

IEEE/ACM CCGrid 2014

14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing

May 26th – 29th, 2014 – Chicago USA

<http://datasys.cs.iit.edu/events/CCGrid2014/>

MESSAGE FROM THE GENERAL CHAIRS

It is our great pleasure to welcome you to the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2014) and to the great global city of Chicago. CCGrid is a forum for all distributed computing technologies and for all technology stakeholders. The inaugural CCGrid conference was held in Brisbane, Australia, in 2001. Three years later, in 2004, it was held in Chicago. Ten years have passed since then. In 2004, Chicago was a leader in Grid computing. Today, it is a center of Cloud technology innovations. CCGrid 2014 is just like the city of Chicago, having naturally blended and mixed with the historical glories and modern stallers into its unique identity. Come and join us to meet preeminent computer scientists, leading creators of Grid and Cluster technologies, and new shining stars, the significant contributors of today's Cloud and big data management advances. We promise you a memorable experience unequalled at other conferences.

Late May is the perfect time to visit Chicago. Flowers are booming, fountains are sprinkling, and sail boats are in the bay; but most of the summer tours are not in yet. The conference hotel, the Hyatt Regency Chicago, is in an ideal location. It is one block off The Magnificent Mile and near Millennium Park, Lake Michigan, and the Chicago River. Even if you are not interested in shopping, simply walking along Michigan Ave or the Chicago River, in the day time or in the evening, is a pleasure. Plus the special hotel rate for CCGrid attendees is very attractive. The conference banquet will be held on a lake cruise in the evening of Wednesday. That night Chicago will have a half hour long firework show on the lake. Of course, the most exciting part is the outstanding programs offered by CCGrid 2014. From Keynotes to Technical Papers, CCGrid14 has first class programs in all categories. We heartily appreciate the committee members and volunteers who have put the wonderful programs together.

CCGrid 2014 features keynote talks, tutorials, workshops, poster sessions and demos, competition, student travel awards, a panel, as well technical papers. We are pleased to have this year's prestigious IEEE Medal for Excellence in Scalable Computing award winner, Professor Yves Robert at Ecole Normale Supérieure de Lyon, as a keynote speaker of CCGrid 2014. Professor Robert is an authority in algorithm design and analysis.

CCGrid 2014 received 302 paper submissions from 40 countries. After administrative filtering, 283 papers received full reviews. In total, 1089 reviews were conducted and 54 papers were accepted, with an acceptance rate of 19% (54/283). There are eight workshops on Monday, May 26, and five concurrent tutorials in the afternoon of Thursday,

IEEE/ACM CCGRID 2014

14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing
May 26th - 29th, 2014 – Chicago USA



May 29. Eight posters and two demos are selected for the poster session in the evening of Tuesday, May 27, and ten papers are accepted for the Doctoral Symposium program on Wednesday, May 28. Dr. Farnam Jahanian, the National Science Foundation Assistant Director for the Computer and Information Science and Engineering (CISE), will deliver the opening keynote. The keynote speaker of the second day, Prof. Ion Stoica, is a professor of Computer Science at University of California at Berkeley, and is known by his current research projects, Mesos and Spark. The IEEE SCALE challenge competition will be held live, with finalists judged by their demonstrations on Tuesday and the winner announced on Thursday. CCGrid 2014 has received generous sponsorship from the U.S. National Science Foundation and the IEEE Technical Committee on Scalable Computing to assist 17 students to attend this conference. A greeting and round-table session are also arranged for the student awardees.

The success of CCGrid 2014 is due to the dedicated efforts and high standards of numerous international volunteers. Our long thank you list starts with the two excellent Program Chairs: Kirk W. Cameron and Dimitrios S. Nikolopoulos, and the Program Committee. Special thanks to the Program Committee Area Chairs who braved the coldest Chicago day in decades to run the program committee meeting on January 27, 2014. We thank Workshops Co-Chairs Zhiling Lan and Matei Ripeanu, the chairs and PC committees of the various workshops, and our Publicity Chairs for getting the word out about the conference. We thank Tutorials Co-Chairs Kate Keahey and Radu Prodan, Poster and Research Demo Co-Chairs Borja Sotomayor and Hui Jin, Doctoral Symposium Chair Judy Qiu, the Student Awards Chair Yong Chen, and the SCALE Challenge Coordinator Douglas Thain. The Cyber Co-chairs, Ge Rong and Wei Tang, did a wonderful job with the conference website. Dr. Pavan Balaji, the Proceedings Chair, ensured the publication of the conference proceedings. We are especially grateful to the Local Organizing Chairs, Ioan Raicu and Kyle Chard, who did a tremendous job on innumerable tasks, from identifying the hotel to negotiating the price of the banquet. Thanks are also due to our sponsors, namely, IEEE, ACM, TCSC, and the organizational supporters at Illinois Institute of Technology and the University of Chicago. Ultimately, however, the success of the conference will be judged by the attendees' experience. We hope that the conference will provide you with a valuable opportunity to share ideas, communicate, learn, and network.

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



Ian Foster

University of Chicago
Argonne National Laboratory



Xian-He Sun

Illinois Institute of Technology



ORGANIZING COMMITTEE

Honorary General Chair

- Daniel A. Reed, University of Iowa, USA

General Chairs

- Xian-He Sun, Illinois Institute of Technology, USA
- Ian T. Foster, University of Chicago & Argonne National Laboratory, USA

Program Committee Chairs

- Kirk W. Cameron, Virginia Tech, USA
- Dimitrios S. Nikolopoulos, Queen's University of Belfast, UK

Program Committee Area Chairs

- **APPLICATIONS**
 - Esmond G. Ng, Lawrence Berkeley National Laboratory, USA
- **ARCHITECTURE & ACCELERATORS**
 - Suren Byna, Lawrence Berkeley National Laboratory, USA
- **AUTONOMIC COMPUTING AND CYBERINFRASTRUCTURE**
 - Vladimir Getov, University of Westminster, UK
 - Xiaobo Zhou, University of Colorado, USA
- **CLOUD COMPUTING**
 - Ali R. Butt, Virginia Tech, USA
 - Xiaosong Ma, Qatar Computing Research Institute, Qatar; North Carolina State University, USA
- **MODELING AND EVALUATION OF PERFORMANCE AND ENERGY**
 - David Lowenthal, University of Arizona, USA
- **PROGRAMMING MODELS, SYSTEMS, AND FAULT-TOLERANT COMPUTING**
 - Kasidit Chanchio, Thammasat University, Thailand
 - Christian Engelmann, Oak Ridge National Laboratory, USA
- **SCHEDULING AND RESOURCE MANAGEMENT**
 - Matthew Grove, Rackspace, USA
 - Rizos Sakellariou, University of Manchester, UK
- **STORAGE AND I/O SYSTEMS**
 - Dan Feng, Huazhong University of Science and Technology, China

Workshops Co-Chairs

- Zhiling Lan, Illinois Institute of Technology, USA
- Matei Ripeanu, University of British Columbia, Canada

Tutorials Co-Chairs

- Kate Keahey, University of Chicago & Argonne National Laboratory, USA
- Radu Prodan, University of Innsbruck, Austria

Doctoral Symposium Chair

- Judy Qiu, Indiana University, USA

Poster and Research Demo Chairs

- Borja Sotomayor, University of Chicago, USA
- Hui Jin, Oracle, USA

SCALE Challenge Coordinator

- Doug Thain, Notre Dame, USA

Student Award Chair

- Yong Chen, Texas Tech University, USA

Publicity Chairs

- Bruno Schulze, LNCC, Brazil
- Cho-Li Wang, The University of Hong Kong, China
- Darren J. Kerbyson, Pacific Northwest National Laboratory, USA
- Toni Cortes UPC/BSC, Spain

Cyber Co-Chairs

- Rong Ge, Marquette University, USA
- Wei Tang, Argonne National Laboratory, USA



Proceedings Chair

- Pavan Balaji, Argonne National Laboratory, USA

Local Organizing Committee

- Ioan Raicu, Illinois Institute of Technology, USA
- Kyle Chard, University of Chicago, USA

Finance Chair

- Dawn DeBartolo, Illinois Institute of Technology, USA

Steering Committee

- Rajkumar Buyya, University of Melbourne, Australia (*Chair*)
- Craig Lee, The Aerospace Corporation, USA (*Co-Chair*)
- Henri Bal, Vrije University, The Netherlands
- Pavan Balaji, Argonne National Laboratory, USA
- Franck Capello, University of Paris-Sud, France
- Ewa Deelman, University of Southern California/ISI, USA
- Jack Dongarra, University of Tennessee & ORNL, USA
- Dick Epema, Technical University of Delft, The Netherlands
- Thomas Fahringer, University of Innsbruck, Austria
- Ian Foster, University of Chicago, USA
- Wolfgang Gentzsch, DEISA, Germany
- Hai Jin, Huazhong University of Science & Technology, China
- Laurent Lefevre, INRIA, France
- Geng Lin, Dell Inc., USA
- Manish Parashar, Rutgers: The State University of New Jersey, USA
- Shikharesh Majumdar, Carleton University, Canada
- Satoshi Matsuoaka, Tokyo Institute of Technology, Japan
- Omer Rana, Cardiff University, UK
- Paul Roe, Queensland University of Technology, Australia
- Bruno Schulze, LNCC, Brazil
- Xian-He Sun, Illinois Institute of Technology, USA
- Nalini Venkatasubramanian, University of California, USA
- Carlos Varela, Rensselaer Polytechnic Institute, USA

Technical Program Committee

APPLICATIONS

- Hasan Metin Aktulga, Lawrence Berkeley National Laboratory
- Bill Barth, The University of Texas at Austin
- Costas Bekas, IBM Research - Zurich
- Sanjukta Bhowmick, University of Nebraska, Omaha
- Sunita Chandrasekaran, University of Houston
- Yifeng Chen, Peking University
- Olivier Coulaud, INRIA
- Alfredo Cuzzocrea, ICAR-CNR and University of Calabria
- Frederic Desprez, INRIA
- Daniel Martin, Lawrence Berkeley National Laboratory
- Kengo Nakajima, The University of Tokyo
- Hai Ah Nam, Oak Ridge National Laboratory
- Esmond Ng, Lawrence Berkeley National Laboratory
- Dana Petcu, West University of Timisoara
- Judy Qiu, Indiana University
- Ashok Srinivasan, Florida State University
- Gerhard Wellein, Erlangen Regional Computing Center
- Rio Yokota, KAUST
- Rui Zhang, IBM Research - Almaden
- Yunquan Zhang, Institute of Software, Chinese Academy of Sciences

ARCHITECTURE & ACCELERATORS

- Michela Becchi, University of Missouri
- Suren Byna, Lawrence Berkeley National Laboratory
- Srihari Cadambi, NEC Labs America Inc.
- Wenguang Chen, Tsinghua University
- Cong Du, Arista Networks
- Natalie Enright Jerger, University of Toronto
- Michael Ferdman, Stony Brook University
- Dong Li, Oak Ridge National Lab
- Sonia Lopez, Rochester Institute of Technology
- Kamesh Madduri, Pennsylvania State University
- Jiayuan Meng, Argonne National Lab
- Ahmad Samih, Intel Corporation
- Vinod Tipparaju, AMD
- Didem Unat, Lawrence Berkeley National Laboratory
- Sathish Vadhiyar, Indian Institute of Science – Bangalore
- Abhinav Vishnu, Pacific Northwest National Laboratory
- Yili Zheng, Lawrence Berkeley National Laboratory

AUTONOMOUS COMPUTING AND CYBERINFRASTRUCTURE

- Marco Aldinucci, University of Torino
- Francoise Baude, Universite de Nice Sophia-Antipolis
- Alexander Bolotov, University of Westminster
- Marian Bubak, AGH Krakow PL and University of Amsterdam
- Abhishek Chandra, University of Minnesota
- Yong Chen, Texas Tech University
- Yuan Chen, HP Labs
- Marco Danelutto, University of Pisa
- Xiaoning Ding, New Jersey Institute of Technology
- Jose Fortes, University of Florida
- Paraskevi Fragopoulou, FORTH-ICS
- Vladimir Getov, University of Westminster
- Sergei Gorlatch, University of Muenster
- Palden Lama, University of Texas at San Antonio
- Zhiqiang Lin, University of Texas at Dallas
- Cristian Lumezanu, NEC Laboratories
- Ningfang Mi, Northeastern University
- Christine Morin, INRIA
- Jia Rao, University of Colorado, Colorado Springs
- Ian Taylor, Cardiff University and Louisiana State University
- Thomas Weigold, IBM Research
- Timothy Wood, George Washington University
- Weikuan Yu, Auburn University
- Ming Zhao, Florida International University
- Xiaobo Zhou, University of Colorado

CLOUD COMPUTING

- Gagan Agrawal, Ohio State University
- Raouf Boutaba, University of Waterloo
- Ali R. Butt, Virginia Tech
- Claris Castillo, Renaissance Computing Institute
- Zhihui Du, Tsinghua University
- Dick Epema, Delft University of Technology
- Renato Figueiredo, University of Florida
- Geoffrey C. Fox, Indiana University
- Chuntao Hong, Microsoft Research
- Howie Huang, George Washington University
- Dejun Jiang, Institute of Computing Technology Chinese Academy of Sciences
- Peter Kilpatrick, Queen's University Belfast
- Jong Kim, HPC Lab
- Youngjae Kim, Oak Ridge National Laboratory
- Eren Kursun, JPMorgan Chase
- Shan Lu, University of Wisconsin at Madison



- Xiaosong Ma, Qatar Computing Research Institute; North Carolina State University
- Kostas Magoutis, ICS-FORTH
- Gaurav Makkar, NetApp
- M. Mustafa Rafique, IBM Research
- Ioan Raicu, Illinois Institute of Technology
- Lavanya Ramakrishnan, Lawrence Berkeley National Lab
- Kai Sachs, SAP AG
- Prasenjit Sarkar, IBM
- Radu Sion, Stony Brook University
- Yang Song, IBM
- Osamu Tatebe, University of Tsukuba
- Douglas Thain, University of Notre Dame
- Devesh Tiwari, Oak Ridge National Laboratory
- Dimitrios Tsoumakos, Ionian University
- Guanying Wang, Google
- Jon Weissman, University of Minnesota
- Gianluigi Zanetti, CRS4
- Jianfeng Zhan, Institute of Computing Technology - Chinese Academy of Sciences
- Rui Zhang, IBM Research
- Zhe Zhang, IBM Research

MODELING AND EVALUATION OF PERFORMANCE AND ENERGY

- Sadaf Alam, Swiss National Supercomputing Centre
- Dorian Arnold, University of New Mexico
- Frank Bellosa, Karlsruhe Institute of Technology
- Laura Carrington, San Diego Supercomputing Center
- Peter Dinda, Northwestern University
- Todd Gamblin, Lawrence Livermore National Laboratory
- Rong Ge, Marquette University
- Judit Gimenez, Barcelona Supercomputing Center
- John Hartman, University of Arizona
- Lizy John, University of Texas at Austin
- Karen L. Karavanic, Portland State University
- Darren Kerbyson, Pacific Northwest National Lab
- James Laros, Sandia National Labs
- David Lowenthal, University of Arizona
- Naoya Maruyama, RIKEN AICS
- Wolfgang E. Nagel, ZIH, TU Dresden
- Vignesh Ravi, AMD
- Barry Rountree, Lawrence Livermore National Laboratory
- Sameer Shende, University of Oregon
- Valerie Taylor, Texas A&M University
- Hans Vandierendonck, Queen's University Belfast
- Ana Varbanescu, University of Amsterdam
- Yonghong Yan, University of Houston
- Xin Yuan, Florida State University

PROGRAMMING MODELS, SYSTEMS, AND FAULT-TOLERANT COMPUTING

- Vassil Alexandrov, BSC
- Christos Antonopoulos, University of Thessaly
- Pavan Balaji, Argonne National Laboratory
- David E. Bernholdt, Oak Ridge National Laboratory
- Filip Blagojevic, Hitachi Research
- George Bosilca, University of Tennessee
- Patrick Bridges, University of New Mexico
- Greg Bronevetsky, Lawrence Livermore National Laboratory
- Franck Cappello, INRIA and University of Illinois at Urbana Champaign
- Kasidit Chanchio, Thammasat University
- Zizhong Chen, University of California at Riverside
- Andrew Chien, University of Chicago
- Nathan Debardeleben, Los Alamos National Laboratory
- Christian Engelmann, Oak Ridge National Laboratory
- Kurt Ferreira, Sandia National Laboratories
- Ada Gavrilovska, Georgia Institute of Technology
- Cecile Germain, LRI
- Michael Gerndt, Technische Universität München
- William Gropp, University of Illinois at Urbana-Champaign
- Thilo Kielmann, Vrije Universiteit
- Bettina Kramer, Université de Versailles St-Quentin-en-Yvelines (UVSQ)
- Dieter Kranzlmüller, Ludwig-Maximilians-Universität München
- Sriram Krishnamoorthy, Pacific Northwest National Lab
- Hatem Ltaief, KAUST
- Satoshi Matsuoka, Tokyo Institute of Technology
- Celso Mendes, University of Illinois
- Kathryn Mohror, LLNL
- Frank Mueller, NCSU
- Wolfgang E. Nagel, ZIH, TU Dresden
- Bogdan Nicolae, IBM Research (Ireland)
- Dhabaleswar Panda, Ohio State University
- Prapaporn Rattanamrong, Thammasat University
- Alexander Reinefeld, Zuse Institute Berlin
- Rolf Riesen, IBM Research
- Scott Schneider, IBM Research
- Martin Schulz, Lawrence Livermore National Laboratory
- Stephen Scott, Oak Ridge National Laboratory
- Shuaiwen Leon Song, Pacific Northwest National Lab
- Wei Tang, Argonne National Laboratory
- Michela Tauber, University of Delaware
- Putchong Uthayopas, Kasetsart University

- Felix Wolf, German Research School for Simulation Sciences

SCHEDULING AND RESOURCE MANAGEMENT

- Carl Albing, Cray Inc. and University of Reading
- Henri Bal, Vrije Universiteit
- Anne Benoit, ENS Lyon - LIP
- Luiz F. Bittencourt, University of Campinas
- Richard Boakes, University of Portsmouth
- Ivona Brandic, TU Wien
- Julita Corbalan, Barcelona Supercomputing Center
- Ewa Deelman, USC Information Sciences Institute
- Erik Elmroth, Umeå University
- Thomas Fahringer, University of Innsbruck
- Matthew Grove, Rackspace
- Alexandru Iosup, Delft University of Technology
- Emmanuel Jeannot, Inria
- Gideon Juve, USC Information Sciences Institute
- Helen Karatza, Aristotle University of Thessaloniki
- Alexey Lastovetsky, University College Dublin
- Charles Lively, DeVry University
- Maciej Malawski, AGH
- Nandini Mukherjee, Jadavpur University
- Jarek Nabrzyski, University of Notre Dame
- Thomas Naughton, ORNL
- George Pallis, University of Cyprus
- Rizos Sakellariou, University of Manchester
- Uwe Schwiegelshohn, TU Dortmund University
- Domenico Talia, University of Calabria

- Matthew E. Tolentino, Intel Corporation
- Denis Trystram, Grenoble Institute of Technology
- Kurt Vanmechelen, University of Antwerp
- Carlos A. Varela, Rensselaer Polytechnic Institute
- Vladimir Vlassov, Royal Institute of Technology (KTH)
- Ramin Yahyapour, University of Göttingen
- Wolfgang Ziegler, Fraunhofer Institute SCAI

STORAGE AND I/O SYSTEMS

- Andre Brinkmann, Johannes Gutenberg-Universität Mainz
- Toni Cortes, Barcelona Supercomputing Center
- Yafei Dai, Peking University
- Dan Feng, Huazhong University of Science and Technology
- Chris Gniady, University of Arizona
- Xubin He, Virginia Commonwealth University
- Yu Hua, Huazhong University of Science and Technology
- Hong Jiang, University of Nebraska Lincoln
- Xiao Qin, Auburn University
- Josef Spillner
- Fang Wang, Huazhong University of Science and Technology
- Chunxiao Xing, Tsinghua University
- Zhenquan Xu, Wuhan University
- Lu Xu, Institute of Computing Technology - Chinese Academy of Sciences
- Yinliang Yue
- Zhao Zhang, Iowa State University
- Yifeng Zhu, University of Maine
- Zhichun Zhu, University of Illinois at Chicago



PROGRAM OVERVIEW

Monday May 26 (Workshops)					
07:30 - 17:30	Registration				
08:30 - 12:30	ExtremeGreen Room: Acapulco	DPMSS Room: Hong Kong	C4Bio and CCGrid- Health Room: Toronto	DV&UV Room: New Orleans	WACC Room: Regency D
13:30 - 17:30				SCRAMBL Room: New Orleans	C4BIE Room: Regency D
18:00 - 19:30	Student Travel Grant Awardees Meet-and-Greet Session Room: Hong Kong				

Tuesday May 27				
07:30 - 17:30	Registration			
08:30 - 08:45	Conference Opening			
08:45 - 10:00	Keynote: <u>Dr. Farnam Jahanian</u>			
10:00 - 10:30	Break			
10:30 - 12:00	Session 1: <u>Best Papers</u>			
12:00 - 13:30	Lunch			
13:30 - 15:30	Session 2A: <u>MapReduce</u>	Session 2B: <u>Energy and the Environment</u>	Session 2C: <u>Resource Management</u>	Session 2D: <u>SCALE Challenge</u>
15:30 - 16:00	Break			
16:00 - 17:30	Session 3A: <u>Apps I</u>	Session 3B: <u>Message Passing</u>	Session 3C: <u>Elasticity and Adaptation</u>	Session 3D: <u>Poster Oral</u>
18:30 - 20:30	<u>Posters</u> + Conference Reception			

Rooms:


Plenary Sessions	Regency C/D
Track A	Regency C
Track B	Regency D
Track C	Acapulco
Track D	Toronto (Doctoral Forum) and Hong Kong (Student Session)
Panel	Regency C/D
Lunch	Big Bar
Poster Session and Reception	Regency C/D

Wednesday May 28				
08:00 - 17:30	Registration			
08:30 - 08:45	Opening Remarks			
08:45 - 10:00	Keynote: <u>Dr. Ion Stoica</u>			
10:00 - 10:30	Break			
10:30 - 12:00	Session 4A: <u>Big Data</u>	Session 4B: <u>Storage and I/O Systems I</u>	Session 4C: <u>Algorithm</u>	Session 4D: <u>Doctoral Forum - I</u>
12:00 - 13:30	Lunch			
13:30 - 15:30	Session 5A: <u>Scheduling</u>	Session 5B: <u>Virtual Machines</u>	Session 5C: <u>Datacenters and Distributed Computing</u>	Session 5D: <u>Doctoral Forum - II</u>
15:30 - 16:00	Break			
16:00 - 17:30	Panel			
18:30 - 22:00	Banquet Dinner Cruise (Busses Depart Hyatt Foyer at 17:45)			

Thursday May 29				
08:00 - 17:30	Registration			
08:30 - 08:45	Opening Remarks			
08:45 - 10:00	Keynote: <u>Dr. Yves Robert</u>			
10:00 - 10:30	Break			
10:30 - 12:00	Session 6A: <u>Apps II</u>	Session 6B: <u>Architecture</u>	Session 6C: <u>Storage and I/O Systems II</u>	Session 6D: <u>Student Travel Grant Awardees Round-Table</u>
12:00 - 12:30	Conference Closing with Award Ceremony			
12:30-13:30	Lunch			
13:30 - 17:30	<u>Tutorial: Autonomic Clouds</u>	<u>Tutorial: Eco4Cloud</u>	<u>Tutorial: Accelerating Big Data Processing</u>	<u>Tutorial: Globus</u> <u>Tutorial: Airavata</u>

Rooms:

Plenary Sessions	Regency C/D
Track A	Regency C
Track B	Regency D
Track C	Acapulco
Track D	Toronto (Doctoral Forum) and Hong Kong (Student Session)
Panel	Regency C/D
Lunch	Big Bar
Poster Session and Reception	Regency C/D

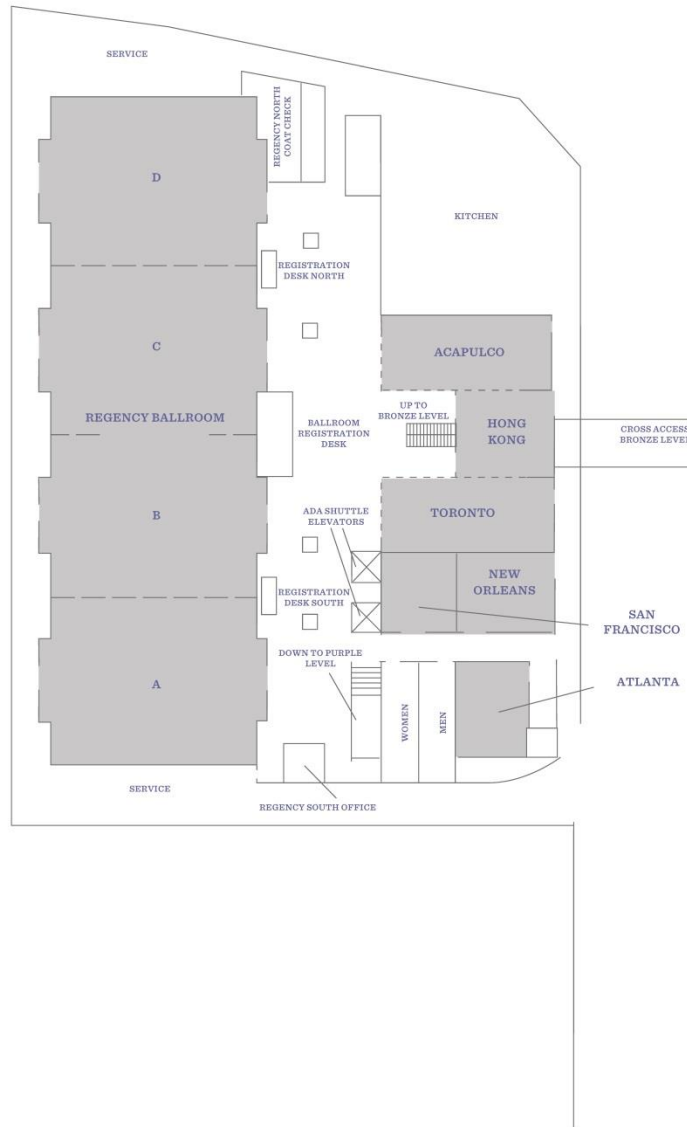


HYATT REGENCY CHICAGO
151 East Wacker Drive
Chicago, IL 60601
USA

T +1 312 565 1234
F +1 312 239 4541
chicagoregency.hyatt.com



FLOOR PLAN
Gold Level (West Tower)



IEEE/ACM CCGRID 2014
14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing
May 26th -29th, 2014 – Chicago USA



HYATT REGENCY CHICAGO

GUEST MAP

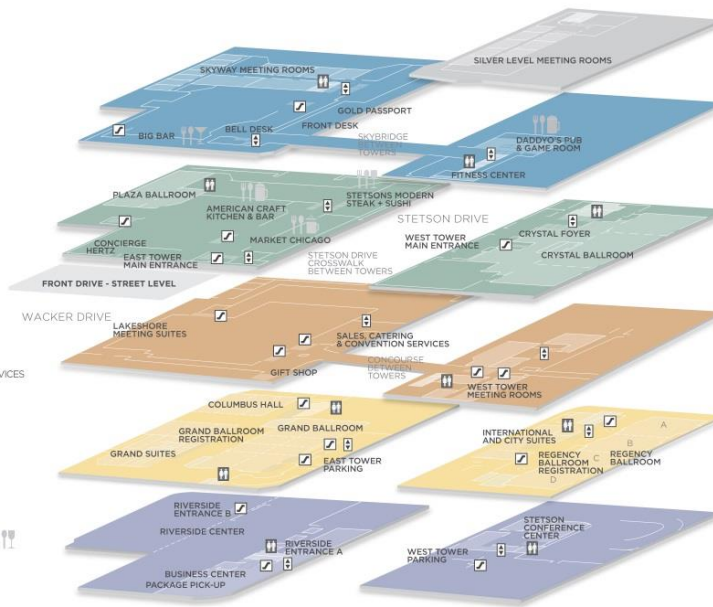


WELCOME TO HYATT REGENCY CHICAGO. Meeting rooms, ballrooms, restaurants and guest amenities are listed in alphabetical order and color coded by floor. For help, dial Guest Services at Extension 4460.

ACAPULCO West Tower, Gold Level	HAYMARKET West Tower, Bronze Level
ADDAMS West Tower, Silver Level	HERTZ East Tower, Green Level
AMERICAN CRAFT KITCHEN & BAR East Tower, Green Level	HONG KONG West Tower, Green Level
ATLANTA West Tower, Gold Level	HORNER West Tower, Silver Level
BELL DESK East Tower, Blue Level	LAKE SHORE MEETING SUITES East Tower, Bronze Level
BIG BAR East Tower, Blue Level	MARKET CHICAGO East Tower, Silver Level
BUCKINGHAM West Tower, Bronze Level	MCCORMICK West Tower, Silver Level
BURNHAM West Tower, Silver Level	NEW ORLEANS West Tower, Gold Level
BUSINESS CENTER East Tower, Purple Level	OGDEN West Tower, Silver Level
COLUMBIAN West Tower, Bronze Level	PACKAGE PICK-UP East Tower, Purple Level
COLUMBUS HALL (ROOMS A-L) East Tower, Gold Level	PICASSO West Tower, Bronze Level
COMISKEY West Tower, Bronze Level	PLAZA BALLROOM East Tower, Green Level
CONCIERGE East Tower, Green Level	REGENCY BALLROOM East Tower, Gold Level
CRYSTAL BALLROOM West Tower, Green Level	RIVERSIDE CENTER East Tower, Purple Level
DADDY O'S PUB & GAME ROOM West Tower, Blue Level	SALES, CATERING & CONVENTION SERVICES West Tower, Bronze Level
DUSABLE West Tower, Silver Level	SAN FRANCISCO West Tower, Silver Level
EAST TOWER MAIN ENTRANCE East Tower, Green Level	SANDBURG West Tower, Silver Level
EAST TOWER PARKING East Tower, Gold Level	SKYWAY MEETING ROOMS East Tower, Blue Level
FIELD West Tower, Silver Level	SOLDIER FIELD West Tower, Bronze Level
FITNESS CENTER West Tower, Blue Level	STETSON CONFERENCE CENTER West Tower, Purple Level
FRONT DESK East Tower, Blue Level	STETSONS MODERN STEAK + SUSHI East Tower, Green Level
GIFT SHOP East Tower, Blue Level	TORONTO West Tower, Gold Level
GOLD COAST West Tower, Bronze Level	WATER TOWER West Tower, Bronze Level
GOLD PASSPORT East Tower, Blue Level	WEST TOWER PARKING West Tower, Purple Level
GRAND BALLROOM East Tower, Gold Level	WRIGHT West Tower, Silver Level
GRAND BALLROOM REGISTRATION East Tower, Gold Level	WRIGLEY West Tower, Bronze Level
GRAND SUITES East Tower, Gold Level	

EAST TOWER

WEST TOWER



ESCALATORS, ELEVATORS AND RESTROOMS are indicated on each floor. Elevators are conveniently located throughout the hotel for guests with disabilities or where no escalator is present.
CROSSING BETWEEN TOWERS: Cross between towers via the Blue Level Skybridge or the Concourse on the Bronze Level. You may also cross on the Green Level via the crosswalk on Stetson Drive.



ACCOMMODATIONS

- 2,019 guestrooms, including 119 suites / parlors, 873 kings, 826 queen / queens, 98 double / doubles, 113 Regency Club®, 103 accessible rooms and all non-smoking rooms;
- AAA four-diamond rated
- Standard Guestroom Amenities:
 - Hyatt Grand Bed™
 - Wireless Internet access (fee)
 - Plasma television with in-room pay movies
 - Digital thermostat with occupancy sensor
 - In-room safety-deposit box
 - Full bath amenities, hair dryer
 - Coffeemaker with complimentary coffee
 - Iron / ironing board
 - iHome® alarm clock radio

SERVICES & FACILITIES

- Hyatt Express Check-In Kiosks
- Hyatt Fast Board™
- 1-800-CHECK-IN®
- In-room dining
- Full-service business center
- 24-hour Hyatt Stay Fit gym
- Wireless Internet in lobby
- Multilingual staff
- Assistive devices for persons with disabilities
- Currency exchange
- Safe-deposit boxes at front desk
- Laundry / dry cleaning
- Valet parking
- Gift shop
- Regency Club®
- Car rental desk

RESTAURANTS & BARS

- Stetson's Chop House and Bar—a traditional Chicago steakhouse at its finest
- BIG Bar—signature BIG cocktails, appetizers and desserts
- Daddy O's Irish Pub—a cozy Irish bar with pub fare
- The BISTRO At 151—breakfast and lunch in the atrium lobby
- The BAR At 151—atrium lounge for cocktails and a bite to eat
- J's Express—open 24-hours. Starbucks® coffee drinks, snacks, a quick meal on the go

CONFERENCES & BANQUETS

- 228,000 square feet of contemporary and flexible function space can be customized for the smallest to the largest groups
- Grand Ballroom, Regency Ballroom, Crystal Ballroom and Plaza Ballroom complement 63 meeting rooms and the 70,000 square foot Riverside Exhibition Center
- Five meeting spaces designed exclusively for small meetings
- Experienced staff, creative cuisine and unmatched event technology and design
- Meeting Concourse at your side from start to finish
- Reservations made easy through Passkey

POINTS OF INTEREST

Hyatt Regency Chicago is centrally located to give you the chance to see and experience as much as possible during your visit to our great city. The hotel is located in the heart of downtown, just steps away from the Magnificent Mile and close to many of the city's premier business, the theatre district, acclaimed shopping, and famous attractions. Here are just some of the things nearby:

- Magnificent Mile
- Navy Pier
- Grant Park and Millennium Park
- Buckingham Fountain
- Michigan Avenue and State Street Shopping
- Theatre and Financial District
- Museum of Science and Industry
- Shedd Aquarium
- Field Museum
- Adler Planetarium
- Art Institute of Chicago
- Museum of Contemporary Art
- Chicago History Museum
- Willis Tower (formerly Sears Tower)
- Soldier Field (Bears)
- Wrigley Field (Cubs)
- U.S. Cellular Field (White Sox)
- United Center (Bulls and Blackhawks)

IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing
14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing
May 26th - 29th, 2014 - Chicago USA

Conference Keynotes

The Transformative Impact of Computing and Communication in a Data-Driven World

Tuesday May 27th, 9:00 – 10:00

Dr. Farnam Jahanian,

The National Science Foundation Assistant Director for the Computer and Information Science And Engineering (CISE)

Abstract: We are witnessing unprecedented growth of scientific and social data, deep integration of the cyber and physical worlds, wireless connectivity at broadband speeds, and seamless access to resources in the cloud. These advances are transforming the way we work, play, communicate, learn, and discover. Investments in ambitious, long-term research and infrastructure, as well as in the development of a workforce empowered by computation and data, enable these advances and are a national imperative. This talk will focus on the technological advances and emerging trends that are shaping our future and accelerating the pace of discovery and innovation across all science and engineering disciplines. It will also describe how these trends inform priorities and programs at National Science Foundation.



Short Biography: Dr. Farnam Jahanian leads the National Science Foundation Directorate for Computer and Information Science and Engineering (CISE). He guides CISE in its mission to uphold the Nation's leadership in scientific discovery and engineering innovation through its support of fundamental research in computer and information science and engineering and of transformative advances in cyberinfrastructure. Dr. Jahanian is on leave from the University of Michigan, where he holds the Edward S. Davidson Collegiate Professorship and served as Chair for Computer Science and Engineering from 2007 - 2011 and as Director of the Software Systems Laboratory from 1997 - 2000. His research on Internet infrastructure security formed the basis for the Internet security company Arbor Networks, which he co-founded in 2001 and where he served as Chairman until its acquisition in 2010. He has testified before Congress on a broad range of topics, including cybersecurity and Big Data. Dr. Jahanian holds a master's degree and a Ph.D. in Computer Science from the University of Texas at Austin. He is a Fellow of ACM, IEEE and AAAS.

Taming Big Data with Berkeley Data Analytics Stack (BDAS)

Wednesday May 28th, 9:00 – 10:00

Dr. Ion Stoica,

Professor, Computer Science Division, University of California, Berkley

Abstract: Today, big and small organizations alike collect huge amounts of data, and they do so with one goal in mind: extract "value" through sophisticated exploratory analysis, and use it as the basis to make decisions as varied as personalized treatment and ad targeting. Unfortunately, existing data analytics tools are slow in answering queries, as they typically require to sift through huge amounts of data stored on disk, and are even less suitable for complex computations, such as machine learning algorithms. These limitations leave the potential of extracting value of big data unfulfilled. To address this challenge, we are developing BDAS, an open source data analytics stack that provides interactive response times for complex computations on massive data. To achieve this goal, BDAS supports efficient, large-scale in-memory data processing, and allows users and applications to trade between query accuracy, time, and cost. In this talk, I'll present the architecture, challenges, results, and our experience with developing BDAS, with a focus on Apache Spark, an in-memory cluster computing engine that provides support for a variety of workloads, including batch, streaming, and iterative computations. In a relatively short time, Spark has become the most active big data project in the open source community, and is already being used by over one hundred companies and research institutions.



Short Biography Dr. Ion Stoica is a Professor in the EECS Department at University of California at Berkeley. He received his PhD from Carnegie Mellon University in 2000. He does research on cloud computing and networked computer systems. Past work includes the Dynamic Packet State (DPS), Chord DHT, Internet Indirection Infrastructure (i3), declarative networks, replay-debugging, and multi-layer tracing in distributed systems. His current research focuses on resource management and scheduling for data centers, cluster computing frameworks, and network architectures. He is an ACM Fellow and has received numerous awards, including the SIGCOMM Test of Time Award (2011), and the ACM doctoral dissertation award (2001). In 2006, he co-founded Conviva, a startup to commercialize technologies for large scale video distribution, and in 2013, he co-founded Databricks as startup to commercialize technologies for Big Data processing.

Fault-Tolerance Techniques for Computing at Scale

Thursday May 29th, 9:00 – 10:00

2014 IEEE TCSC Award Keynote

Dr. Yves Robert,


Professor, Ecole Normale Supérieure de Lyon

Abstract: Resilience is a critical issue for large-scale platforms. This talk will survey fault-tolerant techniques for high-performance computing:

- Overview of failure types and typical probability distributions
- Brief discussion of application-specific techniques, such as ABFT
- The standard general-purpose technique, checkpoint and rollback recovery
- Recent extensions with replication, prediction and silent error detection
- Relevant execution scenarios, evaluated and compared through quantitative models.
- The talk includes several illustrative examples and targets a general audience.



Short Biography: Dr. Yves Robert received the PhD degree from Institut National Polytechnique de Grenoble. He is currently a full professor in the Computer Science Laboratory LIP at ENS Lyon. He is the author of 7 books, 130 papers published in international journals, and 195 papers published in international conferences. He is the editor of 11 book proceedings and 13 journal special issues. He is the advisor of 26 PhD theses. His main research interests are scheduling techniques and resilient algorithms for large-scale platforms. Yves Robert served on many editorial boards, including IEEE TPDS. He was the program chair of HiPC'2006 in Bangalore, IPDPS'2008 in Miami, ISPDC'2009 in Lisbon, ICPP'2013 in Lyon and HiPC'2013 in Bangalore. He is a Fellow of the IEEE. He has been elected a Senior Member of Institut Universitaire de France in 2007 and renewed in 2012. He holds a Visiting Scientist position at the University of Tennessee Knoxville since 2011.



Detailed Conference Program

Tuesday, May 27

08:30-08:45 Conference Opening

09:00-10:00 Keynote (Regency C/D)

The Transformative Impact of Computing and Communication in a Data-Driven World

Farnam Jahanian, the National Science Foundation Assistant Director for the Computer and Information Science and Engineering (CISE)

10:00-10:30 Break

10:30-12:00 Session 1: Best Papers (Regency C/D)

Chair: Kirk Cameron (Virginia Tech), Dimitrios S. Nikolopoulos (Queen's University of Belfast)

Improving I/O Throughput of Scientific Applications using Transparent Parallel Compression

Tekin Bicer, Jian Yin and Gagan Agrawal.

Performance models for CPU-GPU data transfers

Ben Van Werkhoven, Jason Maassen, Frank Seinstra and Henri Bal.

A User-level Infiniband-based File System and Checkpoint Strategy for Burst Buffers

Kento Sato, Kathryn Mohror, Adam Moody, Todd Gambelin, Bronis R. de Supinski, Naoya Maruyama and Satoshi Matsuoka.

12:00-13:00 Lunch (Big Bar)

13:30-15:30 Session 2

Session 2A: MapReduce (Regency C)

Chair: Ali R. Butt (Virginia Tech)

Tagged-MapReduce: A General Framework for Secure Computing with Mixed-sensitivity Data on Hybrid Clouds

Chunwang Zhang, Ee-Chien Chang and Roland H.C. Yap

Toward Detecting Compromised MapReduce Workers through Log Analysis

Eunjung Yoon and Anna Squicciarini

MapReduce Analysis for Cloud-archived Data

Balaji Palanisamy, Aameek Singh, Nagapramod Mandagere, Gabriel Alatorre and Ling Liu

Adaptive MapReduce Scheduling in Shared Environments

Jordà Polo, Yolanda Becerra, David Carrera, Jordi Torres, Eduard Ayguadé and Malgorzata Steinder

Session 2B: Energy and the Environment (Regency D)

Chair: Rong Ge (Marquette University)

Enabling Efficient Power Provisioning for Enterprise Applications

Balaji Subramaniam and Wu-Chun Feng

Analytical/ML Mixed Approach for Concurrency Regulation in Software Transactional Memory

Diego Rughetti, Pierangelo Di Sanzo, Bruno Ciciani and Francesco Quaglia

Bridging Data in the Clouds: An Environment-Aware System for Geographically Distributed Data Transfers

Radu Tudoran, Alexandru Costan, Rui Wang, Luc Bougé and Gabriel Antoniu

Cost-Efficient, Reliable, Utility-Based Session Management in the Cloud

Benjamin Byholm and Ivan Porres

Session 2C: Resource Management (Acapulco)

Chair: Kasidit Chanchio (Thammasat University)

A workflow-inspired, modular and robust approach to experiments in distributed systems

Tomasz Buchert, Lucas Nussbaum and Jens Gustedt

Resource Usage Control In Multi-Tenant Applications

Rouven Krebs, Simon Spinner, Nadia Ahmed and Samuel Kounev

Mapping Algorithms Optimizing the Overall Manhattan Distance for pre-occupied Cluster Computers in SLA-based Grid environments

Barry Linnert, Joerg Schneider and Lars-Olof Burchard

SLA-based Profit Optimization in Cloud Bursting PaaS

Djawida Dib, Nikos Parlavantzas and Christine Morin

Session 2D: SCALE Challenge (Toronto)

Chair: Douglas Thain (Norte Dame)

V for Vicissitude: The Challenge of Scaling Complex Big Data Workflows

Bogdan Ghit, Mihai Capotă, Tim Hegeman, Jan Hidders, Dick Epema, Alexandru Iosup

ToMaR - A Data Generator for Large Volumes of Content

Rainer Schmidt, Matthias Rella, Sven Schlarb

Exploring Infiniband Hardware Virtualization in OpenNebula towards Efficient High-Performance Computing

Tiago Pais Pitta de Lacerda Ruivo, Gerard Bernabeu, Gabriele Garzoglio, Steve Timm, Hyunwoo Kim, Seo-Young Noh, Ioan Raicu

Emulation at Very Large Scale with Distem

Tomasz Buchert, Emmanuel Jeanvoine, Lucas Nussbaum

15:30-16:00 Break

16:00-17:30 Session 3

Session 3A: Apps I (Regency C)

Chair: Esmond Ng (Lawrence Berkeley National Laboratory)

Evaluating Streaming Strategies for Event Processing across Infrastructure Clouds

Radu Tudoran, Kate Keahey, Gabriel Antoniu and Sergey Panitkin.

CUDAlign 3.0: Parallel Biological Sequence Comparison in Large GPU Clusters

Edans F. De O. Sandes, Guillermo Miranda, Alba Cristina M. A. de Melo, Xavier Martorell and Eduard Ayguade.

Opportunistic High Energy Physics Computing in User Space with Parrot

Paul Brenner, Doug Thain, Kevin Lannon and Dillon Skeeahan.

Session 3B: Message Passing (Regency D)

Chair: Satoshi Matsuoka (Tokyo Institute of Technology)

Towards an MPI-like framework for the Azure cloud platform

Dinesh Agarwal, Sara Karamati, Satish Puri and Sushil Prasad.

NoCMsg: Scalable NoC-Based Message Passing

Christopher Zimmer and Frank Mueller.

Modeling and Optimizing Large-Scale Wide-Area Data Transfers

Rajkumar Kettimuthu, Gayane Vardoyan, Gagan Agrawal and P. Sadayappan.

Session 3C: Elasticity and Adaptation (Acapulco)

Chair: Ioan Raicu (Illinois Institute of Technology)

A Language Support for Cloud Elasticity Management

Yousri Kouki, Frederico Alvares, Simon Dupont and Thomas Ledoux.

Elastic MapReduce Execution

Wei Xiang Goh and Kian-Lee Tan.

JCatascopia: Monitoring Elastically Adaptive Applications in the Cloud

Demetris Trihinas, George Pallis and Marios D. Dikaiakos.

Session 3D: Poster Oral Presentation (Regency C/D)

Room: Regency C/D

18:30-20:30 Posters + Conference Reception (Regency C/D)

Wednesday, May 28

09:00-10:00 Keynote (Regency C/D)

Taming Big Data with Berkeley Data Analytics Stack (BDAS)
Ion Stoica, University of California, Berkeley

10:00-10:30 Break

10:30-12:00 Session 4

Session 4A: Big Data (Regency C)

Chair: Yanlong Yin (Illinois Institute of Technology)

Towards a Collective Layer in the Big Data Stack
Thilina Gunarathne, Judy Qiu and Dennis Gannon.

Flexpath: Type-Based Publish/Subscribe System for Large-scale Science Analytics

Jai Dayal, Drew Bratcher, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Xuechen Zhang, Hasan Abbasi, Scott Klasky and Norbert Podhorszki.

Toward Transparent In Situ Data Transformations in ADIOS

David Boyuka li, Sriram Lakshminarasimhan, Xiaocheng Zou, Zhenhuan Gong, Jonathan Jenkins, Eric Schendel, Norbert Podhorszki, Qing Liu, Scott Klasky and Nagiza Samatova.

Session 4B: Storage and I/O Systems I (Regency D)

Chair: Frank Mueller (North Carolina State University)

HyCache+: Towards Scalable High-Performance Caching Middleware for Parallel File Systems
Dongfang Zhao, Kan Qiao and Ioan Raicu

A Flexible Framework for Asynchronous In Situ and In Transit Analytics for Scientific Simulations
Matthieu Dreher and Bruno Raffin

Iteration Based Collective I/O Strategy for Parallel I/O Systems
Zhixiang Wang, Xuanhua Shi, Hai Jin, Song Wu and Yong Chen

Session 4C: Algorithm (Acapulco)

Chair: Alfredo Cuzzocrea (ICAR-CNR and University of Calabria)

Efficiently Handling Skew in Outer Joins on Distributed Systems
Long Cheng, Spyros Kotoulas, Tomas Ward and Georgios Theodoropoulos

A PGAS Execution Model for Efficient Stencil Computation on Many-Core Processors
Mitsuru Ikei and Mitsuhsato Sato

A Branch-and-Bound Algorithm for Autonomic Adaptation of Multi-Cloud Applications
André Almeida, Francisco Dantas, Everton Cavalcante and Thais Batist.

Session 4D: Doctoral Forum (Toronto)

Chair: Judy Qiu (Indiana University)

12:00-13:00 Lunch (Big Bar)

13:30-15:30 Session 5

Session 5A: Scheduling (Regency C)

Chair: Leon Sung (Pacific Northwest National Laboratory)

Decentralized Scheduling and Load Balancing for Parallel Programs
Gary Jackson, Alan Sussman and Pete Keleher.

Multi-Objective Scheduling for Heterogeneous Server Systems with Machine Placement
Hongyang Sun, Patricia Stolf, Jean-Marc Pierson and Georges Da Costa.

PLASTiCC: Predictive Look-Ahead Scheduling for Continuous dataflows on Clouds
Alok Gautam Kumbhare, Yogesh Simmhan and Viktor K. Prasanna.

Link-Heterogeneous Work Stealing
Trong Tuan Vu and Bilel Derbel.

Session 5B: Virtual Machines (Regency D)

Chair: Christian Engelmann (Oak Ridge National Laboratory)

Time-Bound, Thread-Based Live Migration of Virtual Machines
Kasidit Chanchio and Phithak Thaenkaew.

A Reference Model for Virtual Machine Launching Overhead
Hao Wu, Shangping Ren, Gabriele Garzoglio, Steve Timm, Gerard Bernabeu and Seo-Young Noh.



Controlling the deployment of virtual machines on clusters and clouds for scientific computing in CBRAIN

Tristan Glatard, Marc-Etienne Rousseau, Pierre Rioux, Reza Adalat and Alan Evans.

MIMP: Deadline and Interference Aware Scheduling of Hadoop Virtual Machines

Wei Zhang, Sundaresan Rajasekaran, Timothy Wood and Mingfa Zhu.

Session 5C: Datacenters and Distributed Computing (Acapulco)

Chair: Pavan Balaji (Argonne National Laboratory)

Achieving Efficient Distributed Scheduling with Message Queues in the Cloud for Many-Task Computing and High-Performance Computing

Iman Sadooghi, Palur Sandeep, Ajay Anthony, Isha Kapur, Karthik Belagodu, Pankaj Purandare, Kiran Ramamurty, Ke Wang and Ioan Raicu.

Efficient Checkpointing of Virtual Machines using Virtual Machine Introspection

Ferrol Aderholdt, Fang Han, Stephen Scott and Thomas Naughton.

Analysis of Labor Efforts and Their Impact Factors to Solve Server Incidents in Datacenters

Ioana Giurgiu, Jasmina Bogojeska, Sergii Nikolaiev, George Stark and Dorothea Wiesmann.

CMcloud: Cloud Platform for Cost-Effective Offloading of Mobile Applications

Dongju Chae, Jihun Kim, Jangwoo Kim, Jong Kim, Seungjun Yang, Yeongpil Cho, Yongin Kwon and Yunheung Paek.

Session 5D: Doctoral Forum (Toronto)

Chair: Judy Qiu (Indiana University)

15:30-16:00 Break

16:00-17:30 Panel: Architect Cloud and HPC for the Big Data Era (Regency C/D)

Moderator: Xian-He Sun (Illinois Institute of Technology)

Panelist:

Andrew Baptist, VP of Technology, Cleversafe

Rajkumar Buyya, Professor, The University of Melbourne

Kirk Cameron, Professor, Virginia Tech

Ian Foster, Arthur Holly Compton Distinguished Service Professor, University of Chicago

Satoshi Matsuoka, Professor, Tokyo Institute of Technology

Ion Stoica, Professor, University of California, Berkeley

17:45-22:00 Banquet Dinner Cruise (Busses Depart from the Hyatt Foyer at 17:45)

Thursday, May 29

09:00-10:00 Keynote (Regency C/D)

Chair: Manish Parashar (Rutgers: The State University of New Jersey)

Fault-Tolerance Techniques for Computing at Scale

Yves Robert, Professor, Winner of the 2014 IEEE TCSC Award for Excellence in Scalable Computing, Ecole Normale Supérieure de Lyon

10:00-10:30 Break

10:30-12:00 Session 6

Session 6A: Apps II (Regency C)

Chair: Kyle Chard (University of Chicago)

A Credential Store for Multi-Tenant Science Gateways

Thejaka Amila Kanewala, Suresh Marru, Jim Basney and Marlon Pierce.

Cluster-based SNP Calling on Large-Scale Genome Sequencing Data

Mucahid Kutlu and Gagan Agrawal.

Platform calibration for load balancing of large simulations: TLM case

Cristian Ruiz, Mihai Alenxandru, Olivier Richard, Thierry Monteil and Herve Aubert.

Session 6B: Architecture (Regency D)

Chair: Suren Byna (Lawrence Berkeley National Laboratory)

Enhancing locality via caching in the GMU protocol

Hugo Pimentel, Paolo Romano, Sebastiano Peluso and Pedro Ruivo.

Energy-Efficient Collective Reduce and Allreduce Operations on Distributed GPUs

Lena Oden, Benjamin Klenk and Holger Fröning.

Network Topology Optimization For Data Aggregation

Soham Das and Sartaj Sahni.

Session 6C: Storage and I/O Systems II (Acapulco)

Chair: Zhihui Lu (Fudan University)

hatS: A Heterogeneity-Aware Tiered Storage for Hadoop

Krish K. R., Ali Anwar and Ali R. Butt.

A Study of Effective Replica Reconstruction Schemes at Node Deletion for HDFS

Asami Higai, Atsuko Takefusa, Hidemoto Nakada and Masato Oguchi.

A Novel Zero-Knowledge Scheme for Proof of Data Possession in Cloud Storage Applications

Nesrine Kaaniche, Ethmane El Moustaine and Maryline Laurent.

Session 6D: Student Travel Grant Awardees Round-Table (Hong Kong)

12:00-12:30 Conference Closing with Award Ceremony (Regency C/D)

12:30-13:30 Lunch Break (Big Bar)

13:30-15:30 Tutorials

15:30-16:00 Break

16:00-17:30 Tutorials

Poster Program

Tuesday May 27th, 16:00 – 20:30

An High Efficient Disk Scheduling Framework with QoS Mechanism in Xen-Based Cloud Platforms

Tseng-Yi Chen, Hsin-Wen Wei, Ying-Jie Chen, Nia-Yuan Chang, Tsan-Sheng Hsu and Wei-Kuan Shih

Automatic Provisioning, Deploy and Monitoring of Virtual Machines based on Security Service Level Agreement in the Cloud

Kazi Wali Ullah and Abu Shohel Ahmed

RESeED: A Tool for Regular Expression Search over Encrypted Data in Cloud Storage

Mohsen Amini Salehi, Thomas Caldwell, Alejandro Fernandez, Emmanuel Mickiewicz, Eric Rozier, Saman Zonouz and David Redberg

Wave: Trigger Based Synchronous Data Process System

Kun Lu, Mingming Sun and Xuehai Zhou

Network Traffic-Aware Virtual Machine Placement with Availability Guarantees Based on Shadows

Qian Zhang, Mingyu Li and Xiaohui Hu

An Architecture for Orchestrating Hadoop Applications in Hybrid Cloud

Carlos Roberto Senna, Luis Guilherme Cordioli Russi and Edmundo Roberto Mauro Madeira

Performance Evaluation of an IaaS Opportunistic Cloud Computing

Cesar Diaz, German Sotelo, Johnatan E. Pecero, Harold Castro, Mario Villamizar and Pascal Bouvry

Expanding Tasks of Logical Workflows into Independent Workflows for Improved Scalability Nicholas Hazekamp

Olivia Choudhury, Sandra Gesing, Douglas Thain and Scott Emrich

Provenance-Based Prediction Scheme for Object Storage System in HPC

Dong Dai, Yong Chen, Dries Kimpe and Robert Ross

An Adaptive Separation-Aware FTL for Improving the Efficiency of Garbage Collection in SSDs

Wei Xie and Yong Chen

Doctoral Symposium

Wednesday May 28th, 10:30 – 15:30

Towards Generic Metadata Management in Distributed Science Gateway Infrastructures

Richard Grunzke, Sandra Gesing, René Jäkel and Wolfgang E. Nagel

Compiler Optimization for Extreme-Scale Scripting

Timothy Armstrong, Justin Wozniak, Michael Wilde and Ian Foster

Metrics, Models and Methodologies for Energy-Proportional Computing

Balaji Subramaniam and Wu-Chun Feng

Supporting Queries and Analyses of Large-Scale Social Media Data with Customizable and Scalable Indexing

Techniques over NoSQL Databases

Xiaoming Gao and Judy Qiu

Parallel Computing with P2P Desktop Grids

Gary Jackson

Advanced Virtualization Techniques for High Performance Cloud Cyberinfrastructure

Andrew Younge and Geoffrey Fox

Proactive Workload Consolidation for Reducing Energy Cost over a Given Time Horizon

Milan de Cauwer, Deepak Mehta, Barry O'Sullivan, Helmut Simonis and Hadrien Cambazard

Towards a Multi-Cloud SLA Management Framework

Soodeh Farokhi

Strategy-proof Mechanisms for Resource Management in Clouds

Lena Mashayekhy and Daniel Grosu

Runtime Adaptation for Autonomic Heterogeneous Computing

Tom Scogland



Workshop Program

Monday May 26th, 8:30-17:30

1. **ExtremeGreen: Extreme Green & Energy Efficiency in Large Scale Distributed Systems (Acapulco, All Day)**

Laurent Lefevre, Marcos Dias de Assuncao, Wu-Chun Feng, Anne-Cecile Orgerie, Manish Parashar, Ivan Rodero

Improving the energy efficiency of large-scale distributed systems (e.g. data centres and Clouds) is a key challenge for both academic and industrial organizations. Although the topic has gained lots of attention over the past years, some of the proposed solutions often seem conservative and not easily applicable to large-scale systems. Their impact at large-scale remains to be proved. Hence, this workshop aims to provide a venue for discussion of ideas that can demonstrate "more than small % solution" to energy efficiency and their applicability to "real world". After the success of ExtremGreen2013 workshop, the ExtremeGreen2014 workshop will focus on scientific and industrial approaches and solutions that could have a large impact in terms of energy savings and energy efficiency. Clean-slate approaches and innovative solutions breaking conventional approaches are welcome. The workshop also welcomes submissions of work-in-progress papers on ideas that can have a large impact on improving the energy efficiency of large-scale distributed systems. The papers must provide preliminary results that demonstrate the originality and possible impact of the proposed solutions. Submissions will be reviewed by an international group of experts in distributed systems and energy efficiency.

2. **DPMSS: The 3rd Workshop on Data-Intensive Process Management in Large-Scale Sensor Systems (Hong Kong, All Day)**

Alfredo Cuzzocrea, Giancarlo Fortino, Omer Rana

The aim of DPMSS 2013 is to capture the new research trends and results in terms of design, architecture and applications for the management of processes and data in large-scale systems based on sensors, optimization of the processing of sensor data streaming; definition and use of innovative paradigms to develop applications in large-scale sensor networks, and their integration with Grid and Cloud infrastructures. This workshop will also identify potential research directions and technologies that will drive innovations within this domain. We anticipate this workshop to establish a pathway for the development of future-generation large-scale sensor-based systems. Specifically, the areas of interest are the following (but are not limited to), and, with particular emphasis, how these areas are investigated in the context of the integration of large-scale SNs with Grid and Cloud infrastructures.

3. **C4Bio: Cloud for Bio (Toronto, Morning)**

Jesus Carretero, Javier Garcia Blas, Manuel Desco, Dana Petcu

In the last 20 years, computational methods have become an important part of developing emerging technologies for the field of bioinformatics and biomedicine. As applications require more and more computational resources every day in order to increase the knowledge of biological and health processes, bio disciplines create new challenges of scale for computation, storage, and interpretation of petascale data that are difficult to solve with single clusters. Cloud computing has the potential to help solve these problems by offering a utility model based on highly flexible computing and storage capabilities and abstraction layers that allows overcoming many of the constraints present in dedicated systems. It may also increase research productivity by allowing sharing applications, tools, and algorithms in an easy way. The aim of the workshop is to promote the synergies of the cloud computing and bio communities by exploring applications developed in clouds, cloud systems enhancements for bio, and new trends and needs shared between both communities. Thus, this workshop will feature articles that discuss the application of cloud computing to bioinformatics and biomedicine, including solutions to area problems and architectural adaptation of cloud systems to fit those problems.

4. **CCGrid-Health (Toronto, Afternoon)**

Dagmar Krefting, Johan Montagnat, Sandra Gesing, Aaron Golden

This workshop aims at bringing together developers of medical applications and researchers in the field of distributed IT systems. On the one hand, it addresses researchers who are already employing distributed infrastructure techniques in medical applications, in particular scientists developing data- and compute-intensive medical applications that include multi-data studies, large-scale parameter scans or complex analysis pipelines. On the other hand, it addresses computer scientists working in the field of distributed systems interested in bringing new developments into medical applications. The goals are to exchange and discuss existing solutions and latest developments in both fields, and to gather an overview of challenges (technologies, achievements, gaps, roadblocks). The workshop further intends to identify common requirements to lead future developments in collaboration between Health and Computing Sciences, and to collaboratively explore new ideas and approaches to successfully apply distributed IT-systems in translational research, clinical intervention, and decision-making.

5. **DV&UV: Data Vitalization and Universal Village: Extended Version of Smart Cities (New Orleans, Morning)**

Ichiro Masaki, Bertrand David, Jun Zhang, Hao Sheng, Yajun Fang

Utilizing data from heterogeneous sources to support intelligent services for citizens, government and enterprises need to address several important issues, including (1) data sensing: how to obtain the diverse data in the scale of a city operation; how to effectively and efficiently extract entities from unstructured data; (2) data vitalization: how to connect data by identifying hidden relationships between data, how to form the giant knowledge base to facilitate better information seeking and provide high quality service; (3) event driven service: how users' personalized request in an event which normally integrated several domain smartness can be met by data oriented service, how users' behaviors can be discovered from heterogeneous data sources, what are the fundamental structures underlying the service provision. The workshop will provide a forum for researchers to showcase their efforts. Provision of intelligent service for users based on fully analyzing heterogeneous and unstructured data will foster a cross-disciplinary forum to further enhance existing bounds and create new connections among these communities. This workshop solicits contributions on researches and practices of data oriented solutions for smartness request in a smart city.

6. **WACC: The 2nd International Workshop on Assured Cloud Computing (Regency D, Morning)**

Roy H. Campbell, Andrew Martin, Aad van morsel

The workshop on Assured Cloud Computing brings together a community of researchers and practitioners from academia and industry that share an interest in designing dependable cloud architectures that can provide assurances with respect to properties such as security, reliability, and the timeliness of computations (or services). Some new "assured" target applications include, but are not limited to, real-time computations for monitoring, control of cyber-physical systems such as power systems, and mission critical computations for rescue and recovery.

7. **SCRAMBL: The 1st Workshop on Scalable Computing for Real-Time Big Data Applications (New Orleans, Afternoon)**

Charalampos Chelmiss, Marc Frincu, Bogdan Nicolae

The continuous growth of our society has led to complex systems of behavior and also to the need to optimize certain aspects of our day to day activities. Time sensitive applications such as real time power management for smart grids, traffic control or network monitoring require on demand large scale information processing and real time responses. The data these applications gather on a regular basis from monitoring sensors exceeds the normal storage and capacity power of normal machines or even clusters. In addition, the complexities arising from handling large relational datasets include but are not limited to data heterogeneity (i.e. variability), data quality (missing/approximate), data temporality (i.e. high-velocity), data volume or data streaming. Traditional computing platforms and even storage/processing models cannot simultaneously address these efficiently. Advances in the state of art in scalable computing platforms such as clouds, offer an ideal environment for showcasing and advancing methods for modeling, management, mining and analysis of real-time big data. At the same time novel models that take into account all these complexities need to be designed on top of these scalable systems. This workshop aims at providing a venue for designers, practitioners, researchers, developers, and industrial/governmental partners to come together, present and discuss leading research results, use cases, innovative ideas, challenges, and opportunities that arise from real-time big data applications.

8. **C4BIE: Cloud for Business, Industry and Enterprises (Regency D, Afternoon)**

G. Subrahmanya VRK Rao, Ezendu Ariwa, Shikharesh Majumdar, Raj Bala, Chander Kaher

Cloud computing is progressing towards more ubiquity as greater demands from customers and multifaceted capabilities from providers have started to unfold. While cloud computing has started to get global acceptance, dependability on cloud is yet to mature. C4BIE 2014 (Cloud for Business, Industry and Enterprises 2014) intends to gather the preeminent minds of the field to discuss and envision the path ahead for Cloud Computing, which would make sense for Business, Industry and Enterprises. Target audience include technologists from industries such as Financial, Retail, Hospitality, Engineering, Manufacturing and Logistics, Healthcare, Life Sciences, Telecom, Media & Entertainment etc. including CIOs, CTOs, VPs of technology, IT directors, Engineers, Architects, Specialists, managers and business executives. The workshop is also of interest to researchers & scientists from academia and other research institutes who are currently working in the field of Cloud Computing.

ExtremeGreen Workshop – Monday, May 26 Full Day

- 09:00-10:00 **Keynote**
Is Energy-Efficient Exascale and Oxymoron
 Wu Feng, Virginia Tech
- 10:00-10:30 **Break**
- 10:30-12:00 **Session 1**
Energy-Aware Profit Maximizing Scheduling Algorithm for Heterogeneous Computing Systems
 Kyle Tarplee, Anthony Maciejewski and Howard Jay Siegel
- Performance and Power Consumption Evaluation of MHD Simulation for Magnetosphere on Parallel Computer System with CPU Power Capping**
 Keiichiro Fukazawa, Tomonori Tshata, Kyohei Yoshida, Aruta Uehara, Masakazu Kuze, Masatsugu Ueda, Yuichi Inadomi, Koji Inoue and Mutsumi Aoyagi
- A Game-Theoretic Approach to Coalition Formation in Green Cloud Federations"**
 Marco Guazzone, Cosimo Anglano and Matteo Serenob
- 12:00-13:30 **Lunch**
- 13:30-14:30 **Keynote**
Green Computing meets Big Data
 Ivan Rodero, Rutgers University
- 14:30-15:30 **Session 2**
Energy-Aware Data Transfer Tuning
 Ismail Alan, Engin Arslan and Tefvik Kosar
- Energy Consumption of Photo Sharing in Online Social Networks**
 Fatemeh Jalali, Chrispin Gray, Arun Vishwanath, Robert Ayre, Tansu Alpcan, Kerry Hinton and Rodney Tucker

DPMSS Workshop – Monday, May 26 Full Day

- 08:50-09:00 **Welcome**
- 09:00-10:00 **Session 1 : Big Data Processing**
Fuzzy assisted event driven data collection from sensor nodes in Sensor-Cloud infrastructure
 Suman Sankar Bhunia, Nandini Mukherjee and Jayita Pal
- Enforcing QoS on OpenNebula-based Shared Clouds for Highly Dynamic, Large-Scale Sensing Data Streams**
 Rafael Tolosana-Calasanz, Jose' A. Banares, Omer Rana, Congduc Pham, Erotokritos Xydias, Charalampos Marmaras, Panagiotis Papadopoulos and Liana Cipcigan
- 10:00-10:30 **Break**
- A Schema-Transformation Method for Effective and Efficient OLAP over Clouds in Big Data Environments**
 Alfredo Cuzzocrea and Rim Moussa
- 10:30-11:15 **Discussion**
- 11:15-12:30 **Session 2: Applications: Buliding Data Analytics**
An Evaluation Framework for Buildings-oriented Wireless Sensor Networks
 Giancarlo Fortino, Antonio Guerrieri and Wilma Russo
- Cloud based Building Data Analytics**
 Ioan Petri, Omer Rana, Tom Beach, Mengsong Zou, Javier Diaz-Montes, Manish Parashar and Yacine Rezgui
- Discussion**
- 12:30-13:30 **Lunch**
- 13:30-14:45 **Session 3: Wireless Sensor Networks**
Multiobjective Communication Optimization for Cloud-assisted Wireless Body Area Networks
 Dung Phan, Junichi Suzuki, Shingo Omura, Katsuya Oba and Athanasios Vasilakos
- A New Parallelism-Aware Clustering Algorithm for Wireless Sensor Networks**
 Alireza T. Bolorchi, M. H. Samadzadeh and Nazanin Rahnavard
- 14:45-15:00 **Discussion and closing**

C4Bio/CCGrid–Health Workshop – Monday, May 26 Full Day

08:45-09:00	Welcome
09:00-10:00	Keynote (CCGrid-Health) Biomedical Clouds and Biomedical Data Commons and Why they are Changing Health Care Research Robert Grossman, University of Chicago
10:00-10:30	Keynote (C4Bio) Sustainable Ultrascale Computing Systems: Applications in Biomedicine and Bioengineering Jesus Carretero, Universidad Carlos III de Madrid
10:30-11:00	Break
11:00-12:30	CCGrid-Health Session 1 Global Initiative for Sentinel e-Health Network on Grid (GINSENG), Medical data integration and semantic developments for epidemiology Sébastien Cipièrre, Guillaume Ereteo, Alban Gaignard, Nouha Boujelben, Sébastien Gaspard, Vincent Breton, Frédéric Cervenansky, David Hill, Tristan Glatard, David Manset, Johan Montagnat, Jérôme Revillard and Lydia Maigne. Extending XNAT towards a Cloud-based Quality Assessment Platform for Retinal Optical Coherence Tomographies Jie Wu, Christoph Jansen, Maximilian Beier, Michael Witt and Dagmar Krefting. Distributed Detection of Cancer Cells in Cellular Spike Streams Abdul Hafeez, M. Mustafa Rafique and Ali R. Butt.
12:30-13:30	Lunch
13:30-15:00	C4Bio Session 1 Accelerating Comparative Genomics Workflows in a Distributed Environment with Optimized Data Partitioning Olivia Choudhury, Nicholas Hazekamp, Douglas Thain and Scott Emrich Integration of Clustering and Multidimensional Scaling to Determine Phylogenetic Trees as Spherical Phylograms Visualized in 3 Dimensions Ruan Yang, Geoffrey L. House, Saliya Ekanayake, Ursel Schütte, James D. Bever, Haixu Tang and Geoffrey C. Fox A Survey of Approaches and Frameworks to Carry out Genomic Data Analysis on the Cloud Philip Church and Andrzej Goscinski A storage policy for a hybrid federated cloud platform executing bioinformatics applications Deric Lima, Breno Moura, Gabriel Oliveira, Edward Ribeiro, Aleteia Araujo, Maristela Holanda, Roberto Togawa and Maria Emilia Walter
15:00-15:30	Break
15:30-16:30	C4Bio Session 2 A Performance Evaluation of Sequence Alignment Software in Virtualized Environments Zachary J. Estrada, Zachary Stephens, Cuong Pham, Zbigniew Kalbarczyk and Ravinshandar K. Iyer From scripted HPC-based NGS pipelines to workflows on the cloud Jacek Cala, Yaobo Xu, Eldarina Azfar Wijaya and Paolo Missier Evaluation of the feasibility of making large-scale X-ray tomography reconstructions on clouds Estefania Serrano, Guzman Bermejo, Javier Garcia Blas and Jesus Carretero
16:30-17:30	Closing



DV&UV Workshop – Monday, May 26 Morning

08:30-08:40	Welcome
08:40-10:30	Session 1
	Geospatial Sensor Web Resource Management System for Smart City: Design and Implementation Jia Li, Nengcheng Chen and Wei Wang.
	A Robust and Fast Reconstruction Framework for Noisy and Large Point Cloud Data Xiang Feng and Wanggen Wan
	Data vitalization's perspective towards smart city: a reference model for data service oriented architecture Zhang Xiong, Yanwei Zheng, Wenge Rong, Hao Sheng and Chao Li
10:30-11:00	Break
11:00-12:30	Session 2
	Pragmatic Oriented Data Interoperability for Smart Healthcare Information Systems Shixiong Liu, Weizi Li and Kecheng Liu
	Intelligent Congestion Avoidance Algorithm and System--Application of Data Vitalization Yan Huang, Hao Sheng, Chao Li and Zhang Xiong
	A Survey on Workflow Management and Scheduling in Cloud Computing Li Liu and Miao Zhang
	Variable window for outlier detection and impulsive noise recognition in range images Jian Wang and Lin Mei

WACC Workshop – Monday, May 26 Morning

08:30-08:45	Welcome
08:45-10:15	Session 1
	High Performance QoS-enabled S3 based Object Store Yusuke Tanimura, Seiya Yanagita and Takahiro Hamanishi
	On the Evaluation of VM provisioning time in Cloud platforms for Mission_Critical Infrastructures Antonio Marotta, Stefano Avallone, Gabriella Carrozza, Vittorio Manetti, Roberto Canonico and Luigi Battaglia
	Towards Cloud, Service and Tenent Classification for Cloud Computing Sebastian Jeuk, Jakub Szefer and Shi Zhou
10:15-10:45	Break
10:45-12:15	Panel: Assured Clouds for Critical Applications
12:15-12:30	Closing

SCRAMBL Workshop – Monday, May 26 Afternoon

- 13:30-15:00 Welcome & Keynote
C-AMAT: A Concurrent Model for Scalable Data Access
Xian-He Sun, Illinois Institute of Technology
- 15:00-15:30 Break
- 15:30-17:30 Session 1
Scalable Infrastructures for Data in Motion
David Ediger, Rob McColl, Jason Poovey and Dan Campbell
HDMQ: Towards In-Order and Exactly-Once Delivery using Hierarchical Distributed Message Queues
Dharmit Patel, Faraj Khasib, Iman Sadooghi and Ioan Raicu
A Scalable System for Community Discovery in Twitter during Hurricane Sandy
Yin Huang, Han Dong, Yelena Yesha and Shujia Zhou
A Scalable Real-time Photometric System for Automatic Astronomical Observations on Dome A
Lianmeng Li, Ce Yu, Sun Jizhou, Jian Xiao, Zhaohui Shang and Jiajun Li
- 17:30-17:45 Closing

C4BIE Workshop – Monday, May 26 Afternoon

- 13:30-14:35 Welcome & Keynote
Experiences with cloud computing as a discovery accelerator
Ian Foster, University of Chicago and Argonne National Laboratory, Chicago
- 14:35-15:30 Keynote
Exploring the "A" of "SMAC" - Tales from the Data Trenches of Display Advertising
Claudia Perlich, Distillery
- 15:30-16:00 Break
Improving Resource Matchmaking Through Feedback Integration"
Christian Haas, Ioan Petri and Omer Rana
- 16:00-17:30 **MO-BIZZ: Fostering Mobile Business through Enhanced Cloud Solutions"**
Alexander Stanik, Odej Kao, Rui Martins, Antonio Cruz and Dimitrios Tektonidis
A comparative study of predictive models for cloud infrastructure management"
Mahesh Balaji, G Subrahmanya Vrk Rao and Ch. Aswani Kumar



Tutorial Program

Time: 13:30PM-17:30PM Thursday May 29th.

1. **Autonomic Clouds (Regency D)**

Omer Rana and Manish Parashar

Cloud computing continues to increase in complexity due to a number of factors: (i) increasing availability of configuration options from public Cloud providers (Amazon, for instance, offers over 4000 different configuration options) (ii) increasing variability and types of application instances that can be deployed over such platforms, such as tuning options in hypervisors that enable different virtual machine instances to be associated with physical machines, storage, compute and I/O preferences that offer different power and price, to operating system configurations that provide differing degrees of security. This complexity can also be seen in enterprise-scale datacenters that dominate computing infrastructures in industry, which are growing in size and complexity, and are also enabling new classes and scales of complex business applications. Autonomic computing offers self-* capabilities that enable self-management of systems. Proposed as a vision by IBM Research, the concepts in Autonomic Systems are much older and can be applied to each component within a Cloud system (resource manager/scheduler, power manager etc), or could be applied within an application that makes use of such a Cloud system. Understanding where such capability can be most effectively used is a decision variable often hard to fully appreciate and explored in this tutorial.

2. **Algorithms and tools to increase the efficiency of data centers: the case of Eco4Cloud (Acapulco)**

Carlo Mastroianni and Ian Taylor

Improving the efficiency of data centers, with a focus on power consumption and carbon emission, is a topical theme on which we are witnessing a huge and increasing amount of academic and industrial research. Despite the notable progress in reducing the PUE (Power Usage Effectiveness), which measures physical efficiency, there is much room for improvement in computational efficiency. The IT resources of data centers are not used efficiently as most servers exploit only a fraction of their computational power. A viable solution is to dynamically consolidate virtual machines on as few servers as possible, and either put the remaining servers in a sleep state or reuse them to accommodate additional workload. The tutorial will give the rationale and the mathematical basis of the problem, and will review some state-of-the-art algorithms. Then, the tutorial will introduce Eco4Cloud (www.eco4cloud.com), a bio-inspired approach for the adaptive and dynamic assignment of virtual machines to servers, as well as the tools used to assess the performance: mathematical models, simulation and physical deployment. A demo will then show how the Eco4Cloud integrated dashboard can be used to monitor the utilization of resources and consolidate the load by dynamically assigning virtual machines to physical servers.

3. **Accelerating Big Data Processing with Hadoop and Memcached on Modern Clusters (Regency C)**

Dhabaleswar Panda and Xiaoyi Lu

Apache Hadoop is gaining prominence in handling Big Data and analytics. Similarly, Memcached in Web 2.0 environment is becoming important for large-scale query processing. These middleware are traditionally written with sockets and do not deliver best performance on clusters with modern high performance networks. In this tutorial, we will provide an in-depth overview of the architecture of Hadoop components (HDFS, MapReduce, RPC, HBase, etc.) and Memcached. We will examine the challenges in re-designing the networking and I/O components of these middleware with modern interconnects, protocols (such as InfiniBand, iWARP, RoCE, and RSocket) with RDMA and storage architecture. Using the publicly available RDMA for Apache Hadoop (<http://hadoop-rdma.cse.ohio-state.edu>) software package, we will provide case studies of the new designs for several Hadoop components and their associated benefits. Through these case studies, we will also examine the interplay between high performance interconnects, storage systems (HDD and SSD), and multi-core platforms to achieve the best solutions for these components.



4. **Globus: Scalable Research Data Management Infrastructure for Campuses and High-Performance Computing Facilities (Toronto)**

Steve Tuecke and Rajkumar Kettimuthu

Rapid growth of data in science is placing massive demands on campus computing centers and high-performance computing (HPC) facilities. Computing facilities must provide robust data services built on high-performance infrastructure, while continuing to scale as needs increase. Traditional research data management (RDM) solutions are difficult to use and error-prone, and the underlying networking and security infrastructure is often complex and inflexible, resulting in user frustration and sub-optimal resource usage. An approach that is increasingly common in HPC facilities includes software-as-a-service (SaaS) solutions like Globus for moving, syncing, and sharing large data sets. SaaS approach allows HPC resource owners and systems administrators to deliver enhanced RDM services to end-users at optimal quality of service, while minimizing administrative and operations overhead associated with traditional software. Usage of Globus has grown rapidly, with more than 14,500 registered users and over 37 petabytes moved. Globus's reliable file transfer and data sharing services are key functionalities for bridging campus and external resources, and is enabling scientists to easily scale their research workflows. Tutorial attendees will be introduced to RDM functions of Globus, learn how scientists and resource owners are using Globus, and have the opportunity for hands-on interaction with Globus at various levels of technical depth.

5. **Developing Computational Science Gateways using Apache Airavata (New Orleans)**

Suresh Marru, Marlon Pierce and Don Kushan Wijeratne

This tutorial is aimed at science gateway developers and program managers who are responsible for delivering gateway capabilities to scientific communities. Science gateways, or Web portals, are an important mechanism for broadening and simplifying access to computational grids, clouds, and campus resources. Gateways provide science-specific user interfaces to end users who are unfamiliar with or need more capabilities than provided by command-line interfaces. Gateways share many core functionalities beneath their domain-specific user interfaces, which has led us to develop the Apache Airavata system. Airavata provides the services needed to translate science application-centric gateway requirements into resource-centric middleware APIs, as well as adding value by tracking a user's project metadata, storing provenance information for later reproducibility and sharing, and coordinating multi-staged workflow tasks. In this tutorial, we will introduce the audience to Airavata through introductory demonstrations and walk through hands-on exercises. Participants will learn how to use Airavata to manage computations on XSEDE supercomputers and campus cluster resources. Participants will be given hand-on exercises using Airavata through a Javabased reference implementation portal, which they can use after the tutorial as the basis for their own gateway. Participants who already operate gateways and have preferred Web frameworks will learn how to use the Airavata API through language-specific (PHP, Java, Python, JavaScript, etc) client development kits enabled by Airavata's Apache Thrift-based API.