Aware Mobile Crowdsensing in Vehicular Social Networks .....	749
<i>Xiping Hu and Victor C.M. Leung</i>	
Towards a Realistic Scheduler for Mixed Workloads with Workflows .....	753
<i>Alexey Ilyushkin and Dick Epema</i>	
Predicting the Performance of Parallel Computing Models Using Queuing System .....	757
<i>Shen Chao, Tong Weiqin, and Samina Kausar</i>	
Cloud Service Recommendation: State of the Art and Research Challenges .....	761
<i>Lantian Guo, Xianrong Zheng, Chen Ding, Dejun Mu, and Zhe Li</i>	
Cross-Layer SLA Management for Cloud-hosted Big Data Analytics Applications .....	765
<i>Xuezhi Zeng, Rajiv Ranjan, Peter Strazdins, Saurabh Kumar Garg, and Lizhe Wang</i>	
SWAP-Assembler 2: Scalable Genome Assembler towards Millions of Cores—Practice and Experience .....	769
<i>Jintao Meng, Yanjie Wei, Sangmin Seo, and Pavan Balaji</i>	
FlexiMod: Flexible Coexistence Support for Programming Models .....	773
<i>Luna Xu, R. Ali, and Butt</i>	
A Reliable Distributed Convolutional Neural Network for Biology Image Segmentation .....	777
<i>Xiuxia Zhang, Guangming Tan, and Mingyu Chen</i>	
A Method to Accelerate GROMACS in Offload Mode on Tianhe-2 Supercomputer .....	781
<i>Haiqiang Wang, Shaoliang Peng, Xiaoqian Zhu, Chengkun Wu, Xin Liu, Qian Chen, Weiliang Zhu, Jinan Wang, and Huaiyu Yang</i>	
Majority Quorum Protocol Dedicated to General Threshold Schemes .....	785
<i>Theodore Jean Richard Relaza, Jacques Jorda, and Abdelaziz M'zoughi</i>	
HAGP: A Hub-Centric Asynchronous Graph Processing Framework for Scale-Free Graph .....	789
<i>Tao Gao, Yutong Lu, and Baida Zhang</i>	
Cost-Efficient High-Performance Internet-Scale Data Analytics over Multi-cloud Environments .....	793
<i>Shigeru Imai, Stacy Patterson, and Carlos A. Varela</i>	
Big Data Provenance Analysis and Visualization .....	797
<i>Peng Chen and Beth A. Plale</i>	
A Novel Approach To Classify Cloud Entities: Universal Cloud Classification (UCC) .....	801
<i>Sebastian Jeuk, Gonzalo Salgueiro, and Shi ZHou</i>	

## **SCALE Challenge**

Scalable In-Memory Computing .....	805
<i>Alexandru Uta, Andreea Sandu, Stefania Costache, and Thilo Kielmann</i>	
Scaling NWChem with Efficient and Portable Asynchronous Communication in MPI RMA .....	811
<i>Min Si, Antonio J. Peña, Jeff Hammond, Pavan Balaji, and Yutaka Ishikawa</i>	
Accurate Scoring of Drug Conformations at the Extreme Scale .....	817
<i>Boyu Zhang, Trilce Estrada, Pietro Cicotti, Pavan Balaji, and Michela Taufer</i>	
The Challenge of Scaling Genome Big Data Analysis Software on TH-2 Supercomputer .....	823
<i>Shaoliang Peng, Xiangke Liao, Canqun Yang, Yutong Lu, Jie Liu, Yingbo Cui, Heng Wang, Chengkun Wu, and Bingqiang Wang</i>	
Real-Time Analytics for Fast Evolving Social Graphs .....	829
<i>Charith Wickramaarachchi, Alok Kumbhare, Marc Frincu, Charalampos Chelmis, and Viktor K. Prasanna</i>	

## **Workshop on Clusters, Clouds and Grids for Life Sciences—CCGrid-Life**

Analysing Cancer Genomics in the Elastic Cloud .....	835
<i>Christopher Smowton, Crispin Miller, Wei Xing, Andoena Balla, Demetris Antoniadis, George Pallis, and Marios D. Dikaiakos</i>	
SparkSW: Scalable Distributed Computing System for Large-Scale Biological Sequence Alignment .....	845
<i>Guoguang Zhao, Cheng Ling, and Donghong Sun</i>	
A Comparative Analysis of Scheduling Mechanisms for Virtual Screening Workflow in a Shared Resource Environment .....	853
<i>Bui The Quang, Jik-Soo Kim, Seungwoo Rho, Seoyoung Kim, Sangwan Kim, Soonwook Hwang, Emmanuel Medernach, and Vincent Breton</i>	
Classifications of Computing Sites to Handle Numerical Variability .....	863
<i>Tristan Glatard and Alan C. Evans</i>	
Scaling Machine Learning for Target Prediction in Drug Discovery using Apache Spark .....	871
<i>Dries Harnie, Alexander E. Vapirev, Jörg Kurt Wegner, Andrey Gedich, Marvin Steijaert, Roel Wuyts, and Wolfgang De Meuter</i>	
Multicenter Data Sharing for Collaboration in Sleep Medicine .....	880
<i>Maximilian Beier, Christoph Jansen, Geert Mayer, Thomas Penzel, Andrea Rodenbeck, René Siewert, Jie Wu, and Dagmar Krefting</i>	

## **Fourth International Workshop on Data-intensive Process Management in Large-Scale Sensor Systems (DPMSS 2015)**

DBH-CLUS: A Hierarchical Clustering Method to Identify Pick-up/Drop-off Hotspots .....	890
<i>XueJin Wan, Jiong Wang, Yong Du, and Yuan Zhong</i>	
Experience Based Sink Placement in Mobile Wireless Sensor Network .....	898
<i>Subhra Banerjee, Suman Sankar Bhunia, and Nandini Mukherjee</i>	
Cloud-Based Machine Learning Tools for Enhanced Big Data Applications .....	908
<i>Alfredo Cuzzocrea, Enzo Mumolo, and Pietro Corona</i>	
Fast Replica Placement and Update Strategies in Tree Networks .....	915
<i>Xu Wang, Jigang Wu, Guiyuan Jiang, Stew Kei Lam, and Thambipillai Srikanthan</i>	
Cloud-Based OLAP over Big Data: Application Scenarios and Performance Analysis .....	921
<i>Alfredo Cuzzocrea, Rim Moussa, Guandong Xu, and Giorgio Mario Grasso</i>	
A Data Placement Strategy for Data-Intensive Scientific Workflows in Cloud .....	928
<i>Qing Zhao, Congcong Xiong, Xi Zhao, Ce Yu, and Jian Xiao</i>	

## **ExtremeGreen 2015: Extreme Green and Energy Efficiency in Large Scale Distributed Systems**

Computing Heaters—An Energy-Efficient Way to Provide Computing Services .....	935
<i>Jukka K. Nurminen, Johan Strandman, and Tapio Niemi</i>	
Quantifying the Energy Efficiency Challenges of Achieving Exascale Computing .....	943
<i>Jason Mair, Zhiyi Huang, David Eyers, and Yawen Chen</i>	

## **Cloud for Business, Industry and Enterprises 2015—C4BIE 2015**

Spy: A QoS-Aware Anonymous Multi-Cloud Storage System Supporting DSSE .....	951
<i>Pengyan Shen, Kai Guo, Mingzhong Xiao, and Quanqing Xu</i>	
Security-Oriented Cloud Platform for SOA-Based SCADA .....	961
<i>Thar Baker, Michael Mackay, Amjad Shaheed, and Bandar Aldawsari</i>	
An Auto-Scaling Framework for Controlling Enterprise Resources on Clouds .....	971
<i>Anshuman Biswas, Shikharesh Majumdar, Biswajit Nandy, and Ali El-Haraki</i>	

## **2015 Workshop on Data Vitalization and Universal Village—DV&UV**

High Dynamic Range Saturation Intelligence Avoidance for Three-Dimensional Shape Measurement .....	981
<i>Bin Zhang, Yuanxin Ouyang, and Shuo Zhang</i>	
Traffic Management and Forecasting System Based on 3D GIS .....	991
<i>Xiaoming Li, Zhihan Lv, Jinxing Hu, Baoyun Zhang, Ling Yin, Chen Zhong, Weixi Wang, and Shengzhong Feng</i>	

Sociality Analysis in Wireless Networks .....	999
<i>Wenhao Chen and Guangtao Xue</i>	
PCAH: A PCA-Based Hierarchical Clustering Method for Visual Words Construction .....	1009
<i>Ying He, Jian Wang, Xue-xia Zhong, Lin Mei, and Zhi-zong Wu</i>	
Visualizing City Events on Search Engine: Tword the Search Infrustration for Smart City .....	1019
<i>Wenbo Li, Peixia Wang, and Kaifei Yang</i>	
A Structured Light 3D Measurement System Based on Heterogeneous Parallel Computation Model .....	1027
<i>Xiaoyu Liu, Hao Sheng, Yang Zhang, and Zhang Xiong</i>	
<b>SCRAMBL 2015</b>	
File Multicast Transport Protocol (FMTP) .....	1037
<i>Jie Li, Malathi Veeraraghavan, Steve Emmerson, and Robert. D. Russell</i>	
Towards Elasticity in Distributed File Systems .....	1047
<i>Cyril Séguin, Gaël Le Mahec, and Benjamin Depardon</i>	
HPC-ABDS High Performance Computing Enhanced Apache Big Data Stack .....	1057
<i>Geoffrey C. Fox, Judy Qiu, Supun Kamburugamuve, Shantenu Jha, and Andre Luckow</i>	
<b>Parallel Programming Model for the Masses (PPMM2015)</b>	
Programming Heterogeneous Systems with Array Types .....	1067
<i>Xiang Cui, Xiaowen Li, and Yifeng Chen</i>	
Characterizing MPI and Hybrid MPI+Threads Applications at Scale: Case Study with BFS .....	1075
<i>Abdelhalim Amer, Huiwei Lu, Pavan Balaji, and Satoshi Matsuoka</i>	
Implementation and Evaluation of MPI Nonblocking Collective I/O .....	1084
<i>Sangmin Seo, Robert Latham, Junchao Zhang, and Pavan Balaji</i>	
An Evaluation of Unified Memory Technology on NVIDIA GPUs .....	1092
<i>Wenqiang Li, Guanghao Jin, Xuwen Cui, and Simon See</i>	
Analyzing MPI-3.0 Process-Level Shared Memory: A Case Study with Stencil Computations .....	1099
<i>Xiaomin Zhu, Junchao Zhang, Kazutomo Yoshii, Shigang Li, Yunquan Zhang, and Pavan Balaji</i>	

## Poster Short Talks

A Task-Type-Based Algorithm for the Energy-Aware Profit Maximizing Scheduling Problem in Heterogeneous Computing Systems .....	1107
<i>Weidong Li, Xi Liu, Xuejie Zhang, and Xiaobo Cai</i>	
On Energy- and Cooling-Aware Data Centre Workload Management .....	1111
<i>Danuta Sorina Chisca, Ignacio Castineiras, Deepak Mehta, and Barry O'Sullivan</i>	
Energy Profiling Using IgProf .....	1115
<i>Kashif Nizam Khan, Filip Nybäck, Zhonghong Ou, Jukka K. Nurminen, Tapio Niemi, Giulio Eulisse, Peter Elmer, and David Abdurachmanov</i>	
Boosting GPU Performance by Profiling-Based L1 Data Cache Bypassing .....	1119
<i>Yijie Huangfu and Wei Zhang</i>	
Lessons Learned Implementing User-Level Failure Mitigation in MPICH .....	1123
<i>Wesley Bland, Huiwei Lu, Sangmin Seo, and Pavan Balaji</i>	
A Multilevel Fault-Tolerance Technique for the DAG Data Driven Model .....	1127
<i>Hao Fu, Ce Yu, Jizhou Sun, Jun Du, and Mengmeng Wang</i>	
Astrophysics Simulation on RSC Massively Parallel Architecture .....	1131
<i>Igor Kulikov, Igor Chernykh, Boris Glinsky, Dmitry Weins, and Alexey Shmelev</i>	
Adaptive Load-Balancing for Consistent Hashing in Heterogeneous Clusters .....	1135
<i>Lakshminarayanan Srinivasan and Vasudeva Varma</i>	
Improving TLB Performance by Increasing Hugepage Ratio .....	1139
<i>Taowei Luo, Xiaolin Wang, Jingyuan Hu, Yingwei Luo, and Zhenlin Wang</i>	
Understanding Data Access Patterns Using Object-Differentiated Memory Profiling .....	1143
<i>Antonio J. Pena and Pavan Balaji</i>	
Assessing Memory Access Performance of Chapel through Synthetic Benchmarks .....	1147
<i>Engin Kayraklioglu and Tarek El-Ghazawi</i>	
Accelerating Machine Learning Kernel in Hadoop Using FPGAs .....	1151
<i>Katayoun Neshatpour, Maria Malik, and Houman Homayoun</i>	
Parallel DC3 Algorithm for Suffix Array Construction on Many-Core Accelerators .....	1155
<i>Gang Liao, Longfei Ma, Guangming Zang, and Lin Tang</i>	
A Deep Learning Prediction Process Accelerator Based FPGA .....	1159
<i>Qi Yu, Chao Wang, Xiang Ma, Xi Li, and Xuehai Zhou</i>	
Locality-Aware Stencil Computations Using Flash SSDs as Main Memory Extension .....	1163
<i>Hiroko Midorikawa and Hideyuki Tan</i>	
On the Design of a Demo for Exhibiting rCUDA .....	1169
<i>Carlos Reaño, Ferrán Pérez, and Federico Silla</i>	
Toward Implementing Robust Support for Portals 4 Networks in MPICH .....	1173
<i>Ken Raffanetti, Antonio J. Pena, and Pavan Balaji</i>	

BigDataDIRAC: Deploying Distributed Big Data Applications .....	1177
<i>Victor Fernández, Víctor Méndez, and Tomás F. Pena</i>	
A Software Workbench for Interactive, Time Critical and Highly Self-Adaptive Cloud Applications (SWITCH) .....	1181
<i>Zhiming Zhao, Arie Taal, Andrew Jones, Ian Taylor, Vlado Stankovski, Ignacio Garcia Vega, Francisco Jesus Hidalgo, George Suciú, Alexandre Ulisses, Pedro Ferreira, and Cees de Laat</i>	
Enumeration System on HBase for Low-Latency .....	1185
<i>Hang Chen, Kun Lu, MingMing Sun, ChangLong Li, Hang Zhuang, and XueHai Zhou</i>	
Augmenting Performance For Distributed Cloud Storage .....	1189
<i>Matthew B. Hancock and Carlos A. Varela</i>	
Discriminative Model for Google Host Load Prediction with Rich Feature Set .....	1193
<i>Peijie Huang, Dashu Ye, Ziwei Fan, Peisen Huang, and Xuezhen Li</i>	
Self Protecting Data Sharing Using Generic Policies .....	1197
<i>Shiping Chen, Danan Thilakanathan, Donna Xu, Surya Nepal, and Rafael Calvo</i>	
Highly Available Cloud-Based Cluster Management .....	1201
<i>Dmitry Duplyakin, Matthew Haney, and Henry Tufo</i>	
Publish/Subscribe Middleware for Resource Discovery in MANET .....	1205
<i>Malihe Saghian and Reza Ravanmehr</i>	
Hamiltonian Path Strategy for Deadlock-Free and Adaptive Routing in Diametrical 2D Mesh NoCs .....	1209
<i>Poona Bahrebar and Dirk Stroobandt</i>	
Design of a More Scalable Database System .....	1213
<i>Hang Zhuang, Kun Lu, Changlong Li, Mingming Sun, Hang Chen, and Xuehai Zhou</i>	
A Novel Efficient Approach for Screen Image Classification in Remote Display Protocol .....	1217
<i>Pham Xuan Qui, Nguyen Tien Dung, Huynh Cong Thinh, Pham Phuoc Hung, Nguyen Huu Quoc, and Eui-Nam Huh</i>	
A QoS Assured Network Service Chaining Algorithm in Network Function Virtualization Architecture .....	1221
<i>Taekhee Kim, Siri Kim, Kwonyong Lee, and Sungyong Park</i>	
Study of the KVM CPU Performance of Open-Source Cloud Management Platforms .....	1225
<i>F. Gomez-Folgar, A.J. Garcia-Loureiro, T.F. Pena, J.I. Zablah, and N. Seoane</i>	
Modeling the Task of Google MapReduce Workload .....	1229
<i>Xiaoyang Lin, Piyuan Lin, Peijie Huang, Linxiao Chen, Ziwei Fan, and Peisen Huang</i>	
Eliminating the Redundancy in MapReduce-Based Entity Resolution .....	1233
<i>Cairong Yan, Yalong Song, Jian Wang, and Wenjing Guo</i>	
On-Demand Self-Adaptivity of Service Availability for Cloud Multi-tier Applications .....	1237
<i>Jin Yang, Jianmin Pang, Ning Qi, and Tao Qi</i>	

## Author Index