

**2013 13th IEEE/ACM
International Symposium on
Cluster, Cloud, and Grid
Computing**

(CCGrid 2013)

**Delft, Netherlands
13 – 16 May 2013**



**IEEE Catalog Number: CFP13276-PRT
ISBN: 978-1-4673-6465-2**

2013 13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing

CCGrid 2013

Table of Contents

Message from the CCGrid 2013 General Chair.....	xiv
Message from the CCGrid 2013 Program Chair.....	xvi
CCGrid 2013 Organization.....	xviii
Steering Committee.....	xx
Committees.....	xxi
External Reviewers.....	xxvi

Session 1: Keynote 1

Programming Models for High-Performance Computing	1
<i>Marc Snir</i>	

Session 2A: Performance

Exploiting Per User Information for Supercomputing Workload Prediction Requires Care	2
<i>Tuan V. Dinh, Lachlan L.H. Andrew, and Philip Branch</i>	
Building an Elastic Cloud out of Small Datacenters	10
<i>Indra Widjaja, Sem Borst, and Iraj Saniee</i>	
Simulating Application Workflows and Services Deployed on the European Grid Infrastructure	18
<i>Sorina Camarasu-Pop, Tristan Glatard, and Hugues Benoit-Cattin</i>	
Optimizing Large Data Transfers over 100Gbps Wide Area Networks	26
<i>Anupam Rajendran, Parag Mhashilkar, Hyunwoo Kim, Dave Dykstra, Gabriele Garzoglio, and Ioan Raicu</i>	

Session 2B: Clouds I

Automated, Elastic Resource Provisioning for NoSQL Clusters Using TIRAMOLA	34
<i>Dimitrios Tsoumakos, Ioannis Konstantinou, Christina Boumpouka, Spyros Sioutas, and Nectarios Koziris</i>	
Autoflex: Service Agnostic Auto-scaling Framework for IaaS Deployment Models	42
<i>Fábio Jorge Almeida Morais, Francisco Vilar Brasileiro, Raquel Vigolvino Lopes, Ricardo Araújo Santos, Wade Satterfield, and Leandro Rosa</i>	
Towards QoS-Oriented SLA Guarantees for Online Cloud Services	50
<i>Damián Serrano, Sara Bouchenak, Yousri Kouki, Thomas Ledoux, Jonathan Lejeune, Julien Sopena, Luciana Arantes, and Pierre Sens</i>	
Optimal Cloud Resource Auto-Scaling for Web Applications	58
<i>Jing Jiang, Jie Lu, Guangquan Zhang, and Guodong Long</i>	

Session 2C: Doctoral Symposium: Parallel and Distributed Systems

Cost-Efficient Virtual Machine Provisioning for Multi-tier Web Applications and Video Transcoding	66
<i>Adnan Ashraf</i>	
Sesame: A User-Transparent Optimizing Framework for Many-Core Processors	70
<i>Jianbin Fang, Ana Lucia Varbanescu, and Henk Sips</i>	
Reproducing Network Conditions for Tests of Large-Scale Distributed Systems	74
<i>Robert Lübke, Daniel Schuster, and Alexander Schill</i>	
Fast Wide Area Live Migration with a Low Overhead through Page Cache Teleportation	78
<i>Soramichi Akiyama, Takahiro Hirofuchi, Ryousei Takano, and Shinichi Honiden</i>	
Towards an Optimized Big Data Processing System	83
<i>Bogdan Ghiț, Alexandru Iosup, and Dick Epema</i>	

Session 3A: Programming Models

Toward Asynchronous and MPI-Interoperable Active Messages	87
<i>Xin Zhao, Darius Buntinas, Judicael Zounmevo, James Dinan, David Goodell, Pavan Balaji, Rajeev Thakur, Ahmad Afsahi, and William Gropp</i>	
Swift/T: Large-Scale Application Composition via Distributed-Memory Dataflow Processing	95
<i>Justin M. Wozniak, Timothy G. Armstrong, Michael Wilde, Daniel S. Katz, Ewing Lusk, and Ian T. Foster</i>	
Scalable PGAS Metadata Management on Extreme Scale Systems	103
<i>Daniel Chavarría-Miranda, Khushbu Agarwal, and T.P. Straatsma</i>	

SYBL: An Extensible Language for Controlling Elasticity in Cloud Applications	112
<i>Georgiana Copil, Daniel Moldovan, Hong-Linh Truong, and Schahram Dustdar</i>	

Session 3B: Accelerators

GPU-TLS: An Efficient Runtime for Speculative Loop Parallelization on GPUs	120
<i>Chenggang Zhang, Guodong Han, and Cho-Li Wang</i>	
Efficient Intra-node Communication on Intel-MIC Clusters	128
<i>Sreeram Potluri, Akshay Venkatesh, Devendar Bureddy, Krishna Kandalla, and Dhabaleswar K. Panda</i>	
CUDA vs OpenACC: Performance Case Studies with Kernel Benchmarks and a Memory-Bound CFD Application	136
<i>Tetsuya Hoshino, Naoya Maruyama, Satoshi Matsuoka, and Ryoji Takaki</i>	
Evaluation of Inter- and Intra-node Data Transfer Efficiencies between GPU Devices and their Impact on Scalable Applications	144
<i>Antonio J. Pena and Sadaf R. Alam</i>	

Session 3C: Doctoral Symposium: Cloud Computing

Massivizing Multi-player Online Games on Clouds	152
<i>Siqi Shen, Alexandru Iosup, and Dick Epema</i>	
Data Usage Control for the Cloud	156
<i>Florian Kelbert</i>	
Extending the Capabilities of Mobile Devices for Online Social Applications through Cloud Offloading	160
<i>Alexandru-Corneliu Olteanu, Nicolae Tapus, and Alexandru Iosup</i>	
Adapting Scientific Applications to Cloud by Using Distributed Computing Frameworks	164
<i>Pelle Jakovits and Satish Narayana Srirama</i>	
Negotiation-Based Flexible SLA Establishment with SLA-driven Resource Allocation in Cloud Computing	168
<i>Seokho Son and Sung Chan Jun</i>	

Session 4: Poster Presentations

Cloud Computing for High Performance Image Analysis on a National Infrastructure	172
<i>D. Wang, T. Bednarz, Y. Arzhaeva, J. Taylor, P. Szul, S. Chen, N. Burdett, A. Khassapov, and T. Gureyev</i>	
An MPI-IO Compliant Java Based Parallel I/O Library	174
<i>Ammar Ahmad Awan, Muhammad Bilal Amin, Shujaat Hussain, Aamir Shafi, and Sungyoung Lee</i>	

Specialized File Transfer Service for Large Oil and Gas Datasets	176
<i>Fabio de Souza, Marcus Salles, Fabio Campos, Silvio Costa, Myrian Costa, and Nelson Ebecken</i>	
TlaaS: Secure Cloud-assisted Traffic Information Dissemination in Vehicular Ad Hoc Networks	178
<i>Rasheed Hussain, Fizza Abbas, Junggab Son, and Heekuck Oh</i>	
A Novel Checkpointing Scheme for Amazon EC2 Spot Instances	180
<i>Sunirmal Khatua and Nandini Mukherjee</i>	
HDFS+: Concurrent Writes Improvements for HDFS	182
<i>Kun Lu, Dong Dai, and Mingming Sun</i>	
Interference-aware Incoming Message Detection for MPI Threaded Progression	184
<i>Masahiro Miwa, Kohta Nakashima, and Akira Naruse</i>	
Accounting Federated Clouds Based on the JiTCloud Platform	186
<i>Francisco Airtton Silva, Paulo Neto, Vinicius Garcia, Fernando Trinta, and Rodrigo Assad</i>	
Predictive Caching in Computer Grids	188
<i>Efstratios Rappos and Stephan Robert</i>	
Secure Storage Service for IaaS Cloud Users	190
<i>Jinho Seol, Seongwook Jin, and Seungryoul Maeng</i>	
Towards Assurance of Availability in Virtualized Cloud System	192
<i>Seongwook Jin, Jinho Seol, and Seungryoul Maeng</i>	
Isolated Mini-domain for Trusted Cloud Computing	194
<i>Jaewon Choi, Jongse Park, Jinho Seol, and Seungryoul Maeng</i>	
User-Level Remote Memory Paging for Multithreaded Applications	196
<i>Hiroko Midorikawa, Yuichiro Suzuki, and Masatoshi Iwaida</i>	
An Autonomic and Scalable Management System for Private Clouds	198
<i>Matthieu Simonin, Eugen Feller, Anne-Cécile Orgerie, Yvon Jégou, and Christine Morin</i>	
Evaluation of Cloud Providers for VPH Applications	200
<i>Marian Bubak, Marek Kasztelnik, Maciej Malawski, Jan Meizner, Piotr Nowakowski, and Susheel Varma</i>	

Session 6: Keynote 2

The Astronomical Multipurpose Software Environment and the Ecology of Star Clusters	202
<i>Simon Portegies Zwart</i>	

Session 7A: Scheduling

Multi-objective Workflow Scheduling: An Analysis of the Energy Efficiency and Makespan Tradeoff	203
<i>Juan J. Durillo, Vlad Nae, and Radu Prodan</i>	
Hierarchical I/O Scheduling for Collective I/O	211
<i>Jialin Liu, Yong Chen, and Yi Zhuang</i>	
Stretch Out and Compact: Workflow Scheduling with Resource Abundance	219
<i>Young Choon Lee and Albert Y. Zomaya</i>	
Scheduling Transactions in Replicated Distributed Software Transactional Memory	227
<i>Junwhan Kim and Binoy Ravindran</i>	
Automated SLA Negotiation Framework for Cloud Computing	235
<i>Linlin Wu, Saurabh Kumar Garg, Rajkumar Buyya, Chao Chen, and Steve Versteeg</i>	

Session 7B: MapReduce

Bi-Hadoop: Extending Hadoop to Improve Support for Binary-Input Applications	245
<i>Xiao Yu and Bo Hong</i>	
Non-intrusive Slot Layering in Hadoop	253
<i>Peng Lu, Young Choon Lee, and Albert Y. Zomaya</i>	
Resilin: Elastic MapReduce over Multiple Clouds	261
<i>Anca Iordache, Christine Morin, Nikos Parlavantzas, Eugen Feller, and Pierre Riteau</i>	
A Lightweight Continuous Jobs Mechanism for MapReduce Frameworks	269
<i>Trong-Tuan Vu and Fabrice Huet</i>	
A Scalable Implementation of a MapReduce-based Graph Processing Algorithm for Large-Scale Heterogeneous Supercomputers	277
<i>Koichi Shirahata, Hitoshi Sato, Toyotaro Suzumura, and Satoshi Matsuoka</i>	

Session 7C: Applications I

V-BOINC: The Virtualization of BOINC	285
<i>Gary A. McGilvary, Adam Barker, Ashley Lloyd, and Malcolm Atkinson</i>	
Experiences Applying Data Staging Technology in Unconventional Ways	294
<i>Jay Lofstead, Ron Oldfield, and Todd Kordenbrock</i>	
Versioned File Backup and Synchronization for Storage Clouds	302
<i>Shuang Qiu, Jingyu Zhou, and Tao Yang</i>	
Exploring Dynamic Enactment of Scientific Workflows Using Pilot-Abstractions	311
<i>Mark Santcroos, Barbera DC van Schaik, Shayan Shahand, Silvia Delgado Olabariaga, Andre Luckow, and Shantenu Jha</i>	

Evaluating Cloud Computing Techniques for Smart Power Grid Design Using Parallel Scripting	319
<i>Ketan Maheshwari, Ken Birman, Justin Wozniak, and Devin Van Zandt</i>	

Session 9A: Data Management

Understanding Data Characteristics and Access Patterns in a Cloud Storage System	327
<i>Songbin Liu, Xiaomeng Huang, Haohuan Fu, and Guangwen Yang</i>	
Supporting a Light-Weight Data Management Layer over HDF5	335
<i>Yi Wang, Yu Su, and Gagan Agrawal</i>	
PARLO: PARallel Run-Time Layout Optimization for Scientific Data Explorations with Heterogeneous Access Patterns	343
<i>Zhenhuan Gong, David A. Boyuka II, Xiaocheng Zou, Qing Liu, Norbert Podhorszki, Scott Klasky, Xiaosong Ma, and Nagiza F. Samatova</i>	
Consistency in the Cloud: When Money Does Matter!	352
<i>Houssem-Eddine Chihoub, Shadi Ibrahim, Gabriel Antoniu, and María S. Pérez</i>	

Session 9B: Multicore

Partially Separated Page Tables for Efficient Operating System Assisted Hierarchical Memory Management on Heterogeneous Architectures	360
<i>Balazs Gerofi, Akio Shimada, Atsushi Hori, and Yutaka Ishikawa</i>	
SLOAVx: Scalable LOGarithmic AlltoallV Algorithm for Hierarchical Multicore Systems	369
<i>Cong Xu, Manjunath Gorentla Venkata, Richard L. Graham, Yandong Wang, Zhuo Liu, and Weikuan Yu</i>	
Optimizing Burrows-Wheeler Transform-Based Sequence Alignment on Multicore Architectures	377
<i>Jing Zhang, Heshan Lin, Pavan Balaji, and Wu-Chun Feng</i>	
SR-IOV Support for Virtualization on InfiniBand Clusters: Early Experience	385
<i>Jithin Jose, Mingzhe Li, Xiaoyi Lu, Krishna Chaitanya Kandalla, Mark Daniel Arnold, and Dhabaleswar K. Panda</i>	

Session 10: Keynote 3

Clusters, Grids and Clouds: A Look from Both Sides	393
<i>Daniel A. Reed</i>	

Session 11A: Load Balancing

Gang Migration of Virtual Machines Using Cluster-wide Deduplication	394
<i>Umesh Deshpande, Brandon Schlinker, Eitan Adler, and Kartik Gopalan</i>	
Improving HPC Application Performance in Cloud through Dynamic Load Balancing	402
<i>Abhishek Gupta, Osman Sarood, Laxmikant V. Kale, and Dejan Milojicic</i>	

Automatic Performance Prediction for Load-Balancing Coupled Models	410
<i>Daihee Kim, J. Walter Larson, and Kenneth Chiu</i>	
Collocating CPU-only Jobs with GPU-assisted Jobs on GPU-assisted HPC	418
<i>Jiadong Wu and Bo Hong</i>	
Session 11B: Clouds II	
A Chemistry-Inspired Middleware for Self-Adaptive Service Orchestration and Choreography	426
<i>Chen Wang and Jean-Louis Pazat</i>	
Computation Certification as a Service in the Cloud	434
<i>Safwan Mahmud Khan and Kevin W. Hamlen</i>	
Security Risk Assessment of Cloud Carrier	442
<i>Swetha Reddy Lenkala, Sachin Shetty, and Kaiqi Xiong</i>	
Efficient Use of Geographically Spread Cloud Resources	450
<i>Yossi Kanizo, Danny Raz, and Alexander Zlotnik</i>	
Session 11C: Applications II	
Large Scale Parallel Solution of Incompressible Flow Problems Using Uintah and Hypr	458
<i>John Schmidt, Martin Berzins, Jeremy Thornock, Tony Saad, and James Sutherland</i>	
Case Studies in Designing Elastic Applications	466
<i>Dinesh Rajan, Andrew Thrasher, Badi' Abdul-Wahid, Jesus A. Izaguirre, Scott Emrich, and Douglas Thain</i>	
High-Performance Quantum Computing Simulation for the Quantum Geometric Machine Model	474
<i>Adriano Maron, Renata Reiser, and Maurício Pilla</i>	
Stream-Based Admission Control and Scheduling for Video Transcoding in Cloud Computing	482
<i>Adnan Ashraf, Fareed Jokhio, Tewodros Deneke, Sebastien Lafond, Ivan Porres, and Johan Lilius</i>	
Session 12A: Architectures	
Flexible Capacity Partitioning in Many-Core Tiled CMPs	490
<i>Ahmad Samih, Xiaowei Jiang, Liang Han, and Yan Solihin</i>	
On Achieving High Message Rates	498
<i>Holger Fröning, Mondrian Nüssle, Heiner Litz, Christian Leber, and Ulrich Brüning</i>	

Simfrastructure: A Flexible and Adaptable Middleware Platform for Modeling and Analysis of Socially Coupled Systems	506
<i>Keith R. Bisset, Suruchi Deodhar, Hemanth Makkapati, Madhav V. Marathe, Paula Stretz, and Christopher L. Barrett</i>	

Session 12B: Energy and Fault Tolerance

SHStream: Self-Healing Framework for HTTP Video-Streaming	514
<i>Carlos Augusto Cunha and Luis Moura e Silva</i>	
ECOFIT: A Framework to Estimate Energy Consumption of Fault Tolerance Protocols for HPC Applications	522
<i>Mohammed el Mehdi Diouri, Olivier Glück, Laurent Lefèvre, and Franck Cappello</i>	
DUAL: Reliability-Aware Power Management in Data Centers	530
<i>Xin Xu, Kayo Teramoto, Allan Morales, and H. Howie Huang</i>	

Workshop Papers

DPMSS: Second International Workshop on Data-intensive Process Management in Large-Scale Sensor Systems (DPMSS): From Sensor Networks to Sensor Clouds

A Distributed In-Transit Processing Infrastructure for Forecasting Electric Vehicle Charging Demand	538
<i>Rafael Tolosana-Calasanz, José Ángel Banares, Liana Cipcigan, Omer Rana, Panagiotis Papadopoulos, and Congduc Pham</i>	
A Holistic Architecture for the Internet of Things, Sensing Services and Big Data	546
<i>David Tracey and Cormac Sreenan</i>	
A Multi-platform Sensor Coordinated Earth Observing Missions Scheduling Method for Hazard Monitoring	554
<i>Li Jun, Jing Ning, Hu Weidong, and Chen Hao</i>	
Empowering the Invulnerability of Wireless Sensor Networks through Super Wires and Super Nodes	561
<i>Xiuwen Fu, Wenfeng Li, and Giancarlo Fortino</i>	
Implementing BFS-based Traversals of RDF Graphs over MapReduce Efficiently	569
<i>Mirel Cosulschi, Alfredo Cuzzocrea, and Roberto De Virgilio</i>	
A Development and Execution Environment for Early Warning Systems for Natural Disasters	575
<i>Bartosz Balis, Tomasz Bartynski, Marian Bubak, Grzegorz Dyk, Tomasz Gubala, and Marek Kasztelnik</i>	
Managing Data and Processes in Cloud-Enabled Large-Scale Sensor Networks: State-of-the-Art and Future Research Directions	583
<i>Alfredo Cuzzocrea, Giancarlo Fortino, and Omer Rana</i>	

C4BIE: Cloud for Business, Industry and Enterprises

Operating Cost Reduction for Distributed Internet Data Centers	589
<i>Yangyang Li, Hongbo Wang, Jiankang Dong, Junbo Li, and Shiduan Cheng</i>	
CIVSched: Communication-aware Inter-VM Scheduling in Virtual Machine Monitor Based on the Process	597
<i>Bei Guan, Yanjun Wu, Liping Ding, and Yongji Wang</i>	
HPISecure: Towards Data Confidentiality in Cloud Applications	605
<i>Eyad Saleh and Christoph Meinel</i>	
A Framework for Automatic Resource Provisioning for Private Clouds	610
<i>Jose Orlando Melendez, Anshuman Biswas, Shikharesh Majumdar, Biswajit Nandy, Marzia Zaman, Pradeep Srivastava, and Nishith Goel</i>	
Energy-Saving Virtual Machine Placement in Cloud Data Centers	618
<i>Jiankang Dong, Xing Jin, Hongbo Wang, Yangyang Li, Peng Zhang, and Shiduan Cheng</i>	
An Adaptive Implementation Case Study of Apriori Algorithm for a Retail Scenario in a Cloud Environment	625
<i>Mahesh Balaji and G. Subrahmanya VRK Rao</i>	

WACC: First International Workshop on Assured Cloud Computing

BitDeposit: Deterring Attacks and Abuses of Cloud Computing Services through Economic Measures	630
<i>Jakub Szefer and Ruby B. Lee</i>	
Formal Analysis of Fault-tolerant Group Key Management Using ZooKeeper	636
<i>Stephen Skeirik, Rakesh B. Bobba, and Jose Meseguer</i>	
Tenant-ID: Tagging Tenant Assets in Cloud Environments	642
<i>Sebastian Jeuk, Shi Zhou, and Miguel Rio</i>	

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

Energy Efficient VM Placement Supported by Data Analytic Service	648
<i>Dapeng Dong and John Herbert</i>	
Heterogeneity: The Key to Achieve Power-Proportional Computing	656
<i>Georges Da Costa</i>	
Energy-aware VM Allocation on an Opportunistic Cloud Infrastructure	663
<i>Cesar O. Diaz, Harold Castro, Mario Villamizar, Johnatan E. Pecero, and Pascal Bouvry</i>	
Energy Efficient VM Scheduling for Cloud Data Centers: Exact Allocation and Migration Algorithms	671
<i>Chaima Ghribi, Makhlof Hadji, and Djamal Zeglache</i>	

Author Index	679
---------------------------	-----