
ACM HPDC 2010 Program

hpdc2010.eecs.northwestern.edu

Sunday, June 20 through Tuesday, June 22

Open Grid Forum (OGF29) [details available separately via ogf.org]

Cloud demos on Monday evening, June 21.

Monday, June 21

Breakfast (8-9am)

Workshops (9am-6pm)

- W1: ScienceCloud: Workshop on Scientific Cloud Computing (8:45am-6:15pm)
Room: Superior 1
- W2: Emerging Computational Methods for the Life Sciences (8:30am-5:00pm)
Room: Superior 3
- W3: MDQCS: Managing Data Quality for Collaborative Science (9:00am-12:30pm)
Room: Superior 2
- W4: LSAP: Large-Scale System and Application Performance (1:30-6:00pm)
Room: Superior 2

Workshop programs are available on symposium web site

Lunch (12:30pm-1:30pm, Room: LaSalle II)

Coffee Breaks (10:30-11am and 3:30-4pm)

Workshop Social Event (7-10pm)

HPDC 2010 Workshop Social Event Sponsored by Microsoft Research and the Digital Science Center at Indiana University. Event is open to all workshop and conference+workshop registrants, including students. Nature of event is TBD, please visit web site for more information.

Tuesday, June 22

Breakfast (8-9am)

Workshops (9am-5pm)

- W5: CLADE: Challenges of Large Applications in Distributed Environments (9:30am-5pm)
Room: Michigan
- W6: DIDC: Data Intensive Distributed Computing (9am-5pm)
Room: Superior III
- W7: MAPREDUCE: MapReduce and its Applications (9am-5pm)
Room: Superior I
- W8: VTDC: Virtualization Technologies for Distributed Computing (9am-5pm)
Room: Superior II

Workshop programs are available on symposium web site

Lunch (12:30pm-1:30pm, Room: LaSalle II)

Coffee Breaks (10:30-11am and 3:30-4pm)

Reception (7pm-10pm)

Wednesday, June 23

Breakfast (7-8am, Room: LaSalle II)

Opening Remarks from Conference Organizers (8-8:30am, Room: LaSalle I)

K1: Keynote by Guy Steele, Sun/Oracle (8:30-9:30am, Room: LaSalle I)
HPDC 2010 Keynote Series Sponsored by Intel
(Introducer: Peter Dinda)

Coffee Break (9:30-10am)

S1: Best Papers (10-11:30am, Room: LaSalle I) (Session Chair: Peter Dinda)

- *Acknowledgement of best paper, best student paper, best short paper, and best student short paper*
- Horizon: Efficient Deadline-Driven Disk I/O Management for Distributed Storage Systems
Anna Povzner (UC Santa Cruz), Darren Sawyer (NetApp), Scott Brandt (UC Santa Cruz)
- Run-time Optimizations for Replicated Dataflows on Heterogeneous Environments
George Teodoro (Universidade Federal de Minas Gerais), Timothy Hartley (Ohio State University), Umit Catalyurek (Ohio State University), Renato Ferreira (Universidade Federal de

Minas Gerais)

HPDC 2010 Best Paper and Best Student Paper Awards sponsored by Google

HPDC 2010 Best Short Paper and Best Student Short Paper Awards sponsored by NVIDIA

Lunch (11:30am-1pm, Room: LaSalle II)

S2: Workflows (1-2:30pm, Room: LaSalle I) (Session Chair: Ioan Raicu)

- DataSpaces: An Interaction and Coordination Framework for Coupled Simulation Workflows
Ciprian Docan (Rutgers), Manish Parashar (Rutgers), Scott Klasky (Oak Ridge National Lab)
- ParaTrac: A Fine-Grained Profiler for Data-Intensive Workflows
Nan Dun (University of Tokyo), Kenjiro Taura (University of Tokyo), Akinori Yonezawa (University of Tokyo)
- Performance Analysis of Dynamic Workflow Scheduling in Multicenter Grids
Ozan Sonmez (Delft University of Technology), Nezih Yigitbasi (Delft University of Technology), Saeid Abrishami (Delft University of Technology), Alexandru Iosup (Delft University of Technology), Dick Epema (Delft University of Technology)

Coffee Break (2:30-3pm)

S3: Resources and Clouds (3-4:30pm, Room: LaSalle I) (Session Chair: Dongyan Xu)

- Software Architecture Definition for On-demand Cloud Provisioning
Clovis Chapman (University College London), Wolfgang Emmerich (University College London), Fermin Galan Marquez (Telefonica I+D), Stuart Clayman (University College London), Alex Galis (University College London)
- High Occupancy Resource Allocation for Grid and Cloud systems, a Study with DRIVE
Kyle Chard (Victoria University of Wellington), Kris Bubendorfer (Victoria University of Wellington), Peter Komisarczuk (Victoria University of Wellington)
- Highly Available Component Sharing in Large-Scale Multi-Tenant Cloud Systems
Juan Du (North Carolina State University), Xiaohui Gu (North Carolina State University), Douglas Reeves (North Carolina State University)

Coffee Break (4:30-5pm)

P1: Posters (papers given below) (5-6:30pm, LaSalle II) (Session Chair: Matei Ripeanu)

Banquet (7pm-10pm, Room: Superior)

Thursday, June 24

Breakfast (7-8:30am, Room: LaSalle II)

K2: Keynote by Randal Bryant, Carnegie Mellon University (8:30-9:30am, Room: LaSalle I)

HPDC 2010 Keynote Series Sponsored by Intel
(Introducer: Wu-chun Feng)

Coffee Break (9:30-10am)

S4: Map Reduce and Debugging (10am-11:30am, Room: LaSalle I) (Session Chair: Henri Bal)

- **MOON: MapReduce On Opportunistic eNvironments**
Heshan Lin (Virginia Tech), Xaisong Ma (North Carolina State University and Oak Ridge National Lab), Jeremy Archuleta (Virginia Tech), Wu-chun Feng (Virginia Tech), Mark Gardner (Virginia Tech), Zhe Zhang (Oak Ridge National Lab)
- **MRAP: A Novel MapReduce-based Framework to Support HPC Analytics Applications with Access Patterns**
Saba Sehrish (University of Central Florida), Grant Mackey (University of Central Florida), Jun Wang (University of Central Florida), John Bent (Los Alamos National Lab)
- **Data Centric Highly Parallel Debugging**
David Abramson (Monash University), Minh Ngoc Dinh (Monash University), Donny Kuniawan (Monash University), Bob Moench (Cray), Luiz DeRose (Cray)

Lunch (11:30am-1pm, Room: LaSalle II)

S5: Data Centers and Virtualization (1-2:30pm, Room: LaSalle I) (Session Chair: Kate Keahey)

- **Thermal Aware Server Provisioning For Internet Data Centers**
Zahra Abbasi (Arizona State University), Georgios Varsamopoulos (Arizona State University), Sandeep Gupta (Arizona State University)
- **I/O Scheduling Model of Virtual Machine Based on Multi-core Dynamical Partitioning**
Yanyan Hu (Beihang University), Xiang Long (Beihang University), Jiong Zhang (Beihang University), Jun He (Beihang University), Li Xia (Beihang University)
- **A Practical Way to Extend Shared Memory Support Beyond a Motherboard at Low Cost**
Hector Montaner (Universitat Politècnica de València), Federico Silla (Universitat Politècnica de València), Jose Duato (Universitat Politècnica de València)

Coffee Break (2:30-3pm)

S6: Storage and I/O (3-4:30pm, Room: LaSalle I) (Session Chair: Karsten Schwan)

- **A GPU Accelerated Storage System**
Samer Al-Kiswany (University of British Columbia), Abdullah Gharaibeh (University of British Columbia)

Columbia), Sathish Gopalakrishnan (University of British Columbia), Matei Ripeanu (University of British Columbia)

- **Computation Mapping for Multi-Level Storage Cache Hierarchies**
Mahmut Kandemir (Pennsylvania State University), Sai Muralidhara (Pennsylvania State University), Mustafa Karakoy (Pennsylvania State University), Seung Woo Son (Argonne National Lab)
- **Caching in on Hints for Better Prefetching and Caching in PVFS and MPI-IO**
Christina Patrick (Pennsylvania State University), Mahmut Kandemir (Pennsylvania State University), Mustafa Karaköy (Imperial College), Seung Woo Son (Argonne National Lab), Alok Choudhary (Northwestern University)

I1: Industry Session (4:30-6:30pm, Lasalle I) (Session Organizer: Kate Keahey)

Dinner Cruise on the *Odyssey* (7-10pm)

The HPDC 2010 dinner cruise is included for all conference+workshop registrants, including students. You will receive a ticket with your registration. Additional tickets can also be purchased. Navy Pier, the departure point for the cruise, is within easy walking distance of the conference hotel. It is important to be on time since ship departure is punctual.

The cruise is sponsored by the Department of Electrical Engineering and Computer Science at Northwestern University and the NSF Center for Autonomic Computing Lab at the University of Arizona.

Friday, June 25

Breakfast (7-8:30am, Room: Fairbanks)

K3: Keynote by Robert Harrison, Oak Ridge National Lab (8:30-9:30am, Room: Ohio)

*HPDC 2010 Keynote Series Sponsored by Intel
(Introducer: Geoffrey Fox)*

Coffee Break (9:30-10am)

S7: Applications and Provenance (10-11:30am, Room: Ohio) (Session Chair: Jack Lange)

- **Dimension Reduction and Visualization of Large High-Dimensional Data via Interpolation**
Seung-Hee Bae (Indiana University), Jong Youl Choi (Indiana University), Xiaohong Qiu (Indiana University), Geoffrey Fox (Indiana University)
- **New Caching Techniques for Web Search Engines**
Mauricio Marin (Yahoo! Research Latin America), Veronica Gil-Costa (Yahoo! Research Latin

America), Carlos Gomez-Pantoja (Yahoo! Research Latin America)

- *Mendel: Efficiently Verifying the Lineage of Data Modified in Multiple Trust Domains*
Ashish Gehani (SRI International), Minyoung Kim (SRI International)

Lunch (11:30am-1pm, Room: Market Place)

D1: Panel : Expanding Parallel Programming to a Wider Audience (1-2:30pm, Room: Ohio)
(Moderator: Peter Dinda)

Guy Steele, Randal Bryant, Robert Harrison, John Reppy, Robert Findler, TBD

Coffee Break (2:30-3pm)

S8: Communication and Scheduling (3-4:30pm, Room: Ohio) (Session Chair: Thomas Stricker)

- **PV-EASY: A Strict Fairness Guaranteed and Prediction Enabled Scheduler in Parallel Job Scheduling**
Yulai Yuan (Tsinghua University), Guangwen Yang (Tsinghua University), Yongwei Wu (Tsinghua University)
- **XCo: Explicit Coordination to Prevent Network Fabric Congestion in Cloud Computing Cluster Platforms**
Vijay Shankar Rajanna (SUNY Binghamton), Smit Shah (SUNY Binghamton), Anand Jahagirdar (SUNY Binghamton), Kartik Gopalan (SUNY Binghamton)
- **Scalability of Communicators and Groups in MPI**
Humaira Kamal (University of British Columbia), Seyed Mirtaheri (University of British Columbia), Alan Wagner (University of British Columbia)

Coffee Break (4:30-5pm)

W1: Wild and Crazy Ideas (5-6:30pm, Room: Ohio) (Session chair: Fabian Bustamante)

Conference Ends

Posters

- **Impact of Sub-optimal Checkpoint Intervals on Application Efficiency in Computational Clusters**
William Jones (Coastal Carolina University), John Daly (Department of Defense), Nathan DeBardleben (Los Alamos National Lab)
- **Application Live Migration with GPU Acceleration in Personal Cloud**
Ben Lin (Intel), Vince Wang (Intel), Kyle Gui (Intel)
- **A Service Composition Framework for Market-Oriented High Performance Computing Cloud**
Tran Vu Pham (Ho Chi Minh City University of Technology), Hani Jamjoom (IBM Research), Kirk

Jordan (IBM Research), Zon-Yin Shae (IBM Research)

- GPU-based Parallel Householder Bidiagonalization
Fangbin Liu (Universteit van Amsterdam), Frank J. Seinstra (Vrije Univeriteit)
- Authorization within Grid-Computing Using Certificateless Identity-Based Proxy Signature
Mohamed Amin Jabri (Tokyo Institute of Technology), Satoshi Matsuoka (Tokyo Institute of Technology)
- A Global Address Space Framework for Irregular and Dynamic Applications
Kentaro Hara (University of Tokyo), Kenjiro Taura (University of Tokyo)
- Improving Host Swapping Using Adaptive Prefetching and Paging Notifier
Wenzhi Chen (Zhejiang University), Huijun Chen (Zhejiang University), Wei Huang (Zhejiang University), Xiaoqian Chen (Zhejiang University), Dapeng Huang (Zhejiang University)
- Resource Provisioning with Budget Constraints for Adaptive Applications in Cloud Environments
Qian Zhu (Ohio State University), Gagan Agrawal (Ohio State University)
- SoftPower: Fine-Grain Power Estimations Using Performance Counters
Min Yeol Lim (Renaissance Computing Institute), Allan Porterfield (Renaissance Computing Institute), Robert Fowler (Renaissance Computing Institute)
- Simultaneous Performance Exploration and Optimized Search with Volunteer Computing
L. Richard Moore Jr. (Lockheed Martin), Matthew Kopala (L3 Communications), Thomas Mielke (Boeing), Michael A. Krusmark (L3 Communications), Kevin A. Gluck (Air Force Research Lab)
- Lessons learned from moving Earth System Grid data sets over a 20 Gbps wide-area network
Raj Kettimuthu (Argonne National Lab and University of Chicago), Alex Sim (Lawrence Berkeley National Lab), Dan Gunter (Lawrence Berkeley National Lab), Bill Allcock (Argonne National Lab), Peer-Timo Bremer (Lawrence Livermore National Lab), John Bresnahan (Argonne National Lab and University of Chicago), Andrew Cherry (Argonne National Lab), Lisa Childers (Argonne National Lab and University of Chicago), Eli Dart (Energy Sciences Network), Ian Foster (Argonne National Lab and University of Chicago), Kevin Harms (Argonne National Lab), Jason Hick (Lawrence Berkeley National Lab), Jason Lee (Lawrence Berkeley National Lab), Michael Link (University of Chicago and Argonne National Lab), Jeff Long (Lawrence Livermore National Lab), Keith Miller (Data Direct Networks), Vijaya Natarajan (Lawrence Berkeley National Lab), Valerio Pascucci (University of Utah), Ken Raffanetti (Argonne National Lab), David Ressman (Argonne National Lab), Dean Williams (Lawrence Livermore National Lab), Loren Wilson (Argonne National Lab), Linda Winkler (Argonne National Lab)
- Energy-Aware Online Provisioning for Virtualized Clouds and Data Centers
Ivan Rodero (Rutgers University), Juan Jaramillo (Rutgers University), Andres Quiroz (Rutgers University), Manish Parashar (Rutgers University) and Francesc Guim (Intel Labs and Technical University of Catalonia)
- Brown Dwarf: Distributing the Power of OLAP
Katerina Doka (National Technical University of Athens), Dimitrios Tsoumakos (National Technical University of Athens), Nectarios Koziris (National Technical University of Athens)

- Morco: Middleware Framework for Long-running Multi-component Applications on Batch Grids
Sivagama Sundari Murugavel (Indian Institute of Science), Sathish S Vadhiyar (Indian Institute of Science), Ravi S Nanjundiah (Indian Institute of Science)
- Lightning: Self-Adaptive, Energy-Efficient, Storage-Efficient, Tiered, Automated Policy-Driven Green Cloud Storage System
Rini Kaushik (University of Illinois, Urbana-Champaign), Roy Campbell (University of Illinois, Urbana-Champaign), Klara Nahrstedt (University of Illinois, Urbana-Champaign), Ludmila Cherkasova (HP Labs)
- Early Defense - Enabling Attribute-Based Authorization in Grid Firewalls
Jan Wiebelitz (Leibniz University), Michael Brenner (Leibniz University), Christopher Kunz (Leibniz University), Matthew Smith (Leibniz University)
- Design and Evaluation of a Self-healing Kepler for Scientific Workflows
Arjun Hary (University of Arizona), Ali Akoglu (University of Arizona), Youssif AlNashif (University of Arizona), Salim Hariri (University of Arizona), Darrel Jenerette (UC Riverside)
- ParaText: Scalable Text Modeling and Analysis
Daniel Dunlavy (Sandia National Labs), Timothy Shead (Sandia National Labs), Eric Stanton (Sandia National Labs)
- DiscFinder: A Data-Intensive Scalable Cluster Finder for Astrophysics
Bin Fu (Carnegie Mellon University), Julio Lopez (Carnegie Mellon University), Kai Ren (Carnegie Mellon University), Eugene Fink (Carnegie Mellon University), Garth Gibson (Carnegie Mellon University)
- Improving MPI Communication Via Data Type Fission
Benjamin Perry (University of Delaware), Martin Swany (University of Delaware)
- Pwrake: A Parallel and Distributed Flexible Workflow Management Tool for Wide-area Data Intensive Computing
Masahiro Tanaka (University of Tsukuba), Osamu Tatebe (University of Tsukuba)
- Layout-Aware Collective I/O
Yong Chen (Illinois Institute of Technology), Huaiming Song (Illinois Institute of Technology), Rajeev Thakur (Illinois Institute of Technology), Xian-He Sun (Illinois Institute of Technology)