

2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing

(CCGrid 2010)

**Melbourne, Australia
17 – 20 May 2010**



**IEEE Catalog Number: CFP10276-PRT
ISBN: 978-1-4244-6987-1**

2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing

CCGrid 2010

Table of Contents

Message from the General Chair	xvi
Message from the Program Chair	xviii
Message from the Workshops Chair	xx
Organising Committee	xxii
Program Committee Members	xxv
CCGrid 2010 Sponsors	xxvii

KEYNOTES

Enabling the Next Generation of Scalable Clusters	3
<i>William D. Gropp</i>	
Sky Computing: When Multiple Clouds Become One	4
<i>José A.B. Fortes</i>	

REGULAR PAPERS

Session R1A: Algorithms—Cloud Computing and Grids

Dynamic Load-Balanced Multicast for Data-Intensive Applications on Clouds	5
<i>Tatsuhiko Chiba, Mathijs den Burger, Thilo Kielmann, and Satoshi Matsuoka</i>	
Profit-Driven Service Request Scheduling in Clouds	15
<i>Young Choon Lee, Chen Wang, Albert Y. Zomaya, and Bing Bing Zhou</i>	
Availability Prediction Based Replication Strategies for Grid Environments	25
<i>Brent Rood and Michael J. Lewis</i>	
EGSI: TGKA Based Security Architecture for Group Communication in Grid	34
<i>Rajesh Ingle and G. Sivakumar</i>	

Session R1B: Middleware/Runtime—Resource Management

Elastic Site: Using Clouds to Elastically Extend Site Resources	43
<i>Paul Marshall, Kate Keahey, and Tim Freeman</i>	
ConnectX-2 InfiniBand Management Queues: First Investigation of the New Support for Network Offloaded Collective Operations	53
<i>Richard L. Graham, Steve Poole, Pavel Shamis, Gil Bloch, Noam Bloch, Hillel Chapman, Michael Kagan, Ariel Shahar, Ishai Rabinovitz, and Gilad Shainer</i>	
Distributed Diskless Checkpoint for Large Scale Systems	63
<i>Leonardo Arturo Bautista Gomez, Naoya Maruyama, Franck Cappello, and Satoshi Matsuoka</i>	
Enabling Instantaneous Relocation of Virtual Machines with a Lightweight VMM Extension	73
<i>Takahiro Hirofuchi, Hidemoto Nakada, Satoshi Itoh, and Satoshi Sekiguchi</i>	

Session R2A: Applications—Clouds

A Map-Reduce System with an Alternate API for Multi-core Environments	84
<i>Wei Jiang, Vignesh T. Ravi, and Gagan Agrawal</i>	
An Analysis of Traces from a Production MapReduce Cluster	94
<i>Soila Kavulya, Jiaqi Tan, Rajeev Gandhi, and Priya Narasimhan</i>	
An Effective Architecture for Automated Appliance Management System Applying Ontology-Based Cloud Discovery	104
<i>Amir Vahid Dastjerdi, Sayed Gholam Hassan Tabatabaei, and Rajkumar Buyya</i>	

Session R2B: Middleware/Runtime—Program Optimization and Scheduling

Region-Based Prefetch Techniques for Software Distributed Shared Memory Systems	113
<i>Jie Cai, Peter E. Strazdins, and Alistair P. Rendell</i>	
Granularity-Aware Work-Stealing for Computationally-Uniform Grids	123
<i>Vladimir Janjic and Kevin Hammond</i>	
SAGA BigJob: An Extensible and Interoperable Pilot-Job Abstraction for Distributed Applications and Systems	135
<i>André Luckow, Lukasz Lacinski, and Shantenu Jha</i>	

Session R3A: Programming Models and Systems—HPC and Accelerators

Remote Process Execution and Remote File I/O for Heterogeneous Processors in Cluster Systems	145
<i>Masaaki Shimizu and Akinori Yonezawa</i>	
An Adaptive Data Prefetcher for High-Performance Processors	155
<i>Yong Chen, Huaiyu Zhu, and Xian-He Sun</i>	
Designing Accelerator-Based Distributed Systems for High Performance	165
<i>M. Mustafa Rafique, Ali R. Butt, and Dimitrios S. Nikolopoulos</i>	

Efficient On-Demand Connection Management Mechanisms with PGAS Models over InfiniBand	175
<i>Abhinav Vishnu and Manoj Krishnan</i>	
Session R3B: Performance Modeling and Evaluation—Scheduling and Resource Management	
Methodology for Efficient Execution of SPMD Applications on Multicore Environments	185
<i>Ronal Muresano, Dolores Rexachs, and Emilio Luque</i>	
An Evaluation of the Benefits of Fine-Grained Value-Based Scheduling on General Purpose Clusters	196
<i>Ruben Van den Bossche, Kurt Vanmechelen, and Jan Broeckhove</i>	
The Effects of Untruthful Bids on User Utilities and Stability in Computing Markets	205
<i>Sergei Shudler, Lior Amar, Amnon Barak, and Ahuva Mu'alem</i>	
FIRE: A File Reunion Based Data Replication Strategy for Data Grids	215
<i>Abdul Rahman Abdurrah and Tao Xie</i>	
Session R4A: Algorithms—Scheduling and Resource Allocation	
SAQA: A Self-Adaptive QoS-Aware Scheduling Algorithm for Real-Time Tasks on Heterogeneous Clusters	224
<i>Xiaomin Zhu, Jianghan Zhu, Manhao Ma, and Dishan Qiu</i>	
Bandwidth Allocation for Iterative Data-Dependent E-science Applications	233
<i>Eun-Sung Jung, Sanjay Ranka, and Sartaj Sahni</i>	
A Bi-criteria Algorithm for Scheduling Parallel Task Graphs on Clusters	243
<i>Frédéric Desprez and Frédéric Suter</i>	
Low-Cost Tuning of Two-Step Algorithms for Scheduling Mixed-Parallel Applications onto Homogeneous Clusters	253
<i>Sascha Hunold</i>	
Session R4B: Middleware/Runtime—Service Management and Workflows	
ERGOT: A Semantic-Based System for Service Discovery in Distributed Infrastructures	263
<i>Giuseppe Pirrò, Paolo Trunfio, Domenico Talia, Paolo Missier, and Carole Goble</i>	
Towards Autonomic Service Provisioning Systems	273
<i>Michele Mazzucco</i>	
WORKEM: Representing and Emulating Distributed Scientific Workflow Execution State	283
<i>Lavanya Ramakrishnan, Dennis Gannon, and Beth Plale</i>	

Experiments with Memory-to-Memory Coupling for End-to-End Fusion Simulation Workflows	293
<i>Ciprian Docan, Fan Zhang, Manish Parashar, Julian Cummings, Norbert Podhorszki, and Scott Klasky</i>	
Session R5: Programming Models and Systems—Streams	
Streamflow—Programming Model for Data Streaming in Scientific Workflows	302
<i>Chathura Herath and Beth Plale</i>	
Representing Eager Evaluation in a Demand Driven Model of Streams on Cloud Infrastructure	312
<i>Paul N. Martinaitis and Andrew L. Wendelborn</i>	
An MPI-Stream Hybrid Programming Model for Computational Clusters	323
<i>Emilio P. Mancini, Gregory Marsh, and Dhabaleswar K. Panda</i>	
Session R6: Applications	
High Performance Dimension Reduction and Visualization for Large High-Dimensional Data Analysis	331
<i>Jong Youl Choi, Seung-Hee Bae, Xiaohong Qiu, and Geoffrey Fox</i>	
Exploring the Potential of Using Multiple E-science Infrastructures with Emerging Open Standards-Based E-health Research Tools	341
<i>M. Riedel, B. Schuller, M. Rambadt, M.S. Memon, A.S. Memon, A. Streit, Th. Lippert, S.J. Zasada, S. Manos, P.V. Coveney, Felix Wolf, and Dieter Kranzlmüller</i>	
Efficient Runtime Environment for Coupled Multi-physics Simulations: Dynamic Resource Allocation and Load-Balancing	349
<i>Soon-Heum Ko, Nayong Kim, Joohyun Kim, Abhinav Thota, and Shantenu Jha</i>	
On-demand Overlay Networks for Large Scientific Data Transfers	359
<i>Lavanya Ramakrishnan, Chin Guok, Keith Jackson, Ezra Kissel, D. Martin Swamy, and Deborah Agarwal</i>	
Session R7A: Algorithms and Applications—Energy	
Towards Energy Aware Scheduling for Precedence Constrained Parallel Tasks in a Cluster with DVFS	368
<i>Lizhe Wang, Gregor von Laszewski, Jay Dayal, and Fugang Wang</i>	
Runtime Energy Adaptation with Low-Impact Instrumented Code in a Power-Scalable Cluster System	378
<i>Hideaki Kimura, Takayuki Imada, and Mitsuhsisa Sato</i>	
Linear Combinations of DVFS-Enabled Processor Frequencies to Modify the Energy-Aware Scheduling Algorithms	388
<i>Nikzad Babaii Rizvandi, Javid Taheri, Albert Y. Zomaya, and Young Choon Lee</i>	

Session R7B: Performance Modeling and Evaluation—Tracing and Communication

The Failure Trace Archive: Enabling Comparative Analysis of Failures in Diverse Distributed Systems	398
<i>Derrick Kondo, Bahman Javadi, Alexandru Iosup, and Dick Epema</i>	
Scalable Communication Trace Compression	408
<i>Sriram Krishnamoorthy and Khushbu Agarwal</i>	
FaReS: Fair Resource Scheduling for VMM-Bypass InfiniBand Devices.....	418
<i>Adit Ranadive, Ada Gavrilovska, and Karsten Schwan</i>	

Session R8A: Algorithms—Self-Organizing and Peer-to-Peer Systems

A Proximity-Based Self-Organizing Framework for Service Composition and Discovery	428
<i>Agostino Forestiero, Carlo Mastroianni, Giuseppe Papuzzo, and Giandomenico Spezzano</i>	
Dynamic TTL-Based Search in Unstructured Peer-to-Peer Networks	438
<i>Imen Filali and Fabrice Huet</i>	
Enhanced Paxos Commit for Transactions on DHTs	448
<i>Florian Schintke, Alexander Reinefeld, Seif Haridi, and Thorsten Schütt</i>	
Cache Performance Optimization for Processing XML-Based Application Data on Multi-core Processors	455
<i>Rajdeep Bhowmik and Madhusudhan Govindaraju</i>	

Session R8B: Performance Modeling and Evaluation—Workload Modeling and Prediction

A Realistic Integrated Model of Parallel System Workloads	464
<i>Tran Ngoc Minh, Lex Wolters, and Dick Epema</i>	
Discovering Piecewise Linear Models of Grid Workload	474
<i>Tamás Élteto, Cécile Germain-Renaud, Pascal Bondon, and Michèle Sebag</i>	
Identification, Modelling and Prediction of Non-periodic Bursts in Workloads	485
<i>Mario Lassnig, Thomas Fahringer, Vincent Garonne, Angelos Molfetas, and Miguel Branco</i>	
On the Use of Machine Learning to Predict the Time and Resources Consumed by Applications	495
<i>Andréa Matsunaga and José A.B. Fortes</i>	

SHORT PAPERS

Session S1: Cloud Computing and Applications

On the Origin of Services—Using RIDDL for Description, Evolution and Composition of RESTful Services	505
<i>Juergen Mangler, Peter Paul Beran, and Erich Schikuta</i>	
A Categorisation of Cloud Computing Business Models	509
<i>Victor Chang, David Bacigalupo, Gary Wills, and David De Roure</i>	
Dynamic Resource Pricing on Federated Clouds	513
<i>Marian Mihailescu and Yong Meng Teo</i>	
Unibus-managed Execution of Scientific Applications on Aggregated Clouds	518
<i>Jaroslav Slawinski, Magdalena Slawinska, and Vaidy Sunderam</i>	

Session S2: Grid and E-science Applications

File-Access Characteristics of Data-Intensive Workflow Applications	522
<i>Takeshi Shibata, SungJun Choi, and Kenjiro Taura</i>	
Overdimensioning for Consistent Performance in Grids	526
<i>Nezih Yigitbasi and Dick Epema</i>	
Topology Aggregation for E-science Networks	530
<i>Eun-Sung Jung, Sanjay Ranka, and Sartaj Sahni</i>	
Handling Recoverable Temporal Violations in Scientific Workflow Systems: A Workflow Rescheduling Based Strategy	534
<i>Xiao Liu, Jinjun Chen, Zhangjun Wu, Zhiwei Ni, Dong Yuan, and Yun Yang</i>	
A Fair Decentralized Scheduler for Bag-of-Tasks Applications on Desktop Grids	538
<i>Javier Celaya and Loris Marchal</i>	
A Heuristic Query Optimization Approach for Heterogeneous Environments	542
<i>Peter Paul Beran, Werner Mach, Ralph Vigne, Jürgen Mangler, and Erich Schikuta</i>	

Session S3: Data Management in Grids

Planning Large Data Transfers in Institutional Grids	547
<i>Fatiha Bouabache, Thomas Herault, Sylvain Peyronnet, and Franck Cappello</i>	
Framework for Efficient Indexing and Searching of Scientific Metadata	553
<i>Chaitali Gupta and Madhusudhan Govindaraju</i>	
High Performance Data Transfer in Grid Environment Using GridFTP over InfiniBand	557
<i>Hari Subramoni, Ping Lai, Raj Kettimuthu, and Dhabaleswar K. Panda</i>	
Data Injection at Execution Time in Grid Environments Using Dynamic Data Driven Application System for Wildland Fire Spread Prediction	565
<i>Roque Rodríguez, Ana Cortés, and Tomás Margalef</i>	

POSTER PAPERS

Expanding the Cloud: A Component-Based Architecture to Application Deployment on the Internet	569
<i>Mark Wallis, Frans Henskens, and Michael Hannaford</i>	
Fine-Grained Profiling for Data-Intensive Workflows	571
<i>Nan Dun, Kenjiro Taura, and Akinori Yonezawa</i>	
Supporting OFED over Non-InfiniBand SANs	573
<i>Devesh Sharma</i>	
The Lightweight Approach to Use Grid Services with Grid Widgets on Grid WebOS	575
<i>Yi-Lun Pan, Chang-Hsing Wu, Chia-Yen Liu, Hsi-En Yu, and Weicheng Huang</i>	
Energy Efficient Allocation of Virtual Machines in Cloud Data Centers	577
<i>Anton Beloglazov and Rajkumar Buyya</i>	
SciCloud: Scientific Computing on the Cloud	579
<i>Satish Srirama, Oleg Batrashev, and Eero Vainikko</i>	
Rigel: A Scalable and Lightweight Replica Selection Service for Replicated Distributed File System	581
<i>Yuan Lin, Yang Chen, Guodong Wang, and Beixing Deng</i>	
In Search of Visualization Metaphors for PlanetLab	583
<i>Andrew J. Zaliwski</i>	
Design and Implementation of an Efficient Two-Level Scheduler for Cloud Computing Environment	585
<i>R. Jeyarani, R. Vasanth Ram, and N. Nagaveni</i>	
Cluster Computing as an Assembly Process: Coordination with S-Net	587
<i>Clemens Grellck, Jukka Julku, Frank Penczek, and Alex Shafarenko</i>	
Dynamic Job-Clustering with Different Computing Priorities for Computational Resource Allocation	589
<i>Masnida Hussin, Young Choon Lee, and Albert Y. Zomaya</i>	
Dynamic Auction Mechanism for Cloud Resource Allocation	591
<i>Wei-Yu Lin, Guan-Yu Lin, and Hung-Yu Wei</i>	
Policy-Based Management of QoS in Service Aggregations	593
<i>Mohan Baruwal Chhetri, Bao Quoc Vo, and Ryszard Kowalczyk</i>	
Feedback-Guided Analysis for Resource Requirements in Large Distributed System	596
<i>Madhulina Sarkar, Sarbani Roy, and Nandini Mukherjee</i>	
TOPP goes Rapid—The OpenMS Proteomics Pipeline in a Grid-Enabled Web Portal	598
<i>Sandra Gesing, Jano van Hemert, Jos Koetsier, Andreas Bertsch, and Oliver Kohlbacher</i>	

Second International Symposium on Cloud Computing (Cloud 2010)

TrustStore: Making Amazon S3 Trustworthy with Services Composition	600
<i>Jinhui Yao, Shiping Chen, Surya Nepal, David Levy, and John Zic</i>	
Polyphony: A Workflow Orchestration Framework for Cloud Computing	606
<i>Khawaja S. Shams, Mark W. Powell, Tom M. Crockett, Jeffrey S. Norris, Ryan Rossi, and Tom Soderstrom</i>	
Virtual Resources Allocation for Workflow-Based Applications Distribution on a Cloud Infrastructure	612
<i>Tram Truong Huu and Johan Montagnat</i>	
Applying Software Engineering Principles for Designing Cloud@Home	618
<i>Vincenzo D. Cunsolo, Salvatore Distefano, Antonio Puliafito, and Marco Scarpa</i>	
User Requirements for Cloud Computing Architecture	625
<i>Roger Clarke</i>	
D-Cloud: Design of a Software Testing Environment for Reliable Distributed Systems Using Cloud Computing Technology	631
<i>Takayuki Banzai, Hitoshi Koizumi, Ryo Kanbayashi, Takayuki Imada, Toshihiro Hanawa, and Mitsuhsa Sato</i>	

Fourth Workshop on Desktop Grids and Volunteer Computing Systems (PCGrid 2010)

Towards Trust in Desktop Grid Systems	637
<i>Yvonne Bernard, Lukas Klejnowski, Jörg Hähner, and Christian Müller-Schloer</i>	
Decentralized Resource Availability Prediction for a Desktop Grid	643
<i>Karthick Ramachandran, Hanan Lutfyya, and Mark Perry</i>	
Predicting the Quality of Service of a Peer-to-Peer Desktop Grid	649
<i>Marcus Carvalho, Renato Miceli, Paulo Ditarso Maciel Jr., Francisco Brasileiro, and Raquel Lopes</i>	
Generalized Spot-Checking for Sabotage-Tolerance in Volunteer Computing Systems	655
<i>Kan Watanabe and Masaru Fukushi</i>	
UnaGrid: On Demand Opportunistic Desktop Grid	661
<i>Harold Castro, Eduardo Rosales, Mario Villamizar, and Artur Jiménez</i>	
Integration of Heterogeneous and Non-dedicated Environments for R	667
<i>Gonzalo Vera and Remo Suppi</i>	
A High-Level Interpreted MPI Library for Parallel Computing in Volunteer Environments	673
<i>Troy P. LeBlanc, Jaspal Subhlok, and Edgar Gabriel</i>	
mPlogP: A Parallel Computation Model for Heterogeneous Multi-core Computer	679
<i>Liang Li, Xingjun Zhang, Jinghua Feng, and Xiaoshe Dong</i>	

Extending the EGEE Grid with XtremWeb-HEP Desktop Grids	685
<i>Haiwu He, Gilles Fedak, Peter Kacsuk, Zoltan Farkas, Zoltan Balaton, Oleg Lodygensky, Etienne Urbah, Gabriel Caillat, Filipe Araujo, and Ad Emmen</i>	
Resiliency in High Performance Computing (Resilience 2010)	
Hard Data on Soft Errors: A Large-Scale Assessment of Real-World Error Rates in GPGPU	691
<i>Imran S. Haque and Vijay S. Pande</i>	
Team-Based Message Logging: Preliminary Results	697
<i>Esteban Meneses, Celso L. Mendes, and Laxmikant V. Kalé</i>	
Using Cloud Constructs and Predictive Analysis to Enable Pre-Failure Process Migration in HPC Systems	703
<i>James Brandt, Frank Chen, Vincent De Sapio, Ann Gentile, Jackson Mayo, Philippe Pébay, Diana Roe, David Thompson, and Matthew Wong</i>	
Selective Recovery from Failures in a Task Parallel Programming Model	709
<i>James Dinan, Arjun Singri, P. Sadayappan, and Sriram Krishnamoorthy</i>	
Fifth International Workshop on Content Delivery Networks (CDN 2010)	
Mobility Support Through Caching in Content-Based Publish/Subscribe Networks	715
<i>Vasilis Sourlas, Georgios S. Paschos, Paris Flegkas, and Leandros Tassioulas</i>	
Multi-criteria Content Adaptation Service Selection Broker	721
<i>Mohd Farhan Md Fudzee, Jemal Abawajy, and Mustafa Mat Deris</i>	
User Provided Cloud Computing	727
<i>Cláudio Teixeira, Ricardo Azevedo, Joaquim Sousa Pinto, and Tiago Batista</i>	
Challenges for the Application of Grids in Healthcare (CCGrid-Health 2010)	
Gridifying a Diffusion Tensor Imaging Analysis Pipeline	733
<i>Matthan W.A. Caan, Frans M. Vos, Antoine H.C. van Kampen, Silvia D. Olabarriaga, and Lucas J. van Vliet</i>	
Overview of Medical Data Management Solutions for Research Communities	739
<i>Sorina Camarasu-Pop, Frederic Cervenansky, Yonny Cardenas, Jean-Yves Nief, and Hugues Benoit-Cattin</i>	
Development and Support of Platforms for Research into Rare Diseases	745
<i>Richard O. Sinnott, Jipu Jiang, Anthony Stell, and John Watt</i>	
Performance Analysis of Diffusion Tensor Imaging in an Academic Production Grid	751
<i>Dagmar Krefting, Ralf Luetzkendorf, Kathrin Peter, and Johannes Bernarding</i>	

CCGrid-Multicore 2010 “Frontiers of GPU, Multi- and Many-core Systems”

Programming Challenges for the Implementation of Numerical Quadrature in Atomic Physics on FPGA and GPU Accelerators	757
<i>C.J. Gillan, T. Steinke, J. Bock, S. Borchert, I. Spence, and N.S. Scott</i>	
Asynchronous Communication Schemes for Finite Difference Methods on Multiple GPUs	763
<i>Daniel Peter Playne and Kenneth Arthur Hawick</i>	
Solving k-Nearest Neighbor Problem on Multiple Graphics Processors	769
<i>Kimikazu Kato and Tikara Hosino</i>	
Cooperative Multitasking for GPU-Accelerated Grid Systems	774
<i>Fumihiko Ino, Akihiro Ogita, Kentaro Oita, and Kenichi Hagihara</i>	
Multi-FFT Vectorization for the Cell Multicore Processor	780
<i>Jacob Barhen, Travis Humble, Pramita Mitra, and Michael Traweek</i>	
High Resolution Program Flow Visualization of Hardware Accelerated Hybrid Multi-core Applications	786
<i>Daniel Hackenberg, Guido Juckeland, and Holger Brunst</i>	
Running the NIM Next-Generation Weather Model on GPUs	792
<i>Mark W. Govett, Jacques Middlecoff, and Tom Henderson</i>	
Accelerating Climate and Weather Simulations Through Hybrid Computing	797
<i>Shujia Zhou, Carlos Cruz, Daniel Duffy, Robert Tucker, and Mark Purcell</i>	
A Memory Centric Kernel Framework for Accelerating Short-Range, Interactive Particle Simulation	802
<i>Ian Stewart and Shujia Zhou</i>	
From Sparse Matrix to Optimal GPU CUDA Sparse Matrix Vector Product Implementation	808
<i>Ahmed H. El Zein and Alistair P. Rendell</i>	
Performance of Windows Multicore Systems on Threading and MPI	814
<i>Judy Qiu, Scott Beason, Seung-Hee Bae, Saliya Ekanayake, and Geoffrey Fox</i>	
Doctoral Symposium	
Service Oriented Approach to High Performance Scientific Computing	820
<i>Jaison Paul Mulerikkal and Peter Strazdins</i>	
Energy Efficient Resource Management in Virtualized Cloud Data Centers	826
<i>Anton Beloglazov and Rajkumar Buyya</i>	
SLA-Driven Dynamic Resource Management for Multi-tier Web Applications in a Cloud	832
<i>Waheed Iqbal, Matthew N. Dailey, and David Carrera</i>	
On Economic and Computational-Efficient Resource Pricing in Large Distributed Systems	838
<i>Marian Mihalescu and Yong Meng Teo</i>	

A Capabilities-Aware Programming Model for Asymmetric High-End Systems	844
<i>M. Mustafa Rafique</i>	
Author Index	850