

**2011 11th IEEE/ACM
International Symposium
on Cluster, Cloud and
Grid Computing**

(CCGrid 2011)

**Newport Beach, California, USA
23 – 26 May 2011**



**IEEE Catalog Number: CFP11276-PRT
ISBN: 978-1-4577-0129-0**

2011 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing

CCGrid 2011

Table of Contents

Message from the General Cochairs.....	xii
Message from the Program Committee Chair.....	xiv
Conference Chairs.....	xvi
Steering Committee.....	xviii
Program Committee Members.....	xix
Reviewers.....	xxii
Keynotes.....	xxiv

Virtual Machines

Characterizing the Performance of Parallel Applications on Multi-socket Virtual Machines	1
<i>Khaled Z. Ibrahim, Steven Hofmeyr, and Costin Iancu</i>	
CloudSpider: Combining Replication with Scheduling for Optimizing Live Migration of Virtual Machines across Wide Area Networks	13
<i>Sumit Kumar Bose, Scott Brock, Ronald Skeoch, and Shrisha Rao</i>	
Optimized Management of Power and Performance for Virtualized Heterogeneous Server Clusters	23
<i>Vinicius Petrucci, Enrique V. Carrera, Orlando Loques, Julius C.B. Leite, and Daniel Mossé</i>	

GPU-Based Computing

Small Discrete Fourier Transforms on GPUs	33
<i>S. Mitra and A. Srinivasan</i>	
A Parallel Rectangle Intersection Algorithm on GPU+CPU	43
<i>Shih-Hsiang Lo, Che-Rung Lee, Yeh-Ching Chung, and I-Hsin Chung</i>	

GPGPU-Accelerated Parallel and Fast Simulation of Thousand-Core Platforms	53
<i>Christian Pinto, Shivani Raghav, Andrea Marongiu, Martino Ruggiero, David Atienza, and Luca Benini</i>	

Programming Models and Runtime Systems

Assertion Based Parallel Debugging	63
<i>Minh Ngoc Dinh, David Abramson, Donny Kurniawan, Chao Jin, Bob Moench, and Luiz DeRose</i>	
Cheetah: A Framework for Scalable Hierarchical Collective Operations	73
<i>Richard Graham, Manjunath Gorentla Venkata, Joshua Ladd, Pavel Shamis, Ishai Rabinovitz, Vasily Filipov, and Gilad Shainer</i>	
Enabling Multi-physics Coupled Simulations within the PGAS Programming Framework	84
<i>Fan Zhang, Ciprian Docan, Manish Parashar, and Scott Klasky</i>	
Multiple Services Throughput Optimization in a Hierarchical Middleware	94
<i>Eddy Caron, Benjamin Depardon, and Frédéric Desprez</i>	

Grid and Cloud Computing Performance

On the Performance Variability of Production Cloud Services	104
<i>Alexandru Iosup, Nezhir Yigitbasi, and Dick Epema</i>	
The Grid Observatory	114
<i>Cécile Germain-Renaud, Alain Cady, Philippe Gauron, Michel Jouvin, Charles Loomis, Janusz Martyniak, Julien Nauroy, Guillaume Philippon, and Michèle Sebag</i>	
Grid Global Behavior Prediction	124
<i>Jesús Montes, Alberto Sánchez, and María S. Pérez</i>	

Volunteer Computing

A Robust Communication Framework for Parallel Execution on Volunteer PC Grids	134
<i>Eshwar Rohit, Hien Nguyen, Nagarajan Kanna, Jaspal Subhlok, Edgar Gabriel, Qian Wang, Margaret S. Cheung, and David Anderson</i>	
Non-cooperative Scheduling Considered Harmful in Collaborative Volunteer Computing Environments	144
<i>Bruno Donassolo, Arnaud Legrand, and Cláudio Geyer</i>	
Towards Real-Time, Volunteer Distributed Computing	154
<i>Sangho Yi, Emmanuel Jeannot, Derrick Kondo, and David P. Anderson</i>	

Distributed Systems and Applications

GeoServ: A Distributed Urban Sensing Platform	164
<i>Jong Hoon Ahnn, Uichin Lee, and Hyun Jin Moon</i>	
Building an Online Domain-Specific Computing Service over Non-dedicated Grid and Cloud Resources: The Superlink-Online Experience	174
<i>Mark Silberstein</i>	
Techniques for Fine-Grained, Multi-site Computation Offloading	184
<i>Kanad Sinha and Milind Kulkarni</i>	

Resource Scheduling on the Cloud

SLA-Based Resource Allocation for Software as a Service Provider (SaaS) in Cloud Computing Environments	195
<i>Linlin Wu, Saurabh Kumar Garg, and Rajkumar Buyya</i>	
Improving Utilization of Infrastructure Clouds	205
<i>Paul Marshall, Kate Keahey, and Tim Freeman</i>	
Resource and Revenue Sharing with Coalition Formation of Cloud Providers: Game Theoretic Approach	215
<i>Dusit Niyato, Athanasios V. Vasilakos, and Zhu Kun</i>	
Self-Healing Distributed Scheduling Platform	225
<i>Marc E. Frincu, Norha M. Villegas, Dana Petcu, Hausi A. Müller, and Romain Rouvoy</i>	

Data Streaming

Towards Reliable, Performant Workflows for Streaming-Applications on Cloud Platforms	235
<i>Daniel Zinn, Quinn Hart, Timothy McPhillips, Bertram Ludäscher, Yogesh Simmhan, Michail Giakkoupis, and Viktor K. Prasanna</i>	
A Sketch-Based Architecture for Mining Frequent Items and Itemsets from Distributed Data Streams	245
<i>Eugenio Cesario, Antonio Grillo, Carlo Mastroianni, and Domenico Talia</i>	

Caching and Shared Memory

APP: Minimizing Interference Using Aggressive Pipelined Prefetching in Multi-level Buffer Caches	254
<i>Christina M. Patrick, Nicholas Voshell, and Mahmut Kandemir</i>	
PAC-PLRU: A Cache Replacement Policy to Salvage Discarded Predictions from Hardware Prefetchers	265
<i>Ke Zhang, Zhensong Wang, Yong Chen, Huaiyu Zhu, and Xian-He Sun</i>	

Contention Modeling for Multithreaded Distributed Shared Memory Machines: The Cray XMT	275
<i>Simone Secchi, Antonino Tumeo, and Oreste Villa</i>	
Data-Driven Computing	
Predictive Data Grouping and Placement for Cloud-Based Elastic Server Infrastructures	285
<i>Juan M. Tirado, Daniel Higuero, Florin Isaila, and Jesús Carretero</i>	
BAR: An Efficient Data Locality Driven Task Scheduling Algorithm for Cloud Computing	295
<i>Jiahui Jin, Junzhou Luo, Aibo Song, Fang Dong, and Runqun Xiong</i>	
Fault Tolerance and Checkpointing	
On the Scheduling of Checkpoints in Desktop Grids	305
<i>Mohamed Slim Bouguerra, Derrick Kondo, and Denis Trystram</i>	
High Performance Pipelined Process Migration with RDMA	314
<i>Xiangyong Ouyang, Raghunath Rajachandrasekar, Xavier Besseron, and Dhabaleswar K. Panda</i>	
Failure Avoidance through Fault Prediction Based on Synthetic Transactions	324
<i>Mohammed Shatnawi and Matei Ripeanu</i>	
Communication and Network Management	
A Scalable Method for Signalling Dynamic Reconfiguration Events with OpenSM	332
<i>Wei Lin Guay and Sven-Arne Reinemo</i>	
On the Relation between Congestion Control, Switch Arbitration and Fairness	342
<i>Ernst Gunnar Gran, Eitan Zahavi, Sven-Arne Reinemo, Tor Skeie, Gilad Shainer, and Olav Lysne</i>	
Network-Friendly One-Sided Communication through Multinode Cooperation on Petascale Cray XT5 Systems	352
<i>Xinyu Que, Weikuan Yu, Vinod Tipparaju, Jeffrey S. Vetter, and Bin Wang</i>	
Distributed Hash Tables	
Evaluating and Optimizing Indexing Schemes for a Cloud-Based Elastic Key-Value Store	362
<i>David Chiu, Apeksha Shetty, and Gagan Agrawal</i>	
Sophia: Local Trust for Securing Routing in DHTs	372
<i>Raúl Gracia-Tinedo, Pedro García-López, and Marc Sánchez-Artigas</i>	
The Benefits of Estimated Global Information in DHT Load Balancing	382
<i>Nico Kruber, Mikael Höggvist, and Thorsten Schütt</i>	

I/O and File Systems 1

DHTbd: A Reliable Block-Based Storage System for High Performance Clusters	392
<i>George Parisis, George Xylomenos, and Theodore Apostolopoulos</i>	
Adaptive QoS Decomposition and Control for Storage Cache Management in Multi-server Environments	402
<i>Ramya Prabhakar, Shekhar Srikantaiah, Rajat Garg, and Mahmut Kandemir</i>	
A Segment-Level Adaptive Data Layout Scheme for Improved Load Balance in Parallel File Systems	414
<i>Huaiming Song, Yanlong Yin, Xian-He Sun, Rajeev Thakur, and Samuel Lang</i>	

QoS

Classification and Composition of QoS Attributes in Distributed, Heterogeneous Systems	424
<i>Elisabeth Vinek, Peter Paul Beran, and Erich Schikuta</i>	
Autonomic SLA-Driven Provisioning for Cloud Applications	434
<i>Nicolas Bonvin, Thanasis G. Papaioannou, and Karl Aberer</i>	
A Flexible Policy Framework for the QoS Differentiated Provisioning of Services	444
<i>Mohan Baruwal Chhetri, Bao Quoc Vo, and Ryszard Kowalczyk</i>	

Data Intensive Computing and MapReduce

DELMA: Dynamically ELastic MapReduce Framework for CPU-Intensive Applications	454
<i>Zacharia Fadika and Madhusudhan Govindaraju</i>	
Cloud MapReduce: A MapReduce Implementation on Top of a Cloud Operating System	464
<i>Huan Liu and Dan Orban</i>	
Ex-MATE: Data Intensive Computing with Large Reduction Objects and Its Application to Graph Mining	475
<i>Wei Jiang and Gagan Agrawal</i>	

I/O and File Systems 2

ASDF: An Autonomous and Scalable Distributed File System	485
<i>Chien-Ming Wang, Chi-Chang Huang, and Huan-Ming Liang</i>	
Managing Distributed Files with RNS in Heterogeneous Data Grids	494
<i>Yutaka Kawaj, Go Iwai, Takashi Sasaki, and Yoshiyuki Watase</i>	
DDFTP: Dual-Direction FTP	504
<i>Jameela Al-Jaroodi and Nader Mohamed</i>	

Efficient Support for MPI-I/O Atomicity Based on Versioning	514
<i>Viet-Trung Tran, Bogdan Nicolae, Gabriel Antoniu, and Luc Bougé</i>	

Security

Implementing Trust in Cloud Infrastructures	524
<i>Ricardo Neisse, Dominik Holling, and Alexander Pretschner</i>	
Detection and Protection against Distributed Denial of Service Attacks in Accountable Grid Computing Systems	534
<i>Wonjun Lee, Anna C. Squicciarini, and Elisa Bertino</i>	
Dealing with Grid-Computing Authorization Using Identity-Based Certificateless Proxy Signature	544
<i>Mohamed Amin Jabri and Satoshi Matsuoka</i>	

Social Network (SN4CCGridS) Workshop

Open Social Based Collaborative Science Gateways	554
<i>Wenjun Wu, Hui Zhang, and ZhenAn Li</i>	
Social Networks of Researchers and Educators on nanoHUB.org	560
<i>Gerhard Klimeck, George B. Adams III, Krishna P.C. Madhavan, Nathan Denny, Michael G. Zentner, Swaroop Shivarajapura, Lynn K. Zentner, and Diane L. Beaudoin</i>	
A Trustworthiness Fusion Model for Service Cloud Platform Based on D-S Evidence Theory	566
<i>Rong Hu, Jianxun Liu, and Xiaoqing Frank Liu</i>	
Engineering Incentives in Social Clouds	572
<i>Christian Haas, Simon Caton, and Christof Weinhardt</i>	

Clouds for Business, Industry, and Enterprise (C4BIE) Workshop

Utilizing “Opaque” Resources for Revenue Enhancement on Clouds and Grids	576
<i>Jose Orlando Melendez and Shikharesh Majumdar</i>	
Debunking Real-Time Pricing in Cloud Computing	585
<i>Sewook Wee</i>	
Unifying Cloud Management: Towards Overall Governance of Business Level Objectives	591
<i>Mina Sedaghat, Francisco Hernández, and Erik Elmroth</i>	
Defining a Cloud Reference Model	598
<i>Teresa Tung</i>	

Poster Abstracts

Inferring Network Topologies in Infrastructure as a Service Cloud	604
<i>Dominic Battré, Natalia Frejnik, Siddhant Goel, Odej Kao, and Daniel Warneke</i>	
Addressing Resource Fragmentation in Grids through Network-Aware Meta-scheduling in Advance	606
<i>Luis Tomás, Carmen Carrión, Blanca Caminero, and Agustín Caminero</i>	
Performance under Failures of MapReduce Applications	608
<i>Hui Jin, Kan Qiao, Xian-He Sun, and Ying Li</i>	
MPI-IO/Gfarm: An Optimized Implementation of MPI-IO for the Gfarm File System	610
<i>Hiroki Kimura and Osamu Tatebe</i>	
Diagnosing Anomalous Network Performance with Confidence	612
<i>Bradley W. Settlemyer, Stephen W. Hodson, Jeffery A. Kuehn, and Stephen W. Poole</i>	
A Performance Goal Oriented Processor Allocation Technique for Centralized Heterogeneous Multi-cluster Environments	614
<i>Po-Chi Shih, Kuo-Chan Huang, Che-Rung Lee, I-Hsin Chung, and Yeh-Ching Chung</i>	
A Hybrid Shared-Nothing/Shared-Data Storage Architecture for Large Scale Databases	616
<i>Huaiming Song, Xian-He Sun, and Yong Chen</i>	
EZTrace: A Generic Framework for Performance Analysis	618
<i>François Trahay, François Rue, Mathieu Faverge, Yutaka Ishikawa, Raymond Namyst, and Jack Dongarra</i>	
Supporting Federated Multi-authority Security Models	620
<i>John Watt and Richard O. Sinnott</i>	
Author Index	622