Proceedings

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

CCGrid 2010

Proceedings

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

Melbourne, Australia
17–20 May 2010

Editors
Manish Parashar and Rajkumar Buyya

Sponsored/Supported by
IEEE Computer Society
Association for Computing Machinery (ACM)
IEEE Technical Committee on Scalable Computing
Cloud Computing and Distributed Systems (CLOUDS) Laboratory
ARC Research Network on Intelligent Sensors,
Sensor Networks and Information Processing (ISSNIP)
The University of Melbourne, Australia
Manjrasoft Pty Ltd, Melbourne, Australia
NICTA Victoria Lab, Australia
Amazon Web Services, USA
NSF Center for Autonomic Computing
Rutgers, the State University of New Jersey, USA
Victorian State Government, Australia

Los Alamitos, California

Washington • Tokyo

IEEE Computer Society Order Number P4039
BMS Part Number CFP10276-CDR
ISBN 978-0-7695-4039-9

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1 800 272 6657
Fax: + 1 714 821 4641
http://computer.org/cspress
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1 732 981 0060
Fax: + 1 732 981 9667
http://shop.ieee.org/store/
customer-service@ieee.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81 3 3408 3118
Fax: + 81 3 3408 3553
tokyo.ofc@computer.org

Individual paper REPRINTS may be ordered at: <reprints@computer.org>

Editorial production by Lisa O'Conner
Cover art production by Joe Daigle

IEEE Computer Society
***Conference Publishing Services* (CPS)**
http://www.computer.org/cps

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

CCGrid 2010

Table of Contents

Session R1B: Middleware/Runtime—Resource Management

Session R2A: Applications—Clouds

Session R2B: Middleware/Runtime—Program Optimization and Scheduling

Session R3A: Programming Models and Systems—HPC and Accelerators

Session R3B: Performance Modeling and Evaluation—Scheduling and Resource Management

Session R4A: Algorithms—Scheduling and Resource Allocation

Session R4B: Middleware/Runtime—Service Management and Workflows

Session R5: Programming Models and Systems—Streams

Session R6: Applications

Session R7A: Algorithms and Applications—Energy

Session R7B: Performance Modeling and Evaluation—Tracing and Communication

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

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

SHORT PAPERS

Session S1: Cloud Computing and Applications

Session S2: Grid and E-science Applications

Session S3: Data Management in Grids

POSTER PAPERS

Second International Symposium on Cloud Computing (Cloud 2010)

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

Resiliency in High Performance Computing (Resilience 2010)

Fifth International Workshop on Content Delivery Networks (CDN 2010)

Challenges for the Application of Grids in Healthcare (CCGrid-Health 2010)

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

Doctoral Symposium

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 Gropp 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/

Message from the Program Chair

On behalf of the vice chairs and the program committee, it is my pleasure to welcome you to the 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing being held in Melbourne, Australia, the city where it was founded a decade ago. CCGrid continues to focus on bringing together international researchers, developers, and users and to provide an international forum to present leading research activities and results on a broad range of topics related to Cluster, Cloud and Grid platforms, paradigms and their applications.

By the standards of the field, CCGrid is a well established conference. Nonetheless, as the research landscape has changed, we have evolved the focus and format of the conference to match it. The most obvious change is the explicit recognition of Cloud computing in our title. Cloud computing is emerging as a dominant computing paradigm alongside Cluster and Grid computing, and as a result, it is fitting that the scope of CCGrid be expanded to accommodate research in this area. Additionally, the CCGrid 2010 call had a special focus on three important and immediate issues that are significantly influencing all aspects of Cluster, Cloud and Grid computing: *Economics*, *Environment* and *Autonomics*.

We have, of course, retained – and indeed strengthened – CCGrid's focus on presenting strong papers on exciting science. This year we accepted 51 full papers from an exceptionally strong field of 219 reviewed full paper submissions, resulting in an acceptance rate of ~23%. For this I would like to acknowledge the dedication and tremendous efforts of the vice program chairs: **Geoffrey Fox**, Indiana University, USA (Applications and Experience), **David Bader**, Georgia Tech, USA (Algorithms), **Carlos Varela**, RPI, USA (Programming Models and Systems), **Thomas Fahringer**, University of Innsbruck, Austria (Middleware/Runtime), and **Dick Epema**, Delft University, Netherland (Performance Modeling and Evaluation). I would also like to thank the program committee and reviewers, who gave their time and expertise as we handled the large volume of submissions.

Several events complement and strengthen the main CCGrid technical program. We are delighted to welcome two excellent and internationally renowned keynote speakers – the first keynote will be given by Professor William Gropp, recipient of the IEEE Medal of Excellence in Scalable Computing, and the second keynote will be given by Professor Jose Fortes, a world leader in research in Cloud Computing. CCGrid will also host the 3rd IEEE TCSC International Scalable Computing Challenge (SCALE 2010). This year the challenge will be organized by Shantenu Jha and Daniel Katz, and will feature live demonstrations showcasing real-world problem solving using computing that scales.

The conference will also include a dedicated industry track on Cloud computing, a poster session (coordinated by Pavan Balaji) presenting the latest breakthroughs in Cluster, Grid and Cloud technologies, multiple satellite workshops (coordinated by Omer Rana) addressing important related areas of research, tutorials (coordinated by Sushil Prasad), as well as the 3rd IEEE TCSC Doctoral Symposium (coordinated by Rajiv Ranjan and Hyunjoo Kim).

An event such as CCGrid is not possible without the coordinated efforts of multiple dedicated individuals who volunteer their time and expertise. I would like to acknowledge the leadership and untiring efforts of the conference general chair, Rajkumar Buyya, and the guidance provided by the steering committee. The publicity chairs (Cho-Li Wang and Masoud Sadjadi), local organizing chairs (James Broberg and Mukaddim Pathan), and cyber chair (Suraj Pandey) also deserve special mention.

Most of all, I am grateful to the CCGrid community for providing high-quality papers and presentations, and for showing how dynamic the field is becoming.

I do hope that you will find this program interesting and thought provoking, and that CCGrid 2010 will provide you with a valuable opportunity to share ideas with researchers and practitioners from academia and industry from around the world.

Cheers!

Professor Manish Parashar,

National Science Foundation & Rutgers, The State University of New Jersey, USA

Program Chair, CCGrid 2010

CCGrid 2010 – Message from the Workshops Chair

A number of workshop proposals were received for CCGrid 2010, some were successful workshops hosted along previous CCGrid events, whereas others demonstrated the emergence of particular research areas in Grid, Cluster and Cloud computing over recent years. Workshops continue to play an important role in the overall CCGrid programme, as they help identify hot topics of research, stimulate research in emerging areas of interest to the community, and enable participants to discuss and establish collaborative links. As work being presented within a workshop is often at an early stage of maturity, but indicative of significant potential, such events should also allow participants to engage with the authors and generate discussion. If you are attending CCGrid this year, please try to participate in at least one workshop and present your views. It would be wonderful to see an enthusiastic and active community being represented at the workshops, helping identify research directions and challenges for subsequent years.

To ensure that good quality contributions were included, stringent acceptance criteria were adhered to by all workshop organizers. Six (out of the 8 submitted) workshop proposals were accepted this year. The choice was based on the strength of the proposals, the quality of submissions, the experience of the organizers and importance of these emerging areas to Grid, Cluster and Cloud computing research. The following workshops have been accepted for CCGrid 2010:

- **5th International Workshop on Content Delivery Networks (CDN 2010)**
  This workshop focuses on “Content Delivery in the Cloud”, with an emphasis on research trends and results in terms of design, architecture, and applications for content and service delivery in the Internet and Clouds; optimization for Cloud-based content delivery; and performance measurement methodologies.

- **4th Workshop on Desktop Grids and Volunteer Computing Systems (PCGrid 2010)**
  Desktop grids and volunteer computing systems utilize the free resources available in Intranet and Internet environments for supporting large-scale computation and storage. The purpose of this workshop is to provide a forum for discussing recent advances and identifying open issues for the development of scalable, fault-tolerant, and secure desktop Grid systems. This year's PCGrid workshop has special emphasis on the interaction of Clouds and desktop Grids. This workshop has been organized in collaboration with the highly successful European CoreGRID Research Group working in this area.

- **2nd International Symposium on Cloud Computing (Cloud 2010)**
  This workshop has a special theme of “Applied Cloud Technologies for Business and Consumer Applications” and is targeted at researchers and practitioners involved in Cloud computing technologies in addition to those harnessing Clouds for their applications in various fields to maximise performance, minimise cost and improve the scale of their endeavours.

- **Resiliency in High Performance Computing (Resilience 2010)**
  This workshop is based on the premise that High Performance Computing (HPC) carried out over Grid, Cluster and Clouds must utilize large numbers of resources and hence effective HPC in any of these paradigms must address the issue of resiliency at large-scale. The substantial growth in system scale, and the resulting increase in component count, poses a challenge for HPC system and application software with respect to fault tolerance and resilience.

- **Challenges for the Application of Grids in Healthcare (CCGrid-Health 2010)**
  With the increasing interest in the use of Grid-based technologies in Healthcare (and the very active “HealthGrid” community), this workshop aims to encourage discussion about the challenges for the construction and deployment of Grids in Healthcare, offering a contact opportunity between HealthGrid application developers and contributing to reduce the gap between the research and production Grid communities. The workshop has the goals of obtaining an overview of ongoing efforts

in health-related Grid applications; obtaining an overview of challenges (technologies, achievements, gaps, roadblocks); and identifying common requirements to encourage collaboration between Health and Computing Sciences.

- **CCGrid-Multicore 2010 (Frontiers of GPU, Multi- and Many-Core Systems)** There has recently been an increasing use of multi- and many-core microprocessors within Clusters, Clouds and Grids. Both conventional multi- and many-core processors, such as Intel Nehalem and IBM Power7 processors, and unconventional many-core processors, such as NVIDIA Tesla and AMD FireStream GPUs. The aim of this workshop is to discuss issues such as: how to optimize applications for conventional multi- and many-core processors? How does one re-engineer applications to take advantage of the tremendous computing power of a GPU in a reasonable cost-benefit ratio?, and What are effective ways of using GPUs as accelerators?

The workshops also include invited and keynote speakers – from active researchers such as Geoffrey Fox, David Abramson and others. It is also interesting to see a combination of university, industry and national laboratory participation on the programme and organizing committees of the workshops mentioned above.

I would like to thank Rajkumar Buyya and his team (James Broberg, Suraj Pandey and Mukaddim Pathan) with their near zero-latency email responses (!) to my queries, and all the workshop organizers. It has been a pleasure to work with all of you this year.

Enjoy the workshops this year – and be sure to participate actively!

Professor Omer F. Rana, *Cardiff University, UK*

CCGrid 2010 Workshops Chair

Organising Committee

Chairs and Committees

General Chair
Rajkumar Buyya, *University of Melbourne and Manjrasoft Pty Ltd, Australia*

Program Committee Chair
Manish Parashar, *Rutgers University, USA*

Vice Chairs
Applications and Experiences
Geoffrey Fox, *Indiana University, USA*
Algorithms
David Bader, *Georgia Tech, USA*
Programming Models and Systems
Carlos Varela, *RPI, USA*
Middleware/Runtime
Thomas Fahringer, *University of Innsbruck, Austria*
Performance Modeling and Evaluation
Dick Epema, *Delft University, Netherland*

Workshop Chair
Omer F. Rana, *Welsh eScience Center and Cardiff University, UK*

Industry Track Chair
Geng Lin, *Cisco Systems, USA*

Posters Chair
Pavan Balaji, *Argonne National Laboratory, USA*

Publicity Chair
Cho-Li Wang, *University of Hong Kong, China*
Masoud Sadjadi, *FIU, USA*

Tutorials Chair
Sushil K. Prasad, *Georgia State University, USA*

Research Demos/Competitions Chair (SCALE Challenge)
Shantenu Jha, *LSU, USA/eSI, UK*
Daniel S. Katz, *University of Chicago/Argonne National Laboratory, USA*

Local Organising Co-Chairs
Mukaddim Pathan, *University of Melbourne, Australia*
James Broberg, *University of Melbourne, Australia*
Suraj Pandey, *University of Melbourne*, *Australia*

Local Organising Committee Members

Dr. Jinjun Chen, *Swinburne University of Technology*
Suraj Pandey, *University of Melbourne*, *Australia*
Anton Beloglazov, *University of Melbourne*, *Australia*
William Voorsluys, *University of Melbourne*, *Australia*

Cyber Chair

Suraj Pandey, *University of Melbourne, Australia*

Chairs of Associated CCGrid 2010 Workshops:

Third IEEE TCSC Doctoral Symposium
Co-Chairs:
Rajiv Ranjan, *University of New South Wales, Sydney, Australia*
Hyunjoo Kim, *Rutgers University, USA*

Fifth International Workshop on Content Delivery Networks (CDN 2010)
Co-Chairs:
Giancarlo Fortino, *University of Calabria, Italy*
George Pallis, *University of Cyprus, Cyprus*
Mukaddim Pathan, *University of Melbourne, Australia*
Swami Sivasubramanian, *Amazon.com, Inc., USA*

Fourth Workshop on Desktop Grids and Volunteer Computing Systems (PCGrid 2010)
Co-Chairs:
Gilles Fedak, *INRIA, France*
Derrick Kondo, *INRIA, France*
Bahman Javadi, *INRIA, France*

Second International Symposium on Cloud Computing (Cloud 2010)
Co-Chairs:
James Broberg, *The University of Melbourne, Australia*
Bruno Schulze, *National Laboratory for Scientific Computing, Brazil*
Rajkumar Buyya, *The University of Melbourne, Australia*

Workshop on Resiliency in High Performance Computing (Resilience 2010)
Co-Chairs:
Stephen L. Scott, *Oak Ridge National Laboratory, USA*
Chokchai (Box) Leangsuksun, *Louisiana Tech University, USA*
Christian Engelmann, *Oak Ridge National Laboratory, USA*

Workshop on Challenges for the Application of Grids in Healthcare (CCGrid-Health 2010)
Co-Chairs:
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*

Workshop on Frontiers of GPU, Multi- and Many-Core Systems (CCGrid-Multicore 2010)
Co-Chairs:

Shujia Zhou, *NASA, USA*
Judy Qiu, *Indiana University, USA*
Ken Hawick, *Massey University, New Zealand*

Program Committee Members

Kenneth Hawick, *Massey University – Albany, New Zealand*
Cécile Germain-Renaud, *Université Paris-Sud, France*
Carlo Mastroianni, *ICAR-CNR, Italy*
Bo Hong, *Georgia Institute of Technology, USA*
Kamesh Madduri, *Lawrence Berkeley National Laboratory, USA*
Tiffani Williams, *Texas A&M University, USA*
Jakub Kurzak, *University of Tennessee at Knoxville, USA*
Albert Zomaya, *University of Sydney, Australia*
Jon Berry, *Sandia National Laboratory, USA*
Zhihui Du, *Tsinghua University, China*
Viktor Prasanna, *University of Southern California, USA*
Gul Agha, *University of Illinois, Urbana-Champaign, USA*
Paolo Ciancarini, *University of Bologna, Italy*
John Field, *IBM T.J. Watson Research Lab, USA*
Peter van Roy, *Catholic University of Louvain, Belgium*
David Anderson, *University of California, Berkeley, USA*
Marty Humphrey, *University of Virginia, USA*
Henri Bal, *Vrije Universiteit, Amsterdam, the Netherlands*
Franck Cappello, *LRI/INRIA, France*
Christophe Cérin, *University of Paris XII, France*
Seif Haridi, *KTH Stockholm, Sweden*
Thilo Kielmann, *Vrije Universiteit, Amsterdam, the Netherlands*
Michela Taufer, *University of Delaware, USA*
Henri Casanova, *University of Hawaii, USA*
Kenjiro Taura, *University of Tokyo, Japan*
Radu Prodan, *University of Innsbruck, Austria*
Matei Ripeanu, *University of British Columbia, Canada*
Hong-Linh Truong, *Technical University of Vienna, Austria*
Achim Streit, *Jülich Supercomputing Centre, Germany*
Ramin Yahyapour, *University of Dortmund, Germany*
Yutaka Isikawa, *University of Tokyo, Japan*
Ewa Deelman, *Universiyy of Southern California, USA*
Tevfik Kosar, *Louisiana State University, USA*
Erwin Laure, *Royal Institute of Technology (KTH), Sweden*
Ian Taylor, *Cardiff University, United Kingdom*
David Abramson, *Monash University, Australia*
Frédéric Desprez, *LIP/INRIA, France*
Thomas Herault, *INRIA, France*
Alexandru Iosup, *Delft University of Technology, the Netherlands*
Emmanuel Jeannot, *LORIA/INRIA, France*
Derrick Kondo, *INRIA*, *Grenoble, France*
Tiejian Luo, *Graduate University of the Chinese Academy of Sciences, China*
Martin Swany, *University of Delaware, USA*
Douglas Thain, *Notre-Dame University, USA*
Jon Weissman, *University of Minnesota, USA*
Lex Wolters, *Leiden University, the Netherlands*
Daniel S. Katz, *Argonne National Laboratory, USA*
Omer Rana, *Cardiff University, United Kingdom*

CCGrid 2010 Sponsors