

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

Program


CCGrid 2010

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

17-20 May 2010
Melbourne, Australia

Edited by Manish Parashar and Rajkumar Buyya

May 17-20, 2010, Melbourne, Victoria, Australia



<http://www.manjrasoft.com/ccgrid2010>

Message from the General Chair

I am pleased to welcome you to the 10th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid 2010) sponsored by the IEEE Computer Society, IEEE Technical Committee on Scalable Computing (TCSC), and Association for Computing Machinery (ACM).

CCGrid is an important conference for the international community as it provides a forum for all cluster, cloud, and Grid computing researchers, developers and users, and those who are just curious to see the project results and become aware of the progress made in these areas.

The inaugural CCGrid conference was held in Brisbane, Australia in 2001. Since then, the conference has successfully been hosted around the world and has emerged as a truly global event. From 2002 to 2009, CCGrid annual events were held in Germany, Japan, USA, UK, Singapore, Brazil, France and China. Returning back to its originating country, we are honored to host the 10th anniversary of the CCGrid conference in Melbourne, Australia in 2010.

CCGrid has been featuring original and outstanding research work in Cluster, Grid, and Cloud Computing. In fact, many emerging research trends and associated publications are featured “first” in CCGrid and their follow-up papers have appeared in other conferences later. This demonstrates emergence of CCGrid as a “first” class venue for presenting original and ground breaking works. For instance, CCGrid has been featuring Cloud computing actively during the last few years, which emerged as one of the major themes. Hence, from 2010, we explicitly recognized this growing trend in CCGrid by including “Cloud computing” in the conference title.

This 10th anniversary, CCGrid 2010 conference offers outstanding technical program featuring keynote talks, tutorials, workshops, mini-symposiums, posters sessions, industry track, research exhibits and demos, and IEEE SCALE competition.

CCGrid has been extremely fortunate to serve as a venue for presentation of prestigious “**IEEE Medal for Excellence in Scalable Computing**” award offered annually by the IEEE Technical Committee on Scalable Computing. This year, we are pleased to host the 2010 Medal winner Professor William Groppe from the University of Illinois Urbana-Champaign, USA as the opening keynote speaker. We are also fortunate to host a keynote by Professor José Fortes from the University of Florida, USA.

The continued success and leadership of CCGrid requires dedicated and high standard efforts from numerous international volunteers. As the Chair of CCGrid conference series and General Chair of this year’s event, I would like to express my sincere gratitude to the members of the Steering Committee and the Program Committee chaired by Professor Manish Parashar. The Program Committee Chair and his Vice chairs (Professors Geoffrey Fox, David Bader, Carlos Varela, Thomas Fahringer, Dick Epema) have coordinated peer-reviews of all submitted “full” papers and selected top quality research papers for presentation at the conference.

The CCGrid 2010 conference received 219 submissions from 37 countries around the world: USA, China, Australia, Germany, France, Spain, India, Brazil, Japan, United Kingdom, Canada, The Netherlands, Iran, Korea, Italy, Austria, Israel, Serbia, Taiwan, Singapore, Belgium, Egypt, Malaysia, Colombia, Turkey, Sweden, Thailand, Switzerland, UAE, Pakistan, Hong Kong, Russia, New Zealand, Algeria, Greece, Tunisia, and Cyprus. After peer-review of all these “full” papers, the Program Committee accepted 51 papers, resulting in an acceptance rate of ~23%.

I thank Professor Omer Rana for coordinating the organisation of 8 satellite workshops/mini-symposiums on hot topics such as MultiCore Clusters, and Clouds for Business. We appreciate the efforts of the chairs of various workshops and their PC members for attracting and selecting top quality papers for presentation at the conference.

I thank Dr. Pavan Balaji for organising and managing the poster session, Suraj Pandey for the excellent management of the conference website, and publicity coordinators, Dr. Cho-Li Wang and Dr. Masoud Sadjadi, for helping us reach a broader community. I thank tutorials chair Professor Sushil K. Prasad and SCALE Challenge chairs, Dr. Shantenu Jha and Dr. Daniel S. Katz for their efforts in enhancing the conference program with interesting tutorials and demos. I thank Lisa O’Conner for her support in ensuring the publication of the conference proceedings in record time.

As we all know, the local arrangements are a key aspect of any event. I would like to offer my special appreciation to leading volunteers of local organizing committee Mukaddim Pathan, James Broberg, and Suraj Pandey for their dedicated work during the last one year. I would like to thank Kim Stevenson for managing registrations and Dushy Wanigatunga for his friendly services as Catering and Conventions Manager of The Langham Hotel.

Thanks are also due to our sponsors, namely, IEEE, ACM, and TCSC and organization supporters Melbourne University’s CLOUDS Lab, ISSNIP, NICTA Victoria Lab, NSF Center for Autonomic Computing at Rutgers University, Victoria Government (Australia), and Amazon. I also like to thank HPCWire, our media sponsor and Manjrasoft for sponsoring awards.

Ultimately, however, the success of the conference will be judged by how well the delegates have participated, learnt, interacted and established contacts with other researchers in different fields. The Committees and the sponsors have provided the funding, the venue, and the environment to allow these objectives to be achieved. It is now up to all of us to ensure that the conference is an outstanding success.

Finally, I wish everyone a successful, stimulating and rewarding meeting and look forward to seeing you again at future conferences.

Enjoy your visit to multicultural Melbourne and beautiful Australia!



Professor Rajkumar Buyya
Director, Cloud Computing and Distributed Systems (CLOUDS) Lab
Melbourne School of Engineering
The University of Melbourne, Australia
<http://www.cloudbus.org/>

CEO, Manjrasoft Pty Ltd, Melbourne, Australia
<http://www.manjrasoft.com/>

Keynote 1

Enabling the Next Generation of Scalable Clusters

William D. Gropp
Paul and Cynthia Saylor Professor of Computer Science
Computer Science Department
University of Illinois Urbana-Champaign
Urbana, Illinois, USA

Abstract:

Clusters revolutionized computing by making supercomputer capabilities widely available. But one of the main drivers of that revolution, the rapid doubling of processor clock rates, ran out of steam several years ago. To maintain (or even increase) the historic rate of improvement in computing power, processor designs are rapidly increasing parallelism at all levels, including more functional units, more cores, and ways to share resources among threads. Heterogeneous designs that use more specialized processors such as GPGPUs are becoming common. The scale of high-end systems is also getting larger, with 1000-core systems becoming commonplace and systems with over 300,000 cores planned for 2011. However, the software and algorithms for these systems are still basically the same as when the cluster revolution began. Drawing on experiences with the sustained PetaFLOPS system, called Blue Waters, to be installed at Illinois in 2011, and with exploratory work into Exascale system designs, this talk will discuss some of the challenges facing the cluster community as scalability becomes increasingly important and reviews some of the developments in algorithms, programming models, and software frameworks that must complement the evolution of cluster hardware.

Biography:

William Gropp is the Paul and Cynthia Saylor Professor in the Department of Computer Science and Deputy Directory for Research for the Institute of Advanced Computing Applications and Technologies at the University of Illinois in Urbana-Champaign. After receiving his Ph.D. in Computer Science from Stanford University in 1982, he held the positions of assistant (1982-1988) and associate (1988-1990) professor in the Computer Science Department of Yale University. In 1990, he joined the Numerical Analysis group at Argonne, where he held the positions of Senior Scientist (1998-2007) and Associate Division Director (2000-2006). His research interests are in parallel computing, software for scientific computing, and numerical methods for partial differential equations. He is a co-author of "Using MPI: Portable Parallel Programming with the Message-Passing Interface", and is a chapter author in the MPI-2 Forum. His current projects include the design and implementation of MPICH, a portable implementation of the MPI Message-Passing Standard, the design and implementation of PETSc, a parallel, numerical library for PDEs, and research into programming models for parallel architectures. He is a Fellow of ACM and IEEE and a member of the National Academy of Engineering.

Keynote 3

Sky Computing: When Multiple Clouds Become One

José A.B. Fortes
BellSouth Eminent Scholar
Director, NSF Center for Autonomic Computing
The University of Florida, Gainesville, USA

Abstract:

The growing number of announced commercial and scientific clouds strongly suggests that in the near future these providers will be differentiated according to the types of their services, their cost, availability and quality. Users will be able to use these and other criteria to determine which clouds best suit their needs, a plausible

scenario being the case when users need to aggregate capabilities provided by different clouds. In such scenarios it will be essential to provide virtual networking technologies that enable providers to support cross-cloud communication and users to deploy cross-cloud applications. This talk will describe one such technology, its salient features and remaining challenges. It will also put forward the idea of virtual clouds, i.e. providers of computing services overlaid on more than one cloud. A virtual cloud spans across multiple cloud providers and presents the view of a single logical cloud. Virtual clouds would enable high-level computing services to be provided by third parties who do not own physical resources, could be short or long-lived and highly dynamic. Enabling technologies, challenges and examples of sky computing will be presented.

Biography:

Jose Fortes is Professor and BellSouth Eminent Scholar at the University of Florida where he founded and directs both the Advanced Computing and Information Systems laboratory and the NSF Industry/University Cooperative Center for Autonomic Computing. His research interests are in the areas of distributed computing, autonomic computing, computer architecture, parallel processing and fault-tolerant computing. He has lead the development and deployment of Grid-computing software used in several cyber infrastructures for e-Science and digital government. They include In-VIGO, which was the first grid-computing system to use virtualization technologies, and PUNCH, which was an early example of a software-as-a-service provider. His research has been funded by, among others, the AT&T Foundation, Army Research Office, Citrix, General Electric, IBM, Intel, National Science Foundation, Northrop-Grumman, NASA, Office of Naval Research and the Semiconductor Research Corporation. Jose Fortes is a Fellow of the IEEE and has authored or coauthored over 180 technical papers.

CCGrid on site Registration Timing: Monday and Tuesday: 8.15am - 6.30pm Wed-Thurs: 8.30am - 4.00pm				
Day 1: May 17, 2010				
9:00 AM – 10:30 AM	CLOUD	PC Grid	Resilience	
10:30 AM – 11:00 AM	Morning Tea			
11:00 AM – 1:00 PM	CLOUD	PC Grid	Resilience	
1:00 PM – 2:00 PM	LUNCH			
2:00 PM – 3:30 PM	Content delivery Networks over Clouds (CDN)	PC Grid	CCGrid-Health	Tutorial Market-Oriented Cloud Computing Speaker: Prof. Dr. Rajkumar Buyya, Director of CLOUDS Lab, The University of Melbourne, Australia, CEO, Manjrasoft Pty Ltd, Melbourne, Australia
3:30 PM – 4:00 PM	Afternoon Tea			
4:00 PM – 6:00 PM	CDN	PC Grid	CCGrid-Health	Tutorial Continues...
Day 2: May 18, 2010				
9:00 AM – 10:30 AM	Opening & Keynote 1 Keynote 1 – Enabling the Next Generation of Scalable Clusters* *Professor William D. Gropp University of Illinois Urbana-Champaign, USA Winner of IEEE Scalable Computing Medal 2010			
10:30 AM – 11:00 AM	Morning Tea			
11:00 AM – 1:00 PM	Regular Papers	Regular Papers	MultiCore	
1:00 PM – 2:00 PM	LUNCH			
2:00 PM – 3:30 PM	Regular Papers	Regular Papers	MultiCore	
3:30 PM – 4:00 PM	Afternoon Tea			
4:00 PM – 6:00 PM	Regular Papers	Regular Papers	MultiCore	
Day 3: May 19, 2010				
9:00 AM – 10:45 AM	Industry Track – Cloud Panel 9:00 AM – 9:45 AM: Keynote 2 (Industry) “Cloud Computing with Amazon Web Services” - Simone Brunozzi, Amazon.com (AWS) 9:45 AM - 10:35 AM: Panel on “Cloud Deployment Trajectories for National Goals” – Chair Craig Lee, Open Grid Forum (OGF)			
10:35 AM – 11:00 AM	Morning Tea			
11:00 AM – 1:00 PM	Regular Papers	Regular Papers	Meeting	
1:00 PM – 2:00 PM	LUNCH			
2:00 PM – 3:30 PM	Regular Papers	Short Papers	SCALE	
3:30 PM – 4:00 PM	Afternoon Tea			
4:00 PM – 6:00 PM	Posters & Research Demos (Venue: Foyer) *SCALE" Public demos will be held during this period.			
6:30 PM – 9:30 PM	River Cruise & Conference Banquet			
Day 4: May 20, 2010				
9:00 AM – 10:00 AM	Keynote 3 Sky Computing: When Multiple Clouds Become One Professor José A.B. Fortes, University of Florida, USA			
10:00 AM – 10:30 AM	Morning Tea			
10:30 AM – 12:30 PM	Regular Papers	DocSymp	Short Papers	
12:30 PM – 1:30 PM	LUNCH			
1:30 PM – 3:00 PM	Regular Papers	Regular Papers	Short Papers	
3:00 PM – 3:30 PM	Afternoon Tea			
3:30 PM – 5:30 PM	Regular Papers	Regular Papers		
5:30 PM – 6:00 PM	Valedictory & Closing			

Day 1: May 17, 2010

9:00 AM – 10:30 PM	Parallel Sessions – Workshops (CLOUD, PC Grid, Resilience)
Venue: Flinders Room	Tutorial 1
10:30 AM – 11:00 AM	Morning Tea
11:00 AM – 1:00 PM	Parallel Sessions – Workshops (CLOUD, PC Grid, Resilience)
Venue: Ballroom A	<p style="text-align: center;">CLOUD Workshop</p> <p style="text-align: center;">Chairs - James Broberg, The University of Melbourne, Australia Bruno Schulze, National Laboratory for Scientific Computing, Brazil Rajkumar Buyya, The University of Melbourne, Australia</p>
<ul style="list-style-type: none"> • Invited Talk – Prof. Kai Hwang, University of Southern California, USA • TrustStore: Making Amazon S3 Trustworthy with Services Composition <i>Jinhui Yao, Shiping Chen, Surya Nepal, David Levy and John Zic</i> • Polyphony: A Workflow Orchestration Framework for Cloud Computing <i>Khawaja Shams, Mark Powell, Tom Crockett, Jeffrey Norris and Tom Soderstrom</i> • Virtual resources allocation for workflow-based applications distribution on a cloud infrastructure <i>Tram Truong Huu and Johan Montagnat</i> • Applying software engineering principles for designing Cloud@Home <i>Salvatore Distefano, Vincenzo Daniele Cunsolo, Antonio Puliafito and Marco Scarpa</i> • User Requirements for Cloud Computing Architecture <i>Roger Clarke</i> • D-Cloud: Design of a Software Testing Environment for Reliable Distributed Systems Using Cloud Computing Technology <i>Takayuki Banzai, Hitoshi Koizumi, Ryo Kanbayashi, Takayuki imada, Toshihiro Hanawa and Mitsuhsa Sato</i> 	
Venue: Ballroom B	<p style="text-align: center;">PC Grid Workshop (Session 1)</p> <p style="text-align: center;">Chairs - Gilles Fedak, INRIA, France Derrick Kondo, INRIA, France Bahman Javadi, INRIA, France</p>
<ul style="list-style-type: none"> • Towards Trust in Desktop Grid Systems <i>Yvonne Bernard, Lukas Klejnowski, Jörg Hähner and Christian Müller-Schloer</i> • Decentralized Resource Availability Prediction for a Desktop Grid <i>Karthick Ramachandran, Hanan Lutfiyya and Mark Perry</i> • Predicting the Quality of Service of a Peer-to-Peer Desktop Grid <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 <i>Kan Watanabe and Masaru Fukushi</i> • UnaGrid: On Demand Opportunistic Desktop Grid <i>Harold Enrique Castro Barrera, Edgar Eduardo Rosales Rosero, Mario José Villamizar Cano and Artur Miller Jimenez Castro</i> 	
Venue: Ballroom C	<p style="text-align: center;">Resilience Workshop</p> <p style="text-align: center;">Chairs - Stephen L. Scott, Oak Ridge National Laboratory, USA Chokchai (Box) Leangsuksun, Louisiana Tech University, USA Christian Engelmann, Oak Ridge National Laboratory, USA</p>
<ul style="list-style-type: none"> • Hard Data on Soft Errors: A Large-Scale Assessment of Real-World Error Rates in GPGPU <i>Imran S. Haque and Vijay S. Pande</i> • Team-based Message Logging <i>Esteban Meneses, Celso Mendes and Laxmikant Kale</i> • Using Cloud Constructs and Predictive Analysis to Enable Pre-Failure Process Migration in HPC Systems <i>James Brandt, Ann Gentile, Frank Chen, Vincent De Sapio, Jackson Mayo, Philippe Pebay, Diana Roe, David Thompson and Matthew Wong</i> • Selective Recovery From Failures In A Task Parallel Programming Model <i>James Dinan, Sriram Krishnamoorthy, Arjun Singri and P. Sadayappan</i> 	
1:00 PM – 2:00 PM	LUNCH

2:00 PM – 3:30 PM	Parallel Sessions – Workshops (CDN, PC Grid, Health)	Tutorial
3:30 PM – 4:00 PM	Afternoon Tea	
4:00 PM – 6:00 PM	Parallel Sessions – Workshops (CDN, PC Grid, Health)	Tutorial
Venue: Yarra Room 2:00pm to 6:00pm	<p style="text-align: center;">Tutorial Market-Oriented Cloud Computing Speaker: Prof. Dr. Rajkumar Buyya, Director of CLOUDS Lab, The University of Melbourne, Australia CEO, Manjrasoft Pty Ltd, Melbourne, Australia</p>	
Venue: Ballroom A	<p style="text-align: center;">CDN Workshop Chair – Adam Barker, University of Melbourne, Australia</p>	
<ul style="list-style-type: none"> • Invited Talk – "Cloud Technologies and Their Applications", Dr. Judy Qiu, Indiana University, USA • Mobility support through caching in content-based publish/subscribe networks <i>Vasilis Sourlas, Georgios Paschos, Paris Flegkas and Leandros Tassioulas</i> • Multi-criteria Content Adaptation Service Selection Broker <i>Jemal Abawajy and Mohd Farhan Md Fudzee</i> • User Provided Cloud Computing <i>Claudio Teixeira, Ricardo Azevedo, Joaquim S. Pinto and Tiago Batista</i> 		
Venue: Ballroom B	<p style="text-align: center;">PC Grid Workshop (Session 2) Chairs - Gilles Fedak, INRIA, France Derrick Kondo, INRIA, France Bahman Javadi, INRIA, France</p>	
<ul style="list-style-type: none"> • Integration of heterogeneous and non-dedicated environments for R <i>Gonzalo Vera and Remo Suppi</i> • A High-Level Interpreted MPI Library for Parallel Computing in Volunteer Environments <i>Troy LeBlanc, Jaspal Subhlok and Edgar Gabriel</i> • mPlogP: a Parallel Computation Model for Heterogeneous Multi-core Computer <i>Liang Li, Xingjun Zhang, Jinghua Feng and Xiaoshe Dong</i> • Extending the EGEE Grid with XtremWeb-HEP Desktop Grids <i>Haiwu he and Gilles Fedak</i> 		
Venue: Ballroom C	<p style="text-align: center;">Health Workshop Chair - Christophe Blanchet, IBCP, CNRS, FR Silvia D. Olabarriaga, University of Amsterdam, The Netherlands Tony Solomonides, University of the West of England, Bristol, UK Tristan Glatard, Creatis, CNRS, FR</p>	
<ul style="list-style-type: none"> • Invited Talk – Marianne Hibbert, BioGrid Australia, Melbourne Health, Australia • Gridifying a Diffusion Tensor Imaging Analysis Pipeline <i>Matthan W.A. Caan, Silvia D. Olabarriaga, Frans M. Vos, Lucas J. van Vliet, Antoine H.C. van Kampen</i> • Overview of Medical Data Management Solutions for Research Communities <i>Sorina Camarasu-Pop, Frédéric Cervenansky and Hugues Benoit-Cattin</i> • Development and Support of Platforms for Research into Rare Diseases <i>Richard Sinnott, Jipu Jiang, Anthony Stell and John Watt</i> • Runtime and Failure Analysis of Diffusion Tensor Imaging in an Academic Production Grid <i>Dagmar Krefting, Ralf Luetzkendorf, Kathrin Peter and Johannes Bernarding</i> 		

Day 2: May 18, 2010

9:00 AM – 10:30 AM	Conference Opening by Chairs (Buyya and Parashar) Keynote 1 – Enabling the Next Generation of Scalable Clusters* *Professor William D. Gropp University of Illinois Urbana-Champaign, USA Winner of IEEE Scalable Computing Medal 2010	
10:30 AM – 11:00 AM	Morning Tea	
11:00 AM – 1:00 PM	Parallel Sessions – Regular Papers (R1A & R1B)	MultiCore
Venue: Ballroom A	Session R1A: Algorithms – Cloud computing and Grids Chair - Cécile Germain-Renaud, Université Paris-Sud, France	
<ul style="list-style-type: none"> • Dynamic Load-Balanced Multicast for Data-Intensive Applications on Clouds <i>Tatsuhiko Chiba , Mathijs den Burger , Thilo Kielmann , Satoshi Matsuoka</i> • Profit-driven Service Request Scheduling in Clouds <i>Young Choon Lee , Chen Wang , Albert Zomaya , Bing-Bing Zhou</i> • Availability Prediction Based Replication Strategies for Grid Environments <i>Brent Rood , Michael Lewis</i> • EGS: TGKA based Security Architecture for Group Communication in Grid <i>Rajesh Ingle , G. Sivakumar</i> 		
Venue: Ballroom B	Session R1B: Middleware/Runtime – Resource Management Chair - Manish Parashar, Rutgers University, USA	
<ul style="list-style-type: none"> • Elastic Site: Using Clouds to Elastically Extend Site Resources <i>Paul Marshall , Kate Keahey , Timothy Freeman</i> • ConnectX2 InfiniBand Management Queues: New support for Network Offloaded Collective Operations <i>Richard Graham , Stephen Poole , Pavel Shamis , Gil Bloch , Noam Boch , Hillel Chapman , Michael Kagan , Ariel Shahar , Ishai Rabinovitz , Gilad Shainer</i> • Distributed Diskless Checkpoint for Large Scale Systems <i>Leonardo Bautista Gomez , Naoya Maruyama , Franck Cappello , Satoshi Matsuoka</i> • Enabling Instantaneous Relocation of Virtual Machines with a Lightweight VMM Extension <i>Takahiro Hirofuchi , Hidemoto Nakada , Satoshi Itoh , Satoshi Sekiguchi</i> 		
Venue: Ballroom C	Multicore Workshop (Session 1) Chair - Shujia Zhou, NASA, USA Judy Qiu, Indiana University, USA Ken Hawick, Massey University, New Zealand	
<ul style="list-style-type: none"> • Programming challenges for the implementation of numerical quadrature in atomic physics on FPGA and GPU accelerators submission information <i>Charles Gillan</i> • Asynchronous Communication Schemes for Finite Difference Methods on Multiple GPUs submission information <i>Daniel Playne</i> • Solving k-Nearest Neighbor Problem on Multiple Graphics Processors submission information <i>Kimikazu Kato</i> • Cooperative Multitasking for GPU-Accelerated Grid Systems submission information <i>Fumihiko Ino</i> 		
1:00 PM – 2:00 PM	LUNCH	
2:00 PM – 3:30 PM	Parallel Sessions – Regular Papers (R2A & R2B)	MultiCore
Venue: Ballroom A	Session R2A: Applications – Clouds Chair - Bruno Schulze, National Laboratory for Scientific Computing – LNCC, Brazil	
<ul style="list-style-type: none"> • A Map-Reduce System with an Alternative API for Multi-Core Environments <i>Wei Jiang , Vignesh Ravi , Gagan Agrawal</i> • An Analysis of Traces from a Production MapReduce Cluster <i>Soila Pertet , Jiaqi Tan , Rajeev Gandhi , Priya Narasimhan</i> • An Effective Architecture for Automated Appliance Management System Applying Ontology-Based Cloud Discovery <i>Amir Vahid Dastjerdi , Sayed Gholam Hassan Tabatabaei , Rajkumar Buyya</i> 		
Venue: Ballroom B	Session R2B: Middleware/Runtime – Program Optimization and Scheduling Chair - Xian-He Sun, Illinois Institute of Technology, USA	
<ul style="list-style-type: none"> • Region-Based Prefetch Techniques for Software Distributed Shared Memory Systems 		

<p><i>Jie Cai , Peter Strazdins , Alistair Rendell</i></p> <ul style="list-style-type: none"> Granularity-Aware Work-Stealing for Computational Grids <i>Vladimir Janjic , Kevin Hammond</i> SAGA BigJob: An Extensible and Interoperable Pilot-Job Abstraction for Distributed Applications and Systems <i>Shantenu Jha , Andre Luckow</i> 		
Venue: Ballroom C	<p>Multicore Workshop (Session 2) Chair - Shujia Zhou, NASA, USA Judy Qiu, Indiana University, USA Ken Hawick, Massey University, New Zealand</p>	
<ul style="list-style-type: none"> Multi - FFT Vectorization for the Cell Multicore Processor submission information <i>Jacob Barhen</i> High Resolution Program Flow Visualization of Hardware Accelerated Hybrid Multi-Core Applications submission information <i>Daniel Hackenberg</i> Running the NIM Next-Generation Weather Model on GPUs submission information <i>Mark Govett</i> 		
3: 30 PM – 4.00 PM	Afternoon Tea	
4:00 PM – 6:00 PM	Parallel Sessions – Regular Papers (R3A & R3B)	MultiCore
Venue: Ballroom A	<p>Session R3A: Programming Models and Systems – HPC and Accelerators Chair - Pavan Balaji, Argonne National Laboratory, USA</p>	
<ul style="list-style-type: none"> Remote Process Execution and Remote File I/O for Heterogeneous Processors in Cluster Systems <i>Masaaki Shimizu , Akinori Yonezawa</i> An Adaptive Data Prefetcher for High-Performance Processors <i>Yong Chen , Xian-He Sun</i> Designing Accelerator-Based Distributed Systems for High Performance <i>M. Mustafa Rafique , Ali Butt , Dimitrios Nikolopoulos</i> Efficient On-demand Connection Management Mechanisms with PGAS Models on InfiniBand <i>Abhinav Vishnu , Manojkumar Krishnan</i> 		
Venue: Ballroom B	<p>Session R3B: Performance Modeling and Evaluation – Scheduling and Resource Management Chair - Dick Epema, Delft University, Netherland</p>	
<ul style="list-style-type: none"> Efficient Runtime Environment for Coupled Multi-Physics Simulations: Dynamic Resource Allocation and Load-Balancing <i>Shantenu Jha , Soon-Heum Ko</i> An evaluation of the benefits of fine-grained value-based scheduling on general purpose clusters <i>Ruben Van den Bossche , Kurt Vanmechelen , Jan Broeckhove</i> The Effects of Untruthful Bids on User Utilities and Stability in Computing Markets <i>Sergei Shudler , Lior Amar, Amnon Barak, Ahuva Mu'alem</i> FIRE: A File Reunion Based Data Replication Strategy for Data Grids <i>Abdul Abdurrab , Tao Xie</i> 		
Venue: Ballroom C	<p>Multicore Workshop (Session 3) Chair - Shujia Zhou, NASA, USA Judy Qiu, Indiana University, USA Ken Hawick, Massey University, New Zealand</p>	
<ul style="list-style-type: none"> Accelerating Climate and Weather Simulations Through Hybrid Computing submission information <i>Shujia Zhou</i> A Memory-Centric Kernel Framework for Accelerating Short-Range, Interactive Particle Simulation submission information <i>Ian Stewart</i> From Sparse Matrix to Optimal GPU CUDA Sparse Matrix Vector Product Implementation submission information <i>Ahmed El Zein</i> Performance of Windows Multicore Systems on Threading and MPI submission information <i>Judy Qiu</i> 		

Day 3: May 19, 2010

9:00 AM – 10:45 AM	Industry Track – Cloud Panel 9:00 AM – 9:45 AM: Keynote 2 (Industry) “Cloud Computing with Amazon Web Services” – Simone Brunozzi, Amazon.com (AWS) Chair: Dr. Rajkumar Buyya, University of Melbourne 9:45 AM - 10:35 AM: Panel on “Cloud Deployment Trajectories for National Goals” – Panel Chair/Moderator: Craig Lee, Open Grid Forum (OGF) Panel Members: <ul style="list-style-type: none"> • Greg Stone, Chief Technology Officer, Microsoft Australia • Satoshi Matsuoka, TiTech/Global Scientific Information and Computing Center • Rhys Francis, Executive Director, Australian eResearch Infrastructure Council • Manish Parashar, NSF Office of Cyberinfrastructure and Rutgers University 		
10:35 AM – 11:00 AM	Morning Tea		
11:00 AM – 1:00 PM	Parallel Sessions – Regular Papers (R4A & R4B)		Meeting
Venue: Ballroom A	Session R4A: Algorithms – Scheduling and Resource Allocation Chair - Shantenu Jha, LSU, USA/eSI, UK		
<ul style="list-style-type: none"> • SAQA: A Self-Adaptive QoS-Aware Scheduling Algorithm for Real-Time Tasks on Heterogeneous Clusters <i>Xiaomin Zhu</i> • Bandwidth Allocation for Iterative Data-dependent e-Science Applications <i>Eun-Sung Jung , Sanjay Ranka , Sartaj Sahni</i> • A Bi-Criteria Algorithm for Scheduling Parallel Task Graphs on Clusters <i>Frederic Desprez , Frederic Suter</i> • Low-Cost Tuning of Two-Step Algorithms for Scheduling Mixed-Parallel Applications onto Homogeneous Clusters <i>Sascha Hunold</i> 			
Venue: Ballroom B	Session R4B: Middleware/Runtime – Service Management and Workflows Chair – Professor Yanchun Zhang, Victoria University, Australia		
<ul style="list-style-type: none"> • ERGOT: A Semantic-based System for Service Discovery in Distributed Infrastructures <i>Giuseppe Pirro, Paolo Trunfio , Domenico Talia , Paolo Missier , Carole Goble</i> • Towards Autonomic Service Provisioning Systems <i>Michele Mazzucco</i> • WORKEM: Representing and Emulating Distributed Scientific Workflow Execution State <i>Lavanya Ramakrishnan , Dennis Gannon , Beth Plale</i> • Experiments with Memory-to-Memory Coupling for End-to-End Fusion Simulation Workflows <i>Ciprian Docan , Fan Zhang , Manish Parashar , Julian Cummings , Norbert Podhorszki , and Scott Klasky</i> 			
Venue: Ballroom C	Meeting		
1:10 PM – 2:00 PM	LUNCH		
2:00 PM – 3:30 PM	Regular Papers (R5)	Short Papers (S1)	SCALE
Venue: Ballroom A	Session R5: Programming Models and Systems – Streams Chair - Judy Qiu, Indiana University, USA		
<ul style="list-style-type: none"> • Streamflow – Programming Model for Data Streaming in Scientific Workflows <i>Chathura Herath , Beth Plale</i> • Representing eager evaluation in a demand driven model of streams on cloud infrastructure <i>Paul Martinaitis , Andrew Wendelborn</i> • An MPI-Stream Hybrid Programming Model for Computational Clusters <i>Emilio Mancini , Gregory Marsh , Dhableswar Panda</i> 			
Venue: Ballroom B	Session S1: Cloud Computing and Applications Chair - Shikharesh Majumdar, Carleton University, Canada		
<ul style="list-style-type: none"> • On the Origin of Services – Using RIDDl for Description, Composition and Evolution of RESTful Services <i>Juergen Mangler, Erich Schikuta, and Christoph Witzany</i> • A Categorisation of Cloud Computing Business Models 			

Victor Chang

- **Dynamic Resource Pricing on Federated Clouds**
Marian Mihailescu and Yong-Meng Teo
- **Unibus-managed Execution of Scientific Applications on Aggregated Clouds**
Jaroslav Slawinski, Magdalena Slawinska, and Vaidy Sunderam

Venue: Ballroom C	SCALE (Presentation)
3:30 PM – 4:00 PM	Afternoon Tea
4:00 PM – 6:00 PM	Posters & Research Demos (Venue: Foyer) "SCALE" Public demos will be held during this period itself.

Posters:

- **Expanding the Cloud: A component-based architecture to application deployment on the Internet**
Mark Wallis, Frans Henskens and Michael Hannaford
- **Fine-Grained Profiling for Data-Intensive Workflows**
Nan Dun, Kenjiro Taura and Akinori Yonezawa
- **Supporting OFED over Non-InfiniBand SANs**
Devesh Sharma
- **The Lightweight Approach to Use Grid Services with Grid Widgets on Grid WebOS**
Yi Lun Pan and Chang Hsing Wu
- **Energy Efficient Allocation of Virtual Machines in Cloud Data Centers**
Anton Beloglazov and Rajkumar Buyya
- **SciCloud: Scientific Computing on the Cloud**
Satish srirama, Oleg Batrashev and Eero Vainikko
- **Rigel: A Scalable and Lightweight replica selection Service for Replicated Distributed File System**
Yuan Lin, Yang Chen, Guodong Wang and Beixing Deng
- **In Search Of Visualization Metaphors For Planetlab**
Andrew Zaliwski
- **Design and Implementation of an efficient Two-level Scheduler for Cloud Computing Environment**
Jeyarani Rajarathinam and Vasanth Ram Rajarathinam
- **Cluster Computing as an Assembly Process: Coordination with S-Net**
Alex Shafarenko, Clemens Grellck, Frank Penczek and Jukka Julku
- **Dynamic Job-Clustering with Different Computing Priorities for Resource Allocation**
Masnida Hussin, Young Choon Lee and Albert Y. Zomaya
- **Dynamic Auction Mechanism for Cloud Resource Allocation**
Wei-Yu Lin, Guan-Yu Lin and Hung-Yu Wei
- **Improving Grid Systems Reliability with a Failure History Service**
Catalin Leordeanu, Thomas Ropars, Valentin Cristea and Christine Morin
- **Energy-Aware Scheduling using Dynamic Voltage-Frequency Scaling**
Nikzad Babaii Rizvandi
- **Policy-based Management of QoS in Service Aggregations**
Mohan Baruwal Chhetri, Bao Quoc Vo and Ryszard Kowalczyk
- **Feedback-guided Analysis for Resource Requirements in Large Distributed System**
Madhulina Sarkar, Sarbani Roy and Nandini Mukherjee
- **TOPP goes Rapid - The OpenMS Proteomics Pipeline in a Grid-enabled Web Portal**
Sandra Gesing, Jano van Hemert, Jos Koetsier, Andreas Bertsch and Oliver Kohlbacher

Research Demos:

- **Satellite Data Product Generation Using Aneka Cloud**
Raghavendra K, Akilan A , Ravi N, Pramod Kumar K, Geeta varadan
- **XtreemOS, a GRID Operating System**
Yvon J'egou, Christine Morin
- **Digital Elevation Model generation using GP/GPU**
Ramakrishna Reddy V, Pramod Kumar K and Geeta Varadan
- **Presentation of the Grisu framework and some client implementations**
ARCS
- **Enabling Greener Cloud Datacenters with Advanced Virtualization Technology**
Takahiro Hirofuchi, Hidemoto Nakada, Satoshi Itoh, and Satoshi Sekiguchi
- **Using the Any Schedulability Criterion for Matchmaking on Clouds and Grids**
Jose Orlando Melendez, Shikharesh Majumdar, Umar Farooq
- **Parallel Computing with MATLAB® on Amazon Elastic Compute Cloud (EC2)**
Bobby Nedelkovski (The MathWorks Australia Pty Ltd)
- **Browsing Large Scale Cheminformatics Data with Dimension Reduction**
Judy Qiu, Jong Youl Choi, Seung-Hee Bae, Thilina Gunarathne, Geoffrey Fox, Bin Cao, David Wild
- **Service Oriented CARE Resource Broker**
Thamarai Selvi Somasundaram, Balakrishnan Ponnuraman, Kumar Rangasamy, Rajendar Kandan, Kannan Govindarajan, Rajiv Rajaian, Mahendran Ellappan

6:30 PM – 9:30 PM

River Cruise & Conference Banquet

Day 4: May 20, 2010

9:00 AM – 10:00 AM	Keynote 3 – Sky Computing: When Multiple Clouds Become One Professor José A.B. Fortes, University of Florida, USA Chair: Dr. Rajkumar Buyya		
10:00 AM – 10:30 AM	Morning Tea		
10:30 AM – 12:30 PM	Regular Papers (R6)	DocSymp	Short Papers (S2)
Venue: Ballroom A	Session R6: Applications Chair - Satoshi Matsuoka, Tokyo University of Technology, Japan		
<ul style="list-style-type: none"> • High Performance Dimension Reduction and Visualization for Large High-dimensional Data Analysis <i>Jong Youl Choi , Seung-Hee Bae , Xiaohong Qiu , Geoffrey Fox</i> • Exploring the Potential of Using Multiple e-Science Infrastructures with Emerging Open Standards-based e-Health Research Tools <i>Morris Riedel , Bernd Schuller , Michael Rambadt , Shahbaz Memon , Shiraz Memon , Achim Streit , Felix Wolf , Thomas Lippert , Dieter Kranzlmoller , Stefan Zasada , Steven Manos , Peter Coveney</i> • Methodology for Efficient Execution of SPMD Applications on Multicore Environments <i>Ronal Muresano , Dolores Rexachs , Emilio Luque</i> • On-demand Overlay Networks for Large Scientific Data Transfers <i>Lavanya Ramakrishnan , Chin Guok , Keith Jackson , Ezra Kissel , Martin Swany , Deborah Agarwal</i> 			
Venue: Ballroom B	Doctoral Symposium Chairs - Rajiv Ranjan, University of New South Wales, Sydney, Australia Hyunjoo Kim, Rutgers University, USA		
<ul style="list-style-type: none"> • Service Oriented Approach to High Performance Scientific Computing <i>Jaison Mulerikkal</i> • Energy Efficient Resource Management in Virtualized Cloud Data Centers <i>Anton Beloglazov</i> • SLA-Driven Dynamic Resource Management for Multi-tier Web Applications in a Cloud <i>Waheed Iqbal</i> • On Economic and Computational-efficient Resource Pricing in Large Distributed Systems <i>Marian Mihailescu</i> • A Capabilities-Aware Programming Model for Asymmetric High-End Systems <i>M. Mustafa Rafique</i> 			
Venue: Ballroom C	Session S2: Grid and e-Science Applications Chair - Dr.S.Thamarai Selvi, Anna University Chennai		
<ul style="list-style-type: none"> • File-Access Characteristics of Data-intensive Workflow Applications <i>Takeshi Shibata, Kenjiro Taura</i> • Overdimensioning for Consistent Performance in Grids <i>Nezih Yigitbasi, Dick Epema</i> • Topology Aggregation for e-Science Networks <i>Eun-Sung Jung, Sanjay Ranka, and Sartaj Sahni</i> • Handling Recoverable Temporal Violations in Scientific Workflow Systems: A Workflow Rescheduling Based Strategy <i>Xiao Liu, Jinjun Chen, Zhangjun Wu, Zhiwei Ni, Dong Yuan, and Yun Yang</i> • A Fair Distributed Scheduler for Bag-of-tasks Applications on Desktop Grids <i>Javier Celaya, Loris Marchal</i> • A Heuristic Query Optimization Approach for Heterogeneous Environments <i>Peter Beran, Werner Mach, Ralph Vigne, and Erich Schikuta</i> 			
12:30 PM – 1:30 PM	LUNCH		
1:30 PM – 3:00 PM	Parallel Sessions – Regular Papers (R7A & R7B)	Short Papers (S3)	
Venue: Ballroom A	Session R7A: Algorithms and Applications – Energy Chair - Adam Barker, University of Melbourne		
<ul style="list-style-type: none"> • Towards Energy Aware Scheduling for Precedence Constrained Parallel Tasks in a Cluster <i>Lizhe Wang , Jai Dayal</i> • Runtime Energy Adaptation with Low-impact Instrumented code in Power-scalable Cluster System <i>Hideaki Kimura , Takayuki Imada , Mitsuhsa Sato</i> • Energy Minimization of DVFS-enabled processors in HPCS by Linear Combinations of Processor Frequencies <i>Nikzad Babai</i> 			

Venue: Ballroom B	Session R7B: Performance Modeling and Evaluation – Tracing and Communication
<ul style="list-style-type: none"> • The Failure Trace Archive: Enabling Comparative Analysis of Failures in Diverse Distributed Systems <i>Derrick Kondo , Bahman Javadi , Alexandru Iosup , Dick Epema</i> • Scalable communication trace compression <i>Sriram Krishnamoorthy , Khushbu Agarwal</i> • FaReS: Fair Resource Scheduling for VMM-bypass Infiniband Devices <i>Adit Ranadive , Ada Gavrilovska , Karsten Schwan</i> 	
Venue: Ballroom C	Session S3: Data Management in Grids Chair - Laurent Lefevre, INRIA
<ul style="list-style-type: none"> • Planning Large Data Transfers in Institutional Grids <i>Fatiha Bouabache, Thomas Herault, Sylvain Peyronnet, and Franck Cappello</i> • Framework for Efficient Indexing and Searching of Scientific Datasets <i>Chaitali Gupta, Madhusudhan Govindaraju</i> • High Performance Data Transfer in Grid Environment Using GridFTP over InfiniBand <i>Hari Subramoni, Ping Lai, Rajkumar Kettimuthu, Dhableswar Panda</i> • Data Injection at Execution Time in Grid environment using Dynamic Data Driven Application System for Wildland Fire Spread Prediction <i>Roque Rodriguez</i> 	
3:00 PM – 3:30 PM	Afternoon Tea
3:30 PM – 5:30 PM	Parallel Sessions – Regular Papers (R8A & R8B)
Venue: Ballroom A	Session R8A: Algorithms – Self-Organizing and Peer-to-Peer Systems Chair - Kris Bubendorfer, Victoria University Wellington, Australia
<ul style="list-style-type: none"> • A Proximity-Based Self-Organizing Framework for Service Composition and Discovery <i>Carlo Mastroianni , Agostino Forestiero , Giandomenico Spezzano , Giuseppe Papuzzo</i> • Dynamic TTL-Based Search In Unstructured Peer-to-Peer Networks <i>Imen Filali , Fabrice Huet</i> • Enhanced Paxos Commit for Transactions on DHTs <i>Florian Schintke , Alexander Reinefeld , Seif Haridi , Thorsten Schuett</i> • Cache Performance Optimization for Processing XML-based Application Data on Multi-core Processors <i>Rajdeep Bhowmik , Madhusudhan Govindaraju</i> 	
Venue: Ballroom B	Session R8B: Performance Modeling and Evaluation – Workload Modeling and Prediction Chair - Alexandru Iosup, Delft University of Technology, the Netherlands
<ul style="list-style-type: none"> • A Realistic Integrated Model of Parallel System Workloads <i>Minh Tran , Lex Wolters , Dick Epema</i> • Discovering Linear Models of Grid Load <i>Tamas Elteto , Cecile Germain-Renaud , Pascal Bondon</i> • Identification, modelling and prediction of non-periodic bursts in workloads <i>Mario Lassnig , Thomas Fahringer , Vincent Garonne , Angelos Molfetas , Miguel Branco</i> • On the use of machine learning to predict the time and resources consumed by applications <i>Andrea Matsunaga , Jose Fortes</i> 	
5:30 PM – 6:00 PM	Valedictory & Closing

The Third IEEE International Scalable Computing Challenge (SCALE 2010)

1. WebPIE: a Web-scale Parallel Inference Engine

Jacopo Urbani, Spyros Kotoulas, Jason Maassen, Niels Drost, Frank Seinstra, Frank van Harmelen, Henri Bal
Department of Computer Science, Vrije Universiteit Amsterdam

2. Computing at the Petabyte scale with the WLCG

Andrea Sciaba, CERN

3. Scaling-out CloudBLAST: Combining Technologies to BLAST on the Sky

Andréa Matsunaga, Maurício Tsugawa and José Fortes
Advanced Computing and Information Systems Laboratory
Department of Electrical and Computer Engineering, University of Florida

4. A System for Efficient Execution of Bags of Tasks in Multiple Grids and Clouds Mark Silberstein,

Artyom Sharov, Dan Geiger, Assaf Schuster
Technion—Israel Institute of Technology

5. Scalable Personal Health Monitoring based on Mobile and Cloud Computing

Suraj Pandey, Sheng Niu, Ahsan Khandoker, and Rajkumar Buyya
Clouds Lab, University of Melbourne.